siam 2025®

Searchable Abstracts Document

SIAM Conference on Computational Science and Engineering (CSE25)

March 3–March 7, 2025

This document was current as of February 25, 2025. Abstracts appear as submitted.

siam®
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP1

Strong Stability Preserving Methods for Evolution of Hyperbolic Pdes

Strong stability preserving (SSP) RungeKutta methods are desirable when evolving in time problems that have discontinuities or sharp gradients and require nonlinear non-inner-product stability properties to be satisfied. However, SSP method suffer from many barriers and bounds. This talk will present the state of the art in SSP theory, including results for various types and classes of numerical methods, such as multi-step, multi-stage, and multi-derivative methods, as well as additive methods and integrating factor methods. This overview of the topic will allow users to determine if SSP methods are suitable for their own needs, and highlight open problems in the area.

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

IP2

To the Exascale and Beyond: Computing Challenges in Hpc

As we stand at the precipice of the exascale era, the future of traditional High-Performance Computing (HPC) remains shrouded in uncertainty. The landscape is evolving as practitioners navigate the complexities of heightened parallelism, memory bandwidth constraints, and the intricate hierarchies of modern memory architectures. Yet, the horizon brings both challenges and opportunities. The meteoric rise of artificial intelligence presents a tantalizing array of lower-precision computational methods, promising breakthroughs if we can adapt our algorithms. Simultaneously, the conclusion of Moores Law ushers in a new era where power efficiency becomes paramount, necessitating a shift towards programming that is both power-conscious and performance-optimized. In this presentation, I will explore these transformative challenges through the lens of lattice quantum chromodynamics (LQCD) calculationsa top HPC application. I will delve into strategies for navigating this evolving landscape and optimizing computational deployment in this new frontier, offering insights into how we can harness emerging paradigms to drive the next wave of computational excellence.

Kate Clark
NVIDIA
mclark@nvidia.com

IP3

Virtual Lungs in Respiratory Medicine: Multiscale Pulmonary Models and Clinical Applications

Mechanical ventilators proved crucial in the clinical management of acute respiratory failure during the recent COVID-19 pandemic. However, mortality rates of COVID-19 ventilated patients reached up to 90% during the most critical times. While the cause is multifactorial, such high mortality underscores the need for personalized solutions to improve clinical outcomes. One frequent problem is the heterogeneity found in the patients response to mechanical ventilation, which is hard to predict using current medical technology. Can multiscale models of the lungs elucidate whole-lung behavior from alveolar-level microstructural features? Can image-informed numerical simulations of the respiratory system explain variability during mechanical ventilation? In this talk, I will explore some of these questions by presenting a finite-deformation homogenization framework for lung poromechanics. I will also show how non-linear finite-element simulations based on this framework predict key clinical variables that characterize the response of human lungs. Finally, I will discuss our current work toward creating personalized virtual lungs and how we use them to simulate the lung response during mechanical ventilation.

Daniel Hurtado
Pontificia Universidad Católica de Chile
daniel.hurtado@uc.cl

IP4

Physics-Based Model Reduction in the Age of Digital Twins

Digital twins rely on the two-way interaction between a physical system and its digital counterpart. Data from the physical system is assimilated into the digital model, and predictions from the digital twin inform decisions and actions onto the physical system. Fundamental to this bidirectional interaction is the ability to solve the forward model rapidly and accurately. However, the solution of forward problems to a desired accuracy often comes with great computational cost, particulary for systems governed by partial differential equations. One way to address this problem is through projection-based, parametric model order reduction (MOR). MOR has been shown to provide physics-based, accurate yet inexpensive predictions for a large class of problems. In this talk, we will begin with a brief introduction to projection-based MOR. We will then delve into some recent developments in projection-based model order reduction for parametrized partial differential equations, highlighting the limitations, challenges, and opportunities in their use for digital twins.

Karen Veroy-Grepl
TU/e Eindhoven University of Technology
k.p.veroy@tue.nl

IP5

What Happens to a Dream Deferred? Chasing Language-Based Parallel Programming for HPC and AI

In 1951, Harlem Renaissance poet Langston Hughes asked this talk's titular question at the outset of a poem entitled "Harlem." Six years later, IBM mathematician John Backus developed Fortran, the world's first widely used high-level programming language. Backus later explored functional programming and highlighted the functional style in his Turing Award lecture in 1977, a year that also demarcates what one might consider the end of the classical era of Fortran. Building on a vision the presenter first conceived around the turn of the 21st century while teaching in Harlem, this talk will demonstrate how Fortran 2023 can finally deliver on Backus's functional programming dream in traditional high-performance computing (HPC) domains such as partial differential equation (PDE) solvers and in emerging domains such as artificial intelligence (AI). For PDE solvers, the talk will describe language facilities for asynchronously evaluating expressions that apply discrete, parallel, purely-functional differential operators to software abstractions that model continuous mathematical abstractions. For AI, the talk will demonstrate that Fortran's native concurrent loop iterations can combine with side-effect-free, pure procedures to facilitate automatically par-

allelizing deep-learning inference and training algorithms on processors and accelerators. The talk will provide updates on an ongoing effort by Berkeley Lab's Fortran team to realize this dream by through our work at multiple levels in the software stack, including applications, compiler runtime libraries, and networking middleware. Along the way, the talk will highlight ways in which programs promoting inclusivity in science facilitated significant aspects of the presented work.

Damian W. Rouson
Lawrence Berkeley National Laboratory, U.S.
rouson@lbl.gov

IP6

Computational Modelling of Coupled Thermo-Poro-Elastic Deformation of Fractured Rocks in the Context of the Energy Transition

The processes governing fracture growth in complex media and their interactions with smaller-scale and larger-scale discontinuities and material variations are crucial to various natural and engineered systems. This talk presents numerical modeling results of the deformation in three-dimensional fractured rock masses, which feature non-planar, geomechanically-grown fractures at multiple scales with spatially-varying apertures and fluid flow through both the matrix and fractures. The simulations are based on the solution of fully coupled finite element-based discretization of the poro-thermo-elasticity equations. Fractures are represented numerically as lower-dimensional manifolds, resulting in a mixed-dimensional system that capture the coupled hydromechanical deformation of both the matrix and fractures. Fracture geometry is modeled using NURBS surfaces, discretized with triangles or quadrilaterals, while the matrix is discretized with isoparametric tetrahedra and/or hexahedra. Friction along contacting fracture surfaces is resolved using an Augmented Lagrangian approach with the Uzawa iteration method. Matrix heterogeneities are incorporated as material property variations within the mesh, and the fracture geometry discretization adapts as the fracture grows. Examples include applications to nuclear waste disposal, induced seismicity quantification during cyclic hydrogen storage, and thermal shock fracturing. The complex three-dimensional fracture geometries generated by geomechanical processes in our simulations resemble field-observed and experimentally observed patterns. Our results indicate that fracture interactions rapidly lead to growth across scales, driven by intersection and coalescence. Fractures hydromechanically interact during growth, functioning as larger features before intersection occurs.

Adriana Paluszny
Imperial College
apaluszn@imperial.ac.uk

IP7

Artificial Intelligence: Applications in Scientific and Domain-Rich Fields

Everyone has had the opportunity to interact with AI through chat interfaces, but how can it be applied to accelerate scientific innovation and automate domain-rich processes? This talk will cover some of the common concepts in modern AI applications, including retrieval augmented generation, fine tuning, synthetic data generation, and chain of thought. It will then explore potential use cases in materials science, innovation mapping, and building design that show how to go beyond simple chat interfaces to a deeper collaboration between domain experts and AI.

Julia Ling
Google X
julialing11@gmail.com

IP8

Randomized Linear Algebra in Scientific Computing

Randomized algorithms are becoming increasingly popular in matrix computations. Recent software efforts, such as RandLAPACK, demonstrate that randomization is on the verge of replacing existing deterministic techniques for several large-scale linear algebra tasks in scientific computing. The poster child of these developments, randomized SVD, is now one of the state-of-the-art approaches for performing low-rank approximation. In this talk, we will go beyond the randomized SVD and illustrate the great potential of randomization to not only speed up existing algorithms, but to also yield novel and often simple algorithms for solving notoriously difficult problems. Examples covered in this talk include reduced order modeling, acceleration of Krylov subspace methods, joint diagonalization, large null space computation, and spectral density estimation. A common theme of these developments is that randomization helps to transform linear algebra results that only hold generically into robust and reliable numerical algorithms.

Daniel Kressner
EPFL, Switzerland
daniel.kressner@epfl.ch

SP1

SIAG/CSE Best Paper Prize Lecture: Tensor-Tensor Algebra for Optimal Representation and Compression of Multiway Data

To Come

Misha E. Kilmer
Mathematics Department
Tufts University
misha.kilmer@tufts.edu

SP2

SIAG/CSE Early Career Prize Lecture: Solving High-Dimensional Partial Differential Equations using Deep Learning: Original Insights and Recent Progress

For a long time, solving high-dimensional partial differential equations (PDEs) has posed significant computational challenges, primarily due to the notorious "curse of dimensionality." Since the introduction of the Deep BSDE (backward stochastic differential equation) method in 2017, deep learning-based algorithms have made substantial progress in overcoming this obstacle, providing promising solutions for a variety of applications, including physics, engineering, machine learning, economics, and finance. This talk will review the original insights behind the Deep BSDE method, which connects the nonlinear Feynman-Kac formula involving SDEs with stochastic optimization in deep learning. In recent years, neural networks have become a routine tool for approximating functions in PDEs. Nevertheless, the original insight into linking stochastic meth-

ods with PDE solutions remains instrumental, as it has inspired various methods that adopt flexible loss functions and tailored sampling strategies. These aspects are often key to the success of deep learning for high-dimensional PDEs. The talk will also introduce recent developments in this direction, such as the Deep Picard Iteration method, which reformulates neural network training objectives into regression tasks involving function values and gradients, improving both accuracy and scalability.

Jiequn Han
Flatiron Institute
jhan@flatironinstitute.org

SP3

Ivo Renata Babuka Prize Lecture

To Come

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

SP4

James H. Wilkinson Prize in Numerical Analysis and Scientific Computing: Mixed Precision Numerical Linear Algebra

Support for arithmetic in multiple precisions and number formats is becoming increasingly common in emerging architectures. Mixed precision capabilities are already included in many machines on the TOP500 list and will be a crucial hardware feature going forward. From a computational scientist's perspective, our goal is to determine how and where we can safely exploit mixed precision computation in our codes to improve performance. This requires both an understanding of performance characteristics as well as a rigorous understanding of the theoretical behavior of algorithms in finite precision arithmetic. We discuss the challenges of designing mixed precision algorithms and give three cases where low precision can often safely be used to improve performance. One such case, common in computational science, is when there are already other significant sources of "inexactness" present, e.g., discretization error, measurement error, or algorithmic approximation error. In this instance, analyzing the interaction of these different sources of inexactness can give insight into how the finite precision number formats should be chosen in order to "balance" the errors, potentially improving performance without a noticable decrease in accuracy. We present a few recent examples of this approach, which demonstrate the potential for the use of mixed precision in numerical linear algebra.

Erin C. Carson
Charles University MFF
carson@karlin.mff.cuni.cz

SP5

SIAM/ACM Prize in Computational Science Engineering Lecture

To Come

Scott Hansen
Iowa State University
shansen@iastate.edu

MS1

Nonlinear Feedback Control of the Fluidic Pinball

We develop and present a range of feedback control laws for the fluidic pinball control problem. This control problem seeks to control the vortex shedding behind three cylinders where cylinder rotation is the actuation mechanism. This benchmark problem has been used to demonstrate several machine learning control strategies. In this talk, we present an approach that uses interpolatory model reduction to build a polynomial approximation to the perturbation of the Navier-Stokes flow from the steady-state solution that maps the three control inputs to twenty-four output measurements taken downstream of the cylinders. Using this model, we use polynomial approximations to Hamilton-Jacobi-Bellman equations to create a quadratic feedback control law. Numerical simulations of this feedback law (a closed-loop simulation performed using FEniCS) demonstrate that we can completely stabilize the steady-state solution (i.e. no vortex shedding) over a range of low Reynolds number flows. We will comment on the sensitivity of the controller to boundary conditions.

Ali Bouland, Jeff Borggaard
Virginia Tech
bouland@vt.edu, jborggaard@vt.edu

MS1

Nonlinear Feedback Control Design Using Polytopic Autoencoders in Approximative LPV Embeddings

Polytopic autoencoders ensure that latent variables and reconstructed states reside within a polytope, which allows their latent variables to be directly used as the time-varying parameters in polytopic linear parameter-varying (LPV) systems. In this work, we design nonlinear feedback control systems using polytopic autoencoders, approximate LPV embeddings, and second-order series expansions of the solution to the state-dependent Riccati equation. We explain each part of the design and demonstrate that this approach generalizes and outperforms the standard linear-quadratic design in a numerical case study of fluid flows.

Yongho Kim
MPI for Dynamics of Complex Technical Systems
ykim@mpi-magdeburg.mpg.de

Jan Heiland
TU Berlin, Germany
heiland@mpi-magdeburg.mpg.de

Steffen W. R. Werner
Virginia Tech
steffen.werner@vt.edu

MS1

Data-Driven Control with Reduced-Order Modeling

High-dimensional, complex systems are often difficult to model and sometimes necessitate convoluted codes, posing a challenge for model-based control. To tackle the challenges of controlling these complex systems, techniques in data-driven control allow engineers to bypass modeling and design control directly through data from the system. One

such technique is semidefinite programming (SDP), which allows for the design of a convex optimization problem using system data to yield a stabilizing feedback controller. Although SDPs can be solved in polynomial time, they often exhibit numerical issues impeding their use in many practical applications for complex problems. To address this issue, we present a discussion of how model reduction can be used to reduce the problem size without sacrificing stability guarantees.

Steven C. Nguyen
UC San Diego
stn030@ucsd.edu

MS1

Data Efficient Low-Dimensional Controller Inference Via Adaptive Sampling

Stabilizing dynamical systems in science and engineering poses a significant challenge, especially in scenarios where only limited data are available. The quality and quantity of information within the data related to the task at hand are crucial factors to consider. This is particularly true when dealing with unstable dynamics, as the presence of instabilities typically leads to redundant or destructive system behavior, making it difficult to collect large amounts of informative data. In our new approach, we address the issue of generating small yet informative data sets for the stabilization of dynamical systems. The key lies in the adaptive construction of suitable input signals for the data generation via intermediate low-dimensional controllers that stabilize the system dynamics over limited subspaces. Numerical experiments with chemical reactors and fluid dynamics behind obstacles demonstrate that the proposed approach stabilizes systems reliably after observing about ten data samples even though the dimension of states is orders of magnitude higher.

Steffen W. R. Werner
Virginia Tech
steffen.werner@vt.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS1

Data-Driven Frameworks for Global Storm Surge Prediction

Coastal hazards pose significant societal risks, particularly storm surges, which are abnormal rises in sea levels caused by hurricanes and typhoons. While high-fidelity models like the ADvanced CIRCulation (ADCIRC) model predict storm surges accurately, they are computationally expensive. To address this, we are developing data-driven surrogate models for storm surge prediction. Initially, our models consisted of simple feed-forward networks trained on a dataset of 446 synthetic hurricanes in the Texas coast region, and predicted storm surge levels to within 30 centimeters. We have now generated a global dataset with over 48,000 synthetic hurricanes, including more than 13,000 in the North Atlantic basin. This expanded dataset enables our models to predict storm surges worldwide. Preliminary results with the expanded dataset for the North Atlantic region indicate a reduced error of 10 centimeters. The training framework has also been upgraded from single-GPU to a distributed approach. This shift has the potential to provide a unified model applicable to coastal regions globally, particularly those lacking robust storm surge prediction systems. Our current focus is on developing a multi-node operator learning model on unstructured mesh grids for this dataset. We will compare the global and regional predictive accuracy, as well as the computational cost of training and evaluating the operator learning framework, the current feed-forward approach, and the high-fidelity model.

Jinpai Zhao, Benjamin Pachev
Oden Institute - The University of Texas at Austin
max.zhao@utexas.edu, benjaminpachev@utexas.edu

Prateek Arora
Civil and Urban Engineering, New York University
pa2178@nyu.edu

Carlos del-Castillo-Negrete
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
cdelcastillo21@gmail.com

Eirik Valseth
Oden Institute, UT Austin
eirik@utexas.edu

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

MS2

Modeling Heterogeneous Piezo1 Activity in Collective Keratinocyte Migration

During wound reepithelialization, the mechanically-activated ion channel PIEZO1 has been identified as an inhibitory factor in healing, regulating wound edge retraction and disrupting coordinated directionality in the collective migration of keratinocytes. PIEZO1 exhibits heterogeneous channel activity among cells during this process. However, the impact of cellular mixtures with varying levels of PIEZO1 activity on wound healing remains elusive. In this study, we developed an integrative two-dimensional mathematical model to investigate the wound closure dynamics of mixtures containing two distinct cell types. Each cell type is governed by its own set of partial differential equations and parameters, reflecting its specific PIEZO1 activity level and interacting through mechanisms including cell-cell adhesion, volume-filling effects, and wound edge retraction. Simulations with various cell mixtures reveal that cells with higher PIEZO1 activity are generally underrepresented at the wound edge and are associated with wound edge retraction. Moreover, the model shows that mixing mutually repulsive cells enhances wound closure more effectively than homogeneous populations, with this effect being amplified by the heterogeneity of the mixture.

Jinghao Chen
Institute for the Advanced Study of Human Biology
Kyoto University
chen.jinghao.3n@kyoto-u.ac.jp

John Lowengrub
Department of Mathematics
University of California, Irvine

lowengrb@math.uci.edu

Medha Pathak
Department of Physiology and Biophysics
University of California, Irvine
medhap@uci.edu

MS2

Simulation and Control of Inward Solidification in Cryobiology

Mathematical models that predict a cells response to encroaching ice have, for many years, played an important role in developing cryopreservation protocols. It is clear that information about the cellular state as a function of cooling rate can be used to improve cryopreservation protocols and explain reasons for cell damage during freezing. However, previous work has ignored the interaction between the important solutes, the effects on the state of the cell being frozen and encroaching ice fronts. In this talk, I will introduce and summarize our work on this problem and examine the cryobiologically relevant setting of a spherically symmetric model of a biological cell separated by a ternary fluid mixture from an encroaching solidliquid interface. I will illustrate our work on a simplified 1-D problem and demonstrate how the thermal and chemical states inside the cell is influenced and can potentially be controlled by altering cooling protocols at the external boundary.

Dan Anderson
GMU
danders1@gmu.edu

James Benson
University of Saskatchewan
james.benson@usask.ca

Anthony J. Kearsley
Applied and Computational Mathematics Division
National Institute for Standards and Technology
anthony.kearsley@nist.gov

Jessica Masterson
George Mason University
jmaster6@gmu.edu

MS2

Optimal Bandwidth Selection in Biosensor Field-Effect Transistor Measurements

The use of stochastic regression to separate signal from noise produced by Bio-FETs will be discussed in this talk. The noise realized by BioFETs interferes with quantitative and qualitative analysis, thus determining optimal bandwidth associated with experimental Bio-FET data measurements is an important task. Presented results suggest consistent across aspect rations and a choice of stochastic regression kernel function and yield what appear to be good results.

Luis Melara
Shippensburg University
lamelara@ship.edu

MS2

Incorporating Tumor Size in Multistage Clonal Expansion Model of Cancer Incidence

Recent studies report a rise in early-onset cancer cases (diagnosed under 50), raising questions about whether this trend reflects apparent effects, like over-diagnosis, or true increases in risk. Traditional Multi-Stage Clonal Expansion (MSCE) models assume cancer is detected at first malignancy. SEER data show a decreasing trend in tumor size at diagnosis, suggesting an apparent effect. To consider this trend, we extend the MSCE model by incorporating tumor size at diagnosis to account for detection advancements and distinguish between apparent and true effects on cancer incidence. We developed a model with a time-dependent probability function based on a birth process for malignant cells, resolving non-identifiability issues in the classic model and improving parameter estimation. This approach was applied to different cancer types allowing for an evaluation of cancer risk changes while considering detection improvements. Findings indicate that recent cohorts show accelerated carcinogenesis in early and mid-life, while the traditional model is less sensitive to these changes. Our analysis also highlights distinct screening impacts: colorectal cancer screening reduces incidence in those over 50 by removing precancerous polyps, while breast cancer screening increases early detection in women over 40. These results emphasize the importance of accounting for tumor size at diagnosis in cancer modeling, supporting a true increase in early-onset cancer risk in recent years.

Navid Mohammad Mirzaei
Columbia University
nm3519@cumc.columbia.edu

Wan Yang
Columbia University
Department of Epidemiology
wy2202@cumc.columbia.edu

MS2

Optimization Strategies for Immune Checkpoint Blockade in Cancer Treatment

Many cancers progress despite our best treatment efforts. One way cancer cells escape treatment, leading to progressive disease, is by leveraging immune checkpoints. Cancer cells preferentially bind to inhibitory receptors on the immune cell surface, turning off any immune response that would otherwise attack them. Immune checkpoint blockade is designed to work through competitive hindrance of inhibitory checkpoints. Namely, therapeutic antibodies prevent receptors on the cancer cell surface from interacting with receptors on the immune cell surface. However, there is room for improvement in the clinical success rates of standard-of-care treatments utilizing this strategy. Here we frame improving immunotherapy as a mathematical optimization problem. We discuss a mathematical model for immune checkpoint inhibition and immune cell exhaustion, i.e., overstimulation by cancer cells. We first take the approach of a combined system leveraging ordinary differential equations and agent-based modeling to draw conclusions regarding the blockade efficiency necessary to shift a patients most likely treatment outcome from progressive disease to stable disease or remission. We then discuss simulated testing of dosing strategies that allow us to achieve this desired blockade efficiency. This combined model system illustrates the design and utility of multiple modeling strategies for investigating anticipated cellular behaviors

and treatment strategies in the tumor microenvironment.

Anne M. Talkington
University of Buffalo
annetalk@buffalo.edu

Anthony J. Kearsley
Applied and Computational Mathematics Division
National Institute for Standards and Technology
anthony.kearsley@nist.gov

MS3

Efficient Physics-Based Anderson Acceleration for the Robust Solution of Multiphysics Problems

The scope of this work is twofold. On one hand, we provide an efficient implementation of the Alternating Anderson-Richardson method with focus on performance, robustness, and memory reuse. We do this through a specialized implementation of the QR factorization. On the other hand, we propose a physics-based acceleration strategy that reduces numerical artifacts arising from solving multiple physics monolithically, and also yields a smaller acceleration step. This methodology is rigorously analyzed, and we provide simply computable error estimates, which we studied numerically. Surprisingly, these theoretical results are consistent with other estimates arising in randomized sketching algorithms. The overall solution strategy results in a very efficient solver, which we compare against other state-of-the-art solvers in several numerical tests, both linear and nonlinear, to validate our claims.

Nicolas A. Barnafi
Universidad de Chile
nicolas.barnafi@uc.cl

Massimiliano Lupo Pasini
Computational Sciences and Engineering Division
Oak Ridge National Laboratory
lupopasinim@ornl.gov

MS3

Convergence Analysis of the Alternating Anderson-Picard Method for Nonlinear Fixed-Point Problems

Anderson Acceleration (AA) has been widely used to solve nonlinear fixed-point problems due to its rapid convergence. This work focuses on a variant of AA in which multiple Picard iterations are performed between each AA step, referred to as the Alternating Anderson-Picard (AAP) method. Despite introducing 'slow' Picard iterations, this method has been shown to be efficient and robust in both linear and nonlinear cases. However, there is a lack of theoretical analysis for AAP in the nonlinear case, which this paper aims to address. We show the equivalence between AAP and a multisecant-GMRES method that uses GMRES to solve a multisecant linear system at each iteration. More interestingly, the incorporation of Picard iterations and AA establishes a deep connection between AAP and the Newton-GMRES method. This connection is evident in terms of the multisecant matrix, the approximate Jacobian inverse, search direction, and optimization gain – an essential factor in the convergence analysis of AA. We show that these terms converge to their corresponding terms in the Newton-GMRES method as the residual approaches zero. Consequently, we build the convergence analysis of AAP. To validate our theoretical findings, numerical examples are provided.

Xue Feng
UC Davis
xffeng@ucdavis.edu

Paul Laiu
Oak Ridge National Laboratory
laiump@ornl.gov

Thomas Strohmer
University of California,Davis
Department of Mathematics
strohmer@math.ucdavis.edu

MS3

Improved Convergence Rates of Windowed Anderson Acceleration for Symmetric Fixed-Point Iterations

Our presentation discusses the commonly utilized windowed Anderson acceleration (AA) algorithm for fixed-point methods, $x^{(k+1)} = q(x^{(k)})$, and showcases the first proof that when the operator q is linear and symmetric the windowed AA algorithm, which uses a sliding window of prior iterates, improves the root-linear convergence factor over the fixed-point iterations. When q is nonlinear, yet has a symmetric Jacobian at a fixed point, a slightly modified AA algorithm is proved to have an analogous root-linear convergence factor improvement over the fixed-point method. Simulations verify our observations. Furthermore, experiments with different data models demonstrate AA is significantly superior to the standard fixed-point methods for Tylers M-estimation.

Casey Garner
University of Minnesota
garne214@umn.edu

Gilad Lerman
University of Minnesota, Minneapolis
School of Mathematics
lerman@umn.edu

Teng Zhang
University of Central Florida
teng.zhang@ucf.edu

MS3

Leveraging Anderson Acceleration to Improve Iterative Solvers for Navier-Stokes Equations

In this talk, we investigate enhancements of fixed-point type iterative methods, focusing on the iterated penalty Picard (IPP) and ArrowHurwicz (AH) iterations, for solving the incompressible steady Navier-Stokes equations by incorporating Anderson acceleration (AA). We examine these methods in terms of regularity properties of the methods which allow the AA theory in acceleration to be applied. We show AA enables the fast convergence of the IPP with O(1) penalty parameter, even in large scale NSE problems where the penalty parameter usually needed to be chosen very small, causing difficulties with linear solvers. Similarly, we study AH iteration, which is originally designed for general saddle point linear systems and extended to NavierStokes iterations in the 1970s, and employ recently developed divergence-free finite element methods which allows us to connect AH to IPP alongside Ander-

son acceleration to improve AH convergence. Analytical and numerical results demonstrate that this combination yields an efficient and effective solver that is competitive with more commonly used approaches.

Duygu Vargun
Oak Ridge National Laboratory
vargund@ornl.gov

MS4

Energy Conserving, Higher Order Time Discretization of Maxwell's Equations

In this work, our primary contribution is a semi explicit higher order time discretization of Maxwell's equations based only on elementary Taylor series expansion. We formulate Maxwells equations as a system for (p, E, H) where p is a fictitious electrical pressure variable that helps impose discretely exactly the divergence of the electric field intensity E, and H is the magnetic flux density. Our method labelled LF_R is an extension of the classical leapfrog scheme to an arbitrary (even) order R time accurate method. Consequently, we use two staggered time grids, and locate (p, E) and H alternately on them. Under appropriate initial and boundary conditions, we show that our discrete energy is conserved and the scheme is stable. We refer to LF_R as being semi explicit in the sense that p and E are obtained together at a given time step while H is computed at the next on the other grid explicitly. We use LF_R in conjunction with finite element exterior calculus for a full discretization of Maxwell's equations. We then provide a variation of our time discretization that is fully time implicit, and characterize its stability and energy conservation. For our LF_R schemes, we also provide proofs of appropriate notions of error convergence for the computed solutions. Finally, we validate our methods through computations for model problems in two and three dimensions, and we do so for (p, E, H) as well as (E, H) systems of Maxwell's equations.

Archana Arya
Indraprastha Institute of Information Technology Delhi, New Delhi, India
archanaa@iiitd.ac.in

MS4

Shape and Eigenvalue Optimization of Microstructures Governed by Maxwell's Equations

This talk is concerned with a class of shape optimization problems involving optical metamaterial comprised of periodic nanoscale inclusions. We will first summarize the underlying microscale model, the corresponding homogenization theory, and the eigenvalue representation that will serve as a basis for the shape optimization problems. We will then introduce a deformation field on the cell problem and the eigenvalue problem. Finally, we will formulate and solve the shape optimization problems using an adjoint approach.

Mansi Bezbaruah
Texas A&M university
bezba004@tamu.edu

MS4

Computational Inverse and Optimal Design for Topological Photonic Crystals

This work is concerned with inverse design of the grating metasurface over hyperbolic metamaterials (HMMs) in order to enhance spontaneous emission (SE). We formulate the design problem as a PDE-constrained optimization problem and employ the gradient descent method to solve the underlying optimization problem. The adjoint-state method is applied to compute the gradient of the objective function efficiently. Computational results show that the SE efficiency of the optical structure with the optimized metasurface increases by 600% in the near field compared to the bare HMM layer. In particular, an optimized double-slot metasurface obtained by this design method enhances the SE intensity by a factor of over 100 in the observation region.

Junshan Lin
Auburn University
jzl0097@auburn.edu

MS4

A High-Order Spectral Algorithm for the Numerical Simulation of Layered Media with Uniaxial Hyperbolic Materials

The scattering of electromagnetic waves by threedimensional periodic structures is important for many problems of crucial scientific and engineering interest. Due to the complexity and threedimensional nature of these waves, fast, accurate, and reliable numerical simulation of these are indispensable for engineers and scientists alike. For this, HighOrder Spectral methods are frequently employed and here we describe an algorithm in this class. Our approach is perturbative in nature where we view the deviation of the permittivity from a constant value as the deformation and we pursue regular perturbation theory. This work extends our previous contribution regarding the Helmholtz equation to the full vector Maxwell equations, by providing a rigorous analyticity theory, both in deformation size and spatial variable (provided that the permittivity is, itself, analytic)

David P. Nicholls
University of Illinois at Chicago
davidn@uic.edu

MS4

A High-Order Perturbation of Envelopes (hope) Method for Vector Electromagnetic Scattering by Periodic Inhomogeneous Media

The scattering of electromagnetic waves by threedimensional periodic structures is important for many problems of crucial scientific and engineering interest. Due to the complexity and three-dimensional nature of these waves, fast, accurate, and reliable numerical simulation of these are indispensable for engineers and scientists alike. For this, High-Order Spectral methods are frequently employed and in this talk, I present the HOPE method which is in this class. Our approach is perturbative in nature where we view the deviation of the permittivity from a constant value as the deformation and we pursue regular perturbation theory. More specifically, we expand the three-dimensional, vector-valued electric field in a Taylor series in this small deformation parameter, derive recursions that each term in this series must satisfy, invoke a novel elliptic theory to establish bounds on the size of each correction, and thereby show that the purported Taylor series does, in fact, converge. Beyond this, we show that each of these terms in the Taylor series is jointly analytic in all three spatial variables by estimating solutions of gov-

erning equations for derivatives of these terms. Numerical experiments are also provided to validate the theoretical results.

Liet Vo
University of Illinois Chicago
lietvo@uic.edu

MS5

Discontinuous Galerkin Solutions of Moving Boundary Problems Using Arbitrary Lagrangian-Eulerian Method on GPUs

Flow problems with moving boundaries and interfaces appear naturally in many applications, such as turbine flow, flapping flight, etc. Numerical modeling of these flows poses additional challenges since the numerical grid moves with time. The arbitrary Lagrangian-Eulerian method (ALE) is one of the solution strategies for numerically dealing with moving boundary problems. In this talk, we present an arbitrary Lagrangian-Eulerian method on GPUs for moving boundary problems and its implementation to an open-source discontinuous Galerkin library, libParanumal.

Atakan Aygun, Ali Karakus
Middle East Technical University
atakana@metu.edu.tr, akarakus@metu.edu.tr

MS5

Introduction to Compressible Computational Fluctuating Hydrodynamics

In this talk we consider fluctuating hydrodynamics for multicomponent compressible ideal gas mixtures. Fluctuating hydrodynamics augments the compressible Navier-Stokes equations with stochastic flux terms that represent thermal fluctuations. We first review the deterministic formulation then discuss the construction of the stochastic fluxes. We then examine the equilibrium covariance of the resulting system of stochastic partial differential equations. Next we discuss some of the issues that need to be considered in developing numerical methods for the system and introduce a staggered-grid discretization based on those considerations. We present numerical results that demonstrate that the resulting numerical methodology accurately captures the statistics of fluctuations at equilibrium. Finally, we illustrate some of the phenomena induced by fluctuations in systems that are out of equilibrium.

John B. Bell
Lawrence Berkeley National Laboratory, U.S.
jbbell@lbl.gov

MS5

An Introduction to Computational Fluctuating Hydrodynamics

The intrinsic thermal motion of the molecules in a fluid result in microscopic fluctuations that are well-understood at thermodynamic equilibrium. For a fluid in a nonequilibrium state, such as in the presence of gradients of temperature or concentration, these hydrodynamic fluctuations are qualitatively different, potentially becoming macroscopic in amplitude and length scale. Thermal fluctuations can be modeled accurately using a modification of the Navier-Stokes conservation equations that includes a stochastic forcing, as was originally proposed by Landau and Lifshitz. In this fluctuating hydrodynamics (FHD) formulation, a stochastic flux is added to each dissipative flux associated with the transport of species mass, momentum and energy densities in a manner that satisfies the fluctuation-dissipation balance. This talk will briefly review the background for the fluctuating hydrodynamic equations. The basic theory and numerical methods will be illustrated using the one-dimensional stochastic heat equation.

Alejandro L. Garcia
San Jose State University
Alejandro.Garcia@sjsu.edu

MS5

Analyzing a Mesoscopic Hydrodynamic Simulation Method Via Its Structure Factors and Correlation Functions

Unlike macroscopic computational fluid dynamic simulation methods, mesoscopic hydrodynamic simulation methods have a stochastic nature due to thermal fluctuations. As a result, the systematic analysis of the latter methods requires different mathematical tools. In this talk, I introduce the structure factors and correlation functions as physical quantities that capture the spatio-temporal fluctuational behaviors of a mesoscopic system and discuss stochastic analysis techniques based on these quantities. First, I consider the continuum case and demonstrate how to analyze the time evolution equations given in the form of stochastic partial differential equations (SPDEs). For example, by linearizing these SPDEs, I show that the correct thermodynamic equilibrium can be attained as predicted by equilibrium statistical mechanics. Then, I consider the case where the governing SPDEs are spatially discretized to form a large set of ordinary differential equations and discuss how to analyze the resulting numerical scheme. For example, I demonstrate how to analyze the numerical settings to impose boundary conditions. Finally, I discuss how to construct the optimal time-integrator based on the structure factor analysis.

Changho Kim
University of California, Merced
Applied Mathematics
ckim103@ucmerced.edu

MS5

Introduction to Low Mach Number Computational Fluctuating Hydrodynamics

Fluctuating hydrodynamics (FHD) is an emerging technique for highly efficient simulation of mesoscopic fluids, where thermal fluctuations drive microscopic dynamics that can have profound impacts on macroscopic behavior. This formulation, originally proposed by Landau and Lifshitz, requires the augmentation of each dissipative flux in a deterministic continuum fluid model with a stochastic flux. Some recent advances in FHD involve the development and implementation of low Mach number models that exploit the separation of scales between acoustic and bulk convective time scales, offering potentially orders of magnitude computational speedup over compressible FHD models, particularly for liquids. The primary challenge in developing such models is the imposition of a low Mach number velocity constraint within a spatiotemporal framework that satisfies fluctuation-dissipation balance. In this talk I give an overview of the low Mach number FHD formulation and numerical implementation. I will briefly describe applications in multispecies diffusive mixing, chemically reacting flow, electrolytes, and multiphase flow that

effectively leverage this formulation.

Andrew J. Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

MS6

CUQIpy: Computational Uncertainty Quantification for Inverse Problems in Python

In this talk we present CUQIpy (pronounced cookie pie) - a new computational modelling environment in Python that uses uncertainty quantification (UQ) to access and quantify the uncertainties in solutions to inverse problems. The overall goal of the software package is to allow both expert and non-expert (without deep knowledge of statistics and UQ) users to perform UQ related analysis of their inverse problem while focusing on the modelling aspects. To achieve this goal the package utilizes state-of-the-art tools and methods in statistics and scientific computing specifically tuned to the ill-posed and often large-scale nature of inverse problems to make UQ feasible. We showcase the software on problems relevant to imaging science such as computed tomography and partial differential equation-based inverse problems. CUQIpy is developed as part of the CUQI project at the Technical University of Denmark and is available at https://cuqi-dtu.github.io/CUQIpy .

Jakob S. Jorgensen
Technical University of Denmark
jakj@dtu.dk

Amal Alghamdi
Department of Applied Mathematics and Computer Science
Technical University of Denmark
amaal@dtu.dk

MS6

Uncertainty Quantification for the Koopman Operator

Modeling and analyzing dynamical systems can be challenging, especially when the evolution is nonlinear. One approach is through the Koopman op- erator, which models the evolution of an infinite dimensional functional space instead of the phase space itself. This turns the problem from a nonlinear one to a linear, albeit infinite dimensional, one. Linearity implies that tools from spectral analysis are available to analyze the underlying dynamical system. A popular method to approximate the Koopman operator numerically is Extended Dynamic Mode Decomposition (EDMD). This method requires to select a set of functions as a truncated function space basis, and then represents the operator as a matrix in this basis. In kernel-EDMD, kernel functions are used to implicitly define this basis set. In this contribution, we consider learning these kernels from snapshots of states of a dynamical system. This leads to a flexible and more accurate approximation of the Koopman operator. After the optimal kernel is constructed, we then interpret it as the covariance of a Gaussian process over the state space. This allows us to perform uncertainty quantification (UQ) on the data from the original system, and more generally to understand the Koopman operator and the underlying dynamical system through a UQ lense.

Erik Bolager
Technical University of Munich
erik.bolager@tum.de

MS6

UM-Bridge: User-Friendly, Scalable UQ

Treating uncertainties in models is essential in many fields of science and engineering. When dealing with complex and computationally costly numerical models this necessitates a combination of efficient model solvers, advanced UQ methods and HPC-scale resources. The resulting technical complexities and software engineering challenges are holding back many interesting UQ applications. I will introduce UM-Bridge, a flexible, language-agnostic software interface designed to integrate UQ methods with numerical models. UM-Bridge can be understood purely as an interface which resolves issues with linking software, it does not itself implement any UQ algorithms. In this talk, I will introduce UM-Bridge itself and show how to use it to integrate UQ methodologies into complex numerical models. Then, I will present realistic applications from different areas, which scale from laptops to all the way to large clusters. These include modeling the propagation of the 2011 Tohoku tsunami using shallow water equations. The primary objective is to derive precise initial displacement parameters from data collected by two buoys near the Japanese coast. I introduce a fully parallelized MLMCMC method implemented in MUQ and connected to the model code via UM-Bridge.

Anne Reinarz
Durham University
anne.k.reinarz@durham.ac.uk

Linus Seelinger
Karlsruhe Institute of Technology (KIT)
Scientific Computing Center
linus.seelinger@kit.edu

MS7

An Inexact Trust-Region Algorithm for Nonsmooth Risk-Averse Optimization

Many practical problems require the optimization of systems (e.g., differential equations) with uncertain inputs such as noisy problem data, unknown operating conditions, and unverifiable modeling assumptions. In this talk, we formulate these problems as infinite-dimensional, risk-averse stochastic programs for which we minimize a quantification of risk associated with the system performance. For many popular risk measures, the resulting risk-averse objective function is not differentiable, significantly complicating the numerical solution of the optimization problem. Unfortunately, traditional methods for nonsmooth optimization converge slowly (e.g., sublinearly) and consequently are often intractable for problems in which the objective function and any derivative information is expensive to evaluate. To address this challenge, we introduce a novel trust-region algorithm for solving large-scale nonsmooth risk-averse optimization problems. This algorithm is motivated by the primal-dual risk minimization algorithm and employs smooth approximate risk measures at each iteration. In addition, this algorithm permits and rigorously controls inexact objective function value and derivative (when available) computations, enabling the use of inexpensive approximations such as adaptive discretizations. We discuss convergence of the algorithm under mild assumptions and demonstrate its efficiency on various ex-

amples from PDE-constrained optimization.

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS7

Modeling Uncertainties in Large-scale Density-based Topology Optimization Problems

This work presents the utilization of Matrn random fields in thermal and solid mechanics topology optimization problems, with a focus on delivering practical realizations of realistic models for the imperfections and variations in external excitations and their effects on the optimized designs. The optimization objectives are the mean values of the thermal and mechanical compliances, which are subject to constraints on the amount of the distributed solid material. The objectives are approximated with the help of Monte Carlo sampling, where every sample requires the realization of a spatially varying random field throughout the design domain. Traditionally, these realizations rely heavily on series expansion techniques based on the classical KarhunenLove decomposition or convoluting white noise with a selected filter function. The series expansion approach requires a prohibitively large number of terms for Gaussian random fields with small correlation lengths. White noise filtering resolves the above issue by evaluating the convolution in Fourier space, which limits the application to simple rectangular domains and prohibits the precise control of the random fields around the boundaries. In contrast, our proposed approach involves the solution of fractional Stochastic Partial Differential Equations (SPDE), which is highly efficient and can generate random fields on complex domains with arbitrary boundary conditions.

Tobias Duswald
CERN, Geneva, Switzerland
tobias.duswald@cern.ch

Brendan Keith
Brown University
brendan_keith@brown.edu

Boyan Lazarov
Lawrence Livermore National Laboratory
lazarov2@llnl.gov

Socratis Petrides
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
petrides1@llnl.gov

B. Wohlmuth
Technical University of Munich, Germany
wohlmuth@ma.tum.de

MS7

Scenario Approximation for PDE-Constrained Optimization under Uncertainty

We study statistical guarantees for the scenario approximation method for PDE-constrained optimization with chance constraints. This sample-based technique replaces the original chance constraint with computationally tractable constraints derived from random samples. For example, when a chance constraint requires that a parameterized PDE state constraint be satisfied with high probability, the scenario approximation reformulates it into a standard PDE-constrained optimization problem, where the number of state constraints equals the number of samples. We derive estimates for the sample size needed to ensure, with high confidence, that feasible solutions to the original problem can be obtained through the scenario approximation. We then use these results to establish optimality guarantees for solutions to scenario-based PDE-constrained optimization problems. Our analysis is applicable to both linear and nonlinear PDEs with random inputs.

Johannes Milz
Georgia Institute of Technology
H. Milton Stewart School of Industrial & Systems Engineering
johannes.milz@isye.gatech.edu

MS8

Efficient Execution of Multiphysics Fem Assembly Using `Kokkos::Graph`

The computation of elemental system matrices and right-hand-side vectors and their assembly into sparse linear algebra data structures is a key component of FEM codes. Multiphysics simulations can involve multiple types of governing equations that might also change by subdomain (heterogeneous coefficients, different source terms, different types of boundary conditions,). Performing the assembly within a single kernel can thus be inefficient, both in terms of memory footprint and computational cost. Therefore, one might want to specialize the kernels, thus generating many kernels that need to be created and efficiently scheduled, potentially many times (e.g., iterative solver). These asynchronous kernels must also observe dependencies, thus leading naturally to a graph-based implementation. In this talk, we will present how we realised such an implementation in an in-house FEM code using `Kokkos::Graph`. In particular, we will discuss how we designed a polymorphic hierarchy of functors for performing the assembly on device and how we map these functors to nodes in the graph, while avoiding polymorphic calls on device. We will illustrate the proposed approach and evaluate the performance in the context of a computational electromagnetism simulation relevant to diffraction gratings.

Maarten Arnst, Romin Tomasetti
Universite de Liege
maarten.arnst@uliege.be, romin.tomasetti@uliege.be

MS8

Flecsi and Ristra: Task-Based Parallelism Designed for Heterogeneous Computing from the Start

The Flexible Computational Science Infrastructure (FleCSI) is an open-source C++ framework to assist in the development of performance portable multiphysics applications. It can leverage different task-based run-times, such as Legion or HPX, and make use of Kokkos to offload work to CPU threads and GPUs of all major vendors. It supports various types of discretizations through so-called specializations making application development easy. Ristra is a broad effort as part of the US Department of Energy's Advanced Simulation and Computing (ASC) program to develop modular and portable physics packages for simulations in a high-performance computing environment. A large part of its efforts are based on

FleCSI. We will discuss the lessons we learned running as one of the first users on the Grace-Hopper-based Venado system at Los Alamos National Laboratory and the MI300A-based systems precursor systems of El Capitan at Lawrence Livermore National Laboratory. In particular this will cover the benefits and challenges of the separation of concerns in our complex software stack and how we solved problems in interaction with HPC staff and the vendors.

Philipp V. Edelmann, Ben Bergen
Los Alamos National Laboratory
pedelmann@lanl.gov, bergen@lanl.gov

MS8

Effectively Leveraging Multi-Institution and Multi-Vendor Relationships to Address Advanced Software Needs

This talk will discuss challenges app teams at Sandia National Labs have faced when preparing for a new system, and how we approach triaging and resolving issues across the DOE complex as well as directly with vendors providing the system. We will discuss the infrastructure we use, as well as the general approach that has been effective in 1) identifying issues, 2) tracking and triaging issues amongst institutions, 3) understanding the issues and 4) coming to a resolution. Each step in this process is challenging. DOE's NNSA complex spans multiple labs (the "Tri-labs"), and each procurement entails multiple vendors. We will discuss what has worked and discuss existing challenges.

James Elliot
Sandia National Laboratories
jjellio@sandia.gov

MS8

Scaling El Capitan: Early Results from Doe Nnsas First Exascale Computer

El Capitan will be the DOE National Nuclear Security Administrations first exascale supercomputer. In collaboration with our vendor partners HPE and AMD through the El Capitan Center of Excellence (COE), we have spent several years preparing a variety of scientific applications such that they are scalable and performant on El Capitan on Day 1. In this talk, we will overview the application readiness efforts for El Capitan, discuss the modular software strategy employed by LLNL applications to ensure both performance and portability across modern architectures, and provide an early look at the achievements El Capitan will enable.

Judith Hill, Ramesh Pankajakshan
Lawrence Livermore National Laboratory
hill134@llnl.gov, pankajakshan1@llnl.gov

MS8

Understanding Grace+Hopper Specificities and Their Impact on the Development of CEA Codes

Adapting applications to a new supercomputer is always a challenge. In the June 2024 Top500 list, CEA unveiled its new supercomputer based on Nvidia Grace+Hopper superchips. The Grace+Hopper superchip offers new ways to handle memory data between the CPU and the GPU at the hardware level.To fully maximize the potential and performance of this type of chip, it is essential to understand the new mechanisms and their impact on applications. This presentation will detail our approach to analyzing Grace+Hopper behaviors, particularly focusing on memory management and data transfers. Additionally, we will explore their influence on our applications and strategies for adaptation.

Julien C. Jaeger
CEA
julien.jaeger@cea.fr

MS9

GPU Accelerated Linear Solvers for Implicit Time Integrators in the SUNDIALS Library

The implicit time integration methods in the SUNDIALS library require a nonlinear system to be solved at every time step. One of the most robust options for solving this nonlinear system is some form of Newton iteration, which, in turn, requires solving a linear system at every iteration. As such, throughout the integration, a linear solver will be used repeatedly to solve systems with at least the same general structure. The repeated solves can present challenges when using linear solver libraries that were not designed with this use case in mind. With GPUs in the mix, these challenges can be even harder to overcome efficiently without changes to the linear solver library. In this presentation, I will discuss experiences interfacing SUNDIALS with linear solver libraries and highlight ways these libraries can create interfaces that are easier to use and more efficient in the time integration context. LLNL-ABS-868851.

Cody J. Balos, David J. Gardner
Lawrence Livermore National Laboratory
balos1@llnl.gov, gardner48@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Steven B. Roberts, Carol S. Woodward
Lawrence Livermore National Laboratory
roberts115@llnl.gov, woodward6@llnl.gov

MS9

Rethinking Numerical Linear Algebra for Execution on GPUs in Re::Solve Library

Recent developments of computational technology, most notably general purpose graphical processing units (GPUs), have been strongly driven by the explosion of artificial intelligence applications (AI). Since most AI algorithms use dense linear algebra, the computational hardware has been optimized for such operations. This trend has created new computational landscape where traditional sparse linear algebra algorithms, which are essential for many scientific computing applications, do not perform well. These challenges go beyond implementation techniques and require rethinking mathematical algorithms. To address these challenges, we designed Re::Solve, a portable sparse linear algebra library, from the ground up. The librarys features include GPU-resident direct and iterative linear solvers, as well as randomized algorithms, among others. In this talk we describe the co-design of mathematical algorithms and their implementations in Re::Solve, which led to efficient performance on heterogeneous hardware. We provide performance analysis results and discuss differences when evaluating linear solver performance

in standalone examples as opposed within full application stack.

Slaven Peles
Oak Ridge National Laboratory
peless@ornl.gov

Katarzyna Swirydowicz
Pacific Northwest National Laboratory
kasia.swirydowicz@pnnl.gov

Maksudul Alam
ORNL
alamm@ornl.gov

Nicholson Koukpaizan
Oak Ridge National Laboratory
koukpaizannk@ornl.gov

MS9

Can Matrix Cores Help Accelerate Iterative Solvers?

In the last few years, AMD's matrix cores and Nvidia's tensor cores have been extensively used to accelerate dense linear algebra operations such as dense matrix-matrix multiplication. These operations are the fastest if low precision, such as FP16, is used. However, it remains an open question to what extent are matrix/tensor cores useful for sparse numerical linear algebra and in a mixed setting, where both sparse and dense operations are performed. In those cases, the (overall) performance is typically memory bound rather than compute bound, which begs a question whether using matrix/tensor cores can lead to any performance improvement. Moreover, in a practical setting, transferring the data (between lower and higher precision) or aligning the data to to access matrix/tensor cores efficiently are not free operations, so while performance gains are expected, there are trade-offs that are hardly ever considered. We explore practical use of these hardware-accelerated operations in a complex, realistic, primarily memory bound application, analyze all the incurred additional costs, profile its GPU performance and discuss where and when matrix/tensor cores can (and should) be used to improve performance.

Katarzyna Swirydowicz
Pacific Northwest National Laboratory
kasia.swirydowicz@pnnl.gov

MS9

Performance Advantages of Using Krylov-Based Iterative Methods With Sketching

The use of randomized methods in numerical linear algebra has gained popularity over the past several decades as matrix sizes continue to increase. These methods enhance the speed and reliability of algorithms for problems such as approximating solutions to linear systems and approximating the eigenpairs of matrices. One such technique, "sketching,' uses dimensionality reduction on least-squares problems to lower their computational cost with minimal loss of accuracy. The Rayleigh-Ritz method, which approximates eigenpairs of a matrix A from a basis V, can also be reformulated as a least-squares problem, enabling it to be used with sketching. One known challenge of Krylov-based iterative methods is that in floating point arithmetic, as the basis is being constructed, numerical error and repeated directions can result its condition number growing exponentially. Reorthogonalizing the basis can mitigate this issue, but doing this is expensive and may lead to a computational bottleneck, particularly with large bases. Sketching allows for the extraction of accurate solutions from a non-orthonormal basis, provided that its condition number is below $\epsilon_{\texttt{mach}}^{-1}$. This work explores the computational and performance trade-offs of using sketching with two popular Krylov methods, Lanczos and Generalized Davidson, leveraging the high-performance C99-based software library PRIMME.

Heather M. Switzer
College of William & Mary
Department of Computer Science
hswitzer2@washcoll.edu

Andreas Stathopoulos
William & Mary
Department of Computer Science
andreas@cs.wm.edu

MS10

Adaptively Regularised Nonlinear Ensemble Transport Filtering and Smoothing with P-Splines

Most contemporary data assimilation algorithms occupy opposing extreme ends on a line that trades computational efficiency (linear methods: e.g., EnKF) against statistical power (fully nonlinear methods: e.g., particle filter). A way to bridge this divide may be found in ensemble transport filters and smoothers, which leverage triangular measure transport to generalize the efficient-yet-simplistic EnKF/EnKS and thereby permit nonlinear updates of varying complexity. These transport methods construct a map from an unknown, potentially non-Gaussian target distribution - represented only through samples - to a simple reference distribution, often a standard multivariate Gaussian distribution. Inverting this map permits sampling from the targets conditional distributions, including the posterior. The parameterization of this triangular transport map is a flexible but critical choice for the performance of the resulting data assimilation algorithm. More complex maps may capture increasingly complex distributional features but risk unfavourable bias-variance trade-offs. In this presentation, we explore transport maps parameterized by P-splines. We leverage information theoretic metrics to optimize the smoothness of these functions, ensuring an ideal compromise between expressiveness and simplicity. We demonstrate the performance of the resulting algorithm in a nonlinear setting.

Max Ramgraber
Delft University of Technology
Department of Geoscience & Engineering
m.ramgraber@tudelft.nl

Berent Lunde
Equinor
berl@equinor.com

MS10

Long-time Accuracy of Ensemble Kalman Filters for Chaotic and Machine-learned Dynamical Systems

This talk will show theory for long-time accuracy of ensemble Kalman filters. We introduce conditions on the dynamics and the observations under which the estimation error remains small in the long-time horizon. Our theory

covers a wide class of partially-observed chaotic dynamical systems, which includes the Navier-Stokes equations and Lorenz models. In addition, we prove long-time accuracy of ensemble Kalman filters with surrogate dynamics, thus validating the use of machine-learned forecast models in ensemble data assimilation.

Daniel Sanz-Alonso, Nathan Waniorek
University of Chicago
sanzalonso@uchicago.edu, waniorek@uchicago.edu

MS10

Change of Measure for Bayesian Field Inversion with Hierarchical Hyperparameters Sampling

In this talk, we are interested in estimating a scalar field from noisy indirect observations. In order to obtain a full posterior distribution of the quantity of interest, the inverse problem is formulated in a Bayesian framework along with Markov chain Monte Carlo sampling. Inferring a field is expensive with regard to its infinite dimension and to the forward model computational cost. Our approach relies on a parametrization based on the Karhunen-Love decomposition. Although attractive since it provides a representation with a finite number of terms, the KL decomposition depends on the hyperparameters of its autocovariance function that are difficult to choose a priori due to lack of knowledge. Instead of selecting deterministic values, we propose to deal with hyperparameters prior distributions in order to test various field shapes during the sampling. For that purpose, we present a novel approach based on a change of measure of the decomposition coordinates that allows a smooth exploration of the hyperparameter space. Our sampling procedure is accelerated by means of polynomial chaos expansions to replace the data computed from the forward model. A seismic traveltime tomography case shows that exploring the hyperparameters space improves the uncertainty estimation compared to the approach with fixed hyperparameters.

Pierre Sochala
CEA (Atomic Energy and Alternative Energies Commission)
pierre.sochala@cea.fr

Nadège Polette
CEA, DAM, DIF, F-91297 Arpajon
nadege.polette@minesparis.psl.eu

Olivier Le Matre
CNRS/INRIA/CMAP, École Polytechnique, IPP, Route de Saclay
91128, Palaiseau, France
olivier.le-maitre@polytechnique.edu

Alexandrine Gesret
MINES Paris, PSL University, Centre de Géosciences
alexandrine.gesret@minesparis.psl.eu

MS10

Multiplicative Score-Based Generative Model for Fluid Dynamics

Accurately estimating the state of a physical system requires integrating numerical model predictions with noisy observational data. Learning these models is a challenging task as they must account for the physics of the state under study. In particular, numerical prediction of turbulent flows should preserve the structure and conserve energy. We propose a new score-based generative model for turbulent flows, using a skew-symmetric multiplicative noise for both noising and denoising steps. Inspired from physics and naturally conserving energy, this transport noise is the heart of the Kraichnan and Location Uncertainty models. The noising step transforms the distribution p0 of the input data into a rotation-invariant distribution in state space thanks to the skew-symmetric multiplicative noise and then into a Gaussian distribution. The denoising step reconstructs first the rotation invariant distribution from a Gaussian distribution and then the initial distribution p0 using a backward SDE and a score function learned by a neural network. This model aims to improve the accuracy of turbulent flow predictions by effectively integrating physical uncertainties.

Merveille Cyndie TALLA MAKOUGNE
INRAE
merveille.talla@inrae.fr

Valentin Resseguier
INRAE
SCALIAN DS
valentin.resseguier@inrae.fr

Robert Gruhlke
Freie Universitaet Berlin
Germany
r.gruhlke@fu-berlin.de

Dominique Heitz
INRAE
dominique.heitz@inrae.fr

Etienne Memin
Inria Rennes-Bretagne Atlantique
etienne.memin@inria.fr

MS11

Reduced-Basis Methods for Uncertainty Quantification with application to High-Speed Flows

The design and analysis of high-speed flight vehicles is reliant on computational fluid dynamics (CFD) due to the expense and difficulty of flight tests and experiments. Workhorse CFD models are subject to various sources of modeling uncertainties. Therefore, approaches that efficiently propagate and reduce uncertainties are important for a robust design cycle. Unfortunately, CFD is computationally expensive for high-speed flows due to high grid resolution requirements. Also, forward uncertainty propagation and model calibration are many-query problems requiring many model evaluations with a wide range of input parameters. This talk discusses reduced-basis accelerated uncertainty propagation and parameter calibration of high-speed flows solved with the ReynoldsAveraged NavierStokes equations. We discuss a greedy algorithm that leverages novel, easily computable, error estimates for constructing a reduced basis. We embed our reduced-order model with multi-fidelity uncertainty quantification and parameter-calibration workflows. We assess the utility of the reduced-basis method for more efficiently propagating parametric uncertainties in the SpalartAllmaras turbulence model and reducing model-form error. Results are presented on several high-speed turbulent flows. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525. SAND2024-11411A

Patrick J. Blonigan

Sandia National Laboratories
pblonig@sandia.gov

Eric Parish
Sandia National Lab
eric.parish9@gmail.com; ejparis@sandia.gov

MS11

An Evolve-Filter-Relax Stabilized Reduced Order Model for the Boussinesq Equations

The classical Galerkin ROM (G-ROM) generally yields spurious numerical oscillations leading to non-physical solutions in under-resolved or marginally-resolved simulations of convection-dominated flows. Numerical stabilization is often used to alleviate these oscillations. In this study, we investigate an evolve-filter-relax regularized reduced order model (EFR-ROM) to add numerical stabilization to proper orthogonal decomposition (POD) ROMs for the Boussinesq equations. The stability and the error analysis of the full discretization of the model are presented. Numerical examples illustrate the theoretical results.

Medine Demir
Postdoctoral Researcher
demir@wias-berlin.de

MS11

Physics-Informed Active Learning with Simultaneous Weak-Form Latent Space Dynamics Identification

The parametric greedy latent space dynamics identification (gLaSDI) framework has demonstrated promising potential for accurate and efficient modeling of high-dimensional nonlinear physical systems. However, it remains challenging to handle noisy data. To enhance robustness against noise, we incorporate the weak-form estimation of nonlinear dynamics (WENDy) into gLaSDI. In the proposed weak-form gLaSDI (WgLaSDI) framework, an autoencoder and WENDy are trained simultaneously to discover intrinsic nonlinear latent-space dynamics of high-dimensional data. Compared to the standard sparse identification of nonlinear dynamics (SINDy) employed in gLaSDI, WENDy enables variance reduction and robust latent space discovery, therefore leading to more accurate and efficient reduced-order modeling than the strong form. Furthermore, the greedy physics-informed active learning in WgLaSDI enables adaptive sampling of optimal training data on the fly for enhanced modeling accuracy. The effectiveness of the proposed framework is demonstrated by modeling various nonlinear dynamical problems, including viscous and inviscid Burgers' equations, time-dependent radial advection, and the Vlasov equation for plasma physics. With data that contains 5-10% Gaussian white noise, WgLaSDI outperforms gLaSDI by orders of magnitude, achieving 1-7% relative errors and 121 to 1,779x speed-up compared with the high-fidelity models.

Xiaolong He
ANSYS
xiaolong.he@ansys.com

April Tran, David M. Bortz
University of Colorado at Boulder
chi.tran@colorado.edu, dmbortz@colorado.edu

Youngsoo Choi
Lawrence Livermore National Laboratory
choi15@llnl.gov

MS11

Bridging Large Eddy Simulation and Reduced Order Modeling of Convection-Dominated Flows through Spatial Filtering

The presence of flow phenomena over large spatial and temporal ranges in convection-dominated flows poses significant challenges to traditional ROMs: a large number of modes may be required to accurately describe the fluid dynamics, which limits the computational efficiency. On the other hand, if one chooses to reduce the number of modes to improve efficiency, a severe loss of information compromises the solution accuracy. We propose to recover stability for classical ROMs through closures and stabilizations that are inspired by Large Eddy Simulation (LES). A key ingredient for the construction of these ROMs, which we call LES-ROMs, is spatial filtering, i.e., the same principle used to build classical LES models. This ensures a modeling consistency between LES-ROMs and the approaches that generated the data used to train them. We will show that LES-ROMs are extremely easy to implement, very efficient, and, when carefully tuned, accurate in capturing the average physical quantities of interest in challenging convection-dominated flows.

Annalisa Quaini
Department of Mathematics, University of Houston
quaini@math.uh.edu

MS11

A Parametric Rom Framework for Inverse Problems in High Speed Flows

Addressing inverse problems in high-speed flow simulations is critical due to the uncertainties that can impact outcomes. These problems are central to real-world applications, where accurate reconstruction or estimation of system parameters from observed data is vital for safety, performance, and design in high-speed scenarios. The need to solve parametric or stochastic partial differential equations in computational fluid dynamics further complicates these tasks, requiring significant computational resources and advanced modeling techniques. Reduced-order modeling (ROM) has emerged as a key approach, simplifying the solution of inverse problems without sacrificing accuracy. This study introduces a ROM framework for inverse problems that utilizes an autoencoder (AE) to reduce dimensionality while preserving essential dynamics. The research compares AE based ROM's performance with a full-order model, using Markov Chain Monte Carlo methods to explore the solution space and quantify uncertainties. Applying both to the Sod shock tube problem, this study conducts a parametric analysis to examine how variations in snapshot numbers impact the ROMs robustness and accuracy. The findings offer valuable insights into the effectiveness of ROM techniques for solving inverse problems in high-speed flow simulations.

Bipin Tiwari
University of Tennessee at Knoxville
btiwari1@vols.utk.edu

Omer San
Mechanical, Aerospace, and Biomedical Engineering

osan@utk.edu

MS12

Mimetic Metrics for the Dgsem

Free-stream preservation is an essential property for numerical solvers on curvilinear grids. Key to this property is that the metric terms of the curvilinear mapping satisfy discrete metric identies, i.e., have zero divergence. We present an alternative approach for discontinuous Galerkin spectral element methods (DGSEM) that guarantees such divergence free metric terms. Divergence free metric terms are furthermore essential for entropy stability on curvilinear grids. Our proposed mimetic approach uses projections that are compatible with the de Rham Cohomology.

Daniel Bach, Andrés Rueda-Ramírez
University of Cologne
daniel.bach@uni-koeln.de, aruedara@uni-koeln.de

David A. Kopriva
Department of Mathematics
The Florida State University
kopriva@math.fsu.edu

Gregor Gassner
Mathematical Institute
University of Cologne
ggassner@uni-koeln.de

MS12

Hybrid Learning of Spatiotemporal Neural Operators for Turbulent Flows

Recent advancements in operator-type neural networks have shown promising results in approximating the solutions of spatiotemporal PDEs. However, these neural networks often entail considerable training expenses, and may not always achieve the desired accuracy required in many scientific and engineering disciplines. In this paper, we propose a new Spatiotemporal Fourier Neural Operator (SFNO) that learns maps between Bochner spaces, and a new learning framework to address these issues. This new paradigm leverages wisdom from traditional numerical PDE theory and techniques to refine the commonly used end-to-end pipeline. Specifically, in the learning problems for the turbulent flow modeling by the Navier-Stokes Equations (NSE), the proposed architecture initiates the training with a few epochs for SFNO, and then the last linear spectral convolution layer is fine-tuned without the frequency truncation. The optimization uses a negative Sobolev norm for the first time as the loss in operator learning, defined through a reliable functional-type a posteriori error estimator. This design allows the neural operators to effectively tackle low-frequency errors while the relief of the de-aliasing filter addresses high-frequency errors. Numerical experiments on commonly used benchmarks for the 2D NSE operator learning demonstrate a usage of 1% of FLOPs training cost of traditional end-to-end pipeline and achieve up 100000 times more accurate solutions.

Shuhao Cao
Washington University in St. Louis
scao@umkc.edu

Francesco Brarda
Emory University
francesco.brarda@emory.edu

Ruipeng Li
Lawrence Livermore National Laboratory
li50@llnl.gov

Yuanzhe Xi
Department of Mathematics
Emory University
yuanzhe.xi@emory.edu

MS12

Siac Filtering for the Stochastic Galerkin Method

In this talk, we present an innovative application of the SIAC family of filters to enhance a Polynomial Chaos method for solving wave equations with uncertain coefficients. By evolving the chaos expansion coefficients using the discontinuous Galerkin method, we implement SIAC filtering in the finite element framework. Our theoretical analysis and numerical experiments reveal significant improvements in solution accuracy and reduction of statistical noise, showcasing the effectiveness of the SIAC filter.

Andres Galindo-Olarte
University of Texas at Austin
afgalindo@utexas.edu

MS12

High-Order Adaptive Rank Integrators for Multi-Scale Linear Kinetic Transport Equations in the Hierarchical Tucker Format

In this talk, I will present a new adaptive rank method for the linear kinetic transport equation that leverages a macro-micro decomposition, as well as a discrete ordinates method that uses tensor product structure in the angular domain. To address the challenges associated with the curse of dimensionality, the proposed low-rank method is cast in the framework of the hierarchical Tucker decomposition.The adaptive rank integrator is built upon high-order discretizations for both time and space, and it applies high-order singular value decomposition (HOSVD) type truncation to therepresentation of the time-dependent kinetic function in a dimension tree. The methods are applied to several test problems from the literature, and we compare the low-rank solutions against their corresponding full-rank implementations.

William A. Sands
Michigan State University
wsands@udel.edu

Wei Guo
Texas Tech University
weimath.guo@ttu.edu

Jing-Mei Qiu
University of Delaware
jingqiu@udel.edu

Tao Xiong
University of Science and Technology of China
taoxiong@ustc.edu.cn

MS13

Extended Galerkin Neural Network Approximation of Singular Variational Problems with Error

Control

We present extended Galerkin neural networks (xGNN), a variational framework for approximating general boundary value problems (BVPs) with error control. The main contributions of this work are (1) a rigorous theory guiding the construction of new weighted least squares variational formulations suitable for use in neural network approximation of general BVPs (2) an "extended' feedforward network architecture which incorporates and is even capable of learning singular solution structures, thus greatly improving approximability of singular solutions. Numerical results are presented for several problems including steady Stokes flow around re-entrant corners and in convex corners with Moffatt eddies in order to demonstrate efficacy of the method.

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

Justin Dong
Lawrence Livermore National Laboratory
dong9@llnl.gov

MS13

Enhancing Latent Space Bayesian Optimization with Latent Data Augmentation for Efficient De Novo Design

Latent Space Bayesian optimization (LSBO) is emerging as a powerful approach for de novo design, with applications in the design of chemical compounds and crystal materials. By leveraging the expressive power of generative models, LSBO navigates the complex and high-dimensional design space more efficiently than traditional methods. In this framework, a VAE is trained to encode chemical structures or crystal configurations into a continuous latent space, capturing essential features and relationships. BO is then applied within this latent space, using a surrogate model to iteratively suggest new candidates that maximize a target property or objective function. This approach demonstrates significant improvements in optimization efficiency, offering promising directions for de novo discovery. However, LSBO faces challenges due to the mismatch between the objectives of VAE and BO, resulting in poor exploration capabilities and difficulties in ensuring the diversity and novelty of generated structures. In this study, we introduce the concept of latent consistency/inconsistency as a crucial problem in LSBO, arising from the VAE-BO mismatch. To address this issue, we propose a novel usage of data augmentations in the latent space, which resolves the problem of latent inconsistencies while accelerating the de novo design process. We showcase the performance of our methodologies using both chemical compound and crystal material design objectives.

Onur Boyar, Ichiro Takeuchi
Nagoya University
boyaronur@gmail.com, ichiro.takeuchi@mae.nagoya-u.ac.jp

MS13

Epistemic Uncertainty Analysis Using Physics Informed Surrogates and Dempster-Shafer Theory

In this talk, we introduce the numerical strategy for epistemic uncertainty quantification using Dempster-Shafer (DS) theory, and for sensitivity analysis in the framework of DS theory. Specifically, the epistemic uncertainty in the system is represented using non-probabilistic uncertain variables with belief functions. Using the introduced numerical methods, we quantify the uncertainty in the output quantities of interest and study their sensitivity with respect to the uncertain inputs. When observational data is available, the mathematical representation of the uncertainty in inputs can be updated. To reduce the computational cost, physics informed surrogates are adopted to serve as approximations for the full simulation.

Yanyan He
University of North Texas
yanyan.he@unt.edu

Lindsay Waldrop
Chapman University
Biological Sciences
waldrop@chapman.edu

MS13

Efficient Uncertainty Quantification for Large-Scale Scientific Machine Learning Via Ensemble Kalman Inversion

Uncertainty quantification in scientific machine learning, particularly for neural network-based methods, has garnered significant attention. However, current inference techniques often face challenges such as high computational costs for high-dimensional posterior inference or inadequate uncertainty estimates. In this talk, we introduce an efficient uncertainty quantification framework for physics-informed neural networks and deep operator learning using Ensemble Kalman Inversion (EKI). Our findings demonstrate that the proposed method delivers uncertainty estimates that are as informative as those obtained through Hamiltonian Monte Carlo (HMC)-based approaches, but with significantly lower computational costs. Additionally, we shall discuss the promising ways of extending this approach to larger-scale networks.

Xueyu Zhu
Department of Mathematics, University of Iowa
xueyu-zhu@uiowa.edu

MS14

Physics-Preserving AI-Accelerated Simulations of Plasma Turbulence

Countless phenomena in laboratory and astrophysical plasmas involve turbulence. Its inherent complexity, involving the nonlinear dynamics of a large number of degrees of freedom in an open (i.e., non-isolated) physical system, cannot be tackled computationally in a brute-force style even on emerging exascale platforms. Therefore, efficient models of turbulent flows are indispensable, and one of the most popular of these is the Large Eddy Simulation (LES). Here, only the often most relevant large scales are retained explicitly, while small-scale dynamics are described by sub-grid models. We generalize this approach via Machine Learning (ML) techniques and apply our ideas to self-driven plasma turbulence described by the Hasegawa-Wakatani equations, which can be reduced to the Navier-Stokes equations in a certain limit. The combination of a potential-based ML subgrid model with temporally consistent training allows us to remove large parts of the inertial range and to reduce the computational effort by three orders of magnitude while retaining the statistical physical properties of

the turbulent system. We also explore the key question if the learning can be restricted to selected points in parameter space, which would further enhance the method's value and open new doors for future applications.

Frank Jenko
Max Planck Institute for Plasma Physics
The University of Texas at Austin
frank.jenko@ipp.mpg.de

MS14

Variational Auto-Encoders for Surrogate Models of 2D Turbulence

Turbulent transport represents one of the major topics in plasma physics, especially taking into account its impact on the performance of nuclear fusion devices. Since modelling turbulence requires long-time simulations, the use of surrogate models might represent a good compromise between computational cost and physical accuracy. We present a Generative Artificial Intelligence Turbulence (GAIT) surrogate model for 2D plasma turbulence described by the Hasegawa-Wakatani (HW) set of equations. The proposed GAIT model is able to perform fast and long-time turbulent transport computations. It employs a combination of a convolutional variational autoencoder (CVAE) and a densely connected deep neural network (DNN). The CVAE is used to encode snapshots of computed HW turbulence states into a reduced latent space. The DNN is trained to reproduce the time evolution of turbulence in the latent space. Once the GAIT model is trained, new turbulence states are obtained by decoding the latent space dynamics generated by the DNN. The AI generation process is about 500 times faster than the direct numerical integration of the HW model. Several tests, based on both Eulerian and Lagrangian metrics, are presented to show the excellent fidelity of the model. To better understand the encoding of the CVAE, the topology of the turbulent training dataset in the latent space is explored, emphasizing the impact of its dimension on the time evolution of the GAIT-generated turbulence.

David Zarzoso, Benoit Clavier
Aix-Marseille Université, CNRS. Marseille, France
david.zarzoso-fernandez@univ-amu.fr,
benoit.clavier@univ-amu.fr

Diego Del-Castillo-Negrete
Oak Ridge National Laboratory
delcastillod@ornl.gov

Emmanuel Frenod
Universite Bretagne Sud, Centre Yves Coppens, Vannes, France
emmanuel.frenod@univ-ubs.fr

MS14

Latent Space Mapping: Revolutionizing Predictive Models for Exhaust Control in Fusion Plasmas

The inherent complexity of boundary plasma at the boundary region of fusion devices, characterized by multi-scale dynamics and nonlinear multi-physics coupling, has historically limited time-consuming, high-fidelity models to scientific research. Operational challenges in tokamak (a magnetic fusion device) control and scenario development have necessitated reliance on over-simplified empirical methods. This work introduces a transformative machine-learning strategy that bridges this gap by developing rapid, precise surrogate models that encapsulate complex plasma and neutral physics. Utilizing latent space mapping, we efficiently represent complex divertor plasma states in a low-dimensional space, streamlining predictive model construction. We have developed specialized surrogate models through high-fidelity simulations. These models provide quasi-real-time predictions (about 100 us versus hours of direct simulation) with exceptional accuracy (less than 20% relative error), forecasting crucial plasma parameters such as electron density, temperature, heat loads, and peak radiation location for effective plasma exhaust control. Their ability to precisely predict detachment phases (i.e., a desired operational state where impurities and neutral gas cool down hot plasma exhaust before it reaches the vessel wall) marks a significant advancement in plasma physics, enabling new model-predictive control strategies for current experiments and future fusion reactors.

Ben Zhu, Menglong Zhao, Xue-Qiao Xu
Lawrence Livermore National Laboratory
zhu12@llnl.gov, zhao17@llnl.gov, xu2@llnl.gov

Xinxing Ma, KyuBeen Kwon, Anchal Gupta, David Eldon
General Atomics
maxinxing@fusion.gat.com, kwonk@fusion.gat.com, guptaa@fusion.gat.com, eldond@fusion.gat.com

MS14

Probabilistic Sensitivity Analysis of Fusion Plasma Transport Simulations for Integrated Modeling and Machine Learning

Future magnetic confinement fusion reactors will operate in new plasma regimes where a predictive capability to ensure performance within safe operational limits is required. High fidelity gyrokinetic codes of turbulent transport, inhibiting confinement, are computationally prohibitive to span the number of degrees of freedom present at even the experimental scale. Integrated modeling of reduced plasma profiles utilize lower fidelity quasilinear theory that approximates the nonlinear turbulence to achieve computations on the order of seconds for each radial position. Machine learning of surrogates has been proposed as a framework to further advance transport predictions to near real-time for scenario control and design optimization. Recent success with active learning motivates the application of physics-based constraints to the normalization of the large parameter space of nearly 30 dimensions facilitating surrogate correspondence between current and future devices. We employ so-called natural units and normalizing flows to construct multilayer perceptron model ensembles with sampling distributions quantified by the divergence and uncertainty estimated from the relative sensitivity across independent parameter dimensions. Up to an order of magnitude reduction in training data size for similar predictive performance to previous efforts is found, promoting efficient simulation acquisition strategies to the generation of surrogate models for integrated modeling.

Sebastian de Pascuale, Alberto Bocchinfuso, Minglei Yang, Mark Cianciosa
Oak Ridge National Laboratory
depascuales@ornl.gov, bocchinfusoa@ornl.gov, yangm@ornl.gov, cianciosamr@ornl.gov

Juan M. Restrepo
Oak Ridge National Lab
restrepojm@ornl.gov

Diego Del-Castillo-Negrete, Jin Myung Park, Gary Staebler
Oak Ridge National Laboratory
delcastillod@ornl.gov, parkjm@ornl.gov, staeblergm@ornl.gov

MS15

Spatio-Temporal Energy Cascade from Magnetohydrodynamic to Kinetic Scale in Plasma Turbulence

Plasma turbulence is known to produce fluctuations with broadband wavenumber and frequency energy spectra. Numerical and observational studies have shown that turbulent astrophysical and space plasmas are anisotropic with respect to local magnetic fields, and exhibit non-trivial temporal properties, with most of the turbulent energy stored in low frequency fluctuations, with wavenumbers almost perpendicular to the ambient magnetic field. In our work, we present a new approach to study the spatio-temporal properties of plasma turbulence. We introduce a new set of scale-filtered magnetohydrodynamic (MHD) equations, to study the transfer of turbulent energy in wavenumber-frequency space. We test our method on MHD simulations of solar wind turbulence, showing that the turbulent cascade is frequency dependent, and low frequency fluctuations arise from an inverse cascade in frequency space. Low frequency modes act as a stable energy reservoir that supports the overall energy cascade. Our new method represents a new way to investigate turbulence and its spatio-temporal properties. Extensions of this technique to kinetic turbulence are discussed.

Giuseppe Arro
Los Alamos National Laboratory
garro@lanl.gov

MS15

Suppressing Kinetic Plasma Instabilities by PDE Constraint Optimization

Maintaining the stability and shape of a plasma is a crucial task in many technological applications ranging from beam shaping to fusion energy. This is often challenging as plasma systems tend to be naturally unstable and kinetic effects can play an important role in the behavior of the instabilities. Our goal is to find external fields and beam profiles that are able to stabilize the plasma system. This is a challenging global optimization problem as the landscape of the objective function is oscillatory and has many local minima. We propose to use a global optimization algorithm based on genetic evolution that is complemented by a local gradient-based scheme to speed up convergence. We, in particular, observe that nonlinear effects can be used to stabilize plasma systems even if the system is linearly unstable. We also discuss high-performance computing aspects. In particular, optimization problems require the solution of many similar problems for different parameters. In this context, we will discuss our batched Vlasov that is able to utilize GPUs via the Kokkos framework.

Lukas Einkemmer
University of Innsbruck
lukas.einkemmer@gmail.com

MS15

Application of Mesh Refinement to Relativistic Magnetic Reconnection

Relativistic magnetic reconnection is a non-ideal process where strong antiparallel magnetic fields undergo a rapid change in topology releasing large amounts of energy in the form of non-thermal particle acceleration. Reconnection is often invoked to explain high-energy emissions in astrophysical systems. While much progress has been made in understanding the physics of reconnection, especially in 2D, the application of advanced algorithms to improve computational efficiency has been limited. In this talk, we will present our work on applying mesh refinement to 2D reconnection simulations to efficiently model the inherent disparity in length-scales. We use the ultrahigh-order pseudo spectral analytical time-domain (PSATD) Maxwell solver as it can mitigate numerical dispersion that occurs with the traditionally used finite-difference time domain (FDTD) method. Using a mesh refinement ratio of 8 for a 2D reconnection system, we obtained good agreement with the high resolution baseline simulations, using only 36% of the macroparticles and 71% of the node hours needed for the baseline.

Revathi Jambunathan
Lawrence Berkeley National Laboratory
rjambunathan@lbl.gov

Henry Jones, Lizzette Corrales
LBNL
n.henryjones@gmail.com, lc965@cornell.edu

Hannah Klion
Lawrence Berkeley National Laboratory
klion@lbl.gov

Michael Rowan
AMD
michael.rowan@amd.com

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
WeiqunZhang@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

MS15

On the Parallel Performance of a CFD High-Order Spectral Element Code on GPU Based Supercomputers

This work presents the parallel performance of the CFD code SOD2D in different brand-new HPC platforms. SOD2D is a Continuous Galerkin High-Order Spectral Element Code designed to solve simulations of both turbulent compressible and incompressible flows. SEM is selected since it does not require the reconstruction of fluxes between the elements, which could introduce numerical dissipation. Moreover, using SEM is advantageous when using hexahedra for discretizing the spatial domain. The proposed scheme is stabilized by employing a modified version of the Entropy Viscosity model proposed by Guer-

mond et al.. Different integration schemes are implemented: a fully 4th order Runge-Kutta explicit scheme and an Implicit-Explicit Runge-Kutta (IMEX-RK) for the compressible solver, and the velocity-correction integration scheme BFD/EXT-3 proposed by Karniadakis et al. for the incompressible solver. The parallel performance of the SOD2D software will be thoroughly analysed and assessed in different HPC platforms. Specifically, in the present work results for the following HPC platforms will be presented: i) the MareNostrum 5, placed at the Barcelona Supercomputing Center and equipped with GPUs NVIDIA H100; ii) the Leonardo HPC system, hosted by CINECA, based on NVIDIA A100 GPUs with all the nodes interconnected through an Nvidia Mellanox network, iii) and the Karolina supercomputer, located a the IT4Innovations at the Technical University of Ostrava, based on NVIDIA A100 GPUs.

Oriol Lehmkuhl, Jordi Muela, Lucas Gasparino
BSC
oriol.lehmkuhl@bsc.es, jordi.muela@bsc.es, lucas.gasparino@bsc.es

MS16

Modeling Unobserved Variables in Dynamical Systems

Consider a physical phenomenon described by two models: experimental and simulation. The first measures QoI or a solution of the model, but it is expensive to obtain data from it, the second model is an approximation of the ground truth model, but it is cheap to compute. The experimental model depends on known and unknown variables, and the simulation depends on the known variables. There are several challenges. First, there is no clear way of how many unknown variables are or to know what is the distribution of these variables. Next, obtaining samples from the experimental is expensive. Third, the measures of the experimental data is very sensitive to its inputs. Since we know a reduce number of these variables, for a given set of known variables, you can have a family of measurements for this known set. Recent works have used Bayesian modeling and invertible neural networks, to model the influence of unknown parameters, but have struggled to learn a map from the simulation to the experimental model. We propose a general framework that models the discrepancy between these two models using the same input parameters. To this end, we add random variables drawn from a known prior distribution to construct a probabilistic map between the models. To construct this map we use DNNs trained on the empirical and a suitable loss to learn probability distribution on the data. Our numerical results have shown that we can measure the discrepancy between these two models.

Juan Cardenas
University of Colorado Boulder
juca7891@colorado.edu

MS16

Interpolative Decomposition for Bi-fidelity Approximation in a Stochastic System

The interpolative decomposition has previously been applied to non-intrusively approximate a high fidelity data matrix by exploiting the low-rank structure of a cheap, low fidelity data matrix, which may be useful in uncertainty quantification. Focusing on the problem of pressure trace prediction in a laser-ignited combustor, we demonstrate the breakdown of this method on stochastic multi-modal data, and propose a straightforward extension to alleviate this.

Murray C. Cutforth, Tiffany Fan
Stanford University
mcc4@stanford.edu, tiffan@stanford.edu

Eric F. Darve
Institute for Computational and Mathematical Engineering
Stanford University
darve@stanford.edu

Tony Zahtila
Stanford University
tzahtila@stanford.edu

Juan Cardenas
University of Colorado Boulder
juca7891@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS16

Model-free Quantification of Completeness, Uncertainties, and Outliers in Atomistic Machine Learning Using Information Theory

An accurate description of information is relevant for a range of problems in atomistic machine learning (ML), such as crafting training sets, performing uncertainty quantification (UQ), or extracting physical insights from large datasets. However, atomistic ML often relies on unsupervised learning or model predictions to analyze information contents from simulation or training data. In this talk, I will introduce a theoretical framework that provides a rigorous, model-free tool to quantify information contents in atomistic simulations. I will describe how the information entropy of a distribution of atom-centered environments explains known heuristics in ML potential developments, from training set sizes to dataset optimality. Using this tool, I will showcase how information entropy can produce a model-free UQ method that reliably predicts epistemic uncertainty and detects out-of-distribution samples, including rare events in systems such as nucleation. This method provides a general tool for data-driven atomistic modeling and combines efforts in ML, simulations, and physical explainability.

Daniel Schwalbe-Koda
UCLA
dskoda@ucla.edu

MS16

Uncertainty Quantification in KolmogorovArnold Networks for Solving Differential Equations

Physics-informed KolmogorovArnold networks (PIKANs) have recently emerged as an alternative approach for solving differential equations, offering great performance comparable to physics-informed neural networks (PINNs). However, the exploration of uncertainty quantification in PIKANs remains limited, which hinders their reliable deployment in risk-sensitive applications. In this work, we

integrate PIKANs with a range of UQ methods, both Bayesian and non-Bayesian, to address forward and inverse problems involving differential equations with UQ under noisy data conditions. We conduct a systematic and comprehensive study over various UQ methods for PIKANs and PINNs, focusing on their effectiveness in quantifying epistemic uncertainty. Additionally, we examine different data conditions, including large datasets with high noise, small datasets with low noise, and unevenly distributed data. Our findings reveal that non-Bayesian methods perform as well as Bayesian methods while requiring significantly less computational cost. These results underscore the potential of PIKANs for robust uncertainty quantification in solving differential equations, paving the way for their application in fields where reliability and risk management are critical.

Liang Yue
New York University
ly2182@nyu.edu

MS17

Distributed and Parallel Efficiency of Some Methods Using Large Sparse Matrices on Different Cluster Architectures

While scientific and technical computing often appear to be split into separate branches (High Performance Computing, Data/Graph Analysis, Machine Learning), the underlying algorithms are often linear algebra problems. As a result several efforts have been started to make this more explicit (for example GraphBLAS in high performance graph analysis) and benefit from optimized generalized libraries. However the inherent structure of all of these problems is sparse, meaning that the overall performance of the code will be far from the peak which would be obtained with dense linear algebra on commonly used cache-heavy computer architectures. In this study, we focus on several benchmarks (PageRank, conjugate gradient) and test their performance on different cluster architectures (with different CPU families and multiple generations of interconnect). Our implementation is designed to handle large problems by distributing the vectors across the system.
As sparse linear algebra is fundamentally memory-bound, we measure the performance of different sub-systems (memory, communication layers) to understand their influence on the overall results.

Geraud P. Krawezik
Flatiron Institute
gkrawezik@flatironinstitute.org

Maxence Buisson
Polytech Lille
maxence.buisson@polytech-lille.net

MS17

Parallel Scalable Domain Decomposition Methods and CNN Surrogate Models for Generalized Newtonian Blood Flow

Hemodynamics is fundamental to the functioning of the human organism, driving the advancement of computationally efficient simulation methodologies to accurately predict blood flow dynamics. In the first part, we demonstrate simulations of blood flow at the macroscopic level, employing a generalized Newtonian constitutive model. The data-driven viscosity curves are based on dissipative particle dynamics virtual rheometer simulations of blood. The models are incorporated into our in-house FEM solver, FEDDLIB (Finite Element and Domain Decomposition Library), a C++ library designed for large-scale simulations and which provides an interface to the highly scalable implicit domain decomposition solver FROSch (Fast and Robust Overlapping Schwarz), a solver package in Trilinos. We emphasize the utilization of monolithic overlapping Schwarz preconditioners with GDSW/RGDSW coarse spaces and present scalability results. In the second part, we present initial efforts toward achieving accelerated blood flow predictions using surrogate models. Simulation data from the aforementioned method is used to train a CNN-based surrogate model employing an overlapping domain decomposition strategy. Artery geometries are divided into subdomains, which serve as the training basis for the CNN. Challenges faced and initial results obtained are discussed. The viscosity data was provided by the group of Prof. Gerhard Gompper, Dr. Dmitry Fedosov, and Alper Topuz (FZ Jlich).

Axel Klawonn
Department of Mathematics and Computer Science and Center
for Data Centand Simulation Science, University of Cologne
axel.klawonn@uni-koeln.de

Natalie Kubicki
University of Cologne
nkubicki@uni-koeln.de

Martin Lanser
Universitaet zu Koeln
Mathematisches Institut
martin.lanser@uni-koeln.de

Takashi Shimokawabe
Information Technology Center
The University of Tokyo
shimokawabe@cc.u-tokyo.ac.jp

MS17

Road to "AI for Science": Exploring Software Sustainability through "Couplers"

"Coupler" is originally a tool for coupling multiple simulation models such as atmosphere and ocean, structure and fluid. In recent years, computer systems and workloads have become more diverse, and the role of couplers in supercomputing has become more important. In this talk, we focus on the "history" of couplers and consider what software sustainability means. We briefly describe three projects, ppOpen-HPC (2011-2018), h3-Open-BDEC (2019-2024), and JHPC-quantum (2023-2028), and will introduce how couplers have evolved and what role they have been playing in supercomputing.

Kengo Nakajima
Information Technology Center, The University of Tokyo
RIKEN Center for Computational Science (R-CCS)
nakajima@cc.u-tokyo.ac.jp

Shinji Sumimoto
The University of Tokyo
sumimoto@cc.u-tokyo.ac.jp

Yao Hu
KOIBUCHI Laboratory
National Institute of Informatics

hu@cc.u-tokyo.ac.jp

Kazuya Yamazaki
Information Technology Center, The University of Tokyo
kyamazaki@cc.u-tokyo.ac.jp

Takashi Arakawa
CliMTech
arakawa@climtech.jp

Hisashi Yashiro
National Institute for Environmental Studies
yashiro.hisashi@nies.go.jp

MS17

Very Sparse and Very Large Parallel and Distributed Matrix Computing

Exascale machines are now available, based on several different arithmetic types and using different architectures (with on-chip network processors and/or with accelerators). Brain-scale applications, in machine learning and AI for example, manipulate huge graphs that lead to sparse linear algebra problems. Many supercomputers were designed primarily for numerical simulations, not for machine learning and artificial intelligence. The new applications that are maturing after the convergence of Big Data and HPC towards machine learning and AI are being developed using different programming and execution paradigms. Nevertheless, performance often depends on the parallel and distributed efficiency of the linear algebra method. Several methods, including some for linear algebra, generate computational sequences of very large non-Hermitian matrices. In this talk, we present the results of such calculations on the Japanese supercomputer Fugaku, No. 1 for several semesters on the HPCG list. We present some results for the sequencing of very sparse and very large matrices by vector products, which often lead to the distribution of these vectors themselves, with respect to the sparsity, the matrix size, on the one hand, and the number of processes and nodes, as well as network topologies, on the other. Next, we discuss the potential effectiveness of important iterative and restarted linear algebra methods that are relevant to AI and machine learning on such supercomputers.

Serge G. Petiton
University of Lille, and CNRS
serge.petiton@univ-lille.fr

Maxence Vandromme
RIKEN
maxence.vandromme@riken.jp

Simon Huaume, Yves Ananga-Ananga
University of Lille
simon.huaume.etu@univ-lille.fr,
yves.anangaananga.etu@univ-lille.fr

MS18

Analysis of Pancreas on Computed Tomography by Machine Learning and Shape Analysis

Pancreatic cancer remains one of the most lethal forms of cancer, with early detection being crucial for patient survival. Computed tomography (CT) is one of the primary tools for identifying abnormal changes in the pancreas before other examinations. However, correctly identifying these changes requires the expertise of experienced radiologists. The reliance on experienced radiologists is human labor-intensive. Also, some subtle changes in the pancreas can be easily overlooked. Due to these issues, there is a pressing need for automatic AI tools to assist radiologists in analyzing CT images more accurately and efficiently. In this talk, we will present our development of several AI workflows designed to distinguish pancreatic cancer from unremarkable findings on CT scans. Our workflow utilizes deep learning for pancreas segmentation and feature analysis within the segmented region, including shape-based and non-shape-based characteristics. These models have been trained and validated using CT data from the National Taiwan University Hospital, annotated by radiologists, to assess their potential for further use in clinical routine.

Weichung Wang
National Taiwan University
Institute of Applied Mathematical Sciences
wwang@ntu.edu.tw

Dawei Chang
Data Science Degree Program
National Taiwan University and Academia Sinica
davidzan830@gmail.com

MS18

Denoising Autoencoder Neural Networks for Low-Dose Ct Image Enhancement

Noise suppression in medical imaging is crucial to improve accuracy in diagnosis and treatment. Computed Tomography images, in particular, are often affected by noise due to various technical limitations, in addition to the noise produced by reducing the radiation dose used, as well as the noise produced by the reconstruction method. In this context, convolutional neural networks have proven to be powerful tools for noise removal, providing significant improvements in image quality. This study focuses on the application of convolutional neural networks (CNNs) for denoising medical images. An Autoencoder-like architecture has been developed to improve the quality of images obtained by low-dose and/or few-projection Computed Tomography, which eliminates all noise by taking the final full-dose images as ground-truth images. This network will be integrated into a CT image reconstruction process based on an iterative method for solving systems of equations. In this way, starting from the solutions obtained with few iterations of the resolution method, the developed network will be able to improve the images and serve as a regularization process, thus reducing the iterations and the total time needed to obtain the images, as well as improving the final quality.

Anthony Drummond
Universsity of California Berkeley
lanthonydrummond@gmail.com

Vicent Vidal, Mónica Chillarón
Polytechnic University of Valencia
vvidal@upv.es, mnichipr@upv.es

MS18

A Parallel and Scalable Approach for High Performance Learning

In this talk, we will show how to apply the Unite and Conquer approach, used in linear algebra to improve the convergence of iterative methods, to machine and deep learn-

ing techniques. The important features of this inherently parallel and scalable approach make the machine and deep learning techniques thus defined very well suited to multi-level and heterogeneous parallel and distributed architectures. Experimental results demonstrating the interest of these methods for efficient data analysis in the case of clustering, anomaly detection and road traffic simulation will be presented.

Nahid Emad
LI-PaRAD laboratory and Maison de la Simulation
University of Paris Saclay / Versailles
Nahid.Emad@uvsq.fr

Sara Moukir, Quentin Petit, Zineb Ziani
University of Paris Saclay
sara.moukir@universite-paris-saclay.fr, petitquentin41@gmail.com, ziani.zineb.zz@gmail.com

MS18

A Dimensionality Reduction Method Based on Eigenvalue Computation of Modularity Matrices and Its Application to Materials Informatics

In recent years, studies on dimensionality reduction have predominantly focused on nonlinear methods. However, nonlinear approaches often suffer from interpretability issues when generating low-dimensional representations, making them less preferable as preprocessing tools for exploratory data analysis. Therefore, we focus on linear methods, which offer a more intuitive understanding of the process for creating low-dimensional representations. One well-known linear dimensionality reduction method is the Locality Preserving Projections (LPP), which is based on the computation of eigenvalues of the graph Laplacian matrix. In this study, we propose a novel linear dimensionality reduction method based on the computation of eigenvalues of the modularity matrix, which can be seen as an extension of LPP. The proposed method possesses the ability to emphasize the internal structure of dense clusters within high-dimensional spaces, making it particularly effective for applications such as subcluster detection. We demonstrate the effectiveness of our method by applying it for exploratory data analysis on a high-dimensional dataset in materials informatics.

Ibuki Masuda, Yasunori Futamura
University of Tsukuba
s2320645@u.tsukuba.ac.jp, futamura@cs.tsukuba.ac.jp

Jianbo Lin, Anh Khoa A Lu, Ryo Tamura, Tsuyoshi Miyazaki
National Institute for Materials Science
lin.jianbo@nims.go.jp, lu.augustin@cello.t.u-tokyo.ac.jp, tamura.ryo@nims.go.jp, miyazaki.tsuyoshi@nims.go.jp

Tetsuya Sakurai
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

MS19

Assessing Trilinos Linear Solver Stack Performance Across the DOE Complex

The open source Trilinos project provides a collection of interdependent, high performance numerical software libraries that are used by applications at Sandia and elsewhere in the DOE. One of the project goals is to deliver scalable algorithms that will run well on new and emerging supercomputers. In this presentation we examine performance of Trilinos solver algorithms across a variety application spaces and architectures (CPUs and accelerators). We discuss recent advances as they pertain to performance, challenges and opportunities, and future directions.

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

Christopher Siefert
Sandia National Laboratories
csiefer@sandia.gov

MS19

Neko: A Modern, Portable, and Scalable Framework for Extreme-Scale Computational Fluid Dynamics

Recent trends and advancements in including more diverse and heterogeneous hardware in High-Performance Computing (HPC) are challenging scientific software developers in their pursuit of efficient numerical methods with sustained performance across a diverse set of platforms. As a result, researchers are today forced to refactor their codes to leverage these powerful new heterogeneous systems. We present our work on addressing the extreme-scale computing challenges in computational fluid dynamics, ensuring exascale readiness of turbulence simulations. Focusing on Neko, a high-fidelity spectral element code, we outline the optimisation and algorithmic work necessary to ensure scalability and performance portability across a wide range of platforms. Finally, we present performance measurements on a wide range of accelerated computing platforms, including the EuroHPC pre-exascale system LUMI and Leonardo, where Neko achieves excellent parallel efficiency for an extreme-scale direct numerical simulation (DNS) of turbulent thermal convection using up to 80% of the entire LUMI supercomputer.

Niclas Jansson
PDC Center for High Performance Computing
KTH Royal Institute of Technology
njansson@kth.se

Martin Karp
KTH Royal Institute of Technology
makarp@kth.se

Stefano Markidis
KTH Royal Institute of Technology, Sweden
markidis@kth.se

Philipp Schlatter
FriedrichAlexander-Universität Erlangen-Nürnberg (FAU)
philipp.schlatter@fau.de

MS19

Investigating Openmp Offloading for Porting An Established CFD solver to GPUs

As heterogeneous architectures have become the norm in HPC, one may begin to wonder how, or even if, older codes can be updated to better utilize these modern HPC architectures. In this talk we discuss the performance of and our experiences with using OpenMP to port an established CFD solver to the GPU. The code in question is CUBE, a

CFD framework developed in-house at RIKEN-RCCS for large-scale industrial CFD simulations. The code is written in Fortran and was designed and optimized to run on the supercomputer Fugaku, an ARM CPU machine. The code has been ported to run on GPUs using OpenMP in a minimally invasive way, ie. only the addition of OpenMP pragmas and some minor restructuring of loops. The performance of the GPU ported CUBE framework is evaluated over a wide variety of HPC hardware and compared with the original CPU version.

Peter Ohm
RIKEN Center for Computational Science
peter.ohm@riken.jp

MS19

Performance Analysis of Scientific and Machine Learning Applications on Vista's Grace-Grace and Grace-Hopper Architectures

The NSF-funded Frontera system has supported crucial scientific applications in high-performance computing (HPC), Big Data, and Machine Learning (ML) over the last five years. As the fastest US academic supercomputer, Frontera 's petascale capabilities empowers hundreds of researchers daily to tackle complex challenges, enabling breakthroughs in science and engineering. With Frontera nearing the end of its production life, an intermediate system, Vista, bridges the gap to its successor, Horizon. This work presents a comprehensive performance analysis of Vista's two node configuration types Grace-Grace (CPU-CPU) and Grace-Hopper (CPU-GPU). We evaluate Vista's performance across two key areas: traditional HPC simulations and ML workloads, including large language model (LLM) training and inference. Our study compares the Grace-Grace nodes to Intel and AMD CPU architectures and the Grace-Hopper nodes with Intel's Ponte Vecchio GPUs and NVIDIA's H100 and A100 GPUs. Our findings indicate significant performance improvements for applications that effectively utilize both CPU and GPU resources, and in particular, the Grace-Hopper architecture shows exceptional promise for ML workloads. We also discuss optimization challenges and opportunities, sharing best practices for leveraging this heterogeneous system's unique features. We conclude with insights on the implications for future supercomputing infrastructure, providing valuable guidance for researchers and practitioners.

Carlos del-Castillo-Negrete
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
cdelcastillo21@gmail.com

Amit Ruhela, John Cazes
Texas Advanced Computing Center
University of Texas Austin
aruhela@tacc.utexas.edu, cazes@tacc.utexas.edu

MS19

Keeping Trilinos Running Performantly Everywhere Every Night

The Trilinos scientific software library is a key enabling technology for application codes in a variety of physics and engineering areas across the US Department of Energy and beyond. This presentation will "pull back the curtain" and describe the process by which the Trilinos Tpetra/Performance team (a) performs nightly testing across five DOE laboratories, (b) how we identify and remedy performance regressions, (c) how we make the validated build scripts available to Trilinos developers, and (d) how we interface with early users and vendors to ensure that vendor library updates work as advertised.

Christopher Siefert
Sandia National Laboratories
csiefer@sandia.gov

MS20

Compressing Structured Matrices Using Matrix-to-Tensor Maps

Compressing large data sets can lead to improvements in both storage and computational speed. By utilizing a map between structured matrices and tensors, we can exploit tensor methods in the analysis of such matrices. With this map, tensor decompositions can be utilized to reduce storage costs for such matrices, as well as revealing additional structure. This talk will explore such methods and how they can be leveraged to reduce storage costs for structured matrices.

Allison Fuller
Arizona State University
tjfuller@asu.edu

MS20

Log-Sum Regularized Kaczmarz Algorithms for High-Order Tensor Recovery

Sparse and low rank tensor recovery has emerged as a significant area of research with applications in many fields such as computer vision. However, minimizing the ℓ_0-norm of a vector or the rank of a matrix is NP-hard. Instead, their convex relaxed versions are typically adopted in practice due to the computational efficiency, e.g., log-sum penalty. In this work, we propose novel log-sum regularized Kaczmarz algorithms for recovering high-order tensors with either sparse or low-rank structures. We present block variants along with convergence analysis of the proposed algorithms. Numerical experiments on synthetic and real-world data sets demonstrate the effectiveness of the proposed methods.

Katherine J. Henneberger
University of Kentucky
k.henneberger@uky.edu

Jing Qin
Department of Mathematics
University of Kentucky
jing.qin@uky.edu

MS20

Projected Tensor-Tensor Product for Multidimensional Data Compression

Tensor-tensor products have demonstrated success in wide-ranging applications involving multidimensional data. These tensor operations look and feel like matrix multiplication, thereby extending desirable properties of standard matrix algebra to tensors. The underlying multiplication relies on an invertible matrix, which can be computationally demanding and memory intensive for large tensors. In this work, we propose a projected tensor-tensor product which replaces the expensive invertible matrix with a

cheap orthogonal projection onto a low-dimensional subspace. We illustrate that projected products still preserve many linear algebraic properties and demonstrate their utility in imaging applications, such as video and hyperspectral image compression.

Katherine Keegan, Elizabeth Newman
Emory University
kkeega3@emory.edu, elizabeth.newman@emory.edu

MS22
CASS Workforce Working Group

The Consortium for the Advancement of Scientific Software (CASS) is dedicated to building a skilled scientific computing software workforce and fostering an innovative culture within the Department of Energy (DOE) computing sciences community. This minisymposium will highlight CASSs key workforce development priorities for 2025 including launching targeted initiatives to strengthen the community, and recognizing contributors who advance professional growth and innovative culture. Through efforts such as the HPC-Workforce Community Groupwhich organizes professional development webinars and quarterly meetings to share progress and best practicesCASS aims to cultivate a workforce prepared to address the evolving demands of scientific computing and software development. Join us to exchange ideas and insights!

Suzanne Parete-Koon
Oak Ridge National Laboratory, U.S.
paretekoonst@ornl.gov

MS22
CASS UDX Working Group

Scientific discoveries increasingly depend on leveraging computation and data synergistically at scale. Software has become a cornerstone to enable seamless, interactive, searchable, collaborative, and reproducible science. CASS was established to foster collaboration across a diverse collection of Software Stewardship Organizations (SSOs) that are each stewarding and advancing a portion of the scientific software ecosystem, e.g., math libraries, data and viz libraries and tools, and programming systems. In the context of the Consortium for the Advancement of Scientific Software (CASS), we are exploring the notion and work needed to ensure user developer experience. This talk will present the background, motivation, and ongoing work in the space.

Lavanya Ramakrishnan
Lawrence Berkeley National Laboratory
lramakrishnan@lbl.gov

MS22
The CASS Impact Framework and Metrics Working Groups

The Impact Framework and Metrics Working Groups (WGs) are two of the several groups recently launched by the Consortium for the Advancement of Scientific Software (CASS) in support of its mission to steward and advance the current and future ecosystem of scientific computing software. The Metrics WG provides a neutral forum to enable a coordinated effort to curate and develop metrics for measuring software sustainability from the perspective of scientific and research software. The Impact Framework WG is focused on developing a flexible framework to gather information from scientific software projects to understand what theyre doing, the progress theyve made, and the impact their work has had. The Impact Framework will rely significantly, but not exclusively, on metrics, identified and implemented in collaboration with the Metrics WG. This talk will present the activities and plans of the two groups and discuss how they work together. It will also talk about ways that individuals and organizations can become active participants in the working group activities.

Terece L. Turton
Los Alamos National Laboratory
tlturton@lanl.gov

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

Berk Geveci
Kitware
berk.geveci@kitware.com

Addi Malviya Thakur, Gregory Watson
Oak Ridge National Laboratory
malviyaa@ornl.gov, watsongr@ornl.gov

MS22
CASS Integration Working Group

The multi-institutional CASS Integration Working Group engages in efforts at the levels of scientific software ecosystem and individual product communities (for example, data and visualization software, math libraries software, etc) with the goal to improve quality, delivery, and interoperability. These efforts target the needs of user communities and the improvement of the Spack and Extreme-Scale Scientific Software Stack (E4S) continuous integration testing. A common objective across Integration Working Group initiatives is to move activities to higher levels where possible to improve efficiency (for example, to move certain kinds of testing from the product community level to the ecosystem level). This talk will introduce the objectives of the CASS Integration Working Group in greater detail, as well as the groups initial plans for the coming years.

James Willenbring
Sandia National Laboratories
jmwille@sandia.gov

MS23
Fluid Moment Models and Particle-in-Cell Monte Carlo Collision Simulations for Low Temperature Plasmas

Low temperature plasmas (LTPs) play an important role in industrial applications, such as semiconductor manufacturing, spacecraft propulsion, and material processing. In rarefied flows, the particle velocity distribution functions (VDFs) become a non-Maxwellian distribution, requiring kinetic models such as grid- and particle-based kinetic methods. One of the most used kinetic methods is the particle-in-cell (PIC) simulation that is coupled with Monte Carlo collision (MCC) models. Depending on the collisionality of the flow, the VDFs can be treated to be a perturbation to an equilibrium VDF, e.g., a Maxwellian distribution. Moment equations can be derived up to an arbitrary order from the kinetic equation; however, any fluid

moment model requires a closure model. In the LTP community, the most popular fluid model employs the so-called drift-diffusion (DD) approximation, which essentially neglects the inertia term in conservation of momentum, allowing for a linear relationship between the electric field, pressure gradient, and collisional drag. In this talk, I will present examples and challenges of fluid and kinetic models for LTP applications. A 5-moment fluid moment model is developed and captures velocity gradient (shear) in cross-field plasmas, which cannot be captured using a DD model. In addition, a PIC-MCC simulation is used to model both DC and RF plasmas.

Kentaro Hara
Stanford University
kenhara@stanford.edu

MS23

Efficient Computational Algorithms for Magnetic Equilibrium in a Fusion Reactor

In magnetic confinement fusion reactors like Tokamaks, hydrogen isotopes are heated to form a plasma that must be confined by magnetic fields to prevent contact with reactor walls. These fields are generated by external coils, but uncertainties in the current, due to factors like temperature fluctuations and material impurities, can affect plasma confinement and overall stability. This study investigates the impact of stochastic current intensities on plasma confinement and estimates the expected behavior of the magnetic field. This work addresses the computational challenges in quantifying uncertainties by proposing a series of methodologies. Firstly, we develop a surrogate function using stochastic collocation on a sparse grid, enabling significant reductions in sampling costs while preserving accuracy relative to traditional nonlinear methods. Secondly, we explore the multilevel Monte Carlo method, which achieves substantial cost efficiency by performing computations on coarser grids. Thirdly, we combine surrogate models with multilevel Monte Carlo to reduce sampling costs, while maintaining accuracy in plasma boundary and geometric descriptors. Lastly, we investigate the multi-fidelity Monte Carlo approach as an additional strategy to further enhance computational efficiency.

Jiaxing Liang
University of Maryland College Park
jliang18@umd.edu

Howard C. Elman
University of Maryland, College Park
helman@umd.edu

Tonatiuh Sanchez-Vizuet
The University of Arizona
tonatiuh@math.arizona.edu

Matthias Heinkenschloss
Computational Applied Mathematics and Operations Research
Rice University
heinken@rice.edu

MS23

Pic-Mcc Simulations Coupled with External Circuitry for Low Temperature Plasma Reactors

Low temperature plasma reactors used for chip fabrication involve complex physical processes which affect the quality of pattern-transfer on the semiconductor substrate. Detailed investigations of the near-surface plasma environment require fully-kinetic approaches. We use a particle-in-cell/Monte Carlo collision (PIC/MCC) approach coupled with external circuitry to model realistic operating conditions and capture the kinetic effects and measure ion energy and angle distribution functions. We build on the exascale-capable PIC/MCC code, WarpX, and leverage python interfaces to couple external circuitry. In this talk, we will demonstrate this inexpensive, non-iterative coupled strategy using a two-dimensional Gaseous Electronics Conference reference cell.

Richard L. Lombardini
St. Marys University
rlombardini@stmarytx.edu

Revathi Jambunathan
Lawrence Berkeley National Laboratory
rjambunathan@lbl.gov

MS23

Merging Particles into a Fluid in Hybrid Fluid-Pic Plasma Simulations

Hybrid fluid/particle-in-cell schemes combine the relative speed of fluid methods, which assume the plasma is near equilibrium, with the wide applicability of kinetic methods, which make no assumptions on the shape of plasma distributions. In particular, hybrid schemes can model the thermalized bulk of a plasma species as a fluid while bumps-on-tails or other far-from-equilibrium effects are modeled kinetically. However, over simulation timescales, collisions may drive a significant portion of the kinetic population toward equilibrium. When this occurs, kinetic schemes become overly expensive, and a fluid treatment is more appropriate. In this work, we outline several computational advantages to minimizing the kinetic portion of a hybrid plasma simulation. We then present progress toward developing methods that actively identify thermalized subsets of particles and incorporate them into an underlying fluid. These methods reduce the fraction of the plasma simulated kinetically, affording considerable computational benefits when applied to appropriate simulations. This work was supported by the DOE NNSA Laboratory Residency Graduate Fellowship DE-NA0003960. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energys National Nuclear Security Administration under contract DE-NA0003525.

Logan T. Meredith
Sandia National Laboratory
lmeredi@sandia.gov

Christopher Moore, William McDoniel
Sandia National Laboratories
chmoore@sandia.gov, wmcdoni@sandia.gov

Davide Curreli
University of Illinois at Urbana-Champaign
dcurreli@illinois.edu

MS23

Algorithms and High-Performance Computing for Kinetic Low-Temperature Plasma Modeling

The particle-in-cell algorithm combined with Monte-Carlo

collisions (PIC-MCC) is the de-facto technique for performing kinetic simulations of low-temperature plasmas. However, due to strict numerical resolution requirements, PIC simulations are currently too costly to perform fast and accurate prototyping of plasma devices. Recently low-temperature plasma PIC codes have been updated to take advantage of modern heterogeneous hardware, including GPUs. Despite these advances, high computational cost still places full scale modeling of plasma devices out of reach. This motivates a return to the fundamentals of the PIC method, and exploration of algorithms to overcome the strict numerical requirements to significantly reduce simulation cost and realize accurate and rapid computational prototyping of plasma devices. We will focus on overcoming the restriction on cell size via the energy-conserving PIC algorithm, including the use of non-uniform grids. A deeper analysis into the required number of particle-per-cell will allow us to define a more rigorous criterion for suitable resolution. These algorithms will be demonstrated on simulations of radio-frequency capacitively coupled plasma discharges relevant for silicon etching. Finally, a sub-cycled hyperbolic electromagnetic field solver will be explored as a faster and more scalable alternative to solving the elliptical Poisson equation. The advantages and disadvantages of each of these techniques will be discussed in depth.

Andrew Tasman Powis, Sierra Jubin, Stephane Ethier, Igor D. Kaganovich
Princeton Plasma Physics Laboratory
apowis@pppl.gov, sjubin@pppl.gov, ethier@pppl.gov, ikaganov@pppl.gov

MS24

Fast Adaptive Integration for Spectral Densities of Gaussian Processes

The specification of a covariance function is of paramount importance when employing Gaussian process models, but the requirement of positive definiteness severely limits those used in practice. Designing flexible stationary covariance functions is, however, straightforward in the spectral domain, where one needs only to supply a positive and symmetric spectral density. We introduce an adaptive integration framework for efficiently and accurately evaluating covariance functions and their derivatives at irregular locations directly from any continuous, integrable spectral density. In order to make this approach computationally tractable, we employ high-order panel quadrature, the nonuniform fast Fourier transform, and a Nyquist-informed panel selection heuristic, and derive novel algebraic truncation error bounds which are used to monitor convergence. As a result, we demonstrate several orders of magnitude speedup compared to naive uniform quadrature approaches, allowing us to evaluate covariance functions from slowly decaying, singular spectral densities at millions of locations to a user-specified tolerance in seconds on a laptop. We discuss extensions to multiple dimensions using a novel nonuniform fast Hankel transform, and apply our methodology to perform gradient-based maximum likelihood estimation for spectral models of environmental spatial data.

Paul G. Beckman
Courant Institute, New York University
paul.beckman@cims.nyu.edu

MS24

Numerical Methods for the SPDE Approach in Space-Time

Most environmental data sets contain measurements collected over space and time. It is the purpose of spatiotemporal statistical models to adequately describe the underlying uncertain spatially explicit phenomena evolving over time. In this talk I will present a new class of spatiotemporal statistical models which is based on stochastic partial differential equations (SPDEs) involving fractional powers of parabolic operators. In particular, I will discuss the efficient approximation of the covariance operators corresponding to the spatiotemporal latent Gaussian processes. The numerical methods will be based on space-time finite element discretizations of parabolic operators, which are supported by a rigorous error analysis. Furthermore, I will address the motivation for employing this class of SPDEs in statistical applications and give an outlook on the computational benefits for statistical inference from spatiotemporal data.

Kristin Kirchner
Delft University of Technology, The Netherlands
KTH Royal Institute of Technology, Sweden
k.kirchner@tudelft.nl

MS24

Estimation and Inference for Change Points in Functional Regression Time Series

In this work, we study the estimation and inference of change points under a functional linear regression model with changes in the slope function. We present a novel Functional Regression Binary Segmentation (FRBS) algorithm which is computationally efficient as well as achieving consistency in multiple change point detection. This algorithm utilizes the predictive power of piece-wise constant functional linear regression models in the reproducing kernel Hilbert space framework. We further propose a refinement step that improves the localization rate of the initial estimator output by FRBS, and derive asymptotic distributions of the refined estimators for two different regimes determined by the magnitude of a change. To facilitate the construction of confidence intervals for underlying change points based on the limiting distribution, we propose a consistent block-type long-run variance estimator. Our theoretical justifications for the proposed approach accommodate temporal dependence and heavy-tailedness in both the functional covariates and the measurement errors. Empirical effectiveness of our methodology is demonstrated through extensive simulation studies and an application to the Standard and Poor's 500 index dataset.

Shivam Kumar
Department of ACMS
University of Notre Dame
skumar9@nd.edu

Haotian Xu
Department of Statistics
University of Warrwick
haotian.xu.1@warwick.ac.uk

Haeran Cho
School of Mathematics
University of Bristol
haeran.cho@bristol.ac.uk

Daren Wang
Assistant Professor
Department of ACMS, University of Notre Dame

dwang24@nd.edu

MS24

Gaussian Random Fields on Riemannian Manifolds: Applications to Geostatistics

Many applications in spatial and spatio-temporal statistics require data to be modeled by Gaussian processes on non-Euclidean domains, or with non-stationary properties. Using such models generally comes at the price of a drastic increase in operational costs (computational and storage-wise), rendering them hard to apply to large datasets. In this talk, we propose a solution to this problem, which relies on the definition of a class of random fields on Riemannian manifolds. These fields extend ongoing work that has been done to leverage a characterization of the random fields classically used in Geostatistics as solutions of stochastic partial differential equations. The discretization of these generalized random fields, undertaken using a finite element approach, then provides an explicit characterization that is leveraged to solve the scalability problem. Indeed, matrix-free algorithms, in the sense that they do not require to build and store any covariance (or precision) matrix, are derived to tackle for instance the simulation of large Gaussian fields with given covariance properties, even in the non-stationary and/or spatio-temporal setting, or on surfaces.

Mike Pereira
Mines Paris
PSL University
mike.pereira@minesparis.psl.eu

MS24

Removing Bias from Fourier Methods in Signal Processing and Stochastic Processes

The Fourier Transform is a much-used tool across science and engineering. Within signal and image processing, it is often used to understand the frequency or wavenumber content of signals via the power spectral density, as well as speed up computation of statistical operations such as matrix inversion, and perform signal compression. With increasing big data sources, the variance of Fourier-based methods can be reduced by smoothing across multiple measurements, thus reducing the error in stochastic environments with noisy signals. In this talk I will show, however, that there is a real issue of statistical bias which will not always reduce with increasing signal sizes, and indeed will become the dominant source of error meaning that if the bias is ignored then one would become increasingly confident in something that is increasingly wrong as the signal lengths grow! I will show the prevalence and significance of this bias, and how it can be removed, in two contexts: 1) estimating the power spectral density in nonparametric estimates of the power spectral density, and 2) estimating parameters of stochastic processes in the Fourier domain using the Whittle Likelihood. The bias correction is designed to work with both spatial and time series data, as well as data that might be non-Gaussian, subject to missingness, or non-stationary. I will present motivating applications in oceanography and geosciences.

Adam Sykulski
Department of Mathematics
Imperial College London
adam.sykulski@imperial.ac.uk

MS25

The Akhiezer Iteration for Matrix Functions and Sylvester Equations

We present an iterative method for solving indefinite linear systems and computing matrix functions via orthogonal polynomial expansions. In particular, we utilize orthogonal polynomials for weight functions supported on multiple intervals that roughly correspond to the eigenvalues of the matrix in question. Such polynomials can be efficiently computed and utilized due to formulae of Akhiezer and a Riemann–Hilbert-based numerical method. The iterative method applies in settings where classical polynomial approximations behave poorly and are therefore not applicable. Applied to the matrix sign function, the method yields a fast inverse-free iteration for solving Sylvester matrix equations.

Cade Ballew
University of Washington
ballew@uw.edu

Thomas Trogdon
University of Washington
Department of Applied Mathematics
trogdon@uw.edu

Heather D. Wilber
University of Washington, U.S.
hdw27@uw.edu

MS25

Polynomial Preconditioning for the (Inverse) Matrix Square Root

While preconditioning is a long-standing concept to accelerate iterative methods for linear systems, generalizations to matrix functions are still in their infancy. We go a further step in this direction, introducing polynomial preconditioning for Krylov subspace methods which approximate the action of the matrix square root and inverse square root on a vector. Preconditioning reduces the subspace size and therefore avoids the storage problem together with—for non-Hermitian matrices—the increased computational cost per iteration that arises in the unpreconditioned case. Polynomial preconditioning is an attractive alternative to current restarting or sketching approaches since it is simpler and computationally more efficient. We demonstrate this for several numerical examples.

Andreas J. Frommer
Bergische Universitaet Wuppertal
Fachbereich Mathematik und Naturwissenschaften
frommer@uni-wuppertal.de

MS25

Polynomial Preconditioning for Indefinite Linear Equations Interior Eigenvalues

Indefinite spectra occur for linear equations in many applications. Examples include Helmholtz equations such as the wave equation and in quantum chromodynamics. Indefinite problems can be very difficult for Krylov iterative methods. We investigate adding polynomial preconditioning for such problems and show it can give tremendous improvement. Several difficulties can arise in polynomial

preconditioning for indefinite matrices. These are mentioned along with some algorithmic solutions. Eigenvalue problems will also be explored.

Hayden Henson
Baylor University
hayden_henson1@baylor.edu

MS25

Polynomial Approximation of the Inverse of a Matrix with Application to Multilevel Monte Carlo for QCD

A polynomial of A can give a good approximation to the inverse of A. For sparse matrices, this in some sense gives a sparse version of the inverse. For ill-conditioned matrices, a high degree polynomial may be needed. GMRES and Lanczos iterations can be used to develop this polynomial. Application is given to linear equations with multiple right-hand sides. Once a polynomial is developed, it can be applied to each right-hand side, and eigenvalue deflation can be included. The second application is for the trace of the inverse for large quantum chromodynamics matrices. Polynomial approximations are used to develop a Multilevel Monte Carlo method.

Ron Morgan
Baylor University
Ronald_Morgan@baylor.edu

MS25

Simple Cg and Bicg for Multiple Right-Hand Sides

A very simple approach to solving multiple right-hand side systems is proposed. For symmetric problems, the conjugate gradient method is a very efficient way to solve linear equations. We will use the same parameters from solving the first system for other systems. This corresponds to applying a polynomial approximation to the inverse of the matrix, and it requires no dot products. Deflation of eigenvalues using seeding can be included. The natural stability control of symmetric Lanczos will be discussed along with the possibility of additional stability control. For nonsymmetric problems, a similarly simple version of BiCG can be used, and it is naturally more stable than BiCGStab in this context.

Abby Williams
Baylor University
abby_williams5@baylor.edu

MS26

Optimization Approaches to Confidence Intervals in Inverse Problems

Confidence intervals defined as optimization programs date back to the 1960s, initially as tools to address the ill-posedness of inverse problems without introducing prior regularization. I will discuss the recently uncovered connection between those techniques and classical test inversion constructions and how it can help calibrate intervals for the original goal of ill-posed inverse problems. This insight has been vital in disproving the Burrus conjecture (1965) and opening the door for future improvements of optimization-based constructions

Pau Batlle
Caltech
pbatllef@caltech.edu

MS26

Bayesian Inference for Large Scale Inverse Problems Governed by Hyperbolic Dynamical Systems

We present effective numerical schemes for Bayesian inference in initial value control problems for diffeomorphic image registration. Our formulation is governed by a transport equation for the image intensities and the Euler-Poincaré equation associated with the group of diffeomorphisms. We present effective implementation for the evaluation of forward and adjoint operators. We use a Laplace approximation of the covariance matrix of the posterior distribution for uncertainty quantification. Our applications are in biomedical imaging.

Andreas Mang
University of Houston
Department of Mathematics
amang@Central.UH.EDU

MS26

Advancing Model Reduction Through Bayesian Inference

The high-dimensionality of the uncertain parameters and the high computational cost of forward models are two of the main challenges in solving large-scale Bayesian inverse problems. In this talk, Ill introduce parameter and state dimension reduction techniques and demonstrate, with a motivating system of non-Newtonian Stokes equations, how they can be used to address both of these difficulties. Parameter dimension is reduced by exploiting the problems structuresuch as data sensitivity and prior covarianceto identify a likelihood-informed parameter subspace that shows where the change from prior to posterior is most significant. For state dimension reduction, proper orthogonal decomposition (POD) and the discrete empirical interpolation method (DEIM) are used to approximate nonlinear terms in the forward model efficiently.

Preskella Mrad
University of Texas at Dallas
preskella.mrad@utdallas.edu

Sue Minkoff
University of Texas - Dallas
sminkoff@utdallas.edu

MS26

Bayesian Inversion of PDE-Governed Problems Using Integrated Nested Laplace Approximations

The core computational bottleneck in Bayesian inference is the cost of evaluating high-dimensional integrals. For problems which can be expressed as latent Gaussian models, the method of Integrated Nested Laplace Approximations (INLA) has recently shown promise as an alternative to sampling-based approaches such as MCMC, particularly in the area of spatial statistics. INLA has yet to be applied to inverse problems derived from PDE-governed systems, which pose additional difficulties: the forward operator becomes an expensive, non-local solve and precision matrices cannot be built and stored explicitly. We will demonstrate a proof-of-concept application of INLA to one such PDE-

governed inverse problem.

Sonia Reilly
Courant Institute for Mathematical Sciences
sonia.reilly@nyu.edu

Lisa Gaedke-Merzhäuser
Università della Svizzera italiana
lisa.gaedke.merzhaeuser@usi.ch

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS26

Bayesian Inference on Structural Integrity of Spent Nuclear Fuels

This paper addresses the challenges of efficient Bayesian inverse analysis for structural integrity assessments of spent nuclear fuels, focusing on high-dimensional parameter spaces and quantities of interest (QoIs). Key challenges include the computational burden of extensive forward model evaluations and the complexity of exploring high-dimensional parameter spaces. We propose a probabilistic surrogate model using polynomial chaos expansions (PCE), but PCE struggles with the curse of dimensionality and convergence issues in high-dimensional contexts. To overcome these limitations, we introduce an integrated approach that applies dimension reduction techniques to both input parameters and output QoIs. Specifically, we use truncated Karhunen-Love expansion (KLE) for QoIs and an accelerated basis adaptation algorithm for parameters, enhancing computational efficiency. Additionally, a nonparametric stochastic approach is incorporated to manage uncertainties from modeling errors. A block-update Markov Chain Monte Carlo (MCMC) algorithm is implemented to improve the acceptance rate in posterior sampling. The approach is validated through case studies on boiling water reactor spent nuclear fuel assemblies and fully-loaded spent nuclear fuel canisters, demonstrating its effectiveness in high-dimensional problems.

Xiaoshu Zeng, Bora Gencturk, Olivier Ezvan
University of Southern California
xiaoshuz@usc.edu, gencturk@usc.edu, ezvan@usc.edu

Roger Ghanem
University of Southern California
Civil and Environmental Engineering
ghanem@usc.edu

MS27

Fast Fourier Scheme for History Compression in Wave Scattering Problems

We propose a high-order accurate, fast algorithm to solve time-domain boundary integral equations (BIE) for the wave equation in scattering problems. Specifically, we consider problems where the spatial size can become sub-wavelength necessitating small time-steps a case where finite-difference time-domain (FDTD) methods may become inefficient. Splitting the solution representation into a history and a local part, we treat the history part using Fourier representation. We apply a Fast-Fourier transform (FFT) filtering method using the Kaiser-Bessel window function to compress the history computation up to $\mathcal{O}(n^d \log n)$, where d is the dimension and n is the temporal Fourier content of data. We present numerical examples where we consider the scattering from multiple springs on a 1D string, as well as 2D and 3D scattering from complicated geometry.

Nour G. Al Hassanieh
Rensselaer Polytechnic Institute
nalhassanieh@flatironinstitute.org

MS27

A Fast Solver for Band Structure Calculations

Band structure calculations are a ubiquitous task in materials engineering. Mathematically, the problem amounts to the efficient solution of the Schrdinger equation with a periodic potential. In this talk, I will describe a new direct solver to efficiently compute band structures.

Abinand Gopal
UC Davis
gopal@ucdavis.edu

MS27

Finding Scattering Resonances via Generalized Colleague Matrices

We present a scheme for finding all roots of an analytic function in a square domain in the complex plane. The scheme can be viewed as a generalization of the classical approach to finding roots of a function on the real line, by first approximating it by a polynomial in the Chebyshev basis, followed by diagonalizing the so-called colleague matrices. Our extension of the classical approach is based on several observations that enable the construction of polynomial bases in compact domains that satisfy three-term recurrences and are reasonably well-conditioned. This class of polynomial bases gives rise to "generalized colleague matrices," whose eigenvalues are roots of functions expressed in these bases. We also introduce a special-purpose QR algorithm for finding the eigenvalues of generalized colleague matrices, which is a straightforward extension of the recently introduced structured stable QR algorithm for the classical cases. Furthermore, by coupling this scheme to integral equation solvers, we demonstrate its effectiveness in locating resonances and reflectionless states in wave scattering.

Hanwen Zhang
Yale University
hanwen.zhang@yale.edu

MS28

Sensitivity-Driven Surrogate Model Refinement for Optimization of Dynamical Systems with Expensive Black-Box Functions

In many applications one must compute an optimal trajectory for a dynamical system where some components of the dynamics depend on a computationally expensive high-fidelity model. Rather than solving these problems directly, it is much faster and cheaper to solve a perturbed optimization problem where the high-fidelity model is replaced by a computationally inexpensive surrogate model constructed from evaluations of the high-fidelity model. The difference between the solution of the original problem and that of the perturbed problem depends on surrogate errors and the sensitivity of the solution of the optimization problem to perturbations in the model. To make this difference smaller, one may perform additional evaluations

of the high-fidelity model to obtain a more accurate surrogate, but these evaluations must be chosen intelligently. In this talk, I propose a framework for solving optimal control problems with expensive black-box functions using adaptively refined surrogate models constructed via interpolation in reproducing kernel Hilbert spaces, where the model refinement procedure leverages pointwise error bounds for surrogates and sensitivity analysis of the optimization problem to select the best locations to evaluate the high-fidelity model. I use this framework to solve a trajectory optimization problem for a hypersonic vehicle where the lift and drag coefficients are computed from expensive computational fluid dynamics simulations.

Jonathan Cangelosi
Rice University
jrc20@rice.edu

MS28

Dimension-Free Surrogate Models of Partial Differential Equations

Partial differential equations (PDEs) are essential when it comes to modeling physical phenomena such as wave propagation, heat diffusion and fluid dynamics. Most of these equation are defined in one or more spatial dimensions and can be expressed in a dimension-free form, that is, in a form that does not explicitly depend on the number of dimensions. We propose a surrogate-modeling technique that is–by design–independent from the dimensionality of the underlying PDE.

Hans Harder
Uni Paderborn
hans.harder@uni-paderborn.de

Sebastian Peitz
Paderborn University
sebastian.peitz@uni-paderborn.de

MS28

Asymptotic Stability of Koopman-Based Data-Driven Predictive Control

In this talk, we illustrate how data-driven techniques for learning dynamical control systems, such as Extended Dynamic Mode Decomposition (EDMD) embedded in the Koopman framework, may be used in optimization-based controllers while preserving well-established guarantees such as closed-loop asymptotic stability. To this end, we will leverage novel proportional bounds on the learning error of EDMD, that is, error bounds that vanish in the origin, to show (practical) asymptotic stability of EDMD-based Model Predictive Control. Therein, the proportional bounds are key to show that important system-theoretic properties, such as cost controllability, carry over to the data-driven model, if a sufficient amount of samples is chosen. This talk is based on joint work with Lea Bold, Friedrich Philipp and Karl Worthmann (TU Ilmenau) and Lars Grne (University of Bayreuth).

Manuel Schaller
Technische Universität Ilmenau, Germany
manuel.schaller@mathematik.tu-chemnitz.de

Lea Bold
Technische Universtität Ilmenau
lea.bold@tu-ilmenau.de

Lars Grüne
Mathematical Institute
Universitaet Bayreuth
lars.gruene@uni-bayreuth.de

Friedrich Philipp
Technische Universität Ilmenau
friedrich.philipp@tu-ilmenau.de

Karl Worthmann
University of Bayreuth, Germany
karl.worthmann@tu-ilmenau.de

MS29

A Reaction-Diffusion Model of Sperm-Egg Interaction for Fertilization Success

Fertilization consists of spermatozoa motion towards an oocyte, followed by acrosome reaction and fusion. Acrosome reaction occurs between sperm head and the zona pellucida (ZP). As a sperm head moves through the ZP, acrosome reactions hydrolyze proteins. Sperm head drives concentration gradients through the ZP leading to eventual fusion of sperm nucleus with egg cytoplasm and fertilization. We construct a reaction-diffusion-decay model yielding in-silico insight into the underlying competing mechanisms. We derive an exact analytic solution for the purely reaction model and appropriate bounds on the full model. We demonstrate the model performance using numerical simulations and outline the next steps to incorporate glycan interactions.

Prajakta Bedekar
Johns Hopkins University
National Institute of Standards and Technology
pbedeka1@jhu.edu

Anthony Kearsly
National Institute of Standards and Technology
Applied and Computational Mathematics Division
anthony.kearsley@nist.gov

Daniel M. Anderson
Department of Mathematical Sciences
George Mason University
danders1@gmu.edu

MS29

A Novel Computational Method for Maximizing Sensitivity in Early Disease Detection

In the cancer early detection field, logistic regression is a frequently used approach to establish a combination rule that differentiates cancer from non-cancer. However, the application of logistic regression relies on a maximum likelihood approach, which may not yield optimal combination rules for maximizing sensitivity at a clinically desirable specificity and vice versa. Here, we have developed an improved regression framework, Sensitivity Maximization At a Given Specificity, SMAGS, for binary classification that finds the linear decision rule yielding the maximum sensitivity for a given specificity or the maximum specificity for a given sensitivity. We additionally expand the framework for feature selection that satisfies sensitivity and specificity maximizations. We compare our SMAGS method with normal logistic regression using two synthetic datasets and reported data for colorectal cancer (CRC) from the 2018 CancerSEEK study. In the CRC CancerSEEK dataset, we

report 14% improvement in sensitivity at 98.5% specificity (0.31 vs 0.57; p-value:¡0.05). The SMAGS method provides an alternative to logistic regression for modeling combination rules for biomarkers and early detection applications.

Seyyed Mahmood Ghasemi, Chunhui Gu, Johannes Fahrmann, Samir Hanash, Kim-Anh Do, James Long, Ehsan Irajizad
MD Anderson Cancer Center
smahmoodghasemi@gmail.com, chunhui.gu@uth.tmc.edu, jffahrmann@mdanderson.org, shanash@mdanderson.org, kimdo@mdanderson.org, jplong@mdanderson.org, eirajizad@mdanderson.org

MS29

Topological Data Analysis of Molecular Dynamics Simulations of Monoclonal Antibodies

A variety of medical treatments depend on artificially produced monoclonal antibodies. Measuring monoclonal antibody structure is very important due to their widespread use in biomedical applications. Similar amino acid sequences can generate biomolecules which adopt different shapes in 3-dimensional space. Molecular dynamics simulations are a valuable tool for revealing potential arrangements of the atoms in these proteins. By quantifying the structure present in each frame of a simulation, we can search for variations in the configuration of the molecule. Topological data analysis detects and quantifies structural features which are not easily measured by classical data analysis techniques. We will discuss our results from using topological data analysis to explore molecular dynamics simulations of monoclonal antibodies, with a focus on the NIST monoclonal antibody (NISTmAb) reference material.

Melinda Kleczynski, Christina Bergonzo
National Institute of Standards and Technology
melinda.kleczynski@nist.gov, christina.bergonzo@nist.gov

Anthony J. Kearsley
Applied and Computational Mathematics Division
National Institute for Standards and Technology
anthony.kearsley@nist.gov

MS29

Exploring Hepatitis B Dynamics: A Homotopy Approach to Fractional Modeling with Uncertainty Quantification

Mathematical models are frequently used to better understand the dynamics governing viral infection spread. This research introduces a fuzzy fractional-order model of the hepatitis B virus (HBV) to address the inherent uncertainty in biological systems. In this model, the dynamics of infection and the removal of infected cells are studied employing a double parametric-based homotopy approach. The existence and uniqueness of the model are also investigated. Results indicate that incorporating a fractional-order derivative with fuzzy logic reduces peak viral load and minimizes cellular damage, mainly when early treatment is applied, although complete eradication may take longer. These findings contribute to our comprehension of the progression of HBV through various infection stages, as well as to the prediction of disease outcomes and the optimization of clinical management strategies.

Ajay Kumar
Sardar Vallabhbhai National Institute of Technology
ajaykhator123@gmail.com

Ramakanta Meher
Sardar Vallabhbhai National Institute of Technology
Surat-395007, Gujarat, India
meher_ramakanta@yahoo.com

MS29

Time-Dependent Antibody Kinetics of Infected and Vaccinated Individuals: a Graph-Theoretic Approach

Modeling the deterioration of antibody levels is paramount to understanding the time-dependent viral response to infections or vaccinations. These events have been studied experimentally, but also benefit from a rigorous mathematical underpinning. Disease or vaccination prevalence in the population and time-dependence on a personal scale simultaneously affect antibody levels, interact non-trivially, and pose considerable modeling challenges. We propose a time-inhomogeneous Markov chain model for event-to-event transitions coupled with a probabilistic framework for post-event antibody kinetics. This approach is ideal to model sequences of infections and vaccinations, or personal trajectories in a population. We demonstrate the modeling process as well as estimation of transition probabilities. This work is an important step towards a comprehensive understanding of antibody kinetics that will allow us to simulate and predict real-world disease response scenarios.

Rayanne Luke
George Mason University
rluke@gmu.edu

Prajakta Bedekar
Johns Hopkins University
National Institute of Standards and Technology
pbedeka1@jhu.edu

Anthony J. Kearsley
Applied and Computational Mathematics Division
National Institute for Standards and Technology
anthony.kearsley@nist.gov

MS30

Transport-Based Variational Bayesian Inference

We are investigating transport-based methods for variational Bayesian inference. Our approach transforms particles from a simple prior distribution to match a target posterior distribution by minimizing the Kullback-Leibler divergence. We establish the theoretical foundations, deriving the optimal transport map through variational principles. Our implementation uses advanced numerical techniques to efficiently move particles toward the target distribution. We provide detailed derivations and visualize the particle distribution's evolution over multiple iterations, illustrating the convergence process. To enhance efficiency, we explore various numerical optimization techniques. These improve convergence rates and overall algorithm performance. Our empirical study evaluates the method's performance on test problems of varying complexity, analyzing convergence rates, iteration counts, approximation errors, and computational run times. This research contributes to developing efficient variational Bayesian inference methods, offering insights into practical implementation and performance characteristics of transport-based

approaches.

Pegah Amiri
University of Houston
pamiri@CougarNet.UH.EDU

Andreas Mang
University of Houston
Department of Mathematics
amang@Central.UH.EDU

MS30

Bayesian Variational System Identification

We focus on the problem of system identification that seeks to determine, from observation data, operators present in a governing partial differential equation (PDE) and values of their coefficients. More specifically, we use Variational System Identification (VSI) that approaches the task through a weak-form formulation of the PDE. An important challenge to VSI, and system identification in general, is the presence of uncertainty arising from observation noise and model misspecification—all hindering an accurate system identification result. It is therefore important to incorporate uncertainty quantification (UQ) into VSI. We introduce Bayesian VSI, a framework to perform parameter inference and model selection under uncertainty to achieve system identification. We enable parameter inference for a given candidate PDE, where we (1) establish an effective likelihood by propagating the observation noise to the PDE residual and thereby sidesteps the need for repeated PDE solves, and then (2) employ Stein Variational Gradient Descent (SVGD) to efficiently find an approximate parameter posterior distribution. Then, we perform model selection to compare and select among different PDE candidates via Bayesian Information Criterion and a decision-theoretic posterior-expected utility. We demonstrate these methods on examples including a linear diffusion-advection problem and a non-linear Poisson equation..

Chengyang Huang
University of Michigan
chengyah@umich.edu

Siddhartha Srivastava
University of Southern California
Aerospace and Mechanical Engineering
ss79128@usc.edu

Krishna Garikipati
Aerospace & Mechanical Engineering
University of Southern California
garikipa@usc.edu

Xun Huan
University of Michigan
xhuan@umich.edu

MS30

A Control-Oriented Approach to Optimal Sensor Placement

We propose a control-oriented optimal experimental design (cOED) approach for linear PDE-constrained Bayesian inverse problems. In particular, we consider optimal control problems with uncertain parameters that need to be estimated by solving an inverse problem, which in turn requires measurement data. Specifically, we consider the case where data is collected at a set of sensors. While classical Bayesian OED techniques provide experimental designs (sensor placements) that minimize the posterior uncertainty in the inversion parameter, these designs are not tailored to the demands of the optimal control problem. In the present control-oriented setting, we prioritize the designs that minimize the uncertainty in the state variable being controlled or the control objective. We propose a mathematical framework for uncertainty quantification and cOED for parameterized PDE-constrained optimal control problems with linear dependence to the control variable and the inversion parameter. We also present scalable computational methods for computing control-oriented sensor placements and for quantifying the uncertainty in the control objective. Additionally, we present illustrative numerical results in the context of a model problem motivated by heat transfer applications.

Madhusudan Madhavan
North Carolina State University
mmadhav@ncsu.edu

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Bart G. Van Bloemen Waanders, Rebekah White
Sandia National Laboratories
bartv@sandia.gov, rebwhit@sandia.gov

MS30

Probabilistic Forecasting of Multivariate Time-series with Attentive Transport Maps

The stochastic control problems that arise throughout science, engineering, and finance provide a powerful mechanism for mapping forecasts of a system's response into optimal decisions. However, the quality of these decisions is directly related to our ability to predict the system's response. Within the stochastic control setting, this work considers the problem of probabilistically forecasting multivariate time series given other time-dependent quantities and deterministic forecasts. We develop a hybrid approach that blends the ability of polynomial-based transport maps to parsimoniously characterize probability distributions with the ability of transformer architectures to capture complex spatio-temporal relationships. We evaluate this approach in the context of multiple energy applications and give particular attention to the approach's data-efficiency and ability to quickly generate forecast samples in a stochastic control setting.

Matthew Parno
Solea Energy
parnomd@gmail.com

MS30

Data-Driven Parameterization Refinement for the Structural Optimization of Cruise Ship Hulls

The shipbuilding industry is facing multiple challenges due to the rise in environmental consciousness, such as reducing operational costs, adopting new engine technologies and reducing building materials, while adhering to complex structural stability regulations and manufacturing constraints.

We present a computational pipeline for the structural optimization of cruise ships [1, 2], with multi-objective optimization based on genetic algorithms and a hierarchical parameterization refinement procedure based on clustering of hull responses. The pipeline is highly automated and able to efficiently optimize ship hulls starting from very coarse parameterizations. The optimization has been tested on a simplified hull and a full ship, achieving considerable reduction of the total mass and proving effective in streamlining the design phase. [1] Tezzele, M., Fabris, L., et al. A multifidelity approach coupling parameter space reduction and nonintrusive POD with application to structural optimization of passenger ship hulls. International Journal for Numerical Methods in Engineering, 124(5), 1193-1210, 2023. [2] Fabris, L., Tezzele, M., et al. Data-driven parameterization refinement for the structural optimization of cruise ship hulls. In preparation.

Lorenzo Fabris
SISSA, International School for Advanced Studies
Trieste, Italy
lfabris@sissa.it

Marco Tezzele
Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
marco.tezzele@emory.edu

Marco Tezzele
Emory University
marco.tezzele@emory.edu

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS31

Optimization of the Boltzmann Equation

The kinetics of rarefied gases and plasmas are described by the Boltzmann equation and numerically approximated by the Direct Simulation Monte Carlo (DSMC) method. We present an optimization method for DSMC, derived from an augmented Lagrangian. After a forward (in time) solution of DSMC, adjoint variables are found by a backwards solver. They are equal to velocity derivatives of an objective function, which can then be optimized. This is joint work with Yunan Yang (Cornell) and Denis Silantyev (U Colorado, Colorado Springs).

Russel Caflisch
New York University
caflisch@courant.nyu.edu

MS31

A Multiscale Hybrid Maxwellian-Monte-Carlo Coulomb Collision Algorithm for Particle Simulations

In many plasma systems, Coulomb collisions often exhibit multiscale temporal behavior. Traditional Monte Carlo (MC) methods [Takizuka & Abe, JCP, (1977) 25: 205-219] have a timestep accuracy constraint $\nu\Delta t \ll 1$ to resolve the collision frequency (ν) effectively [Dimits et al, JCP (2009) 228: 4881-4892]. This constraint becomes particularly stringent with high-Z species or inter-species collisions with substantial mass disparities, rendering such simulations extremely expensive or impractical. We propose a multiscale hybrid particle-Maxwellian collisional algorithm that utilizes Maxwellians (i.e., isotropic Gaussians) [Echim et al., Surveys in Geophysics (2011) 32:1-70] for highly collisional kinetic species and/or fluid components (e.g., in hybrid codes). We employ an enhanced Lemons method [Lemons et al. JCP, (2009) 228(5), pp.1391-1403] for particle-Maxwellian collisions. Additionally, we introduce a new scheme that extends the standard TA method for arbitrary particle weights without compromising conservation properties. The overall scheme is strictly conservative and may significantly outperform standard MC methods, with orders of magnitude improvement in computational efficiency. We will substantiate the accuracy and performance of the proposed method through several examples of varying complexity, encompassing both relaxation and transport problems.

Luis Chacon, Guangye Chen, Adam Stanier, Steven Anderson, Bobby Philip
Los Alamos National Laboratory
chacon@lanl.gov, gchen@lanl.gov, stanier@lanl.gov, andeste@lanl.gov, bphilip@lanl.gov

MS31

Stochastic Interacting Particle Field Methods in the Computation of Chemotaxis and Haptotaxis

In this talk, I will present the most recent developments of the stochastic interacting particle field (SIPF) methods. The methodology originates from the Lagrangian framework in the computation of asymptotic behaviors of PDEs, for instance the effective diffusivities and KPP front speed. Then we turns to the interacting particle methods for elliptic type Keller Segel equations. Very recently, we further extend our research to SIPF which facilitates the computation of parabolic type Keller Segel and haptotaxis equations. If time permitted, I will also discuss how these results 'interacts' with the field of generative modeling.

Boyi Hu
The University of Hong Kong
huby22@connect.hku.hk ¡huby22@connect.hku.hk¿

Zhongjian Wang
Nanyang Technological University
zhongjian.wang@ntu.edu.sg

Jack Xin
UC, Irvine
jackx@uci.edu

Zhiwen Zhang
Department of Mathematics,
The University of Hong Kong,
zhangzw@hku.hk

MS32

The Picard Newton Iteration for the Boussinesq Equations

We study the Picard-Newton iteration for the incompressible Boussinesq equations, which is a two-step iteration resulting from the composition of the Picard and Newton iterations. We prove that this iterative method retains Newtons quadratic convergence but has less restrictive sufficient conditions for convergence than Newton and also is unconditionally stable under a small data condition. In this sense, Picard-Newton can be considered as a Newton iter-

ation that is nonlinearly preconditioned with Picard. Our numerical tests illustrate this quadratic convergence and stability on benchmark problems. Furthermore, the tests show convergence for significantly higher Rayleigh number than both Picard and Newton, which illustrates the larger convergence basin of Picard-Newton that the theory predicts. We also introduce Ander- son acceleration into the Picard step in our Picard-Newton numerical tests, and this enables convergence for even higher Rayleigh number.

Elizabeth Hawkins
Clemson University
evhawki@clemson.edu

MS32

A Numerical Study on Anderson Acceleration and GMRES Applied to the Stokes-Darcy Equations

In this talk, we consider the numerical solution of the Stokes-Darcy equations. The model is discretized using marker-and-cell scheme, leading to a large system with a saddle-point structure. In our previous work, we proposed several block-triangular preconditioners for Krylov subspace methods. Here, we consider using Anderson acceleration with finite window size to solve the preconditioned system. We compare the performance of GMRES and Anderson acceleration for different values of physical parameters. This will help us better understand in which situations Anderson acceleration could outperform GMRES. We also investigate restarted GMRES and restarted Anderson acceleration. Numerical results show that for small values of physical parameters, Anderson acceleration outperforms both GMRES and restarted Anderson acceleration in terms of iteration numbers.

Yunhui He
Department of Mathematics
University of Houston
yhe43@central.uh.edu

MS32

Accelarating with Crop Method

Anderson Acceleration is a well-established method that allows to speed up or encourage convergence of fixed-point iterations. It has been successfully used in a variety of applications, in particular within the Self-Consistent Field (SCF) iteration method for quantum chemistry and physics computations. In recent years, the Conjugate Residual with OPtimal trial vectors (CROP) algorithm was introduced and shown to have a better performance than the classical Anderson Acceleration with less storage needed. This talk explores the intricate connections between the classical Anderson Acceleration method and the CROP algorithm. Our objectives include a comprehensive study of their convergence properties, explaining the underlying relationships, and substantiating our findings through some numerical examples. Through this exploration, we contribute valuable insights that can enhance the understanding and application of acceleration methods in practical computations, as well as the development of new and more efficient acceleration schemes.

Ning Wan, Agnieszka Miedlar
Virginia Tech
wning@vt.edu, amiedlar@vt.edu

MS33

High-Order IFGF Green function acceleration with optimized kernel factorizations and geometric adaptivity

We present a fast high-order integral equation algorithm for the solution of scattering problems. Based on the recently developed IFGF (Interpolated Factored Green Function) acceleration method, this presentation demonstrates a number of significant algorithmic innovations, including high-order Green function integration at essentially vanishing memory cost, high-order interpolated acceleration, modified Green-function factorizations, and geometry adaptivity. The character and overall performance of the new algorithm, in terms of accuracy, speed and effective parallelization, will be demonstrated by means of a variety of numerical results for realistic complex engineering structures.

Sebastian Lamas, Oscar P. Bruno
Caltech
slamas@caltech.edu, obruno@caltech.edu

MS33

Finite Element Operator Network for Solving Forward and Inverse Parametric Pdes

This talk introduces the Finite Element Operator Network (FEONet), a novel deep learning framework designed to solve both forward and inverse parametric partial differential equations (PDEs). By integrating finite element methods with operator learning, FEONet effectively handles complex parametric dependencies, providing accurate and efficient solutions. The proposed approach is tested on various benchmark problems, demonstrating its potential to significantly outperform traditional numerical methods in terms of computational efficiency and generalization capability across different parameter spaces.

YoungJoon Hong
Seoul National University
hongyj@snu.ac.kr

Jaeyong Lee
Chung-Ang University, Korea
jaeyong@cau.ac.kr

MS33

Compatible Energy Preserving Discretizations for Nonlinear Optical Wave Propagation: The Maxwell-Duffing Approach

This talk explores the modeling and numerical discretization of Maxwells equations in nonlinear optical media, specifically focusing on the Maxwell- Duffing model. We present the constitutive laws governing electromag- netic wave propagation in non-magnetic, non-conductive media, describ- ing the materials response using a nonlinear cubic Duffing model coupled with Maxwells equations. The presentation includes the derivation of energy relations for the one-dimensional nonlinear Maxwell model. We introduce a high-order spatial discretization method based on fully discrete leap-frog finite-difference time-domain (FDTD) methods, designed for the accurate and stable simulation of nonlinear wave propa- gation. Numerical simulations highlight the effectiveness of these methods in capturing

the complex dynamics of electromagnetic waves in nonlinear media. Special attention is given to the implementation of the Second or- der in time and higher order in space leap-frog scheme and its application to traveling wave solutions. We prove Energy Stability of the Higher Order Yee FDTD Schemes for the cubic Maxwell-Duffing Model and demonstrate these results through Numerical Simulations. This work provides critical insights into the mathematical and compu- tational challenges of modeling nonlinear optical materials, offering robust techniques for advancing research in nonlinear photonics.

Victory Obieke
Oregon State University, U.S.
obiekev@oregonstate.edu

Vrushali Bokil
Oregon State University
vrushali.bokil@oregonstate.edu

MS33

Manipulating the Propagation of Plasmons on a Graphene Sheet

In this talk, we discuss the propagation of plasmons on graphene with a space-time perturbation in the Drude weight. The graphene is modeled as a 1-dimensional conductive sheet in a 2-D medium where the EM field obeys Maxwells equations. The current density on the sheet is governed by Drudes law. The system of equations and boundary conditions can be rewritten as a single integro-differential equation. The equation gives the current density on the graphene as a function of the position and time parameters. We discuss methods to obtain a closed-form solution to the integro-differential equation, as well as numerical experiments to present a perturbation analysis for various wave types.

Fadil Santosa
Johns Hopkins University, U.S.
fsantos9@jhu.edu

MS34

A Structure-Preserving Discretization for Open Fluids

Many real fluid dynamical systems are open: they have boundaries where mass, momentum, energy and other quantities are exchanged with an environment. Correctly capturing these exchanges is critical to accurate simulation of such systems. Recently, a geometric mechanics (GM) formulation for open fluids was developed, that is capable of treating the general case of arbitrary boundary conditions. In this talk, I will discuss a structure-preserving discretization of this new GM formulation, specifically a discrete exterior calculus (DEC) approach applied to the Hamiltonian version. In particular, the new discretization has discrete analogues of key properties of the continuous formulation, such as involution constraints and conservation laws; along with physically realistic transport operators. The new method will be demonstrated for a variety of examples, including the (thermal) shallow water equations and the compressible Euler equations.

Chris Eldred
Sandia National Laboratory
celdred@sandia.gov

MS34

Metriplectic Dynamics: Using the 4-bracket for Constructing Thermodynamically Consistent Models and Structure Preserving Numerical Algorithms

Metriplectic dynamics originates in my early works of the 1980s, which in recent times led to the discovery of the metriplectic 4-bracket. The metriplectic 4-bracket is akin to the Poisson bracket in that it is a multilinear operator that satisfies the Leibniz rule and acts on phase space functions, which is used to generate dynamics. However, it acts on four such functions rather than two, possess different symmetry properties, and requires a Hamiltonian and an entropy function to generate dynamics. The symmetry properties are reminiscent of those of a Riemannian curvature tensor; from these symmetries conservation of energy and the production of entropy are guaranteed. In this talk I will describe how the 4-bracket formalism leads to an algorithm for constructing complicated thermodynamically consistent systems, like those for the Cahn-Hilliard-Navier-Stokes system for two phase flow and collision operators on a noncanonical phase space. The formalism also provides a way of obtaining semi-discrete projections of complicated PDEs to finite-dimensional thermodynamically consistent systems. In this was we can obtain metriplectic integrators. Recent progress on numerical results will be described.

Philip J. Morrison
Physics Department
University of Texas at Austin
morrison@physics.utexas.edu

MS34

Coupled Lie-Poisson Neural Networks (CLPNets): Data-Based Computing of Coupled Hamiltonian Systems

Physics-Informed Neural Networks (PINNs) have received much attention recently due to their potential for high-performance computations for complex physical systems. The idea of PINNs is to approximate the equations and boundary and initial conditions through a loss function for a neural network. PINNs combine the efficiency of data-based prediction with the accuracy and insights provided by the physical models. However, applications of these methods to predict the long-term evolution of systems with little friction, such as many systems encountered in space exploration, oceanography/climate, and many other fields, need extra care as the errors tend to accumulate, and the results may quickly become unreliable. We provide a solution to the problem of data-based computation of Hamiltonian systems utilizing symmetry methods, paying special attention to systems that come from the discretization of continuum mechanics systems. For example, for simulations, a continuum elastic rod can be discretized into coupled elements with dynamics depending on the relative position and orientation of neighboring elements. For data-based computing of such systems, we design the Coupled Lie-Poisson neural networks (CLPNets). We consider the Poisson bracket structure primary and require it to be satisfied exactly, whereas the Hamiltonian, only known from physics, can be satisfied approximately. By design, the method preserves all special integrals of the bracket (Casimirs) to machine precision. We present applications of CLPNets applications for several particular cases, such as coupled rigid bodies or elastically connected elements.

CLPNets yield surprising robustness for increasing the dimensionality of the system, enabling the computing of dynamics for a high number of dimensions (up to 18) using networks with a small number of parameters (one to two hundred) and only one to two thousand data points used for learning. Joint work with Chris Eldred (Sandia National Laboratory) and Francois Gay-Balmaz (NTU Singapore).

Vakhtang Putkaradze
University of Alberta
putkarad@ualberta.ca

MS34

Hamiltonian Interaction of Closed Vortex Sheets and Vortex Patches

The talk will describe a Hamiltonian model of the complex dynamic interaction between a closed vortex sheet, separating an ideal fluid with two different densities, and a vortex patch located either inside or outside the closed sheet. The model, without the Hamiltonian formalism, was presented in a previously published journal paper co-authored with Rangachari Kidambi. Following an introduction to vortex sheets and vortex patches, the talk will place this model in the framework of Hamiltonian fluid mechanics, incorporating Zakharovs canonical Possion brackets for the free surface water wave problem and the Lie-Poisson brackets for incompressible fluid flows.

Banavara Shashikanth
New Mexico State University
Department of Mechanical Engineering
shashi@nmsu.edu

MS35

Dean Kawasaki Equation: Analysis and Hybrid Methods

The dynamics of a system of non-interacting random-walker particles can, under an appropriate mathematical framework, be formally described by the Dean-Kawasaki equation. We will introduce nonlinear and non-Gaussian models that approximate the evolution of the empirical density of a given particle system, and we will discuss their well-posedness along with properties such as positivity and mass preservation. This is a joint work with N. Perkowski and H. Kremp. Additionally, we aim to study the numerical approximation of these types of nonlinear SPDEs. However, the standard finite volume approximation implicitly requires a sufficiently large number of particles to ensure both the positivity of the solution and an accurate approximation of the stochastic flux. To address this challenge, we extend hybrid algorithms for particle systems to cases where the density is low. We develop criteria for determining the threshold by comparing higher-order statistics from the finite volume method with particle simulations. We then demonstrate the use of these criteria for dynamic adaptation in both two- and three-dimensional spatial settings. This is a joint work with J. Bell and A. Almgren.

Ana Djurdjevac
Freie Universität Berlin
adjurdjevac@zedat.fu-berlin.de

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
Lawrence Berkeley National Laboratory, U.S.
jbbell@lbl.gov

Nicolas PerKowski
Institut f ur Mathematik
Freie Universit at Berlin
perkowski@math.fu-berlin.de

Helena Kremp
TU Vienna
helena.kremp@asc.tuwien.ac.at

MS35

Incorporating the Langmuir Adsorption Model into Compressible Fluctuating Hydrodynamics

For the computational modeling of gas-solid interfacial systems (e.g. heterogeneous catalysts), which plays an important role in various fields, such as energy and environmental sciences, the accurate and computationally efficient description of both the reactive dynamics on the catalytic surface and the transport dynamics in the gas phase is required. Furthermore, since fluctuations in both phases are significant at the mesoscale, the correct description of fluctuations in each phase and across the interface is critically important to develop a robust and reliable mesoscopic simulation methodology. For a gas-solid interfacial system, where some chemical species undergo reversible adsorption, we propose a modeling approach for incorporating the Langmuir adsorption model into compressible fluctuating hydrodynamics (FHD). To this end, we derive a thermodynamically-consistent energy update scheme. By performing a stochastic analysis of covariances for the linearized system, we analytically confirm that our update scheme can attain the thermodynamic equilibrium predicted by equilibrium statistical mechanics. Using the simulation codes "FHDeX", we also validate our overall scheme and implementations. As a numerical example, we consider an ideal gas mixture of CO and Ar with CO undergoing reversible adsorption. We also perform simulation studies using two non-equilibrium situations: an initially clean surface and two walls with different temperatures.

Hyun Tae Jung
Korea Advanced Institute of Science and Technology
chemtae@kaist.ac.kr

Changho Kim
University of California, Merced
Applied Mathematics
ckim103@ucmerced.edu

Hyungjun Kim
Korea Advanced Institute of Science and Technology
Department of Chemistry
linus16@kaist.ac.kr

Ishan Srivastava
Lawrence Berkeley National Laboratory
isriva@lbl.gov

Alejandro L. Garcia
San Jose State University
Alejandro.Garcia@sjsu.edu

Andy Nonaka
Lawrence Berkeley National Laboratory

Center for Computational Sciences and Engineering
ajnonaka@lbl.gov

John B. Bell
Lawrence Berkeley National Laboratory, U.S.
jbbell@lbl.gov

MS35

Modelling Rare Rupture and Adhesion of Nanoscale Thin Films with Fluctuating Hydrodynamics

Nanoscale thin liquid films are crucial in applications like solar cell manufacturing, insulation layer coating, and mediating cell-cell interactions. At this scale, random particle movements (Brownian motion) become significant, requiring fluctuating hydrodynamics to describe fluids. To simplify modeling, the high aspect ratio of thin films is exploited, and the lubrication approximation is applied to derive stochastic thin film equations (STFE), describing film height evolution. In this talk, we demonstrate STFE usage in two scenarios: i) rare rupture of thin films on solid substrates, ii) adhesion and phase separation of cell membrane adhesion patches. We solve STFE numerically with free surface boundary conditions and Van der Waals forces, observing film ruptures in the linearly stable regime induced by thermal fluctuations. The STFE can be rearranged into a gradient flow form, allowing for the application of rare event theory to predict average rupture time theoretically. Molecular dynamics simulations show good agreement with numerical and theoretical results. For cell membranes modeled as elastic bending sheets, STFE describes the fluid flow between them. Numerical solutions of STFE reveal that for membrane adhesion, thermal fluctuations must bring membranes close enough for protein bonds to form, aligning with rare event theory. Coarsening is observed during adhesion patch phase separation, revealing a different coarsening rate power law compared to active model B.

Jingbang Liu, Vira Dhaliwal
University of Oslo
jingbanl@math.uio.no, vira@math.uio.no

Duncan Lockerby
University of Warwick
duncan.lockerby@warwick.ac.uk

Tobias Grafke
Warwick Mathematics Institute
University of Warwick
T.Grafke@warwick.ac.uk

James E. Sprittles
University of Warwick
j.e.sprittles@warwick.ac.uk

Andreas Carlson
University of Oslo
acarlson@math.uio.no

MS35

Exploratory Computation of Statistical Solutions to Incompressible Fluid Flows

We present stochastic lattice Boltzmann methods (LBMs) for efficiently approximating statistical solutions to the incompressible Navier–Stokes and Euler equations in three spatial dimensions. Space-time adaptive kinetic relaxation frequencies are used for stable and consistent numerical solutions with decreasing viscosity. With single level Monte Carlo LBMs and stochastic Galerkin LBMs we approximate responses from random perturbations of the initial flow field. The schemes are implemented in the parallel C++ data structure OpenLB and executed on high-performance computing machines. In exploratory computations, we observe the expected scaling of the energy spectra and the structure functions in terms of Kolmogorov's K41 theory. We compute along the proven inviscid limit of statistical solutions of the incompressible Navier–Stokes equations toward weak-strong unique statistical solutions of the incompressible Euler equations in three dimensions. Convergence in the Wasserstein metric approves the validity of our new approach to computing statistical solutions of incompressible fluid flow models. Extensions of the methodology with machine learning and toward initial-boundary value problems of turbulent wall-bounded fluid flows are discussed.

Stephan Simonis, Mathias J. Krause
Karlsruhe Institute of Technology
stephan.simonis@kit.edu, mathias.krause@kit.edu

MS35

Fluctuating Hydrodynamics of Homogeneous Isotropic Compressible Turbulence

Using a recently developed staggered-grid method for fluctuating hydrodynamics of compressible fluids, we investigate the effects of thermal fluctuations on homogeneous, isotropic compressible turbulence in the nonlinear subsonic regime. We show that molecular fluctuations not only modify energy spectra at wavelengths larger than the Kolmogorov length scales in compressible turbulence, but also inhibit spatio-temporal intermittency across the entire dissipation range. Using large-scale direct numerical simulations of computational fluctuating hydrodynamics, we demonstrate that the extreme intermittency characteristic of turbulence models is replaced by nearly-Gaussian statistics in the dissipation range. Our results are in good agreement with recent and rediscovered theoretical predictions on the role of thermal noise in turbulence, as well as recent molecular gas dynamics simulations of decaying compressible turbulence.

Ishan Srivastava
Lawrence Berkeley National Laboratory
isriva@lbl.gov

Daniel R. Ladiges
Center for Computational Science and Engineering
Lawrence Berkeley National Laboratory
drladiges@lbl.gov

Andrew J. Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

Alejandro L. Garcia
San Jose State University
Alejandro.Garcia@sjsu.edu

John B. Bell
Lawrence Berkeley National Laboratory, U.S.

jbbell@lbl.gov

MS36

Inversion of Non-linear Parameters in Regional-scale 3D Dynamic Rupture Simulations using Multi-level Delayed Acceptance Markov Chain Monte Carlo Sampling with Optimized Speed-up

Physics-based regional-scale 3D dynamic rupture simulations are essential for understanding the mechanisms behind earthquakes. Properly modeling a real earthquake typically requires more than hundreds of CPU hours, presenting a significant challenge in quantitatively constraining the non-linear parameters in 3D dynamic rupture simulations based on geophysical observations. We propose a method to accelerate the non-linear inversion problem, leveraging a Multi-level Delayed Acceptance (MLDA) Markov Chain Monte Carlo (MCMC) algorithm and modern computing infrastructure. The MLDA code interfaces with the dynamic rupture simulations via UM-Bridge. This method accelerates the inversion through three layers of parallelization: optimally parallelized 3D dynamic rupture forward simulations with SeisSol, concurrent progression of multiple Markov chains, and prefetching-based parallel evaluation within each chain. We demonstrate how to explore the optimal configuration of these parallelization layers to minimize the time to solution, given an upper bound on the available computing resource and a fixed number of MCMC samples required for evaluating the posterior probability distribution of the non-linear parameters. We demonstrate the MLDA algorithm can be optimized for the non-linear inversion of parameters in dynamic rupture models of the Ridgecrest earthquake sequence in 2019, utilizing seismic and geodetic observations.

Zihua Niu
Ludwig-Maximilians-Universität München
zihua.niu@lmu.de

Maximilian Kruse
Karlsruhe Institute of Technology
maximilian.kruse@kit.edu

Linus Seelinger
Karlsruhe Institute of Technology (KIT)
Scientific Computing Center
linus.seelinger@kit.edu

Sebastian Wolf
Technical University Munich
s.wolf@tum.de

Vikas Kurapati, Michael Bader
Technical University of Munich
vikas.kurapati@tum.de, bader@in.tum.de

Heiner Igel
Ludwig-Maximilians-Universität München
heiner.igel@geophysik.uni-muenchen.de

Alice Gabriel
UC San Diego
liese.gabriel@gmail.com

MS36

Gaussian Process Surrogates and Uncertainty Quantification for Complex Industrial Simulations

Complex numerical simulations often suffer from long runtimes and need to be replaced by cheaper surrogate models, particularly in the context of uncertainty quantification where several thousand model runs are needed to compute reliable estimates for certain quantities of interest. In this talk, we introduce and analyze Gaussian process-based approaches like standard Gaussian processes (GPs), multi-output GPs, and non-stationary deep kernel GPs to approximate complex simulations. The output of such models is often high-dimensional or even function-valued. We therefore exploit projection methods like principal component analysis (PCA) to expand and decompose the outputs. We build surrogate models that predict only a few coefficients of the PCA basis, retaining high accuracy in the final model approximation with reasonable computational effort. Finally,we show numerical results where we apply the surrogate models to braking pressure simulations from the Robert BOSCH GmbH, illustrating the performance and applicability of the proposed methods in actual industrial workflows.

Julian Nürk
Heidelberg University & Bosch Research
mailto:julian.nuerk@de.bosch.com

MS36

Characterizing Aquifer Properties Through a Sparse Grid-Based Bayesian Framework and Insar Measurements: A Basin-Scale Application to Alto Guadalentn, Spain

Aquifer characterization is essential for predicting aquifer responses and ensuring sustainable groundwater management. In this study we develop a sparse-grids-based Bayesian framework to infer the hydraulic conductivity and the soil compressibility of over-exploited aquifer systems using Interferometric Synthetic Aperture Radar (InSAR) ground displacement datasets and piezometric records. The phenomenon is described by a poroelastic model (consisting of a set of coupled, nonlinear PDEs) that describes the interplay between groundwater depletion and soil deformation through the explicit quantification of the porosity change. The Bayesian inversion approach enables to compute the posterior distribution of the uncertain model parameters. However, exploring this posterior using Markov Chain Monte Carlo is computationally prohibitive due to the substantial cost of solving the poroelastic model. To overcome this issue, we use sparse-grids surrogate models to approximate such solutions. The methodology is applied to the Alto Guadalentn basin, Spain, where long-term aquifer exploitation has led to a lowering of the water table larger than 100 m causing impressive land subsidence, with rates up to 15 cm/yr as evidenced by InSAR. The results demonstrate that integrating InSAR data significantly enhances the characterization of the aquifer properties, with the resulting numerical simulations aligning well with available observations.

Yueting Li, Claudia Zoccarato
University of Padova
yueting.li@phd.unipd.it, claudia.zoccarato@unipd.it

Chiara Piazzola
Technische Universität München
chiara.piazzola@tum.de

Lorenzo Tamellini
CNR-IMATI

tamellini@imati.cnr.it

Guadalupe Bru, Carolina Guardiola Albert
Geological Survey of Spain
g.bru@igme.es, c.guardiola@igme.es

Pietro Teatini
University of Padova
pietro.teatini@unipd.it

MS37

On Scalable Preconditioning of an Implicit Vms Resistive MHD Formulation with Application to MCF

A scalable multiphysics block preconditioner for an implicit low Mach number compressible resistive magnetohydrodynamics (MHD) formulation based on a variational multiscale (VMS) finite element (FE) method is presented. The VMS formulation stabilizes the strongly convective flow effects, and saddle point structure of the magnetic field elliptic divergence cleaning term and the low Mach number nearly incompressible flow limit. The MHD model is intended to simulate macroscopic plasma instabilities and disruptions in complex 3D Tokamak devices used for exploring magnetic confinement fusion (MCF). In this talk the block structure of the Newton linearized discrete system is presented along with a brief description of the approximate block factorization and Schur complement operators that encode the critical cross-coupling physics of the system. The discussion includes a few simulation results on MCF relevant problems in ITER geometry. The numerical results will then present representative strong and weak scaling of the block preconditioner along with the Lundquist number robustness of the method. This work is in collaboration with J. Hu, and X. Tang. This work was partially supported by the U.S. DOE, Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics Program. It has also been partially supported by the DOE, ASCR and Fusion Energy Sciences SciDAC Partnership program.

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Jesus Bonilla
Los Alamos National Laboratory
jbonilla@lanl.gov

Peter Ohm
RIKEN Center for Computational Science
peter.ohm@riken.jp

Edward G. Phillips, Michael Crockatt
Sandia National Laboratories
egphill@sandia.gov, mmcrock@sandia.gov

Raymond S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

MS37

Analysis of Wave Interactions of Incompressible Hall MHD Turbulence

The incompressible Hall MHD system exhibits two branches of waves, each as positive and negative curl eigenstates, which acquire whistler and cyclotron frequencies at large wavenumbers. To study properties of wave interactions and turbulence, we decompose the results of a pseudospectral Hall MHD simulationinto these four waves. Strang splitting is applied to exactly integrate the forcing and dissipation terms, and the implicit midpoint rule is used to integrate the ideal fluid terms in a way that conserves the quadratic invariants - the energy, magnetic helicity, and canonical helicity. Upon integration, the total energy in each wave is shown over the course of the simulation to visualize wave interactions. Energy spectra, longitudinal and transverse structure functions, and the Goldreich-Sridhar nonlinearity parameter are obtained for each type of wave.

Erik Hansen
University of Texas - Austin
ehansen99@utmail.utexas.edu

Akash Shukla, David Hatch
University of Texas at Austin
akashukla@utexas.edu, drhatch@austin.utexas.edu

Prerana Sharma
Ujjain Engineering College
preranaiit@gmail.com

Justin Walker
University of Wisconsin, Madison
justin.walker@wisc.edu

Philip J. Morrison
Physics Department
University of Texas at Austin
morrison@physics.utexas.edu

Swadesh Mahajan
University of Texas at Austin
Shiv Nadar University
mahajan@mail.utexas.edu

MS37

Preconditioning and Verification of Spectral-Element MHD for Stellarators

The NIMSTELL code has been developed to model macroscale plasma dynamics in stellarators. Non-axisymmetric geometry, equilibrium and perturbed fields are represented using spectral elements over the poloidal plane and Fourier series in a generalized toroidal coordinate. The discretized linear MHD equations form a large sparse system for Fourier components of the perturbed fields, which are coupled due to asymmetry. The linear system is preconditioned by LU factorization of matrices for subsets of Fourier components, called "bands." Bands of consecutive Fourier harmonics are inefficient for stellarators with multiple field periods, where the strongest coupling is between non-consecutive harmonics. Therefore, a new capability was added to form the preconditioning bands using user-defined sets of harmonics. Also, since the sparse matrix structure is determined solely by the 2D mesh and not the Fourier harmonics, a reduced matrix is used to minimize the memory required by ParMETIS [Karypis, 10.1007/978-0-387-09766-4_500] for calculating a fill-reducing permutation. A tearing mode benchmark of NIMSTELL for a set

of W7-A cases [Nikulsin, et al, Phys Plasmas 29, 063901] shows maximum deviations of 1.5% and 11.0% with respect to CASTOR3D and the reduced-MHD JOREK, respectively. Results for tearing modes in quasiaxisymmetric cases and pressure-driven modes in quasihelically symmetric cases are also presented. Work supported by US DOE grants DE-SC0024548 and DE-FG02-99ER54546.

Sanket Patil
University of Wisconsin - Madison
sapatil@wisc.edu

Carl Sovinec
University of Wisconsin-Madison
csovinec@wisc.edu

MS39

Rabbits and Their Warrens Plus Other Co-Design Tales

As part of the design and development of the El Capitan system for LLNL, HPE (and formerly Cray prior to acquisition by HPE) worked closely with both AMD and LLNL to design a solution best equipped to meet LLNL's needs. Over the course of several years since the contract was signed, many different design meetings were held, focused on a broad range of the hardware and software aspects being developed for deployment. This talk will cover several of these co-design topics. Rabbit is a "near-node-local storage solution" that was the result of significant co-design activities. Various integration strategies for other compute node components were also discussed as part of the co-design efforts.

Larry Kaplan
HPE
laurence.kaplan@hpe.com

MS39

Grace-Hopper Innovations for Science

Grace-Hopper systems are delivering significant scientific capabilities for HPC and AI supercomputing workloads, including Venado at Los Alamos and EXA-HE at CEA. In this talk, well discuss some of the power efficient design choices made, including the first use of LPDDR in a data center processor and the high-speed low-power bus connecting Grace and Hopper. We will also discuss the collaborative work with computer centers to enable their mission applications to run well on these systems and lessons learned from these engagements.

Ian Karlin, Scott Halverson, Giridhar Chukkapali
NVIDIA
ikarlin@nvidia.com, shalverson@nvidia.com, gchukkapalli@nvidia.com

MS39

Porting and Optimizing a Large-Scale CFD Solver for Exascale Systems

In this presentation we summarize our experiences porting and optimizing a large-scale CFD solver developed at NASA for use on Frontier and Aurora, recent exascale systems based on AMD and Intel GPUs, respectively. The solver is an unstructured-grid formulation incorporating various turbulence treatments and finite-rate chemistry for applications across the speed range. Several performance challenges were encountered when migrating CUDA-optimized kernels to SYCL, including high register spills, memory latency, and poor vectorization. These issues were successfully addressed through the use of ESIMD, a low-level extension to the Intel oneAPI programming framework.

Eric Nielson
NASA Langley Research Center
eric.j.nielsen@nasa.gov

Aaron Walden, Gabriel Nastac
NASA Langley
aaron.c.walden@nasa.gov, gabriel.c.nastac@nasa.gov

Mohammad Zubair
Old Dominion University
zubair@cs.odu.edu

Christoph Bauinger
Intel Corporation
christoph.bauinger@intel.com

MS39

Harnessing the Capabilities of An APU: Software and Hardware Co-Design for Emerging Supercomputers

The El Capitan supercomputer at Lawrence Livermore National Laboratory (LLNL) is expected to be the largest supercomputer in the world, a development representing a transition in scientific computing: early progress in the era of Exascale computing. Notably, El Capitan leverages innovations of the AMD Instinct MI300A APU, its underlying accelerator which was designed to enable state-of-the art capabilities in raw compute, high-bandwidth memory, programming models, and energy efficiency. While the El Capitan supercomputer and the underlying MI300A APU demonstrate a feat of engineering and innovation in high-performance computing, how will such raw computational capability translate into improved modeling and scientific simulation? This question is at the forefront of research and development efforts at the El Capitan Center of Excellence. In this talk, we discuss not only innovative aspects in the MI300A APU architecture, but how these are translating into realized performance gains in scientific simulations, drawing examples from a diverse array of applications represented in the El Capitan Center of Excellence and other close collaborations with application developers.

Michael Rowan, Nicholas Malaya
AMD
michael.rowan@amd.com, nicholas.malaya@amd.com

MS39

Co-Design of the Cerebras Wafer-Scale Engine for Physical Simulation

The recent explosion in the AI market has led to the development of specialized hardware for performing training and inference on neural networks. The Cerebras CS-3 is one such hardware innovation: the system is built from the ground up to accelerate large AI training and inference workloads. The CS-3 is built around the wafer-scale engine (WSE), a single chip consisting of 900,000 processing elements (PEs) laid out in a 2D mesh, with each PE containing 48 kB of memory accessible in a single cycle. While this chip has explicitly been the product of co-design with AI workloads, many of its features also address bottlenecks

present in more traditional HPC workloads. In this talk we discuss recent results on several HPC workloads in molecular dynamics and seismic imaging which have benefited from these architectural innovations.

Leighton Wilson
Cerebras
leighton.wilson@cerebras.net

MS40

Impact of Preconditioning, Reordering, and Overlap on Solving Stiff Linear Systems with Gmres for Low Dissipation CFD Schemes

The development of fast, robust, and scalable preconditioning strategies remains a challenge for compressible fluid flow problems, particularly on GPU-based computing platforms. In strongly hyperbolic regimes algebraic multigrid methods exhibit severe degradation and the focus has been on additive Schwarz style preconditions, predominantly with sparse direct solvers for local domain blocks. This talk examines GPU-enabled solvers within the cuSOLVER, Trilinos, SuperLU, and Ginkgo linear algebra packages with a focus on tradeoffs between the efficiency of the preconditioner setup (factorization) and apply (triangular solves). A preconditioner reuse approach is employed to mitigate expensive factorization over multiple nonlinear iterations/solves. Additionally, the impact of different parallel mesh partitioning strategies on the efficacy of additive Schwarz preconditioning will be discussed. Numerical results for two and three dimensional test problems will be presented.

Ryan S. Glasby
University of Tennessee
glasbyrs@ornl.gov

Steven Hamilton, Taylor Erwin
ORNL
hamiltonsp@ornl.gov, erwinjt@ornl.gov

MS40

Unifying the MPI Communication Capabilities in Hypre

At large scales, parallel applications often encounter significant communication costs, particularly in those that rely on stencil operations and linear solvers. Stencil-based applications are typically dominated by boundary exchanges, while linear solvers face bottlenecks due to irregular communication patterns. When data resides in GPU memorieswhether device-only or unifiedadditional layers of cost, performance challenges, and software complexity arise. In this work, we present a unified software approach in hypre that addresses various MPI communication scenarios. We also showcase performance improvements achieved through optimizations such as neighborhood collectives, locality-aware irregular communication, and persistent and partitioned communication in the AMG solve.

Ruipeng Li
Lawrence Livermore National Laboratory
li50@llnl.gov

Gerald Collom
University of New Mexico
Albuquerque NM, USA
geraldc@unm.edu

MS40

Improvement of the Strong Scalability of a Block Tridiagonal Solver for Exascale Architectures

One of the most efficient solvers used in Sandia Parallel Aerodynamics Reentry Code (SPARC) and MPAS-Albany Land Ice is the block tridiagonal solver. The solver is a line preconditioner that divides the mesh into lines: sequences of interconnected mesh nodes. The solver consists in a stationary iterative method that requires the applications of the inverse of the matrix where only the interactions between different lines are dropped. It relies on a block tridiagonal LU per line and its application per line. Both of them are parallelized using a batched strategy using a MPI+Kokkos strategy for the loop over the lines. The lines are distributed over different MPI ranks and every MPI ranks use the Kokkos programming model to loop over the associated lines. While this solver exhibited almost perfect strong scaling on CPU-based supercomputers, the scalability can be significantly deteriorated on GPU-based supercomputers. This is due to the fact that while strong scaling, the number of lines assigned per MPI ranks decreases and that number can becomes to small to fully utilize the available computational power per rank when using GPUs. In order to improve the strong scalability, we investigated the usage of a Schur complement approach per line to divide lines into sublines and to expose more parallelism per GPU. In this presentation, we describe the algorithm, the Kokkos based implementations, and present results on NVIDIA V100 and on AMD MI300.

Kim Liegeois, Luc Berger-Vergiat, Victor Brunini
Sandia National Laboratories
knliege@sandia.gov, lberge@sandia.gov,
vebruni@sandia.gov

Travis Fisher
Sandia National Labs
tcfishe@sandia.gov

Sivasankaran Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

MS40

New Techniques for Improving Sparse Triangular Solvers for Distributed Memory GPUAccelerated Systems

Sparse triangular solvers are key components in many sparse solvers used in direct and iterative preconditioning methods. However, their poor scalability remains a bottleneck in scaling solver applications, especially on highly parallel systems. Notably, computing sparse triangular solvers is problematic even for dense matrices. Improving sparse triangular solver efficiency on GPU-accelerated distributed memory systems is challenging. We present approaches to reduce communication and improve the scalability of sparse triangular solvers on these architectures. We introduce new techniques to boost their performance on GPUs and distributed memory. The main idea is to view a solution of triangular systems as traversing directed acyclic graphs. By analyzing the graph structure, we apply node contraction, shortcut creation, extra edge addition, and node/path duplication transformations to optimize perfor-

mance in different parallel and distributed scenarios.

Piyush Sao
Oak Ridge National Laboratory
saopk@ornl.gov

MS41

Transformers for Data Assimilation

Attention mechanisms and their use in transformer architectures have been widely successful at modeling nonlocal correlations in data. Recent interest in applying attention for operator learning motivates a formulation of the methodology in the function space setting. In this talk we outline the construction of an attention mechanism in the continuum. We show how this formulation can be leveraged to design transformer neural operators, neural network architectures mapping between infinite-dimensional spaces of functions, and discuss relevant universal approximation theory. We show that these architectures are competitive in cost and accuracy for operator learning tasks and demonstrate their application to Data Assimilation problems.

Edoardo Calvello
California Institute of Technology
e.calvello@caltech.edu

MS41

Dynamics-Aware Sampling and Bayesian Inference

Due to nonlinearities and symmetries in the underlying dynamical system, the invariant measure, and consequently, the Bayesian filtering and smoothing distributions possess a low-dimensional structure. In this talk, we explore data assimilation algorithms that exploit such low-dimensional structure that arise from certain classes of chaotic dynamical systems. Specifically, we leverage absolute continuity on submanifolds to reduce the dimensionality of score learning and score-based sampling algorithms. We discuss how to use Jacobian information from numerical models of the dynamics and likelihood for efficient score learning and score-based sampling.

Nisha Chandramoorthy
Georgia Institute of Technology
nishac@uchicago.edu

MS41

Data Assimilation for State and Parameter Estimation in Plasma Physics

Plasma dynamics are governed by physical processes that span a vast range of spatial and temporal scales. Plasma applications often require an accurate parameterization of the multi-scale physics. However, real-time control and optimization in applications demands computationally efficient models that can rapidly predict and adapt to changing conditions. Therefore, various modeling approximations are made, introducing numerous sources of uncertainties. In this work, we evaluate how spatially and temporally sparse data can be used to systematically estimate the underlying parameters and their uncertainty for range of plasma applications. For online estimation, we focus on sequential data assimilation (DA) techniques in two classes of problems. First, we employ data assimilation (DA) techniques to infer computationally expensive electron kinetics from sparse nonlinear optical intensity measurements in a system with stiff chemical reaction networks. Our objective is to identify key spectroscopic observables and chemical pathways, enabling efficient determination of electron properties relevant to semiconductor manufacturing. For the second application, we estimate spatio-temporally varying coefficients in reaction-advection PDEs from sparse data. Our results highlight the necessity of incorporating spatially correlated parameterization to accurately capture model and measurement uncertainty in spatially distributed models encountered in applications involving plasma propulsion.

Anubhav Dwivedi, Kentaro Hara
Stanford University
anubhavd91@gmail.com, kenhara@stanford.edu

MS41

Data Assimilation and Dynamical Closure from the Perspective of Quantum Mechanics

In recent years, a fruitful approach for data-driven analysis and modeling of dynamical systems has been to consider the action of (nonlinear) dynamics in state space on linear spaces of observables. These methods leverage the linearity of the associated evolution operators, namely the Koopman and transfer operators, to carry out tasks such as mode decomposition, forecasting, and uncertainty quantification using linear operator techniques. Mathematically, the operator-theoretic approach has close connections with representations of nonlinear transformations (the state space dynamics) into spaces of functions (the observables) with a commutative algebraic structure. In this talk we discuss generalizations of this framework to the setting of non-commutative algebras of operators using ideas from quantum theory. Central to our approach is a representation of observables and probability densities through multiplication operators and density operators ("quantum satates"), respectively. Using these objects, and the dynamical operators governing their evolution, we formulate data assimilation and dynamical closure in an operator-theoretic language. We discuss how the operator-theoretic approach leads to positivity-preserving computational schemes which are also amenable to data-driven implementation using kernel methods for operator learning.

Dimitrios Giannakis
Dartmouth College
Dimitrios.Giannakis@dartmouth.edu

MS42

Maximum Likelihood Discretization of the Transport Equation

Traditional Galerkin methods can struggle to maintain positivity when solving the transport equation, a fundamental local mass conservation law in continuum fluid and gas dynamics. This results in nonphysical negative densities, hindering downstream computations like sampling, computing pressure or reaction rates. Viewing the mass density as a non-negative probability distribution turns the problem into a statistical inference task, enabling the development of Galerkin methods that intrinsically preserve positivity. We propose treating the Galerkin basis as a parametric probability distribution and using maximum likelihood estimation to infer the density distribution as the particles are advected. Thus, we derived a Fisher-Rao Galerkin Semidiscretization of the transport equation. This scheme implicitly preserves positivity and is more sensitive to relative magnitude changes, as measured by KL

divergence, due to the Fisher-Rao metrics invariance under re-parameterization.

Brooks Eyob
Georgia Institute of Technology, U.S.
brook@gatech.edu

Florian Schaefer
Georgia Institute of Technology
florian.schaefer@cc.gatech.edu

MS42

Cohomology-Aware Reduced-Order Modeling

Projection-based techniques for the reduced-order modeling of partial differential equations (PDEs) frequently rely on approximation spaces which maximize the variance in a set of snapshot data. However, this procedure does not respect the topology of the domain on which these PDEs are solved, leading to poor results on solution domains with nontrivial connectedness. We show that a straightforward conceptual modification based on the Hodge decomposition can effectively resolve this issue at little-to-no additional overhead, ensuring that the reduced-order space of solutions has the same topological character as its full-order counterpart and leading to large improvements in accuracy over topologically uninformed methods.

Anthony Gruber
Sandia National Laboratories
adgrube@sandia.gov

MS42

Intelligent Attractors for Singularly Perturbed Dynamical Systems

Singularly perturbed dynamical systems, commonly known as fast-slow systems, play a crucial role in various applications such as plasma physics. They are closely related to reduced order modeling, closures, and structure-preserving numerical algorithms for multiscale modeling. A powerful and well-known tool to address these systems is the Fenichel normal form, which significantly simplifies fast dynamics near slow manifolds through a transformation. However, the Fenichel normal form is difficult to realize in conventional numerical algorithms. In this work, we explore an alternative way of realizing it through structure-preserving machine learning. Specifically, a fast-slow neural network (FSNN) is proposed for learning data-driven models of singularly perturbed dynamical systems with dissipative fast timescale dynamics. Our method enforces the existence of a trainable, attracting invariant slow manifold as a hard constraint. Closed-form representation of the slow manifold enables efficient integration on the slow time scale and significantly improves prediction accuracy beyond the training data. We demonstrate the FSNN on several examples that exhibit multiple timescales, including the Grad moment system from hydrodynamics, two-scale Lorentz96 equations for modeling atmospheric dynamics, and Abraham-Lorentz dynamics modeling radiation reaction of electrons in a magnetic field.

Daniel A. Serino
Los Alamos National Lab
dserino@lanl.gov

Allen Alvarez Loya
Department of Applied Mathematics
University of Colorado, Boulder, CO
aalvarezloya@gmail.com

Joshua W. Burby
UT Austin
joshua.burby@austin.utexas.edu

Ioannis G. Kevrekidis
Johns Hopkins University
yannisk@jhu.edu

Qi Tang
Los Alamos National Lab
qtang@gatech.edu

MS43

Physics-Based Learning Through Nested Operator Inference

We introduce the data-driven nested Operator Inference (OpInf) method for learning projection-based reduced-order models (ROMs) from snapshot data of high-dimensional dynamical systems. Projection-based ROMs exploit the intrinsic low-dimensionality of a full-order solution manifold to 1) achieve significant computational savings, 2) guarantee approximation accuracy through established error theory, and 3) remain interpretable through the governing equations. However, constructing ROMs via projection requires access to the full-order operators – a significant shortcoming for applications with legacy codes or commercial solvers. OpInf circumvents this requirement by learning the intrusive ROM from available full-order data and the structure of the governing equations. However, meeting its data requirements in practice can be challenging, especially for highly non-linear operators. In contrast, our nested OpInf approach partitions the learning problem into multiple regression problems, each provably better conditioned than when all reduced-order operators are learned together. The partition is based upon a nested structure in the projection-based reduced-order matrices. It exploits a hierarchy in the reduced space's basis vectors to guarantee that the ROM's dominant dynamics are learned accurately. Nested OpInf is particularly applicable to higher-order polynomial systems, which we demonstrate for the shallow ice equations with eighth-order polynomial operators.

Nicole Aretz
The University of Texas at Austin
nicole.aretz@austin.utexas.edu

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

MS43

Closing the Gap Between Scientific Foundation Models and Real-World Applications

In the era of GPT models, one gets notoriously confronted with the question of where we stand with applicability of large-scale deep learning models within scientific or engineering domains. The discussion starts by reiterating on recent triumphs in weather and climate modeling. Further, we discuss recent breakthroughs in fluid dynamics and related engineering fields, and subsequently extrapolate insights applicable to relatively untouched scientific and engineering fields. Finally, we outline challenges and potential solutions when it comes to scalability beyond traditional

numerical schemes and discuss the respective impact on industry and scientific environments.

Johannes Brandstetter
Microsoft Research
Johannes Kepler University Linz
brandstetter@ml.jku.at

MS43

Machine Learning for Sparse Nonlinear Modeling and Control

Accurate and efficient reduced-order models are essential to understand, predict, estimate, and control complex, multiscale, and nonlinear dynamical systems. These models should ideally be generalizable, interpretable, and based on limited training data. This work describes how machine learning may be used to develop accurate and efficient nonlinear dynamical systems models for complex natural and engineered systems. We explore the sparse identification of nonlinear dynamics (SINDy) algorithm, which identifies a minimal dynamical system model that balances model complexity with accuracy, avoiding overfitting. This approach tends to promote models that are interpretable and generalizable, capturing the essential physics of the system. We also discuss the importance of learning effective coordinate systems in which the dynamics may be expected to be sparse. This sparse modeling approach will be demonstrated on a range of challenging modeling problems, for example in fluid dynamics, and we will discuss how to incorporate these models into existing model-based control efforts.

Steven Brunton
University of Washington
sbrunton@uw.edu

MS43

Physics-Informed Graph Neural Network Framework for Solving Spatiotemporal PDEs with Irregular Boundaries

We introduce a physics-informed framework that integrates an encoder-decoder architecture with a Graph Neural Network (GNN) to address spatiotemporal partial differential equations (PDEs) under irregular boundary conditions without relying on observed data. The encoder utilizes a fully connected neural network to convert input features into a graph structure, encoding spatial information into nodes and edges. This structure enables the network to learn and represent spatial interactions by embedding node states and their pairwise relationships into the graph. The processor, implemented as a GNN, iteratively updates node and edge attributes through message-passing, refining the graph state to capture intricate spatial relationships and interactions. The decoder, leveraging Gated Recurrent Units (GRUs), models the temporal dynamics by extracting relevant information from the final graph state. The GRUs manage the temporal evolution of the PDE solution, ensuring that critical information from the initial conditions is retained throughout time integration. This method effectively handles the challenges posed by irregular boundaries by integrating both spatial and temporal dependencies. The proposed approach is validated on various PDEs with irregular domains, demonstrating its robustness in managing complex boundary conditions.

Jie Long, Abdul Khaliq
Middle Tennessee State University
jl7u@mtmail.mtsu.edu, abdul.khaliq@mtsu.edu

Monther Rashed Alfuraidan
King Fahd University of Petroleum and Minerals, Saudi Arabia
monther@kfupm.edu.sa

MS43

Formal Grammars in Scientific Discovery

The automated discovery of physics forms an emerging research area, that focuses on automatically obtaining symbolic expressions describing the behavior of physical systems from experimental or synthetic data. Existing symbolic/sparse regression methods rely on the availability of libraries of models, which are typically hand-designed by a human expert using known models as reference, or deploy algorithms with exponential complexity which are only applicable for systems described by simple mathematical expressions. In this talk, we discuss a novel approach to physics discovery relying on formal grammars as an automated and systematic tool to generate physics expressions. Compliance with physics constraints is enforced through the use of language semantics. We deploy the approach for two tasks: i) Automatically generating a library of valid mathematical expressions that describe material laws and dynamic equations of motion; ii) Performing data-driven discovery of these models from measurements with different noise levels. For the task of automatic library generation, we demonstrate the flexibility and efficiency of the proposed methodology in avoiding hand-crafted features and human intervention. For the data-driven discovery task, we demonstrate the accuracy, robustness and significant generalizability of the proposed methodology.

Eleni Chatzi
Institute of Structural ENgineering
ETH Zurich
chatzi@ibk.baug.ethz.ch

Georgios Kissas, Karin Yu
ETH Zurich
gkissas@ai.ethz.ch, yuk@ethz.ch

Orestis Oikonomou
University of Zurich
orestis.oikonomou@ai.ethz.ch

MS44

Approximate Deconvolution Reduced Order Models: Analysis and Numerical Simulation

In the field of reduced order modeling, filtering methods have received significant recent attention for treating poorly performing models as a method of regularization. However, filtering by itself can be susceptible to oversmoothing of generated solutions. We propose and prove properties of new reduced order models that integrate the methods of approximate deconvolution with a goal of balancing accuracy and smoothing properties. The models and theorems are tested with particular application to fluid flows simulated by the Navier-Stokes equations.

Ian Moore
Department of Mathematics
Virginia Tech
ianm9123@vt.edu

Anna Sanfilippo
Dipartimento di Matematica
Università di Trento
anna.sanfilippo@unitn.it

Francesco Ballarin
Catholic University of the Sacred Heart, Italy
francesco.ballarin@unicatt.it

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS44

Entropy Conserving Discountinous Galerkin Method

Reduced order models are sought after to minimize the total number of variables required to solve for a full-order solution. In turn, this theoretically reduces the computational cost to achieve similar results. One type of reduced-order model is the projection-based model. Projection-based methods have been used to generate reduced order models that are conservative, and potentially stable [Carlberg, Barone, and Antil, Galerkin v. least-squares PetrovGalerkin projection in nonlinear model reduction, 2017]. Two approaches, the Galerkin or Petrov-Galerkin are used due to simplicity. While these approaches are not typically guaranteed to be stable, they can be modified to prove the conservation of a conservative quantity [Kalashnikova et al., Construction of Energy-Stable Galerkin Reduced Order Models, 2013]. While entropy-stable Galerkin reduced-order modeling has been established for finite volume [Chan, Entropy stable reduced order modelling of nonlinear conservation laws, 2020], we expand the approach to a discontinuous Galerkin (DG) model. This approach works by modifying the projection operator typically seen in an entropy-conservative DG scheme to account for the projection into the reduced-order space. This is then expanded to a hyper-reduction using an empirical cubature for point selection and weighting. Results are shown for test cases for the Euler equations in multiple dimensions that display that the convective entropy is conserved.

Tyson Ofstie
McGill University
tyson.ofstie@mail.mcgill.ca

MS44

Numerical Analysis of Filter-Based Stabilization for Reduced Order Models of Convection-Dominated Flows

In recent times, there have been several developments in ways to utilize spatial filtering for fluid flow problems. This talk will mainly involve two different variations in which spatial filtering can be applied: a more involved Large Eddy Simulation (LES) approach with closure modeling, and a simpler stabilization approach that leads to several regularized ROMs (Reg-ROMs). LES-ROMs and Reg-ROMs are strategies that leverage spatial filtering to mitigate the spurious numerical oscillations typically seen in classical Galerkin ROMs (G-ROMs) during under-resolved numerical simulations of convection-dominated flows. We prove results such as verifiability, stability, and convergence with several numerical simulations, showcasing agreement with theory.

Jorge Reyes
Virginia Tech
reyesj@vt.edu

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS44

Data-Driven Reduced Order Modeling of Two-Dimensional Stratified Rotating Fluid

In the current study, we investigate rotational and stratified turbulent systems within the framework of the Boussinesq equations. The effects of rotation and stratification give rise to a variety of dynamical regimes in atmospheric and oceanic turbulence. To develop a better understanding of the fundamental flow properties through a more straightforward mathematical framework, we commence our research with a simplified two-dimensional Boussinesq model that incorporates rotation and stratification effects. We explore two distinct limiting dynamics, considering pure rotation and pure stratification separately. A data-driven Reduced order model with deep neural networks will be introduced to find the low-dimensional latent dynamics of the Boussinesq system.

Xuping Xie
Old Dominion University
xxie@odu.edu

MS45

A Fourth-Order Exponential Time Differencing Scheme with a Real and Distinct Poles Rational Approximation for Solving Nonlinear Systems of Reaction Diffusion Equations

Reaction-diffusion systems are mathematical models that describe the spatiotemporal dynamics of chemical substances as they diffuse and react. Variety of time discretization schemes have been developed to solve the stiff ODE system resulting from the spatial-discretization of the PDE. The quest to develop more efficient and accurate schemes to handle very stiff and non-smooth problems is ever increasing. In this work, we develop a fourth-order, L-stable, parallelizable exponential time differencing scheme (ETD) by approximating the matrix exponentials in the class of ETD Runge-Kutta schemes with a fourth-order non-Pad rational function having real distinct poles (RDP). A variety of non-linear reaction-diffusion systems having Dirichlet, Neumann and periodic boundary conditions are used to empirically validate the order of convergence of the scheme and compare its performance with existing fourth order schemes.

Wisdom K. Attipoe, Emmanuel Asante-Asamani
Clarkson University
attipowk@clarkson.edu, easantea@clarkson.edu

MS45

Nonoverlapping, Localized Exponential Time Differencing Methods for Evolution Partial Differential Equations

Exponential integrators, among them the Exponential

Time Differencing (ETD) methods, have been widely used for solving stiff evolution equations due to their accuracy, stability and ability of maintaining exponential behavior. The cost of these methods is dominated by the computing of matrix exponentials and their products with vectors. To overcome this challenge, one possibility is to use domain decomposition techniques and solve a sequence of smaller-sized subdomain problems, in which the matrix exponentials are computed locally and in parallel. In this talk, we develop nonoverlapping, localized ETD methods for hyperbolic conservation laws and shallow water equations with discontinuous Galerkin (DG) discretization in space. Both mathematical analysis and numerical performance of the proposed methods will be investigated.

Thi Thao Phuong Hoang, Cao-Kha Doan, Phuoc-Toan Huynh
Auburn University
tzh0059@auburn.edu, kcd0030@auburn.edu, tph0017@auburn.edu

MS45

Low-rank Implicit Integrators Based on Reduced Augmentation for Solving 3D Advection-diffusion Equations in the Tucker Tensor Format

Many classical methods for solving partial differential equations (PDEs) suffer from the curse of dimensionality, that is, the number of degrees of freedom grows exponentially with the number of dimensions. To remedy this problem, many recent works have focused on exploiting low-rank structure in the PDE solution to reduce the storage and computational complexities. For high-dimensional problems, this often involves using low-rank tensor decompositions to store the solution. However, achieving high-order accuracy in time using implicit and implicit-explicit discretizations in the low-rank framework remains a challenge. We propose a low-rank scheme inspired by the popular dynamical low-rank (DLR) class of methods that incorporates high-order implicit time discretizations for solving three-dimensional advection-diffusion equations.

Joseph Nakao
Swarthmore College
jnakao1@swarthmore.edu

MS45

OEDG: Oscillation-Eliminating Discontinuous Galerkin Method for Hyperbolic Conservation Laws

Controlling spurious oscillations is crucial for designing reliable numerical schemes for hyperbolic conservation laws. In this talk, we propose a novel, robust, and efficient oscillation-eliminating discontinuous Galerkin (OEDG) method on general meshes, motivated by the damping technique in [Lu, Liu, and Shu, SIAM J. Numer. Anal., 59:12991324, 2021]. The OEDG method incorporates an OE procedure after each RungeKutta stage, and it is devised by alternately evolving the conventional semidiscrete DG scheme and a damping equation. A novel damping operator is designed to possess both scale-invariant and evolution-invariant properties. We rigorously prove optimal error estimates of the fully discrete OEDG method for smooth solutions of linear scalar conservation laws. The OEDG method exhibits many notable advantages. It effectively eliminates spurious oscillations for problems spanning various scales and wave speeds without problem-specific parameters. Furthermore, it retains the key properties of the conventional DG method, such as conservation, optimal convergence rates, and superconvergence. Moreover, the OEDG method maintains stability under the normal CFL condition, even in the presence of strong shocks associated with highly stiff damping terms. The OE procedure is non-intrusive, facilitating seamless integration into existing DG codes as an independent module. Extensive numerical results confirm the analysis and validate the effectiveness of the OEDG method.

Manting Peng
Southern University of Science and Technology
na

Zheng Sun
The University of Alabama
zsun30@ua.edu

Kailiang Wu
Southern University of Science and Technology
wukl@sustech.edu.cn

MS45

Accurate Discretizations and Efficient AMG Solvers for Extremely Anisotropic Diffusion Via Hyperbolic Operators

Heat conduction in magnetic confinement fusion can reach anisotropy ratios of 10^9-10^{10}, and in complex problems the direction of anisotropy may not be aligned with (or is impossible to align with) the spatial mesh. Such problems pose major challenges for both discretization accuracy and efficient implicit linear solvers. Although the underlying problem is elliptic or parabolic in nature, we argue that the problem is better approached from the perspective of hyperbolic operators. The problem is posed in a directional gradient first order formulation, introducing a directional heat flux along magnetic field lines as an auxiliary variable. We then develop novel continuous and discontinuous discretizations of the mixed system, using stabilization techniques developed for hyperbolic problems. The resulting block matrix system is then reordered so that the advective operators are on the diagonal, and the system is solved using AMG based on approximate ideal restriction (AIR), which is particularly efficient for upwind discretizations of advection. Compared with traditional discretizations and AMG solvers, we achieve orders of magnitude reduction in error and AMG iterations in the extremely anisotropic regime.

Golo Wimmer, Ben Southworth
Los Alamos National Laboratory
gwimmer@lanl.gov, southworth@lanl.gov

Qi Tang
Los Alamos National Lab
qtang@gatech.edu

Xianzhu Tang
Los Alamos National Laboratory
xtang@lanl.gov

MS46

Scientific Machine Learning for Surrogate Modeling, Parameter Identification and Transfer Learning of Multiphase Flows in Porous Media: Applica-

tion to the FluidFlower Co2 Injection Experiment

Scientific machine learninga modeling approach in which a deep neural network representation of a process is informed by both data and physics has received increased attention in the field of geophysics, from subsurface flows to earthquake forecasting and coupled thermo-hydro-mechanical processes. Here, we develop scientific machine learning approaches to build surrogate models of CO2 injection and migration in porous media in realistic geologic settings. In contrast with most explorations of ML techniques, here we rely on physical experiments conducted in the FluidFlower rig. We explore and compare different ML approaches in their ability to robustly: (1) emulate, via training with full-physics simulations, the complex fluid flow dynamics of the experiments in terms of the full state variables; (2) determine the sensitivity of quantities of interest to the parameters of the full-physics model; (3) solve the inverse problem to identify the underlying parameters in the ground-truth physical experiments; and (4) transfer learning from the detailed surrogate model building in one geometry of the geologic layers to a different geometry and injection location. We comment on the power, current limitations, and opportunities of machine learning approaches for modeling, forecasting and uncertainty quantification of subsurface CO2 sequestration and energy storage.

Hannah Lu
University of Texas at Austin
hannah.lu@austin.utexas.edu

Lluis Salo-Salgado
MIT
lsalo@mit.edu

Ehsan Haghighat
Department of Civil and Environmental Engineering
Massachusetts Institute of Technology
ehsanh@mit.edu

Ruben Juanes
MIT
Civil and Environmental Engineering
juanes@mit.edu

MS46

Pushing the Boundaries of Surrogate Model Accuracy: Physics-Informed Deep Operators for Complex PDEs

We introduce a novel multistage training approach for physics-informed operator learning using deep neural networks (DNNs) to capture high-frequency components in parametrized partial differential equation (PDE) solutions. Our method addresses the spectral bias limitation of DNNs through a sequential training process. The approach comprises two stages: (1) training an initial Solution DeepONet to predict the quantity of interest, and (2) sequentially training multiple Corrector DeepONets to minimize residual errors of combined preceding networks. Corrector DeepONet training involves scaling weights of preceding networks using factors derived from Fast Fourier Transform (FFT) analysis of residual errors. This FFT-based sequence enables capturing high-frequency solution components, significantly reducing generalization errors. We demonstrate our method's efficacy on 1D and 2D benchmark problems, including Burgers' equation, Euler-Bernoulli beam equation, and Darcy's equation. Our approach achieves near machine-precision accuracy, with generalization errors ranging from $O(10^{-12})$ to $O(10^{-16})$ in modeling these complex physical systems. These results highlight the potential of our multistage training approach for improving accuracy and generalization capabilities of physics-informed neural networks in solving parametrized PDEs.

Dibakar Roy Sarkar
Department of Civil and Systems Engineering
Whiting School of Engineering, Johns Hopkins University
droysar1@jh.edu

Somdatta Goswami
Johns Hopkins University
sgoswam4@jhu.edu

MS46

Predictive Digital Twins for Congestive Heart Failure Patients

We present a digital twin (DT) methodology for clinical decision-making to aid patients with congestive heart failure (CHF). CHF is a cardiovascular condition that gradually progresses to a chronic state. Severe CHF stages are a public health concern, estimated to affect 64 million people globally. Adopting a DT approach to monitor CHF patients could enable anticipatory personalized treatment, prevent acute stages, slow disease progression, and improve survival rates. The evolution of CHF patients is modeled using a dynamic Bayesian network, which governs the observations-to-decisions flow and quantifies uncertainties in identifying pathology states and treatments responses. Observational data, including biomarker measurements and lifestyle, socio-economic, and comorbidity indices, are continuously collected and assimilated for diagnostics using random forest classifiers. A digital state, reflecting the risk of CHF hospitalization, is continuously updated through sequential Bayesian inference, enabling the forecasting of future hospitalization risks and guiding treatment decisions. A preliminary offline phase involves training the random forest classifier, computing a control policy, and forming prior beliefs about control-dependent transition dynamics using a cohort of CHF patients from the UK Biobank database. Simulation results on experimental data demonstrate the DT ability to monitor patients conditions and suggest the most appropriate course of treatment.

Matteo Torzoni
Politecnico di Milano
matteo.torzoni@polimi.it

MIchela C. Massi
Human Technopole
michela.massi@fht.org

Davide Carnevali
Politecnico di Milano
davide1.carnevali@mail.polimi.it

Emanuele Di Angelantonio
Human Technopole
emanuele.diangeloantonio@fht.org

Francesca Ieva
Politecnico di Milano
Human Technopole
francesca.ieva@polimi.it

Andrea Manzoni
Politecnico di Milano
andrea1.manzoni@polimi.it

MS47

Reduced Order Modeling for Sh Wave Propagation

Exploring the influence of various seismic parameters on ground motion typically requires computationally intensive high-fidelity simulations, such as those using the Finite Element or Finite Difference methods. To overcome these computational challenges, we employ a reduced-order method (ROM) that approximates the solution in a lower-dimensional space, significantly cutting computational costs while preserving the critical dynamics of the problem. This study introduces a reduced-order modeling approach for solving the time-dependent wave equation within the context of seismic ground motion analysis. Our focus is on shear wave (SH) propagation, where shear velocity, a key varying parameter, critically influences the seismic response. We compute displacement as a function of time across different SH wave velocities. This approach enables efficient simulation across varying shear velocities, facilitating rapid assessment of ground motion under diverse seismic scenarios. The accuracy and efficiency of the ROM are validated by comparison with full-order finite element simulations, highlighting its potential as a powerful tool for seismic analysis in scenarios requiring fast and reliable predictions.

Shafqat Ali
University of L'Aquila Italy
shafqat.ali@univaq.it

Federica Di Michele
Gran Sasso Science Institute L'Aquila, Italy
federica.dimichele@gssi.it

Andriy Styahar
Ivan Franko National University of Lviv
astyahar@gmail.com

Bruno Rubino
Department of Mathematics, University of L'Aquila
Italy
bruno.rubino@univaq.it

Pierangelo Marcati
Gran Sasso Science Institute L'Aquila, Italy
pierangelo.marcati@gssi.it

MS47

Machine Learning, Reduced Order Modelling and Data Assimilation for Spatial Temporal Forecast

Numerical simulations of fluid dynamics have been indispensable in many applications relevant to physics and engineering. For improving predictive capability, numerical algorithms have become increasingly sophisticated and complicated by using more spatial and temporal resolution of datasets. Deep learning techniques applied to fluid flow modelling have gained significant attention in recent years. Advanced deep learning techniques achieve great progress in rapidly predicting fluid flows without prior knowledge of the underlying physical relationships. In this talk, an overview of deep learning techniques in fluidity dynamics is provided. Focus will be on Recurrent Neutral Network (RNN) (Long short-term memory LSTM), Convolutional Neutral Network (CNN) and Generative Adversarial Network (GAN). Reduced order modelling (ROM) and data assimilations techniques will be introduced for real-time operational modelling and uncertainty analysis. Having the compatibility of machines learning and data assimilation will be nothing short of revolutionary for a large number of disciplines. Application examples of large data-driven modelling to fluid flow problems will be presented, for example Ozone forecast in China and flooding prediction in Denmark.

Fangxin Fang
Department of Earth Science and Engineering
Imperial College London, U.K.
f.fang@imperial.ac.uk

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

MS47

Domain Decomposition for Surrogate Modelling

When conducting surrogate modeling for complex PDE systems, a whole global reduced model can be inefficient when the underlying problem involves high frequency modes or large physical domains. Based on the idea of physical domain decomposition, we propose local surrogates on subdomains and assemble a whole global surrogate through iterations.

Qifeng Liao
ShanghaiTech University
liaoqf@shanghaitech.edu.cn

MS47

Reduced-Order Modelling Approach Based on Large-Language Model

In the field of natural language processing, the advent of Transformer neural networks has led to the widespread application of highly parallel neural network architectures based on self-attention mechanisms in pre-training models, with BERT, GPT-2, and other improved pre-trained models establishing the "pre-training and fine-tuning" learning paradigm. Recent research has expanded large language models (LLMs) from natural language processing to time series and spatiotemporal tasks, such as weather forecasting, energy demand prediction, traffic flow forecasting, and disease trend prediction in healthcare. Despite their strong learning and representation capabilities, LLMs face challenges in time series and spatiotemporal tasks, including dependence on large historical datasets and generalization across different contexts. In frozen pre-training modules, we retain self-attention layers and feed-forward layers, as they contain most of the knowledge from pre-trained language models, providing better results when frozen during fine-tuning. Fine-tuning is applied to position encoding layers and the final output projection linear layers to ensure effectiveness with minimal modifications. In addition to the 2D cylindrical flow case, we also validated the outstanding performance of this large language model in time-series prediction tasks through a highly complex 3D air flow case in the South Bank University area.

Pin Wu
School of Computer Engineering and Science
Shanghai University

wupin@shu.edu.cn

MS47

Domain Decomposition for Physics Data Combined Neural Network Based ROM

We propose a domain decomposition (DD) method for the physics-data dual driven reduced-order model (ROM). The computational domain is divided into a number of subdomains, and a Physics-Data dual driven ROM is constructed for each subdomain, which not only represents the dynamics within its own region but also considers interactions with surrounding subdomains. The capabilities of this method are compared against the model without domain decomposition method and traditional model reduction method. The performance of this domain decomposition method is tested using three nonlinear problems, such as the Korteweg-de Vries (KdV) equation, the two-dimensional Kovasznay flow and the two-dimensional NavierStokes equation.

Dunhui Xiao
School of Mathematical Sciences
Tongji University
xiaodunhui@tongji.edu.cn

Xinyu Pan
School of Mathematical Sciences, Tongji University
panxy1019@tongji.edu.cn

MS48

Scalable Uncertainty Quantification with Approximate Priors

This talk studies Bayesian posterior distributions with approximate priors. Prior approximations often arise due to computational reasons, such as the need to discretize random fields. The focus of the talk is on quantitative rates of convergence of the approximate priors and posteriors. The analysis proceeds at the infinite-dimensional process level, which widens the applicability of the approach and guarantees its scalability. The theoretical results are applied in several uncertainty quantification settings, including priors learned from data and approximate Gaussian process regression.

Nicholas H. Nelsen
Massachusetts Institute of Technology
nnelsen@mit.edu

MS48

Multifidelity Linear Regression for Scientific Machine Learning from Scarce Data

Machine learning (ML) methods have garnered significant interest as potential methods for learning surrogate models for complex engineering systems for which traditional simulation is expensive. However, in many scientific and engineering settings, training data are scarce due to the cost of generating data from traditional high-fidelity simulations. ML models trained on scarce data have high variance and are sensitive to vagaries of the training data set. We propose a new multifidelity training approach for scientific machine learning that exploits the scientific context where data of varying fidelities and costs are available; for example high-fidelity data may be generated by an expensive fully resolved physics simulation whereas lower-fidelity data may arise from a cheaper model based on simplifying assumptions. We use the multifidelity data to define new multifidelity control variate estimators for the unknown parameters of linear regression models, and provide theoretical analyses that guarantee accuracy and improved robustness to small training budgets. Numerical results show that multifidelity learned models achieve order-of-magnitude lower expected error than standard training approaches when high-fidelity data are scarce.

Elizabeth Qian
Georgia Institute of Technology
eqian@gatech.edu

Dayoung Kang
Georgia Tech
School of Aerospace Engineering
dkang339@gatech.edu

Vignesh Sella
University of Texas Austin
vsella@utexas.edu

Anirban Chaudhuri
MIT
anirbanc@oden.utexas.edu

Anirban Chaudhuri
University of Texas Austin
anirbanc@oden.utexas.edu

MS48

Machine Learning Data-Driven Closure Models

Computational simulations of dynamical systems often require constitutive or closure models to represent unresolved phenomena, enhance computational efficiency, or correct model form error in the simulation. Data-driven closure models (DDCMs) employ machine learning (ML) to learn the closure term as a function of dynamic system states and have the potential to increase the predictive capability of computational simulations. Our approach decouples the estimation of the closure term from training the DDCM for increased scalability. The closure term is estimated from experimental data using a particle filter, and then training the DDCM with machine learning occurs offline, which increases efficiency. For trustworthiness, accurate uncertainty quantification (UQ) for DDCM predictions is essential. Uncertainties that arise from sparse or noisy experimental data are mapped to the estimated closure term using the particle filter, and then probabilistic ML algorithms propagate uncertainties to the predicted quantity. We demonstrate our methodology on an exemplar from dynamic ecological modeling and provide preliminary results using sparse machine learning regression for the DDCM. *SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.*

Toryn Schafer, Dan Drennan
Texas A&M University
tschafer@tamu.edu, danjdrennan@tamu.edu

Wolfgang Brightenburg
University of Washington
wbrigh@uw.edu

Kyle Neal, Teresa Portone, Moe Khalil
Sandia National Laboratories
kneal@sandia.gov, tporton@sandia.gov,

mkhalil@sandia.gov

MS48

Physics-Informed Score-Based Generative Modeling

Score-based generative models, such as diffusion models, leverage the reversibility of a diffusion process to produce strong results in typical generative tasks, such as image data. This work aims to develop a framework for extending these models to incorporate known dynamics of time-dependent data, specifically data governed by a dynamical system which is at least partially known. We propose the use of a physics-informed neural network (PINN) and accompanying loss to guide the training of score-based models using the models' estimated probability density function and governing dynamics' Liouville equations. We accompany this framework with an implementation of denoising diffusion probabilistic modelling (DDPM) within it.

Patrick Wyrod, Ashesh Chattopadhyay
UC Santa Cruz
pwyrod@ucsc.edu, aschatto@ucsc.edu

Daniele Venturi
Department of Applied Mathematics
University of California Santa Cruz
venturi@ucsc.edu

MS49

Adaptive Importance Sampling for Enhancing Offshore Wind Turbine Reliability

In this work, we build on the reliability study initiated by Y. Liu and M. Chertkov [2024], focusing on the extreme values of key mechanical characteristics – pitch, surge, and heave – of a floating offshore wind turbine (FOWT). Leveraging a comprehensive set of wind and wave patterns that cause significant deviations in FOWT performance, initially identified through brute-force Markov Chain Monte Carlo (MCMC) simulations, we have developed an efficient Adaptive Importance Sampling (AIS) MCMC approach. This method allows us to explore the tails of the probability distributions, uncovering higher and potentially more damaging values of pitch, surge, and heave that are beyond the reach of standard MCMC techniques. Enhanced modeling of large-scale wind fluctuations has enabled us to identify and analyze both known and new rare but dangerous regimes. Notably, using AIS-MCMC, we have pinpointed and examined a surge anomaly driven by rare coherent wind patterns with relatively low mean values, coupled with wave interactions that interfere with turbine control. We conclude by discussing how this approach can be integrated with alternatives, such as Deep Reinforcement Learning and Physics-Informed Diffusion Models, to further explore rare but potentially hazardous mechanical extremes in FOWT operations.

Yihan Liu
Virginia Tech
yihanl21@vt.edu

Misha Chertkov
University of Arizona
chertkov@math.arizona.edu

MS50

Benchmarking Offline Implementations of Deep Material Networks

The availability of high-resolution CT images of materials has allowed creation of detailed virtual material characterization workflows. However, direct multiscale simulations are computationally expensive, and treating inelastic material behaviour requires specific techniques like model-order reduction or neural network models. Such strategies require tedious calibration and are not guaranteed to provide accurate predictions away from the training set. As an alternative, Deep Material networks(DMNs) approximate the complex geometries of a microstructure using neural networks(NNs) built based on simple mechanistic building blocks, i.e., laminates, and linear elastic training only. This DMN microstructure is then solved with the same applied load and identical material models that were used for complex geometries to obtain inelastic predictions with reasonable accuracies. Thermodynamic consistency of the DMNs during inference allows predictions away from the training set. In this talk, we will discuss different implementations of the training process with DMNs and benchmark this technology for industrial use cases.

Pavan Bhat Keelanje Srinivas
Fraunhofer ITWM
pavan.bhat.keelanje.srinivas@itwm.fraunhofer.de

MS50

A combination of OPAL and extreme sparsification for ML-enabled UQ in materials modeling

Towards the development of frameworks for the discovery of machine learning-enabled consti- tutive models, model selection becomes a non-trivial challenge. Providing a structured approach, this study introduces a method for identifying the best predictive computational model from a range of possible mechanistic models for a physical system, in our case, for a constitutive model. To achieve this, the Occam-Plausibility Algorithm for optimal model selection is expanded to the selection of neural networks and network structures. The framework adaptively balances the trade-off between complexity and validity of the candidate models while taking into account hyperparameters associated with neural network structures and the convergence characteristics of the individual optimization schemes. Addition to previous application of this framework to tradi- tional phenomenological models, here we account for complexity of the data, and of the training, especially as we explore transfer learning paradigm. We combine the above with our previously developed extreme-sparsification paradigm for the discovery of interpretable constitutive models.

Jingye Tan
Cornell University
jt894@cornell.edu

Nikolaos Bouklas
Sibley School of Mechanical and Aerospace Engineering
Cornell University
nb589@cornell.edu

MS50

Deep Material Network for Solving Mechanical and Thermal Problems in Heterogeneous Media

Reduced Order Models (ROM) of Direct Numerical Simulations (DNS) capture essential features for homogenization (e.g., constituent properties, microstructure layout) through reduced degrees of freedom while being computa-

tionally cheaper than full DNS. Using machine learning, ROMs enable faster predictions of microstructure properties under various conditions. However, they rely on specific datasets with limited generalizability. In this talk, I will present recent development on the Deep Material Network (DMN), a ROM for mechanical and thermal homogenization problems. DMN integrates micromechanics principles directly into its architecture, achieving faster and more accurate predictions than traditional methods. I will demonstrate how DMN efficiently learns complex relations better than standard machine-learning approaches that do not account for physics. Additionally, I will show how DMN can perform uncertainty quantification and sensitivity analysis for complex heterogeneous structures more rapidly than DNS. This talk highlights DMN's potential in solving multi-physics problems and offers a new perspective on material simulations. SAND2024-10960A. SNL is managed and operated by NTESS under DOE NNSA contract DE- NA0003525.

Remi Dingreville
Sandia National Laboratories
rdingre@sandia.gov

Dongil Shin
Myongji University
Dept. of Chemical Engineering
dongil@mju.ac.kr

Ricardo Lebensohn
Los Alamos National Laboratory
lebenso@lanl.gov

MS50

Synergizing Theory, Machine Learning, and Experiments in Functional Materials Dynamics

Functional materials continue to attract significant research interest for their applications in the areas of tunable multifunctional spintronics, magnetoelectric random access memory, and optoelectronic devices. Notable candidates in this field include the perovskite oxide family, two-dimensional (2D) layered ferroelectrics, and heterostructures. These materials offer considerable flexibility in terms of compositional variations, supercell periodicity, and cation radius mismatches, which are crucial for materials design. Deriving the atomistic level mechanisms behind formation of domains, corresponding dynamics dependent on external fields such as temperature and electric field, are critical to develop a comprehensive understanding behind exhibition as well as tunability of spontaneous polarization in these systems. This presentation will focus on such mechanisms driving functionalities from atomistic simulations in combination with analyzing underlying cause-effect relations. Machine learning-based approaches to establish direct feedback loop between experiments and simulations to study tunable parameters for optimizing target properties will also be discussed. Acknowledgments: This research is sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy.

Ayana Ghosh
Oak Ridge National Laboratory
ghosha@ornl.gov

MS50

Investigating Grain Growth Using a Physics-Regularized Interpretable Machine Learning Model

Physics-based mesoscale models of grain growth have been unable to accurately represent the grain growth behavior of real materials. We are developing the Physics Regularized Interpretable Machine Learning Microstructure Evolution (PRIMME) model learns to predict grain growth directly from grain growth data. The PRIMME algorithm uses a multi-level neural network to predict grain growth in a voxelated domain. It uses a regularization function that encourages evolution that never increases the number of nearest neighbor voxels assigned to different grains. PRIMME helps to interpret and understand its learned grain growth behavior by determining the likelihood of a voxel changing to the grain of neighboring voxels. It was originally trained using data from 2D isotropic simulation results and has been extended to 3D isotropic and 2D anisotropic behavior. We are now training it using experimental data rather than just simulation results.

Michael R. Tonks, Vishal Yadav, Joseph Melville, Lin Yang
University of Florida
michael.tonks@ufl.edu, v.yadav@ufl.edu, joseph.melville@ufl.edu, lin.yang@ufl.edu

Michael Kesler
Oak Ridge National Laboratory
keslerms@ornl.gov

Amanda R. Krause
Carnegie Mellon Universtiy
amandakr@andrew.cmu.edu

Joel Harley
University of Florida
joel.harley@ufl.edu

MS51

A Data-driven Approach to Kernel Learning in PNP Equations and Keller-Segel Model

The Keller-Segel model and Poisson-Nernst-Planck (PNP) equations are foundational mathematical frameworks used to describe chemotaxis in biological systems and ionic transport in electrochemical systems, respectively. In this work, we propose a data driven approach for learning the kernels that forms the foundation of the PDE systems. We demonstrate the effectiveness of our approach through a series of computational experiments, comparing the learned kernels to the true kernels within the Keller-Segel model and the PNP equations.

Chi-An Chen
Illinois Institute of Technology
cchen156@hawk.iit.edu

MS51

A Data-driven Structure Preserving PINN for Learning PDEs

We introduce a novel structure preserving Physics Informed Neural Networks (PINNs) to stabilize the training of PINNs by closely fitting the initial conditions. Through extensive numerical experiments for a wide range of time-dependent PDEs with periodic boundary, particularly diffusion-reaction PDEs, we demonstrate the effectiveness of our approach. Additionally, we present some

real world applications: the estimation of kinetic parameters from medial data and social data to provide data-driven crime modeling. Our findings illustrate the robustness and applicability of our training scheme in complex and real world scenarios.

Baoli Hao
Illinois Institute of Technology
bhao2@hawk.iit.edu

MS51

Efficient Clustering on Riemannian Manifolds by Fréchet Mappings

Symmetric Positive Definite (SPD) matrices, in particular correlation matrices, appear in various applications from applied mathematics and engineering, most notably in neuroimaging applications such as Diffusion Tensor Imaging and Functional Magnetic Resonance Imaging, where they are employed to model the strength of neural connections between different brain sites. Mathematically, the set of SPD matrices is not a vector subspace of the Euclidean space under standard matrix addition and scalar multiplication, however it possesses a smooth manifold structure that can be endowed with a Riemannian metric. Hence the proper measure of similarity between SPD matrices is not an Euclidean distance but a Riemannian distance that can capture the intrinsic geometrical structure of the underlying space. Unfortunately, computation of distances in the Riemannian manifold of SPD matrices becomes very expensive as the matrix size increases so that computing k-means of large SPD matrices directly in the Riemannian setting can be prohibitively expensive. Here we present and demonstrate a novel approach to efficiently cluster data on the space of SPD matrices taking advantage of a specially designed Frchet mapping.

Ji Shi
University of Houston
shijibest@gmail.com

MS52

Trapped Acoustic Waves and Raindrops: High-Order Accurate Integral Equation Method for Localized Excitation of a Periodic Staircase

We present a high-order boundary integral equation (BIE) method for the frequency-domain acoustic scattering of a point source by a singly-periodic, infinite, corrugated boundary. We apply it to the accurate numerical study of acoustic radiation in the neighborhood of a sound-hard two-dimensional staircase modeled after the El Castillo pyramid. Such staircases support trapped waves which travel along the surface and decay exponentially away from it. We use the array scanning method (Floquet-Bloch transform) to recover the scattered field as an integral over the family of quasiperiodic solutions parameterized by on-surface wavenumber. Each such BIE solution requires the quasiperiodic Green's function, which we evaluate using an efficient integral representation of lattice sum coefficients. We avoid the singularities and branch cuts present in the array scanning integral by complex contour deformation. For each frequency, this enables a solution accurate to around 10 digits in a few seconds. We propose a residue method to extract the limiting powers carried by trapped modes far from the source. Finally, by computing the trapped mode dispersion relation, we use a simple ray model to explain an acoustic chirp-like time-domain response that is referred to in the literature as the "raindrop effect".

Fruzsina Agocs
Center for Computational Mathematics
Flatiron Institute
fruzsina.agocs@colorado.edu

MS52

Integration of Singular Functions over Deformable Surfaces Corrected Quadratures, Regularized Quadratures, and Rational Approximation Quadratures

Many elliptic PDEs can be recast as integral equations involving surface integrals $S(x) = \int_\Gamma G(x,y)q(y)\,dS(y)$ of a Green's function $G(x,y)$ multiplying a layer density $q(y)$. Such an integral, called a layer potential, must often be evaluated on the surface itself, $x \in \Gamma$, in which case the Green's function is singular (but $S(x)$ itself is finite). This calls for specialized methods to be used on-surface. We discuss and compare three such specialized methods. The first method is based on local corrections to a standard quadrature method (the trapezoidal rule), in a few grid points around the singularity. For instance, correcting at nine points around the singularity leads to a 5th order scheme. The second method is based on regularizing the Green's function so that a standard quadrature method can be used. For a suitable choice of the regularization, the method can be observed to have 3rd or even 5th order convergence on-surface. The third method is based on evaluating the layer potential along a line off-surface, and extrapolating onto the surface. In particular, we consider extrapolation using rational functions, and compare with polynomials. All methods are discussed in the context of simulating deformable capsules in Stokes flow. The fact that the surface deforms rules out any precomputation that depends on its particular shape. A partition of unity discretization is used to avoid the second-order boundary error of the trapezoidal rule.

Joar Bagge
UT Austin
joar.bagge@austin.utexas.edu

Bowei Wu
UMass Lowell
bowei_wu@uml.edu

George Biros
University of Texas at Austin
Oden Institute
biros@oden.utexas.edu

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

Anna-Karin Tornberg
KTH
akto@kth.se

Pritpal Matharu
KTH Royal Institute of Technology
pritpal@kth.se

MS52

Adaptive Greens Function Integration over Bloch Wavevectors with Explicit Handling of Complex

Singularities

Computing the linear response of a spatially periodic medium often requires integration with respect to a Bloch (quasi-periodicity) wavevector parameter $\mathbf{k}$, over the dual (reciprocal) torus for the lattice. The integrands are analytic in each coordinate of $\mathbf{k}$, apart from complex plane singularities of known generic type near or on the real axis. Two applications in which such integrals must be repeatedly performed are: 1) In electronic structure calculations, a Green's function $G(\mathbf{k})$ must be integrated over the Brillouin zone in order to compute quantities such as density of states. 2) In time-harmonic wave scattering with a periodic medium but localized source, the Floquet-Bloch method requires integration of Helmholtz (say) solutions over their Bloch wavevector. We present high-order accurate results for 1). The integrand is the trace of the inverse of a small matrix that is analytic in $\mathbf{k}$. The integral is performed iteratively, dimension by dimension. For the innermost integral, we show increased efficiency using a variant of adaptive integration which explicitly handles nearby poles. For the next integral we investigate the explicit handling of the square-root type singularities, via quadratic Pad. In both cases multiple nearby singularities at unknown locations are handled robustly without contour deformations.

Alex H. Barnett, Jason Kaye
Flatiron Institute, Simons Foundation
abarnett@flatironinstitute.org, jkaye@flatironinstitute.org

Lorenzo X. Van Muñoz
Massachusetts Institute of Technology
Department of Physics
lxvm@mit.edu

MS52

The Resolution of Singularities by Rational Functions and Mesh Refinement

Singularities frequently appear in the solution of PDEs when the computational domains has edges, interfaces or corners. A popular technique for 2D problems is to include the singularity, if it is known, into the approximation space. This is less feasible for 3D problems, as it becomes harder to analytically characterise the singularities. An alternative is to resolve the singularities, for which the predominant approaches in finite element methods use mesh refinement. It is customary in hp-methods in particular to use geometrically graded meshes. We first survey known results on the resolution of singularities on graded meshes. Yet, the main topic of the talk is its similarity to approximation by rational functions. We demonstrate that similar mathematics underlies both approximation by piecewise polynomials and approximation by rational functions with clustered poles. This is true beyond univariate functions: singular behaviour of PDE solutions near edges can be resolved to high accuracy using multivariate piecewise polynomials or using multivariate rational functions, with the latter being relatively unexplored in comparison. This talk is based on joined work with Nicolas Boull, Astrid Herremans and Nick Trefethen.

Daan Huybrechs
Department of Computer Science
KU Leuven
daan.huybrechs@kuleuven.be

MS53

Second-Order Parallel Tensor Integrators for Dynamical Low-Rank Approximation

A large number of problems in physics and engineering are governed by solutions that exist in high-dimensional phase spaces. Numerical methods for these problems often require substantial memory and computational resources. To address these challenges, we employ dynamical low-rank approximation (DLRA), which evolves the solution on the manifold of low-rank functions. To manage the stiffness that arises from the geometry of this manifold, several high-order time integrators have recently been developed for matrix-valued solutions. This talk presents recent advancements in high-order integrators for DLRA and their extension to high-dimensional problems encountered in kinetic theory, uncertainty quantification, and quantum physics.

Jonas Kusch
Norwegian University of Life Sciences
jonas.kusch@nmbu.no

Dominik Sulz
University of Tuebingen
dominik.sulz@uni-tuebingen.de

MS53

Block Hashed Randomized Leverage Scores Homogenizers for Low-Rank Approximation on Distributed Architectures

Space embeddings offer a dimension reduction technique for high-dimensional data. In this work we introduce a new class of block structured random matrices, block randomized leverage scores homogenizers (block HLSH). We prove this type of matrices are oblivious subspace embeddings: approximating high dimensional matrices with them can be done with high probability independently from the data. Block HLSH generalizes and expands some of the most widely and recently used sparse sketching matrices: SRFT, HRHT and block SRHT. Block HLSH works by first preconditioning the matrix to embed with a novel class of block structured matrices that allow fast matrix-matrix computations. Because of the block structure, this method parallelizes naturally. Then a second dimension reduction is done with a sparse random matrix. The combination between fast matrix-matrix multiplication, parallelization, and sparsity makes this new approach computationally efficient yet easy to implement. In the context of the J-L lemma, we prove that such approach has optimal sketching dimension just as Gaussian matrices. In combination with Nystrm approximation, numerical experiments illustrate the performance of block HLSH compared to existing sketching matrices.

Mariana G. Martinez Aguilar
EPFL
mariana.martinezaguilar@epfl.ch

Laura Grigori
EPFL and PSI, Switzerland
laura.grigori@epfl.ch

MS53

Sparse Grid Methods for Particle-in-Cell Schemes

The Vlasov model of plasmas is six-dimensional in general. The most prevalent method combatting this high dimensionality is particle-in-cell (PIC), which represents the system via particles interacting with fields on a spatial

mesh. Use of particles introduces slow-converging sampling errors, while the spatial mesh permits only partial mitigation of the curse of dimensionality. We show that using sparse grids with PIC gives complexity depending only logarithmically on dimension, thereby dramatically reducing sampling noise. We report progress combining sparse PIC with symplectic and implicit methods, then discuss ongoing work toward adaptive coordinate selection for sparse grids to reduce astronomically large coefficients in front of favorable asymptotic scalings. *Prepared by LLNL under Contract DE-AC52-07NA27344.

Lee Ricketson
Lawrence Livermore National Laboratory
lee.ricketson@gmail.com

MS53

Sparse-Grid Discontinuous Galerkin Methods for the VlasovPoissonLenardBernstein Model

Sparse-grid methods have recently gained interest in reducing the computational cost of solving high-dimensional kinetic equations. In this talk, we will construct adaptive and hybrid sparse-grid methods for the VlasovPoissonLenard-Bernstein (VPLB) model. This model has applications to plasma physics and is simulated in a 1x3v slab geometry. We use the discontinuous Galerkin (DG) method as a base discretization due to its high-order accuracy and ability to preserve important structural properties of partial differential equations. The method utilizes a multiwavelet basis expansion to determine the sparse-grid basis and the adaptive mesh criteria. We will analyze the proposed sparse-grid methods on a suite of three test problems by computing the savings afforded by sparse-grids in comparison to standard solutions of the DG method. Results of this talk are obtained using the adaptive sparse-grid discretization library ASGarD.

Stefan Schnake
Oak Ridge National Laboratory, U.S.
schnakesr@ornl.gov

Coleman Kendrick
Lawrence Livermore National Laboratory
kendrick6@llnl.gov

Eirik Endeve
Oak Ridge National Laboratory
University of Tennessee, Knoxville
endevee@ornl.gov

Miroslav Stoyanov
Oak Ridge National Lab
mkstoyanov@gmail.com

Steven Hahn
Oak Ridge National Laboratory
hahnse@ornl.gov

Cory Hauck
Oak Ridge National Laboratory and University of Tennessee
hauckc@ornl.gov

David Green
Commonwealth Scientific and Industrial Research Organisation
d.green@csiro.au

Phil Snyder
Oak Ridge National Laboratory
snyderpb@ornl.gov

John Canik
Type One Energy Group
john.canik@typeoneenergy.com

MS53

Time Evolution in the Quantized Tensor Train Format with CUR and SVD Factorization

Quantized tensor trains (QTTs) are a multiscale low-rank ansatz that can potentially enable efficient time integration of initial value problems. Typically, QTTs are decomposed using one of two factorizations: the singular value decomposition, which provides the optimal low-rank approximation in the L2-sense, or the CUR factorization, which exactly captures the values at a sub-selected set of indices. In this talk, we will compare local time integration schemes that arise when using these two different QTT decompositions. We investigate their performance for both linear and nonlinear systems, monitoring numerical stability, accuracy, and efficiency. Finally, we compare the local time integration schemes to global schemes to determine when it may be more advantageous to use one over the other.

Erika Ye
Lawrence Berkeley National Laboratory
erikaye@lbl.gov

MS54

Generative AI Surrogate Models for Particle-Based Kinetic Computations Using Normalizing Flows and Diffusion Models

Particle-based kinetic computations require the numerical integration of stochastic differential equations (SDE) in which the deterministic part describes the Hamiltonian orbit dynamics and the stochastic part collisions and diffusion processes. A potential limitation of this Monte-Carlo approach is that, to avoid statistical sampling errors, the SDE need to be solved for very large ensembles of particles. To overcome this limitation, we present a pseudo-reversible normalizing flow (PR-NF) generative artificial intelligence method for the acceleration of the integration of SDE. [Yang et.al., SIAM journal of Scientific Computing 46, (4) C508-C533 (2024)]. After training, the PR-NF model can directly generate samples of the SDEs final state without simulating trajectories. A convergence analysis using the Kullback-Leibler divergence metric is presented, along with benchmark numerical examples verifying the accuracy and efficiency of the proposed method. Complementing these results we also present an alternative approach based on the use of diffusion models. To illustrate the practical use of these methods we present applications to magnetically confined plasmas for controlled nuclear fusion and to transport in fluids exhibiting chaotic advection in 3D.

Diego Del-Castillo-Negrete, Minglei Yang, Guannan Zhang
Oak Ridge National Laboratory
delcastillod@ornl.gov, yangm@ornl.gov, zhangg@ornl.gov

MS54

Opportunities for ML/AI and System Identifica-

tion in Fusion Control and Design

Designing and controlling a fusion plasma requires predicting and observing complex nonlinear behavior on fast time scales. In this context Machine Learning (ML) methods provide a powerful tool to bridge gaps between current high-fidelity/offline models and online/optimization-relevant models. This talk will focus on two areas relevant to plasma control and design optimization for magnetically-confined fusion concepts: 1) Use of ML techniques to build smart diagnostics and support sensor fusion across distinct diagnostics to enable generation of a complete picture of the plasma state from limited, nonlinearly-convolved diagnostic signals. 2) Application of system identification techniques to build fast reduced-order models of relevant dynamics that are interpretable and support well-defined stability characteristics and bounds both of which are important for control applications and licensing considerations. In these areas, ML has the potential to significantly impact the timeline for commercial fusion power through fast models for design optimization to minimize cost and ML-enabled observers and state models for real time control of optimal operational scenarios.

Chris Hansen
Columbia University. Applied Physics and Applied Mathematics
christopher.hansen@columbia.edu

Ryan Forelli
Lehigh University
rff224@lehigh.edu

Samuel Freiberger
Columbia University
swf2112@columbia.edu

Alan Kaptanoglu
New York University
aak572@nyu.edu

Jeffrey Levesque
Columbia University
jpl2131@columbia.edu

Nhan Tran
Fermilab
ntran@fnal.gov

Yumou Wei
MIT
yumou.wei2013@gmail.com¿

MS54

Neural Network-Based Surrogate Models for Fusion Plasma Instabilities and Transport

Magnetic confinement fusion is poised to achieve high fusion gain, with fusion pilot plants envisioned on a decadal timescale. Turbulent transport is the biggest factor determining the confinement in a fusion device, and by extension, its cost. Turbulent transport can be predicted using gyrokinetic simulations, which are accurate but computationally expensive. Machine Learning (ML) based surrogate models represent a promising avenue for efficient prediction and optimization of transport dynamics. This presentation will discuss ML-based surrogates/metamodels tailored to addressing two critical issues in turbulent transport: (1) transport barriers in tokamaks, and (2) confinement optimization in stellarators. I will discuss a Bayesian framework that integrates experimental and multi-fidelity simulation data to simulate transport barriers, supplemented by surrogates for expedited modeling and uncertainty quantification. For stellarators, I will outline the development of innovative neural networks designed to optimize turbulent transport. Additionally, I will introduce the MGKDB simulation database, a key piece of foundational infrastructure for these activities, being developed to consolidate and curate gyrokinetic simulation data from multiple gyrokinetic codes and activities throughout the fusion community. This database will serve as a crucial infrastructure to underpin the development of robust surrogate models, facilitating the design and optimization of fusion devices.

David Hatch
University of Texas at Austin
drhatch@austin.utexas.edu

Craig Michoski
Sapientai
michoski@sapient-a-i.com

Todd Oliver
The University Of Texas-Austin
oliver@oden.utexas.edu

Aaro Jarvinen, Daniel Jordan
VTT Finland
aaro.jarvinen@vtt.fi, daniel.jordan@vtt.fi

Mike Kotschenreuther
ExoFusion
mtk@fusionenergyassociates.com

Joseph Schmidt, Myriam Hamed, Ping-Yu Li, Cole Stephens
University of Texas at Austin
joseph.schmidt@utexas.edu,
myriam.hamed@austin.utexas.edu,
pingyu.li@austin.utexas.edu,
cole.stephens@austin.utexas.edu

Swadesh Mahajan
University of Texas at Austin
Shiv Nadar University
mahajan@mail.utexas.edu

MS54

A Machine Learning Normalizing Flow Method for Uncertainty Quantification in Fusion Plasmas

Understanding and quantifying uncertainties in physical models, numerical simulations, and experimental data is crucial. We propose a conditional pseudo-reversible normalizing flow (PR-NF) method to construct surrogate models of physical systems in the presence of noise, enabling efficient quantification of forward and inverse uncertainty propagation. The PR-NF model excels in determining conditional distributions for these processes. We apply the PR-NF model to study runaway electron dissipation with impurity injection in tokamaks. Developing a robust disruption mitigation system for ITER, and future machines, is essential to prevent reactor wall damage, yet the optimal impurity type and deposition method remain unresolved. We discuss how the spatiotemporal density profiles of injected impurities impact runaway electron dissipation. Also, to demonstrate the effectiveness of our ap-

proach, we present two validation tests and an application to runaway electron dissipation.

Minglei Yang, Diego Del-Castillo-Negrete, Guannan Zhang
Oak Ridge National Laboratory
yangm@ornl.gov, delcastillod@ornl.gov, zhangg@ornl.gov

MS55

Robust Optimal Design of Bayesian Inverse Problems Governed by PDEs

We consider optimal design of experiments for nonlinear Bayesian inverse problems governed by partial differential equations (PDEs). An optimal design is one that optimizes the statistical quality of the solution of the inverse problem, as measured by some utility. The computed optimal design, however, depends on the modeling assumptions encoded in the governing PDEs or the parameterization of the observation model. In cases where some of these elements are subject to large uncertainties, it is prudent to follow a robust optimal experimental design (ROED) approach. We follow a worst-case scenario approach, and develop a scalable computational framework for robust optimal design of nonlinear Bayesian inverse problems, which incorporates a probabilistic optimization paradigm for the resulting combinatorial max-min optimization problem. In this talk, we focus on Bayesian ROED, where the goal is to maximize information gain in presence of uncertainties in the measurement error model. The proposed approach is illustrated in the context of optimal sensor placement for a coefficient inverse problem governed by an elliptic PDE.

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Abhijit Chowdhary
North Carolina State University
achowdh2@ncsu.edu

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

MS55

Optimization of Total Variation-Regularized Functions using Inexact Proximal Solves

Total Variation (TV) optimization penalizes the gradient of a control variable or state; such regularization has found utility in image processing, inverse problems, and topology optimization. We will address two fundamental challenges with TV optimization: (i) the typical slow convergence of existing TV optimization methods, and (ii) the inexact evaluation of the TV proximity-operator. The proximal operator of the TV function can be just as difficult as solving the original problem. We allow for inexact proximal subproblem solves with an error-tolerance governed by a trust-region globalization scheme. We propose an algorithm for general nonsmooth, nonconvex TV optimization and illustrate our technique on imaging and topology problems.

Robert Baraldi
Sandia National Laboratories
Optimization & Uncertainty Quantification
rjbaral@sandia.gov

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS55

SiMPL method for Topology Optimization

We present a rigorous convergence analysis of a new method for topology optimization: Sigmoidal Mirror descent with a Projected Latent variable. SiMPL provides point-wise bound preserving design updates and faster convergence than other popular first-order topology optimization methods. Due to the strong bound preservation, the method is exceptionally robust, as demonstrated in numerous examples. Furthermore, it is easy to implement with clear structure and analytical expressions for the updates. Our analysis covers two versions of the method, characterized by the employed line search strategies. We consider a modified Armijo backtracking line search and a Bregman backtracking line search. Regardless of the line search algorithm, SiMPL is guaranteed to deliver a minimizing sequence of designs that converges to a stationary point of the compliance functional. In addition, the numerical experiments demonstrate apparent mesh-independent convergence of the algorithm.

Dohyun Kim
Brown University
dohyun_kim@brown.edu

Boyan Lazarov
Lawrence Livermore National Laboratory
lazarov2@llnl.gov

Thomas M. Surowiec
Simula Research Laboratory
thomasms@simula.no

Brendan Keith
Brown University
brendan_keith@brown.edu

MS56

Element Learning: a Systematic Approach of Accelerating Finite Element-Type Methods Via Machine Learning, with Applications to Radiative Transfer

Over the past decade, neural networks and machine learning have emerged as transformative technologies, addressing complex challenges across various domains. In scientific computing, particularly for the numerical solution of partial differential equations (PDEs), machine learning has shown promise in accelerating computations. However, the high training costs for large problems and the loss of key advantages of traditional numerical methods, such as interpretability, reliability, and applicability to complex geometries, remain significant challenges. This talk introduces a systematic approach, termed "element learning", aimed at accelerating finite element-type methods through machine learning. This approach retains the desirable features of finite element methods while substantially reducing training costs. It draws on principles from hybridizable discontinuous Galerkin (HDG) methods, replacing HDG's local solvers with machine learning models. We demonstrate the

efficacy of this approach through numerical tests on the parametric radiative transfer equation, relevant to shortwave radiation calculations for clouds. Our results show a significant speed-upranging from 5 to 10 timescompared to standard finite element methods, with training completed in minutes on a single RTX 30 series GPU.

Shukai Du
Syracuse University
sdu113@syr.edu

Samuel N. Stechmann
University of Wisconsin-Madison
stechmann@wisc.edu

MS56

Space-Angle Discontinuous Galerkin Method for 2D Rte with Diffusively and Specularly Reflective Boundary Conditions

A space-angle discontinuous Galerkin (saDG) method for solving the Radiative Transfer Equation (RTE) in arbitrary 2D domains is proposed. The space-angle domain is fully discretized by the DG method. An angular decomposition approach, resulting in an iterative solution process, is adopted to accelerate the solution and reduce memory usage. The performance of the method is analyzed in a square geometry with different types of boundary conditions. For nonscattering and open systems, the number of iterations each time increases by one with one additional reflective boundary. For other cases, we show that the number of iterations or convergence rate of the iterative scheme for the specular reflection problems is not only affected by the number of the reflective surfaces, but also by the scattering and extinction properties of the media. Uniformly and directionally diffuse reflective boundary conditions are compared in the same medium. Finally, the localization at the focal point of a parabolic reflector is studied.

Hang Wang
New Jersey Institute of Technology
hang.wang@njit.edu

Reza Abedi
University of Tennessee Space Institute
University of Tennessee Space Institute
rabedi@utsi.edu

Saba Mudaliar
Air Force Research Laboratory, Sensors Directorate, 2241 Avi
saba.mudaliar@us.af.mil

MS58

Adaptive Surrogate Modeling of Coupled Hall Thruster Plasma and Plume Simulations

Fluid models of plasma in a Hall thruster enable rapid prediction of system performance across a range of operating conditions and uncertainties. Due to the multiscale nature of the plasma, ad hoc closures for small-scale kinetic effects must be incorporated into the fluid models, thus introducing uncertain parameters which must be tuned to fit experimental data. Furthermore, the discharge plasma couples to the expanding thruster exhaust plume, which couples to downstream surfaces such as the spacecraft body or the test facility. The result is a highly coupled multidisciplinary system with significant uncertainty due to both aleatoric and epistemic sources. To this end, we propose a multifidelity, multidisciplinary approach for Bayesian calibration of the integrated Hall thruster system. We approximate the coupled system model with an adaptive surrogate method based on sparse grids to enable computationally expensive outer-loop analyses. We additionally perform uncertainty quantification over the calibrated model to assess the effects of epistemic and aleatoric uncertainty. Based on our results and a global sensitivity analysis, we make several recommendations for future model refinement.

Josh Eckels, Thomas Marks, Gabriel Mora, Alex Gorodetsky
University of Michigan
eckelsjd@umich.edu, marksta@umich.edu, morag@umich.edu, goroda@umich.edu

MS58

Partitioned Coupling of Multifidelity, Multiphysics Models Using Optimization-Based Coupling

Problems involving coupling over a nonoverlapping interface abound, with common examples including virtual interfaces introduced for the purpose of domain decomposition and physical interfaces such as are found in, e.g., multimaterial solid interaction and fluid-structure interaction. Partitioned approaches for solution of coupled problems enable the use of subdomain solvers developed by domain experts and are amenable to a plug-and-play style of framework. As model reduction techniques continue to mature, it is desirable to use cheaper, more efficient models where possible. This fits closely with the partitioned approach to solving coupled problems. It is towards this goal of solving multifidelity coupled problems involving reduced order models (ROMs) and full order models (FOMs) that we extend a partial differential equation (PDE) constrained optimization technique, developed for elliptic FOMs and later extended to parabolic FOMs and multiphysics, to ROM-ROM coupling. The use of a ROM in one or more subdomains presents a challenge with respect to reduced adjoint systems being required to be solved as part of the technique. We present numerical studies demonstrating the accuracy of the approach along with an investigation of snapshot generation techniques with which to generate a reduced basis for the adjoint systems.

Paul Kuberry
Sandia National Laboratories
pakuber@sandia.gov

Elizabeth Hawkins
Clemson University
evhawki@clemson.edu

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

MS58

Optimizing Coupled Systems: Insights from Co-Design Imaging and Optical Chemistry

Coupling interacting components in computational modeling presents significant opportunities for end-to-end optimization. This talk highlights advances in tightly-coupled system optimization, focusing on the inverse design of optical devices for computational imaging and chemical sensing. I will also discuss how physics-enhanced surrogates for PDEs and multi-fidelity modeling could capture

low-dimensional feasible spaces in end-to-end optimization. This work exemplifies the transformative potential of end-to-end approximate modeling in scientific applications and underscores future directions for advancing co-design in coupled systems.

Raphaël Pestourie
Georgia Institute of Technology, U.S.
rpestourie2@gatech.edu

Wenchao Ma
Massachusetts Institute of Technology
mawc@mit.edu

Steven Johnson
MIT
stevenj@math.mit.edu

Lu Lu
Yale University
lu.lu@yale.edu

Zin Lin
Virginia Tech
zinlin@vt.edu

MS59

A Sparse-Grids Filter for Structure-Preserving Electromagnetic Pic Methods

In the past decade, the idea of using the sparse-grid recombination technique as a means to reduce noise in electrostatic particle-in-cell (PIC) methods has generated substantial interest. Likewise, structure-preserving electromagnetic (EM) PIC methods based on discretizing the Lagrangian or Hamiltonian structure of the Vlasov-Maxwell system have been a prominent direction of inquiry for their ability to automatically conserve known invariants of the continuous dynamics (e.g. energy and Gausss laws). This work unites these two ideas to provide a structure-preserving EM PIC method with a noise-reducing sparse-grids filter. Specially designed filter matrices, which are sparse operators with a convenient Kronecker structure, are inserted into the discrete Lagrangian of a general variational EM PIC method in such a manner that the symmetries of the Lagrangian are preserved. This yields a filtered, variational EM PIC method which retains all the structure-preserving properties of a usual variational EM PIC method, which enjoys the noise-reduction properties of a sparse-grids PIC method, and whose implementation involves only a small, inexpensive modification of existing structure-preserving EM PIC methods. The strategy is general and should be amenable to a broad class of variational EM PIC methods. Numerical examples are provided using the GEMPIC method.

William Barham, Josh Burby
University of Texas at Austin
william.barham@utexas.edu,
joshua.burby@austin.utexas.edu

Eric Sonnendrücker
Max-Planck-Institut für Plasmaphysik, NMPP division, and
Technische Universität München, Department of Mathematics
eric.sonnendruecker@ipp.mpg.de

MS59

PEPC: an Open-source Multi-physics Framework for Mesh-free Simulation of N-body Systems

PEPC is a multi-physics community tree-code developed and hosted by JSC over the last two decades [P. Gibbon et al. (2024), DOI: 10.5281/zenodo.11035167.]. At its core the code utilises a hierarchical tree structure originally based on the Salmon-Warren Hashed Oct Tree scheme to perform rapid force summation of a dynamical N-body system interacting via long-range potentials on heterogeneous superomputer architectures. Depending on the physical choice of potential, the algorithm can be applied to diverse fields such as plasma, gravitational systems and vortex fluids [Durante et al., Math. Comp. in Simulation 225, 528 (2024)]. In this talk we highlight some of the recent key developments in the PEPC framework, covering algorithmic and physical aspects and such as GPU porting, time-integration schemes, management of dynamical particle populations [Chew et al., Plasma Phys. Contr. Fusion 63, 045012 (2024)] and inclusion of self-generated magnetic fields [Siddi et al., Phys. Plasmas 24, 082103 (2017)] in charged-particle systems.

Paul Gibbon, Dirk Broemmel, Junxian Chew, Robert Speck
Forschungszentrum Juelich GmbH
Juelich Supercomputing Centre
p.gibbon@fz-juelich.de, d.broemmel@fz-juelich.de, j.chew@fz-juelich.de, r.speck

MS59

A Particle Method with Adaptive Refinement and Remeshing for the Vlasov-Poisson Equations

A forward semi-Lagrangian scheme called FARSIGHT is presented for collisionless electrostatic plasmas described by the 1D1V Vlasov-Poisson equations. The phase space distribution is represented by Lagrangian quadrilateral panels having a hierarchical tree structure, where each panel is a 3×3 particle grid that is tracked by 4th order Runge-Kutta time stepping. The electric field is expressed as a convolution integral of the charge density with a regularized Green's function. The particles are remeshed every time step using biquadratic interpolation on each panel, and the panels are adaptively refined to resolve the phase space distribution. The electric field integral is discretized by the trapezoid rule and the discrete sums are computed by a GPU-accelerated barycentric Lagrange treecode. Numerical results are presented for Landau damping, two-stream instability, and halo formation in a mismatched thermal sheet beam.

Ryan Sandberg
Lawrence Berkeley National Laboratory
rsandberg@lbl.gov

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

Alexander Thomas
University of Michigan

agrt@umich.edu

MS60

Numerical Analysis and Acceleration of Particle-Particle Particle-Mesh Method Using Cabana

Particle-in-cell (PIC) methods describe advective processes via particle discretizations for partial differential equations. Our objectives are two-fold: accelerate the Methods of Local Corrections (MLC) with Cabana, a performance portable library for particle applications, to solve vortex methods in 3D and develop formal numerical analysis for PIC methods via numerical experiments using MLC. MLC is a particle-particle particle-mesh method where the velocity field induced by the vortices is calculated in two steps separating the near and far field. MLC exercises the full capabilities and data choreography of Cabana where N-body calculations and grid to particle interpolations are accelerated. Colella and his collaborators at LBNL, following the ideas of Cottet, have developed a formal, numerical analysis framework for PIC methods. They have demonstrated its potential for the design of PIC methods with improved accuracy and efficiency for 1+1D kinetic problems. For 1+1D kinetics problems, the deformation matrix is carried as an auxiliary variable and used to remap adaptively when particles have deformed significantly from the original grid which causes large interpolation errors. We applied this approach to 2D vortex methods and evaluated potential error indicators. The eigenvalue of R from QR decomposition shows correlation with regions of high error. Using MLC, we explore adaptive remapping in 3D with the eigenvalue of R as a remapping error indicator.

Zoe Barbeau
Stanford University
zbarbeau@stanford.edu

Het Y. Mankad
Carnegie Mellon University
mankadhy@ornl.gov

Het Y. Mankad, Sam Reeve
Oak Ridge National Laboratory
mankadhy@ornl.gov, reevest@ornl.gov

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

MS60

Randomized Preconditioning Based on Approximate Range Deflation

We introduce a novel randomized preconditioning technique for solving large-scale linear systems. We mostly focus on regularized symmetric positive definite systems, although our technique also works for systems with unknown regularization parameters and non-symmetric systems. Our preconditioner can be categorized as a deflation-based preconditioner. Theoretical analysis of the preconditioner, when paired with an iterative solver like preconditioned conjugate gradient, suggests this method can outperform other competing preconditioners when the trailing spectral decay is slow. We validate the effectiveness of this method by applying it to linear systems arising in data science and numerical optimization, and we also compare against state-of-the-art methods.

Oleg Balabanov
INRIA
olegbalabanov@gmail.com

Caleb Ju
Georgia Institute of Technology
calebju4@gatech.edu

Kaiwen He
Purdue University
he788@purdue.edu

Michael Mahoney
University of California, Berkeley
mmahoney@stat.berkeley.edu

MS60

Investigating the Effect of Wind Shear on Wind Turbine Rotor Aerodynamics in Large Eddy Simulations

Variations in wind speed and direction in the atmosphere can affect wind turbine structural loading and power production. The effects of these variations increase as turbines become larger and extend farther into the atmospheric boundary layer, where wind conditions can be more complex than those near the surface. Conventional wind turbine power models assume that incident wind conditions do not affect airfoil efficiency or induced rotor velocity, and are therefore limited in their ability to account for the aerodynamic effects of shear that modify turbine power production. In this study, we investigate an actuator disk in large eddy simulations to resolve the aerodynamic interactions between the disk and inflow wind profiles. These simulations demonstrate that rotor-averaged induction (i.e., decreases in wind speed due to fluid blockage at the turbine) increases monotonically as the magnitude of direction shear over the rotor increases, lowering the disk velocity and power production. For wind conditions near uniform, standard 1D momentum theory overpredicts induction on the rotor by 2%, and underpredicts it by as much as 3% as the average direction shear over the rotor increases. Since power production scales cubically with wind speed, these results point to potentially large errors in estimated power production with current models. We investigate the aerodynamic effects of shear with the goal of more accurately modeling turbine performance in complex wind conditions.

Storm Mata, Kirby Heck, Michael Howland
Massachusetts Institute of Technology
storm22@mit.edu, kheck@mit.edu, mhowland@mit.edu

MS60

Causal Inference for Spatio-Temporal Interventions

In our interconnected world, many critical decisions, from climate policy to public health interventions, must be made in complex environments where both space and time play crucial roles. Understanding the true effects of these decisionssuch as how a new policy might influence regional climate patterns or how an intervention could impact disease spread over timerequires more than correlational analysis; it demands robust causal inference. This talk will explore the key role of causal inference in evaluating policies and interventions within spatio-temporal contexts. I will begin by highlighting the unique challenges posed by spatio-temporal data, including issues of high dimensionality and

complex dependencies across space and time. These challenges make it difficult to apply traditional causal inference methods effectively. To address these issues, I will explore how advances in machine learning, particularly spatio-temporal modeling and counterfactual reasoning with neural networks, can serve as powerful tools for uncovering causal relationships. These techniques allow us to accurately and reliably estimate causal effects, even when faced with the complexities inherent in spatio-temporal data. Through this discussion, I aim to provide a framework for researchers and practitioners to better understand and apply causal inference methods to spatio-temporal data, ultimately leading to more informed and effective decision-making in complex, dynamic environments.

Antonia M. Oprescu
Cornell University
amo78@cornell.edu

Shinjae Yoo
Brookhaven National Lab
sjyoo@bnl.gov

MS61

Oscillation-Free Numerical Schemes for Biot's Model and Their Iterative Coupling Solution

Single-phase flow problems on deformable porous media are modeled by means of the so-called Biots model. Several challenges appear in the numerical solution of this model. On the one hand, it is important to choose appropriate discretization schemes that avoid the appearance of spurious oscillations in the pressure approximation when low permeabilities and/or small time steps are considered. On the other hand, the efficient solution of the large-sparse systems of equations arising after discretization also is challenging. In this work, for different finite-element discretizations of Biots model, we propose a new stabilized scheme that provides numerical solutions that are free of non-physical oscillations, and that, at the same time, allows us to iterate the fluid and mechanic problems in a convergent way to obtain the solution of the whole coupled system. We present numerical results illustrating the robust behavior of both the stabilization and iterative solver with respect to the physical and discretization parameters of the model.

James H. Adler
Tufts University
james.adler@tufts.edu

MS61

Monolithic Multigrid Using Point Smoothers for Systems with Saddle Point Structure

Systems of linear equations with saddle point structure arise in naturally during the simulation of fluid flows. When discretizing the associated PDEs on structured meshes and when suitable boundary conditions are applied the blocks have Toeplitz structure. Multigrid methods for Toeplitz matrices have been studied extensively in the past. They appear in many applications, most notably scalar PDEs. Recently, we focussed on the analysis of block matrices with Toeplitz blocks [Bolten, Donatelli, Ferrari and Furci, SIMAX 2022]. Building on these findings and using results from Notay [Notay, Numer. Math. 2016] we started working on monolithic multigrid methods for systems with saddle point structure [Bolten, Donatelli, Ferrari and Furci, LAA 2023; Bolten, Donatelli, Ferrari and Furci, APNUM 2023]. Using a suitable transformation and plain point smoothers convergence of the resulting two-level method can be shown. Further, the technique can be applied recursively to obtain a multigrid method. We present the method and the techniques used, as well as the obtained results and numerical examples that demonstrate the efficiency in the case of fluid flows.

Matthias Bolten
University of Wuppertal
bolten@math.uni-wuppertal.de

Marco Donatelli
University of Insubria
marco.donatelli@uninsubria.it

Paola Ferrari
University of Wuppertal
University of Wuppertal
ferrari@uni-wuppertal.de

Isabella Furci
University of Genoa
isabella.furci@dima.unige.it

MS61

Coupled Ordering Techniques for Schur Complement Based Preconditioners of Saddle Point Problems

Fluid flow problems can be modeled by the Navier-Stokes or, after linearization, by the Oseen equations. Their discretization results in discrete saddle point problems. These systems of equations are typically very large and need to be solved iteratively. Standard (block) preconditioning techniques rely on an approximation of the (inverse) Schur complement. In this talk, we discuss coupled ordering techniques for the pressure and velocity unknowns that facilitate the computation of such Schur complement approximations. In particular, such orderings can be combined with block clustering strategies in the construction of hierarchical matrices and accelerate the construction of hierarchical LU factorizations of the Schur complement. Numerical results illustrate the performance of the resulting saddle point preconditioners.

Sabine Le Borne
Hamburg University of Technology
Department of Mathematics
leborne@tuhh.de

Jonas Grams
Hamburg University of Technology
jonas.grams@tuhh.de

MS61

Accelerating Computation of Unstable Eigenmodes for Flows

In this talk we will look at recent results in accelerating eigenvector computations to determine unstable modes for Couette flow. We will apply a recently developed adaptive momentum technique to accelerate an implicitly defined shifted power iteration. We will introduce the problem, the methods, and the highlights of the theory for the momentum-type acceleration.

Sara Pollock
University of Florida

s.pollock@ufl.edu

MS61

A Computational Study on Concentrated Vortex Lines in Transitional Pipe Flow

Concentrated vortex lines appear in many flows, such as aerodynamics, turbulence, and weather systems. Although significant past work has focused on understanding the behavior of isolated vortices, the interaction of vortices with walls and shear flows also plays an important role in many systems. To study these interactions, we have developed tools to simulate an experiment which generates streamwise vortices in pipe flow at the laminar/turbulent transition regime. Our simulations employ a second-order projection method with adaptive octree grids and supra-convergent finite difference schemes to solve the incompressible Navier-Stokes equations across a wide range of Reynolds numbers. We use a level-set formulation and efficient tree-traversal algorithms to represent the realistic solid obstructions that are used in these experiments. A key feature of our simulation toolbox is the ability to augment and fine-tune the solver to match the experimental apparatus, allowing us to visualize and compare our results with experimental data easily.

Scott R. West
Department of Applied Mathematics
University of California, Merced
swest10@ucmerced.edu

Cole Cooper, Dustin Kleckner
Department of Physics
University of California, Merced
ccooper10@ucmerced.edu,

Maxime Theillard
University of California, Merced
mtheillard@ucmerced.edu

MS62

Lie Group Variational Collision Integrators for a Class of Hybrid Systems

Hybrid systems are dynamical systems that exhibits both continuous and discrete dynamics. The state of a hybrid system changes either continuously by the flow described by differential equations or discretely following some jump conditions. A canonical example of a hybrid system is the bouncing ball, imagined as a point mass, over a horizontal plane. The extension of this problem to 3-dimensions, wherein the bouncing body is rigid and convex, is rather complex, especially in the case of sharp corner impacts; in fact, these systems have unilateral constraints that describe the collision surface. We study such problems with perfectly elastic collisions and the Lie group variational collision integrators (LGVCI) are derived. The advantage of these frameworks is that they yield a global description of the system, in contrast to local representations such as Euler angles. Furthermore, in high-precision physics engine and graphics dynamics, the integrator becomes a foundation, and its extensions with inelastic collisions and friction can be derived to fully actualize the engine. This is also naturally applicable to problems in optimal control with similar nonlinear manifold constraints. In particular, these constraints and optimal control problems arise in robotics and multi-body dynamics.

Melvin Leok
University of California, San Diego
Department of Mathematics
mleok@ccom.ucsd.edu

Khoa Tran
University of California, San Diego
ktran@ccom.ucsd.edu

MS62

Fluid Implicit Particles on Coadjoint Orbits

We propose Coadjoint Orbit FLIP (CO-FLIP), a high order accurate, structure preserving fluid simulation method in the hybrid Eulerian-Lagrangian framework. We start with a Hamiltonian formulation of the incompressible Euler Equations, and then, using a local, explicit, and high order divergence free interpolation, construct a modified Hamiltonian system that governs our discrete Euler flow. The resulting discretization, when paired with a geometric time integration scheme, is energy and circulation preserving (formally the flow evolves on a coadjoint orbit) and is similar to the Fluid Implicit Particle (FLIP) method. CO-FLIP enjoys multiple additional properties including that the pressure projection is exact in the weak sense, and the particle-to-grid transfer is an exact inverse of the grid-to-particle interpolation. The method is demonstrated numerically with outstanding stability, energy, and Casimir preservation. We show that the method produces benchmarks and turbulent visual effects even at low grid resolutions.

Sina Nabizadeh
University of California, San Diego
mnabizad@ucsd.edu

MS62

Stokes Flow of an Evolving Fluid Film with Arbitrary Shape and Topology

The dynamics of evolving fluid films in the viscous Stokes limit is relevant to various applications, such as the modeling of lipid bilayers in cells. While the governing equations were formulated by Scriven in 1960, solving for the flow of a deformable viscous surface with arbitrary shape and topology has remained a challenge. In this study, we present a straightforward discrete model based on variational principles to address this long-standing problem. We replace the classical equations, which are expressed with tensor calculus in local coordinate, with a simple coordinate-free, differential-geometric formulation. The formulation provides a fundamental understanding of the underlying mechanics and directly translates to discretization. We construct a discrete analogue of the system using the Onsager variational principle, which, in a smooth context, governs the flow of a viscous medium. In the discrete setting, instead of term-wise discretizing the coordinate-based Stokes equations, we construct a discrete Rayleighian for the system and derive the discrete Stokes equations via the variational principle. This approach results in a stable, structure-preserving variational integrator that solves the system on general manifolds.

Cuncheng Zhu
University of California San Diego
cuzhu@ucsd.edu

MS63

Aerodynamic Sensitivities Over Separable Shape

Tensors

We present a sensitivity analysis of wind-turbine airfoil representations informed by separable shape tensors. The shape representation uniquely benefits the airfoil design process by isolating various well-studied shape characteristics, such as general linear scale variations, and facilitates improved regularization of nonlinear shape deformations. The design domain is informed by a principal geodesic analysis of separable shape tensors given a curated database containing tens of thousands of suitable engineering airfoils. Over the pushforward of this submanifold domain, maximum mean discrepancy of the joint distribution of aerodynamic quantities informs a dramatic dimensionality reduction of the learned parameter space. This simple numerical experiment showcases a novel approach for retaining design effectiveness while promoting regularity of the shape representations. Finally, we generate novel reduced-dimension airfoil designs and use the HAM2D RANS solver to inform subsequent sensitivity analyses to assert consistency of parameter influence on the aerodynamic quantities. We also explore low-dimensional polynomial ridge approximations to motivate physical intuitions and offer explanations of the approximated sensitivities.

Zach Grey
National Institute of Standards and Technology
Applied and Computational Mathematics Division
zachary.grey@nist.gov

Olga Doronina, Bumseok Lee
National Renewable Energy Laboratory
olga.doronina@nrel.gov, bumseok.lee@nrel.gov

Andrew Glaws
National Renewable Energy Laboratory
15013 Denver West Parkway Golden, CO 80401-3305
andrew.glaws@nrel.gov

MS63

Multi-Index Stochastic Collocation for Pdes with Imperfect Solvers

This talk considers the construction of surrogate models (response surfaces) for parametric PDEs using multi-fidelity collocation methods, namely Multi-Index Stochastic Collocation (MISC). In some scenarios, in addition to standard discretization errors, the PDE approximations used to build a MISC response surface are affected by noise (e.g. due to iterative method tolerances, pre-asymptotic meshes, time-stepping). This noise is particularly problematic in low fidelity models; it might be parameter-dependent and hard to estimate and control a priori. Noise is interpolated by MISC and corrupts the approximated response surface (loss of monotonicity, spurious high-frequency oscillations), spoiling any subsequent UQ analysis. We propose an improved version of MISC that can detect such phenomena. Within our updated adaptive algorithm, at each iteration, we inspect the spectral content of the response surface and consequently stop exploring fidelities once the decay of their spectral coefficients stagnates due to such noise. Numerical experiments show the effectiveness of our approach in preventing the MISC approximation from becoming corrupted.

Benjamin M. Kent
The University of Manchester
kent@imati.cnr.it

Lorenzo Tamellini
CNR-IMATI
tamellini@imati.cnr.it

Matteo Giacomini, Antonio Huerta
Universitat Politècnica de Catalunya
matteo.giacomini@upc.edu, antonio.huerta@upc.edu

MS63

Multilevel Active Subspaces

The Active Subspace (AS) method is a powerful technique for identifying the most influential directions in high-dimensional input spaces that affect the output of a computational model. These can be then exploited in order to ease the construction of surrogate models, such as polynomails on a reduced number of variables. However, the standard AS algorithm requires a large number of gradient evaluations (samples) of the input output map to achieve quasi-optimal reconstruction of the Active Subspace, which can lead to a significant computational cost if the samples include numerical discretization errors which have to be kept sufficiently small. To address this issue, we propose a multilevel version of the Active Subspaces method (MLAS) that utilizes samples computed with different accuracies, which are often available in scientific computing models. The MLAS method yields different Active Subspaces for the model outputs across accuracy levels, which can match the accuracy of single-level Active Subspace with reduced computational cost. Then, one can build different surrogates at each level on a reduced number of variables, and combine these to obtain a cheaper global surrogate of the computational model. We demonstrate the practical viability of the MLAS method through numerical experiments based on random partial differential equations (PDEs) simulations.

Matteo Raviola
EPFL
matteo.raviola@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Raul F. Tempone
RWTH Aachen University & KAUST
tempone@uq.rwth-aachen.de

MS63

A Comparison of Multi-Fidelity Architectures for Neural Emulators

Outer loop tasks like optimization, uncertainty quantification, or inference become impractical when the underlying *high-fidelity* model is computationally expensive. Moreover, data-driven architectures often require large datasets to achieve sufficient predictive accuracy. A common strategy to address these issues is to develop emulators that can be evaluated at a lower cost. To this end, recent research has focused on creating multi-fidelity emulators that can utilize potentially biased and approximate *low-fidelity* information sources to reduce the total computationally cost. However, the performance of these emulators may be influenced by factors such as the dimensionality of the high-fidelity model, the presence of sharp gradients or discontinuities that low-fidelity models might not accurately capture. Additionally, high- and low-fidelity models may require different numbers of inputs or even involve multiple

low-fidelity models with varying accuracy across the input space. Our study aims to explore the conditions that lead to an optimal performance of multi-fidelity emulators, examining factors such as coordinate encoders, architecture complexity, spectral bias, and sampling levels across a wide range of test cases.

Cristian J. Villatoro
University of Notre Dame
cvillat2@nd.edu

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

MS64

Efficient Krylov Subspace Methods for Large-Scale Hierarchical Bayesian Inverse Problems

Uncertainty quantification for large-scale inverse problems remains a challenging task. For linear inverse problems with additive Gaussian noise and Gaussian priors, the posterior is Gaussian but sampling can be challenging, especially for problems with a very large number of unknown parameters (e.g., dynamic inverse problems) and for problems where computation of the square root and inverse of the prior covariance matrix are not feasible. Moreover, for hierarchical problems where several hyperparameters that define the prior and the noise model must be estimated from the data, the posterior distribution may no longer be Gaussian, even if the forward operator is linear. Performing large-scale uncertainty quantification for these hierarchical settings requires new computational techniques. In this work, we consider generalized Golub-Kahan based methods for large-scale, hierarchical Bayesian inverse problems. We consider a hierarchical Bayesian framework and exploit generalized Golub-Kahan based methods to efficiently sample from the posterior distribution. Numerical examples from dynamic photoacoustic tomography and atmospheric inverse modeling demonstrate the effectiveness of the described approaches.

Julianne Chung
Department of Mathematics
Emory University
julianne.mei-lynn.chung@emory.edu

MS64

Efficient Multilevel Methods for Material Properties Inversion in Heat Transfer Problems

In this work we extend multilevel methods that were originally developed for linear-quadratic distributed optimal control of elliptic equations to a different type of inverse problems regarding heat transfer in solids. More precisely, given a set of synthetic experiments that involve applying various heat sources and measuring boundary temperatures and heat fluxes, the goal is to identify the diffusivity tensor. Due to its ill-posedness, this is formulated as a regularized least-squares problem, which is nonlinear and non-convex. We use a hierarchy of geometric grids and multilevel methods both for the purpose of identifying a good initial guess at every level/grid, and for accelerating the non-linear iteration. We explore these strategies both in the context of interior point and semismooth Newton methods, in order to ensure the diffusivity is positive, for the isotropic case, and a symmetric and positive definite tensor, for the - more challenging - anisotropic case.

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

Liviu Marin
Department of Mathematics
University of Bucharest, Romania
liviu.marin@fmi.unibuc.ro

Alex Sheranko
Department of Mathematics and Statistics
University of Maryland, Baltimore County
asheran1@umbc.edu

MS64

Bayesian Inference Using a Constrained Conditional GAN

Generative models have been very useful in solving Bayesian inference problems. In particular, a novel conditional generative adversarial network (cGAN) framework [Ray et al., 2023] was developed to learn and sample from the conditional posterior distribution of the field of inference, given specific measurements. This framework has proven successful in various real-world applications, including medical imaging tasks and predicting wildfire spread. However, training such a cGAN typically requires thousands of labeled training samples, which are often generated in advance through experiments or forward model simulations. This requirement poses a challenge for many practical problems, where such extensive datasets may not be readily available. Moreover, the traditional cGAN approach does not explicitly incorporate known relationships between the inferred field and the measurements or any other physical constraints of the system. In this talk, we present an enhanced cGAN framework that integrates the systems physical relationships and constraints directly into the training process. This integration accomplishes two key objectives: (i) it guides the predictions towards the correct solution manifold associated with the underlying system, and (ii) it reduces the data complexity required for training the cGAN. We will explore several methods for imposing these physical constraints and demonstrate the effectiveness of this approach through a series of numerical experiments.

Deep Ray
University of Maryland, College Park
deepray@umd.edu

MS64

Extending Neural Network-Based Inference to Uncertainty Quantification

This talk considers deep neural networks to solve inverse problems, where we seek to estimate uncertain parameters of physical models, such as deterministic differential equations and stochastic processes. We train the neural networks to approximate inverse maps. The network is mapping from observational data to parameters. The point estimates from networks are accurate (provided sufficient training data), but they lack information about uncertainties in the predicted parameters. We discuss approaches to extend neural network-based inference to sta-

tistical settings to computationally solve inverse problem in a Bayesian framework.

Johann Rudi
Virginia Tech
jrudi@vt.edu

Andreas Mang
University of Houston
Department of Mathematics
amang@Central.UH.EDU

MS64

Efficient Sampling in Linear Inverse Problems with Hierarchical Matrices

Computing solutions to Bayesian inverse problems remains challenging, particularly when dealing with a large number of unknown parameters. Traditional methods often struggle due to the computational infeasibility of exact storage and operations with the prior covariance matrix, which are either impossible or computationally expensive to manage. This work investigates the use of hierarchical matrices to approximate the prior covariance matrix, a technique that significantly reduces memory requirements while maintaining accuracy. Hierarchical matrices not only allow for memory-efficient storage but also facilitate computationally efficient operations, such as matrix-vector product and Cholesky factorization, which are essential for handling large-scale problems. By integrating the hierarchical matrix-approximated prior covariance into the randomize-then-optimize (RTO) method, we can efficiently draw samples from the posterior distribution. This integration leverages the strengths of both hierarchical matrices and the RTO method, providing a robust framework for tackling high-dimensional inverse problems. Numerical experiments are conducted on deconvolution and a source estimation problem. These experiments demonstrate the effectiveness and efficiency of the proposed approach, showcasing significant improvements in computational performance.

Chao Zhang, Jakob S. Jorgensen
Technical University of Denmark
chaz@dtu.dk, jakj@dtu.dk

MS65

Compatible Discretizations for Hall Magnetohydrodynamics

The magnetohydrodynamic (MHD) equations are the coupling of Maxwells equations, which govern how electromagnetic waves behave, with the equations of fluid dynamics, and are used to model electrically conducting fluids. We present two novel finite difference schemes for the kinematic Hall MHD equations, which includes the nonlinear Hall term in Ohms Law. The first scheme uses averaging techniques and discrete derivative operators to calculate a second-order accurate numerical solution, while the second scheme makes use of operator-splitting to handle the nonlinear term. The contributed talk will give a brief overview of finite-difference methods for Maxwells equations, the MHD equations and present some energy estimates for the Hall model. For comparison, simulation results will be presented for both the resistive MHD model and the Hall MHD model.

Fernando Angulo Barba, Vrushali Bokil
Oregon State University
angulobf@oregonstate.edu,
vrushali.bokil@oregonstate.edu

MS65

Using Voigt-Regularized MHD to Generate Equilibria with Islands

The computation of 3D MHS equilibria without the assumption of nested flux surfaces is a longstanding challenge for magnetic confinement fusion. Recent work (Constantin & Pasqualotto, 2022) on Voigt regularization of MHD has proven the existence of time-asymptotic, regular solutions of Voigt-MHD. The addition of non-ideal Voigt terms to the MHD equations allows the magnetic topology to break, forming islands. We demonstrate that Voigt-MHD can be used to efficiently compute 2D equilibra with islands, and investigate how the choice of Voigt parameters affects computation time and eventual equilibrium. We discuss strategies for generalization to 3D equilibrium calculations. This work is supported by the Simons Foundation and the DoE SciDAC project HifiStell.

Andrew Brown
Princeton University
aobrown@princeton.edu

Yi-Min Huang
University of Maryland at College Park
NASA Goddard Space Flight Center
yiminh@princeton.edu

Wrick Sengupta, Amitava Bhattacharjee
Princeton University
wsengupta@princeton.edu, amitava@princeton.edu

MS65

A Comprehensive Exploration of Quasisymmetric Stellarators and Their Coil Sets

We augment the 'QUAsi-symmetric Stellarator Repository' (QUASR) to include vacuum field stellarators with quasi-helical symmetry using a globalized optimization workflow. The database now has around 370,000 quasisaxisymmetry and quasihelically symmetric devices along with coil sets, optimized for a variety of aspect ratios, rotational transforms, and discrete rotational symmetries. We also outline a couple of ways to explore and characterize the data set. We plot devices on a near-axis quasisymmetry landscape, revealing close correspondence to this predicted landscape. We also use principal component analysis to reduce the dimensionality of the data so that it can easily be visualized in two or three dimensions. Principal component analysis also gives a mechanism to compare the new devices here to previously published ones in the literature. We are able to characterize the structure of the data, observe clusters, and visualize the progression of devices in these clusters. These techniques reveal that the data has structure, and that typically one, two or three principal components are sufficient to characterize it. The data set is archived at https://zenodo.org/doi/10.5281/zenodo.10050655 and can be explored online at quasr.flatironinstitute.org.

Andrew Giuliani
Flatiron Institute
agiuliani@flatironinstitute.org

Eduardo Rodríguez
Max Planck Institute of Plasma Physics at Greifswald
eduardo.rodriguez@ipp.mpg.de

Marina Spivak
Flatiron Institute
mspivak@flatironinstitute.org

MS65

Recent Progress in the DESC Stellarator Optimization Code

We present recent work in the DESC stellarator optimization code suite [Dudt 2020, Dudt 2023, Panici 2023, Conlin 2023]. The code is written in Python and uses the JAX package for its automatic-differentiation capabilities, as well as for just-in-time compilation and GPU-capabilities for improving code performance. The DESC code is written with modern software practices that enable flexibility in its approach to the equilibrium and optimization problems. One example is near-axis constrained equilibria in DESC, which connects near-axis expansion theory to global 3D ideal magnetohydrodynamics (MHD) in a natural way, by constraining the DESC equilibrium behavior order-by-order in radius to match near-axis expansion (NAE) theory [Garren 1991]. This allows for global solutions which retain the desirable NAE properties but avoid relying on the NAE far from axis, where the NAE is less reliable. Other recent works include the implementation of a high-order singular integral algorithm [Malhotra 2019, Malhotra 2020] to efficiently evaluate the plasma field at the boundary needed for free-boundary equilibrium solving, and the implementation of a general omnigenous field model [Dudt 2024], allowing for optimization targeting omnigenity types beyond the usual quasisymmetric or quasi-isodynamic varieties. This work is supported by US DoE DE-AC02-09CH11466, DESC0022005, and Field Work Proposal No. 1019

Dario Panici, Rory Conlin, Daniel Dudt
Princeton University
dpanici@princeton.edu, wconlin@princeton.edu, ddudt@princeton.edu

Eduardo Rodriguez
Max Planck Institute for Plasma Physics
edurodriurreta@gmail.com

Rahul Gaur, Kian Orr, Yigit Gunsur Elmacioglu, Egemen Kolemen
Princeton University
rg6256@princeton.edu, kianorr@princeton.edu, ye2698@princeton.edu, ekolemen@princeton.edu

MS65

An Introduction to Solving the Magnetohydrostatic Equations for Stellarators

Stellarators are magnetic plasma confinement devices with potential use for fusion energy. Unlike the axisymmetric tokamaks, stellarators are genuinely three-dimensional, meaning they satisfy a set of nonlinear 3D PDEs often referred to as the magnetohydrostatic equations (MHS). Because numerical MHS solvers are often used within optimization loops, there are many choices that scientists make to balance speed, reliability, and the ability to describe the relevant physics. In the first half of this talk, I will introduce the MHS equations with a focus on the tension between pressure, magnetic geometry, and solver speed. Then, in the second half, I will describe a numerical approach to solving MHS based on an asymptotic expansion in small distances about the stellarator's axis. This so-called near-axis expansion can be used to quickly explore the space of stellarators, but it is outperformed by full 3D solvers far from the axis. I will describe recent efforts in improving its convergence and expanding the physics it can describe.

Maximilian Ruth, David Bindel
Cornell University
maximilian.ruth@austin.utexas.edu, bindel@cornell.edu

Rogerio Jorge
University of Wisconsin-Madison
rogerio.jorge@wisc.edu

MS66

Summation-by-Parts Approach for Kinetic Equations

Summation-by-parts (SBP) operators are discrete approximations of differ-ential operators derived from various numerical methods such as finite differ-ences, finite volumes and discontinuous Galerkin methods. They are extensively used in the approximation of various partial differential equations (PDEs) due to their remarkable stability properties. Their structure makes it easy to mimic continuous stability estimates for semi-discrete and fully-discrete approxima-tions of a range of problems. The focus of this talk is on space-time SBP discretisations of kinetic equations, which are used in the study of rarefied gasses, plasma physics and neutron transport. In particular, we consider a linear kinetic transport equation under a diffusive scaling with periodic boundary conditions. The goal is to propose a stable discretisation of the problem that is also asymptotic preserving, such that the numerical scheme converges to a consistent approximation of the limiting macroscopic PDE.

Anita Gjesteland
University of Waterloo
Department of Applied Mathematics
agjesteland@uwaterloo.ca

MS66

Curvilinear Hexahedral Cell Complexes for Swept Trivariate Splines via Foliations

Isogeometric analysis (IGA) is a computational technique that uses smooth spline discretizations for finite element and other analyses. Although IGA offers significant advantages in per degree of freedom accuracy, numerical conditioning, and representation of the domain, generating smooth spline discretizations for real-world geometries remains a challenge. Most notably, robust techniques for automatically creating the hexahedral cell complexes needed for trivariate locally tensor product splines are still lacking. In this work, we introduce a method to construct trivariate cell complexes over topologically swept manifolds. Our technique builds these trivariate cell complexes based on bivariate cell complexes defined over the source of the sweep. We apply bijective parameterization techniques from the bivariate setting on a lower-dimensional foliation of the manifold to extend the parameterization from the sweep source to the entire domain. The resulting trivariate parameterization inherits certain bijectivity guarantees from the bivariate setting. We demonstrate the effectiveness of our method by generating unstructured splines over hexahedral layouts for various geometries.

Caleb B. Goates
Brigham Young University

Department of Civil and Construction Engineering
calebgoates@gmail.com

Kendrick Shepherd
Brigham Young University
kendrick_shepherd@byu.edu

MS66

Discontinuous Galerkin Methods for Kinetic Equations with Flexible Coordinates: Conservation Relations, L^2 stability, and More

Kinetic equations describe, from near first-principles, the physics of gases, plasmas, radiation, neutrino, and neutron transport. The robust, accurate, and cost-effective discretization of these fundamental equations poses unique challenges though, owing to these various kinetic equations high dimensionality and multi-scale nature. Discontinuous Galerkin (DG) methods have become an increasingly popular approach to the numerical solution of different flavors of kinetic equations, such as the Boltzmann equation for neutral gases and the Vlasov equation for plasmas, because with the proper choice of solution space, conservative, L^2 stable DG discretizations can be derived, even for nonlinear kinetic equations. However, because kinetic equations are high-dimensional, up to six dimensions plus time, the extension of these methods to kinetic equations in general geometry so that optimized coordinates can be utilized for a given problem is of vital importance. In this presentation, we will discuss how the current cutting edge of DG methods for kinetic equations must be modified for certain favorable coordinate transformations. In particular, we focus on velocity-space coordinate transformations such as a transformation to the local hydrodynamic flow frame and non-uniform velocity-space mappings, which permit optimized deformations of the velocity space coordinates to minimize the necessary resolution required for the problems of interest. We will review why DG methods produce conservative, L^2 stable discretizations of the kinetic equation, and how these proofs are modified in the presence of more general coordinates. Finally we demonstrate on a few classical problems in plasma physics the utility of this approach and empirically verify some of the properties of our new DG schemes for kinetic equations.

Jimmy Juno
Princeton Plasma Physics Laboratory
jjuno@pppl.gov

Jason TenBarge
Princeton Plasma Physics
tenbarge@princeton.edu

Ammar Hakim
Princeton Plasma Physics Laboratory
ahakim@pppl.gov

MS66

Summation-by-Parts Finite-Difference Operators for Singular Coordinate Transforms

We present a general scheme for using existing summation-by-parts (SBP) finite-difference (FD) operators with singular coordinate transformations. The scheme preserves the accuracy and SBP properties of the original operators and permits simple implementation into existing codes. The scheme allows taking advantage of the many previously constructed SBP-FD operators and developments when solving problems involving coordinate singularities. This greatly simplifies the design of the numerical method by avoiding re-constructing operators for the given coordinate system. The operators are modified by viewing them in a weak form and eliminating the degrees of freedom associated with the coordinate singularities. By then returning to a strong form formulation an operator for the reduced grid is achieved which can be handled as any other SBP-FD operator. Using the scheme we derive a stable and high-order accurate finite-difference method for underwater acoustic wave propagation in an axisymmetric domain. The method handles range and depth-dependent material properties, including discontinuous jumps. The accuracy and stability properties of the method are proven and corroborated using numerical experiments.

Jonatan Verpers
FOI
jonatan.verpers@foi.se

MS66

High-Order Symmetric Positive Weight Quadrature Rules for Construction of Summation-by-Parts Operators on Simplices

The summation-by-parts (SBP) property can be used to develop provably-stable discretizations of partial differential equations. The existence of a degree p SBP operator on a given domain depends on the existence of a quadrature rule of degree $2p-1$ or larger. In this talk, we present a technique to derive very high order quadrature rules on simplices. Novel fully-symmetric positive-interior quadrature rules of degrees up to 84 on triangles and 40 on tetrahedra are derived. Most of the rules have e?ciencies of more than 95% and 80% on triangles and tetrahedra, respectively, where effciency is the ratio of the minimum number of quadrature nodes predicted theoretically to the number of nodes we obtain in this work. We demonstrate the accuracy of the quadrature rules using numerical examples.

Zelalem Arega Worku
University of Waterloo
zelalem.worku@uwaterloo.ca

MS67

Model Enrichments for Uncertain, Reduced Reactions in Ablation Models for Hypersonic Flight

During hypersonic flight, air reacts with a planetary entry vehicles thermal protection system (TPS), creating reaction products that deplete the TPS. Reliable assessment of TPS performance depends on accurate ablation models. New finite-rate gas-surface chemistry models are advancing state-of-the-art in TPS ablation modeling, but model reductions that neglect chemical species may be necessary in some cases for computational tractability. This work develops a theory-informed stochastic enrichment operator to improve the predictive capability and quantify uncertainties in such reduced models while maintaining computational tractability. Specifically, modeling a subset of chemical species causes discrepancies in predicted carbon monoxide production. We propose an enrichment operator, embedded in the reduced model, to quantify the effect of neglected species. The enrichment operator is theory-informed and calibrated with Bayesian inference. Numerical results show the enrichment operator improves the reduced models predictions without significantly increasing computational cost. SNL is managed and operated by NT-

ESS under DOE NNSA contract DE-NA0003525.

Rileigh Bandy, Rileigh Bandy, Teresa Portone
Sandia National Laboratories
rjbandy@sandia.gov, rjbandy@sandia.gov, tporton@sandia.gov

MS67

Embedded Model Error Quantification for Climate Modeling

There exist several models for simulating climate change, but each relies upon simplifying physical assumptions and empirical calibrations. As such, a central goal for practitioners is to not only quantify the uncertainty in the predictions obtained from these models, but also explicitly delineate the contributions from modeling error, noise, and limited observational data. This study explores a non-intrusive, data-driven method to quantify uncertainty induced by model errors in the E3SM models land component, ELM. Specifically, a surrogate modeling approach using domain decomposition and the Karhunen-Love Decomposition (KLD) is leveraged to build an accurate, low-cost approximation for ELM. Then, an additive correction in the form polynomial chaos expansions is employed to embed model error directly within the surrogate. A Bayesian framework is then used with low-cost samples from the surrogate and observational data from select FLUXNET sites across the US to calibrate both the model error parametrization and ELM input parameters. The calibrated predictions are consequently endowed with uncertainty that can be decomposed into contributions due to model error and the quality of training data. This study not only provides insight into the parameters that contribute most to prediction variability in ELM, but also showcases an approach for efficiently quantifying model error, while remaining consistent with the physical constraints inherent within complex models.

Nikhil Iyengar
Georgia Institute of Technology
niyenga@sandia.gov

Daniel Ricciuto
Oak Ridge National Laboratory
ricciutodm@ornl.gov

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

MS67

Multi-Fidelity Uncertainty Quantification of Embedded Model Error in Turbulence Modeling

Reynolds-averaged Navier-Stokes (RANS) models are very popular in the computational fluid dynamics community for describing turbulent flows. However, the predictive capability of these models is hindered by errors stemming from their linear eddy viscosity assumptions, particularly in separated flow regimes such as slow, high-lift flight at high angles of attack. Recently, embedded, internal model corrections have been explored to quantify structural errors together with model parameters, within a Bayesian inference context. Model parameters are modified through an additive error, enabling a non-intrusive approach to quantify model uncertainty. This framework necessitates advanced inverse UQ methodologies that leverage multiple fidelities to reduce computational burden. In this context, we will use a simulation-based likelihood approach to model the probabilistic likelihood and implement a novel coupling strategy to synchronize Markov chains within a multi-level Markov Chain Monte Carlo (ML-MCMC) estimator. To validate our methodology, numerical experiments will be carried out for test cases that undergo high separation for certain boundary conditions.

Sanjan Muchandimath, Alex Gorodetsky
University of Michigan
sanjancm@umich.edu, goroda@umich.edu

MS68

Coupling Volume Integral Equations and Boundary Element Methods to Simulate Ultrasound Propagation

Focused ultrasound is one of the most effective modalities for neuromodulation and non-invasive cancer treatments. Its application to a clinical environment depends on the ultrasound instrument's capacity to guide sufficient energy toward the targeted region and ablate the tumor while sparing healthy tissue and organs in the beam path. Computational methods aid in safety guideline assessment and patient-specific treatment planning. We have already achieved realistic simulations with the Boundary Element Method for the Helmholtz equation implemented in our open-source Python library, OptimUS. However, our previous approach is limited to harmonic wave propagation through piecewise homogeneous materials. In practice, localized heterogeneities in materials such as bone may aberrate the focus. For this purpose, we designed a Volume Integral Equation for heterogeneous materials that naturally couples with the Boundary Element Method for unbounded domains. This approach allows us to take the speed of sound and density values from biomedical images. We extensively verified our proposed technique by benchmarking against analytical solutions, the Finite Element Method, and FEM-BEM coupling. Our Volume-Surface Integral Equations show a high accuracy on a voxel grid and fast convergence at various frequencies. Our computational techniques apply to diverse simulations in computational acoustics beyond our specific objective of modeling focused ultrasound propagation in the human body.

Danilo Aballay
Institute for Mathematical and Computational Engineering
Pontificia Universidad Católica de Chile
danilo.aballay@uc.cl

Elwin Van 't Wout
Pontificia Universidad Católica de Chile
e.wout@uc.cl

MS68

FFT-Accelerated Inverse Media Scattering in Three Dimensions Using Continuation in Frequency

In this work we describe a fast, robust, and accurate algorithm to solve an inverse problem for reconstructing the sound speed profile of a three-dimensional variable media using multifrequency scattered data. The inverse problem is first recast as a collection of PDE-constrained optimization problems, one for each frequency of interest. In order to ensure stabilit and achieve a computational speedup over the full multifrequency optimization problem, the solver takes advantage of a continuation in frequency

algorithm which solves a series of single-frequency inverse scattering problems in order of increasing frequency. Each single-frequency inverse problem is ill-posed and nonlinear. The ill-posedness is addressed by explicitly constraining the search space to a band-limited representation of the sound speed profile, while the non-linearity is treated using an iterative method. We present numerical results to demonstrate the effectiveness of our solver in reconstructing both obstacles and smooth sound profiles.

Carlos Borges
University of Central Florida
carlos.borges@ucf.edu

Michael O'Neil
Courant Institute
New York University
oneil@cims.nyu.edu

MS68

Boundary Integral Approaches for Spectral Shape Optimization

We suggest a robust and spectrally converging numerical method for the spectrum of the Steklov-Helmholtz eigenvalue problem for smooth domains in the plane. Our method involves layer potentials and the fundamental solution of the Helmholtz equation, which allows us to reconstruct the eigenfunctions as well. The method can be used to investigate many properties of the eigenpairs. In particular, we use our method to study certain shape optimization problems for the Steklov-Helmholtz eigenvalues.

Kshitij A. Patil
Simon Fraser University
kap15@sfu.ca

Nilima Nigam
Dept. of Mathematics
Simon Fraser University
nigam@math.sfu.ca

MS68

Integral Formulation of Dirac Singular Waveguides

This talk concerns a boundary integral formulation for the two-dimensional massive Dirac equation. The mass term is assumed to jump across a one-dimensional interface, which models a transition between two insulating materials. This jump induces surface waves that propagate outward along the interface but decay exponentially in the transverse direction. After providing a derivation of our integral equation, we establish that it has a unique solution for almost all choices of parameters using holomorphic perturbation theory. We then implement a fast numerical method for solving our boundary integral equation and present several numerical examples of solutions and scattering effects.

Guillaume Bal
University of Chicago
Departments of Statistics and Mathematics
guillaumebal@uchicago.edu

Jeremy Hoskins
University of Chicago
jeremyhoskins@uchicago.edu

Solomon Quinn
University of Minnesota
srquinn@umn.edu

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS68

A 2D Simulation Framework for Nonlinear Electromagnetic Wave Interactions in Cellular Structures

In this presentation, we explore the application of integral equation methods to the precise analysis and design of photonic devices, particularly photonic chips interfaced with optical waveguides of various geometries. Integral equation techniques offer significant advantages for modeling complex electromagnetic interactions in photonic systems, providing accurate solutions across a wide range of device configurations. We will discuss the strengths and limitations of these methods, highlighting their effectiveness in different operating regimes and their role in enhancing the accuracy and efficiency of photonic device simulations. Practical examples will be provided to demonstrate the versatility of the approach in addressing real-world engineering challenges in photonics.

Felipe Vico
Universitat Politècnica de València
fevibon@teleco.upv.es

MS70

Expnode: Exponential Integrators in Neural Ordinary Differential Equations

With the abundance of data generated by high-resolution simulations it has become popular to build models derived from data. Of these data-driven models there has been many works on neural ordinary differential equations. While these models have enjoyed success, those based on explicit integrators suffer from limited stability, impeding their efficiency and robustness when encountering stiff problems. In this work we propose using an exponential integrator, which is an explicit integrator with stability properties comparable to implicit methods. We demonstrate that our model has advantages in both learning and deployment over standard explicit neural ordinary differential equation methods, as the favorable stability properties allows us learn from lower-resolution data. Our neural network implementation takes advantage of Higham's algorithm of a matrix-free application of the matrix exponential on a vector. We present a parameterization based on the Hurwitz decomposition that controls the spectrum of the linear operator to be to the left of the complex plane. Examples such as Kuramoto-Sivashinky, Grad-13 moments and quantum computing are presented to demonstrate effectiveness.

Allen Alvarez Loya
Department of Applied Mathematics
University of Colorado, Boulder, CO
aalvarezloya@gmail.com

MS70

Data-Driven Whitney Forms Model Reduction for Maxwell

Maxwell's equations is a coupled partial differential equa-

tion that is used to model the evolution of electromagnetic waves. Typically, these models are discretized using compatible finite element methods with high their computational costs. This talk proposes a reduced model for Maxwell's equations based on a data-driven Whitney form rediscretization approach. This methodology fits a partition of unity based reduced discretization of Maxwell's to simulation data. This allows a reduction of computational costs relative to the full scale simulation with limited error. We will compare the reduction to a classical principle orthogonal decomposition model reduction approach, while highlighting potential advantages of thew new approach to a broader range of nonlinear problems.

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

Jonas Actor
Sandia National Laboratories
jaactor@sandia.gov

MS70

Leveraging Interpolation Models and Error Bounds for Verifiable Scientific Machine Learning

Effective verification and validation techniques for modern scientific machine learning workflows are challenging to devise. Statistical methods are abundant and easily deployed, but often rely on speculative assumptions about the data and methods involved. Error bounds for classical interpolation techniques can provide mathematically rigorous estimates of accuracy, but often are difficult or impractical to determine computationally. We present a best-of-both-worlds approach to verifiable scientific machine learning by demonstrating that (1) multiple standard interpolation techniques have informative error bounds that can be computed or estimated efficiently; (2) comparative performance among distinct interpolants can aid in validation goals; (3) deploying interpolation methods on latent spaces generated by deep learning techniques enables some interpretability for black-box models. We present a detailed case study of our approach for predicting lift-drag ratios from airfoil images. Code developed for this work is available in a public Github repository.

Tyler H. Chang
Argonne National Laboratory
tchang@anl.gov

Andrew Gillette
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
gillette7@llnl.gov

Romit Maulik
Pennsylvania State University
rmm7011@psu.edu

MS70

Construction of Coarse-Grained Molecular Dynamics with Many-Body Non-Markovian Memory

We introduce a machine-learning-based coarse-grained molecular dynamics (CGMD) model that faithfully retains the many-body nature of the inter-molecular dissipative interactions. Unlike the common empirical CG models, the present model is constructed based on the Mori-Zwanzig formalism and naturally inherits the heterogeneous state-dependent memory term rather than matching the mean-field metrics such as the velocity auto-correlation function. Numerical results show that preserving the many-body nature of the memory term is crucial for predicting the collective transport and diffusion processes, where empirical forms generally show limitations.

Liyao Lyu, Huan Lei
Michigan State University
lyuliyao@msu.edu, leihuan@msu.edu

MS70

Reducing Structural Errors of Ensemble Score Filter for Nonlinear Data Assimilation with Partial Observation

The Ensemble Score Filter (EnSF) has emerged as a promising approach to leverage score-based diffusion models for high-dimensional recursive Bayesian inference in nonlinear data assimilation. While initial applications to Lorenz-96 and quasi-geostrophic systems showed potential, the current EnSF methodology faces two key limitations in score function estimation. First, it employs a heuristic weighted sum to combine the likelihood and prior distribution scores, introducing structural errors into the estimation of the posterior distribution. Second, the method tends to underestimate the correlation among different state components, diminishing its effectiveness for data assimilation with partial observations, i.e., only a subset of the state components are observable. This workr addresses both challenges through two innovations: (1) deriving the exact posterior score function and its approximations by assuming the prior distribution is a Gaussian mixture model defined by posterior ensemble from previous Bayesian iteration, and (2) incorporating the composition of observation operator and state equation into the likelihood function to generate non-zero likelihood score components for unobservable states. Numerical experiments demonstrate that these enhancements substantially improve EnSFs accuracy in both posterior distribution estimation and tracking of high-dimensional dynamical systems under partial observations.

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

Siming Liang
Florida State University
sliang@fsu.edu

Zezhong Zhang
Oak Ridge National Laboratory
zhangz2@ornl.gov

Hristo G. Chipilski
Florida State University, Department of Scientific Computing
hchipilski@fsu.edu

MS71

On Learning Nonlinearly State Dependent Port-

Hamiltonian Systems

Port-Hamiltonian systems have received a lot of attention in recent years because of their interesting properties in modeling and control. They are especially useful for system interconnection because they retain their port-Hamiltonian structure. These structured systems, which are also close to physics, allow for a more generic understanding of the underlying dynamics of physical systems. Data-driven modeling of port-Hamiltonian systems remains an open problem, particularly for nonlinear systems. In this talk, we look at the recent emergence of scientific machine learning techniques for learning port-Hamiltonian systems. More specifically, we compare between different architectures for learning port Hamiltonian systems from data and discuss different strategies to enable a more efficient training. Finally, we illustrate our results with different examples of physical systems.

Karim Cherifi
Wuppertal University
cherifi@uni-wuppertal.de

Achraf El Messouadi
Sorbonne university, France
achraf.el_messaoudi@etu.sorbonne-universite.fr

Hannes Gernandt
Fraunhofer IEG, Cottbus, Germany
gernandt@uni-wuppertal.de

MS71

Gaussian-Process-Based Parametric Latent Dynamics Modeling of Shock-Induced Pore Collapse Process

Shock-induced pore collapse is a key mechanism for various applications in science and engineering, such as targeted drug delivery, shock lithotripsy and heterogeneous energetic (HE) materials. Much efforts have been dedicated to establish high-fidelity numerical models for prediction, though they are computationally expensive. This necessitates surrogate reduced-order models for multi-query applications, such as design optimization and uncertainty quantification. In this work, we study the use of a Gaussian-process-based parametric latent dynamics learning method. Learning the pore collapse dynamics is challenging in that the training data is scarce and furthermore the latent dynamics changes depending on the pore shape and location. In order to tackle this challenge, the uncertain latent dynamics at an unknown parameter point is calibrated with Gaussian process based on latent dynamics learned at sample points. Another important quantity of interest is the evolution of the discontinuous pore interface moving along the space. In order to accurately track the pore, the latent-space autoencoder is augmented with pore indicator decoding. Compared to standard multi-layer perceptron, the augmented autoencoder showed its superiority in accurately capturing both the pore interface and HE temperature. On reproductive cases, the trained latent dynamics is able to achieve more than a thousand times speed-up, while maintaining less than 1% error.

Seung Whan Chung, Youngsoo Choi
Lawrence Livermore National Laboratory
chung28@llnl.gov, choi15@llnl.gov

Paul Tranquilli
Lawrence Livermore National Lab
tranquilli1@llnl.gov

Christopher Miller, H. Keo Springer, Kyle Sullivan
Lawrence Livermore National Laboratory
miller279@llnl.gov, springer12@llnl.gov, sullivan34@llnl.gov

MS71

Entropy Consistent Generative Models for Physics Simulations Using Stochastic Interpolants

Despite the massive potentials of generative models, there is quite limited work being done in the area of applying such models to physics simulations. This is due to challenges in achieving efficient inference times while maintaining physical consistency. In this work, we address these challenges by leveraging the stochastic interpolant framework as a generative model with entropy consistency incorporated. In this framework one aims to learn a stochastic differential equation (SDE) that maps samples from one distribution to another by defining a stochastic interpolation between the two distributions, which is then used to train a drift term in the SDE. In contrast to the widely used denoising diffusion probabilistic models, which are limited to Gaussian priors, the stochastic interpolant framework can map samples between arbitrary distributions. This makes the method more computationally efficient as appropriate priors can be chosen. For physics simulations, the target distribution is the conditional distribution of a system's state given an initial state. By applying this approach autoregressively, we can generate complete trajectories, accounting for the inherent uncertainty by producing multiple plausible outcomes from the same initial condition. The incorporation of entropy consistency ensures that all generated trajectories adhere to the laws of physics. We demonstrate the effectiveness of our method with examples from fluid dynamics.

Nikolaj T. Mücke
Centrum Wiskunde & Informatica (CWI)
Amsterdam, NL
nikolaj.mucke@cwi.nl

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cwi.nl

MS71

Rigorous a-Posteriori Error Estimation for Scientific Machine Learning: a Control-Theoretical Perspective

Physics-informed neural networks (PINNs) leverage knowledge by incorporating the underlying dynamical system into the loss function during the learning to enhance the data-driven machine learning approach. Access to the differential equation allows for a rigorous a posteriori error estimation framework. When dealing with a partial differential equation, the error estimator has to reflect three error sources: (i) approximation of the differential equation, (ii) approximation of the initial value, and (iii) approximation of the boundary conditions. While the first two ingredients are standard in residual-based error estimation, particular care has to be given to the latter. In this work, we present a semigroup approach combined with input-to-state stability to derive the error estimator. The required grown constants and growth functions are then determined by a classical approximation step for which we prove convergence. If the boundary error is sufficiently smooth - a reasonable assumption in the context of PINNs - then the

computation can be further simplified via a homogenization argument. We illustrate our theoretical findings on a numerical toy problem and a challenging poroelasticity problem on the geometry of a human brain.

Birgit Hillebrecht
University of Stuttgart
Germany
birgit.hillebrecht@simtech.uni-stuttgart.de

Benjamin Unger
University of Stuttgart, Germany
benjamin.unger@simtech.uni-stuttgart.de

MS71

Applications of Order Value Optimization in Machine Learning

In order value optimization (OVO), given a collection of functions f_i's, $i = 1, ..., n$, one is concerned with minimization of the sum of m point-wise smallest functions. When $m = 1$ and $m = n$, the problem reduces to the commonly studied problems of minimizing the point-wise minimum of functions and the point-wise maximum of functions, respectively. So far, OVO has been applied in several domains, including that of optimization formulations in finance. However, to date, its usage in the domain of machine learning has not been sufficiently explored. OVO exhibits several potential advantages in machine learning, including: 1) cost-effectiveness - the number of calculations is reduced due to the usage of only a subset of the collection of the functions; and 2) amenability to noisy and outlier data scenarios, since the influence of outliers in the data is diminished by considering only m smallest function values. In this talk, we provide an overview of convergence results in OVO. Specifically, the presented OVO method converges to the critical points of objective function. In addition, we point to its relationship with existing machine learning methods such as tilted empirical risk minimization and superquantile minimization. Moreover, we provide several numerical studies that show the potential of OVO in machine learning problems with outlier and noisy data.

Manojlo Vukovic
Faculty of Technical Sciences, University of Novi Sad
manojlo.vukovic@uns.ac.rs

Dusan Jakovetic
University of Novi Sad
dusan.jakovetic@dmi.uns.ac.rs

Natasa Krejic
Faculty of Sciences, University of Novi Sad
natasa@dmi.uns.ac.rs

Natasa Krklec Jerinkic
University of Novi Sad
Novi Sad, Serbia
natasa.krklec@dmi.uns.ac.rs

Luka Rutesic
Faculty of Sciences, University of Novi Sad
luka.rutesic@dmi.uns.ac.rs

MS72

Efficient Component Libraries for the Design of Boundary Integral Equation Solvers

Imagine wanting to write your own bespoke integral equation solver from scratch. It would require management of surface grids, quadrature rules, basis functions, fast direct evaluation of Green's functions (ideally with SIMD on ARM and Intel), sparse and dense linear algebra, all the way up to implementation of Fast Multipole Methods or other fast solver techniques. These libraries should be able to talk to each other and be easy to install and deploy. In this talk we discuss our own experience in splitting up our own integral equation solver into small component libraries that can be independently used, and have well defined C interfaces that allow integration into high-level languages. The goal is to allow users to pick and choose what components they need, and to create a low-level platform that makes it easy to implement novel integral equation formulations and technologies without needing to reinvent the wheel each time.

Timo Betcke, Srinath Kailasa, Ignacia Fierro Piccardo
University College London
t.betcke@ucl.ac.uk, srinathkailasa@gmail.com,
maria.piccardo.19@ucl.ac.uk

Matthew Scroggs
University of Cambridge
matthew.scroggs.14@ucl.ac.uk

MS72

Faster Evaluation of Line Expansions of Layer Potentials Using Recurrence Relations

We present a method to symbolically generate recurrence relations for derivatives of Green's functions under an assumption of radial symmetry. These recurrence relations can be used to evaluate Taylor expansions of the kernel which are relevant in the context of QBX, a method for the evaluation of the singular integrals occurring in layer potentials. Unlike naive symbolic computation, our process results in linear complexity with respect to the expansion order. We provide data on the flop count and efficiency results of our recurrence generator in the context of evaluating Taylor expansions. We then give experimental results when these recurrences are integrated into the already existing QBX method for layer potential evaluation.

Hirish Chandrasekaran
University of Illinois
hirishc2@illinois.edu

MS72

Code-Verification Techniques for Electromagnetic Surface Integral Equations

Code verification is an important step towards establishing the credibility of the results of computational physics codes by assessing whether the underlying numerical methods have been correctly implemented. The discretization of differential, integral, or integro-differential equations necessarily incurs some truncation error, and thus the approximate solutions produced from the discretized equations will incur an associated discretization error. If the solution to the problem is known, a measure of the discretization error may be evaluated directly from the approximate solution. The code may be verified by examining the rate at which the error decreases as the discretization is refined, thereby verifying the observed order of accuracy of the discretization scheme is the expected order of accuracy. Elec-

tromagnetic surface integral equations yield many code-verification challenges due to the various sources of numerical error and their possible interactions. These error sources arise from the use of basis functions to approximate the solution, quadrature to numerically integrate, and planar elements to represent curved surfaces. In this work, we provide approaches to separately measure the numerical errors arising from these different error sources using manufactured solutions. We demonstrate the effectiveness of these approaches for cases with and without coding errors.

Brian A. Freno
Sandia National Laboratories
bafreno@sandia.gov

MS72

Chunkie: a Matlab Integral Equation Toolbox

chunkIE is MATLAB package for prototyping integral equation methods in two dimensions. It includes support for the solution of Laplace, Helmholtz, Stokes, Yukawa, and elasticity boundary value problems. The package also includes support for handling corners, and multiple junction interfaces, coupling to fast multipole methods whenever available, and to fast direct solvers using the Fast linear algebra package in MATLAB (FLAM). In this talk, we will discuss the software architecture of chunkIE and demonstrate its capabilities through several examples.

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS73

Diffusion-Based Image Reconstruction for Position-Blind Ptychography

Ptychography is a type of coherent diffraction imaging which uses a strongly coherent X-ray source from a synchrotron to reconstruct high resolution images. By shifting the illumination source, it exploits redundancy of multiple diffraction patterns to robustly solve the related phase retrieval problem. The shift parameters are often subject to uncertainty and introduce additional ill-posedness in the reconstruction task. Motivated by applications in single particle imaging, we consider the extreme task of entirely unknown shifts, which have to be recovered jointly with the image. The resulting blind inverse problem requires careful regularisation. We explore the ability of learned priors, encoded by a score-based diffusion model, to jointly solve the reconstruction problem and recover the scan positions.

Simon Welker
Universität Hamburg
simon.welker@uni-hamburg.de

Lorenz Kuger
FAU Erlangen-Nürnberg & DESY Hamburg
lorenz.kuger@desy.de

Tim Roith
Deutsches Elektronen-Synchrotron DESY
Hamburg, Germany
tim.roith@desy.de

Berthy Feng
California Institute of Technology
bfeng@caltech.edu

Martin Burger
DESY Hamburg
Universität Hamburg
martin.burger@desy.de

Timo Gerkmann
Universität Hamburg
timo.gerkmann@uni-hamburg.de

MS74

Multi-Scale Latent Space Dynamics Identification

We propose multi-scale parametric Latent Space Dynamics Identification (mLaSDI), a data-driven reduced-order modeling framework for problems which exhibit multi-scale features. LaSDI has successfully been applied to build reduced-order models of many different partial differential equations which have some parametric dependence. Since its recent inception, there have been many variants of LaSDI to make some improvements. These methods typically follow the same general algorithm. First, we compress the data to a low-dimensional latent space using an autoencoder. Then, we use the sparse identification of nonlinear dynamics (SINDy) algorithm to find linear ordinary differential equations (ODEs) with approximate the latent space for each simulation. Given a new parameter, we then interpolate and evolve the ODEs to make rapid predictions. However, the autoencoder may have reconstructing data with multiple scales. Additionally, using linear ODEs in the latent space restricts the behaviors we are able capture while higher order ODEs are less stable and harder to interpolate. To address these issues, we introduce mLaSDI. Simply by using successive linear approximations of the latent space, we can capture significantly more behaviors in the latent space compared to LaSDI. This allows us to extend the LaSDI framework and build accurate reduced-order models for more problems where LaSDI performs poorly.

William Anderson, Youngsoo Choi
Lawrence Livermore National Laboratory
anderson316@llnl.gov, choi15@llnl.gov

MS74

Physically-Consistent Flow Matching – a Paradigm for Data Generation in Physical Systems and Density Estimation

Generative models have shown significant promise in applications such as computer vision, natural language processing, and climate modeling. This work presents a flow-based generative modeling approach for data generation and probability density estimation in physical systems. This technique integrates flow-based models with physical constraints to ensure that generated samples and density estimates respect the underlying laws governing the system. In contrast to the widely used Diffusion models and continuous-time normalizing flows, our approach has the following notable advantages. First, likelihood evaluation is not required, thereby accelerating model training. Second, exact density estimation is possible by solving the probability density transport equation. We demonstrate the framework by generating physically consistent solutions to the Darcy flow problem. We also present an application of our framework to Bayesian optimal experimental design, where highly expressive variational distributions can be constructed using conditional flow matching to estimate

the expected information gain.

Sahil Bhola
University of Michigan, Ann Arbor
sbhola@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS74

Entropy-Consistent Reduced-Order Models for Compressible Flow

In this work we propose a novel method to ensure important entropy inequalities are satisfied semi-discretely when constructing reduced order models (ROMs) on nonlinear reduced manifolds. We are in particular interested in ROMs of systems of nonlinear hyperbolic conservation laws. The so-called entropy stability property endows the semi-discrete ROMs with physically admissible behaviour. The method generalizes earlier results on entropy-stable ROMs constructed on linear spaces. The ROM works by evaluating the projected system on a well-chosen approximation of the state that ensures entropy stability. To ensure accuracy of the ROM after this approximation we locally enrich the tangent space of the reduced manifold with important quantities. Using numerical experiments on some well-known equations (the inviscid Burgers equation, shallow water equations and compressible Euler equations) we show the improved structure-preserving properties of our ROM compared to standard approaches and that our approximations have minimal impact on the accuracy of the ROM. We additionally generalize the recently proposed polynomial reduced manifolds to rational polynomial manifolds and show that this leads to an increase in accuracy for our experiments.

Robin Klein
Centrum Wiskunde & Informatica
rbk@cwi.nl

Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
b.sanderse@cwi.nl

MS74

Entropy Stable Reduced Order Modeling of Nonlinear Conservation Laws Using Discontinuous Galerkin Methods

We generalize the construction of entropy stable reduced order models (ES-ROMs) for nonlinear conservation laws from finite volume methods (FVM) to high order discontinuous Galerkin (DG) methods. This generalization preserves entropy stability while advancing the hyper-reduction steps by introducing new weighted test basis, utilizing the Caratheodory pruning for the hyper-reduction of boundary conditions, and implementing domain decomposition (DD).

Ray Qu, Jesse Chan
Rice University
rq10@rice.edu, jesse.chan@rice.edu

MS75

Multi-Domain EncoderDecoder Neural Networks for Efficient Reduced Order Data Assimilation

The first part of my talk introduces TorchDA, a novel Python package designed for non-explicit observation operators and GPU-accelerated data assimilation. Handling complex, high-dimensional systems is challenging due to the computational intensity and difficulty in defining observation functions. TorchDA integrates deep neural networks into data assimilation, serving as models for state transitions and observations. It includes Kalman Filter, EnKF, 3DVar, and 4DVar algorithms, offering flexible algorithm selection. Testing on the Lorenz 63 model and a 2D shallow water system shows significant performance improvements, effectively mapping between different physical spaces and providing a versatile tool for complex dynamical systems. The second part of my talk focuses on latent data assimilation for high-dimensional systems, where full space assimilation is computationally prohibitive. Latent methods operate in a reduced-order space but struggle with complex, nonlinear mappings. Recent advancements like Generalised Latent Assimilaiton (GLA) and Latent Space Data Assimilation (LSDA) address these issues but add computational costs and uncertainties. To overcome this, I present MEDLA, a deep-learning-based framework that uses a shared latent space to reduce computational demands and improve accuracy by minimizing interpolation errors. MEDLA consistently outperforms state-of-the-art methods in managing multi-scale data and complex mappings.

Sibo Cheng
CEREA, École des ponts, Institut Polytechnique de Paris
sibo.cheng@enpc.fr

Rossella Arcucci
Data Science Institute, Imperial College London
r.arcucci@imperial.ac.uk

Omar K. Matar
Imperial College London
o.matar@imperial.ac.uk

MS75

Local KarhunenLove Approximation of Global Gaussian Samples

For Bayesian inference, such as that applied to modeling incompressible flow, localizing the sampling strategy for the construction of the prior distribution has shown promise constraining the inherent variability of the posterior distribution. Specifically, the physical domain of the problem is decomposed into partitions, with sampling taking place in each subdomain. We investigate the approximation properties of localized Karhunen-Leve (KL) eigenfunctions in representing global Gaussian field samples, which are constructed from the KL expansion (KLE) of the global problem. For a given global sample, we assess two parameters: the number of subdomains (N) and the dimension of the local KLEs (LD). Aiming to design efficient Markov Chain Monte Carlo (MCMC) methods, we focus on optimizing N and LD and employing projection to reduce discrepancies at the interface between adjacent subregions. We assess the quality of the approximation of global samples via two criteria: (i) least squares errors and (ii) the number of iterations for convergence of the multiscale MCMC methods.

Aidan R. Gettemy
University of Texas at Dallas
arg210006@utdallas.edu

Marcio Borges
LNCC/MCTIC
marcio.rentes.borges@gmail.com

Felipe Pereira
The University of Texas at Dallas
luisfelipe.pereira@utdallas.edu

MS75

An Efficient Reduced-Order Model Based on Dynamic Mode Decomposition for Parameterized Pdes

Dynamic mode decomposition (DMD), as a data-driven method, has been frequently used to construct reduced-order models (ROMs) due to its good performance in time extrapolation. However, existing DMD-based ROMs suffer from high storage and computational costs for high-dimensional problems. To mitigate this problem, we develop a new DMD-based ROM, i.e., TDMD-GPR, by combining tensor train decomposition (TTD) and Gaussian process regression (GPR), where TTD is used to decompose the high-dimensional tensor into multiple factors, including parameter-dependent and time-dependent factors. Parameter-dependent factor is fed into GPR to build the map between parameter value and factor vector. For any parameter value, multiplying the corresponding parameter-dependent factor vector and the time-dependent factor matrix, the result describes the temporal behavior of the spatial basis for this parameter value and is then used to train the DMD model. In addition, incremental singular value decomposition is adopted to acquire a collection of important instants, which can further reduce the computational and storage costs of TDMD-GPR. The comparison TDMD and standard DMD in terms of computational and storage complexities shows that TDMD is more advantageous. The performance of the TDMD and TDMD-GPR is assessed through several cases, and the numerical results confirm the effectiveness of them.

Xiang Sun
School of Mathematical Sciences
Ocean University of China
sunxiang@ouc.edu.cn

MS75

A Non-Overlapping Optimization-Based Domain Decomposition Approach to Model Order Reduction of Fluid Flows

We present a model order reduction procedure to efficiently and accurately solve fluid-structure interaction problems. Our approach leverages a non-overlapping optimization-based domain decomposition technique to determine the control variable that minimizes jumps across the interfaces between sub-domains. To solve the resulting constrained optimization problem, we propose a sequential quadratic programming method, which effectively transforms the constrained problem into an unconstrained formulation. Furthermore, we integrate model order reduction techniques into the optimization framework, to speed up computations. Numerical results are presented to demonstrate the validity and effectiveness of the overall methodology.

Lei Zhang
School of Mathematical Sciences, Tongji University
22210@tongji.edu.cn

Tommaso Taddei
INRIA Bordeaux South-West
tommaso.taddei@inria.fr

Xuejun Xu
School of Mathematical Sciences
Tongji University
xxj@lsec.cc.ac.cn

MS76

Bayesian Optimization for Scientific Applications

Bayesian Optimization is a powerful technique for efficient optimization of black box functions, with many applications in the sciences. This talk will provide an overview of Bayesian Optimization, including important recent methodological advancements in areas such as multi-objective, multi-fidelity, and high-dimensional problems, and illustrate those using various scientific applications, including nuclear fusion, particle accelerator design, optics and sustainable concrete. This talk will also discuss software for Bayesian optimization and key open challenges.

Maximilian Balandat
Meta
balandat@meta.com

MS76

Gray-Box Bayesian Optimization in Function Spaces Using NEON - Neural Epistemic Operator Networks

Bayesian Optimization (BO) is a field of machine learning that aims to optimize a hard-to-compute function f in a sample-efficient way without access to gradients. This framework requires building probabilistic surrogate models to learn the behavior of f from data, often in a black-box manner. However, for many problems in science and engineering, there is additional information about the structure of f that can be exploited to produce more accurate and reliable surrogate models, replacing the 'black-box' nature with a 'gray-box' one. In this talk, I will briefly describe the composite BO framework and how it can be helpful for many problems in the real world. I will also talk about my new paper (https://www.nature.com/articles/s41598-024-79621-7), where we exploit the compositional structure of problems in science and engineering to create Neural Epistemic Operator Networks (Neon, for short), which allow for efficient BO in very high dimensions, including potentially infinite-dimensional cases of function spaces.

Leonardo Ferreira Guilhoto
University of Pennsylvania
guilhoto@sas.upenn.edu

MS76

Lucie: A Lightweight Uncoupled Climate Emulator with Long-Term Stability and Physical Consistency for O(1000)-Member Ensembles

We present LUCIE, a data-driven atmospheric emulator that remains stable during autoregressive inference for a thousand of years with minimal drifting climatology. LUCIE was trained using 9.5 years of coarse-resolution ERA5 data, incorporating 5 prognostic variables, 2 forcing variables, and one diagnostic variable (6-hourly total precipitation), all on a single A100 GPU over a two-hour period. LUCIE autoregressively predicts the prognostic variables and outputs the diagnostic variables similar to AllenAIs

ACE climate emulator. Unlike all the other state-of-the-art AI weather models, LUCIE is neither unstable nor does it produce hallucinations that result in unphysical drift of the emulated climate. The low computational requirements of LUCIE allow for rapid experimentation including the development of novel loss functions to reduce spectral bias and improve tails of the distributions. Furthermore, LUCIE does not impose true sea-surface temperature (SST) from a coupled numerical model to enforce the annual cycle in temperature. We demonstrate the long-term climatology obtained from LUCIE as well as subseasonal-to-seasonal scale prediction skills on the prognostic variables. LUCIE is capable of 6000 years of simulation per day on a single GPU, allowing for O(100)-ensemble members for quantifying model uncertainty for climate and ensemble weather prediction.

Haiwen Guan
Penn State University
hzg18@psu.edu

Troy Arcomano
Argonne National Laboratory
tarcomano@anl.gov

Ashesh Chattopadhyay
UC Santa Cruz
aschatto@ucsc.edu

Romit Maulik
Pennsylvania State University
rmm7011@psu.edu

MS76

Estimating Predictive Uncertainties of Neural Operators As Simulators of Oil Reservoirs Flows

Optimizing oil exploitation in underground reservoirs requires previous knowledge of the undersurface geology (leading to uncertainties in the modeling) and high-fidelity computer models. It involves complex multi-phase flows within a subsurface heterogenous medium, and its operational optimization and design rely on very expensive computational simulations. In order to cope with that, Scientific Machine Learning Surrogates [1] have been used, leading to affordable and accurate predictions on proper time. A critical component of reliable predictions is the ability to provide uncertainty quantification. Here, we systematically estimate uncertainties in such predictions due to building surrogates with limited data [2]. We concentrate the study on Fourier Neural Operators tailored for , and some specific geological scenarios are explored, employing different concepts and algorithms for quantifying epistemic uncertainty [2]. References: [1] Gege Wen, Zongy Li, Kamyar Azizzadeneshel, Anima Anandkumar, Sally M. Benson. U-FNOAn enhanced Fourier neural operator-based deep-learning model for multiphase flow. Advances in Water Resources, 163, 104180, 2022. [2] Apostolos F. Psaros, Xuhui Meng, Zongren Zou, Ling Guo, George E. Karniadakis, Uncertainty quantification in scientific machine learning: Methods, metrics, and comparison. Journal of Computational Physics, 477, 111902, 2023

Fernando A. Rochinha
Federal University of Rio de Janeiro
Rua Nascimento Silva 100 401
faro@mecanica.coppe.ufrj.br

Alvaro Coutinho
Dept. of Civil Engineering
Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

Rodolfo Freitas, Romulo Silva, Gabriel Barros
Federal University of Rio de Janeiro
rodolfodasilvamachadofreitas@qmul.ac.uk,
omulo.silva@nacad.ufrj.br, gabriel.barros@coc.ufrj.br,

MS77

AI Techniques for Multi-site Extremes for Use in Energy Systems Assessments

Inspired by recent extreme cold winter weather events in North America caused by atmospheric blocking, we examine several probabilistic generative models for the entire multivariate probability distribution of daily boreal winter surface air temperature. We propose metrics to measure spatial asymmetries, such as long-range anticorrelated patterns that commonly appear in temperature fields during blocking events. The model allows us to produce a conditional sampling approach that can be paired with a risk measure to address the inherent challenge in approximating the risk of low-frequency events within a sampling based approach. As a result, we present a model for spatially correlated, county-specific temperatures and a method to generate both unconditional and conditionally extreme temperature samples from this model efficiently. We demonstrate its consequences in a resource adequacy setting.

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Julie Bessac
US National Renewable Energy Laboratory
julie.bessac@nrel.gov

Mitchell Krock
University of Missouri
mk52n@missouri.edu

MS77

Generative Latent Diffusion Model for Spatiotemporal Inflow Turbulence

Eddy-resolving turbulence simulations often require the generation of stochastic inflow conditions. Traditionally, recycling-based methods are used to obtain turbulent inlets, which involve performing high-fidelity, computationally expensive simulations. Alternatively, synthetic inflow turbulence generators have been developed to synthesize inflow turbulence efficiently on-the-fly without precursor simulations. However, these methods often struggle to accurately replicate the coherent structures of true turbulence. Recent advancements in deep learning (DL) have provided new approaches for turbulent inlet generation. Existing research has utilized sequence models to learn dynamics from simulation data, but these deterministic models tend to suffer from error accumulation over long-term predictions. Probabilistic models, while potentially more robust, are challenging to train and often lack generalizability across different flow conditions. Additionally, most DL models fail to generate inflows that strictly adhere to physical laws, leading to turbulence decay at downstream. To address these challenges, we propose a DL network that leverages divergence-free basis together with

latent diffusion models (LDM). The turbulence statistical distribution is learned by LDM in a latent space encoded by the divergence-free basis. By enforcing the volume-conservative property of incompressible flow, our approach offers an effective turbulence generator applicable to various flow conditions.

Xin-Yang Liu, Meet Parikh, Xiantao Fan, Pan Du, Jian-Xun Wang
University of Notre Dame
xliu28@nd.edu, mparikh2@nd.edu, fanxiantao1@gmail.com, pdu@nd.edu, jwang33@nd.edu

MS77

Modeling Spatio-temporal Extremes with Conditional Variational Autoencoders

Extreme weather events are widely studied in the fields of agriculture, climatology, ecology, and hydrology, to name a few. Enhanced scientific understanding of the spatio-temporal dynamics of extreme events could significantly improve policy formulation and decision-making within these domains. We formulate a novel approach to model spatio-temporal extremes by conditioning on a time series (e.g., the El Nino-Southern Oscillation (ENSO) index) via a conditional variational autoencoder (extreme-CVAE). The prominent alignment of extremal dependences showcases the model's ability to be a spatio-temporal extreme emulator. Along with a decoding path, a convolutional neural network was built to investigate the relationship between the time series dynamics and parameters within the latent space, thereby inheriting the intrinsic temporal dependence structures. An extensive simulation validated the effectiveness and time efficiency of the model. We conducted an analysis of the monthly precipitation data which adequately demonstrates both the time efficiency and model performance of our approach in real-world scenarios.

Xiaoyu Ma, Likun Zhang, Christopher Wilke
University of Missouri
xm3cf@health.missouri.edu, likun.zhang@missouri.edu, wiklec@missouri.edu

MS78

Accelerating Phase Field Simulations Through a Hybrid Adaptive Fourier Neural Operator with U-Net

Computational simulation of phase field dynamics can be prohibitively expensive when using standard numerical solvers. For example, high-fidelity simulations often use very small time steps due to stability considerations, which can become a bottleneck when the target quantities of interest require predictions far out in time. To address this challenge, we employ machine learning-based surrogate models to help predict key dynamics forward in time, enabling predictions at time horizons far beyond what is achievable through traditional methods alone. Specifically, we propose U-AFNO, a novel model which incorporate a special kind of vision transformers, Adaptive Fourier Neural Operators (AFNO), with a U-Nets backbone. We train it to predict future states with much coarser time steps thus encapsulating multiple high-fidelity steps within a single surrogate evaluation. We show that our model is capable of predicting accurately chaotic phase-field dynamics over long time roll-outs, and recovers key quantities of interest describing the global phase field dynamics. Most notably, we also show that our model predictions are statistically indistinguishable from simulations otherwise predicted by the numerical solver. Furthermore, we discuss novel data augmentation strategies using generative models to effectively alleviate the need for numerous high-fidelity training data, while reducing prediction errors.

Christophe Bonneville
Sandia National Laboratories
cpbonne@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

MS78

Model Reduction for Kohn-Sham Density Functional Theory

Density Functional Theory (DFT) is a formulation of the electronics structure problem that offers a practical balance between computational complexity and accuracy. As the predictive tool of choice for materials science applications, DFT codes account for a large fraction of the utilization on many supercomputing systems. However, the high computational cost associated with the solution of the DFT equations poses a significant limitation for practical quantum molecular dynamics (MD) simulations, necessary to obtain meaningful inference for many realistic systems. To address this limitation, we investigate data-driven strategies which combine linear scaling algorithms for DFT and reduced order models (ROMs), constrained by the physics, to enable the exploration of intrinsic low dimensional subspace and manifold through data. Combining ROMs with DFT provides a natural way to incorporate physics through the underlying PDE with favorable implications for error control, yielding low computational cost, yet high fidelity solutions. In this talk, we will present the data-driven strategies and numerical results on prototype problems which exhibit the technical difficulties posed by Kohn-Sham Density Functional Theory.

Siu Wun Cheung, Youngsoo Choi, Seung Whan Chung
Lawrence Livermore National Laboratory
cheung26@llnl.gov, choi15@llnl.gov, chung28@llnl.gov

Jean-Luc Fattebert
Oak Ridge National Laboratory
fattebertj@ornl.gov

Coleman Kendrick, Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
kendrick6@llnl.gov, oseikuffuor1@llnl.gov

MS78

Transformer Models in Continuum Mechanics

Transformer models have gained popularity as solvers of partial differential equations because their attention mechanisms can be easily transferred to spatio-temporal data. We use the vision transformer architecture along with axial and cross attention to learn spatio-temporal dynamics from large scale direct numerical simulation data in continuum solid mechanics. It is of interest to explore the extension of

recent results suggesting that pre-training on multiphysics simulations leads to higher performance upon focused fine tuning, as well as to study the limits of inference with this framework.

Rahul Gulati, Benjamin Jasperson
University of Southern California
rahulgul@usc.edu, bjaspers@usc.edu

Krishna Garikipati
Aerospace & Mechanical Engineering
University of Southern California
garikipa@usc.edu

MS79

Learning Methods in Collective Motion

Interacting particle systems, also known as agent-based models, are utilized to study a wide range of physical phenomena across multiple scales; examples include the collective motion of bacteria, the flocking of birds and other animal species, and the coordination of mobile networks. Most such systems exhibit a form of emergence: local interactions leading to large-scale coordination. A fundamental scientific question is thus discovering local interaction laws which lead to observed behavior. To answer this question we consider a general model structure based on a distance based interaction kernel that can describe a wide variety of interacting agent systems. We present a number of machine learning methods that can be utilized to infer the interaction kernel from trajectory data. Specifically, we compare a number of non-parametric learning methods, including a novel neural ordinary differential equation method. We also extend our framework to stochastic differential equations, and present methods for learning drifts in the presence of both known and unknown noise, and the systems considered take a very general form. Many applications are demonstrated, demonstrating the ability of our methods to learn interaction mechanisms from trajectory data alone; models considered include the flocking of birds, the schooling of fish, and cancer cell migration.

James Greene
Clarkson University
jgreene@clarkson.edu

MS79

Interacting Particle Systems on Networks: Joint Inference of the Network and the Interaction Kernel

Modeling multi-agent systems on networks is a fundamental challenge in a wide variety of disciplines. We jointly infer the weight matrix of the network and the interaction kernel, which determine respectively which agents interact with which others and the rules of such interactions from data consisting of multiple trajectories. The estimator we propose leads naturally to a non-convex optimization problem, and we investigate two approaches for its solution: one is based on the alternating least squares (ALS) algorithm; another is based on a new algorithm named operator regression with alternating least squares (ORALS). Both algorithms are scalable to large ensembles of data trajectories. We establish coercivity conditions guaranteeing identifiability and well-posedness. The ALS algorithm appears statistically efficient and robust even in the small data regime but lacks performance and convergence guarantees. The ORALS estimator is consistent and asymptotically normal under a coercivity condition. We conduct several numerical experiments ranging from Kuramoto particle systems on networks to opinion dynamics in leader-follower models.

Quanjun Lang
Duke University
quanjun.lang@duke.edu

MS79

Radius of Information for Two Intersected Centered Hyperellipsoids and Implications in Optimal Recovery from Inaccurate Data

For objects belonging to a known model set and observed through a prescribed linear process, we aim at determining methods to recover linear quantities of these objects that are optimal from a worst-case perspective. Working in a Hilbert setting, we show that, if the model set is the intersection of two hyperellipsoids centered at the origin, then there is an optimal recovery method which is linear. It is specifically given by a constrained regularization procedure whose parameters can be precomputed by semidefinite programming. This general framework can be applied to several scenarios, including the two-space problem and problems involving l2-inaccurate data. It can also be applied to the problem of recovery from l1-inaccurate data. For the latter, we reach the conclusion of existence of an optimal recovery method which is linear, again given by constrained regularization, under a computationally verifiable sufficient condition.

Chunyang Liao
Texas A&M University
liaochunyang@tamu.edu

MS79

Imaged-based Modeling for Cardiac Physiology

This talk will describe our efforts to construct imaged-based computer models of cardiac physiology. We will focus on an application in pediatric cardiology related to defective coronary artery anatomy. Models are constructed from computed tomography data and boundary conditions are estimated from cardiac catherization data. Attention will be paid to how one might define the "model validation" in this context and how one can use surrogate models for calibrating boundary condition parameters to available clinical data.

Charles Puelz
University of Houston
cepuelz@Central.UH.EDU

MS79

Comparing Image Reconstructions of Different Methods for Simulated Eit Data

Electrical Impedance Tomography (EIT) is a portable, inexpensive, noninvasive imaging system that does not use ionizing radiation. Due to these properties, EIT has been of great interest for many different medical applications. In practice, electrodes are attached to a patient's body where current is applied, and the resulting voltage is measured to reconstruct internal electrical properties. However, the EIT inverse problem is severely ill-posed due to the limited amount of measurements from the partial boundary data, and aggressive changes in the electrical properties inside the body cause regularization difficulties. Here, we will compare both non-iterative and iterative methods, as well

as machine learning methods, to improve image reconstructions. These methods will be compared using simulated data with anomalies placed in varying locations.

Jacob Roarty, Malena I. Espanol
Arizona State University
jroarty@asu.edu, malena.espanol@asu.edu

MS80

An Accurate Method of Fundamental Solutions for Large and Dense Particle Systems in Stokes Flow

The Stokes mobility and resistance problems for rigid particles in viscous fluids have broad applications in both natural processes and industrial contexts and are important in the study of transport or diffusion processes, rheology and nonlinear shear thickening. Hydrodynamic interactions between particles are challenging to determine, as they are simultaneously long-ranged and expensive to resolve for dense suspensions. The latter is caused by near-singular lubrication forces resulting from close-to-touching bodies in relative motion. With the aim of controlling the accuracy for dense suspensions with a computationally cheap method, we present a new technique for the two problems that combines the method of fundamental solutions (MFS) with the method of images. For rigid spheres, we propose to represent the flow using fundamental solutions (Stokeslets) located on interior spheres, augmented by lines of image sources adapted to each near-contact to resolve lubrication. Source strengths are found by a least-squares solve at contact-adapted boundary collocation nodes. In the talk, we also discuss efficient preconditioning. The resulting well-conditioned schemes are competing with state-of-the art solvers tailored for spherical particles in Stokes flow, yet are compatible also with other smooth particle shapes, and scale linearly in the number of objects. For instance, a problem of 10,000 ellipsoids is solved to 5-digit accuracy on a workstation in less than two hours.

Anna Broms
KTH Royal Institute of Technology
annabrom@kth.se

Alex H. Barnett
Flatiron Institute, Simons Foundation
abarnett@flatironinstitute.org

Anna-Karin Tornberg
KTH
akto@kth.se

MS80

Parallel Set-Valued Rational Approximation for PDEs

In the solution of (parametrized) PDEs and non-linear eigenvalue problems from physics and engineering applications, often large sets of functions have to be approximated, containing singularities in either the spatial or frequency domain. This typically results in a prohibitively large computational and memory cost. Both can be reduced by switching to set-valued approximation. Here, we present a fast and parallel set-valued rational approximation algorithm for (parametrized) PDEs. Our method, called PQR-AAA, is a set-valued variant of the Adaptive Antoulas Anderson (AAA) algorithm, accelerated by using local approximate orthogonal bases obtained from a truncated QR decomposition. We demonstrate its connection to interpolative decompositions and present various theoretical insights. We focus in particular on gluing together separate approximations built in parallel, while maintaining prescribed accuracy. We demonstrate both theoretically and numerically this method's accuracy and effectiveness. Various reference problems and real-world 3D applications are presented. This talk is based on joint work with Daan Huybrechs and Karl Meerbergen.

Simon Dirckx
KU Leuven
simon.dirckx@austin.utexas.edu

MS80

Computing Greens Functions Using Recurrences

Singular and nearly sing8lar integrals involving Cauchy and log kernels of weighted orthogonal polynomials solve simple three-term recurrences. By a careful choice between forward recurrence and (F.W.J.) Olver's algorithm we can efficiently and accurately compute these singular integrals, arbitrarily close or even on the element. These techniques extend to log kernel integrals on the square where one needs to simultaneously utilise two bivariate recurrence relationships.

Sheehan Olver
Imperial College London
s.olver@imperial.ac.uk

MS80

Recursive Reduction Quadrature for the Evaluation of Laplace Layer Potentials in Three Dimensions

A high-order quadrature rule is constructed for the evaluation of Laplace single and double layer potentials and their normal derivatives on smooth surfaces in three dimensions. The construction starts from a harmonic approximation of the density on each patch that admits a natural harmonic polynomial extension in a volumetric neighborhood of the patch. Then by the general Stokes theorem, singular and nearly singular surface integrals are reduced to line integrals preserving the singularity of the kernel. These singularity-preserving line integrals can be evaluated semi-analytically by the singularity-swap quadrature involving a complex-valued parameter root finding for the given target and polynomial approximation of the smooth part on the boundary of the patch. In other words, the evaluation of singular and nearly singular surface integrals is reduced to function evaluations on the vertices on the boundary of each patch. The recursive reduction quadrature largely removes adaptive integration that is needed in most existing high-order quadratures for singular and nearly singular surface integrals, leading to improved efficiency and robustness.

Hai Zhu
Flatiron Institute
hzhu@flatironinstitute.org

Shidong Jiang
Center for Computational Mathematics
Flatiron Institute, Simons Foundation

sjiang@flatironinstitute.org

MS81

Low-Rank Methods for Nonlinear Equations

Recently there has been a growing interest in low-rank methods for solving matrix and tensor differential equations. Traditional approaches to computing low-rank solutions require the equation of interest to be expressible in a low-rank format. However, many equations, particularly nonlinear ones, do not fit into this framework. To overcome this challenge, we introduce novel algorithms that leverage sparse sampling techniques to compute low-rank solutions to nonlinear differential equations. We will demonstrate the effectiveness of these methods through various numerical applications.

Alec Dektor
Lawrence Berkeley National Laboratory
adektor@lbl.gov

MS81

The Runge-Kutta Sparse Grid Discontinuous Galerkin Method with Stage-Dependent Mesh for Transport Equations

In this talk, we present a class of Runge-Kutta (RK) sparse grid discontinuous Galerkin (DG) methods with stage-dependent mesh for transport equations. The new method extends beyond the traditional method of lines framework and utilizes stage-dependent sparse grid DG finite element spaces for the spatial discretization operators. It features fewer floating-point operations and may achieve larger time step sizes. Numerical tests for transport problems in high dimensions are provided to demonstrate the performance of the new method.

Juntao Huang
Texas Tech University
juntao.huang@ttu.edu

MS81

On Asymptotic-Preserving Dynamical Low-Rank Approximation for Thermal Radiative Transfer

In recent years, dynamical low-rank approximation (DLRA) has gained popularity for efficiently handling high-dimensional kinetic equations. Especially the flexibility provided by basis-update & Galerkin (BUG) integrators has allowed the construction of numerical schemes that account for a wide range of properties. In this work, we are interested in constructing a low-rank scheme using the parallel BUG integrator that, on a discrete level, preserves the long-timescale dynamics of the thermal radiative transfer (TRT) equations. Numerical schemes that do so consistently, stably, and efficiently are called asymptotic-preserving (AP) schemes. TRT equations are used to model the formation of thermal radiation fronts propagating in a cold medium and observed in phenomena like supernova explosions, star formations, and radiation emitted from a Hohlraum striking a fusion target. On a long time scale, they converge to the non-linear diffusion-type PDE called the Rosseland equation. The key challenges for constructing an AP scheme for the TRT equations, that captures the limiting Rosseland equation, lie in choosing an appropriate low-rank ansatz and ensuring that the truncation step preserves the range space of the substeps. We propose a low-rank AP scheme based on a macro-micro ansatz of the particle density and a range-preserving truncation step. The proposed scheme is shown to be energy-stable with a CFL condition that is independent of the Knudsen number in the limiting case.

Chinmay Patwardhan
Karlsruhe Institute of Technology
chinmay.patwardhan@kit.edu

Jonas Kusch
Norwegian University of Life Sciences
jonas.kusch@nmbu.no

MS81

A Hybrid Amr Tensor-Train Algorithm for Solving the 2D3V Boltzmann-Bgk Equation

The Boltzmann equation describes the time evolution of a particle distribution function in a six-dimensional position-velocity phase space. The exponential growth in computational complexity often challenges a grid-based approach to modeling the Boltzmann equation as the dimensionality grows. Scalable low-rank tensor decomposition techniques have recently been developed with applications to high-dimension PDEs to address this issue. Despite the remarkable progress made in the community, low-rank structures in the phase-space are not evident in realistic engineering systems with complex geometries (e.g., electric propulsion systems and fusion reactors), where discontinuities, shocks, complex boundary conditions, and material-dependent physics (e.g., collisions, fusion reactions, ionization/excitation, charge-exchange processes) pose formidable challenges. In this talk, we propose a novel hybrid algorithm where quad-tree adaptive mesh refinement (AMR) is applied in real space while a low-rank tensor-train approximation is applied in the velocity space. The AMR algorithm efficiently handles challenges pertaining to complex structures in real space, while the tensor-train formulation targets dimensionality challenges in the velocity space. We present preliminary results on the new algorithm applied to challenging multi-dimensional gas kinetics problems.

William T. Taitano
Los Alamos National Laboratory
taitano@lanl.gov

Samuel Araki
Amazon
samuel.j.araki@gmail.com

MS82

GPU-Accelerated Vecchia Algorithms for Scalable Gaussian Process Modeling with Batched Matrix Computations

Gaussian Processes (GPs) are widely used in geospatial analysis but face significant computational challenges when applied to large datasets, primarily due to the need to invert n x n covariance matrices, where n represents the number of spatial locations. Approximation techniques, such as low-rank and sparse methods, have been developed to address this issue, with the Vecchia approximation standing out as one of the most promising approaches. This study introduces parallel GPU implementations of the classic Vecchia method and newly proposed algorithms - Block Vecchia and Scaled Vecchia - leveraging batched matrix computations to enable efficient likelihood estimation in Maximum Likelihood Estimation (MLE). The Block

Vecchia algorithm evaluates multivariate conditional likelihoods for location blocks in parallel, offering improved scalability, efficiency, and reduced storage requirements for large-scale problems compared to the classic Vecchia approach. Meanwhile, the Scaled Vecchia algorithm improves accuracy in high-dimensional problems. This study makes several key contributions. (1) It demonstrates significant speedups for the classic, Block, and Scaled Block Vecchia algorithms on various GPUs; (2) The Block Vecchia algorithm improves modeling accuracy by using larger conditioning sizes and block counts, outperforming the classic Vecchia approach; (3) The Block Vecchia algorithm achieves substantial computational speedups, up to 80 faster than the classic Vecchia method, without any loss in accuracy; (4) The Block Vecchia algorithm exhibits enhanced scalability, effectively handling problems that are 40 larger than those manageable by the classic algorithm; (5) The scaled block Vecchia algorithm addresses high-dimensional problems that were previously infeasible to solve using the classic or Block Vecchia approaches. Experiments were performed to evaluate accuracy using statistical methods on synthetic and real datasets and performance assessments on single and multiple GPUs.

Sameh Abdulah
Extreme Computing Research Center (ECRC)
King Abdullah University of Science and Technology
sameh.abdulah@kaust.edu.sa

MS82

Engineering Fast Kernels for Rotation-Equivariant Deep Neural Networks

Rotation-equivariant deep neural networks, i.e. networks that preserve certain geometric relationships between their inputs and outputs, achieve state of the art performance on interatomic potential calculations. A core calculation in these networks is the Clebsch-Gordon tensor product, a kernel that multiplies a highly-structured sparse matrix against a Kronecker product of two dense feature vectors. Due to its low arithmetic intensity and irregular data access pattern, the kernel remains a significant obstacle to scaling equivariant networks to large workloads. We introduce a sparse CUDA kernel generator for the CG tensor product with performance matching or exceeding the best existing open and closed-source implementations. To accomplish this, we used template metaprogramming to engineer kernels that maximize instruction-level parallelism while minimizing global memory transactions. We generate a computation schedule ahead of time to manage the limited pool of GPU shared memory and offer operators that fuse the CG tensor product with graph convolution. We also provide kernels for the gradients of the CG tensor product along with a novel mathematical identity to compute higher-order partial derivatives. We provide up to 5.4x inference-time speedup over the original version of the MACE foundation model, with planned work to support new autograd frameworks and accelerators.

Vivek Bharadwaj
University of California, Berkeley
vivek_bharadwaj@berkeley.edu

Austin Glover
University of California Berkeley
austin_glover@berkeley.edu

Aydin Buluç
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

James W. Demmel
UC Berkeley
Division of Computer Science
demmel@berkeley.edu

MS82

Mixed Precision HODLR Matrices

In this talk, we demonstrate that off-diagonal blocks of HODLR matrices that are approximated by low-rank matrices can be represented in low precision without degenerating the quality of the overall approximation (with the error growth bounded by a factor of 2). We also present an adaptive-precision scheme for constructing and storing HODLR matrices, and we verify that the use of mixed precision does not compromise the numerical stability of the resulting HOLDR matrix-vector product and LU factorization. That is, the resulting error in these computations is not significantly greater than the case where we use one precision (say, double) for constructing and storing the HOLDR matrix. Our analyses further give insight into how one must choose the working precision in HOLDR matrix computations relative to the approximation error to mitigate the downside of finite precision.

Xinye Chen
Univerzita Karlova
Xinye.Chen@lip6.fr

MS82

Simulating Quantum Clifford Circuits with Gaussian Elimination

Quantum circuits are considered more powerful than classical circuits but require exponential resource to simulate. Clifford circuits are a special class of quantum circuits that can be simulated in polynomial time but still show important quantum effects such as entanglement. In this work, we present an algorithm that simulates Clifford circuits by performing Gaussian elimination on a modified adjacency matrix derived from the circuit structure. By utilizing efficient sparse Gaussian elimination algorithms, we expect to achieve good performance on Clifford circuits with sparse structures that are common in practice. Moreover, by representing general quantum circuits using Clifford gates and T gates, this algorithm could potentially accelerate general quantum circuit simulation.

Yuchen Pang
Department of Computer Science
University of Illinois at Urbana-Champaign
yuchenp2@illinois.edu

MS82

Secure Peer-to-Peer Matrix Multiplication for 100+ Players

We describe secure multiparty computation (MPC) algorithms for matrix multiplication. Our algorithms are motivated by the challenges of large networks of battery-powered drones: many players, vulnerability to fail-stop faults, and severe bandwidth limitations. Given shares of two matrices distributed among a large number of players n, we describe an algorithm to compute shares of their product. Our algorithm is scalable in n: each player sends and receives amortized O(1) messages per matrix multiplication for O(vn) matrix multiplications. The algorithm is private against coalitions of size T(vn) and tol-

erates T(vn) fail-stop faults. We implement our matrix multiplication algorithm using Cicada, a general-purpose open-source MPC software library, and demonstrate that it works well in practice as the engine of a privacy-preserving linear regression method with n = 100.

Cynthia Phillips
Sandia National Laboratories
caphill@sandia.gov

MS83

Multi-Mesh Adaptive Finite Element Methods for Kohn-Sham Eigenvalue Problem

The divide-and-conquer (DAC) method is a popular way to reduce the computational cost in the solution of the Kohn-Sham (KS) eigenvalue problem, which is commonly implemented by the domain decomposition technique and involves breaking down the computational domain into smaller subdomains. In this work, we explore an alternative approach to the DAC method by splitting the KS eigenvalue problem rather than the computational domain. Specifically, we propose a splitting method that divides the eigenpairs into groups and solves them on different approaching spaces. The splitting of the eigenpairs is based on the regularity of the corresponding wavefunctions, and the multi-mesh adaptive finite element method is employed to tailor the discretized space for each group. Numerical experiments validate the effectiveness and efficiency of the proposed method.

Yang Kuang
School of Mathematics and Statistics
Guangdong University of Technology
ykuang@gdut.edu.cn

MS83

Inchworm-Based Simulation of Open Quantum Systems with Tensor Networks

We study the dynamical simulation of open quantum spin chains with nearest-neighbor coupling, where each spin interacts with a harmonic bath. Building on our prior research, we extend the inchworm method and modular path integrals to handle both diagonal and off-diagonal couplings. Our method first decouples the system into single-spin problems, which are solvable via the inchworm method. By structuring the data appropriately, we treat each spins quantities as tensors, allowing us to connect individual spins into a complete system using tensor operators. Additionally, to reduce computational and memory costs in long-time simulations, we apply tensor-train representation to efficiently represent the reduced density matrix of the spin chains and employ the Transfer Tensor Method (TTM) to avoid the exponential growth of computational cost over time. Numerical tests on simple examples will be conducted to validate our method.

Yixiao Sun
Department of Mathematics
National University of Singapore
s.yixiao@u.nus.edu

Geshuo Wang, Zhenning Cai
National University of Singapore
geshuowang@u.nus.edu, matcz@nus.edu.sg

MS83

Tensor Neural Networks and Their Applications to Eigenvalue Problems in Quantum Mechanics

In this presentation, we will introduce our recently proposed Tensor Neural Network architecture and its application to high-dimensional eigenvalue problems, especially to eigenvalue problems in quantum mechanics. Tensor Neural Networks are designed to avoid the dependence of numerical integration on the Monte-Carlo scheme in loss functions. By converting high-dimensional integration into one-dimensional integration, high-dimensional integration can be computed with high precision using classical numerical integration formulas. The characteristics of TNN can simplify computational complexity and significantly improve the accuracy of neural network methods, making TNN far more accurate than other methods, significantly when solving high-dimensional eigenvalue problems. We will show the high accuracy of the TNN method for solving multi-state problems in quantum mechanics and the prospect of application in other issues. In addition, we will introduce other researchers' discussion of TNN. Finally, we will introduce the implementation of the TNN structure efficiency program and the corresponding algorithm.

Yifan Wang
School of Mathematical Sciences
Peking University
wangyifan1994@pku.edu.cn

MS83

A Gradient Flow Model for Wigner Ground State Calculations Based on Density Functional Theory

In this talk, a gradient flow model is proposed for conducting Wigner ground state calculations of many-body system in the framework of density functional theory. Firstly, a gradient flow model is derived to obtain the ground state Wigner function of one-body system. Subsequently, this model is extended to many-body system within the context of density functional theory. To ensure computational efficiency, the Fourier pseudo-spectral method and an operator splitting method are introduced for numerical simulations. Two toy models are considered to validate our approach, including the periodic extension of a one-dimensional delta-interacting system with a local density approximation, and a three-dimensional system with Coulomb interaction. Results from these two models successfully demonstrate the potential of the method toward large-scale systems and the system with defects.

Hongfei Zhan
School of Mathematical Sciences
Peking University
hfzhan@nus.edu.sg

MS84

Sparse, Empirically Optimized Quadrature for Broadband Spectral Integration

Integration over the highly nonlinear electromagnetic spectrum dominates the expense of broadband radiative calculations, and climate and weather models require fast parameterizations to characterize the flow of energy through the atmosphere. Though modern parameterizations are efficient and accurate, they are finely tuned for specific appli-

cations and can be difficult to understand. Here we present a data-driven quadrature scheme that approximates the spectral flux integral by performing a weighted sum of monochromatic calculations at a small subset of points. This subset is chosen combinatorically via simulated annealing, quadrature weights are found by solving a linear program, and both optimize the same cost function. The cost function allows us to flexibly prioritize computing quantities of interest including flux and flux divergence profiles in atmospheres with present day ranges in temperature and gas concentration. Though trained on cloudless profiles, the longwave (thermal) quadrature scheme extrapolates well to cloudy atmospheres. We update the optimization procedure for shortwave (sunlight) calculations, which must be robust to variations in solar zenith angle and surface reflectivity while maintaining a small training dataset, and apply the approach to include variability in all greenhouse gases for flexibility in different climate scenarios. Together the results show data driven quadrature offers a versatile and efficient alternative to traditional methods.

Paulina Czarnecki
Columbia University
pc2943@columbia.edu

Manfred Brath
University of Hamburg
manfred.brath@uni-hamburg.de

Robert Pincus
Lamont Doherty Earth Observatory
robertp@ldeo.columbia.edu

Lorenzo Polvani
Department of Applied Physics and Applied Mathematics
Columbia University
polvani@columbia.edu

MS84

The Benefits of Embracing Complexity and Solving the Inverse 3D Radiative Transfer Equation for Atmospheric Remote Sensing

Satellite remote sensing using shortwave radiation is a powerful tool to measure the microphysical properties of clouds. Operationally, 1D radiative transfer models are employed to invert radiance measurements. These models neglect the interactions of multiple scattering with cloud heterogeneities, resulting in systematic errors in basic quantities such as cloud optical thickness that can reach ~70% in heterogeneous, cumuliform clouds. The solution to this issue is to solve the inverse 3D radiative transfer equation to retrieve the 3D microphysical properties of the atmosphere. This requires an increase in measurement information content. Our approach is to utilize high-resolution, multi-angle imagery to solve the non-linear scattering tomography problem using a fast approximation to the adjoint. We quantify the accuracy with which this problem can be solved and describe the recent advances in the theory of cloud remote sensing that have emerged through exploration of this concept. We find that many of the numerical challenges in our solution of the radiative transfer equation in cloudy atmospheres, whether they are angular singularities in the source or the presence of spatial variability, become our greatest allies in performing cloud tomography. This is because they increase the information content of the measurements in the face of the loss of precision that occurs when inverting multiply scattered radiation from optically thick clouds.

Jesse Loveridge
Colorado State University
jesse.loveridge@colostate.edu

Aviad Levis
Technion Israel Institute of Technology
aviad.levis@gmail.com

Larry Di Girolamo
University of Illinois Urbana-Champaign
gdi@illinois.edu

Vadim Holodovsky
Technion Israel Institute of Technology
dvdholod@gmail.com

Linda Forster, Anthony B. Davis
Jet Propulsion Laboratory
California Institute of Technology
linda.i.forster@jpl.nasa.gov, anthony.b.davis@jpl.nasa.gov

Yoav Schechner
Technion Israel Institute of Technology
yoav@ee.technion.ac.il

MS84

Exploring Aerosol-Meteorology Interactions and Radiative Forcing Through Measurement and Model Integration

Recent changes in the global aerosol burden have significantly altered aerosol direct radiative forcing (DRF), which indicates aerosol influence on climate via perturbation of the radiation budget. The DRFs estimated using the Clouds and the Earth's Radiant Energy System (CERES) data show notable spatial and temporal variations of DRF across the South and Southeast United States over recent decades. The TRacking Aerosol Convection Interactions Experiment (TRACER) provides a unique opportunity to study aerosol processes that affect climate in the coastal environment of the Southern Texas region. The changes in DRF are considered to be primarily driven by changes in anthropogenic activities associated with population growth and urbanization. However, the effect of interactions between the aerosol environment and mesoscale meteorological phenomena on DRF is less studied. A comprehensive suite of measurement data is used to understand these processes, complemented by detailed meteorological observations during the summer over the TRACER environment. DRF is estimated using the measurement-informed Optical Properties of Aerosols and Clouds (OPAC) software package and the Santa Barbara DISORT Atmospheric Radiative Transfer (SBDART) radiative transfer model. Additionally, the WRF-Chem model simulations performed over the southern Texas region further elucidate the impact of meteorological phenomena and aerosol environment on the DRF from local to regional scales.

Tamanna Subba
Brookhaven National Laboratory
tsubba@bnl.gov

MS84

Angular-Spatial Hp-Adaptivity for Radiative

Transfer

Radiative transfer is important for many science and engineering applications, and numerical simulations of radiative transfer can be challenging. For instance, the radiation field is seven-dimensional - three spatial, two angular, one wavelength, and one temporal - and often features steep gradients. Therefore, memory usage is a key issue. To reduce memory, some past work has investigated the use of adaptive mesh refinement (AMR), typically for either the spatial or angular coordinate, and typically for only h-adaptivity. Here, we propose the use of AMR for the spatial and angular coordinates together, and the use of h- and p-adaptivity together as hp-AMR for the potential for further memory savings. We implemented the proposed method for several test cases in two spatial and one angular dimension, with the discontinuous Galerkin spectral element method. Our primary findings from these test cases were: (1) Angular hp-adaptivity can deliver the radiation solution with the same accuracy as, and with much less computational memory than, uniform angular h- or p-refinements, or angular h-adaptivity alone. (2) Full spatial-angular hp-adaptivity is more efficient in solution representation, compared to solely spatial or solely angular hp-adaptivity. These results suggest that adaptive angular-spatial hp-refinement may perform well in large-scale seven-dimensional applications.

Jason L. Torchinsky
Sandia National Laboratories
jltorch@sandia.gov

MS86

Graph p-Laplacian Eigenpairs via Spectral Energy Functions

Several problems in graph theory, such as graph clusetring and core-periphery detection, can be related to appropriate nonlinear eigenvalue problems. These eigenvalue problems are formulated in terms of the critical point equation of some p-Rayleigh quotient, where the p-norms appear in both the numerator and the denominator. In this talk we discuss the particular case of the graph p-Laplacian spectrum, with focus on the limit cases $p = 1$ and $p = \infty$. Indeed the 1-Laplacian and the ∞-Laplacian spectrum provide approximations of the Cheeger cuts and the packing radii of the graph, respectively. Our main goal is to present a new method for computing the p-Laplacian eigenpairs. To this end, we first reformulate the graph p-Laplacian eigenproblem in terms of a constrained weighted 2-Laplacian eigenvalue problem. Then, based on this reformulation, we introduce a class of spectral energy functions whose differentiable saddle points correspond to p-Laplacian eigenpairs. In particular, any such eigenpair is such that the Morse index of the p-Rayleigh quotient in the eigenfunction matches the index of the energy function. Finally, based on gradient flows for the energy functions, we propose numerical methods for the computation of nonlinear p-Laplacian eigenpairs as limit points of sequences of linear generalized eigenpairs.

Piero Deidda
Scuola Normale Superiore
Gran Sasso Science Institute
piero.deidda@sns.it

Nicola Segala
Department of Mathematics
University of Padova
nicola.segala@math.unipd.it

Francesco Tudisco
University of Edinburgh
School of Mathematics
f.tudisco@ed.ac.uk

Martin Burger
DESY Hamburg
Universität Hamburg
martin.burger@desy.de

Mario Putti
Department of Mathematics “Tullio Levi-Civita’
University of Padua, Italy
mario.putti@unipd.it

MS86

Doubling Algorithms for Eigenvalue Problems

Doubling algorithms have demonstrated extreme successes in solving several high profile nonlinear matrix equations of practical significance such as those arising from optimal control theory, applied probability and transportation theory, Markov-modulated fluid queue theory, and quantum transport in nano research. A key point that makes all possible is to embed the solutions of those nonlinear matrix equations into basis matrices of the eigenspaces of certain matrix pencils, but doing so demands that the eigenspaces admit basis matrices having particular structural properties. The latter limits the applicability of existing doubling algorithms for eigenvalue problem in general. In this talk, we will present a unifying framework of doubling algorithms that can disregard the structural property to work for any regular matrix pencils. Existing doubling algorithms fall into the unifying framework.

Ren-Cang Li
University of Texas - Arlington
rcli@uta.edu

MS86

Nonlinear Eigenvector Problems and Optimization of Homogeneous Functions

Eigenvector-dependent nonlinear eigenvalue problems (NEPv) are commonly solved in various fields such as machine learning, signal processing, and computational physics. These problems typically arise as algebraic characterizations for the solutions of optimization problems, where a homogeneous objective function is optimized under an equality constraint involving another homogeneous function. In this talk, we will introduce a unified NEPv characterization for such homogeneous optimization problems. Unlike traditional approaches that rely solely on first-order optimality conditions, our NEPv incorporates additional second-order information, which helps address the eigenvalue position issue in NEPv. Additionally, we demonstrate that the self-consistent-filed iteration exhibits superlinear convergence when applied to these NEPv. We will cover the theoretical foundations, implementation details, and practical applications of this approach.

Ding Lu
University of Kentucky
ding.lu@uky.edu

MS87

COLABS: Collaboration for Better Software (for

Science)

The COLABS project focuses on training, workforce development, and the research software engineering (RSE) community of practice (CoP) as part of DOEs Next-Generation Scientific Software Technologies program and a member of the Consortium for the Advancement of Scientific Software (CASS). COLABS develops and delivers software engineering and software development training, and coordinates training efforts across CASS. Relatedly, COLABSs primary workforce effort entails the development of curricula to facilitate new hires gaining skills they might lack, thus allowing hiring managers to cast a broader net when recruiting staff for software-focused positions. Additionally, we promote and facilitate internship opportunities in software-focused projects for students from underrepresented groups. Finally, we collaborate with the U.S. Research Software Engineer Association to encourage sharing technical and professional knowledge and experiences to build the RSE CoP in the national labs.

Anshu Dubey, Rinku Gupta
Argonne National Laboratory
adubey@anl.gov, rgupta@anl.gov

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

Gregory Watson
Oak Ridge National Laboratory
watsongr@ornl.gov

Dan Gunter
LBL
dan gunter ¡dkgunter@lbl.gov¿

Keith Beattie
Lawrence Berkeley Lab
ksbeattie@lbl.gov

MS87

The CASS Approach: a Consortium for the Advancement of Scientific Software

Many of the scientific discoveries that shape our world are made possible by a sophisticated ecosystem of software technology. This software ecosystem has evolved over decades of research with input from countless stakeholders and application scientists. Sustaining it and ensuring a strong foundation for the future requires a comprehensive and thoughtful approach to software stewardship and advancement. The Consortium for the Advancement of Scientific Software (CASS) was formed by the Department of Energy in 2024 to help address this need. It is comprised of a collection of complementary, agile organizations that support specific scientific software communities and provide crosscutting services. These organizations work together under a unified consortium framework to coordinate effort, pool resources, and engage with the community for maximum impact. The purpose of this talk is to summarize the high-level mission, history, and structure of CASS and explore how it enables the range of activities that will be highlighted in this minisymposium.

Philip Carns
Argonne National Laboratory
carns@mcs.anl.gov

MS87

PESO: Partnering for Scientific Software Ecosystem Stewardship Opportunities

This presentation will introduce the newly established PESO project, which supports software-ecosystem stewardship and advancement, collaborating through the Consortium for the Advancement of Scientific Software (CASS). PESOs vision is that investments by the U.S. Department of Energy (DOE) in software have maximum impact through a sustainable scientific software ecosystem consisting of high-quality libraries and tools that deliver the latest high-performance algorithms and capabilities to serve application needs at DOE and beyond. Key PESO goals are (1) enabling applications to leverage robust, curated scientific libraries and tools, especially in pursuit of improvement in high-end capabilities and energy efficiency by leveraging accelerator (GPU) devices, and (2) emphasizing software product quality, the continued fostering of software product communities, and the delivery of products, while advancing workforce inclusivity and sustainable career paths. PESO delivers and supports software products via Spack and E4S, and PESO provides porting and testing platforms leveraged across product teams to ensure code stability and portability. PESO also facilitates the delivery of other products, such as AI/ML libraries, as needed by the HPC community. PESO collaborates with CASS to transform independently developed products into a portfolio whose total is much more than the sum of its partsestablishing a trusted software ecosystem essential to DOEs mission.

Michael A. Heroux
Sandia National Laboratories
St. John's University
maherou@sandia.gov

Lois C. McInnes
Argonne National Laboratory
curfman@anl.gov

MS87

Towards Software Sustainability: Center for Open-Source Research Software Stewardship and Advancement

The Center for Open-Source Research Software Stewardship and Advancement (CORSA) is a new community of practice that aims to address the long-term sustainability of scientific and research software by fostering collaboration among stakeholders, facilitating partnerships with open-source and not-for-profit software foundations, developing tools for measuring sustainability-related aspects of software projects, and educating the community regarding approaches to the stewardship and advancement of open-source software. Recent events, such as the completion of the 7-year DOE Exascale Computing Project, have highlighted the need for better strategies to ensure the long-term stewardship of critically important scientific software. Many factors impact the sustainability of research software projects, so gaining a better understanding of, and being able to measure these over the long term can provide valuable information when developing a sustainability strategy. The stewardship of this software, in turn, requires the provision of resources and structures that ensure that there is a healthy vibrant community that is able to undertake the necessary work. In this presentation, we will provide an

overview of the activities that CORSA has been undertaking, our achievements during the first year of the project, and upcoming and future plans that the community can get involved in.

Gregory Watson
Oak Ridge National Laboratory
watsongr@ornl.gov

MS88

Scaling Gaussian Processes to Model Stationary Spatial Data

Gaussian processes are ubiquitous as the primary tool for modeling spatial data. However, the Gaussian process is limited by its $O(n^3)$ cost, making direct parameter fitting algorithms infeasible for the scale of modern data collection initiatives. The Nearest Neighbor Gaussian Process (NNGP) was proposed as a scalable approximation to dense Gaussian processes which has been successful for $\sim 10^6$ observations. We introduce the clustered Nearest Neighbor Gaussian Process (cNNGP) which further reduces the computational and storage cost of a stationary NNGP. Our simulations demonstrate significant reductions in the computational time for model fitting while preserving model fit quality. Furthermore, relative to the full NNGP, the reduction in computational requirements improves as the data size grows. The proposed methods were implemented to obtain an uncertainty-equipped, wall-to-wall map for predicted biomass using biomass estimates generated by NASAs Global Ecosystem Dynamics Investigation (GEDI).

Ashlynn Crisp
Portland State University
acrisp@pdx.edu

Andrew Finley
Departments of Forestry and Geography
Michigan State University
finleya@msu.edu

Daniel Taylor-Rodriguez
Portland State University
dantayrod@pdx.edu

MS88

Understanding Microscopic Behavior of Battery Electrolytes Through Molecular Dynamics Simulations

Ionic liquids (ILs) are prime candidates for a broad range of applications, including spacecraft propulsion propellants, water purification media, and potential replacements to traditional, organic solvent battery electrolytes. Unlike molecular solvents, ILs have large liquid cohesive energies due to strong Coulomb interactions between cation/anion pairs, giving rise to their unique thermophysical and transport properties. In the case of battery applications, the fundamental bulk phase and electrochemical, interfacial behavior of ILs has significant implications on energy storage capacity and safety considerations. In this work, we examine how a novel, quaternary ammonium-based IL, methyltrioctylammonium bis(trifluoromethylsulfonyl)imide or [N1888][TFSI], behaves macroscopically as a bulk liquid as well as how it rearranges near a gold, Au(111), electrode surface. We showcase the power of statistical mechanics and advanced computational chemistry methods in interpreting macroscopic properties at the application level via microscopic studies. Agreement between simulations and experiment indicate important theoretical findings regarding the long-range correlations and electrostatic behavior for [N1888][TFSI].

Shehan Parmar
Georgia Tech University
parmar@gatech.edu

Jesse McDaniel
Georgia Institute of Technology
School of Chemistry and Biochemistry
jesse.mcdaniel@chemistry.gatech.edu

MS88

Trans-dimensional Bayesian Model Selection and Parameter Estimation for Characterization of Local Nuclear Spin Environments in NV Centers

Detailed knowledge of the local environments of defects in semiconductors, such as nitrogen-vacancy (NV) centers in diamond, is essential for optimizing control and entanglement protocols in sensing and quantum information applications. However, direct experimental characterization is not scalable, since characterizing an individual sample can take weeks. In this work, we address the ill-posed inverse problem of recovering the atomic positions and hyperfine couplings of random isotopic nuclei surrounding an NV center using relatively sparse experimental coherence signals, which can be obtained in hours for a single sample. A significant challenge is the unknown number of nuclei in each sample, requiring us to determine both the number of nuclear hyperfine parameters and their values. To address this challenge, we employ a trans-dimensional Bayesian approach using Reverse Jump Markov Chain Monte Carlo (RJMCMC) combined with parallel tempering. This approach simultaneously determines the number of nuclei and provides a full posterior distribution of their atomic positions and hyperfine couplings. These posterior distributions enable us to select the promising NV samples for further characterization. Our approach paves the way for high-throughput characterization of the local environment of NV centers, which is crucial for advancing widespread use of defects in semiconductors centers in sensing and quantum information processing applications.

Abigail N. Poteshman
University of Chicago
poteshman@uchicago.edu

Giulia Galli
University of Chicago and ANL
gagalli@uchicago.edu

MS88

A Case Study on Assessing Credibility in Computational Models for Medical Devices

The credibility of computational modeling and simulation (CMS) is critical for advancing its role in medical device development and regulatory decision-making. However, challenges persist due to model uncertainties, complex device behavior, and non-standardized credibility frameworks. This case study focuses on assessing the credibility of a CMS model for an electronic drug delivery system, where a conjugate heat transfer process between heated coils and a fluid was simulated. The model employs finite volume methods to solve the Navier-Stokes and energy equations, simulating fluid flow and heat transfer dynam-

ics. Using the US Food and Drug Administration guidance and recognized standards we designed and implemented a credibility plan that includes a thorough verification study using a Grid Convergence Index and observed order of accuracy, yielding errors on the order of 1e-5. To address validation under small sample conditions, we employed a probabilistic multi-metric approach, incorporating the area metric, confidence interval, and tolerance interval, implemented via a Python script with continuous integration testing. This analysis highlighted that the largest credibility errors arise from validation uncertainties. While proprietary constraints on the commercial solver limit full reproducibility of the model itself, we ensured transparency in the verification, validation, and uncertainty quantification processes through a publicly accessible GitHub repository. This work contributes to establishing robust credibility practices for CMS in medical device applications, ultimately supporting more informed and reliable regulatory decisions.

Paulina Rodriguez
DOE CSGF
The George Washington University
pxrodriguez@gwu.edu

Lorena A. Barba
Department of Mechanical and Aerospace Engineering
George Washington University
labarba@gwu.edu

MS88

Machine-Learning-Based Spectral Methods Toward Reduced-Order Modeling for Partial Differential Equations

Many multiscale physical systems contain too many degrees of freedom to simulate accurately given limited computational resources. Reduced-order modeling techniques reduce the prohibitively large system to a computationally feasible size without sacrificing essential dynamical features. Model reduction which involves coarsening a representation using standard basis functions, e.g. Fourier functions, is well developed. The applicability and effectiveness of spectral methods depend crucially on the choice of basis functions used to expand the solution of a partial differential equation. Deep learning is a strong contender in providing efficient representations of complex functions [Meuris et al, Sci. Rep. 13, 1739, 2023]. Deep neural networks (DNNs) have shown potential in learning continuous operators or complex systems from streams of scattered data. The deep operator network (DeepONet) [Lu et al, Nat. Mach. Intell 3, 2021] consists of a DNN for encoding the discrete input function space (branch net) and another DNN for encoding the domain of the output functions (trunk net). Physics-informed DeepONets [Wang et al, Sci. Adv. 7, 40, 2021] leverage automatic differentiation to impose the underlying physical laws during model training. In this work, we employ physics-informed machine-learning extracted basis functions from DeepONets which are custom-made for the particular system, with the goal of reduced-order modeling with spectral methods for partial differential equations.

Emily J. Williams
Massachusetts Institute of Technology
ejwllms@mit.edu

Amanda Howard, Panos Stinis
Pacific Northwest National Laboratory
amanda.howard@pnnl.gov, panos.stinis@pnnl.gov

MS89

Runge-Kutta Discontinuous Galerkin Methods with Compact Stencils for Hyperbolic Conservation Laws

In this talk, we develop a new type of Runge-Kutta (RK) discontinuous Galerkin (DG) methods for solving hyperbolic conservation laws. Compared with the standard RKDG methods, the new methods feature improved compactness and allow simple boundary treatment. The convergence to weak solution and the accuracy of the numerical solutions are studied. Their connections with the Lax-Wendroff DG schemes and the ADER DG schemes are also investigated. Numerical examples are given to confirm that the new RKDG schemes are as accurate as standard RKDG methods, while being more compact and cost-effective, for a wide range of problems including two-dimensional Euler systems of compressible gas dynamics.

Qifan Chen
The Ohio State University, Math department
chen.11010@osu.edu

Yulong Xing
The Ohio State University
xing.205@osu.edu

Sun Zheng
University of Alabama
zsun30@ua.edu

MS89

Structure-Preserving Finite Element Methods for Ideal Relaxation Problem

The Parker problem remains an open question since it was first proposed in 1972. It states that given a magnetic configuration, the static equilibrium magnetic field will be tangentially discontinuous. This relaxation process can be described by ideal magnetohydrodynamics equations. For such a time-dependent problem, one of the challenges lies in numerical methods which can preserve the topology of the magnetic field, called helicity, which acts as a topological barrier for energy decay. In this talk, we will first outline the computational approaches to investigate the Parker problem, which includes finite element structure-preserving discretisation, initial condition truncation and magnetic potential computation for helicity. Numerical results will be shown. This is a joint work with Patrick Farrell and Kaibo Hu.

Mingdong He
University of Oxford
mingdong.he@queens.ox.ac.uk

MS89

Software and Algorithms for High-Order Implicit Time Stepping

Although high-level software packages have greatly simplified the application of efficient finite element techniques to problems in flow (and beyond), far less attention has been applied to time stepping. Here, we describe the Irksome package, a high-level time stepping library for the Firedrake project. Users describe the semi-discrete variational form of their partial differential equation in an extension of

the Unified Form Language (UFL). Irksome transforms this semi-discrete form into the fully discrete variational problem for a Runge-Kutta method at each time step. Irksome supports a wide range of RK methods, including explicit, diagonally implicit, and even fully implicit methods. Irksome inherits Firedrake's rich interface to PETSc, so we can also deploy novel preconditioners that make fully implicit methods competitive with and even superior to more widely-used schemes.

Robert C. Kirby
Baylor University
robert_kirby@baylor.edu

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

MS89

A Uniform Framework for Fluid Dynamics in Free-Flow and Porous Media

This talk presents a uniform framework for computational fluid dynamics in porous media based on finite element velocity and pressure spaces with minimal degrees of freedom. The velocity space consists of linear Lagrange polynomials enriched by a discontinuous, piecewise linear, and mean-zero vector function per element, while piecewise constant functions approximate the pressure. Since the fluid model in porous media can be seen as a combination of the Stokes and Darcy equations, different conformities of finite element spaces are required depending on viscous parameters, making it challenging to develop a robust numerical solver uniformly performing for all viscous parameters. Therefore, we propose a pressure-robust method by utilizing a velocity reconstruction operator and replacing the velocity functions with a reconstructed velocity. The robust method leads to error estimates independent of a pressure term and shows uniform performance for all viscous parameters, preserving minimal degrees of freedom. We prove well-posedness and error estimates for the robust method while comparing it with a standard method requiring an impractical mesh condition. We finally confirm theoretical results through numerical experiments with two- and three-dimensional examples and compare the methods' performance to support the need for our robust method.

Seulip Lee
Tufts University
USA
seulip.lee@tufts.edu

Lin Mu
University of Georgia
linmu@uga.edu

MS89

One-Way Spatial Marching of the Navier-Stokes Equations

The One-Way Navier-Stokes (OWNS) equations are based on a one-way spatial integration method for linear hyperbolic equations with slowly-varying coefficients in the direction of integration. Linear OWNS has been demonstrated for stability analysis of high- and low-speed boundary layer flows, where the Navier-Stokes equations are linearized about a time-invariant baseflow, Q, and time-varying disturbances, q', to Q are evolved in the marching direction. More recently, the OWNS approach has been extended to support nonlinear effects, yielding the ability to accurately evolve the early stages of laminar-turbulent transition. Although recent work has focused on hydrodynamic stability analysis, this one-way formulation could be applied to other inhomogeneous wave propagation problems, such as those that arise in underwater acoustics, medical imaging, and seismology. Global frequency-domain analysis eliminates the transients associated with time-marching, while the spatial-marching yields a reduced computational cost relative to the global method since the equations can be solved sequentially rather than simultaneously in the marching direction. We will demonstrate the linear and nonlinear OWNS methodology for boundary-layer flows, jets, and mixing layers. In particular, we will show that enables rapid prediction of laminar-turbulent transition for boundary-layer flows.

Michael K. Sleeman
California Institute of Technology
michael.k.sleeman@gmail.com

Omar Kamal
California Institute of Technology
Department of Mechanical and Civil Engineering
okamal@caltech.edu

Matthew Lakebrink
Boeing Research & Technology
matthew.t.lakebrink@boeing.com

Tim Colonius
Division of Engineering and Applied Science
California Institute of Technology
colonius@caltech.edu

MS90

Thornado: a Toolkit for Relativistic Neutrino-radiation Hydrodynamics

The toolkit for high-order neutrino-radiation hydrodynamics (thornado) is developed as an open source library of modules for solving the equations of neutrino-radiation hydrodynamics for relativistic astrophysics applications, such as core-collapse supernovae, with tabulated microphysics provided by WeakLib. Its development has been focused on specialized solvers based on discontinuous Galerkin methods and node-level performance and portability on heterogeneous computing systems, focusing primarily on kernels for updating solution representations organized in logically Cartesian data structures. As such, thornado can be incorporated in large-scale frameworks with relative ease, as exemplified by its instantiation in the multi-physics software Flash-X. Separate from its collection of native modules, thornado maintains separate division for large-scale simulation by leveraging AMReX for adaptive mesh refinement and distributed parallelism. This talk will provide a brief overview of thornado's current capabilities for relativistic neutrino-radiation hydrodynamics, emphasizing recent progress on robust methods for spectral neutrino transport.

Eirik Endeve
Oak Ridge National Laboratory
University of Tennessee, Knoxville
endevee@ornl.gov

Samuel J. Dunham
California Institute of Technology

dunhamsj@caltech.edu

Austin Harris
Oak Ridge National Laboratory
harrisja@ornl.gov

Jessie Buffaloe
University of Tennessee, Knoxville
jbuffal1@vols.utk.edu

Paul Laiu
Oak Ridge National Laboratory
laiump@ornl.gov

Bronson Messer
University of Tennessee, Knoxville
bronson@ornl.gov

Vassilios Mewes
Oak Ridge National Laboratory
mewesv@ornl.gov

Anthony Mezzacappa
University of Tennessee
mezz@utk.edu

Nick Roberts
University of Tennessee, Knoxville
jrober50@vols.utk.edu

MS90

Exascale Core-Collapse Supernova Simulations with Flash-X+thornado

Neutrino-matter coupling via weak interactions is one of the most important physical mechanisms in the evolution of core-collapse supernovae (CCSN). The numerical modeling of these systems in an inherently multi-physics, multi-method and multi-scale problem. In this talk, we will describe the union of three codes to simulate CCSN: Flash-X+thornado+WeakLib, which evolve the fluid and gravity and provide the computational infrastructure (Flash-X), a code that evolves neutrino radiation hydrodynamics in a spectral two-moment model using the discontinuous Galerkin method (thornado), and a library that provides the equation of state and weak interaction opacity tables (WeakLib). We will describe the two-moment spectral neutrino transport in thornado, interfacing fluid data in different representations (finite volume and discontinuous Galerkin), as well as our performance portability strategy using OpenMP/OpenACC offloading to harness heterogeneous exascale machines such as Frontier at the OLCF. We will present results and performance data from fully adaptive mesh refinement CCSN simulations using Flash-X+thornado+WeakLib.

Vassilios Mewes
Oak Ridge National Laboratory
mewesv@ornl.gov

Eirik Endeve
Oak Ridge National Laboratory
University of Tennessee, Knoxville
endevee@ornl.gov

Austin Harris
Oak Ridge National Laboratory
harrisja@ornl.gov

Steven Fromm, Sean Couch
Michigan State University
frommste@msu.edu, scouch@msu.edu

William Raphael Hix
Oak Ridge National Laboratory
raph@ornl.gov

Bronson Messer
University of Tennessee, Knoxville
bronson@ornl.gov

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

MS90

Kharma: Flexible Accretion Simulations at Exascale

KHARMA (an acronym for "Kokkos-based High-Accuracy Relativistic Magnetohydrodynamics with Adaptive mesh refinement") is a new open-source performance-portable code for conducting general-relativistic magnetohydrodynamic (GRMHD) simulations, such as black hole accretion systems, at scale. By heavily leveraging the Parthenon adaptive mesh refinement framework and Kokkos programming model, KHARMA has been written to be not only fast but easily extensible, modular, and readable. Nearly all functionality is split into packages, each representing an algorithmic component or physics extension. This modularity has made it straightforward to adapt the code to tackle the unique stability and accuracy concerns at scale on OLCF Frontier, as well as to add important extra physics relevant for our problems of interest. I will briefly cover the requirements which led to developing KHARMA, its structure as a Parthenon code, and some of the unique methods and other features we've implemented so far.

Benjamin Prather
LANL
bprather@lanl.gov

MS90

Arepo-RT: Moving Mesh Radiation Hydrodynamics with GPU Acceleration

Radiative transfer (RT) is essential for modeling many astrophysical phenomena, but its integration into radiation-hydrodynamics (RHD) simulations is computationally intensive due to the stringent time-stepping and high dimensionality requirements. The emergence of exascale supercomputers, equipped with extensive CPU cores and GPU accelerators, offers new avenues for optimizing these simulations. This talk will outline our progress in adapting Arepo-RT for exascale environments. Key advancements include a new node-to-node communication strategy utilizing shared memory, which significantly reduces intra-node communication overhead by leveraging direct memory access. By consolidating inter-node messages, we increase network bandwidth utilization, improving performance on both large-scale and smaller-scale systems. Additionally, transitioning RT calculations to GPUs has led to a speedup of approximately 15 times for standard benchmarks. As a case study, cosmological RHD simulations of the Epoch of Reionization demonstrate a threefold improvement in efficiency without requiring modifications to the core Arepo codebase. These developments have broad implications for

the scalability and efficiency of future astrophysical simulations, offering a framework for porting similar simulation codes based on unstructured resolution elements to GPU-centric architectures.

Aaron Smith
University of Texas at Dallas
asmith@utdallas.edu

Oliver Zier
Harvard University
ozier@mit.edu

Rahul Kannan
York University
kannanr@yorku.ca

Mark Vogelsberger
MIT
mvogelsb@mit.edu

MS91

Bi-fidelity Design Under Uncertainty Using Neural Networks

With the ubiquitousness of uncertainty in every engineering system, the robust design of these systems requires addressing the presence of any uncertainty. However, in design under uncertainty, the estimation of response and/or gradients of the response with respect to the design parameters is required for multiple realizations of the uncertainty at every iteration. This may lead to a significant computational cost for the design of large and complex structures. To alleviate this computational burden, in this talk, a computationally efficient bi-fidelity approach is proposed, where a simplified low-fidelity model is used along with neural networks to describe the behavior of any nonlinear or design components that represent the complex high-fidelity structures behavior together. Once trained, these neural networks provide an inexpensive method to get the structure's response and its gradients with respect to the design parameters for many realizations of the uncertain parameters, reducing the optimization cost. As a numerical example, the design of the base-isolation layer of an 11-story building is used to show the efficacy of the proposed approach.

Subhayan De
Department of Mechanical Engineering
Northern Arizona University
SUBHAYAN.DE@nau.edu

Patrick Brewick
Civil and Environmental Engineering and Earth Sciences
University of Notre Dame
pbrewick@nd.edu

MS91

Spectral Expansions based Surrogate for Operator Learning and Sensitivity Analysis: Application to flows in Thermochemical Nonequilibrium

This work presents a surrogate modeling framework combining Karhunen-Love Expansion (KLE) and Polynomial Chaos Expansion (PCE) for two applications: (i) operator learning of coarse-grained chemistry in thermochemical nonequilibrium flows and (ii) multifidelity Global Sensitivity Analysis (GSA) of key metrics like maximum heat flux with respect to chemical kinetic parameters in hypersonic flows. By accounting for model errors from simplified physics, the framework delivers robust and reliable probabilistic predictions. For operator learning, KLE is used to capture time dynamics, isolating temporal modes, while PCE maps model errors and input parameters to KLE coefficients. This approach enhances stability when coupled with fluid solvers, fully incorporates parametric and model uncertainty, and facilitates analytical sensitivity analysis through PCE. For GSA, KLE is used to reduce spatial dimensions, and Bayesian PCE is employed to manage data sparsity from high-fidelity models. The operator learning method is validated for the O2-O chemistry system under hypersonic conditions in a 0D reactor and 1D normal shock, demonstrating stable integration, probabilistic consistency, and errors below 10%. The multifidelity GSA framework is validated in a 1D shock scenario and applied to analyze heat flux sensitivity for a spherical reentry vehicle, showing high accuracy even with limited high-fidelity data.

Mridula Kuppa
University of Illinois at Urbana-Champaign
vkuppa2@illinois.edu

Roger Ghanem
University of Southern California
Civil and Environmental Engineering
ghanem@usc.edu

Marco Panesi
University of Illinois at Urbana-Champaign
mpanesi@illinois.edu

MS91

An Application of the Multi-Index Stochastic Collocation Method for the Uncertainty Quantification of a Naval Engineering Problem

The use of Multi-Index Stochastic Collocation (MISC) for forward Uncertainty Quantification (UQ) of PDEs with uncertain parameters is discussed. MISC is the multi-fidelity counterpart of the Sparse Grids Stochastic Collocation method, and can be used to compute a surrogate model for the solution of a PDE with respect to its uncertain parameters, given a hierarchy of solvers. The UQ aims at assessing the performance of the 5415 DTMB model (namely, its resistance to advancement) subject to uncertain draft and advancing speed (i.e., to operational uncertainties). These vary according to a truncated Gaussian and a triangular PDF, respectively. The 5415 model performance is evaluated by means of numerical simulations performed with an in-house linear potential flow solver. The use of MISC with two and three fidelities is investigated and a comparison against the single-fidelity counterpart will be provided.

Riccardo Pellegrini
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
riccardo.pellegrini@cnr.it

Lorenzo Tamellini
CNR-IMATI
tamellini@imati.cnr.it

Andrea Serani
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
andrea.serani@cnr.it

Chiara Piazzola
Technische Universität München
chiara.piazzola@tum.de

Matteo Diez
Consiglio Nazionale delle Ricerche
Istituto di Ingegneria del Mare (CNR-INM)
matteo.diez@cnr.it

MS91

Uncertainty Quantification in Scientific Machine Learning

Neural networks (NNs) are currently changing the computational paradigm on how to combine data with mathematical laws in physics and engineering in a profound way, tackling challenging inverse and ill-posed problems not solvable with traditional methods. However, quantifying errors and uncertainties in NN-based inference is more complicated than in traditional methods. This is because in addition to aleatoric uncertainty associated with noisy data, there is also uncertainty due to limited data, but also due to NN hyperparameters, overparametrization, optimization and sampling errors as well as model misspecification. Although there are some recent works on uncertainty quantification (UQ) in NNs, there is no systematic investigation of suitable methods towards quantifying the total uncertainty effectively and efficiently even for function approximation, and there is even less work on solving partial differential equations and learning operator mappings between infinite-dimensional function spaces using NNs. In this work, we present a comprehensive framework that includes uncertainty modeling, new and existing solution methods, as well as evaluation metrics and post-hoc improvement approaches. To demonstrate the applicability and reliability of our framework, we present an extensive comparative study in which various methods are tested on prototype problems, including problems with mixed input-output data, and stochastic problems in high dimensions.

Zongren Zou
Brown University
zongren_zou@brown.edu

MS92

Distributed Computing for Physics-based Data-driven Reduced Modeling at Scale

High-performance computing (HPC) has revolutionized our ability to perform detailed simulations of real-world processes. A prominent example is from aerospace propulsion, where HPC is used for rotating detonation rocket engine (RDRE) simulations in support of the design of next-generation rocket engines; however, these simulations take millions of core hours even on powerful computers, which makes them impractical for engineering tasks like design exploration and risk assessment. Reduced-order models (ROMs) address this limitation by constructing computationally cheap yet sufficiently accurate approximations that serve as surrogates for the high-fidelity model. In this presentation, we discuss a new distributed algorithm that achieves fast and scalable construction of predictive physics-based ROMs trained from sparse datasets of extremely large state dimension. The algorithm learns structured physics-based ROMs that approximate the dynamical systems underlying those datasets. We demonstrate our algorithm's scalability using up to $2,048$ cores on the Frontera supercomputer at the Texas Advanced Computing Center. We focus on a real-world three-dimensional RDRE. Using a training dataset of $2,536$ snapshots each of state dimension 76 million, our distributed algorithm enables the construction of a predictive ROM in just 13 seconds on $2,048$ cores on Frontera. Distribution Statement A: Approved for Public Release; Distribution is Unlimited. PA# AFRL-2024-1411

Ionut-Gabriel Farcas
Virginia Tech
farcasi@vt.edu

Rayomand Gundevia
Jacobs Engineering Group, Inc.
Edwards Air Force Base
rayomand.gundevia.ctr@afrl.af.mil

Ramakanth Munipalli
Edwards Air Force Base
Air Force Research Laboratory
ramakanth.munipalli@us.af.mil

Karen E. Willcox
University of Texas at Austin, U.S.
kwillcox@utexas.edu

MS92

Efficient Langevin Sampling for Bayesian Classification on Graphs

This work addresses the task of semi-supervised learning for images classification in safety-critical application. Images are represented by a set of nodes in a high-dimensional space, pairwiseconnected by edges, with weights depending on their similarity. Unlabelled images are then injected to the graph, and we aim at correctly classifying the images using graphical learning through the design of an operator such as the Laplace operator. The formulation of this problem into a Bayesian inference one allows to quantify uncertainty and we investigate the use of Langevin dynamics to sample from the label distribution. We also assess how an approximation of the Laplace operator and the terms in the stochastic Langevin dynamics impacts the uncertainty and which sampling methods can be used to create highly efficient Markov chain Monte Carlo algorithms.

Alix Leroy, Desmond Higham
University of Edinburgh
a.leroy@sms.ed.ac.uk, d.j.higham@ed.ac.uk

Jonas Latz
Heriot-Watt University
jonas.latz@manchester.ac.uk

MS92

Analysis of Matched Source Waveform Inversion to Overcome Cycle Skipping in Seismic Inversion

Seismic inverse problems are notoriously difficult to solve because they require repeated solution of the forward problem, in this case the wave equation, as part of an iterative optimization loop. Because of the expense of this pde solve, only local gradient-based optimization methods are viable. Hence, successfully locating the geologically-informative minimizer requires a very accurate initial guess of the parameters to be estimated, for example, wave velocity. Standard least-squares objective function minimization (known as "full waveform inversion" or "FWI") is well-known to stagnate at local minimizers far from the

true solution. To overcome this "cycle skipping" problem, a penalty term is often added to the FWI objective function. "Extension" methods relax physical restrictions on the solution, for example by not requiring the source to be compactly supported or introducing a filter that loosens the bond between predicted and observed data. This filter maps the observed and predicted data trace-by-trace, penalizing deviation of the filter from the Dirac delta function. This method is known as matched source waveform inversion or MSWI. In this talk we consider transmission data and show numerically that the for single arrival transmitted data, MSWI overcomes cycle skipping because it is equivalent to travel-time inversion. Thus we no longer require a good initial guess for the acoustic wave velocity. However, the method fails for multi-arrival data.

Susan Minkoff
Department of Mathematical Sciences
University of Texas at Dallas
sminkoff@utdallas.edu

William Symes
Rice University
symes@caam.rice.edu

Huiyi Chen
Department of Mathematical Sciences
University of Texas at Dallas
hxc180000@utdallas.edu

MS92

Extropy and Entropy As Measures of Information and Uncertainty with Related Applications

This work is devoted to discussing various monotone properties and characterization results associated with the extropy measure as a measure of uncertainty for ranked data. A comparative analysis is conducted between the extropy of the ranked data and its counterpart in simple random sampling. To further contribute to the field, a consistent estimator for the residual extropy of RSS is proposed. The effectiveness of this estimator is demonstrated through an illustrative example, highlighting its performance in practical scenarios.

M. Z. Raqab
Kuwait University
m.raqab@ku.edu.kw

MS92

Mixed-Precision and Gauss-Newton Algorithms for Training Deep Neural Networks

We present ongoing research on enhancing the efficiency of deep neural network training through mixed-precision computation and modified Gauss-Newton algorithms. By adjusting the precision of various computations and parameters, we achieve reduced model sizes and lower computational and communication costs. To minimize the expense of Gauss-Newton approximations, we utilize advanced automatic differentiation techniques. We demonstrate our results across diverse learning tasks, including physics-informed neural networks, supervised learning, and generative modeling via normalizing flows.

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Samy Wu Fung
Colorado School of Mines
swufung@mines.edu

Deepanshu Verma
Clemson University
dverma@clemson.edu

Elizabeth Newman
Emory University
elizabeth.newman@emory.edu

MS93

Parallel Higher Order Fractional Time Step Methods for Embedded Boundary Methods

We introduce a novel approach to constructing high-order time-split methods for the numerical solution of time-dependent PDEs. Time-split methods, also known as fractional step methods, offer an attractive approach by splitting the original problem into simpler sub-problems, ideally easier to solve. However, the computational cost and complexity associated with achieving high-order accuracy have been a significant challenge. We address this challenge by proposing a parallel deferred correction framework. This framework allows for the parallel evaluation of splitting stages, enabling the generation of a high-order splitting method while leveraging parallel computing resources. The effectiveness of our method is demonstrated through numerical experiments, showcasing its ability to achieve high-order accuracy with improved computational efficiency.

Rochishnu Chowdhury, Nate Overton-Katz
Lawrence Berkeley National Laboratory
rochishnu00@lbl.gov, novertonkatz@lbl.gov

Alexander Novotny
Virginia Tech
anovotny@vt.edu

Benjamin W. Ong
Michigan Technological University
ongbw@mtu.edu

Hans Johansen
Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

MS93

Higher-Order Surface Representations for Embedded Boundary Methods

We demonstrate a high-order embedded boundary algorithm that features adaptive refinement of the boundary surface representation, in conjunction with mesh refinement. This algorithm manages both irregular structures and multiple levels of refinement efficiently using brick-based data layouts. This provides both a compact approach for data communication, and portability to GPUs. High-order surface representations are achieved using an application of the divergence theorem with quadrature rules over a surface triangulation. This strategy supports localized refinement of the boundary representation. Our algorithm is validated for both parabolic and hyperbolic PDEs, and is demonstrated on domains characterized by complex geometries. We compare approaches for resolving

the geometries of interest, and the resulting effects these boundary representations have on the underlying physics.

Nate Overton-Katz
Lawrence Berkeley National Laboratory
novertonkatz@lbl.gov

Stephen Guzik
Colorado State University
stephen.guzik@chord-cfd.com

Hans Johansen
Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

MS93

A High-Order Multiresolution Immersed Interface Method for Elliptic Problems in Complex 3D Domains

We develop a high-order immersed interface method for elliptic problems in complex 3D domains. The Laplacian operator is discretized using a high-order immersed finite difference scheme on a structured cartesian grid. In this scheme, we reconstruct finite difference ghost values across the interface or boundary using a high-order least-squares interpolation of field values, together with boundary conditions on the complex geometry. The elliptic solver relies on a low-order multigrid based preconditioner coupled to a high-order iterative Krylov solver on the finest resolution. Our approach leads to fourth order convergence in the infinity norm for Poisson equations with both Dirichlet and Neumann boundary conditions. We implemented the method in our 3D wavelet-based high order adaptive grid solver Murphy. In this talk, we outline the algorithm and implementation, show convergence results, and analyze the iterative solver and parallel performance for complex geometries on multiresolution grids.

Wim van Rees
Department of Mechanical Engineering
Massachusetts Institute of Technology
wvanrees@mit.edu

MS94

An artificial viscosity approach to high order entropy stable discontinuous Galerkin methods

Entropy stable discontinuous Galerkin (DG) methods improve the robustness of high order DG simulations of nonlinear conservation laws. These methods yield a semi-discrete entropy inequality, and rely on an algebraic flux differencing formulation which involves both summation-by-parts (SBP) discretization matrices and entropy conservative two-point finite volume fluxes. However, explicit expressions for such two-point finite volume fluxes may not be available for all systems, or may be computationally expensive to compute. We propose an alternative approach to constructing entropy stable DG methods using an artificial viscosity coefficient based on the local violation of a cell entropy inequality a local entropy dissipation estimate. The resulting method yields the same global semi-discrete entropy inequality satisfied by entropy stable flux differencing DG methods. The artificial viscosity coefficients are parameter-free and locally computable over each cell. The resulting artificial viscosity preserves high order accuracy, improves linear stability, and does not result in a more restrictive maximum stable time-step size under explicit time-stepping.

Jesse Chan
Rice University
jesse.chan@rice.edu

MS94

High-order Summation-by-parts Operators Over Point Clouds

The first half of this talk will summarize our recently proposed algorithm for constructing high-order summation-by-parts (SBP) operators on point clouds. The construction algorithm assembles global SBP operators from local operators that are defined on cells of a background mesh. We refer to these cell-local operators as degenerate SBP operators, because their diagonal mass matrix may not be positive definite, in general. We enforce positive-definiteness of the global mass matrix by solving a linear optimization problem. There is considerable flexibility in how the cell-local operators are defined; thus, the second half of the talk will present two options for the cell-local operators, and numerical experiments will explore the impact these choices have on the global SBP operators.

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

MS94

Space-Time Entropy-Stable Shock-Tracking Summation-by-Parts Discretizations of the Compressible Euler Equations

High-order methods can provide efficient and accurate simulations, however, they offten lack robustness, especially in the presence of shocks, leading to spurious oscillations and instabilities that can cause numerical solvers to crash. To overcome these challenges, several stabilization techniques have been developed. However, many of these methods degrade to first-order accuracy near shocks and increase computational costs. Recently, Persson and Zahr introduced a method to align mesh interfaces with shock surfaces. In this approach, the conservation law is discretized without knowledge of the shock's location. Then, an optimization problem adjusts the mesh to align with the shock, while finding the corresponding solution. This approach is advantageous as it can be applied to any discretization allowing inter-element discontinuities, ensuring that high-order accuracy is retained even for shocked problems. In this work, we employ the SBP framework to develop high-order, provably-stable schemes for problems with shocks, with a particular focus on space-time curvilinear SBP operators for the Burgers' equation. A combination of two-point flux and average flux is utilized to demonstrate the stability of the scheme. We present a series of numerical results that confirm the schemes effectiveness in handling shock problems.

Dongze Li
University of Waterloo
Department of Applied Mathematics
d336li@uwaterloo.ca

MS94

Entropy Stable Discontinuous Galerkin Methods

Based on Artificial Viscosities with Multiple Parameters

Entropy stable discontinuous Galerkin (DG) methods display improved robustness for problems with shocks, turbulence, and under-resolved features by enforcing an entropy inequality. Such methods rely on entropy conservative (EC) fluxes that are computationally expensive to evaluate. An alternative approach for enforcing an entropy inequality is through a minimally dissipative artificially viscosity. We review how to construct such an artificial viscosity and extend this approach to artificial viscosities with multiple parameters (e.g., viscosity and thermal diffusivity). Through the method of Lagrange multipliers, we determine simple analytical expressions for optimal viscosity parameters. We compare this to the case of a single monolithic viscosity parameter, and compare the behavior of the artificial viscosity using different split form fluxes for various 1D and 2D problems.

Raymond Park
Rice University
Department of Computational and Applied Math
raymondpark01@gmail.com

MS94

Entropy-Stable Real Gas Dynamics

A framework for numerical evaluation of entropy-conservative volume fluxes in multi-species gas flows with internal energies is developed, for use with high-order discretization methods. The novelty of the approach lies in the ability to use arbitrary expressions for the internal degrees of freedom of the constituent gas species. The developed approach is implemented in an open-source discontinuous Galerkin code for solving hyperbolic equations. Numerical simulations are carried out for several model high-enthalpy 2-D flows with non-equilibrium chemistry and the results are compared to those obtained with the finite volume-based solver DLR TAU.

Georgii Oblapenko
ACoM, RWTH Aachen
oblapenko@acom.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS95

An Information Field Theory Framework for Model-Form Error Detection

We derive properties of information field theory (IFT) as applied to inverse problems. The results here can be extended to methodologies which can be seen as limiting cases of IFT, such as Gaussian process regression and physics-informed machine learning. We first define the concept of a well-posed inverse problem within the context of IFT, and pose a few useful theorems for conditions in which an inverse problem becomes well-posed. Using the Gaussian random field interpretation of IFT, we show how identifying parameters of a covariance kernel becomes a well-posed inverse problem under certain conditions. An expression for the Hessian of the inverse problem log posterior is derived to construct the results. A specific focus is placed on the inverse problem of detecting model-form error. We provide an example where the physics are assumed to be the Poisson equation and prove conditions for which identifying model-form error in this case becomes a well-posed inverse problem under IFT.

Alexander Alberts, Ilias Bilionis
Purdue University
albert31@purdue.edu, ibilion@purdue.edu

MS95

Robust Bayesian Optimal Experimental Design under Model Misspecification

Bayesian optimal experimental design (BOED) is increasingly used to improve uncertainty quantification by optimizing the way scientists and engineers gather data. However, as with all methods, understanding the impact of assumptions and model discrepancy is critical. In fact, this is particularly important for BOED because it guides the way data is gathered potentially causing BOED not to sample data that could falsify assumptions or sampling data whose errors are poorly quantified due to low-fidelity models. In this talk we discuss a new information criterion, Expected Generalized Information Gain (EGIG), for BOED problems that incorporate model discrepancy. This criterion is intended to augment traditional BOED based on Expected Information Gain by defining a trade off in robustness vs performance of the experimental design. EGIG measures how poorly inference using an incorrect model is expected to perform compared to an appropriate model for the experiment. We will discuss both the theoretical aspects of this new approach to OED and scalable algorithms incorporating it into BOED. Finally, we will describe how we use this approach to analyze and optimize experiments in materials calibration and geophysics.

Tommie Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS95

Learning Latent Space Dynamics with Model-Form Uncertainties

In this talk, we present a stochastic reduced-order modeling approach to represent and quantify model-form uncertainties in the reduced-order modeling of complex systems using operator inference techniques. The proposed method captures these uncertainties by expanding the approximation space through the randomization of the projection matrix. The relevance of the method is demonstrated on canonical problems in fluid mechanics.

Jin Yi Yong
Duke University
jinyi.yong@duke.edu

Rudy Geelen
Texas A&M University
rgeelen@tamu.edu

Johann Guilleminot
Department of Civil and Environmental Engineering
Duke University
johann.guilleminot@duke.edu

MS95

Validation of Displacement Damage Models

As the third pillar of science, computational simulation

has allowed scientists to explore, observe, and test physical regimes previously thought to be unattainable. High-fidelity models are derived from physical principles and calibrated to experimental data. However, missing or unknown physics and measurement, experimental, and numerical errors give rise to uncertainties in the model form and parameter values in even the most trustworthy models. Thus, rigorous calibration and validation of a computational model is paramount to its effective us as a predictive tool. The popularity of the Bayesian paradigm stems from its natural integration of measurement and model uncertainties. A systematic approach to model validation, progressing from parameter and quantity of interest identification to sensitivity analysis, calibration, and validation, is applied to a drift-diffusion simulation code called Charon. Charon allows the computational qualification of semiconductor devices subjected to displacement damage. *Sandia National Laboratories is a multimission laboratory managed and operated by National Technology Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energys National Nuclear Security Administration under contract DE-NA0003525.

Kathryn Maupin
Sandia National Laboratories
kmaupin@sandia.gov

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Larry Musson, Suzey Gao
Sandia National Laboratories
lcmusso@sandia.gov, xngao@sandia.gov

MS96

Boundary Integral Methods for Flexural Wave Scattering

The waves which propagate in thin elastic plates are known as flexural waves; among other applications, flexural waves are important in the study of the behavior and break up of sea ice. The standard time-harmonic model for these waves involves a fourth order differential operator related to the biharmonic operator. It is thus natural to consider boundary integral reformulations of these equations but existing integral representations do not lead to well-conditioned integral equations. Here, we will present integral representations for the clamped plate, free plate, and supported plate boundary conditions which result in well-conditioned integral equations. We will present some numerical results using these representations and discuss the process used to derive them. Some of the representations employ a novel regularization strategy which may be of broader interest.

Travis Askham
New Jersey Institute of Technology
askham@njit.edu

Peter Nekrasov
University of Chicago, U.S.
pn3@uchicago.edu

Zhaosen Su
University of Utah
u1530865@utah.edu

Jeremy Hoskins
University of Chicago
jeremyhoskins@uchicago.edu

MS96

Uniform H-Matrix Compression with Applications to Boundary Integral Equations

Boundary integral equations lead to dense system matrices when discretized, yet they are data-sparse. Using the $\mathcal{H}$-matrix format, this sparsity is exploited to achieve $\mathcal{O}(N \log N)$ complexity for storage and multiplication by a vector. This is achieved purely algebraically, based on low-rank approximations of subblocks, and hence the format is also applicable to a wider range of problems. The $\mathcal{H}^2$-matrix format improves the complexity to $\mathcal{O}(N)$ by introducing a recursive structure onto subblocks on multiple levels. However, in practice this comes with a large proportionality constant, making the $\mathcal{H}^2$-matrix format advantageous mostly for large problems. In this paper we investigate the usefulness of a matrix format that lies in between these two: Uniform $\mathcal{H}$-matrices. An algebraic compression algorithm is introduced to transform a regular $\mathcal{H}$-matrix into a uniform $\mathcal{H}$-matrix, which maintains the asymptotic complexity. Using examples of the BEM formulation of the Helmholtz equation, we show that this scheme lowers the storage requirement and execution time of the matrix-vector product without significantly impacting the construction time.

Kobe Bruyninckx
KU Leuven
Department Computer Science, Division NUMA
kobe.bruyninckx@kuleuven.be

Daan Huybrechs
Department of Computer Science
KU Leuven
daan.huybrechs@kuleuven.be

Karl Meerbergen
KU Leuven
karl.meerbergen@kuleuven.be

MS96

The Method of Fundamental Solutions for Waveguide Scattering Problems

Waveguide couplers are an important building block of many photonic devices. Their simulation is thus an important part of the design process. Their simulation is also challenging, however, because photonic waveguides are leaky and so couple through an unbounded region. In this talk, we describe a numerical method for simulating these couplers. Our method is based on the observation that individual waveguides are simple enough that their fundamental solutions can be computed numerically. These fundamental solutions can then be used to transform the problem into an integral equation on the interface connecting the waveguides. The resulting integral equation formulation is well-posed and can be solved using coordinate complexification. We present numerical simulations of a variety of couplers and show how similar techniques can be applied to a number of other scattering problems.

Charles L. Epstein
Department of Mathematics, University of Pennsylvania
cepstein@flatironinstitute.org

Tristan Goodwill, Jeremy Hoskins

University of Chicago
tgoodwill@uchicago.edu, jeremyhoskins@uchicago.edu

Shidong Jiang
Flatiron Institute, Simons Foundation
Center for Computational Mathematics
sjiang@flatironinstitute.org

Manas N. Rachh
Center for Computational Mathematics
Flatiron Institute
mrachh@flatironinstitute.org

MS96

Error Estimation Based Adaptive Quadrature for Layer Potentials

Accurate numerical evaluation of layer potentials near domain boundaries necessitates specialized quadrature techniques to handle the rapid variations in the integrand. We present the singularity swap surface quadrature (S3Q) method, a recently developed adaptive quadrature approach tailored for complex geometries featuring multiple axisymmetric bodies. The method features automatic parameter selection based on error estimation to reach a preset error tolerance. It combines a trapezoidal rule for the azimuthal angle with a Gauss-Legendre quadrature on an adaptively subdivided grid in the polar angle, utilizing a so-called interpolatory semi-analytical quadrature in conjunction with a singularity swap technique as needed. Building on its applications to the Laplace and Stokes equations, we extend the S3Q method to the Helmholtz equation by introducing a kernel splitting technique, enabling error-controlled evaluation of oscillatory integrals associated with wave propagation. We also explore enhancements to the S3Q methods computational efficiency.

David Krantz
KTH Royal Institute of Technology
davkra@kth.se

Anna-Karin Tornberg
KTH
akto@kth.se

MS96

Some Scaling Limits for Wave Propagation in Random Media

Inherently rich in multi-scale features, wave propagation in random media is a rather interesting and complex phenomenon. While a complete understanding of wave fields propagating in reasonably arbitrary random media remains essentially out of reach, much progress has been made in the setting of paraxial beam propagation. The paraxial approximation aims at considering high-frequency waves propagating over long distances along a privileged direction with negligible backscattering. Here, I will discuss the so-called white noise paraxial scaling, where different asymptotic regimes lead to very different statistical limits for the wave field. In particular, I will focus on a diffusive scaling under which an arbitrary beam profile transitions to a complex Gaussian distribution with the intensity following an exponential law. This provides a mathematical justification for the complex Gaussian conjecture widely acknowledged in the physical literature and used to characterize speckle patterns observed in experiments.

Anjali Nair
University of Chicago
anjalinair@uchicago.edu

Guillaume Bal
University of Chicago
Departments of Statistics and Mathematics
guillaumebal@uchicago.edu

MS97

Towards Digital Twins for Power-to-X: Comparing Surrogate Models for Methanation Reactors

A digital twin (DT) is a digital model of an intended or actual real-world physical product, system, or process (physical twin) that serves as the effectively indistinguishable digital counterpart. Reduced-order modeling (ROM) can be viewed as an enabling tool for the construction of reliable DTs. A key feature of some ROM approaches is the data-driven aspect, which alleviates the need to explicitly access the (complex) model. In this work, we address the challenge of ROM for dynamical systems in the field of process engineering, by employing scientific machine learning (SciML). Our test case is driven by carbon dioxide (CO2) methanation i.e., the conversion of CO2 to methane, which facilitates the recycling of CO2 emissions and enables green carbon processing. The methods covered here are operator inference (OpInf), a non-intrusive data-driven method for learning dynamical systems from data, and SINDy, which enables the discovery of governing equations from data by sparse identification of nonlinear systems. By harnessing the latest developments, stable quadratic surrogate models are inferred. Based on a complex model of the reactor, we apply the proposed methods for real-time collected data and investigate the behavior of the surrogates for different operating conditions, together with their predictive capabilities. We compare results to standard ML methods (based on Graph Neural Networks) and present conclusions towards implementing a reliable DT infrastructure.

Luisa Peterson
Max Planck Institute, Magdeburg, Germany
peterson@mpi-magdeburg.mpg.de

Ali Forootani
Helmholtz Centre for Environmental Research - UFZ
ali.forootani@ufz.de

Edgar Ivan Sanchez Medina
Max Planck Institute, Magdeburg, Germany
sanchez@mpi-magdeburg.mpg.de

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

Kai Sundmacher
Max Planck Institute for Dynamics of Complex Technical Sys.
sundmacher@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany

benner@mpi-magdeburg.mpg.de

MS97

Data-Efficient Kernel Methods for Pde Discovery

For many problems in computational sciences and engineering, observational data exists for which the underlying physical models are not known. PDE discovery methods provide systematic ways to infer these physical models directly from data. We introduce a framework for identifying and solving PDEs using kernel methods. In particular, given observations of PDE solutions and source terms, we employ a kernel-based data-driven approach to learn the functional form of the underlying equation. We prove convergence guarantees and a priori error estimates for our methodology. Through numerical experiments, we demonstrate that our approach is particularly competitive in the data-poor regime where few observations are available.

Yasamin Jalalian
California Institute of Technology
yjalalia@caltech.edu

MS97

Theory and Applications of Approximate Online State-Parameter Estimation

Online inference, which aims to estimate the states and parameters of a dynamical system as it evolves, is critical yet challenging across various fields, and is critical for enabling next-generation digital twins to reason about streaming data. Existing methods often lack robustness or scalability, and little work has explored non-asymptotic accuracy for Bayesian online state-parameter estimation. Our work addresses this gap in two ways. First, we conduct a general accuracy analysis from a Bayesian perspective, quantifying estimation errors by measuring the distance between the approximate posterior and the ground truth. We present theorems that establish a strong type of local Lipschitz continuity for the temporal propagation of posteriors and provide upper bounds on errors accumulated over time. We also outline sufficient conditions for automatic error reduction over time, which can guide the design of more effective online estimation algorithms. Secondly, we introduce an efficient and scalable online variational inference framework for state-parameter estimation. Our method, free from assumptions about the structure of joint posterior distribution for states and parameters, offers accurate posterior approximations in a flexible way. We demonstrate its effectiveness in applications to both low- and high-dimensional, partially observable dynamical systems.

Liliang Wang, Alex Gorodetsky
University of Michigan
liliangw@umich.edu, goroda@umich.edu

MS98

Train Small, Model Big: A Novel Data-Driven Scaleup for Large-Scale Physics Simulations

Traditional finite element methods (FEM) are widely used for large-scale physics problems. However, refining these models significantly increases problem size and computational cost, often reaching exa-scale and causing numerical instability due to CFL limits and mesh distortion. These issues create a bottleneck, hindering important missions. We introduce a data-driven finite element method (DD-FEM) for multi-physics problems, embodying a train small, model big paradigm. DD-FEM learns from small-scale data to enable exa-scale modeling. Recent developments have shown up to 1000x speed-up and scale-up in lattice-type structure design and steady NavierStokes equations. DD-FEM innovates by using a data-driven basis instead of a traditional polynomial basis. Data generation occurs at the element level, allowing extensive data creation. This data constructs linear or nonlinear manifold bases, surpassing traditional bases and enabling larger time steps, accelerating simulations. DD-FEM introduces data-driven refinement (d-refinement), enhancing accuracy faster than traditional methods. It also accommodates arbitrary element shapes, increasing flexibility. This ensures DD-FEM's readiness for exa-scale, time-dependent, nonlinear multi-phase, and multi-physics problems.

Seung Whan Chung
Lawrence Livermore National Laboratory
chung28@llnl.gov

Alejandro N. Diaz
Rice University
andiaz@sandia.gov

Ivan Zanardi
University of Illinois at Urbana-Champaign
zanardi3@illinois.edu

Sean McBane
University of Texas at Austin
sean@ices.utexas.edu

Axel Larsson
Princeton University
a.larsson@princeton.edu

Youngsoo Choi
Lawrence Livermore National Laboratory
choi15@llnl.gov

MS98

Dynamic Mode Decomposition Variants and Extensions for Robust Data-Driven Modeling

The dynamic mode decomposition (DMD) is a data-driven reduced-order modeling technique that decomposes time-varying data sets into coherent spatiotemporal modes. This representation of the data allows for a large variety of tasks, including dimensionality reduction, future-state prediction, and system control, which has led to the methods utilization and success across many scientific disciplines. In this talk, we discuss several major developments regarding DMD and its applications. In particular, we explore two recent methodological extensions of the DMD algorithm: (1) sparse-mode DMD for generating spatially local DMD modes, and (2) mrCOSTS for decomposing multiscale data sets. We demonstrate key insights provided by these different approaches by analyzing synthetic and real-world systems, including sea surface temperature patterns and optical waveguides. In addition to reviewing new DMD variants, we also discuss the use of time-delay embeddings in conjunction with DMD models. Recent findings suggest that time-delays can be used to not only apply DMD to partial measurement data, but also to identify chaotic behavior even in the absence of measurements that embed the true underlying system. We thus explore this phenomenon and discuss how it might inform our understanding of Takens-embeddable coordinates and time-delay

DMD models.

Sara M. Ichinaga
Department of Applied Mathematics
University of Washington, Seattle
sarami7@uw.edu

Karl Lapo
University of Innsbruck
karl-eric.lapo@uibk.ac.at

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Steven Brunton, Aleksandr Aravkin
University of Washington
sbrunton@uw.edu, saravkin@uw.edu

MS98

Better Training Through Structure and Second-Order Information

Deep neural networks (DNNs) have achieved inarguable success as high-dimensional function approximators in countless applications, including numerous scientific applications such as surrogate modeling, operator learning, and model discovery. However, this success comes at a significant hidden cost, notably a long training time. In this talk, we will to make training easier by exploiting commonly-used network structures and incorporating second-order information. Specifically, we will use Gauss-Newton-like techniques and will demonstrate the computational advantages of our approach over traditional stochastic optimizers on several benchmark deep learning tasks.

Deepanshu Verma
Clemson University
dverma@clemson.edu

Elizabeth Newman
Emory University
elizabeth.newman@emory.edu

Samy Wu Fung
Colorado School of Mines
swufung@mines.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

MS98

On the Advancement of Physics-Informed Machine Learning Via Optimization

We present recent efforts advancing the predictive power of physics-informed machine learning through robust training methods. Firstly, we present an effective training method for multivariate neural networks with smooth activation, namely, σ-Active Neuron Least Squares (σ-ANLS). σ-ANLS was developed from a general iterative training framework whose core mechanism is single-neuron optimization (SNO). In particular, it explores the optimum in a pool of vectors constructed by locality and non-locality. Secondly, we present a two-step training method for deep operator networks, which is constructed by two subnetworks: the branch and trunk networks. Typically, the two sub-networks are trained simultaneously, which often fail. The proposed two-step training method trains the trunk network first and then sequentially trains the branch network. The core mechanism is motivated by the divide-and-conquer paradigm and is the decomposition of the entire complex training task into two subtasks with reduced complexity. Extensive computational examples demonstrate the effectiveness of the proposed training methods compared with existing first- and second-order optimization algorithms on various learning tasks ranging from function approximation, dynamical systems, and solving PDEs to operator learning.

Yeonjong Shin
North Carolina State University
yeonjong_shin@ncsu.edu

MS99

Orbit: Oak Ridge Base Foundation Model for Earth System Predictability

Earth system predictability is challenged by the complexity of environmental dynamics and the multitude of variables involved. Current AI foundation models, although advanced by leveraging large and heterogeneous data, are often constrained by their size and data integration, limiting their effectiveness in addressing the full range of Earth system prediction challenges. To overcome these limitations, we introduce the Oak Ridge Base Foundation Model for Earth System Predictability (ORBIT), an advanced vision transformer model that scales up to 113 billion parameters using a novel hybrid tensor-data orthogonal parallelism technique. As the largest model of its kind, ORBIT surpasses the current climate AI foundation model size by a thousandfold. Performance scaling tests conducted on the Frontier supercomputer have demonstrated that ORBIT achieves 684 petaFLOPS to 1.6 exaFLOPS sustained throughput, with scaling efficiency maintained at 41% to 85% across 49,152 AMD GPUs. These breakthroughs establish new advances in AI-driven climate modeling and demonstrate promise to significantly improve the Earth system predictability.

Prasanna Balaprakash
Oak Ridge National Laboratory
pbalapra@ornl.gov

MS99

Improved Subseasonal Forecasting of Extreme Weather Using Machine Learning

The best early warning products for extreme events only run up to 15 days in advance. These predictions are helpful in protecting vulnerable members of the public, but a greater lead time would allow for mobilizing the resources necessary to protect critical energy and defense infrastructure. This leaves an opportunity for improvement in the forecasting of targeted extreme events over the subseasonal (2-8 week) period using machine learning. Our hypothesis is that machine learning can be used to take in intelligently selected precursor features, such as an El Nio index, and predict the occurrence of large-scale temperature extremes over the continental US on subseasonal timescales with greater accuracy then a physics-based weather model. We have had preliminary success in creating a subseasonal forecaster of polar vortex extremes using machine learning methods. We used geometric moments representing the

polar vortex state and 31 scalar precursor variables were calculated on each winter day from MERRA-2 Reanalysis data. Using random forest (RF) we inspected partial dependence plots to identify the precursors with the most predictive strength. Withholding the last year, we trained statistical and machine learning methods on the best precursors to predict the state of the polar vortex 15 days in advance. The best performing algorithm was a RF forecaster that outperformed recently published results but struggled to adequately capture extreme events.

Thomas Ehrmann
SANDIA National Laboratories
tsehrma@sandia.gov

MS99

Aurora: A Foundation Model of the Atmosphere

Deep learning foundation models are revolutionizing many facets of science by leveraging vast amounts of data to learn general-purpose representations that can be adapted to tackle diverse downstream tasks. Foundation models hold the promise to also transform our ability to model our planet and its subsystems by exploiting the vast expanse of Earth system data. Here we introduce Aurora, a large-scale foundation model of the atmosphere trained on over a million hours of diverse weather and climate data. Aurora leverages the strengths of the foundation modelling approach to produce operational forecasts for a wide variety of atmospheric prediction problems, including those with limited training data, heterogeneous variables, and extreme events. In under a minute, Aurora produces 5-day global air pollution predictions and 10-day high-resolution weather forecasts that outperform state-of-the-art classical simulation tools and the best specialized deep learning models. Taken together, these results indicate that foundation models can transform environmental forecasting.

Paris Perdikaris
University of Pennsylvania
pgp@seas.upenn.edu

MS99

A Physics-Enhanced AI Method for Cloud Physics in General Circulation Models

We introduce a novel physics-enhanced deep learning method for resolving cloud physics in general circulation models, also known as hybrid modeling. Our approach addresses key challenges in hybrid modeling and ensures physically accurate predictions as well as stable long-term simulations without loss of precision. The novel deep learning approach significantly enhances the robustness and interpretability of hybrid models (machine learning + physics) for climate prediction, potentially advancing our ability to simulate and predict complex atmospheric processes.

Xin Wang
National University of Singapore
xin.w24@nus.edu.sg

Gianmarco Mengaldo
National University of Singapore
College of Design and Engineering
mpegim@nus.edu.sg

MS100

A Second Kind Formulation for Algebraic Contrained Differential Equations: The Rare Time When Reckless Differentiation Makes Things Better

A key characteristic of incompressible fluids is the disparity in the differential order of the equations. This discrepancy generates numerical difficulties that are often resolved by differentiating the applied force, adding boundary conditions, and/or using different discretization approaches for pressure and velocity. In this work, we introduce a new formulation for the periodic channel, presenting an unexpectedly simple yet effective solution to resolve the mismatch between pressure and velocity in incompressible fluids. By employing this seemingly unconventional approach, we craft a well-conditioned system that eliminates the traditional mismatch problem. Notably, this method's versatility extends beyond incompressible fluids and applies to various differential-algebraic equations, such as inelastic fibers and curl-free fluids.

Vicente J. Gomez Herrera
New York University
vicente.g@nyu.edu

David Stein
Flatiron Institute
Simons Foundation
dstein@flatironinstitute.org

MS100

A Robust Fast Algorithm for the Evaluation of Layer and Volume Potentials

Over the last two decades, several fast, robust, and high-order accurate methods have been developed for solving the Poisson equation in complicated geometry using potential theory. In this approach, rather than discretizing the partial differential equation (PDE) itself, one first evaluates a volume integral to account for the source distribution withinthe domain, followed by solving a boundary integral equation to impose the specified boundary conditions. Here, we present a new fast algorithm which is easy to implement and compatible with virtually any discretization technique, including unstructured domain triangulations, such as those used in standard finite element or finite volume methods. Our approach combines earlier work on potential theory for the heat equation, asymptotic analysis, the nonuniform fast Fourier transform (NUFFT), and the DMK framework. It is insensitive to flaws in the triangulation, permitting not just nonconforming elements, but arbitrary aspect ratio triangles, gaps and various other degeneracies. On a single CPU core, the scheme computes the solution at a rate comparable to that of the FFT in work per gridpoint.

Leslie Greengard
Simons Foundation and Courant Institute
New York University
lgreengard@flatironinstitute.org

Fredrik Fryklund
Flatiron Institute and Courant Institute, NYU
ffryklund@flatironinstitute.org

Shidong Jiang

Flatiron Institute, Simons Foundation
Center for Computational Mathematics
sjiang@flatironinstitute.org

Samuel Potter
Coreform LLC
Courant Institute, New York University
sam@coreform.com

MS100

Quadrature by Matched Asymptotic Expansions (qbmax) for the Evaluation of Layer Potentials with Boundary Layers

Evaluating layer potentials becomes computationally expensive when dealing with rapidly decaying kernels, such as Helmholtz kernels with large imaginary wave numbers. In this work, we propose a high-order method for evaluating layer potentials with such kernels, referred to as Quadrature by Matched Asymptotic Expansions (QBMAX). This method is an extension of the recently developed high-order scheme, Quadrature by Expansion (QBX), which works for singular, weakly singular, or even hyper-singular kernels. The idea behind QBX is to form local expansions around off-surface points and extrapolate layer potentials to points near or on the surface. However, in the case of rapidly decaying kernels, the local expansion converges poorly, requiring high-order expansion and discretization. Instead, QBMAX forms off-surface local expansions scaled by the matched asymptotic expansion of the underlying PDE near the boundary, demonstrating accuracy and efficiency. We include numerical experiments that show QBMAX offers robust results in both 2D and 3D when compared to the QBX method.

Shawn Lin
University of Illinois
chaoqi2@illinois.edu

Xiaoyu Wei
University of Illinois UrbanaChampaign
xywei@illinois.edu

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

MS100

Boundary Integral Methods for Particle Diffusion in Complex Geometries: Shielding, Confinement, and Escape

Many problems in Engineering and Biology necessitate solving the first passage time problem, which addresses questions such as the expected time for a Brownian particle in unbounded space to reach a target. I will present a boundary integral equation method for solving this mean first passage time with complex geometries of absorbing and reflecting bodies. The method applies the Laplace transform to the time-dependent problem, yielding a modified Helmholtz equation which is solved with a boundary integral method. This approach circumvents the limitations of traditional time-stepping methods and effectively handles the long equilibrium timescales associated with diffusion problems in unbounded domains. Returning to the time domain is achieved by applying quadrature along the so-called Talbot contour to calculate the inverse Laplace transform. I will demonstrate the method for various complex geometries formed by disjoint bodies of arbitrary shape on which either uniform Dirichlet or Neumann boundary conditions are applied. The examples include geometries that guide diffusion processes to particular absorbing sites, absorbing sites that are shielded by reflecting bodies, and finding the exits of confining geometries, such as mazes.

Bryan D. Quaife
Scientific Computing
Florida State University
bquaife@fsu.edu

MS101

Identifying Point Sources in Elliptic Problems Using Deep Neural Networks

The inverse problem of recovering point sources represents an important class of applied inverse problems. However, there is still a lack of neural network-based methods for point source identification, mainly due to the inherent solution singularity. In this work, we develop a novel algorithm to identify point sources, utilizing a neural network combined with a singularity enrichment technique. We employ the fundamental solution and neural networks to represent the singular and regular parts, respectively, and then minimize an empirical loss involving the intensities and locations of the unknown point sources, as well as the parameters of the neural network. Moreover, by combining the conditional stability argument of the inverse problem with the generalization error of the empirical loss, we conduct a rigorous error analysis of the algorithm. We demonstrate the effectiveness of the method with several challenging experiments.

Tianhao Hu
The Chinese University of Hong Kong
1155202953@link.cuhk.edu.hk

MS101

Machine Learning for Synthetic Aperture Radar Imaging

We consider the problem of synthetic aperture radar (SAR) imaging. A single transmitter/receiver mounted on a moving platform emits pulses and collects the corresponding scattered field from an area of interest, called the imaging window (IW). We propose two deep learning networks for solving the SAR imaging problem. Assuming labeled data for small point like targets are available we train the two networks and compare their performance. The first network takes as input collected SAR data and generates a reflectivity map that indicates the locations of the scatterers in the IW. We propose an interpretation of this network as a SAR imaging operator plus a deconvolution one. The second network is different in that it takes as input SAR images and using convolutional neural layers it deconvolves the unknown SAR Kernel. Both networks achieve higher resolution than the traditional SAR imaging techniques. This will be demonstrated with numerical simulations in the SAR Gotcha regime.

Chrysoula Tsogka
University of California, Merced
ctsogka@ucmerced.edu

Symeon Papadimitropoulos, Jeffrey Tucker
University of California Merced

symeon@ucmerced.edu, jtucker7@ucmerced.edu

MS102

An Ensemble Score Filter for Tracking High Dimensional Nonlinear Dynamical Systems

We propose an ensemble score filter (EnSF) for solving high-dimensional nonlinear filtering problems with superior accuracy. A major drawback of existing filtering methods is the low accuracy in handling high-dimensional and highly nonlinear problems. EnSF attacks this challenge by exploiting the score-based diffusion model, defined in a pseudo-temporal domain, to characterizing the evolution of the filtering density. EnSF stores the information of the recursively updated filtering density function in the score function, instead of storing the information in a set of finite Monte Carlo samples (used in particle filters and ensemble Kalman filters). Unlike existing diffusion models that train neural networks to approximate the score function, we develop a training-free score estimation that uses a mini-batch-based Monte Carlo estimator to directly approximate the score function at any pseudo-spatial-temporal location, which provides sufficient accuracy in solving high-dimensional nonlinear problems as well as saves a tremendous amount of time spent on training neural networks. High-dimensional Lorenz-96 systems are used to demonstrate the performance of our method. EnSF provides surprising performance, compared with the state-of-the-art Local Ensemble Transform Kalman Filter method, in reliably and efficiently tracking extremely high-dimensional Lorenz systems (up to one million dimensions) with highly nonlinear observation processes.

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

MS102

Data-Driven Long Term Stable Forecasting of Chaotic Systems

Forecasting high-dimensional dynamical systems is a fundamental challenge in various fields, such as geosciences and engineering. Deep Learning has emerged as a promising tool for forecasting complex nonlinear dynamical systems. However, classical techniques used for their training are ineffective for learning chaotic dynamical systems. Secondly the spectral bias of neural networks prohibits the long term predictions to follow the attractor of the underlying dynamics. In this work, we propose a couple of novel techniques for robust learning of such systems. The Multistep Penalty(MP) method introduces penalized discontinuities in the predicted trajectory while training to address the challenges of non-convexity and exploding gradients associated with training networks with underlying chaotic dynamics. We also introduce an energy spectum based loss function that mitigates the spectral bias of the trained networks. These methods are applied to chaotic systems such as the Lorenz system, KuramotoSivashinsky equation, the Kolmogorov flow, ERA5 reanalysis data, and ocean data. We also test them with several architectures like Neural ODE, Fourier Neural Operator, and UNet. We observe that these techniques provide viable performance for chaotic systems, not only for short-term predictions but also for preserving invariant statistics that are hallmarks of physical systems.

Dibyajyoti Chakraborty
Pennsylvania State University
d.chakraborty@psu.edu

Seung Whan Chung
Lawrence Livermore National Laboratory
chung28@llnl.gov

Arvind T. Mohan
Los Alamos National Laboratory
arvindm@lanl.gov

Romit Maulik
Pennsylvania State University
rmm7011@psu.edu

MS102

A Stochastic Precipitating Quasi-Geostrophic Model

Efficient and effective modeling of complex systems, incorporating cloud physics and precipitation, is essential for accurate climate modeling and forecasting. However, simulating these systems is computationally demanding since microphysics has crucial contributions to the dynamics of moisture and precipitation. In this paper, appropriate stochastic models are developed for the phase-transition dynamics of water, focusing on the PQG model as a prototype. By treating the moisture, phase transitions, and latent heat release as integral components of the system, the PQG model constitutes a set of partial differential equations (PDEs) that involve Heaviside nonlinearities due to phase changes of water. Despite systematically characterizing the precipitation physics, expensive iterative algorithms are needed to find a PDE inversion at each numerical integration time step. As a crucial step toward building an effective stochastic model, a computationally efficient Markov jump process is designed to randomly simulate transitions between saturated and unsaturated states that avoids using the expensive iterative solver. The transition rates, which are deterministic, are derived from the physical fields, guaranteeing physical and statistical consistency with nature. Furthermore, to maintain the consistent spatial pattern of precipitation, the stochastic model incorporates an adaptive parameterization that automatically adjusts the transitions based on spatial information.

Nan Chen
University of Wisconsin-Madison
chennan@math.wisc.edu

Changhong Mou
Purdue University
mouc@purdue.edu

Leslie Smith
Mathematics and Engineering Physics
University of Wisconsin, Madison
lsmith@math.wisc.edu

Yeyu Zhang
Shanghai University of Finance and Economics
zhangyeyu@mail.shufe.edu.cn

MS104

High-Dimensional Hamilton-Jacobi-Bellman PDEs

for Global Optimization

This talk introduces a novel approach to global optimization leveraging solutions of Hamilton-Jacobi-Bellman (HJB) equations, with an application to accelerating consensus-based optimization (CBO) algorithms. We begin by reformulating the global optimization problem as an infinite-horizon optimal control problem. The solution to the associated HJB equation provides a value function that approximates the objective function's landscape. By extracting gradient information from this value function, we obtain a control signal that guides the search towards the global optimum without requiring explicit derivatives of the original objective function. We then demonstrate how this HJB-derived control can be integrated into the CBO framework to significantly enhance its performance. The resulting controlled CBO method exhibits faster convergence rates and improved robustness compared to standard CBO, especially in challenging scenarios with limited particles or poor initialization.

Yuyang Huang
Imperial College London
yuyang.huang21@imperial.ac.uk

MS104

Solving the Rubik's Cube Using Its Local Graph Structure

This research proposes a new approach to solving the Rubiks Cube by using two neural networks that focus on the local graph structure of the cubes current state. Building on previous work in the reinforcement learning community such as DeepCubeA, we extend these contributions and optimize cube-solving methods by utilizing the local graph structure rather than the entire state space of the cube. Our first neural network predicts the next optimal move to make. This network aims to improve accuracy and reduce computation time. Our second neural network estimates the distance from the current state to the solved state of the cube. This network is trained on extensive datasets using supervised learning techniques. This approach reduces the computational challenges that arise from the Rubiks Cube's immense state space. Preliminary results show strong promise for the efficacy of this method, exhibiting the potential to advance cube-solving techniques in the reinforcement learning community. Through continued research and validation, we demonstrate the effectiveness of utilizing neural networks that leverage local graph structures for solving complex puzzles.

Mitchy Lee
Trinity University
mitchy-lee@outlook.com

Shunyu Yao
Chinese University of Hong Kong
yao853836045@gmail.com

MS104

Inverse Reinforcement Learning of Optimal Control and Differential Games

One of the most efficient approaches for autonomous agents to master complex skills and decent behaviors efficiently and safely is to imitate the behaviors of humans or expert agents who already optimally perform those tasks. Inverse reinforcement learning (IRL) control stands out as a principled framework that enables agents to imitate the demonstrated behaviors described by optimal control systems and differential games by reconstructing the underlying cost or objective functions while ensuring the stability and convergence of algorithms. Inverse RL offers model-free properties during reconstruction and reduces uncertainties with unknown dynamics, distinguishing it from the classic inverse optimal control (IOC) that shares similar functionality but relies on explicit system dynamics. The presentation will show state-of-the-art data-driven IRL of optimal control systems and differential games, including zero-sum, non-zero-sum, and multiagent graphical games. In particular, data-driven off-policy integral inverse RL algorithms are developed to reconstruct cost functions and control policy.

Bosen Lian
Auburn University
bzl0098@auburn.edu

MS104

A Physics-Informed Learning Framework for Nonlinear Optimal Control Problems

In this talk we will propose a physics-informed neural networks (PINNs) framework to solve the infinite-horizon optimal control problem of nonlinear systems. In particular, since PINNs are generally able to solve a class of partial di??erential equations (PDEs), they can be employed to learn the value function of the infinite-horizon optimal control problem via solving the associated steadystate Hamilton-Jacobi-Bellman (HJB) equation. However, an issue here is that the steady-state HJB equation generally yields multiple solutions; hence if PINNs are directly employed to it, they may end up approximating a solution that is di??erent from the optimal value function of the problem. We tackle this by instead applying PINNs to a finite-horizon variant of the steady-state HJB that has a unique solution, and which uniformly approximates the optimal value function as the horizon increases. An algorithm to verify if the chosen horizon is large enough is also given, as well as a method to extend it – with reduced computations and robustness to approximation errors – in case it is not. Unlike many existing methods, the proposed technique works well with non-polynomial basis functions, does not require prior knowledge of a stabilizing controller, and does not perform iterative policy evaluations.

Kyriakos Vamvoudakis
Georgia Institute of Technology
kyriakos@gatech.edu

MS104

Sciml for Modeling and Control of Dynamic Systems with Applications

Innovation in deep learning methods, tools, and technology offers an unprecedented opportunity to transform the systems engineering practice and bring much excitement to systems theory research. In this talk, I will introduce recent results in modeling dynamic systems with deep learning representations that embed domain knowledge. I will also discuss differentiable predictive control, a data-driven approach that uses physics-informed deep learning representations to synthesize predictive control policies. Ill motivate and illustrate the concepts with examples of engineered energy systems such as floating offshore wind and intelligent buildings.

Draguna L. Vrabie, Bruno Jacob

Pacific Northwest National Laboratory
draguna.vrabie@pnnl.gov, bruno.jacob@pnnl.gov

Natalie Isenberg
Brookhaven National Laboratory
nisenberg@bnl.gov

Jan Drgona
Pacific Northwest National Laboratory
jdrgona1@jh.edu

MS105

Proportional Asymptotics of Bayesian Neural Network Predictions

Bayesian inference could provide the framework for principled uncertainty quantification for neural networks. Barriers to adoption include the interpretability of predictive distributions and the challenge of fully learning multimodal posterior distributions. We demonstrate that under a discretized prior for the inner layer weights, we can exactly characterize the posterior predictive distribution as a mixture of Gaussian modes. This setting allows us to define equivalence classes on realizations of network parameters which produce the same training error and to relate the elements of these classes to the networks scaling regime: the number of training samples, the size of each layer, and the number of final layer parameters. Of particular interest are distinct parameter realizations which map to low training error and correspond to distinct modes in the posterior predictive distribution. We describe the impact of the scaling regime on multimodality and the inevitability of uncertainty when network size grows with sample size.

Katharine Fisher
MIT
kefisher@mit.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

MS105

Physics Informed IB-UQ: Information Bottleneck Based Uncertainty Quantification for Physics Informed Scientific Computing

Neural networks (NNs) are currently changing the computational paradigm on how to combine data with mathematical laws in physics and engineering in a profound way, tackling challenging inverse and ill-posed problems not solvable with traditional methods. However, quantifying errors and uncertainties in NN-based inference is more complicated than in traditional methods. In this talk, we will present Information Bottleneck based uncertainty quantification method for physics informed neural networks and operator learning. The effectiveness of the model will be demonstrated through several representative examples.

Ling Guo
Shanghai Normal University
lguo@shnu.edu.cn

Tao Zhou
Institute of Computational Mathematics
Chinese Academy of Sciences
tzhou@lsec.cc.ac.cn

Hao Wu
ShanghaiJiao Tong university
hwu81@sjtu.edu.cn

MS105

Multifidelity Neural Networks with UQ for Representing Free Energy Surfaces

We consider high-dimensional distributions representing population densities evolving in $\mathbb{R}^n$. Of interest to us is an efficient approach to infer the dynamics in Fokker-Planck form. We adopt a two step process: In the first step we combine Knothe-Rosenblatt rearrangements with globally C^∞ basis functions. In the second step we use Bayesian neural networks to parameterize the unknown n-dimensional Fokker-Planck equations that govern the time-evolution. The learning problem leads to inference of the Fokker-Planck forms with quantified uncertainty. These methods provide advantages over discretization-based inference of partial differential equations that are afflicted by the curse of high-dimensionality.

Saem Han
University of Michigan
saem@umich.edu

Krishna Garikipati
Aerospace & Mechanical Engineering
University of Southern California
garikipa@usc.edu

MS105

Bayesian Optimal Experimental Design for Adaptive Training of Neural Network Surrogate Models

Multi-query tasks such as uncertainty propagation, inference, and optimization all require repeated evaluations of a forward model. When the forward model involves a high-fidelity physics-based simulation, these tasks can become prohibitive. One approach to accelerate the computation is by training a parametric surrogate model, such as a neural network (NN), from high-fidelity simulation data, and using it to replace the original forward model for all tasks. The surrogate is often trained in a deterministic manner without quantifying the uncertainty in the surrogate models parameters or downstream predictions, and with a single batch of training data that is randomly selected, for instance via space-filling designs. We propose the use of Bayesian Optimal Experimental Design (OED) for training a NN, that quantifies and updates surrogate uncertainty and leveraging it to select locations to acquire new training simulations. We make use of a greedy-sequential OED to iterate between the prior-to-posterior update of NN parameters via variational inference, and identifying and acquiring new high-fidelity simulations that would yield the greatest Expected Information Gain (EIG). In particular, we formulate goal-oriented OED approaches where the EIG does not focus on the abstract NN parameters, but on quantities of interest from downstream tasks that will rely on the completed NN surrogate. We compare our framework to existing active data collection methods on numerical examples.

Aniket Jivani, Xun Huan
University of Michigan

ajivani@umich.edu, xhuan@umich.edu

MS105

Probabilistic Solver for Parametric PDEs Via Neural Operator Induced Gaussian Process

The study of neural operators has significantly advanced the development of efficient approaches for solving partial differential equations (PDEs), offering advantages over traditional methods. However, most of the existing neural operators lack the capability to quantify predictive uncertainty, a crucial aspect, especially in data-driven scenarios with the limited availability of data. In this work, we propose a probabilistic solver for solving parametric partial differential equations (PDEs) which utilizes the probabilistic characteristic of Gaussian Processes (GPs) while leveraging the learning prowess of neural operators. The proposed framework leads to improved prediction accuracy and offers a quantifiable measure of predictive uncertainty. The proposed framework is extensively evaluated through experiments on various PDE examples, including Burgers equation, Darcy flow, non-homogeneous Poisson, and wave-advection equations. Furthermore, a comparative study with wavelet neural operator (WNO), its Bayesian variant along with GP (zero mean) is presented to highlight the advantages of our proposed approach. The results demonstrate superior accuracy and expected uncertainty characteristics, suggesting the promising potential of the proposed framework in applications pertaining to mechanics and engineering.

Sawan Kumar
Ph.D. scholar
amz228095@am.iitd.ac.in

Rajdip Nayek
Indian Institute of Technology, Delhi
rajdipn@am.iitd.ac.in

Souvik Chakraborty
Department of Applied Mechanics
Indian Institute of Technology Delhi
souvik@am.iitd.ac.in

MS106

Numerical Investigation of the SRTD Algorithm for the Oldroyd 3-Parameter Model

The SRTD algorithm was recently developed to solve the steady-state case of a certain 3-parameter subset of the Oldroyd models which contains the upper-convected and lower-convected Maxwell models. It is an iterative algorithm, and it is similar to the original UCM algorithm by Renardy, but SRTD decouples the Oldroyd 3-parameter system into a Navier-Stokes like momentum equation and two transport equations. When the decoupled equations are solved exactly, the iteration is provably convergent for a sufficiently nice domain, tangential boundary conditions, and sufficiently small parameters, but no numerical implementation had been tested. In this talk, we present the results of a numerical investigation of a finite element implementation of the SRTD algorithm and discuss future plans of accelerating and extending the SRTD method.

Christian Austin
University of Florida
christianaustin@ufl.edu

MS106

Semi-Lagrangian Characteristic Reconstruction and Projection for Transport under Incompressible Velocity Fields

We present a novel semi-Lagrangian characteristic reconstruction method that leverages a volume preserving projection to advect quantities under incompressible velocity fields. A key advantage of this framework is to see the traditional semi-Lagrangian scheme as the construction of a diffeomorphism between the deformed and original geometry (reference map). This representation allows us to use the local deformation of the geometry to design a projection for the reference map onto the space of volume preserving diffeomorphisms. In the context of the advection of an implicit surface representation (level set method), this results in significant improvements to the interface precision and mass conservation. We demonstrate the performance of our new method with a variety of two- and three-dimensional examples, compare this new approach to traditional schemes, and show the effectiveness of our characteristic bending scheme is particularly well suited for simulations of the incompressible Navier-Stokes equations.

Matthew Blomquist
University of California Merced
mblomquist@ucmerced.edu

Maxime Theillard
University of California, Merced
mtheillard@ucmerced.edu

MS106

Viscoelastic Phase Separation of Polymer-Solvent Systems: Analysis Structure-Preserving Discretisation

In the context of binary phase separation dynamic symmetry of both phases, i.e. similar relaxation time scales, plays a crucial role in modelling. The generally accepted model is an incompressible Cahn-Hilliard-Navier-Stokes system for the evolution of the volume fraction and the velocity. In viscoelastic phase separation, this dynamic symmetry is broken, since polymer chains and solvent particles have effects on completely different time scales. To accommodate these asymmetrical effects, the model is extended by additional equations. On one hand, the viscoelastic effects arising from chain dynamics are modeled using the Peterlin model for the conformation tensor. This renders the flow Non-Newtonian. On the other hand, mixing effects between solvent and chains are captured through a nonlinearly coupled advection-reaction-diffusion equation, augmenting the Cahn-Hilliard equation. In the absence of the conformation tensor, we employ a structure-preserving approximation using conforming finite elements both in space and time. This method is shown to accurately preserve essential thermodynamic quantities, such as conservation of mass and energy dissipation. Theoretical findings are complemented by a convergence test and an illustrative example drawn from practical applications.

Aaron Brunk, Maria Lukacova
Johannes Gutenberg University
Institute of Mathematics

abrunk@uni-mainz.de, lukacova@uni-mainz.de

MS106

Pressure Robust Discretization of Non-Newtonian Fluids

We present a nonconforming Crouziex-Raviart discretization for the nonlinear Stokes equation modelling a non-Newtonian incompressible fluid. Our a priori estimates are pressure robust, i.e. the convergence rate for the velocity does not depend on the regularity of the pressure.

Lars Diening
Bielefeld University
lars.diening@uni-bielefeld.de

Adrian Hirn
Hochschule Esslingen
adrian.hirn@hs-esslingen.de

Christian Kreuzer
TU Dortmund University
christian.kreuzer@tu-dortmund.de

Pietro Zanotti
Università degli Studi di Milano
pietro.zanotti@unipv.it

MS106

A Survey of Recent Results for Discretization of Rheological Models

There are many models for non-Newtonian fluids. In recent years, mathematical theory about such models and their numerical solution has increased substantially. For example, some models due to Rivlin and Ericksen and to Oldroyd have been proved to be equivalent in the limit of small parameters, even though these models initially appear to be quite different. Since this understanding has improved recently, it is reasonable to ask what future research directions should be explored. We will survey the state of the field, describe some recent results, and suggest some open problems to be considered.

Ridgway Scott
University of Chicago
ridg@uchicago.edu

MS107

Transporting Mass with Low Distortion

Traditionally, stochastic mass transport is studied as an optimization problem. The goal is to push forward a given source measure to a target measure in a way that minimizes distances. According to Caffarelli, a prominent result in Euclidean spaces establishes the Lipschitz regularity of these optimal transport maps in some specific cases. Lipschitz continuity allows the transfer of analytic properties of the source measure to an a priori complicated target, and Caffarellis result has proven to be widely influential. In this talk, we will explore extensions of Caffarellis theorem to various settings. Our main observation is that the optimality conditions mentioned above are not necessary for most applications. We shall thus introduce a general construction of transport maps, based on semigroup interpolation, and analyze its Lipschitz constant via stochastic calculus of variations techniques. In particular, we will go beyond the Euclidean setting and consider Riemannian manifolds as well as infinite-dimensional spaces. Interestingly, these types of problems relate to some modern numerical applications, which we shall also discuss.

Dan Mikulincer
Massachusetts Institute of Technology, U.S.
danmiku@gmail.com

MS107

Tree-Wasserstein Distance for Data with Latent Feature Hierarchy

Finding meaningful distances between high-dimensional data samples is an important scientific task. In this work, we present a new tree-Wasserstein (TW) distance, which we term Hyperbolic Diffusion Tree-Wasserstein (HDTW) distance, for high-dimensional data with a latent feature hierarchy. While the conventional use of the TW distance is to speed up the computation of the Wasserstein distance, we use its inherent tree as a means to learn the latent feature hierarchy. The key idea of our method is to embed the features into a continuous multi-scale hyperbolic space using diffusion geometry, and then, present a new tree decoding algorithm by establishing analogies between the hyperbolic embedding and trees. We show that our HDTW distance provably recovers the TW distance for data with a latent feature hierarchy and that its computation is efficient and scalable. We showcase the usefulness of the proposed HDTW distance in applications to word-document and single-cell RNA-sequencing datasets, demonstrating its advantages over competing baselines based on the TW distances and methods based on pre-trained models.

Ya-Wei Lin
Technion
lin.ya-wei@campus.technion.ac.il

Gal Mishne
UC San Diego
gmishne@ucsd.edu

Ronald Coifman
Yale University
Department of Computer Science
coifman@math.yale.edu

Ronen Talmon
Technion
ronen@ee.technion.ac.il

MS107

Efficient Neural Network Approaches for Conditional Optimal Transport with Applications in Bayesian Inference

We present two neural network approaches that approximate the solutions of static and dynamic conditional optimal transport (COT) problems. Both approaches enable conditional sampling and conditional density estimation, which are core tasks in Bayesian inferenceparticularly in the simulation-based ("likelihood-free") setting. Our methods represent the target conditional distributions as transformations of a tractable reference distribution and, therefore, fall into the framework of measure transport. Although many measure transport approaches model the transformation as COT maps, obtaining the map is computationally challenging, even in moderate dimensions. To improve scalability, our numerical algorithms use neural networks to parameterize COT maps and further exploit

the structure of the COT problem. Our static approach approximates the map as the gradient of a partially input-convex neural network. It uses a novel numerical implementation to increase computational efficiency compared to state-of-the-art alternatives. Our dynamic approach approximates the conditional optimal transport via the flow map of a regularized neural ODE; compared to the static approach, it is slower to train but offers more modeling choices and can lead to faster sampling. We demonstrate both algorithms numerically, comparing them with competing state-of-the-art approaches, using benchmark datasets and simulation-based Bayesian inverse problems.

Zheyu Oliver Wang
Massachusetts Institute of Technology
olivrw@mit.edu

Ricardo Baptista
Caltech
rsb@caltech.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Deepanshu Verma
Clemson University
dverma@clemson.edu

MS107

Stochastic Inverse Problem Based on Push-Forward Maps

Inverse problems in physical sciences are typically formulated as PDE-constrained optimization problems, aimed at identifying unknown parameters by optimizing models to match observed data. While this approach assumes deterministic parameters, many real-world problems involve stochastic parameters, requiring the recovery of their full distribution. This talk addresses the challenge of solving such stochastic inverse problems by examining three key aspects: the stability of the inverse problem, the role of regularization in the resulting variational formulation, and the application of the Wasserstein gradient flow as an effective solver for optimizing over probability distributions, providing new insights into this emerging class of inverse problems.

Yunan Yang, Maria Oprea
Cornell University
yunan.yang@cornell.edu, mao237@cornell.edu

Li Wang
University of Minnesota
wang8818@umn.edu

Qin Li
University of Wisconsin-Madison
qinli@math.wisc.edu

MS108

Data Assimilation with Measure Transport and Generative AI

Measure transport and related generative AI methods have shown impressive performance in image and video generation tasks. Their probabilistic nature also serves as an appealing framework for extending standard data assimilation (DA) methods in the geosciences. The Ensemble Kalman Filter (EnKF) is one of the most popular ensemble DA algorithms with demonstrated stability in high-dimensional settings. However, the Gaussian assumptions made in its formulation lead to biases that eventually deteriorate the model predictions. This talk will highlight two recent DA frameworks which utilize measure transport and generative AI tools to relax the Gaussian assumptions in the EnKF. Specifically, the Ensemble Conjugate Transform Filter (ECTF) generalizes the EnKF to arbitrarily complex distributions based on ideas from normalizing flows, while the Ensemble Score Filter (EnSF) samples the Bayesian posterior with training-free diffusion models. Both methods take advantage of the appealing idea that target distributions in the DA problem can be obtained as successive transformations of much simpler base distributions like the Gaussian. Following a brief mathematical introduction, I will describe recent progress in applying these new DA algorithms to problems in the geosciences.

Hristo G. Chipilski
Florida State University, Department of Scientific Computing
hchipilski@fsu.edu

Feng Bao
Florida State University
bao@math.fsu.edu

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

Siming Liang
Florida State University
sliang@fsu.edu

Jeffrey Whitaker
NOAA/Earth System Research Laboratories
Physical Sciences Laboratory
jeffrey.s.whitaker@noaa.gov

Derek J. Posselt
Univ. of Michgan, Ann Arbor
Derek.Posselt@jpl.nasa.gov

MS108

Non-Gaussian Extensions of the Maximum Likelihood Ensemble Filter

With the resolution of numerical weather and ocean prediction models becoming higher resolution then more nonlinear dynamics need to be approximated which leads to non-Gaussian statistics. Most of the traditional data assimilation systems are based upon assuming Gaussian statistics for the errors involved. Over the last 20 years there has been development of lognormal versions of the variational data assimilation systems but not so much towards the ensemble systems. In this presentation we introduce the non-Gaussian versions of the maximum likelihood ensemble filter based upon the lognormal and the reverse lognormal distributions and show that it can perform better than a Gaussian fits also version of he filter with the cou-

pled Lorenz 1963 model.

Steven Fletcher
CIRA-Colorado State University
steven.fletcher@colostate.edu

Senne Van Loon
Colorado State University
senna.van_loom@colostate.edu

Milija Zupanski
Cooperative Institute for Research in the Atmosphere
Colorado State University
Milija.Zupanski@colostate.edu

MS108

Understanding Bias in Ensemble Filters

Ensemble filtering methods are empirically proven to be effective at addressing filtering problems in high-dimensional settings. However, these methods are inherently biased, and a better theoretical understanding of this bias is interesting. In this talk, we will quantify this bias in the mean-field setting. We will focus on two particular transport-based filtering methods: the standard Ensemble Kalman Filter and the stochastic map filter, a nonlinear analog that can effectively capture non-Gaussian posteriors. Each of these methods has a corresponding class M of dynamical models for which they are unbiased. Interestingly, we can quantify, in W2-distance, how the bias of our ensemble method depends on the proximity of the true dynamics to the unbiased class M. This sheds light on the origin of bias in the investigated ensemble filters and how it interacts with the systems dynamics.

Frederic Jørgensen
Massachusetts Institute of Technology, U.S.
fjorgen@mit.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

Ricardo Baptista, Franca Hoffmann
Caltech
rsb@caltech.edu, franca.hoffmann@caltech.edu

MS108

Local Sequential Mcmc for Data Assimilation with Applications to Drifters Data and Surface Water and Ocean Topography-Like Data

In this talk, I present a new data assimilation (DA) scheme based on a sequential Markov Chain Monte Carlo (SMCMC) DA technique that is provably convergent and has been recently used for filtering, particularly for high-dimensional non-linear, and potentially, non-Gaussian state-space models. Unlike particle filters, which is a convergent method as well, SMCMC does not assign weights to the samples/particles, and therefore, the method does not suffer from the well-known issue of weight-degeneracy when a relatively small number of samples is used. We design a time-dependent localization approach within the SMCMC framework that focuses on regions where observations are located and restricts the transition densities included in the conditional distribution of the state to these regions. This results in immensely reducing the effective degrees of freedom and thus improving the efficiency. We test the new technique on high-dimensional ($d \sim 10^4 - 10^5$) linear Gaussian model and non-linear shallow water models with Gaussian noise with real and synthetic observations. We show superiority in terms of efficiency and accuracy over competing ensemble methods and the SMCMC filter. For two of the numerical examples, the observations mimic the Surface Water and Ocean Topography data. We also use a set of ocean drifters' real observations in which the drifters are moving according the ocean kinematics and assumed to have unknown random locations at the time of assimilation.

Hamza Ruzayqat
King Abdullah University of Science and Technology
hamza.ruzayqat@kaust.edu.sa

Omar Knio
King Abdullah University of Science and Technology (KAUST)
omar.knio@kaust.edu.sa

MS109

Augmenting Subspace Optimization Methods with Linear Bandits

In recent years, there has been a growing interest in developing iterative optimization methods that focus on finding descent restricted to affine subspaces containing an incumbent. Each iteration typically consists of choosing a subspace according to some possibly randomized technique, and then building and minimizing a local model employing derivatives of the function projected onto the chosen subspace. In model-based derivative-free optimization, where gradient approximations essentially require a finite difference (i.e., a number of function evaluations linear in problem dimension), these methods suggest serious opportunities for practical gains in efficiency, since the number of function evaluations necessary to obtain a projected gradient approximation instead scales with the chosen subspace dimension. Motivated by practicality, we consider a simple augmentation of such a generic subspace method. In particular, we consider a sequential optimization framework where actions consist of one-dimensional linear subspaces and rewards consist of projected gradient measurements made along corresponding affine subspaces. This sequential optimization problem can be analyzed through the lens of dynamic regret. We modify an existing upper confidence bound (UCB) linear bandit method to achieve sublinear dynamic regret. We demonstrate the efficacy of employing this UCB method alongside a sketched version of the derivative-free optimization method, POUNDers.

Matt Menickelly
Argonne National Laboratory
mmenickelly@anl.gov

MS109

Solving 10,000-Dimensional Optimization Problems with Inaccurate Function Values: An Old Algorithm

We re-introduce a derivative-free subspace optimization framework originating from Chapter 5 the thesis [Z. Zhang, On Derivative-Free Optimization Methods, PhD thesis, Chinese Academy of Sciences, Beijing, 2012] of the author under the supervision of Ya-xiang Yuan. At each iteration, the framework defines a (low-dimensional) subspace based on an approximate gradient, and then solves a subproblem in this subspace to generate a new

iterate. We sketch the global convergence and worst-case complexity analysis of the framework, elaborate on its implementation, and present some numerical results on solving problems with dimension as high as 10,000. The same framework was presented during ICCOPT 2013 in Lisbon under the title "A Derivative-Free Optimization Algorithm with Low-Dimensional Subspace Techniques for Large-Scale Problems", although it remains nearly unknown to the community until very recently. An algorithms following this framework named NEWUOAs was implemented by Zhang in MATLAB in 2011 (https://github.com/newuoas/newuoas), ported to Module-3 by Nystroem (Intel) in 2017, and included in cm3 in 2019 (https://github.com/modula3/cm3/blob/master/caltech-other/newuoa/src/NewUOAs.m3).

Zaikun Zhang
Sun Yat-sen University
zaikunzhang@gmail.com

MS110

Optimizing Gaussian Processes for Large-Scale Applications with Mixed-Precision Computation

Gaussian Processes (GPs) are a fundamental class of machine learning models widely used for regression, uncertainty quantification, and the estimation of statistical model parameters, particularly in geospatial data analysis. In computational materials science, GPs are employed to model the residual error between machine learning interatomic potentials (MLIPs) and reference quantum mechanical (QM) data, with MLIPs acting as surrogate models for QM calculations. Despite their strengths, GPs face significant challenges when applied to large datasets due to their high computational complexity. This complexity arises primarily from the need to construct and invert large covariance matrices, which are both computationally intensive and memory-demanding tasks. This presentation explores a novel optimization of GP regression models using accuracy-preserving mixed-precision computation across various stages of the process. By redesigning the computation of matrix generation through optimized matrix-matrix multiplication operations, and accelerating the solution of linear systems via mixed-precision Cholesky-based solvers, we achieve substantial improvements in both speed and efficiency. This approach demonstrates the potential for significant computational savings while maintaining the accuracy of GP models, making them more practical for large-scale applications.

Rabab Alomairy
King Abdullah University of Science and Technology
rabab.omairy@kaust.edu.sa

MS110

H2O: Orchestrating Hierarchical Algorithms on Heterogeneous Architectures with the Parsec Runtime System for Complex Scientific Applications

Task-based runtime systems have become pillars of the computational revolution, renowned for their dynamic execution models and strategic resource management. These systems are fundamental to the development of increasingly complex and adaptable algorithms, which are vital to advancing computational science. In this talk, we introduce the PaRSEC runtime system, designed to provide a scalable and efficient framework for deploying and managing complex scientific applications. Our discussion aims to showcase PaRSECs capacity to enhance the efficiency of processing algorithms on hierarchical matrix operations, like mixed-precision computations, in heterogeneous computing environments. The results confirm that runtime systems like PaRSEC can significantly boost the processing of large-scale scientific computations, establishing them as essential tools for researchers in high-performance computing domains.

Qinglei Cao
University of Tennessee Knoxville
qinglei.cao@slu.edu

MS110

Leveraging Sparse Linear Algebra for High-Performance Kernel K-means Clustering

The diverse and non-trivial challenges of parallelism in data analytics require computing infrastructures that go beyond the demand of traditional simulation-based sciences. The growing data volume and complexity have outpaced the processing capacity of single-node machines in these areas, making massively parallel systems an indispensable tool. However, programming on high-performance computing (HPC) systems poses significant productivity and scalability challenges. It is important to introduce an abstraction layer that provides programming flexibility and productivity while ensuring high system performance. As we enter the post-Moore's Law era, effective programming of specialized architectures is critical for improved performance in HPC. As large-scale systems become more heterogeneous, their efficient use for new, often irregular and communication-intensive data analysis computation becomes increasingly complex. In this talk, we will present how sparse linear algebra can be used to achieve performance and scalability on extreme-scale systems while maintaining productivity for emerging data-intensive scientific challenges. In particular, I will present our project on GPU-based Kernel K-means, called Popcorn, where we achieve significant speedups over traditional CPU and GPU implementations by formulating Kernel K-means using sparse- dense matrix multiplication (SpMM) and sparse matrix- vector multiplication (SpMV), demonstrating the effectiveness of sparse matrices for efficient parallel programming.

Giulia Guidi
Cornell University
Lawrence Berkeley National Laboratory
gg434@cornell.edu

MS110

Improving Spgemm Performance Through Reordering and Cluster-Wise Computation

Sparse matrix-matrix multiplication (SpGEMM) is a key kernel in many scientific applications and graph workloads. However, it often performs poorly due to irregular memory accesses. A common strategy for optimizing sparse matrix operations is to reorder the matrix with the goal of improving data locality. In this work, we comprehensively study the impact of reordering on SpGEMM performance by applying 10 reordering algorithms.

Helen Xu
Georgia Tech

hxu615@gatech.edu

MS111

Cumper: Cubical Multiparameter Persistence for Topological Machine Learning

We introduce cubical multipersistence, a topological data analysis technique, to extract robust topological signatures from 2D and 3D images. Deep learning excels in learning data representations but struggles with low-data regimes and challenging modalities like MRI, CT scans, and histopathology. To address these challenges, we integrate cubical multipersistence into machine learning to enhance performance on small, intricate datasets. This approach captures valuable topological details often overlooked by traditional methods. Our pipeline uses handcrafted grayscale features and their filtrations, along with a novel learnable 2D multifiltration pipeline powered by a U-Net backbone. We ensure the validity of neural network outputs as multifiltrations through differentiable constraints. From these cubical multiparameter filtrations, we derive silhouette functional summaries of persistence diagrams as topological features, combining them with latent appearance features for downstream tasks like classification. Our model can be trained end-to-end using straight-through estimators, also learning the filtration thresholds. We demonstrate the efficacy of CuMPer on diverse 2D medical image datasets. Our models, whether data-driven or not, outperform the state-of-the-art in both classification and regression tasks. Additionally, our end-to-end model excels in 2D classification, surpassing the state-of-the-art in various benchmarks.

Brighton Nuwagira
UT Dallas
brighton.nuwagira@utdallas.edu

MS111

Smoothness-Increasing Accuracy-Conserving (siac) Filters and Its Corresponding Probability Density Function

Numerical filters have often been investigated in a Fourier context. They are usually designed to alleviate Gibbs phenomena or aliasing error. They can also extract hidden information in the data in order to achieve a more accurate solution. These so-called accuracy extracting filters, such as the Smoothness-Increasing-Accuracy-Conserving (SIAC) filter, rely on patterns of information contained in the given data. The patterns determine the filter scaling and might arise from the numerical discretization or different sources of noise. Filters such as SIAC rely on moment conditions statistical or mechanical to extract extra information. This motivates us to explore the probability density function (PDF) corresponding to SIAC, which allows for estimating a given state over time. In this talk we will discuss the generalised SIAC filter, its corresponding PDF, and the relation to filters based on Geganbauer polynomials. This work was done in collaboration with A. Edoh, J. Glaubitz, Y. Hristova, T. Li, and Nisha.

Jennifer K. Ryan
KTH Royal Institute of Technology
jryan@kth.se

MS111

Ergodic Properties of Generalized Billiards

We consider the generalization of the classical billiard map via different reflection laws. We study several dynamical properties of the resulting dynamical systems with the emphasis on the ergodic properties of the dynamic. Mathematical billiards are a form of dynamical system used to simulate real-world phenomena, from the movements of microorganisms, to the navigation of robot vacuum cleaners, and beyond. In certain cases, it can be beneficial to modify the law of specular reflection. Generalized billiards have been employed to investigate chaotic systems, as well as to develop algorithms that can solve optimization problems with greater efficiency. We consider the billiard transformation $T(s, \varphi)$ on a convex table Γ defined by the following way:

$$\begin{cases} s_{n+1} = s_n + f(s_n, \varphi_n) \\ \varphi_{n+1} = \lambda \hat{\varphi}_n + \alpha. \end{cases} \quad (1)$$

Georgii Sechkin
University of Texas at Arlington
georgii.sechkin@uta.edu

MS111

A Hybrid Siac - Data-Driven Post-Processing Filter for Discontinuities in Solutions to Numerical Pdes

We introduce a hybrid filter that incorporates a mathematically accurate moment-based filter with a data driven filter for discontinuous Galerkin approximations to PDE solutions that contain discontinuities. Numerical solutions of PDEs suffer from an $\mathcal{O}(1)$ error in the neighborhood about discontinuities, especially for shock waves that arise in inviscid compressible flow problems. While post-processing filters, such as the Smoothness-Increasing Accuracy-Conserving (SIAC) filter, can improve the order of error in smooth regions, the $\mathcal{O}(1)$ error in the vicinity of a discontinuity remains. To alleviate this, we combine the SIAC filter with a data-driven filter based on Convolutional Neural Networks. The data-driven filter is specifically focused on improving the errors in discontinuous regions and therefore *only includes top-hat functions in the training dataset.* For both filters, a consistency constraint is enforced, while the SIAC filter additionally satisfies $r-$moments. This hybrid filter approach allows for maintaining the accuracy guaranteed by the theory in smooth regions while the hybrid SIAC-data-driven approach reduces the ℓ_2 and ℓ_∞ errors about discontinuities. Thus, overall the global errors are reduced. We examine the performance of the hybrid filter about discontinuities for the one-dimensional Euler equations for the Lax, Sod, and sine-entropy (Shu-Osher) shock-tube problems.

Soraya Terrab, Samy Wu Fung
Colorado School of Mines
sterrab@mines.edu, swufung@mines.edu

Jennifer K. Ryan
KTH Royal Institute of Technology
jryan@kth.se

MS111

From Pixels to Patterns: Leveraging Topological

Data Analysis in Histopathology

Topological Data Analysis (TDA) has emerged as a sophisticated framework for uncovering the underlying structure of complex datasets. In histopathology imaging, where conventional methods often struggle with subjectivity and variability, TDAprimarily through Persistent Homologyoffers a robust mechanism for extracting and analyzing topological invariants across multiple scales. By integrating TDA with deep learning, we have identified generic topological patterns in tumor tiles across various cancer types, enabling the development of a versatile model that distinguishes between tumor and healthy tissue. This approach leverages topological signatures to enhance the diagnostic precision of histopathological models, providing a novel pathway for tissue morphology characterization and cancer classification. The talk will detail these methodologies and present experimental results that underscore the effectiveness of TDA in creating a generic, scalable diagnostic tool for diverse histopathological applications.

Ankur R. Yadav
UT Dallas
Meta
ankur.yadav@utdallas.edu

MS112

SWAS: Advancing Scientific Workflows and Application Services

The Center for Stewardship and Advancement of Workflows and Application Services (SWAS) is designed to support the growing scientific workflows and application services community by providing a cohesive, sustainable ecosystem. As workflows become central to scientific research, particularly in DOE science for simulations, real-time data analysis, and AI-driven studies, the demand for robust, interoperable, and scalable software solutions has never been greater. SWAS aims to address this need by stewarding critical workflow software, supporting core capabilities, and fostering a collaborative community focused on innovation, best practices, and sustainability. This talk will explore SWASs multi-faceted approach to building a sustainable software ecosystem that integrates diverse software projects, enables seamless interoperability across platforms, and enhances user experience. Attendees will learn how SWAS is engaging stakeholders from academia, industry, and national labs to co-create a resilient software infrastructure that supports the entire lifecycle of workflow development. We will also highlight SWASs strategic collaborations, including partnerships with DOE facilities and other Software Stewardship Organizations, aimed at driving forward the future of scientific computing.

Rafael Ferreira da Silva
Oak Ridge National Laboratory
silvarf@ornl.gov

MS112

STEP: Software Tools Ecosystem Project

The Software Tools Ecosystem Project (STEP) provides the support and enhancement needed for critical High Performance Computing (HPC) software tools to remain effective, efficient, and relevant in the rapidly evolving field of HPC. We focus on the collection of software packages that can be applied to monitor, analyze, and diagnose performance and behavior of computational science applications and systems. STEPs support for these projects emphasizes the ongoing sustainment and enhancement necessary by identifying key tool gaps and addressing needs through novel software strategies and extending leading HPC capabilities. We are pursuing opportunities for co-design with hardware vendors, application developers, facilities and tool developers. This talk will explore STEP's approach to building a sustainable software tools ecosystem. By providing a forum to improve communication among tool developers, application teams, vendors and facilities, STEP seeks to be a clearing-house for efficiently managing critical HPC technology. We invite application teams and interested parties to join with STEP as we engage stakeholders from academia, industry, and national labs to co-create a robust software tools infrastructure that supports the newest machines and the latest programming environments

Terry Jones
Oak Ridge National Laboratory, USA
trjones@ornl.gov

Phil Carns
Argonne National Laboratory, USA
carns@mcs.anl.gov

John Mellor-Crummey
Computer Science Dept.
Rice University.
johnmc@rice.edu

Heike Jagode
University of Tennessee
jagode@icl.utk.edu

Barton Miller
University of Wisconsin, Madison
bart@cs.wisc.edu

Sameer Shende
Dept. of Computer and Information Science
University of Oregon
sameer@cs.uoregon.edu

Shane Snyder
Argonne National Laboratory
ssnyder@mcs.anl.gov

MS112

Software Stewardship Activities in the FASTMath and RAPIDS SciDAC Institute

The goals of the software stewardship activities in the FASTMath and RAPIDS SciDAC Institutes, which we refer to as OASIS (Organization for the Advancement and Stewardship of Integrated Software), are (1) advance and steward a selected set of production software products for use in leadership computing providing key capabilities to DOE science teams, (2) help bring relevant new software products to a production state, and (3) foster the integration of software tools into effective solutions that help achieve DOE mission objectives. OASIS focuses its software efforts on three technology areas that are central to many science endeavors: applied mathematics; data management, storage, and analysis; and artificial intelligence and machine learning. Working alongside the other SSO teams, we will help ensure that key software underpinnings of DOE mission activities are available, performant, and trustworthy. In addition, we will make our capabilities and software products available to the DOE community and to

the computational science community internationally. In this talk, we will review the OASIS software portfolio, the first-year work plan, and the progress tracking strategies. We will also describe our planned stewardship activities, including our approach to stewardship in the AI/ML area, our integration efforts, and our approach to incubating potential new products.

Xiaoye S. Li
Lawrence Berkeley National Laboratory
xsli@lbl.gov

Rob Ross
Argonne National Laboratory
rross@mcs.anl.gov

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

Leonid Oliker, Esmond G. Ng
Lawrence Berkeley National Laboratory
loliker@lbl.gov, egng@lbl.gov

MS112

S4PST: Advancing the Programming Systems Ecosystem for Future HPC

The successful completion of the US Department of Energy (DOE) Exascale Computing Project (ECP) has highlighted the crucial role of software in maximizing the potential of heterogeneous extreme-scale high-performance computing (HPC) systems. To ensure that future HPC systems benefit the broadest possible user base, we must invest in a robust and adaptable software ecosystem. Programming systems, as the primary interface for most application developers and users, are essential. Therefore, advancing the programming systems ecosystem requires careful attention. In this presentation, we will introduce the Stewardship for Programming Systems and Tools (S4PST) project, funded by the DOE Advanced Scientific Computing Researchs (ASCR) Next Generation of Scientific Software and Tools (NGSST) Initiative. This project aims to advance the development and stewardship of ten software products essential for future HPC systems, while also incubating new products targeting the integration of HPC and AI in programming systems (E.g. LLMs, AI, high-productivity languages). We will provide a brief overview of these products, their stewardship teams, key features, recent advancements, and engagement activities.

Keita Teranishi, Pedro Valero Lara, William F. Godoy
Oak Ridge National Laboratory
teranishik@ornl.gov, valerolarap@ornl.gov, godoywf@ornl.gov

MS113

SpECTRE: An Open-source Code Using Spectral Methods for Compact Binary Merger Simulations

Spectral-type methods, such as finite-element methods, are an efficient alternative to classical finite difference methods commonly employed, particularly for smooth solutions. This makes spectral methods well-suited for the inspiral portion of compact object merger simulations, and the ringdown portion of binary black hole mergers. I will provide a brief review of spectral-type methods, both discontinuous Galerkin and pseudospectral, before discussing their advantages and drawbacks when applied to binary compact object merger simulations. In particular, I will provide a computational cost comparison between two implementations, the closed-source Spectral Einstein Code and the open-source SpECTRE code. Finally, I will highlight the particular advantage these methods have for building large catalogs of long, high-accuracy gravitational waveforms from binary mergers.

N Deppe
Cornell University
nd357@cornell.edu

Kyle Nelli
Caltech
knelli@caltech.edu

Geoffrey Lovelace
Cal State Fullerton
glovelace@fullerton.edu

Lawrence Kidder
Cornell University
kidder@astro.cornell.edu

Nils Vu
Caltech
nilsvu@caltech.edu

William Throwe
Cornell University
wtt6@cornell.edu

Mark Scheel
California Institute of Technology
scheel@caltech.edu

Saul Teukolsky
Cornell University
saul@astro.cornell.edu

Harald Pfeiffer
Max Planck Institut
Albert Einstein Institut
harald.pfeiffer@aei.mpg.de

MS113

XNet: Performance Portable Nuclear Reaction Kinetics for Astrophysical Applications

Nuclear burning processes are an integral component for a variety of astrophysical phenomena and are responsible for the synthesis of elements that enrich the universe. We present the nuclear reaction kinetics code `XNet` that evolves the rate equations which describe these processes using general-purpose reaction networks. `XNet` is primarily used in one of two ways: as a stand-alone application for post-processing nucleosynthesis from pre-calculated thermodynamic histories or as a module for coupling nuclear burning processes in a multiphysics simulation framework. Here, we focus on the computational aspects of the latter, as post-processing calculations are often trivially parallel and relatively inexpensive compared to implementations in coupled multiphysics simulations. In this talk, we present the numerical and computational methods for the integration methods and linear solvers available in `XNet` and describe our approach to performance portability across var-

ious CPU and GPU architectures.

Austin Harris
Oak Ridge National Laboratory
harrisja@ornl.gov

MS113

Astrophysical Plasmas on the Exascale: GPU-Accelerated Particle-in-Cell Methods and Parallel Performance in the Entity Code

The Entity code is a state-of-the-art, open-source particle-in-cell (PIC) framework for efficient modeling of magnetized plasmas and compact object magnetospheres on current and next-generation GPU architectures. This talk will showcase Entity's capabilities to scale the integration of the Vlasov-Maxwell equations effectively in different coordinate systems. We will discuss Entitys grid-agnostic algorithmic structure, and its use of the Kokkos performance portability library, which enables flexible performance and efficient device parallelization across CPU and GPU platforms. This talk will outline Entity's primary scientific objectives and ongoing specializations to relevant physics applications. Recent scientific contributions of Entity include simulations of plasma dynamics around highly magnetized neutron stars with self-consistent quantum electrodynamics (QED) feedback. Entitys exascale capabilities enable advanced models of magnetized turbulence with unprecedented levels of scale separation and long-term dynamic consistency. We will present performance benchmarks on the Frontier supercomputer.

Jens Mahlmann
Columbia University
jfm2231@columbia.edu

Hayk Hakobyan
Princeton Plasma Physics Laboratory
haykh.astro@gmail.com

Sasha Philippov
University of Maryland
sashaph@umd.edu

Benjamin Crinquand
University of Toulouse
benjamin.crinquand@univ-tlse3.fr

Alexander Chernoglazov
University of Maryland
achernog@umd.edu

Arno Vanthieghem
Observatoire de Paris
arno.vanthieghem@obspm.fr

MS113

Gram-X: a New GPU-Accelerated Dynamical Spacetime Grmhd Code for Exascale Computing with the Einstein Toolkit

In this talk, I will present a new GPU-accelerated dynamical-spacetime GRMHD code called GRaM-X which I have developed during my PhD. GRaM-X can perform production-level simulations of core-collapse supernovae in full GR on GPUs on world's fastest supercomputers. I will present the implementation, code tests as well as performance results (scaling) of the code. I will then discuss a jet-driven core-collapse supernova simulation that I have performed using GRaM-X on OLCF Summit and OLCF Frontier. I will present the results of the simulation as well as discuss the performance of the code in the simulation.

Swapnil Shankar
University of Tennessee, Knoxville
sshanka5@utk.edu

Philipp Moesta
University of Amsterdam
p.moesta@uva.nl

Steven Brandt
Louisiana State University
sbrandt@cct.lsu.edu

Roland Haas
NCSA
rhaas@illinois.edu

Erik Schnetter
Perimeter Institute
eschnetter@perimeterinstitute.ca

Yannick de Graaf
University of Amsterdam
yannickdgraaf@gmail.com

MS114

Adaptive Pilot Sampling for Multi-Fidelity Stochastic Optimization

Stochastic optimization can be prohibitively expensive for high-fidelity models, generally requiring sampling-based estimation of model statistics within an outer optimization loop. To expedite estimation of these high-fidelity model statistics, multi-fidelity techniques (e.g., approximate control variate (ACV) and multilevel BLUE) leverage cheaper runs of low-fidelity models, often producing orders-of-magnitude computational savings. However, such estimators generally require knowledge of both model costs and covariances to determine hyperparameters such as the control-variate weights and sample-allocation ratios. The model covariances in particular are usually not known a priori and are thus typically estimated via independent and potentially expensive and wasteful pilot sampling. Furthermore, the specific hyperparameters of these multi-fidelity estimators are not designed with an outer design optimization loop in mind. For instance, covariance may vary smoothly across the design domain, so pilot sampling may only be needed at few key designs. We present a framework for multi-fidelity stochastic optimization with ACVs using probabilistic covariance emulation across the design domain. We provide an active learning strategy for efficient pilot sampling to provide improved covariance estimates, and thus improved ACV hyperparameters, at the most important designs towards the outer optimization loop. We compare different variants of this approach on numerical benchmarks.

Thomas Coons, Aniket Jivani, Xun Huan
University of Michigan
tcoons@umich.edu, ajivani@umich.edu, xhuan@umich.edu

MS114

A Bandit-Learning Multi-Fidelity Estimation

Framework: Removing Pilot Studies

Multi-fidelity methods that use an ensemble of models to compute a Monte-Carlo estimator of the expectation of a high-fidelity model often require orders of magnitude less computational resources than estimators constructed using a single model. Such estimators use oracle statistics, specifically the covariance between models, to optimally allocate samples to each model in an ensemble such that the mean-squared-error of the estimator is minimized. However, in practice the oracle statistics are estimated using a small number of model evaluations, and the computational cost of collecting, and the error introduced by using, these finite number of model evaluations is typically ignored. Consequently, this paper generalizes the bandit-learning procedure in ["A bandit-learning approach to multifidelity approximation," Xu et al., 2022] to optimally balance the cost and accuracy of the estimated oracle statistics with the accuracy and cost of the final multi-fidelity estimator. Our numerical results show that our novel algorithm produces estimators with mean-squared errors commensurate with estimators obtained from state-of-the-art methods even when those methods ignore the cost of computing that oracle information.

Yiming Xu
University of Utah
yxu@math.utah.edu

Thomas Dixon, Alex Gorodetsky
University of Michigan
tdixono@umich.edu, goroda@umich.edu

Akil Narayan
University of Utah
akil@sci.utah.edu

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

MS114

UQ for Compute Intensive Simulation: UM-Bridge and Parallelization of Hierarchical Methods

We showcase recent developments around UM-Bridge, a universal framework that enables rapid development of advanced uncertainty quantification (UQ) methods and applications. It enables even prototype UQ codes to control simulations on thousands of processor cores, and makes advanced UQ available to a wide audience. In addition, we present UQ methods capable of handling compute intensive models: A fully parallelized approach to Multilevel Delayed Acceptance, including fully asynchronous prefetching of future model evaluations, and developments towards a Multiindex extension of Delayed Acceptance. Making use of UM-Bridge, we demonstrate how these methods enable UQ in earthquake research and cosmology at supercomputing scale.

Linus Seelinger
Karlsruhe Institute of Technology (KIT)
Scientific Computing Center
linus.seelinger@kit.edu

Maximilian Kruse, Sharana Shivanand
Karlsruhe Institute of Technology
maximilian.kruse@kit.edu, sharana.shivanand@kit.edu

Abdul-Lateef Haji-Ali
School of Mathematical & Computer Sciences
Heriot-Watt University
a.hajiali@hw.ac.uk

Louise Kluge
Karlsruhe Institute of Technology
louise.kluge@kit.edu

Zihua Niu
Ludwig-Maximilians-Universität München
zihua.niu@lmu.de

Michael Bader
Technical University of Munich
bader@in.tum.de

Alice Gabriel
UC San Diego
liese.gabriel@gmail.com

MS115

Unique Reconstruction for Discretized Inverse Problems, a Random Sampling Approach

Inverse problem theory, through the lens of the PDE-constrained optimization, suggests the unique reconstruction of the parameter function. Unfortunately, such well-posedness result only holds when infinite amount of data is provided. As we are limited to finite amount of measurements due to experimental or economical reasons, the recovery of the smooth function cannot be accomplished in practice. Consequently, one must compromise the inference goal to a discrete approximation of the unknown function. Open questions in the finite-dimensional setting are: What is the reconstruction power of a fixed number of data observations? How many parameters can one reconstruct? Here we describe a probabilistic approach, and spell out the interplay of the observation size (r) and the number of parameters to be uniquely identified (m). The technical pillar is the random sketching strategy, in which the matrix concentration inequality and sampling theory are largely employed. By analyzing randomly sub-sampled Hessian matrix, we attain well-conditioned reconstruction problem with high probability. This is joint work with Qin Li, Anjali Nair and Sam Stechmann.

Ruhui Jin
University of Wisconsin
rjin@math.wisc.edu

MS115

Subspace Method of Moments for Cryo-Em

Cryo-electron microscopy (cryo-EM) is an increasingly popular imaging technique that aims to recover the three-dimensional structures of biological molecules from randomly oriented tomographic projection images, taken at extremely low signal-to-noise-ratio. One mathematical approach to this non-linear inverse problem is the method-of-moments, where low-order moments of the projection images are inverted to recover the 3D structure. However, the storage, sample and computational complexities scale exponentially with the moment order. This talk will discuss a method of reducing these complexities by using randomized tensor sketching techniques to produce compressed Tucker decompositions that can be used in the cryo-EM pipeline.

Oscar Mickelin
Princeton University
hm6655@princeton.edu

MS115

On Applications and Faster Methods for Calculating Tensor Decompositions and Tensor Eigenpairs

As the first talk of the mini-symposium, I am going to introduce tensors and how several of their applications and challenges are relevant in the realm of Computational Science and Engineering. I will introduce CP tensor decomposition and tensor eigensystems, and show how these differ from their matrix counterparts. Finally, I will talk about a new computational method to calculate tensor eigenvectors, and how this method leads to faster CP tensor decomposition methods.

Joao M. Pereira
IMPA - Instituto de Matemática Pura e Aplicada
Rio de Janeiro, Brazil
jpereira@uga.edu

MS115

Efficient and Robust Online Tensor Decomposition Through Global Sketching

In this talk we introduce a new method for scalable tensor decomposition in a streaming setting. The key idea is to store the compressed historical data using sketch operators. In particular, we show the method applied to online CP decomposition is much more robust against outliers and more capable of catching trends in the streaming data. Theoretical guarantees of the new method will also be discussed.

Yifan Zhang
University of Texas at Austin
yf.zhang@utexas.edu

Joe Kileel
University of Texas at Austin, U.S.
jkileel@math.utexas.edu

MS116

Highly Scalable Methods for Adaptive Mesh Refinement of Complex Domains

Cut cell methods are increasingly becoming popular for complex and moving geometries, especially in conjunction with quasi-structured adaptive meshes such as octree-based meshes. This ensures a scalable and efficient approach to grid generation and adaptivity for complex domains and multiphysics problems. However, the special treatment required for cut cells and their distribution within the domain presents some challenges in ensuring load balance and, consequently, parallel scalability. Although spectral partitioning approaches such as metis, ptscotch, etc., are able to effectively partition meshes produced by cut cell methods, the complexity and scalability, especially in parallel, is a major bottleneck to the efficient parallelization of cut cell methods. In this work, we present a new local partitioning method based on graph distance, that is able to effectively partition meshes produced by cut cell methods, for both complex geometries as well as moving boundary scenarios. This new approach can minimize partition boundary sizes, comparable to spectral methods while delivering multiple orders of better performance for the partitioning stage on distributed architectures. This approach is well suited for adaptive repartitioning use cases.

Budvin C. Edippuliarachchi, Hari Sundar
Tufts University
budvin.edippuliarachchi@tufts.edu,
hari.sundar@tufts.edu

MS116

High-Order Wavelet-Based Multiresolution Grid Adaptation Around Immersed Interfaces

We propose a novel adaptive resolution scheme for the immersed interface method. Our approach applies multiresolution analysis based on interpolating wavelets to the immersed interface fields, which are potentially discontinuous across interface. Before performing a wavelet transform, we first apply a polynomial extrapolation step near the interface. With this treatment, the detail coefficients near the interface scale as $O(h^N)$ where N is the order of the wavelet, and h is the grid width, so these values are comparable with the detail coefficients in free space. Moreover, with this strategy, our wavelet transform is still shown to be lossless, meaning that no information is lost during coarsening and refinement processes. We corroborate our claims by applying our strategy to a high order immersed interface solver for parabolic PDEs and the Navier-Stokes equation.

Nigel Shen
Massachusetts Institute of Technology
shen827@mit.edu

MS117

Learning Coarse-Grained Dynamics on Graphs

I will discuss a Graph Neural Network (GNN) non-Markovian modeling framework to identify coarse-grained dynamical systems on graphs. We systematically determine the GNN architecture by inspecting how the leading term of the Mori-Zwanzig memory term depends on the coarse-grained interaction coefficients that encode the graph topology. I will present numerical demonstrations on two examples, a heterogeneous Kuramoto oscillator model and a power system, which shows that the proposed GNN architecture can predict the coarse-grained dynamics under fixed and time-varying graph topologies.

Daning Huang, John Harlim
Pennsylvania State University
dnhuang@umich.edu, jharlim@psu.edu

MS117

Consensus-Based Construction of High-Dimensional Free Energy Surface

One essential problem in quantifying the collective behaviors of molecular systems lies in the accurate construction of free energy surfaces (FESs). The main challenges arise from the prevalence of energy barriers and the high dimensionality. Existing approaches are often based on sophisticated enhanced sampling methods to establish efficient exploration of the full-phase space. On the other hand, the collection of optimal sample points for the numerical approximation of FESs remains largely under-explored, where the discretization error could become dominant for

systems with a large number of collective variables (CVs). We propose a consensus sampling-based approach by reformulating the construction as a minimax problem which simultaneously optimizes the function representation and the training set. In particular, the maximization step establishes a stochastic interacting particle system to achieve the adaptive sampling of the max-residue regime by modulating the exploitation of the Laplace approximation of the current loss function and the exploration of the uncharted phase space; the minimization step updates the FES approximation with the new training set. By iteratively solving the minimax problem, the present method essentially achieves an adversarial learning of the FESs with unified tasks for both phase space exploration and posterior error-enhanced sampling. We demonstrate the method by constructing the FESs of molecular systems with several CVs up to 30.

Huan Lei, Liyao Lyu
Michigan State University
leihuan@msu.edu, lyuliyao@msu.edu

MS117

Optimizing Shot Assignment in Variational Quantum Eigensolver Measurement with Machine Learning

Variational quantum algorithms (VQAs) present a promising approach to finding approximate solutions for computationally demanding problems by combining parameterized quantum circuits with classical optimization techniques. However, estimating probabilistic outcomes on quantum hardware requires repeated measurements (shots), which can be resource-intensive when higher accuracy is needed. The limited quantum measurement budget results in shot noise during the energy objective function evaluation in VQAs, posing challenges for gradient-based quantum optimization. In this talk, I will briefly introduce our recently proposed approaches to reducing shot usage in the variational quantum eigensolver (VQE), an important application of VQAs used to find the ground state energy of a system. Our approaches include (1) Strategic shot allocation for energy objective function evaluation; (2) Reinforcement learning-based shot assignment; (3) Efficient gradient estimation in gradient-based quantum optimization.

Senwei Liang
Lawrence Berkeley National Laboratory
liang339@purdue.edu

Linghua Zhu
University of Washington
linghua8@uw.edu

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

Xiaosong Li
University of Washington
xsli@uw.edu

MS117

Finite Expression Method for Committor Functions

In this talk, we present the finite expression method (FEX, Liang and Yang (2022)) as a tool for computing the committor function for rare events. FEX approximates the committor by an algebraic expression involving a fixed finite number of nonlinear functions and binary arithmetic operations. The optimal nonlinear functions, the binary operations, and the numerical coefficients in the expression template are found via reinforcement learning. The FEX-based committor solver is tested on several high-dimensional benchmark problems. It gives comparable or better results than neural network-based solvers. Most importantly, FEX is capable of correctly identifying the algebraic structure of the solution which allows one to reduce the committor problem to a low-dimensional one and find the committor with any desired accuracy.

Zezheng Song
University of Maryland, College Park
zsong001@umd.edu

Maria K. Cameron
University of Maryland
mariakc@umd.edu

Haizhao Yang
University of Maryland College Park
haizhaoyang@yahoo.com

MS118

Enhanced Data Assimilation Based on the Energy Spectrum of Nonlinear Chaotic Dynamics

In data assimilation, an ensemble provides a way to propagate the probability density of a system described by a nonlinear prediction model. Although a large ensemble size is required for statistical accuracy, the ensemble size is typically limited to a small number due to the computational cost of running the prediction model, which leads to a sampling error. Several methods, such as localization and inflation, exist to mitigate the sampling error, often requiring problem-dependent fine-tuning and design. This work introduces a nonintrusive sampling error mitigation method that modifies the ensemble to ensure a smooth turbulent spectrum. It turns out that the ensemble modification to satisfy the smooth spectrum leads to inhomogeneous localization and inflation, which apply varying localization and inflation levels at different locations. The efficacy of the new idea is validated through a suite of stringent test regimes of the Lorenz 96 turbulent model.

Bosu Choi
Georgia State University
bchoi12@gsu.edu

Yoonsang Lee
Dartmouth College
ylee@math.dartmouth.edu

MS118

Efficiency and Accuracy Trade-Offs: Kans Vs. Mlps in Neural Operator Learning for Pdes

We present a comparative study of Multi-Layer Perceptrons (MLPs) and Kolmogorov-Arnold Networks (KANs) within the deep operator network (DeepONet) framework for learning PDE operators. The Universal Approximation Theorem states that any continuous function defined over a compact subspace can be approximated by a finite-width neural network with a single hidden layer and a fixed activation function, typically implemented as MLPs. The Kolmogorov-Arnold Representation Theorem (KAT) provides specific bounds, showing that a network of width

$2n+1$ (where n is the input dimension) with learnable activation functions can approximate any continuous function, leading to KANs. The promise of reduced parameters and improved interpretability with KANs motivates our objective to compare the cost-accuracy performance of MLPs and KANs in DeepONet. We analyze their approximation qualities through loss landscapes, linearity and homogeneity test errors, Rademacher complexity, and computational complexity (floating point operations, training, and inference time). Preliminary results demonstrate the superior accuracy of KAN-DeepONet with fewer parameters in specific PDE learning scenarios to the difference in activation functions. This advantage is particularly pronounced in cost reduction and robustness to perturbations and parameter uncertainty when the PDE operator class allows for smooth Kolmogorov-Arnold representation.

Hassan Iqbal, Raghav Pant
The University of Texas at Austin
hassan.iqbal@utexas.edu, raghav.pant@utexas.edu

Xingjian Li
Emory University
xingjian.li@austin.utexas.edu

Somdatta Goswami
Johns Hopkins University
sgoswam4@jhu.edu

Krishna Kumar
University of Texas at Austin
krishnak@utexas.edu

MS118

DeepPropNet: A Recursive Deep Neural Network Propagator for Learning Evolutionary PDE Operators

In this talk, we propose a deep neural network approximation to the evolution operator for time dependent PDE systems over long time period by recursively using one single neural network propagator, in the form of POD-DeepONet with built-in causality feature, for a small time interval. The trained DeepPropNet of moderate size is shown to give accurate prediction of wave solutions over the whole time interval.

Lizuo Liu
Dartmouth College
lizuo.liu@dartmouth.edu

MS118

On Some Computational Aspects for Phase Field Models

We explore a class of splitting schemes employing implicit-explicit (IMEX) time-stepping to achieve accurate and energy-stable solutions for thin-film equations and Cahn-Hilliard models with variable mobility. These splitting methods incorporate a linear, constant coefficient implicit step, facilitating efficient computational implementation. We investigate the influence of stabilizing splitting parameters on the numerical solution computationally, considering various initial conditions. Furthermore, we generate energy-stability plots for the proposed methods, examining different choices of splitting parameter values and timestep sizes. These methods enhance the accuracy of the original bi-harmonic-modified (BHM) approach, while preserving its energy-decreasing property and achieving second-order accuracy. We present numerical experiments to illustrate the performance of the proposed methods.

Saulo Orizaga
Department of Mathematics
NMT
saulo.orizaga@nmt.edu

MS118

Machine Learning Enhanced Time Integration Methods for PDEs

The numerical solution of partial differential equations (PDEs) is critical in many scientific and engineering applications, but traditional time integration methods often face significant computational challenges, especially when dealing with complex systems. In this study, we present a machine-learning enhanced time integration method for PDEs that uses neural networks to accelerate classical time integrators. In our approach, a neural network is trained to predict the evolution of the system, allowing the use of larger time steps. The loss function of the neural network is formulated based on the residuals of the fully discrete systems resulting from the discretization of the underlying PDEs. This novel approach eliminates the need for ground truth data during the training process. By learning the underlying dynamics from data generated by the classical solver, the neural network can bypass some of the computationally expensive intermediate steps. This hybrid strategy seamlessly combines the stability and robustness of conventional time integration schemes with the computational efficiency of deep learning. Our results demonstrate the potential of machine learning to improve the efficiency of numerical methods for solving PDEs, making them more practical for large-scale and real-time applications.

Lu Zhang
Rice University
lz82@rice.edu

MS119

On Recent Advancements in the Euler-Lagrange Modeling of Particle-Laden Flows

Euler-Lagrange (EL) modeling has a long history. It was initially applied in the dilute regime where particles are smaller than the Kolmogorov scale. In this limit, its application becomes rigorous with an accurate force model. In recent decades, the EL methodology has been applied to more complex multiphase examples where the particles are larger than the Kolmogorov scale and their volume fraction is appreciable. In this situation, the traditional application of EL methodology with only the feedback force suffers in accuracy, even when the particle force is computed with the best available drag model that depends on the particle Reynolds number and local volume fraction. First, the force on individual particles substantially deviated from the model prediction, both in the drag and lift components. Second, the Reynolds stress arising from subgrid turbulence is important and also varies from particle to particle. Third, if the filter width of the EL governing equations becomes comparable to particle size, force variation around the particle surface cannot be ignored and cannot be assumed to act at the center of the particle. In this talk, we will address how best to account for these effects and arrive at a next-generation EL framework that is far more accurate for large particles at higher concentrations. The use of machine learning in this modeling effort

will also be discussed.

S. Balachandar
University of Florida
bala1s@ufl.edu

MS119

A Filtered Coarse-Grain CFD-DEM Approach for Simulating Fluidized Particles

Resolved CFD-DEM simulations of gas-particle flows are prohibitively expensive at industrial scales. A common approach to reduce the computational cost is to coarse grain the simulation by lumping particles into parcels and coarsening the grid. Traditional approaches neglect particle-particle variation within the parcel and often fail to converge to the underlying deterministic equations in the limit the number of particles within the parcel approaches unity. In this work, a rigorous formulation of the filtered Euler-Lagrange equations are presented. The equations, while exact, give rise to unclosed terms, notably a sub-filter drag force. Parcel collisions are handled using a modified soft-sphere approach. Sub-filter drag force is informed by highly resolved Euler-Lagrange simulations of unbounded, moderately dense gas-solid flows. Variation in particle velocities within each parcel is quantified. The relative contribution of the sub-filtered drag is found to increase with the number of particles per parcel, and depend on the local filtered volume fraction and Reynolds number. Symbolic regression is employed to obtain closed-form algebraic models.

Sathvik Bhat
University of Michigan
sathvikb@umich.edu

Yuan Yao, Yi Fan
Engineering and Process Science, Core R&D,
The Dow Chemical Company, Lake Jackson, TX, USA
yyao9@dow.com, yfan5@dow.com

Jorg Theuerkauf
Engineering and Process Science, Core R&D,
The Dow Chemical Company, Midland, MI, USA
jtheuerkauf@dow.com

Jesse Capecelatro
University of Michigan
jcaps@umich.edu

MS119

Physics Informed Hierarchical Machine Learning Model of Hydrodynamic Forces and Torques in Particle-Laden Flows

This study introduces a novel physics-inspired, data-driven framework for modeling particle-laden flows that overcomes the limitations of conventional methods. By building on the hierarchical model of [**?**], our framework incorporates global and local heterogeneities and integrates higher-order quaternary interactions. We leverage machine learning, to first sequentially train neural networks to predict isolated n-order interactions for $(N = n)$-body systems. We then feed the information obtained from binary $n = 2$, ternary $n = 3$ and quaternary $n = 4$ interactions to a suspension neural network, with an architecture featuring distinct, shared blocks for each n interaction order and linear activation to integrate interactions of varying orders. Our findings demonstrate significant promise in capturing the hydrodynamic disturbances within suspensions of homogeneous and stepwise heterogeneous test cases. Specifically, our study indicates that the reformulated Quaternary Hierarchical Model (QHM), which incorporates four hierarchical structures (based on unary, binary, ternary, and quaternary interactions), provides promising results in predicting the forces and torques experienced by a reference particle. Our results support the generalizability and universality of all proposed model architectures with minimal variation in performance between training and testing across all datasets, as well as for each individual subset.

Layal Jbara, Anthony Wachs, Zihao Cheng
University of British Columbia
lmj44@mail.ubc.ca, multiphase1973f82210, glennn1909@gmail.com

MS119

Recent Advances in Euler-Euler Modeling of Multiphase Polydisperse Flows with Quadrature-Based Moment Methods

Many industrial processes involve multiphase flows made of a continuum phase and other disperse phases, which may experience turbulent mixing and chemical reaction. An established approach to describe these flows is based on tracking the evolution of the joint probability density function (PDF) of the populations of entities composing the disperse phase and of the composition. PDF methods have been used to describe disperse multiphase flows, leading to Eulerian-Eulerian multi-fluid models, as well as to high-order moment methods that account for deviation from equilibrium and for particle trajectory crossing. They have also been applied to model turbulent scalar mixing in reacting flows, where the PDF of the mixture fraction defines the local mixing conditions, and a conditional moment closure (CMC) is used to describe the mixing of reactants. In this presentation, the foundations of multivariate quadrature-based moment methods for disperse multiphase flows are introduced, focusing on two applications: non-equilibrium disperse flows, and chemically reacting flows with a disperse phase involving mass transfer. In both cases, a conditional quadrature method of moments (CQMOM) is formulated. In the first problem, a joint size-velocity number density function is considered. In the second, CQMOM is applied to CMC. The approach is then extended to use the generalized quadrature method of moments, which allows the quadrature discretization to be arbitrarily refined.

Alberto Passalacqua
Iowa State University
Department of Chemical and Biological Engineering
albertop@iastate.edu

Lorenzo Vasquez Giuliano
Department of Mechanical Engineering
Iowa State University
vasquezl@iastate.edu

Rodney O. Fox
Iowa State University
Department of Chemical and Biological Engr.
rofox@iastate.edu

MS120

Sample-Efficient Active Learning Strategies for Deep Learning in Scientific Computing

Active learning is an important topic in machine learning

for scientific computing in which learning algorithms can query ground truth data selectively to enhance model accuracy. This is increasingly vital in science applications where data acquisition is costly. This talk introduces a broad framework for active learning in regression problems that extends beyond traditional pointwise data samples to include various practical data types such as data from transform domains (e.g., Fourier data), vector-valued data (e.g., gradient-augmented data), data along continuous curves, and multimodal data (i.e., involving combinations of different types of measurements). This framework uses random sampling from a finite number of sampling measures and accommodates arbitrary nonlinear approximation spaces (model classes). We then introduce generalized Christoffel functions to optimize these sampling measures and discuss how this leads to near-optimal sampling strategies for various important problems of interest. The focus will be on applications in scientific computing, highlighting the efficacy of this framework in gradient-augmented learning with polynomials, Magnetic Resonance Imaging (MRI) with generative models, adaptive sampling for solving PDEs using Physics-Informed Neural Networks (PINNs), and operator learning.

Nick Dexter
Florida State University, U.S.
nick.dexter@fsu.edu

Ben Adcock
Simon Fraser University
ben_adcock@sfu.ca

Juan Cardenas
University of Colorado Boulder
juca7891@colorado.edu

Sebastian Moraga
Simon Fraser University
sebastian_moraga_scheuermann@sfu.ca

MS120

Explorative in-Situ Analysis of Turbulent Flow Data Based on a Data-Driven Approach

The Proper Orthogonal Decomposition (POD) has been used for several years in the post-processing of highly-resolved Computational Fluid Dynamics (CFD) simulations. While the POD can provide valuable insights into the spatial-temporal behaviour of single transient flows, it can be challenging to evaluate and compare results when applied to multiple simulations. Therefore, we propose a workflow based on data-driven techniques, namely dimensionality reduction and clustering. The aim is to extract knowledge from simulation bundles arising from large scale transient CFD simulations. As an example, we apply this workflow to investigate the flow around two cylinders that contain complex modal structures in the wake region. A special emphasis lies on the introduction of in-situ algorithms to compute suitable data-driven representations efficiently and concurrently to the run of a simulation on the compute cluster. The in-situ data analysis approach reduces the amount of data in- and output, but also enables a simulation monitoring to reduce computational efforts, e.g. a data-driven early stopping or outlier analysis when running several simulations in engineering design studies. Finally, a classifier is trained to predict characteristic physical behaviour in the flow only based on the input parameters, e.g. the relative positions of the two cylinders.

Jochen Garcke
Institut für Numerische Simulation, Universität Bonn
Fraunhofer SCAI
garcke@ins.uni-bonn.de

Christian Gscheidle
Fraunhofer SCAI
christian.gscheidle@scai.fraunhofer.de

MS120

LazyDINO: Fast, Scalable, and Efficiently Amortized Bayesian Inversion via Structure-Exploiting and Surrogate-Driven Measure Transport

We present LazyDINO, a variational inference method for fast, scalable, and efficiently amortized solutions to large-scale nonlinear Bayesian inverse problems with expensive parameter-to-observable (PtO) maps. When given observation data, the method rapidly optimizes for a structure-exploiting transport map with low-dimensional non-linearity, i.e., a lazy map [Brennan et al., NeurIPS, (2020)], using a fast-to-evaluate ridge function surrogate of the PtO map. The transport map then produces approximate posterior samples, given data. We construct an optimized surrogate by numerically minimizing upper bounds on the expected error in the surrogate-driven lazy map optimization and posterior approximation. When the surrogate is parameterized by neural networks, we show that such an optimized surrogate is a derivative-informed neural operator (DINO) [O'Leary-Roseberry et al., J. Comput. Phys., 496 (2024)] learned using joint samples of the PtO map and its Jacobian. Our numerical results show that LazyDINO consistently outperforms existing variational inference methods, including Laplace approximation, lazy map, and neural posterior estimation, for solving challenging infinite-dimensional PDE-constrained nonlinear Bayesian inverse problems. Furthermore, at the same training sample generation cost, LazyDINO is 10–25 times more accurate in amortized posterior approximation compared to LazyNO, which utilizes conventional operator learning methods trained on PtO map samples alone.

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

Joshua Chen
The University of Texas at Austin
joshua@oden.utexas.edu

Michael C. Brennan
MIT, CCSE
mcbrenn@mit.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

Omar Ghattas
The University of Texas at Austin

omar@oden.utexas.edu

MS120

Fast Kronecker tensor operations for sparse grids

Sparse grids is a commonly used approach for constructing surrogate models in high dimensional parameter space and helping accelerate applications of uncertainty quantification. The grids are constructed from hierarchical superposition of multidimensional tensors of single dimensional basis functions and many operations resemble algebraic operations using classical Kronecker matrices. Kronecker theory is a powerful tool for tensor linear algebra that allows us to combine like-terms and reduce the overall computational cost by a factor exponential in the number of dimensions. However, the same like-term relationship is lost in a general sparse context. We present several sufficient conditions that allow us to recover the Kronecker cost-reduction in a hierarchical sparse grid context. We look at completeness of the hierarchy, direction of information transfer (up or down the hierarchy), as well as minimal padding strategies. We demonstrate 100x speedup for even moderate dimensions when computing surplus coefficients for sparse grid interpolation and even more when using sparse grids and discontinuous Gelerkin method for the solution of high-dimensional partial differential equations.

Miroslav Stoyanov
Oak Ridge National Laboratory, U.S.
stoyanovmk@ornl.gov

MS121

Principled Structures in Deep Learning-Based Autoregressive Models of Dynamical Systems

We present a general theoretical framework, based on linear stability analysis and random matrix theory to analyze deep neural networks (DNNs) when applied to nonlinear PDEs, for example, in terms of stability and convergence. Data-driven solvers, which are fast and cheap once trained, have a broad range of applications, e.g., to perform probabilistic weather forecasting or efficient long-term emulation of the climate system. Many recent studies have shown promising results, for short-term forecasts, on canonical PDEs, and even for real-world data, such as atmospheric observations. However, there are also challenges such as long-term instabilities (drifts and blow-ups) and lack of convergence. These challenges have been difficult to address due to the lack of a rigorous (i.e., system- and architecture-agnostic) framework for analyzing DNNs when applied to PDEs. Here, we present such a system- and architecture-agnostic framework that combines numerical analysis, deep learning theory, and random matrix theory. The framework is based on linear stability analysis of DNNs wherein we demonstrate the predictive capabilities of the linear operator to the general behavior of the models in short- and long term using ideas from random matrix theory. We show that this framework can guide the development of stable, convergent neural PDE-solvers for any system and architecture, which has wide-ranging applications on nonlinear geophysics, e.g., to build data-driven weather/climate models.

Conrad Ainslie, Ashesh Chattopadhhyay
University of California Santa Cruz
cainslie@ucsc.edu, aschatto@ucsc.edu

MS121

Quantifying InSAR Coherence Loss Due to Snow Cover over Mt. St. Helens using a Computationally Inexpensive Neural Network

Synthetic Aperture Radar Interferometry (InSAR) is a powerful tool for monitoring ground deformation, but its effectiveness can be compromised by coherence loss due to snow cover, particularly in mountainous regions like Mt. St. Helens. This study introduces a computationally inexpensive neural network designed to quantify InSAR coherence loss caused by snow cover over Mt. St. Helens.

Rebecca Bussard
The Pennsylvania State University
rbussard@uoregon.edu

Guido Cervone
Penn State University
guc18@psu.edu

Christelle Wauthier
The Pennsylvania State University
cuw25@psu.edu

MS121

Conditional Gaussian Koopman Operator for Modeling Complex Systems and Data Assimilation

Koopman theory has been actively explored in the context of data- driven modeling during the past few years. The key advantage is the promise of linear dynamics in the Koopman embedding space, that facilitates the learning and control of complex dynamical systems. However, for many real applications of complex dynamical systems linear dynamics may demand a very high-dimensional Koopman embedding space and could not make full use of some existing knowledge of nonlinear physics. In this work, we propose a conditional Gaussian Koopman network (CGKN) framework that leverages both known nonlinear physics and the Koopman embedding of unknown physics for modeling a complex dynamical system. The framework also facilitates an analytical formula for data assimilation (DA) of unknown physics in both the Koopman embedding space and the original physical space. We demonstrate that the analytical formula of DA can be used in data-driven modeling as an additional computationally affordable loss. This DA loss function promotes the proposed framework to capture the interactions between state variables and thus advances its modeling skills. Numerical experiments based on chaotic systems with intermittency and strong non-Gaussian features indicate that the proposed framework outperforms the standard Koopman operator models that strictly aim at a linear model in the embedding space, and the DA loss of the unknown physics further enhances the modeling skills of the CGKN framework.

Chuanqi Chen, Nan Chen
University of Wisconsin-Madison
cchen656@wisc.edu, chennan@math.wisc.edu

MS122

High-Order Approximation of Smooth Target

Functions in High(er) Dimensions

The easiest representation of an approximation to a function f using only point-wise data has the form $Q(f) = \sum_j f(\boldsymbol{x}_j)\varphi_j$ with suitable *shape-functions* φ_j. Usually, one demands for such *quasi-interpolation processes* Q that they reproduce a certain class of functions, e.g., polynomials of a given degree. One way to construct shape functions that provide polynomial reproduction is by solving weighted least-squares problems, leading to a method commonly referred to as *moving least-squares*. Using compactly supported weight functions yields a fast, high-order reconstruction method using scattered data in arbitrary domains. Furthermore, the specific form of the quasi-interpolation operators allow an easy application of the combination technique to obtain a high-order approximation method of high-dimensional functions on sparse grids that are constructed with nearly arbitrary low-dimensional point sets. In this talk, I will repeat the basics of the low dimensional approximation process, the moving least-squares method. Then I introduce the new Smolyak moving least-squares method and discuss its properties. Additionally, I will confirm the theoretical results with numerical examples.

Rüdiger Kempf
University of Bayreuth, Germany
ruediger.kempf@uni-bayreuth.de

MS122

Optimal Sampling Recovery and Minimal Norm Sampling Projections

This talk is mainly concerned with new developments in understanding the power of different types of information (function values vs. linear measurements, deterministic vs. random) as well as different classes of algorithms (linear vs. nonlinear), that play an important role in the Information Based Complexity and Compressed Sensing. The recent growing attention to this topic is caused by numerous practical applications. I will concentrate on the problem of optimal recovery of functions from their samples and conditions under which linear sampling algorithms are optimal among all (possibly non-linear) algorithms. The main emphasis is made on the uniform recovery of bounded complex-valued functions. Besides, we show how to lift L_2-error bounds to error bounds in general semi-normed spaces using the spectral function. In what follows, we show that there are sampling projections onto arbitrary n-dimensional subspaces of the space of bounded functions with at most $2n$ samples and norm of order $\sqrt{n}$. This result is related to the famous Kadets-Snobar theorem and Auerbach's lemma. Further we discuss a problem of discretization of continuous norms. In the end of the talk, I will present recent findings on the tight MarcinkiewiczZygmund inequalities. This is based on joint work with Felix Bartel (TU Chemnitz), David Krieg (University of Passau), Martin Schfer (TU Chemnitz), Mario Ullrich (JKU Linz) and Tino Ullrich (TU Chemnitz).

Kateryna Pozharska
Institute of Mathematics of the National Academy of Sciences
Chemnitz University of Technology, Germany
pozharska.k@gmail.com

MS122

A Dimension-Adaptive Sparse Grid Algorithm for Stochastic PDEs with Spatial Singularities

In this talk, we present a new adaptive algorithm for the approximation of the solution of a stochastic elliptic PDE with uncertain input. These problems require both an approximation in the high-dimensional stochastic parametric space as well as the spatial discretization. Monte Carlo or stochastic collocation methods together with a finite element method on uniformly refined meshes have shown to work well for smooth solutions. However, uniform refinement becomes less efficient when dealing with PDE problems with spatial singularities. We propose a dimension-adaptive sparse grid algorithm that incorporates adaptivity in the spatial variable for problems with structural singularities. More precisely, we replace the uniform refinement by a sequence of adaptively refined meshes generated using an h-adaptive finite element method and suitable error estimators. We then balance the spatial approximation with the stochastic approximations through a sparse grid formulation which is obtained by a dimension-adaptive algorithm based on the benefit-cost ratio. This strategy optimizes computational cost and accuracy without a priori knowledge of convergence rates. We present numerical results that demonstrate the effectiveness and robustness of the proposed approach and compare it to approximations with uniform refinements. The new approach achieves in practice better computational complexity and convergence rates comparable to those for problems without spatial singularities.

Uta Seidler
University of Bonn
Institute for Numerical Simulation
seidler@ins.uni-bonn.de

Michael Griebel
Institut für Numerische Simulation, Universität Bonn
and Fraunhofer-Institut SCAI
griebel@ins.uni-bonn.de

MS122

Sparse and Multilevel Approximations for PDE-constrained Optimization Under Uncertainty

In this talk we are concerned with the minimization of the expected value of a functional constrained by a random partial differential equation. The solution of such problems requires an extremely high computational cost, and thus motivates a very active research area. Several works proposed multilevel/sparse approximations of the expectation operator, which however leads to negative quadrature weights which may destroy the (possible) convexity of the continuous optimization problem. We here present a novel and alternative framework for using multilevel and sparse quadrature formulae (presented in [Nobile, Vanzan, *A combination technique for optimal control problems constrained by random PDEs*, SIAM/ASA J. on UQ., 2024] and [Nobile, Vanzan, *Multilevel quadrature formulae for the optimal control of random PDEs*, arXiv:2407.06678, 2024]), that still preserves the properties of the original problem. Our approach consists in solving a sequence of optimization problems, each discretized with different levels of accuracy of the physical and probability spaces. The final approximation of the control is obtained in a postprocessing step, by suitably *combining* the adjoint variables computed on the different levels. We will discuss a complete convergence analysis for multilevel quadrature formulae, and present numerical experiments confirming the better computational complexity of our multilevel approach, even

beyond the theoretical assumptions.

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Tommaso Vanzan
Politecnico di Torino, Italy
tommaso.vanzan@polito.it

MS123

Laplace Neural Operator for Solving Differential Equations

Neural operators (NOs) map multiple functions to different functions, possibly in different spaces, unlike standard neural networks. Hence, NOs allow solution of parametric ordinary (ODEs) and partial differential equations (PDEs) for a distribution of boundary/initial conditions and excitations but can also be used for system identification as well as designing various components of digital twins. We introduce the Laplace neural operator (LNO), which incorporates the pole-residue relationship between input-output spaces, leading to better interpretability and generalization for certain classes of problems. LNO is capable of processing non-periodic signals and transient responses resulting from simultaneously zero and non-zero initial conditions, which makes it achieve better approximation accuracy over other neural operators for extrapolation circumstances in solving several ODEs and PDEs. We also highlight LNOs good interpolation ability, from a low-resolution input to high-resolution outputs at arbitrary locations within the domain. To demonstrate the scalability of LNO, we conduct large-scale simulations of Rossby waves around the globe, employing millions of degrees of freedom. Taken together, our findings show that a pre-trained LNO model offers an effective real-time solution for general ODEs and PDEs at scale and is an efficient alternative to existing NOs.

Qianying Cao
Brown University
qianying_cao@brown.edu

Somdatta Goswami
Johns Hopkins University
usa
somdatta.goswami@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS123

Function Encoders and Their Applications in Physics Informed Neural Networks

Physics-Informed Neural Networks (PINNs) has gained significant traction in recent years as an alternative approach to traditional numerical methods for solving PDEs arising in physics, engineering, and other fields. The physics-informed nature of PINNs differentiates it from many other methods, such as DeepONet, by enabling model training without the need for existing data. However, one limitation of PINNs is that the training of a model is usually tied to specific inputs, such as boundary, initial conditions or other source terms. Consequently, PINNs struggles to handle problems with varying input conditions, as retraining is required whenever these conditions change. To address this limitation, we introduce function encoders. In its original form, function encoders use neural networks to generate a set of basis functions that span target functional Hilbert spaces leveraging the inner product. The function encoders method excels in efficiency and accuracy and has seen success in areas such as operator learning and more. In our work, we propose physics-informed function encoders, where a function encoder is used to parameterize the input functional space, with additional PDE loss incorporated into the learning problem to better capture the system's underlying physics. Once trained, these function encoders are able to predict solutions to differential equations under different boundary or input conditions with minimal computation.

Xingjian Li
Emory University
xingjian.li@austin.utexas.edu

Krishna Kumar, Adam J. Thorpe, Tyler Ingebrand
University of Texas at Austin
krishnak@utexas.edu, adam.thorpe@austin.utexas.edu,
tyleringebrand@utexas.edu

Somdatta Goswami
Johns Hopkins University
sgoswam4@jhu.edu

MS123

One-shot Learning for Solution Operators of Partial Differential Equations

Learning and solving governing equations of a physical system, represented by partial differential equations (PDEs), from data is a central challenge in a variety of areas of science and engineering. Traditional numerical methods for solving PDEs can be computationally expensive for complex systems and require the complete PDEs of the physical system. On the other hand, current data-driven machine learning methods require a large amount of data to learn a surrogate model of the PDE solution operator, which could be impractical. Here, we propose the first solution operator learning method that only requires one PDE solution, i.e., one-shot learning. By leveraging the principle of locality of PDEs, we consider small local domains instead of the entire computational domain and define a local solution operator. The local solution operator is then trained using a neural network, and utilized to predict the solution of a new input function via mesh-based fixed-point iteration (FPI), meshfree local-solution-operator informed neural network (LOINN) or local-solution-operator informed neural network with correction (cLOINN). We test our method on diverse PDEs, including linear or nonlinear PDEs, PDEs defined on complex geometries, and PDE systems, demonstrating the effectiveness and generalization capabilities of our method across these varied scenarios.

Lu Lu
Yale University
lu.lu@yale.edu

MS123

Novel Methods for Derivative-Informed Neural Operator Learning

In this talk, we will discuss efficient methods for constructing neural operator approximations of parametric PDE solutions. By exploiting low-dimensional sensitivity informa-

tion within a high-dimensional map (when such features exist), we generate and learn high-dimensional Jacobian with computational costs that are independent of discretization dimensions. We use the sensitivity information to construct a derivative-informed neural operator (DINO) that serves as a powerful surrogate for various applications, notably high-dimensional design and inverse problems involving PDE-constrained optimization. We show that DINO results in significantly improved surrogate accuracy at the same training data generation cost as conventional operator learning methods, leading to high quality solutions in design and inverse problems. Our numerical results cover a range of applications, including fluid flow and nonlinear elasticity.

Boyuan Yao
UT Austin
johnyao@utexas.edu

Lianghao Cao, Nikola Kovachki
California Institute of Technology
lianghao@caltech.edu, nkovachki@caltech.edu

Dingcheng Luo
Oden Institute for Computational Engineering and Sciences,
The University of Texas at Austin
dc.luo@utexas.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS123

Structured and Balanced Multi-Component and Multi-Layer Neural Networks (MMNN)

In this work, we propose a balanced multi-component and multi-layer neural network (MMNN) structure to approximate functions with complex features with both accuracy and efficiency in terms of degrees of freedom and computation cost. The main idea is motivated by a multi-component, each of which can be approximated effectively by a single-layer network, and multi-layer decomposition. While an easy modification to fully connected neural networks (FCNNs) or multi-layer perceptrons (MLPs) through the introduction of balanced multi-component structures in the network, MMNNs achieve a significant reduction of training parameters, a much more efficient training process, and a much improved accuracy compared to FCNNs or MLPs. Extensive numerical experiments are presented to illustrate the effectiveness of MMNNs in approximating high oscillatory functions and its automatic adaptivity in capturing localized features.

Hongkai Zhao
Duke University
zhao@math.duke.edu

Shijun Zhang
The Hong Kong Polytechnic University
shijun.zhang@polyu.edu.hk

Yimin Zhong
Auburn University
yzz0225@auburn.edu

Haomin Zhou
School of Mathematics, Georgia Institute of Technology
GA, USA
hmzhou@gatech.edu

MS124

Generative Adversarial Representation Learning for Stochastic Dynamical Systems: Data-Driven Modeling of Particle Interactions, Force-Laws, and Noise-Driven Dynamics

Recent emerging deep learning methods combined with more traditional numerical analysis are presenting new opportunities for developing more robust data-driven approaches for modeling and simulation. A central challenge is to incorporate physical principles and other prior scientific knowledge to obtain robust models and some level of interpretability, while still retaining flexibility in learning. We discuss Stochastic Dynamic Generative Adversarial Networks (SDYN-GANs) for sample-based MMD learning of probabilistic models from trajectory observations of stochastic systems. SDYN-GANs learns dynamical representations in terms of SDEs and stable m-step stochastic numerical integrators for use in simulations. SDYN-GANs uses a sample-based approach avoiding the need to specify likelihood functions and can be used to provide losses alternative to MLE. We show how SDYN-GANs can be used for inertial stochastic systems arising in statistical mechanics to learn parameters both of the drift and diffusive contributions. We then discuss how SDYN-GANs can be used to learn unknown non-linear force-laws from observations of the trajectories of the stochastic dynamics. The discussed methods and results show a few strategies toward developing more robust and interpretable machine learning methods for scientific simulations.

Paul J. Atzberger
University of California-Santa Barbara
atzberg@ucsb.edu

MS124

Clutter to Clarity: Emergent Neural Operators Via Questionnaire Metrics

Traditional supervised learning models struggle with unlabeled or disorganized real-world datasets, making it difficult to learn underlying dynamics. To address this, we've developed a framework that combines unsupervised organizational learning (using questionnaires) with state-of-the-art deep neural operators to create "structural" generative models, uncover emergent equations, and build effective emulators from shuffled data. We demonstrate the framework's effectiveness on two illustrative systems: a 1D advection-diffusion partial differential equation representing a winding underground pipe, and an ensemble of Stuart-Landau oscillators, an agent-based system of coupled ordinary differential equations. In both cases, we successfully reconstructed meaningful spatial, temporal, and parameter embeddings from scrambled data, leading to accurate predictions of system dynamics. This framework has potential for wider applications, particularly in fields where data is inherently disorganized or where parameter spaces are hidden, enabling data-driven system identification.

Anastasia S. Georgiou, Arjun Manoj

Johns Hopkins University
ageorgi3@jhu.edu, amanoj1@jhu.edu

Pei-Chun Su
Yale University
pei-chun.su@yale.edu

Ronald Coifman
Yale University
Department of Computer Science
coifman@math.yale.edu

Somdatta Goswami, Ioannis G. Kevrekidis
Johns Hopkins University
sgoswam4@jhu.edu, yannisk@jhu.edu

MS125

A New Boundary Condition for the Nonlinear Poisson-Boltzmann Equation in Electrostatic Analysis of Proteins

As a well-established implicit solvent model, the Poisson-Boltzmann equation (PBE) models the electrostatic interactions between a solute biomolecule and its surrounding solvent environment over an unbounded domain. One numerical challenge in solving the nonlinear PBE lies in the boundary treatment. Physically, the boundary condition of this solute solvent system is defined at infinity where the electrostatic potential decays to zero. Computationally, a finite domain has to be employed in grid-based numerical algorithms. However, the Dirichlet boundary conditions commonly used in protein simulations are known to produce unphysical solutions in some cases. This motivates the development of a few asymptotic conditions in the PBE literature, which are global boundary conditions and have to resort to iterative algorithms for calculating volume integrals from the previous step. To overcome these limitations, a simple Robin condition is proposed in this work as a local boundary condition for the nonlinear PBE, which can be implemented in any finite difference or finite element method. The derivation is based on the facts that away from the biomolecule, the asymptotic decaying pattern of the nonlinear PBE is essentially the same as that of the linearized PBE, and the monopole term will dominate other terms in the multipole expansion. Asymptotic analysis has been carried out to validate the application range and robustness of the proposed Robin condition.

Sylvia Amihere
Southern Methodist University
Department of Mathematics
samihere@smu.edu

MS125

A Poisson-Nernst-Planck Single Ion Channel Model and Its Effective Finite Element Solver

A single ion channel is a membrane protein with an ion selectivity filter that allows only a single species of ions (such as potassium ions) to pass through in the open state. Its selectivity filter also naturally separates a solvent domain into an intracellular domain and an extracellular domain. Such biological and geometrical characteristics of a single ion channel are novelly adopted in the construction of a new kind of dielectric continuum ion channel model, called the Poisson-Nernst-Planck single ion channel (PNPSIC) model. In this talk, an effective PNPSIC finite element solver is introduced for a single ion channel with a three-dimensional X-ray crystallographic molecular structure and a mixture of multiple ionic species. Numerical results for a potassium channel confirm the convergence and efficiency of the PNPSIC finite element solver and demonstrate the high performance of the software package. Moreover, the PNPSIC model is applied to the calculation of electric current and validated by biophysical experimental data.

Zhen Chao
Western Washington University
chaoz2@wwu.edu

MS125

Methods and Tools for the Studies of Biomolecular Electrostatic Interactions

This talk reports our recent development of numerical methods and software for the studies of electrostatic interactions of solvated biomolecules. The adopted mathematical models for the physical description of biomolecular electrostatics are the Poisson-Boltzmann model and the Poisson-Nernst-Planck model at various level of details. The solutions to these models are numerically challenging due to long-range pairwise interaction, complex geometry, interface discontinuity, charge singularity, large-scale computing, etc. We designed and developed methods of fast summation, mesh generation, interface treatment, charge regularity, as well as parallel computing and machine learning to address these challenges.

Weihua Geng
Southern Methodist University
wgeng@smu.edu

MS125

An Algebraic Graph Neural Network Model for Protein-Ligand Binding Affinity Prediction

Predicting protein-ligand binding affinity is a fundamental challenge in drug discovery. Recent advances in deep learning have led to the development of numerous models, many of which rely on three-dimensional protein-ligand complex structures and focus primarily on affinity prediction. In this study, we introduce an Algebraic Graph Neural Network (AGNN) model designed to encode molecular structures into a low-dimensional graph representation while preserving critical biochemical interactions. While algebraic graph theory has been widely used in physical modeling and molecular studies, traditional methods often struggle to accurately capture the complexity of biomolecular interactions. To address this limitation, our proposed AGNN model leverages multiscale weighted colored subgraphs to describe molecular interactions through graph neural network. These representations allow the model to effectively learn the geometric and topological features of protein-ligand complexes. The AGNN model integrates graph convolutional layers and attention mechanisms to refine feature extraction and improve the interpretability of learned embeddings. Furthermore, we incorporate gradient boosting decision trees (GBDTs) to enhance the prediction of binding affinities by capturing nonlinear relationships between molecular features. Our approach is extensively validated using benchmark datasets, including PDBBind and CASF-2016, demonstrating superior performance in binding affinity prediction compared to state-of-the-art scoring functions.

Augustine Ouru

University of Arkansas
aouru@uark.edu

MS126

Inference of Utilities and Time Preference in Stochastic Controls

This work introduces a novel stochastic control framework to enhance the capabilities of automated investment managers, or robo-advisors, by accurately inferring clients' investment preferences from past activities. Our approach leverages a continuous-time model that incorporates utility functions and a generic discounting scheme of a time-varying rate, tailored to each client's risk tolerance, valuation of daily consumption, and significant life goals. We address the resulting time inconsistency issue through state augmentation and the establishment of the dynamic programming principle and the verification theorem. Additionally, we provide sufficient conditions for the identifiability of client investment preferences. To complement our theoretical developments, we propose a learning algorithm based on maximum likelihood estimation within a discrete-time Markov Decision Process framework, augmented with entropy regularization. We prove that the log-likelihood function is locally concave, facilitating the fast convergence of our proposed algorithm. Practical effectiveness and efficiency are showcased through two numerical examples, including Merton's problem and an investment problem with unhedgeable risks. Our proposed framework not only advances financial technology by improving personalized investment advice but also contributes broadly to other fields such as healthcare, economics, and artificial intelligence, where understanding individual preferences is crucial.

Haoyang Cao
Johns Hopkins University
hycao@jhu.edu

Zhengqi Wu, Renyuan Xu
University of Southern California
zhengqiw@usc.edu, renyuanx@usc.edu

MS126

Applications of High-Dimensional Hamilton-Jacobi Solutions

This talk presents recent results and future challenges in developing numerical solutions to certain classes of the Hamilton-Jacobi (HJ) equations as they relate to computing reachable sets of dynamic systems and optimal vehicle trajectories. Many interesting problems can be used in this solution framework including multiple vehicle coordination, environment exploration and information collection, and collaborative pursuit-evasion. This talk will pay special attention to safety critical applications such as reactive collision avoidance and autonomous emergency landing of aircraft. Specific HJ equations form a theoretical basis that precisely characterize both safe and unsafe operating conditions, necessitating the rapid computation of these solutions for safety critical operations, where a vehicle must react to unforeseen events in the environment. HJ equations have had limitations in the past for computing usable solutions due to poor scaling with respect to system dimension. This was because of the fact that a spatial grid had to be constructed densely in each dimension. Creating equivalent formulations that no longer require spatial grids is the goal of the research and leads to methods that can be executed in real-time on embedded hardware. Additionally, I will discuss ongoing efforts to expand and generalize these methods to cover a larger field of applications.

Matthew Kirchner
Auburn University
kirchner@auburn.edu

MS126

Analysis and Numerical Approximation of Mean Field Game Partial Differential Inclusions

The Mean Field Game (MFG) system of Partial Differential Equations (PDE), introduced by Lasry & Lions in 2006, models Nash equilibria of large population stochastic differential games of optimal control where the players of the game have unique optimal controls, and the convex Hamiltonian of the underlying optimal control problem is differentiable. In this talk, we introduce a new class of model problems called Mean Field Game Partial Differential Inclusions (MFG PDI), which extend the MFG system of Lasry and Lions to situations where players may have possibly nonunique optimal controls, and the resulting Hamiltonian of the underlying optimal control problem is not required to be differentiable. We prove the existence of unique weak solutions to MFG PDI for a broad class of Hamiltonians that are convex, Lipschitz, but possibly nondifferentiable, under a monotonicity condition similar to one considered previously by Lasry & Lions. Moreover, we introduce a class of monotone finite element discretizations of the weak formulation of MFG PDI and present theorems on the strong convergence of the approximations to the value function in the $L^2(H_0^1)$-norm and the strong convergence of the approximations to the density function in $L^p(L^2)$-norms. We conclude the talk with discussion of numerical experiments involving non-smooth solutions.

Yohance Osborne, Iain Smears
University College London
yohance.a.osborne@durham.ac.uk, i.smears@ucl.ac.uk

MS126

Modeling in Reinforcement Learning for Robust Control

Optimal control designed with reinforcement learning can be sensitive to model mismatch. We demonstrate that designing such controllers in a virtual simulation environment with an inaccurate model is not suitable for deployment in a physical setup. Controllers designed using an accurate model are robust against disturbance and small mismatch between the physical setup and the mathematical model derived from first principles; while a poor model results in a controller that performs well in simulation but terrible in physical experiments. Sensitivity analysis was used to justify these discrepancies. We are particularly interested with a controller that is non-smooth and piecewise-affine with respect to the state variables.

Nikki Xu
North Carolina State University
xxu35@ncsu.edu

Hien Tran
Center for Research in Scientific Computation
North Carolina State University

tran@ncsu.edu

MS127

Modeling Unknown Stochastic Systems Via Generative Model

We present a numerical framework for learning unknown stochastic dynamical systems using measurement data. Termed stochastic flow map learning (sFML), the new framework seeks to approximate the unknown flow map of the underlying system. Technically, it is realized by (conditional) generative models, such as Generative Adversarial networks (GANs), Autoencoders and Normalizing Flows. Once a sFML model is trained, it serves as a stochastic evolution model that is a weak approximation, in term of distribution, of the unknown stochastic system. It allows us to analyze the long-term system behavior under different initial conditions. A comprehensive set of numerical examples are presented to demonstrate the flexibility and effectiveness of the proposed sFML method for various types of stochastic systems. It is capable of handling systems driven by both Gaussian and non-Gaussian noises, even for jump processes, such as systems generated by Gillespie's stochastic simulations.

Yuan Chen
SIAM
chen.11050@osu.edu

Dongbin Xiu
Ohio State University
xiu.16@osu.edu

MS127

Stein Variational Inference and Sparsification for Uncertainty Quantification of Physical Neural Networks

Most scientific machine learning (SciML) applications of neural networks involve hundreds to thousands of parameters, and hence, uncertainty quantification for such models is plagued by the curse of dimensionality. Using physical applications, we show that L_p, $p \in [0, 1]$, sparsification combined with Stein variational gradient descent (SVGD) is a more robust and efficient means of uncertainty quantification, in terms of computational cost and performance than the direct application of SGVD or projected SGVD methods. The talk will show demonstrations of the performance of a few variations on the proposed approach and discuss the applicability of each.

Reese Jones
Sandia National Laboratories
rjones@sandia.gov

Govinda Padmanabha
Cornell University
ga355@cornell.edu

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

Nikolaos Bouklas
Cornell University
nb589@cornell.edu

MS127

Stochastic Neural Ordinary Differential Equations Models of Time-Dependent Processes with Uncertainty

We propose a scalable approximate inference framework in neural ordinary differential equation (NODE) models for time dependent processes. The posterior distribution corresponding to the NODE model parameters (weights and biases) follows a stochastic differential equation with a drift term that is learned jointly with the NODE model parameters. This approach offers flexibility in balancing the computational budget between the NODE model and the approximate posterior density for its parameters. We demonstrate the gradient-based stochastic variational inference in canonical settings to validate its construction. We then apply it to material deformation models with uncertain internal deformation dynamics that are unobservable due to measurement limitations.

Cosmin Safta, Reese Jones, Craig M. Hamel, Dan Bolintineanu, Ravi Patel
Sandia National Laboratories
csafta@sandia.gov, rjones@sandia.gov, chamel@sandia.gov, dsbolin@sandia.gov, rgpatel@sandia.gov

Raelynn Wonnacott
University of Maryland
raelynnw@umd.edu

MS127

A Hybrid Bayesian DeepONet for Data Assimilation and Uncertainty Quantification in Hypersonic Flows

Deep operator networks (DeepONets) have shown significant promise and success in operator regression. Bayesian DeepONets inherit these advantages in addition to quantifying uncertainties in their predictions through their stochastic parameters. The goal of training a Bayesian DeepONet is to learn the posterior distribution of the network parameters. Hamiltonian Monte Carlo (HMC) methods can provide accurate estimates for the posterior distribution of parameters. However, HMC techniques are computationally demanding due to the high dimensionality of the parameter space and the non-convexity of posterior distributions. Therefore, various approximation techniques such as Variational Inference (VI) are employed to infer the posterior distribution of parameters. However, such approximations may result in inaccurate estimates of uncertainties. We propose a hybrid framework combining VI and HMC to infer the posterior distribution of parameters in Bayesian DeepONets. This hybrid framework can accurately predict uncertainties while converging significantly faster than traditional HMC. We use this framework to train a Bayesian DeepONet to act as a surrogate model for the compressible Navier-Stokes equations. We use this surrogate to perform data assimilation to identify oncoming disturbances from wall-pressure measurements in a Mach 6 flow over a cone. We demonstrate the accuracy and efficiency of the data-assimilation process and quantify uncertainties in the predicted disturbances.

Ponkrshnan Thiagarajan, Tamer A. Zaki, Michael D. Shields
Johns Hopkins University

ponkrshnan@jhu.edu, t.zaki@jhu.edu, michael.shields@jhu.edu

MS128

Designing Conservative and Accurately Dissipative Numerical Integrators in Time

Numerical methods for the simulation of transient systems with structure-preserving properties are known to exhibit greater accuracy and physical reliability, in particular over long durations. These schemes are often built on powerful geometric ideas for broad classes of problems, such as Hamiltonian or reversible systems. However, there remain difficulties in devising higher-order-in-time structure-preserving discretizations for nonlinear problems, and in conserving non-polynomial invariants. In this work we propose a new, general framework for the construction of structure-preserving timesteppers via finite elements in time and the systematic introduction of auxiliary variables. The framework reduces to Gauss methods where those are structure-preserving, but extends to generate arbitrary-order structure-preserving schemes for nonlinear problems, and allows for the construction of schemes that conserve multiple higher-order invariants. We demonstrate the ideas by devising novel schemes that exactly conserve all known invariants of the Kepler and Kovalevskaya problems, and arbitrary-order schemes for the compressible Navier-Stokes equations that conserve mass, momentum, and energy, and provably dissipate entropy.

Boris Andrews
University of Oxford
boris.andrews@oriel.ox.ac.uk

Patrick E. Farrell
Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS128

Pressure-robustness for Fluid Flow with Inhomogeneous Boundary Conditions

In various fluid flow applications complex and inhomogeneous boundary conditions appear. Especially for non-Newtonian fluids the boundary conditions play a crucial role. Even in the special case of Newtonian fluids boundary conditions are still under investigation. For pressure robust numerical schemes the velocity error estimates are independent of the pressure. One way to obtain pressure-robustness is to use methods with exactly divergence-free discrete velocity approximations, such as the Scott-Vogelius finite element. The discrete solutions can be approximated by an Uzawa type iteration: the iterated penalty method (IPM). We present a pressure-robust velocity estimate for the IPM for the Stokes problem subject to inhomogeneous Dirichlet boundary conditions. One should note, that it is not possible to impose any approximation of the boundary data while maintaining pressure-robustness. Another challenge arises from the singular vertices affecting the convergence rate.

Franziska Eickmann
Technical University of Darmstadt
eickmann@mathematik.tu-darmstadt.de

Ridgway Scott
University of Chicago
ridg@uchicago.edu

Tabea Tscherpel
Universitat Bielefeld
tscherpel@mathematik.tu-darmstadt.de

MS128

Accelerated Solvers for Bingham and Regularized Bingham Models of Viscoplastic Flow

We develop, analyze and test accelerated solvers for the Bingham and Regularized Bingham models of viscoplastic flow, and the Oldroyd-B model of viscoelastic flow. We consider applying Anderson acceleration to commonly used solvers for these systems. We first analytically study the properties of the solvers' associated fixed point operators to check if they fit the Anderson convergence theory sufficient condition assumptions. We then test the solvers with Anderson acceleration (and variations thereof) in an effort to get solvers to converge faster and for larger parameter ranges (and thus for a wider set of physical problems).

Leo Rebholz
Clemson University
Department of Mathematical Sciences
rebholz@clemson.edu

MS128

A Second-Order-in-Time, Explicit Approach Addressing the Redundancy in the Low-Mach, Variable-Density Navier-Stokes Equations

Low-speed turbulent flows with significant density variation are often modeled with the low-Mach variable-density Navier-Stokes equations. For these equations, there is a redundancy between the mass conservation equation, the equation of state, and the transport equation(s) for the scalar(s) which characterize the thermochemical state. This redundancy is challenging, particularly for explicit time discretizations because inconsistencies in the state variables lead to instabilities and limit the magnitude of density variation that can be simulated. In this presentation, we introduce a novel algorithm for explicit temporal discretization of the low-Mach equations. We demonstrate how to analytically eliminate the redundancy and propose an iterative scheme to solve the resulting transformed scalar equations. The method obtains second-order accuracy in time regardless of the number of iterations, so one can terminate this subproblem once stability is achieved. Hence, flows with larger density ratios can be simulated while still retaining the efficiency, low cost, and parallelizability of an explicit scheme. We demonstrate the algorithm within a pseudospectral direct numerical simulation which extends the method of Kim, Moin, and Moser (1987) for incompressible flow to the variable-density, low-Mach system, achieving stability for density ratios up to ~ 25.7.

Bryan Reuter
Sandia National Laboratories
bwreute@sandia.gov

Todd Oliver
The University Of Texas-Austin
oliver@oden.utexas.edu

Robert Moser
UT Austin
Oden Institute

rmoser@oden.utexas.edu

MS129

Transport Transforms and Wasserstein-like Distances

We will present some recent developments in non-linear transforms derived from optimal transport theory. When data are well-modeled by transport phenomena, these transforms simplify the data structure in the transform domain, thereby simplifying machine learning algorithms. Additionally, Wasserstein-like distances between two measures μ, ν can be derived, which are computationally efficient.

Akram Aldroubi
Vanderbilt University
aldroubi@math.vanderbilt.edu

MS129

Structured Approximations in Wasserstein Space

We consider structured approximation of measures in Wasserstein space $W_p(\mathbb{R}^d)$ for $p \in [1, \infty)$ using general measure approximants compactly supported on Voronoi regions derived from a scaled Voronoi partition of $\mathbb{R}^d$. We show that if a full rank lattice Λ is scaled by a factor of $h \in (0, 1]$, then approximation of a measure based on the Voronoi partition of $h\Lambda$ is $O(h)$ regardless of d or p. We then use a covering argument to show that N-term approximations of compactly supported measures is $O(N^{-\frac{1}{d}})$ which matches known rates for optimal quantizers and empirical measure approximation in most instances. Additionally, we generalize our construction to nonuniform Voronoi partitions, highlighting the flexibility and robustness of our approach for various measure approximation scenarios. Finally, we extend these results to noncompactly supported measures with sufficient decay. Our findings are pertinent to applications in computer vision and machine learning where measures are used to represent structured data such as images.

Keaton Hamm
University of Texas at Arlington
keaton.hamm@uta.edu

Varun Khurana
University of California, San Diego
varun_khurana@brown.edu

MS129

Data-Driven Nonlinear Filtering Algorithms with Optimal Transport Maps

In this talk, I present a new variational formulation of the Bayes' law, that will be used for construction of a new family of nonlinear filtering algorithms. The variational formulation is based on the optimal-transportation (OT) theory, and aims at approximating the Brenier Optimal transport map from the prior to the posterior distribution, as a solution to a stochastic optimization problem. I present error analysis and numerical results that illustrate the performance of the algorithm in comparison with the SIR particle filters, and present an extension of the algorithm that is model-free and only requires recorded data from the state and observation processes.

Amirhossein Taghvaei
University of Washington Seattle
amirtag@uw.edu

Mohammad Al-Jarrah, Bamdad Hosseini
University of Washington
mohd9485@uw.edu, bamdadh@uw.edu

MS129

Nonparametric Distribution Learning Via Neural ODEs

In this talk, we explore approximation properties and statistical aspects of Neural Ordinary Differential Equations (Neural ODEs). Neural ODEs are a recently established technique in computational statistics and machine learning, that can be used to characterize complex distributions. Specifically, given a fixed set of independent and identically distributed samples from a target distribution, the goal is either to estimate the target density or to generate new samples. We first investigate the regularity properties of the velocity fields used to push forward a reference distribution to the target. This analysis allows us to deduce approximation rates achievable through neural network representations. We then derive a concentration inequality for the maximum likelihood estimator of general ODE-parametrized transport maps. By merging these findings, we are able to determine convergence rates in terms of both the network size and the number of required samples from the target distribution.

Youssef M. Marzouk, Youssef M. Marzouk, Zhi Ren
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu, zren@mit.edu

Sven Wang
Department of Mathematics
Humboldt University Berlin
sven.wang@hu-berlin.de

Jakob Zech
Heidelberg University
jakob.zech@uni-heidelberg.de

MS130

A Quadrature Technique for Efficient Kalman Filtering with Model-Space Covariance Localization

In geoscientific applications of ensemble data assimilation, covariance localization is crucial for mitigating the effects of spurious correlations between model variables, as well as for preventing ensemble collapse. Certain observation types require localization to be done in model space, a computationally demanding task which involves various trade-offs between accuracy, runtime, and memory usage. This talk introduces a new algorithm for model space localization which achieves a favorable trade-off between these three factors by combining modern techniques from numerical quadrature, Krylov subspace iteration, and matrix function evaluation. This algorithm is compatible with a wide variety of spatial and spectral covariance localization schemes, is parallelizable, and is built upon linear-algebraic primitives whose accuracy is backed up by strong error analyses. We will present the results of numerical experiments comparing its accuracy and efficiency to existing methods such as the gain-form ensemble transform Kalman filter (GETKF). We will also discuss recent efforts to test this algorithm on large-scale atmospheric forecasting problems using the Joint Effort for Data Assimilation Integra-

tion (JEDI) and the Model for Prediction Across Scales (MPAS).

Robin Armstrong
Cornell
rja243@cornell.edu

Ian Grooms
Department of Applied Mathematics
University of Colorado, Boulder
ian.grooms@colorado.edu

Chris Snyder
Mesoscale and Microscale Meteorology Division
NCAR
chriss@ucar.edu

MS130

Data Assimilation Techniques Using Wasserstein Metric

Data Assimilation aims to produce an optimal prediction of states/parameters by combining sparse observations and a numerical model. Particle filters have gained traction over the last few decades because it has no underlying linear, Gaussian assumptions, something that plagues the "Kalman' based filters. In this study, we develop a particle filter which minimizes the Wasserstein distance to get samples from the posterior distribution. The new filter demonstrates promising results in low dimensional problems. We also explore and discuss the challenges in higher dimensions and propose possible solutions.

Abhinab Bhattacharjee
Virginia Tech
abhinab93@vt.edu

MS130

Riemannian Manifold Hamiltonian Monte Carlo Via Derivative-Informed Neural Operators Towards Digital Twins for Laser Powder Bed Fusion

We will briefly introduce our efforts towards digital twins for laser powder bed fusion—an additive manufacturing process governed by thermoelastoplasticity. The talk will then elaborate on derivative-informed neural operators (DINOs) as a surrogate framework for PDE-constrained parameter-to-observable maps and their parametric derivative (i.e. the Jacobian). We introduce DINO^2 which extends the formulation of the DINO reduced-basis operator learning problem to the topology of the H^2 Sobolev space with a Gaussian measure. The resulting DINO^2 surrogates faithfully estimate the Jacobian and the Hessian of high-dimensional parameter-to-observable maps. In turn, they enable accelerating complex sampling methods, such as Riemannian Manifold Hamiltonian Monte Carlo, to rapidly sample Bayesian posteriors with high acceptance ratios while mitigating the computational cost of the required parametric derivative information. Numerical experiments are provided to compare the proposed approach to other sampling methods.

Blake Christierson
University of Texas at Austin
bechristierson@utexas.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

Nicholas Alger
The University of Texas at Austin
Oden Institute
nalger225@gmail.com

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS131

A Stochastic Trust-Region Method with Adaptive Subspaces for Quantum Approximate Optimization Algorithms

We introduce ANASTAARS, a noise-aware stochastic trust-region algorithm that leverages adaptive random subspace strategies to achieve scalability across low-, moderate-, and potentially high-dimensional problems. The method optimizes random models within low-dimensional affine subspaces, significantly reducing per-iteration function evaluation costs. These subspaces and their dimensions are defined using JohnsonLindenstrauss transforms, such as Haar-distributed orthogonal random matrices. Unlike previous approaches that rely on fixed-dimension subspaces, ANASTAARS introduces an adaptive subspace selection strategy. Instead of generating entirely new interpolation points at each iteration, it selectively updates the model by adding a fewor even a single-new point while reusing past points and function values from lower-dimensional subspaces. This ensures the interpolation set remains poised while avoiding the inefficiencies of constructing models in fixed-dimension subspaces. Additionally, to improve robustness in noisy settings, ANASTAARS integrates noise-aware optimization techniques by estimating the noise level in function evaluations, an approach known to enhance trust-region methods when the signal-to-noise ratio is low. The effectiveness of the proposed method is demonstrated through numerical experiments in the context of the quantum approximate optimization algorithm (QAOA).

Kwassi Joseph Dzahini
Argonne National Laboratory
kdzahini@anl.gov

MS131

Adaptive Computing: Derivative-Free Optimization in the Context of Distributed Computing and Resource Constraints

On-the-fly model training is conducted using simulations, which are executed on various computing resources (e.g., high-performance computing, cloud, and local resources). A Bayesian optimization problem is solved to identify candidate locations in the sample space where simulations could be performed. An integer decision problem is solved to determine which candidate simulations should be run and on which resources. The solution minimizes model uncertainty subject to resource-specific constraints as well as

global constraints such as time to solution.

Kevin P. Griffin
National Renewable Energy Laboratory
kevin.griffin@nrel.gov

Juliane Mueller
National Renewable Energy Lab
juliane.mueller@nrel.gov

Weslley Da Silva Pereira
National Renewable Energy Laboratory
Colorado
weslley.dasilvapereira@nrel.gov

Hilary Egan, Nicholas Wunder, Olga Doronina, Monte Lunacek
National Renewable Energy Laboratory
hilary.egan@nrel.gov, nick.wunder@nrel.gov,
olga.doronina@nrel.gov, monte.lunacek@nrel.gov

Ryan King
NREL
ryan.king@nrel.gov

Marc Day
National Renewable Energy Laboratory
marcus.day@nrel.gov

MS131

Branch-and-Model: A Derivative-Free Global Optimization Algorithm

Derivative-free optimization (DFO) is an important class of optimization algorithms that solve problems based on objective and function evaluations. In this work, we present a novel derivative-free global algorithm Branch-and-Model (BAM). The BAM algorithm features a flexible partition scheme combined with model-based search techniques, which exploit the local trend and accelerate convergence during solution refinement. Our study demonstrates that the timing of local refinement is crucial, with DBSCAN clustering proving effective in improving this timing and in identifying promising areas for further refinement. Extensive computational experiments over 500 publicly open-source test problems show that BAM outperforms state-of-the-art DFO algorithms, particularly in higher-dimension problems. The algorithm holds promise for a wide range of applications.

Kaiwen Ma
Carnegie Mellon University
kaiwenm@alumni.cmu.edu

Luis Miguel Rios
End -to-End Analytics, Lima, Peru
lmrios@gmail.com

Atharv Bhosekar
GAMS Development Corporation, U.S.
atharv.bhosekar@gmail.com

Nikolaos Sahinidis
Georgia Institute of Technology
nikos@gatech.edu

Sreekanth Rajagopalan
The Dow Chemical Company
srajagopalan@dow.com

MS131

Matrix-Free Linear Algebra for Trajectory Optimization

In this talk, we discuss matrix-free linear algebra that scales trajectory optimization to large discretizations. Unlike the traditional notion of sparsity in which matrices are stored as arrays of index-value pairs, we consider linear algebra that avoids the instantiation of large, unwieldy matrices altogether. To illustrate this idea in the context of trajectory optimization, we focus on spectral collocation discretizations. Although these discretizations produce linear operators with dense matrix representations, collocating at the Chebyshev-Lobatto points permits us to apply these operators "faster than we can write them down." That is, with respect to the size of the discretization, applications of the operator scale almost linearly while sizes of the matrices representing the operators scale quadratically. We present a new preconditioner that extends these fast applications to fast matrix-free linear algebra. We then incorporate that linear algebra into an optimization framework to solve a low-thrust orbit transfer problem with millions of unknowns. This solve is faster than its matrix-based counterpart, even when the latter has orders of magnitude fewer unknowns.

John D. Steinman
Rice University
jds27@rice.edu

Aurya Javeed
Sandia National Laboratories
asjavee@sandia.gov

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

I. Michael Ross
Naval Postgraduate School
imross@nps.edu

MS131

Barycenter of Weight Coefficient Region of Least Weighted H^2 Norm Updating Quadratic Models with Vanishing Trust-Region Radius

Derivative-free optimization problems are problems where derivative information is unavailable or extremely difficult to obtain. Least H^2 norm updating quadratic model is a novel kind of model function proposed for model-based derivative-free trust-region algorithms. Different weight coefficients refer to different models, which are important for the algorithm. The radii of trust region in such algorithms will converge to 0 when solving most numerical problems. This talk tries to discuss the weight coefficients in a highly accurate case, which usually comes out when the numerical solution is close to the theoretical local solution. We discuss the selection of the weight coefficients in the objective function of the subproblem for interpolation model, minimizing which provides the least H^2 norm

updating quadratic model, when the trust-region radius is vanishing. KKT matrix distance and KKT matrix error are defined to illustrate the distance or difference between the KKT matrices obtained by minimizing weighted H^2 norms with different weight coefficients. We propose the weight coefficients whose corresponding KKT matrix is the center, among KKT matrices provided by the coefficient region satisfying $0 \leq C_1 + C_2 \leq 1$, in the case where $C_1 + C_2 + C_3 = 1$, where C_1, C_2, C_3 are the weight coefficients. Such weight coefficient is the barycenter of the corresponding coefficient region. Numerical results also show the advantage of using the barycenter.

Pengcheng Xie
Lawrence Berkeley National Laboratory
pxie@lbl.gov

Stefan Wild
Lawrence Berkeley National Laboratory
Applied Mathematics and Computational Research
wild@lbl.gov

MS132

An Efficient Streaming Type Algorithm for the General Gauss-Markov Linear Model

We discuss new methods to solve very large generalized regression problems, including those were data comes in continually, as in the case of satellite data, so that the 'full' problem is never available. Our methods must also be able to handle very large general observation and prior covariance matrices. This requires iterative methods to compute solution search spaces and approximate solutions using variants of the CS decomposition that are easily updated without having to manipulate large, generally dense, orthogonal factors. For problems where data comes in continually, we need methods to efficiently and effectively combine current solution data and search spaces with incoming data to continually update solutions.

Eric De Sturler
Virginia Tech
sturler@vt.edu

Christopher A. Beattie
Virginia Polytechnic Institute and State University
beattie@vt.edu

MS132

Downscaling with Uncertainty using Data-Driven Surrogate Models

Super-resolution of climate models (also known as downscaling) is key in prediction and planning for critical infrastructure. The relationship between fine-scale and coarse-scale information is complex and often not well captured with traditional statistical or machine-learning techniques, especially in the context of extreme events. Incorporating knowledge about the physics of the system can increase accuracy but at a computational cost. The advent of data-driven surrogates, however, has shown great promise in accelerating prediction. We utilize one such model with a data assimilation framework to carry out super-resolution with uncertainty quantification on the fine-scale prediction.

Philip Dinenis
Argonne National Laboratory, USA
pdinenis@anl.gov

MS132

Exploring Use of Machine Learning for Practical Data Assimilation Problems

Data assimilation for numerical weather prediction using the operational system ingests a wide range of spatially and temporary irregular, heterogeneous observations to re-initialize the state through optimization. Both model state obtained by forecast from the last initialization and observations are subject to noise and errors. There are a few common features in data assimilation and machine learning that either can benefit from each other. Conventionally, operational data assimilation involves empirical data selection and quality control of observations in the ingestion process. In this talk, some examples for the integration will be presented from both theoretical and practical point of views.

Kayo Ide
University of Maryland, College Park
ide@umd.edu

MS132

Fast and Scalable FFT-Based GPU-Accelerated Algorithms for Hessian Actions Arising in Linear Inverse Problems Governed by Autonomous Dynamical Systems

We present an efficient and scalable algorithm for performing matrix-vector multiplications ("matvecs") for block Toeplitz matrices. Such matrices arise in the context of solving inverse problems governed by autonomous systems, and time-invariant systems in particular. We consider inverse problems for inferring unknown parameters from observational data of a linear time-invariant dynamical system given in the form of PDEs. Matrix-free Newton-conjugate-gradient methods require solving a pair of linearized forward/adjoint PDE solves per Hessian action, which may be prohibitive for large-scale, hyperbolic inverse problems. However, time invariance of the forward PDE problem leads to a block Toeplitz structure of the discretized parameter-to-observable (p2o) map defining the mapping from inputs (parameters) to outputs (observables) of the PDEs. This block Toeplitz structure enables us to exploit two key properties: compact storage of the p2o map and its adjoint; and efficient fast Fourier transform-based Hessian matvecs. The proposed algorithm is mapped onto large multi-GPU clusters and achieves more than 80 percent of peak bandwidth on an NVIDIA A100 GPU. Excellent weak scaling is shown on TACC's Lonestar6 GPU cluster for up to 48 A100 GPUs. For the targeted problems, the implementation executes Hessian matvecs within fractions of a second, orders of magnitude faster than can be achieved by the conventional matrix-free Hessian matvecs via forward/adjoint PDE solves.

Sreeram R. Venkat
University of Texas Austin
201 E 24th St, Austin, TX 78712
srvenkat@utexas.edu

Milinda Fernando
Oden Institute, UT Austin
milinda@oden.utexas.edu

Stefan Henneking
University of Texas at Austin

stefan@oden.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS133

Superfast Divide-and-conquer SVDs

For matrices that have small off-diagonal numerical ranks, there exist superfast divide-and-conquer methods that can find accurate approximate eigenvalue decompositions in nearly linear complexity if the matrix is symmetric. Now, when the matrix is nonsymmetric, we can also design a superfast divide-and-conquer SVD algorithm. In this work, we propose a structured divide-and-conquer method without the need to symmetrize the matrix. The matrix is converted into a multilevel block broken-arrowhead form that is suitable for the method. Safeguards are also incorportated into the method to ensure orthogonality of singular vectors and to handle clustered singular values. Numerical tests show significant advantages over traditional SVDs in terms of the storage and the cost.

Chenyang Cao, Xiaofeng Ou, Jianlin Xia
Purdue University
cao302@purdue.edu, ou17@purdue.edu, xiaj@purdue.edu

MS133

Performance Portability with OpenMP for Nuclear Structure

Many-Fermion Dynamics–nuclear, or MFDn, is a configuration interaction (CI) code for nuclear structure calculations. The computational effort consists mainly of constructing a large (dimension $\sim 10^{10}$), sparse ($\sim 10^{13}$ nonzero element) matrix, and subsequently finding a few tens of extremal eigenpairs. It uses a hybrid MPI+X programming model in Fortran 90. For CPU-based systems, it has a highly-optimized OpenMP implementation, while an OpenACC implementation targets GPU-accelerated systems. In order to improve the portability of MFDn, in particular in order to use the upcoming Aurora system, MFDn has been updated to use OpenMP for GPU offloading. In this work we describe some key challenges to translating from an OpenACC implementation to an OpenMP implementation.

Patrick J. Fasano
Argonne National Laboratory
patrick.fasano@nextsilicon.com

Pieter Maris
Iowa State University
pmaris@iastate.edu

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

MS133

Accelerating Eigenvalue Computation Via Perturbative Corrections

We present the Subspace Projection with Perturbative Corrections (SPPC) method, a new method for efficiently computing the lowest eigenvalues and corresponding eigenvectors of large-scale symmetric matrices. The SPPC method partitions the matrix into a sum of two components, treating the second as a perturbation to the first. By focusing initially on the first matrix, which allows for more efficient computation, we obtain approximate eigenvalues and eigenvectors. Subsequently, perturbative corrections are derived from a sequence of linear systems, refining these approximations. These correction vectors can be combined with the approximate eigenvectors of the first matrix to construct a subspace from which more accurate approximations of the desired eigenpairs can be obtained. Numerical experiments, including those involving a nuclear many-body Hamiltonian and a 2D Laplacian matrix, demonstrate that the SPPC method can be more efficient than conventional iterative methods, such as Lanczos, block Lanczos, and the locally optimal block preconditioned conjugate gradient (LOBPCG) method. Additionally, the SPPC method can be combined with other methods to avoid convergence stagnation.

Dong Min Roh, Dong Min Roh
Lawrence Berkeley National Laboratory
droh@lbl.gov, droh@lbl.gov

MS133

The Effectiveness of Sketched Rayleigh Ritz in Large Scale Eigenvalue Solvers

Sketched Rayleigh-Ritz (sRR) is a randomized alternative to classical Rayleigh-Ritz that avoids the expensive orthogonalization that plagues Krylov-like iterative methods. We investigate the viability of this method with the Lanczos and the Davidson methods. With unrestarted Lanczos, sRR may provide computational benefits if the number of eigenvalues sought is not too large so that new Krylov vectors provide new information to the basis. For restarted Lanczos sRR does not work. For Davidson, there are computational trade-offs but for very large matrix sizes sRR can provide a speedup up to two.

Andreas Stathopoulos
William & Mary
Department of Computer Science
andreas@cs.wm.edu

MS134

Multi-Scale Modeling of Phenotypic Heterogeneity During Collective Invasion in Cancer.

Collective cell invasion is an emergent property of interactions between cancer cells and their surrounding extracellular matrix (ECM). However, tumor populations invariably consist of cells expressing variable levels of adhesive proteins that mediate such interactions, disallowing an intuitive understanding of how tumor invasiveness at a multicellular scale is influenced by spatial heterogeneity of cell-cell and cell-ECM adhesion. Here, we have used a Cellular Potts model-based multiscale computational framework constructed on the histopathological principles of glandular cancers. Here, we constitute a tumor core of two separate cell subsets showing distinct intra- and inter-subset cell-cell or cell-ECM adhesion strengths. These two subsets are arranged to varying extents of spatial intermingling, which we call the heterogeneity index (HI). We observe that low and high inter-subset cell adhesion favors invasion of high-HI and low-HI intermingled populations with distinct intra-subset cell-cell adhesion strengths, respectively. In addition, for explored values of cell-ECM adhesion strengths, populations with high HI values collectively invade better than those with lower HI values. Our

simulations also reveal how adhesion heterogeneity qualifies collective invasion, when either cell-cell or cell-ECM adhesion type is varied but results in an invasive dispersion when both adhesion types are simultaneously altered, thus driving an epithelial-to-mesenchymal transition.

C V S Prasanna
Indian Institute of Science
chigicherlav@iisc.ac.in

Mohit Jolly
Indian Institute of Science
Department of Bioengineering,
mkjolly@iisc.ac.in

Ramray Bhat
Indian Institute of Science
ramray@iisc.ac.in

MS134

Digital Twin for Nanoparticle-Based Targeted Drug Delivery

The margination of particles, particularly platelets and nanoparticles, is a crucial phenomenon in blood flow, significantly influencing the efficacy of targeted drug delivery systems. Margination refers to the lateral drift and subsequent accumulation of these particles near the vessel walls, facilitating their interaction with the endothelium. Understanding margination is essential because it enhances the targeting precision of nanocarriers, improving therapeutic outcomes. This study introduces a novel computational framework to investigate the margination behavior of nanoparticles within blood flow, emphasizing the critical interactions between red blood cells and particulate matter. By integrating fluid dynamics and statistical mechanics principles, the model employs the Fokker-Planck equation to elucidate the concentration profiles of nanoparticles within the blood vessels. The insights gained from this study enhance our understanding of particle margination and offer valuable guidance for the rational design of nanocarriers. Furthermore, by coupling this spatial model with a traditional PBPK model, we aim to create a digital twin of a patient model. This digital twin incorporates essential physiological features such as blood flow dynamics and vasculature branching. By subjecting various biophysical attributes of the model to traditional sensitivity analysis tools, we can quantify uncertainty for this multiphysics PBPK model.

Sahil Kulkarni
University of Pennsylvania
Department of Bioengineering
sahil16@seas.upenn.edu

Ravi Radhakrish
Department of Bioengineering
University of Pennsylvania
rradhak@seas.upenn.edu

MS134

The Power of Verifiable Data-Driven Digital Twins in Science and Engineering

We demonstrate the transformative potential of Generative AI (GenAI) in creating verifiable digital twins for environmental science and public health. In environmental science, we address the challenge of predicting soil organic carbon (SOC) dynamics in permafrost regions. Using generative adversarial networks (GANs), we generate digital twins of soil carbon profiles to enhance climate model accuracy, validated against data from field collected circumpolar region. This approach reduces uncertainties in SOC stock predictions, crucial for understanding climate change impact on current and future land carbon reserves. Additionally, we demonstrate the use case in public health by demonstrating improved early detection of infections and identify outbreak hotspots during the Covid pandemic using GenAI based synthetic data augmentation. By applying a co-kurtosis-based projection for anomaly detection within multivariate timeseries data from personal wearables, we enhance our ability to detect anomalies in physiological metrics. This technique addresses data availability, bias, maintaining the statistical integrity of realworld data. These applications have great potential for application in improving sensor-based threat detection and pattern recognition in data-scarce environments using GenAI methods. We highlight the importance of digital twins in advancing scientific understanding and technological innovation, showcasing their role in addressing critical global challenges.

Kunal Poorey
Sandia National Laboratories
kpoorey@sandia.gov

Methun Kamruzzaman
Sandia National Laboratories
System Biology Department
mkamruz@sandia.gov

Jorge Salinas
Sandia National Laboratories
7011 East Ave, Livermore, CA 94551.
jsalin@sandia.gov

Uma Balakrishnan, Hemanth Kolla, Umakant Mishra, Kenneth Sale
Sandia National Laboratories
ubalakr@sandia.gov, hnkolla@sandia.gov, umishra@sandia.gov, klsale@sandia.gov

MS134

Reducing Uncertainties in Permafrost Soil Organic Carbon Using Genai Powered Digital Twins

Permafrost-affected soils, primarily located in high latitudes, store substantial quantities of soil organic carbon (SOC) primarily due to unique environmental conditions that prevent organic matter decomposition. Future projections indicate a significant reduction in permafrost areas due to global warming, potentially releasing large amounts of greenhouse gases into the atmosphere over extended periods. Given the critical role of soils in the global carbon cycle and climate regulation, it is essential to improve our understanding about the future permafrost soil carbon dynamics. However, data collection in high-latitude regions is challenging due to geographical and logistical constraints, resulting in limited data availability. To address this, we introduce a generative AI technique using generative adversarial networks (GANs) to create large datasets of digital twins of soil carbon profiles, aiming to reduce uncertainties in SOC stock predictions. Soil samples are classified by key environmental variables (e.g., biome, soil type, mean annual temperature, and precipitation) to generate synthetic profiles that can be assigned to locations with similar characteristics. A rigorous testing procedure is conducted in regions with abundant data (e.g., Alaska and Canada) to

evaluate the accuracy of this generative technique. Our efforts can reduce uncertainties in predicting climate change impacts on permafrost soil carbon and inform developing more reliable climate models.

Jorge Salinas
Sandia National Laboratories
7011 East Ave, Livermore, CA 94551.
jsalin@sandia.gov

Sagar Last Gautam
Sandia National Laboratories
sgautam@sandia.gov

Changpeng Fan
University of Arizona
cfmtok@arizona.edu

Methun Kamruzzaman
Sandia National Laboratories
System Biology Department
mkamruz@sandia.gov

Uma Balakrishnan, Kunal Poorey, Umakant Mishra
Sandia National Laboratories
ubalakr@sandia.gov, kpoorey@sandia.gov, umishra@sandia.gov

MS135

Error Bounds for Waveform Relaxation for Heterogeneous Heat Equations

We consider dynamical coupled problems such as heat transfer and fluid structure interaction, or more specifically PDEs that interact through a lower dimensional interface. Our general goal is a partitioned method that is high order, allows for different and adaptive time steps in the separate models, makes efficient use of hardware resources, is robust, and contains fast inner solvers. A prime candidate to fulfill this wishlist are waveform relaxation methods. In order to design and effectively use such methods, it is important to have error estimates. To this end, we study and analyze Dirichlet-Neumann waveform relaxation for two coupled 1D heterogeneous linear heat equations in the continuous and semidiscrete setting. We present linear and super-linear error bounds that depend on the material parameters, but also the domain sizes, time window length and the discretizations. Numerical experiments show that the length of the time window plays an important role in whether a linear or nonlinear error bound is more accurate. Additionally, it plays a role in whether we are in a regime where the continuous analysis gives an accurate description of the discrete behavior, or where it does not. The analysis also predicts which subdomain should use Dirichlet transmission conditions for fast convergence. Finally, we discuss the validity of the analysis for more realistic nonlinear test cases.

Philipp Birken
Lund University
philipp.birken@math.lu.se

Niklas Kotarsky
Centre for Mathematical Sciences at Lund University
niklas.kotarsky@math.lu.se

Martin J. Gander
University of Geneva
martin.gander@unige.ch

MS135

Multirate Methods in Microelectronics Simulations

Simulating microelectronic circuitry including ferromagnetic and ferroelectric materials requires coupling multiple physical processes across disparate spatial and temporal scales. In micromagnetics models computing the effective magnetic field that drives spintronic devices involves expensive convolution calculations for demagnetization effects that occur on a slower time scale compared to other processes. Similarly, phase-field models for ferroelectric-based devices involve costly linear system solves to compute the electric potential that evolves on a slower time scale relative to other effects. In this talk we present the application of multirate time integration methods from the SUNDIALS library to enable advancing processes in time at different rates to improve the efficiency in the AMReX-based codes, MagneX and FerroX.

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

Prabhat Kumar, Andrew J. Nonaka
Lawrence Berkeley National Laboratory
prabhatkumar@lbl.gov, ajnonaka@lbl.gov

MS135

Strong Stability-Preserving Runge-Kutta Schemes for Accurate Damage Evolution in a Numerical Ice Sheet Model

Accurate calculation of ice damage and its impact on flow and iceberg calving are necessary to predict ice-sheet response to climate forcing. However, numerical ice sheet models generally neglect the impact of ice damage except in simple parameterizations of iceberg calving. We couple ice damage to ice dynamics in the numerical ice sheet model MALI through an enhancement factor on the viscosity to investigate the effect of damage on the retreat of Thwaites Glacier, Antarctica. However, we observe that the modeled damage field is unphysically diffuse when using standard first-order numerical schemes, leading to a relatively uniform moderate value of damage across the entire floating portion of the glacier. This contrasts with observations that indicate highly localized damage features. To allow for accurate advection of damage features, we implemented a third-order flux-corrected transport scheme for thickness and tracers, but we found that a prohibitively small forward Euler time step is required to benefit from it. Therefore, we implemented three Strong Stability-Preserving Runge-Kutta time integration schemes for use with third-order advection. We find that these schemes preserve damage localization, allow for a roughly order-of-magnitude increase in time step length relative to first order, and yield a 50% decrease in simulated sea-level contribution from Thwaites Glacier by 2100 compared with simulations using third-order advection with first-order time stepping.

Trevor Hillebrand
Los Alamos National Laboratory
trhille@lanl.gov

Matthew Hoffman
Los Alamos National Laboratory, Los Alamos, NM, USA
mhoffman@lanl.gov

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

MS135

A Semi-Implicit Lagrangian Method for Radiation-Hydrodynamics

We propose a semi-implicit finite element scheme in the Lagrangian reference frame for the grey diffusion equations of radiation hydrodynamics. For this system, the hydrodynamic and radiation variables are nonlinearly coupled. This coupling does not easily lend itself to a straightforward IMEX partitioning strategy. To address this, we adopt the semi-implicit methods of Boscarino et al., enabling one to derive an efficient, high-order time integration method for this complex system of equations. The spatial discretization employs a staggered-grid continuous finite element scheme. This choice allows a stable update to the mesh nodes without requiring a multi-dimensional Riemann solver as would be required by a discontinuous Galerkin scheme. We validate the proposed method through several numerical tests from the literature.

Steven Walton, Brian Tran, Sam Olivier, Ben Southworth, Ryosuke Park
Los Alamos National Laboratory
stevenw@lanl.gov, btran@lanl.gov, solivier@lanl.gov, southworth@lanl.gov, hkpark@lanl.gov

MS136

Black-box Coupling of Simulation Codes Using preCICE

The coupling library preCICE (https://precice.org/) employs efficient numerical algorithms (such as interface quasi-Newton acceleration and RBF-based space interpolation methods) to couple stand-alone simulation codes into a multi-physics simulation. While originally developed for fluid-structure interaction simulations on HPC, this general-purpose library has been demonstrated for a wide variety of applications, including neuromuscular simulations, glacier simulations, and more. Recent additions to the core library include more efficient RBF mapping based on a Partition-of-Unity approach, time interpolation, as well as experimental support for geometric multiscale mapping and dynamic meshes. On top of that, new tools expand the capabilities of preCICE into different directions: the micro-manager enables two-scale simulations, while fmiprecice allows coupling to models implementing the functional mock-up interface. The preCICE community has developed several more tools and integrations with popular simulation packages, while it has also published simulation setups in various formats and venues. This talk will give an overview of recent advances in preCICE, and introduce the ecosystem standardization efforts stemming from the DFG project preECO.

Gerasimos Chourdakis
Technical University of Munich
gerasimos.chourdakis@ipvs.uni-stuttgart.de

MS136

Coupling Fluid and Kinetic Descriptions for Space Physics Simulations

In heliospheric physics, the physics of our solar system, there are two challenges that we often need to overcome: multiple scales and different environments. We will look at two examples, one where small-scale physics drives large scale dynamics, and a second where a plasma transitions between a collisional and collisionless regime. In the first example, we capture the small-scale physics by embedding kinetic, small scale, techniques into a large-scale fluid model. This will capture the important processes involved with magnetic reconnection on the smallest scale. The result is a generation of large scale flows that we can capture in a fluid description. In the second part we will discuss the work done to couple a highly collisional system to a system of both collisional and non-collisional plasmas. To do this, we make use of both fluid and kinetic models working together to capture the transition between collisional and collisionless regimes.

Lars Daldorff
NASA
larskristenselberg.daldorff@nasa.gov

Alex Glocer
GSFC NASA
alex.glocer-1@nasa.gov

Gabor Toth
Center for Space Environment Modeling
University of Michigan
gtoth@umich.edu

Yuxi Chen, Xiantong Wang
University of Michigan
yuxichen@umich.edu, xtwang@umich.edu

MS136

An Active Flux Method for the Vlasov-Poisson System

In this work we present a positivity-preserving third-order accurate Active Flux method for the 1+1-dimensional Vlasov-Poisson system. The Vlasov-Poisson system, a simplification of the Vlasov-Maxwell system, models the evolution of a collisionless plasma. Active flux methods are finite volume methods that use, in addition to cell averages, point values as degrees of freedom. In the presented work, the point values are evolved in time using the characteristic form of the equation, while cell averages are updated using the conservative form of the equation. The availability of additional point values allows us to compute moments of the Vlasov equation more easily, which enables the construction of an efficient third-order accurate method for the Vlasov-Poisson system. In order to guarantee the non-negativity of the approximated solution, a positivity-preserving flux limiter will be introduced which extends our method.

Yanick Kiechle
Heinrich Heine University Düsseldorf
yanick.kiechle@hhu.de

MS136

Entropy-Stable, Adaptive, High-Order Discontinuous Galerkin Scheme for Modelling Liquid Hydrogen Tank Sloshing Dynamics

One of the relevant physical phenomena for designing liquid hydrogen storage systems in motion is sloshing. It is a complex three-dimensional, non-stationary phenomenon which is difficult to model by simplified methods. This justifies the use for computational multi-phase fluid dynamics. As

part of the inter-disciplinary research project HYTAZER at DLR we use highly scalable Discontinuous Galerkin (DG) methods as a novel alternative to established numerical solution strategies. The objective is to derive, investigate, and implement a high-order, entropy-stable DG scheme [Gassner et. al,2021] for a weakly compressible two-phase fluid model with diffusive gas-liquid interface layers [Dumbser,2011] in the open-source code Trixi.jl [Schlottke-Lakemper et al.,2020]. The numerical solver operates on tree-based time-dynamic, adaptive meshes provided by our open-source software library t8code [Holke et al.,2022]. Next to provable entropy stability we rely on additional stability properties such as well-balancedness and a-posteriori regularization approaches [Markert,2021]. The final goal is to provide an open-source software framework for three-dimensional sloshing simulation on modern CPU and GPU platforms of the upcoming exascale era. In this talk we present the various numerical components, implementation details, and latest numerical results.

Johannes Markert
German Aerospace Center
johannes.markert@dlr.de

MS136

Real-Time Hybrid Physical-Numerical Simulation for Multi-Physic Systems of Model-Scale Floating Offshore Wind Turbine Testing

In physical experiments of small-scale multi-physic systems, it is impossible to completely represent all the relevant physics simultaneously without similitude distortions. Herein, we propose a Real-Time Hybrid Simulation (RTHS) method, a cyber-physical approach that partitions a prototype assembly into model-scale physical and full-scale computational sub-assemblies. These assemblies interact in real time through actuators and sensors, allowing different similitude laws to be applied to each sub-assembly. Our work focuses on the simulation of floating offshore wind turbines (FOWT), a classic multi-dynamics problem in scaled model testing, where incompatible scaling laws for gravitational, viscous, and solid mechanics effects hinder complete representation of aero-hydro-structural dynamic forces. First, we introduce the RTHS framework for a 1:50 modal-scale FOWT simulation in a wave basin, with aerodynamic forces simulated numerically and applied via robotic arm actuation. Then, we discuss the force control design for the real-time application of dynamic forces needed for RTHS. Our results show that the developed RTHS approach can simulate the response of FOWT subject to multiple physical interactions by mitigating similitude distortions and maintaining physical wave-structure interaction.

Yun Ni, Akiri Seki
Stanford University
yni23@stanford.edu, sekia@stanford.edu

Bret Bosma
Oregon State University
bosma@eecs.oregonstate.edu

Barbara Simpson
Stanford University
bsimpson@stanford.edu

Bryson Robertson, Ted Brekken
Oregon State University
bryson.robertson@oregonstate.edu,
brekken@eecs.oregonstate.edu

Andreas Schellenberg
Maffei Structural Engineering
andreas@maffei-structure.com

Pedro Lomonaco
Oregon State University
pedro.lomonaco@oregonstate.edu

MS137

Parallel Adaptive Mesh Refinement Algorithms on GPU for Unstructured Grids

Adaptive mesh refinement (AMR) is a technique used to increase the spatial resolution of specific regions within the spatial domain of a simulation. It is useful in large-scale simulations to achieve a more computationally efficient distribution of meshes. However, large-scale simulations also require the use of parallel computations, utilizing CPUs or GPUs. In this case, implementing an AMR technique in a parallel computation system presents challenges. In this study, we propose algorithms for parallel AMR on GPUs, focusing on unstructured grids in both 2D and 3D domains, composed of triangles and tetrahedrals, respectively. Our refinement strategy ensures the maintenance of a conformal mesh structure during the processes of refinement and coarsening. We want to utilize this strategy to solve flows present discontinouties like compressible flows with shockwaves.

Onur Ata, Ali Karakus
Middle East Technical University
onurata@metu.edu.tr, akarakus@metu.edu.tr

MS137

The Moving Discontinuous Galerkin Method with Interface Condition Enforcement for Shock-Dominated, Viscous Flows

The Moving Discontinuous Galerkin Method with Interface Condition Enforcement (MDG-ICE) is an implicit shock fitting method that simultaneously solves for the flow field and discrete geometry. MDG-ICE is able to compute highly accurate high-order solutions without artificial dissipation as the grid points are automatically adjusted to fit a priori unknown shocks with arbitrary topology and resolve smooth regions of the flow with sharp gradients. We begin with an overview of MDG-ICE for inviscid flows, in which shocks are fit exactly along grid interfaces. We then proceed to the viscous setting, wherein high-aspect-ratio elements form to resolve viscous shocks, which are sharp (yet smooth) features, via anisotropic curvilinear r-adaptivity. Unsteady flows are treated in a space-time manner. A major difficulty encountered in MDG-ICE calculations is frequent element degeneration. To address this issue, we introduce an optimization solver based on the Levenberg-Marquardt algorithm that features an anisotropic, locally adaptive penalty method. Strategies to approximate strong viscous shocks as inviscid discontinuities are also discussed. Finally, we present results for two- and three-dimensional test cases, including hypersonic flows over blunt bodies, which are extremely challenging for conventional schemes on simplicial grids. Even without artificial dissipation, the computed solutions are free from spurious oscillations and yield highly symmetric surface heat-flux profiles.

Eric Ching

US Naval Research Laboratory
eric.j.ching.civ@us.navy.mil

Andrew Kercher, Andrew Corrigan
Naval Research Laboratory
andrew.kercher@nrl.navy.mil, andrew.corrigan@nrl.navy.mil

MS137

High-Order Limiting Methods Using Maximum Principle Bounds Derived from the Boltzmann Equation

The use of limiting methods for high-order numerical approximations of hyperbolic conservation laws generally requires defining an admissible region/bounds for the solution. This talk will present a novel approach for computing solution bounds and limiting for the compressible Euler/Navier-Stokes equations through the kinetic representation provided by the Boltzmann equation, which allows for extending limiters designed for linear advection directly to macroscopic governing equations which can be derived from an underlying kinetic representation. Given an arbitrary set of solution values to compute bounds over (e.g., numerical stencil), the proposed approach yields an analytic expression for the admissible region of particle distribution function values, which may be numerically integrated to yield a set of bounds for the density, momentum, and total energy. These solution bounds are shown to preserve positivity of density/pressure/internal energy and, when paired with a limiting technique, can robustly resolve strong discontinuities while recovering high-order accuracy in smooth regions without any ad hoc corrections (e.g., relaxing the bounds). This approach is demonstrated in the context of explicit unstructured high-order discontinuous Galerkin-type schemes for a variety of difficult problems including cases with extreme shocks and shock-vortex interactions.

Tarik Dzanic
Lawrence Livermore National Laboratory
dzanic1@llnl.gov

Luigi Martinelli
Princeton University
martinel@princeton.edu

MS137

Adaptive Mesh Refinement and Viscosity Regularization for Discontinuous Galerkin Approximation of Hypersonic Flows

We present an optimal transport approach for mesh adaptivity and shock capturing of compressible flows. Shock capturing is based on a viscosity regularization of the governing equations by introducing an artificial viscosity field as solution of the modified Helmholtz equation. Mesh adaptation is based on the optimal transport theory by formulating a mesh mapping as solution of Monge-Ampre equation. The marriage of optimal transport and viscosity regularization for compressible flows leads to a coupled system of the compressible Euler/Navier-Stokes equations, the Helmholtz equation, and the Monge-Ampre equation. We propose an iterative procedure to solve the coupled system in a sequential fashion using homotopy continuation to minimize the amount of artificial viscosity while enforcing positivity-preserving and smoothness constraints on the numerical solution. We explore various mesh monitor functions for computing r-adaptive meshes in order to reduce the amount of artificial dissipation and improve the accuracy of the numerical solution. The hybridizable discontinuous Galerkin method is used for the spatial discretization of the governing equations to obtain high-order accurate solutions. Extensive numerical results are presented to demonstrate the approach on a wide range of hypersonic flows.

Cuong Nguyen
Massachusetts Institute of Technology
cuongng@mit.edu

MS138

Learning the Discrepancy Between Low and High-Fidelity Models to Improve Parameter Estimation for Cardiovascular Simulations

Cardiovascular models offer valuable insight for surgical planning but are still faced with skepticism in the clinical community due to their deterministic nature. However, recent advancements have begun to shift towards uncertainty-aware modeling, including forward propagation, parameter estimation, and end-to-end, data-to-prediction paradigms. In this study, we focus on multi-fidelity inverse problems and analyze the discrepancy between high-fidelity and low-fidelity quantities of interest to improve Bayesian inversion. We propose to learn the discrepancy between low- and high-fidelity models either in the form of a lower dimensional surrogate or as a density estimated through normalizing flow. The derived surrogate model or probability density is then leveraged to significantly reduce the computational cost of evaluating the posterior distribution of the input parameters. We compare our method with the traditional solution of inverse problems based entirely on the high-fidelity model, both in terms of accuracy and overall computational cost. We further demonstrate our approach on test cases of increasing complexity, from lumped-parameter Windkessel models to patient-specific three-dimensional anatomies. Supported by Yansouni Family Stanford Graduate Fellowship and NSF grant 2105345. Computing resources were provided by Stanford Research Computing Center.

Chloe Choi, Andrea Zanoni
Stanford University
chloe1@stanford.edu, azanoni@stanford.edu

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

Alison L. Marsden
Stanford University
amarsden@stanford.edu

MS138

Bayesian Multimodel Inference to Account for Model Uncertainty in Systems Biology

When modeling biological systems, it is often possible to formulate a set of related mathematical models that vary in the simplifying assumptions used to represent the system mathematically. One should account for the uncertainties associated with these assumptions when making predictions. Model selection and discrepancy modeling allow one to choose the best model and account for model uncertainty by considering the mismatch between data and predictions. However, given the limited data in systems biology, these approaches may lead to biases and misrepresentations of

uncertainty due to selecting a single model. Alternatively, multimodel inference (MMI) combines predictions from all specified models to avoid conditioning on a single model. Here, we outline several methods for Bayesian MMI in systems biology. We compare pseudo-Bayesian model averaging, stacking of predictive densities, and Bayesian model averaging to construct robust predictors using the entire set of user-specified models. Further, we demonstrate MMI with examples of mitogen-activated protein kinase signaling, AMP-activated protein kinase signaling, and glucose-driven insulin secretion in diabetes. These examples highlight how MMI increases certainty and predictive quality when multiple models are specified. Our results show how combining information from multiple models is a promising approach to handling model uncertainty within systems biology.

Nathaniel J. Linden-Santangeli, Padmini Rangamani, Boris Kramer
University of California San Diego
nlinden@ucsd.edu, prangamani@ucsd.edu, bmkramer@ucsd.edu

MS138

Multi-Fidelity Uncertainty Quantification for Homogenization Problems in Crystal Plasticity Finite Element Models

Quantifying uncertainty (UQ) in materials science is a challenging computational task due to the multi-scale characteristics of polycrystalline alloys and the inherent randomness of microstructures. This UQ problem requires solving numerous forward PDE problems to consider the random geometry and various constitutive models in solid mechanics. In this presentation, we examine the application of multi-level Monte Carlo and multi-index Monte Carlo methods in crystal plasticity finite element methods, focusing on different mesh resolutions and constitutive models (including phenomenological and dislocation-density-based models) to estimate and quantify uncertainty in the stress-strain curve. Computational results indicate a significant reduction in computational cost, which could facilitate the development of multi-scale materials digital twins for future applications.

Anh Tran, Pieterjan Robbe, Theron M. Rodgers
Sandia National Laboratories
anhtran@sandia.gov, pmrobbe@sandia.gov, trodger@sandia.gov

Hojun Lim
Sandia National Labs
hnlim@sandia.gov

MS138

Multifidelity Uq for Laser-Ignition Reliability Analysis

In a rocket combustor, hydrodynamic ejection of a laser-induced spark placed in a co-flowing fuel-oxidizer jet can facilitate re-ignition capability throughout a mission. The present work aims to understand the various scenarios in which successful ignition may occur and builds probability distributions of ignition itself and associated post-ignition pressure rise in the rocket chamber. In uncertainty space, this is a high-dimensional problem, and the inherited run-to-run variabilities arise from uncertainty in the deposited laser characteristics, and the stochastic nature of instantaneous turbulence. To this extent, a reduced order low-fidelity surrogate model is constructed that achieves a compression ratio of two orders of magnitude compared to the accurate high-fidelity simulations. This is achieved by a reduced chemistry mechanism and reduction in the simulation scale-resolution. Further, the surrogate model is constructed by means of an inverse problem in which deterministic processes from high-fidelity pilot cases are used to estimate the parameters in the low-fidelity space that map to high-fidelity counterparts. Thereafter, multi-fidelity Monte Carlo sampling is employed to obtain realizations of the uncertainty space. Finally, we present analysis of ignition likelihood and sensitivities.

Tony Zahtila
Stanford University
tzahtila@stanford.edu

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

Murray C. Cutforth
Stanford University
mcc4@stanford.edu

MS139

Einsum Trees

A typical computational backend for the einsum language comprises two parts: First, a contraction path algorithm that breaks down an einsum expression into a sequence of binary tensor contractions. Contraction path algorithms aim to minimize the number of floating point operations (flops) for executing the tensor expression. Second, the execution of the binary contractions. For efficient binary contractions, the data layout of the tensors needs to be optimized. So far, computing contraction paths and optimizing the data layout for single, that is, local, binary tensor contractions have been studied in isolation. Here, we introduce einsum trees, an intermediate representation for globally optimizing the data layout for a given contraction path. Thereby, optimizing the overall execution times for einsum expressions. We illustrate the effectiveness of the approach on a state-of-the-art Arm server processor, an x86 desktop system, and a system on a chip (SoC) that is targeting the mobile market.

Alexander Breuer
University of California, San Diego
alex.breuer@uni-jena.de

MS139

Robust Low-Rank Tensor Decomposition with the L2 Criterion

The growing prevalence of tensor data, or multiway arrays, in science and engineering applications motivates the need for tensor decompositions that are robust against outliers. In this work, we present a robust Tucker decomposition estimator based on the L2 criterion, called the Tucker-L2E. Our numerical experiments demonstrate that Tucker-L2E has empirically stronger recovery performance in more challenging high-rank scenarios compared with existing alternatives. The appropriate Tucker-rank can be selected in a data-driven manner with cross-validation or hold-out validation. The practical effectiveness of Tucker-L2E is validated on real data applications in fMRI tensor denoising, PARAFAC analysis of fluorescence data, and

feature extraction for classification of corrupted images.

Qiang Heng
UCLA
qheng@g.ucla.edu

Eric Chi
Rice University
echi@rice.edu

Yufeng Liu
UNC Chapel Hill
yfliu@email.unc.edu

MS139

Tensor-network for Probabilistic Modeling

We present a new perspective on randomized linear algebra, showcasing its usage in estimating a density as a tensor-network, without the curse of dimensionality, and without the use of optimization techniques. Moreover, we illustrate how this can be used for solving high-dimensional PDEs.

Yuehaw Khoo
University of Chicago
ykhoo@uchicago.edu

MS139

Tensor-Based Approach to Synchronization

Synchronization is the problem of fusing noisy local relative information in a network to obtain reliable global information. It applies to structure from motion in computer vision, unmanned navigation in robotics, molecular structuring in cryo-electron microscopy, among others. In this talk I will present a novel mathematical approach to synchronization, based on a formulation involving tensors and a low-rank tensor decomposition algorithm. Results on benchmark datasets show that the method increases the accuracy of camera location estimation in structure from motion over standard approaches. Joint work with Daniel Miao and Gilad Lerman at U Minnesota.

Joe Kileel
University of Texas at Austin, U.S.
jkileel@math.utexas.edu

MS139

Adaptive Hard Thresholding Methods for Low-Rank Tensor Recovery

Tensor-structured random dimension reduction maps have multiple advantages: they can be applied quickly to the data with multi-modal structure, they can preserve vital structural properties of the data such as low-rankness and also require much less memory than generic unstructured random maps. However, the tensor structure makes the distributions of the entries in the compression matrices inherently heavy-tailed, which raises complications related to both practical use and theoretical analysis of the associated algorithms. In this talk, I will discuss several iterative algorithms for the efficient recovery of CP and Tucker low-rank tensors from a few linear measurements that aim to cover, in particular, memory-efficient linear maps with tensorial structure.

Elizaveta Rebrova
Princeton University
elre@princeton.edu

MS140

Quasi-Trefftz Methods with Applications to Maxwell's Equations

Trefftz methods are a class of high-order numerical schemes for solving problems modeled by partial differential equations. They construct the numerical solution starting from functions specifically tailored to the problem considered, with test and trial functions that are solutions of the PDE under consideration on each mesh element. Notably, Trefftz methods achieve a level of accurancy with reduced computational cost compared to traditional approaches like finite elements or standard Discontinuous Galerkin methods. However, their application is limited to linear, homogeneous and piecewise-constant coefficients problems, as exact solutions are often unavailable in more complex cases. To address this limitation, quasi-Trefftz methods use elementwise approximate solutions of the PDE, enabling the analysis of a wider range of problems while maintaining the good properties of Trefftz schemes. Recent studies have demonstrated the convergence and stability of quasi-Trefftz methods for some scalar problems, including the diffusion-advection-reaction equation and acoustic problems. In this work, we extend these techniques to problems governed by vector-valued PDEs, with a focus on variable-coefficients Maxwell's equations in the frequency domain, addressing the questions of construction of the quasi-Trefftz space and converge of the method. We present recent results and novel insights aimed at expanding the applications of quasi-Trefftz methods to complex vector-valued PDEs.

Ilaria Fontana, Lise-Marie Imbert-Gerard
University of Arizona
ifontana@arizona.edu, lmig@math.arizona.edu

MS140

Quasi-Trefftz Method for Solving Aeroacoustic Problem

This study will present a method which is relevant to several Airbus applications, especially for simulating engine noise propagation through airflow to model ground and cabin noise. Variational Trefftz methods are discontinuous Galerkin methods whose basis functions are local solutions of the PDE under consideration. In the context of homogeneous problems, analytical solutions, such as plane waves, are available. Aeroacoustic models involve equations whose physical characteristics depend on the spatial variables. In general, this PDE system cannot be solved analytically. A natural idea is to resort to basis functions that are approximate solutions of the considered PDE, i.e., quasi-Trefftz methods. In this talk we will see how to build three types of quasi-Trefftz function for the heterogeneous Helmholtz equation with constant flow : Amplitude-based Generalized Plane Waves, Phase-based Generalized Plane Waves and polynomials [Imbert-Gerard, Sylvand, (2022). Three types of quasi-Trefftz functions for the 3D convected Helmholtz equation: construction and approximation properties.]. We will also verify some approximation properties of these functions. Then we present a quasi-Trefftz Discontinuous Galerkin variational formulation inspired by hyperbolic system formulation. We will finish with numerical results comparing the solution obtained with this quasi-Trefftz method to those from the

traditional Discontinuous Galerkin method.

Andrea Lagardere
Airbus
MAKUTU - Inria
andrea.lagardere@inria.fr

MS140

Polynomial Quasi-Trefftz DG for PDEs with Smooth Coefficients: Elliptic Problems

Trefftz schemes are high-order Galerkin methods whose discrete spaces are made of elementwise exact solutions of the underlying PDE. Trefftz basis functions can be easily computed for many PDEs that are linear, homogeneous, and have piecewise-constant coefficients. However, if the equation has variable coefficients, exact solutions are generally unavailable. Quasi-Trefftz methods overcome this limitation relying on elementwise "approximate solutions" of the PDE, in the sense of Taylor polynomials. We define polynomial quasi-Trefftz spaces for general linear PDEs with smooth coefficients and source term, describe their approximation properties and, under a non-degeneracy condition, provide a simple algorithm to compute a basis. We then focus on a quasi-Trefftz DG method for variable-coefficient elliptic diffusion-advection-reaction problems, showing stability and high-order convergence of the scheme. The main advantage over standard DG schemes is the higher accuracy for comparable numbers of degrees of freedom. For non-homogeneous problems with piecewise-smooth source term we propose to construct a local quasi-Trefftz particular solution and then solve for the difference. Numerical experiments in 2 and 3 space dimensions show the excellent properties of the method both in diffusion-dominated and advection-dominated problems.

Chiara Perinati, Andrea Moiola
University of Pavia
chiara.perinati01@universitadipavia.it,
andrea.moiola@unipv.it

Lise-Marie Imbert-Gerard
University of Arizona
lmig@math.arizona.edu

Paul Stocker
University of Vienna
paul.stocker@univie.ac.at

MS140

Embedded Trefftz Discontinuous Galerkin Methods

We introduce a new numerical approach aimed at reducing the system size in discontinuous Galerkin (DG) methods, drawing inspiration from Trefftz methods. The central idea of Trefftz methods is to construct optimal discretization spaces that minimize the number of unknowns while retaining optimal approximation properties. The embedded Trefftz discontinuous Galerkin method, is the Galerkin projection of an underlying discontinuous Galerkin method onto a subspace of Trefftz-type. The subspace can be described in a very general way and to obtain it no Trefftz functions have to be calculated explicitly, instead the corresponding embedding operator is constructed. In the simplest cases the method recovers established Trefftz discontinuous Galerkin methods. But the approach allows to conveniently extend to general cases, including inhomogeneous sources and non-constant coefficient differential operators. We introduce the method, discuss implementational aspects and explore its potential on a set of standard PDE problems. As typical for Trefftz-DG methods, we observe a severe reduction of the globally coupled unknowns when compared to standard discontinuous Galerkin methods, reducing the corresponding computing time significantly.

Paul Stocker
University of Vienna
paul.stocker@univie.ac.at

Christoph Lehrenfeld
University of Göttingen
lehrenfeld@math.uni-goettingen.de

MS141

A Blob Method for Mean Field Control

In the present work, we develop a novel particle method for a general class of mean field control problems with source and terminal constraints. Specific examples of the problems we consider include the dynamic formulation of the p-Wasserstein metric, optimal transport around an obstacle, and measure transport subject to acceleration controls. Unlike existing numerical approaches, our particle method is meshfree and does not require global knowledge of an underlying cost function or of the terminal constraint. A key feature of our approach is a novel way of enforcing the terminal constraint via a soft, nonlocal approximation inspired by recent work on blob methods for diffusion equations.

Katy Craig
University of California, Santa Barbara
kcraig@ucsb.edu

MS141

Unsupervised Solution Operator Learning for Mean-Field Games

Recent advances in deep learning have led to many innovative methods that solve high dimensional mean-field games (MFG) accurately and efficiently. These methods, however, are restricted to solving single-instance MFG and demands extensive computational time per instance, limiting practicality. In this talk, I will present our recent framework to learn the MFG solution operator. Our model takes a MFG instances as input and output their solutions with one forward pass. To ensure the proposed parametrization is well-suited for operator learning, we introduce and prove the notion of sampling invariance for our model, establishing its convergence to a continuous operator in the sampling limit. Our method features two key advantages. First, it is discretization-free, making it particularly suitable for learning operators of high-dimensional MFGs. Secondly, it can be trained without the need for access to supervised labels, significantly reducing the computational overhead associated with creating training datasets in existing operator learning methods. We test our framework on synthetic and realistic datasets with varying complexity and dimensionality to substantiate its robustness.

Rongjie Lai
Purdue University

lairj@purdue.edu

MS141

Score-Based Generative Models Through the Lens of Wasserstein Proximal Operators

In this presentation, I will discuss the essence of score-based generative models (SGMs) as entropically regularized Wasserstein proximal operators (WPO) for cross-entropy, elucidating this connection through mean-field games (MFG). The unique structure of SGM-MFG allows the HJB equation alone to characterize SGMs, demonstrated to be equivalent to an uncontrolled Fokker-Planck equation via a Cole-Hopf transform. Furthermore, leveraging the mathematical framework, we introduce an interpretable kernel-based model for the score functions, enhancing the performance of SGMs in terms of training samples and training time. The mathematical formulation of the new kernel-based models, in conjunction with the utilization of the terminal condition of the MFG, unveils novel insights into the manifold learning and generalization properties of SGMs.

Siting Liu
University of California, Los Angeles, U.S.
siting6@math.ucla.edu

MS141

Optimal Transport for Inverse Problems in Ocean Electromagnetics

Internal waves are crucial to understanding ocean mixing, energy distribution, and their impact on marine environments. Although challenging to observe directly due to their subsurface nature, the magnetic fields induced by these waves offer a practical method for monitoring them. By solving the inverse problem of inferring velocity fields from magnetic measurements, we can obtain valuable insights into the complex dynamics of internal waves. Recently, optimal transport (OT) has demonstrated significant advantages in solving inverse problems, particularly in its robustness to noise. In this study, we apply the OT approach to an inverse problem aimed at recovering velocity fields from electromagnetic (EM) field measurements generated by internal ocean waves, where the measurements are taken over a part of the domain or across the entire domain at a specific time T. The model used is the linearized induction equation around the Earths magnetic field. To validate our approach, we will use the Podney benchmark [1], comparing the results obtained through the OT approach with those from the classical L2-norm approach and the analytical solution provided by Podney. The OT method will be tested in scenarios with noisy data, highlighting its potential for robust and accurate inverse problem solutions in geophysical applications. [1] W. Podney, "Electromagnetic fields generated by ocean waves," Journal of Geophysical Research, vol. 80, no. 21, pp. 29772990, 1975.

Christina Frederick, Yacine Mokhtari
New Jersey Institute of Technology
christin@njit.edu, yacine.mokhtari@njit.edu

Yunan Yang
Cornell University
yunan.yang@cornell.edu

MS142

Martingale Neural Networks for BVPs and Eigenvalue Problems of PDEs and Stochastic Optimal Controls

In this talk, we present martingale-based neural networks for solving high-dimensional PDE's BVPs and eigenvalue problems (EVPs) as well Hamilton-Jacobi-Bellman (HJB) equations where no explicit expression is needed for the infimum of the Hamiltonian, and stochastic optimal control problems (SOCPs) with controls on both drift and volatility. For the BVPs and EVPs, the PDEs are reformulated using Varadhan's martingale problem and the neural network solution is trained to meet the martingale properties. For the SOCPs, we reformulate the HJB equations for the value function by training two neural networks, one for the value function and one for the optimal control with the help of two stochastic processes- a Hamiltonian process and a cost process. The control and value networks are trained so the Hamiltonian process is minimized to satisfy the minimum principle of a feedback SOCP, and the cost process becomes a martingale. To enforce the martingale property for the cost process, we employ an adversarial network and construct a loss function characterizing the projection property of the conditional expectation condition of the martingale. Numerical results show that the proposed SOC-MartNet is effective and efficient for solving 2000 dim HJB-type equations and SOCPs in a small number of epochs (less than 20) or stochastic gradient method iterations (less than 2000) for the training.

Wei Cai
Department of Mathematics
Southern Methodist University
cai@smu.edu

MS142

Instabilities in Neural Autoregressive Models of Complex Systems: Evidence from Linear Stability Analysis and Random Matrix Theory

We present a general framework, based on linearization and eigenmode decomposition, to analyze deep neural networks (DNNs) when applied to nonlinear PDEs, for example, in terms of stability and convergence. There is growing interest in using DNNs for building data-driven solvers of nonlinear PDEs. Many recent studies have shown promising results, at least for short-term forecasts, for canonical PDEs, and even for real-world data, such as atmospheric reanalysis datasets. However, there are also challenges such as long-term instabilities (drifts and blow-ups) and lack of convergence (decline in accuracy as time step is reduced). These challenges and other issues have been difficult to address due to the lack of a rigorous and general (i.e., system- and architecture-agnostic) framework for analyzing DNNs when applied to PDEs. Here, we present such a system- and architecture-agnostic framework that combines concepts from numerical analysis and deep learning theory. The framework is based on eigenvalue analysis of the linearized DNN. The eigenvalues provide quantitative information on the stability of the DNN-based integration, as well as the amount of spectral bias, and thus the convergence property. We show that this framework can guide the development of stable, convergent neural PDE-solvers for any system and architecture, which has wide-ranging applications on nonlinear geophysics, e.g., to build data-

driven weather/climate models.

Ashesh Chattopadhyay
UC Santa Cruz
aschatto@ucsc.edu

MS142

Uncertainty Propagation in Input-Output Maps Defined by Neural Operators

We propose a method for studying uncertainty propagation in feed-forward neural network architectures. In particular we investigate the case of a multi-layer perceptron (MLP) neural network subjected to small random perturbations in either their input vectors or network parameters, i.e. their weights or biases. By utilizing perturbation expansions, we derive analytical expressions for the probability distribution functions (PDF) of the individual components of the output vectors, their statistical moments, and a Gaussian copula representation of the joint PDF of the output vector. To validate the derived analytical expressions, we conduct Monte Carlo simulations of an MLP neural network representing a nonlinear integro-differential operator between two polynomial function spaces. Implementations of this method towards random neural operators are presented and discussed.

Jeremy Diamzon
UC Santa Cruz
jdiamzon@ucsc.edu

Daniele Venturi
Department of Applied Mathematics
University of California Santa Cruz
venturi@ucsc.edu

MS142

Learning Kinetic Monte Carlo Using Statistics-Informed Neural Network

The statistic-informed neural network (SINN) is a machine learning framework developed to learn and reproduce general stochastic processes. This approach can be useful as a surrogate model in reproducing kinetic Monte Carlo (KMC) simulations, which are widely employed in both chemical and biological studies. Although the atomistic configuration of the surface is required to simulate elementary surface chemical reactions using KMC, the information one eventually wants to obtain from KMC simulations is aggregated information such as surface coverages of chemical species and the turnover rate of the overall catalytic reaction. In this talk, we apply SINN to learn and generate time trajectories of surface coverages. We will present numerical examples, including the Langmuir adsorption model and the lattice Lotka-Volterra model.

Zihan Xu, Patrick Park
University of California Merced
zxu64@ucmerced.edu, ppark12@ucmerced.edu

Changho Kim
University of California, Merced
Applied Mathematics
ckim103@ucmerced.edu

Andrew J. Nonaka, Yuanran Zhu
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, yzhu4@lbl.gov

MS142

Residual Stress Prediction of Multi-Principal-Element Alloys with Short Range Order on Continuum Scale by Generative Model

Multi-principal-element alloys (MPEAs) are a class of metallic materials composed of at least three primary elements, which have excellent mechanical properties, such as high strength, corrosion resistance, etc. However, both modeling and simulation on continuum scale are extremely difficult tasks for MPEAs, because of nearly random atomic occupancy with short-range order (SRO) and vast compositional space. In this presentation, we propose a generative framework, which combines a conditional variational autoencoder (cVAE) and two concentration/SRO smoothers, to predict the residual stress of MPEAs as a function of atomic concentration and short-range order on continuum scale. A large dataset of samples with initial random atomic occupancy, WC parameters and their virial stress after equilibrium is generated by hybrid MC/MD simulation, and the block-average is used to obtain the concentration, averaged WC parameters and residual stress on continuum scale. The results show this generative model can predict the distribution of residual stress with the mean absolute percentage error below 5%, and the microstructure smoother significantly increases the regularity of the dataset, improving the training efficiency and testing accuracy. Our method demonstrates the excellent ability of generative model in the surrogate modeling and simulation of MPEAs on continuum scale.

Ningyu Yan
the Hong Kong University of Science and Technology
nyanac@connect.ust.hk

Yejun Gu
Institute of High Performance Computing
yejun_gu@ihpc.a-star.edu.sg

Yang Xiang
Hong Kong University of Science and Technology
maxiang@ust.hk

MS143

Randomized Iterative Methods for Deblurring and Tomography

Large scale inverse problems often require massive amounts of memory. To amelioriate this, we find randomized numerical linear algebra (rand NLA), and in particular matrix sketching, reduce this burden of memory constraints. In this talk, we will show the impact of implementing rand NLA strategies on uncertainty quantification in large-scale multiresolution tomography and deblurring problems.

Jimmie Adriazola
Arizona State University
jimmie.adriazola@asu.edu

MS143

Reconstruction of Extended Regions in EIT with a Generalized Robin Transmission Condition.

In this talk, we will discuss an application of the Regularized Factorization Method (RegFM) to a problem coming from Electrical Impedance Tomography (EIT) with a

second-order Robin condition. This method falls under the category of qualitative methods for inverse problems. Qualitative methods are used in non-destructive testing where physical measurements on the surface of an object are used to infer the interior structure. The Robin condition on this boundary asymptotically models delamination. We assume that the Dirichlet-to-Neumann (DtN) mapping is given on the exterior boundary and will be used to recover an unknown, extended region. Using Cauchy data as physical measurements, we can determine if all of the coefficients from the Robin condition are real- or complex-valued. We study these two cases separately and show how RegFM can be used to detect whether delamination has occurred and recover the damaged subregion. Numerical examples will be presented for both cases in two dimensions in the unit circle.

Govanni Granados
The University of North Carolina at Chapel Hill
ggranad@unc.edu

MS143

A Direct Reconstruction Method for Radiating Sources in Maxwells Equations with Single-Frequency Data

This work presents a fast and robust numerical method for reconstructing point-like sources in the time-harmonic Maxwell's equations given Cauchy data at a fixed frequency. This is an electromagnetic inverse source problem with broad applications, such as antenna synthesis and design, medical imaging, and pollution source tracing. We introduce new imaging functions and a computational algorithm to determine the number of point sources, their locations, and associated moment vectors, even when these vectors have notably different magnitudes. The number of sources and locations are estimated using significant peaks of the imaging functions, and the moment vectors are computed via explicitly simple formulas. The theoretical analysis and stability of the imaging functions are investigated, where the main challenge lies in analyzing the behavior of the dot products between the columns of the imaginary part of the Green's tensor and the unknown moment vectors. Additionally, we extend our method to reconstruct small-volume sources using an asymptotic expansion of their radiated electric field. We provide numerical examples in three dimensions to demonstrate the performance of our method.

Thu Le
Kansas State University
tle38@wisc.edu

MS143

Exploring Low Rank Structure for Inverse Scattering

Inverse problems play important roles in various applications, including target identification, non-destructive testing, and parameter estimation. Particularly challenging is the inverse scattering problem in inhomogeneous media, which aims to estimate unknowns based on available measurement data. Given its inherently ill-posed nature, our aim is to address this challenge by exploring the underlying low rank structure that is capable of handling potentially noisy and large-scale measurement data. The unknown is solved in a low-dimensional space comprising disk prolate spheroidal wave functions, which are computed efficiently via a Sturm-Liouville problem. The low rank structure leads to reliable numerical algorithms that demonstrate increasing stability and dimensionality reduction, in the presence of noisy and large-scale measurement data. A stability estimate is proved by leveraging the interplay between a Fourier integral operator and a Sturm-Liouville differential operator. The talk concludes with a discussion of current and future work on exploring low-rank structures.

Shixu Meng
Virginia Tech
sgl22@vt.edu

MS143

Parameter Selection by Gcv and a χ^2 Test Within Iterative Methods for ℓ_1-Regularized Inverse Problems

ℓ_1 regularization is used to preserve edges or enforce sparsity in a solution to an inverse problem. We investigate the Split Bregman and the Majorization-Minimization iterative methods that turn this non-smooth minimization problem into a sequence of steps that include solving an ℓ_2-regularized minimization problem. We consider selecting the regularization parameter in the inner generalized Tikhonov regularization problems that occur at each iteration in these ℓ_1 iterative methods. The generalized cross validation and χ^2 degrees of freedom methods are extended to these inner problems. In particular, for the χ^2 method this includes extending the χ^2 result for problems in which the regularization operator has more rows than columns and showing how to use the $A-$weighted generalized inverse to estimate prior information at each inner iteration. Numerical experiments for image deblurring problems demonstrate that it is more effective to select the regularization parameter automatically within the iterative schemes than to keep it fixed for all iterations. Moreover, an appropriate regularization parameter can be estimated in the early iterations and used fixed to convergence.

Brian F. Sweeney
Arizona State University
bfsweene@asu.edu

Rosemary A. Renaut
Arizona State University
School of Mathematical and Statistical Sciences
renaut@asu.edu

Malena I. Espanol
Arizona State University
malena.espanol@asu.edu

MS144

Quantum Recurrent Neural Networks for Nonlinear Partial Differential Equations

Nonlinear partial differential equations are fundamental in modeling complex phenomena across various fields, yet they pose significant challenges due to their computational complexity in higher dimensions. This study leverages Quantum Recurrent Neural Networks that integrate Variational Quantum Circuits into Gated Recurrent Units and Long-Short Term Memory networks. The algorithms harness the power of qubits and quantum entanglement, allowing the representation and processing of vast amounts of information simultaneously, which enhances their ability to explore larger solution spaces and mitigate the curse of dimensionality. The algorithms are implemented on data-driven PDEs including, Hamilton-Jacobi-Bellman equa-

tion, Burgers' equation, and reaction-diffusion equations in pattern formation, investigating their potential advantages over classical recurrent neural networks in handling nonlinearity, high-dimensional spaces, and ensuring stable solutions.

Yuan Chen, Abdul Khaliq
Middle Tennessee State University
yc3y@mtmail.mtsu.edu, abdul.khaliq@mtsu.edu

Khaled Furati
Department of Mathematics
King Fahd University of Petroleum & Minerals
kmfurati@kfupm.edu.sa

MS144

Preconditioning Strategies for a Nested Primal-Dual Method for Image Deblurring

Variational models for image deblurring problems typically consist of a smooth term and a potentially nonsmooth convex term. A common approach to solving these problems is using proximal gradient methods. Strategies such as variable metric methods have been introduced in the literature to accelerate the convergence of these first-order iterative algorithms. In this talk, we show that, for image deblurring problems, the variable metric strategy can be reinterpreted as a right preconditioning method. Consequently, we explore an inexact left-preconditioned version of the same proximal gradient method. We prove the convergence of the new iteration to the minimum of a variational model where the norm of the data fidelity term depends on the preconditioner. The numerical results show that left and right preconditioning are comparable in terms of the number of iterations required to reach a prescribed tolerance, but left preconditioning needs much less CPU time, as it involves fewer evaluations of the preconditioner matrix compared to right preconditioning. The quality of the computed solutions with left and right preconditioning are comparable. Finally, we propose some nonstationary sequences of preconditioners that allow for fast and stable convergence to solve the variational problem with the classical ℓ_2-norm on the fidelity term.

Stefano Aleotti, Marco Donatelli
University of Insubria
saleotti@uninsubria.it, marco.donatelli@uninsubria.it

Rolf Krause, Rolf Krause
KAUST / UniDistance Switzerland
krause.r@gmx.net, krause.r@gmx.net

Giuseppe Scarlato
University of Insubria
scarlato giuseppe ¡gscarlato@uninsubria.it¿

MS144

A Stochastic PDE Approach for Gaussian Random Field Realizations at Multiple Scales

In this work we present an algorithmically scalable multilevel approach to generate Gaussian random field realizations that is well-suited to incorporate into multilevel sequential algorithms, such as multilevel Markov chain Monte Carlo (MCMC). In such algorithms, generating fine grid samples from a desired distribution is accelerated by leveraging coarse grid simulations to perform the majority of the parameter space exploration (or random-walk as is the case in MCMC). A drawback of the existing multilevel framework is that the multilevel Gaussian random field sampling is not scalable. For large-scale Gaussian random field simulations, e.g., from a 3D finely resolved mesh, we instead consider solving a reaction-diffusion PDE with a spatial white noise source function as the righthand side. In order to achieve matching fine and coarse realizations, we decompose the finite element white noise (source function) into complementary spaces, across multiple levels of discretization. Then, using tools from algebraic multigrid, we form a (scalable) fine grid Gaussian random field realization from a fine grid proposal by combining Gaussian random fields sampled across multiple levels of discretization. In this talk, we describe this approach, corresponding theory, and numerical results.

Hillary Fairbanks
Lawrence Livermore National Laboratory
fairbanks5@llnl.gov

Panayot Vassilevski
Portland State University
panayot@pdx.edu

MS144

Domain Decomposition for Neural Networks

Scientific machine learning (SciML) is a rapidly evolving research field that combines techniques from scientific computing and machine learning. This talk focuses on the application of domain decomposition methods to design neural network architectures and enhance neural network training, Specifically, it explores the use of domain decomposition techniques in neural network-based discretizations for solving partial differential equations with physics-informed neural networks (PINNs) and operator learning, as well as in classical machine learning tasks like semantic image segmentation using convolutional neural networks (CNNs). Computational results show that domain decomposition methods can improve efficiency both in terms of time and memory as well as enhance accuracy and robustness.

Alexander Heinlein
Delft University of Technology
Delft Institute of Applied Mathematics
a.heinlein@tudelft.nl

MS144

Integrating BDDC with Adaptive SDC Time Stepping in HPC Environments

Simulating cardiac electrophysiology on the geometrically resolved microscale leads to large scale problems, exceeding the size of homogenized bidomain simulations by a factor of more than 100. This calls for highly efficient solvers for the arising linear equation systems, which are well-integrated into both their algorithmic and their computational environment. We will cover two aspects of integrating a Balanced Domain Decomposition by Constraints (BDDC) preconditioner with the Extracellular-Membrane-Intracellular (EMI) model: algebraic adaptivity in Spectral Deferred Correction (SDC) time stepping and compressed communication within BDDC. Algebraic adaptivity is a form of spatial adaptivity with little overhead compared to mesh refinement/coarsening. Within the SDC iteration, it reduces problem size on the algebraic level based on estimated local correction size. This puts some pressure on preconditioners, which have to cope with varying problems. We illustrate how algebraic adaptivity can be

transferred to BDDC itself, retaining low overhead. BDDC requires communication between subdomains in every iteration. With a growing gap between CPU power and communication bandwidth in HPC systems, communication is increasingly contributing to overall execution times. We investigate the use of tailored lossy data compression for exchanging boundary values, and select quantization tolerances based on error estimates and work models.

Martin Weiser
Zuse Institute (ZIB)
Berlin, Germany
weiser@zib.de

Fatemeh Chegini, Thomas Steinke
Zuse Institute Berlin
chegini@zib.de, steinke@zib.de

MS145

A Spatially-Correlated Random Walk Model for Capturing Two-Point Statistics in Particle-Laden Turbulence

The multi-scale nature of turbulent particle-laden flows prevents the use of direct numerical simulations. For this reason, large-eddy simulations (LES) and Reynolds-averaged Navier-Stokes (RANS) have become essential for modern computational fluid dynamics (CFD) applications. However, these techniques need subgrid-scale models to recapture the effects of unresolved turbulence on particle dynamics. Stochastic subgrid models are commonly used because of their simplicity, low computational cost, and ability to capture one-point and two-time velocity statistics. Yet, a significant drawback is their inability to capture two-point spatial correlations due to missing relative spatial information. In this talk, we propose a framework that incorporates two-point information into a Lagrangian stochastic model by considering a multi-particle, spatially-correlated system. Initially, we utilize stochastic inference techniques to highlight the need for spatial information in the coefficients of the stochastic model. Next, we introduce a computationally-efficient model, termed the spatially-correlated random walk (SCRW), drawing connections between the SCRW, particle-pair motion theories, and Komogorov's 1941 hypotheses. Preliminary results verify the SCRW's capability to capture spatial heterogeneity and two-point velocity correlations. The talk concludes by discussing remaining challenges for the SCRW and suggesting directions to extend the SCRW to inhomogeneous anisotropic turbulence.

Max Herzog
University of Michigan
maxzog@umich.edu

Andrew Bragg
Duke University
andrew.bragg@duke.edu

Jesse Capecelatro
University of Michigan
jcaps@umich.edu

MS145

A Flexible Interface for Fluid-Particle Reaction Management in MFIX-Exa

MFIX-Exa (`https://mfix.netl.doe.gov`) is a multiphase computational fluid dynamics (CFD) code to aid in the design, optimization and scale-up of emerging technologies that advance the U.S. Department of Energys (DOE) historic clean energy goals. The code, initially developed under DOEs Exascale Computing Project by researchers at the National Energy Technology Laboratory (NETL) and Berkeley National Laboratory (LBL), is built on the AMReX software framework (`https://amrex-codes.github.io/`) and is designed to leverage modern accelerator-based compute architectures. This presentation provides a general overview of MFIX-Exa fluid and particle models, as well as the recent development of a flexible interface for managing fluid-particle chemical reactions.

Jordan Musser
National Energy Technology Laboratory (NETL)
jordan.musser@netl.doe.gov

Roberto Porcu
National Energy Technology Laboratory
roberto.porcu@netl.doe.gov

MS145

Coupling Particles to Multicomponent Lattice Boltzmann Fluids for Simulating Particle-Stabilized Emulsions

Particle-laden emulsions continue to attract interest for their use in applications such as drug delivery systems and slurries for materials processing. A particular interesting class of such emulsions are bicontinuous interfacially jammed emulsion gels (bijels) that emerge during spinodal decomposition of a binary fluid with suspended particles. Computational modeling and simulations can help gain a fundamental understanding of the interplay of phase behavior, fluid dynamics, and interfacial thermodynamics in these systems. In this talk, I will give an overview of hybrid lattice Boltzmann-Molecular Dynamics methods for particle-laden suspensions and multicomponent fluids. In particular, I will discuss distinct approaches for coupling of colloidal particles and polymeric molecules to the lattice Boltzmann fluid. I will illustrate the capabilities of these methods for simulating the formation of bijels stabilized by anisotropic particles. I will further review the fluctuating lattice Boltzmann equation for a single-component, ideal fluid, which can reproduce Brownian motion in suspensions. I will close with an overview of current developments of fluctuating lattice Boltzmann methods for multiphase and multicomponent fluids.

Ulf D. Schiller, Nikhil Karthikeyan
University of Delaware
uschill@udel.edu, nkarthi@udel.edu

MS146

Discriminative Sampling and Active Learning for Model Discovery Using Information Theory

Fisher information and Shannon entropies are crucial concepts in stochastic models from different perspectives. They can characterize unknown parameters in numerical methods by quantifying the amount of information carried in certain random variables. From an information point of view, different initial trajectories and different parts of a single trajectory contribute differently to learning a system. In this study, we firstly provide an overview of the fast-growing applications of Fisher information and Shannon entropies. Secondly, we leverage the Fisher Information Matrix (FIM) in the data-driven method of *sparse identification of nonlinear dynamics* (SINDY). We graph the

information pattern for both single trajectory and multiple initial settings, which enables us to demonstrate how information analysis can improve the sampling efficiency by increasing the model performance from more informative data. The efficacy of utilizing statistical bagging can also be explained by complementary distortions in the feature space. Thirdly, we demonstrate concrete usages of the Fisher information and entropy metrics on promoting data efficiency for an unknown dynamical system in three different cases - when only one trajectory is available, when a control parameter is available for tuning, and when multiple trajectories are available with freely chosen initials.

Yuxuan Bao
Univeristy of Washington
baoyx@uw.edu

Urban Fasel
University of Washington
u.fasel@imperial.ac.uk

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS146

Comparing Dynamical Systems Using Time-Delayed Invariant Measures

In recent years, several works have compared dynamical systems by studying the discrepancy between their invariant measures under a suitable metric or divergence on the space of probability measures. While the robustness of invariant measures to noisy measurements, uncertain initial conditions, and slow sampling makes them appealing tools for comparing dynamical systems in applications, the approach is also limited by the existence of infinitely many distinct systems all admitting the same invariant measure. To overcome this difficulty, we instead propose studying invariant measures in time-delayed coordinate systems. We present theoretical results which show that, up to a topological conjugacy, the time-delayed invariant measure can distinguish between dynamical systems, and we provide numerical examples which demonstrate the utility of our proposed method.

Jonah Botvinick, Yunan Yang
Cornell University
jrb482@cornell.edu, yunan.yang@cornell.edu

MS146

Reduced Models with Memory and Random Dynamical Systems

In the absence of sharp timescale separation, reduced models of many-degree-of-freedom dynamical systems can exhibit memory effects. The Mori-Zwanzig projection operator formalism has been used to construct such non-Markovian reduced models in a variety of settings, often with a focus on deterministic chaotic dynamical systems. In this talk, I will review a data driven approach to model reduction motivated by a discrete-time Mori-Zwanzig formalism, then discuss its extension to dynamical systems subjected to random forcing.

Kevin K. Lin
Department of Mathematics
University of Arizona
klin@math.arizona.edu

Fei Lu
Johns Hopkins University
flu15@jh.edu

MS146

On the Importance of Learning Non-Local Dynamics for Stable Data-Driven Climate Modeling

Machine learning (ML) techniques, especially neural networks (NNs), have shown promise in learning subgrid-scale parameterizations for climate models. However, a major problem with data-driven parameterizations is model instability. Current remedies are often ad-hoc and lack a theoretical foundation. We combine ML theory and climate physics to address a source of instability in NN-based parameterization. We demonstrate the importance of learning non-local dynamics using a 1D model of the quasi-biennial oscillation (QBO) with gravity wave (GW) parameterization as a testbed. While common offline metrics fail to identify shortcomings in learning non-local dynamics, we show that the concept of receptive field (RF) can identify instability a-priori. We find that NN-based parameterizations, which seem to accurately predict GW forcings, cause unstable simulations when RF is too small to capture the non-local dynamics, while NNs of the same size but large-enough RF are stable. We examine three broad classes of architectures, namely convolutional NNs, Fourier neural operators, and fully-connected NNs; the latter two have inherently large RFs. We also demonstrate that learning non-local dynamics is crucial for the stability and accuracy of a data-driven spatiotemporal emulator of the zonal wind field. Given the ubiquity of non-local dynamics in the climate system, we expect the use of effective RF, which can be computed for any NN architecture, to be important for many applications.

Hamid A. Pahalvan
Univeristy of Chicago
pahlavan@nwra.com

Pedram Hassanzadeh
The University of Chicago
pedramh@uchicago.edu

M. Joan Alexander
NorthWest Research Associates
alexand@nwra.com

MS146

Promoting Stable Dynamics by Learning the Invariant Measure of Chaotic Systems

Learning dynamics from dissipative chaotic systems is notoriously difficult due to their inherent instability, as formalized by their positive Lyapunov exponents, which exponentially amplify errors in the learned dynamics. However, many of these systems exhibit ergodicity and an attractor: a compact and highly complex manifold, to which trajectories converge in finite-time, that supports an invariant measure, i.e., a probability distribution that is invariant under the action of the dynamics, which dictates the long-term statistical behavior of the system. In this work, we leverage this structure to propose a new framework that targets learning the invariant measure as well as the dynamics, in contrast with typical methods that only target the misfit between trajectories, which often leads to di-

vergence as the trajectories' length increases. We use our framework to propose a tractable and sample efficient objective that can be used with any existing learning objectives. Our Dynamics Stable Learning by Invariant Measure (DySLIM) objective enables model training that achieves better point-wise tracking and long-term statistical accuracy relative to other learning objectives. By targeting the distribution with a scalable regularization term, we hope that this approach can be extended to more complex systems exhibiting slowly-variant distributions, such as weather and climate models.

Yair Schiff
Cornell University
yairschiff@cs.cornell.edu

Zhong Yi Wan, Jeff Parker, Stephan Hoyer, Fei Sha, Leonardo Zepeda-Núñez
Google
wanzy@google.com, parkerjeff@google.com, shoyer@google.com, fsha@google.com, lzepeda-nunez@google.com

MS147

Generative Learning and Bayesian Conditional Diffusion Models for Forecasting and Parametric Analysis of Complex Spatiotemporal Turbulence Dynamics

We will present a generative modeling framework that addresses two pivotal challenges in high-dimensional systems: accurately capturing their effective dynamics and enabling parametric sampling for diverse applications. By focusing on high-dimensional systems, the framework identifies and learns their underlying dynamics and statistical properties, achieving significant improvements in computational efficiency and accuracy. The parametric sampling capability extends the frameworks utility by generating realistic system behaviors across a broad range of parameter settings. This enables versatile tasks such as forecasting, reconstruction, optimization, and exploration of system dynamics. Our framework demonstrates exceptional robustness and adaptability, making it well-suited for handling the complexities of high-dimensional systems governed by intricate dynamics. This represents a significant step forward in bridging generative modeling with practical scientific and engineering applications.

Han Gao
Harvard University
hgao1@seas.harvard.edu

Sebastian Kaltenbach
Technical University of Munich
skaltenbach@seas.harvard.edu

Petros Koumoutsakos
Harvard University
petros@seas.harvard.edu

MS147

Basis-to-Basis Operator Learning: A Paradigm for Scalable and Interpretable Operator Learning on Hilbert Spaces

We present Basis-to-Basis (B2B) operator learning, a framework for learning nonlinear operators between Hilbert spaces of functions. Using the theory of function encoders, B2B learns a set of basis functions parameterized by deep neural networks to span Hilbert spaces. This allows functions to be represented as linear combinations of these learned basis functions with coefficients efficiently estimated using least-squares. By parameterizing input and output spaces of potentially nonlinear operators using function encoders, B2B learns operators as mappings between basis function coefficients. This approach overcomes limitations of existing methods, such as the need for fixed sensor locations in data representation, as seen in deep operator networks (DeepONet) and Fourier neural operators (FNO), two state-of-the-art operator learning architectures. Our framework offers several advantages: interpretability and amenability to analysis, adaptability to various irregular problem domains, and high generalization capabilities due to a structured representation of input and output domains. We demonstrate the efficacy of B2B in modeling solutions to linear and nonlinear PDEs and provide a comparative analysis with existing neural operator methodologies, focusing on the generalization ability to unseen test samples and the sharpness of the loss landscape. This work contributes a scalable operator learning framework with broad applications in scientific computing and machine learning.

Adam J. Thorpe, Tyler Ingebrand
University of Texas at Austin
adam.thorpe@austin.utexas.edu,
tyleringebrand@utexas.edu

Somdatta Goswami
Johns Hopkins University
sgoswam4@jhu.edu

Krishna Kumar
University of Texas at Austin
krishnak@utexas.edu

Ufuk Topcu
University of Texas
utopcu@utexas.edu

MS148

Truncating FCI Using the Mathematics of Complexity for Solving the Electronic Structure and Nuclear Structure Schrdinger Equations

An important goal of contemporary electronic and nuclear structure is to be able obtain a desired accuracy from approximate solutions to the Schrdinger equation in the least time. The full configuration interaction (FCI) method provides the best possible wave-function within a basis set. The approach prosed here involves selecting a subset of the solutions used in FCI (typically a subset of Hartree-Fock solutions within a basis set). Results from Griebel and others [1-4] in the mathematics of complexity literature show that if the analytic solutions to a given Schrdinger equation have mixed bound derivatives then it is possible to obtain FCI accuracy with polynomial scaling . The theorems from Griebel and others [1-4] indicate which solutions to include based on the number of nodes in each solution. We refer to this approach as the GK-CI method. This truncation of the FCI method does not require any physical intuition of orbitals (e.g. the Hartree-Fock solutions) and is not constructed from the typical excitation-hierarchy approach. In this presentation I will described the method as well as provide preliminary results from both electronic structure and nuclear structure calculations.[5] [1] M. Griebel et al., Constr. Approx. 16, 525 2000. [2] H. Bungartz et al., Acta Numer. 13, 147 2001. [3] G.W. Wasilkowski et al.,

Found. Comput. Math. 5, 451 2005. [4] S.A. Smolyak, Dokl. Akad. Nauk. 4, 240 1963. [5] J.S.M. Anderson, et al. Comput. Theor. Chem. 1142, 66-77 2018.

James S. M. Anderson
Universidad Nacional Autónoma de México
james.anderson@iquimica.unam.mx

MS148

Adaptive Sparse Grid Discontinuous Galerkin Method for High Dimensional PDEs

In this talk, I will introduce adaptive sparse grid DG methods, and discuss their applications in solving high dimensional PDEs, including kinetic equations, Hamilton-Jacobi equations, etc. The methods, which combines the advantages of sparse grid and DG approach, are shown to be particularly effective for high dimensional transport equations. We prove stability and accuracy for model equations. Adaptivity is incorporated for time evolution problems. Benchmark test results are shown.

Wei Guo
Texas Tech University
weimath.guo@ttu.edu

Yingda Cheng
Virginia Polytechnic Institute and State University
yingda@vt.edu

Juntao Huang
Texas Tech Univeresity
juntao.huang@ttu.edu

Yuan Liu
Wichita State University
yuan.liu@wichita.edu

MS148

Efficient Uncertainty Quantification and Global Sensitivity Analysis of Time-Dependent Outputs in Hydrology Modeling

Hydrology models, which conceptualize complex, spatially distributed processes, inevitably result in non-unique input-output relationships and significant predictive uncertainties. The reliable assessment of these uncertainties is crucial for models simulating real-world problems. This study explores computationally efficient methods for Forward Uncertainty Quantification (FUQ) and Global Sensitivity Analysis (SA) of hydrology models that produce time-dependent quantities of interest (QoI). We utilize a polynomial chaos expansion surrogate model, combined with non-intrusive pseudo-spectral projection and sparse grid (SG) strategies, to effectively express uncertainty evolution and perform SA with minimal model runs. To maintain the combination technique's black-box property while focusing on regions of interest adaptively, we utilize a spatially adaptive SG combination technique with dimension-wise refinement algorithm [Obersteiner, M. and Bungartz, H. J., A generalized spatially adaptive sparse grid combination technique with dimension-wise refinement, 2021]. The particular focus of our work is on how to combine time-dependent QoIs with adaptive SG. Our work further incorporates the KarhunenLove expansion as an intermediate surrogate to capture system dynamics, reducing computational demands. This method aims to efficiently identify key uncertain parameters and their interactions, enhancing the practical application of UQ and adaptive SG in real-world hydrology models.

Ivana Jovanovic Buha
Technical University Munich
ivana.jovanovic@tum.de

MS148

Discretization of High Dimensional PDEs on Locally Adaptive Sparse Grids

High-dimensional elliptic partial differential equations with variable coefficients can be discretized on sparse grids. With prewaveletes being L^2-orthogonal, one can apply the Ritz-Galerkin discretization to obtain a linear equation system with $O(N(\log N)^{d-1})$ unknowns. We extend this discretization to locally adaptive sparse grids. These grids are constructed in a way that preserves the optimality properties and the unidirectional approach of the matrix-vector multiplication algorithm developed for regular sparse grids. Furthermore, we discuss how to approximate general variable coefficients. This is done by a local cell-wise approach, which allows the application of randomized quadrature formulae, more precisely the Monte Carlo method, to compute the local stiffness matrices.

Leonie Rumi
FAU Erlangen-Nürnberg, Germany
leonie.rumi@fau.de

MS149

Nonlocal Attention Operator a Foundation Model for Inverse PDE Problems

Learning problems in physical systems are often characterized as discovering operators that map between function spaces based on a few instances of function pairs. This task frequently presents a severely ill-posed PDE inverse problem. In this work, we propose a novel neural operator architecture based on the attention mechanism, which we coin Nonlocal Attention Operator (NAO), and explore its capability towards developing a foundation physical model. In particular, we show that the attention mechanism is equivalent to a double integral operator that enables nonlocal interactions among spatial tokens, with a data-dependent kernel characterizing the inverse mapping from data to the hidden parameter field of the underlying operator. As such, the attention mechanism extracts global prior information from training data generated by multiple systems, and suggests the exploratory space in the form of a nonlinear kernel map. Consequently, NAO can address ill-posedness and rank deficiency in inverse PDE problems by encoding regularization and achieving generalizability. We empirically demonstrate the advantages of NAO over baseline neural models in terms of generalizability to unseen data resolutions and system states. Our work not only suggests a novel neural operator architecture for learning interpretable foundation models of physical systems, but also offers a new perspective towards understanding the attention mechanism.

Siavash Jafarzadeh
Lehigh University
sij222@lehigh.edu

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

Ning Liu

GEM.inc
ningliu@umich.edu

Fei Lu
Johns Hopkins University
flu15@jh.edu

Tian Gao
IBM Thomas J. Watson Research Center
tgao@us.ibm.com

Stewart Silling
Sandia National Laboratories
sasilli@sandia.gov

MS149

Scientific and Engineering Foundation Models: Current State and Future Directions

Foundation models will transform the landscape of scientific research and engineering, offering unprecedented capabilities for solving complex, interdisciplinary challenges. In this talk, I will provide an overview of the current state of the art in scientific and engineering foundation models, focusing on their architecture, capabilities, and emerging trends. A key focus will be on the integration of dynamical systems as a training modality, which enhances the ability of these models to predict and simulate real-world phenomena. I will explore a range of potential applications. The talk will conclude with a discussion on future directions, addressing the need for uncertainty quantification, scalability, and the incorporation of domain-specific knowledge to push the boundaries of what these models can achieve.

Christian Moya
Purdue University
cmoyacal@purdue.edu

MS149

Convergence and Error Control of Consistent Pinns for Elliptic PDEs

We study the convergence in terms of the number of collocation points of Physics-Informed Neural Networks (PINNs) for the solution of elliptic PDEs. Specifically, given Sobolev or Besov space assumptions on the right hand side of the PDE and on the boundary values, we determine the minimal number of collocation points required to achieve a given accuracy. These results apply more generally to any collocation method which only makes use of point values. Next, we introduce a novel PINNs loss function based upon the PDE regularity theory, which we call consistent PINNs. We prove an error bound for consistent PINNs in terms of the number of collocation points and the final loss value achieved. Finally, we present numerical experiments which demonstrate that the consistent PINNs loss results in improved solution error.

Andrea Bonito
Texas A&M University
Department of Mathematics
bonito@math.tamu.edu

Ronald DeVore
Department of Mathematics
Texas A&M University
rdevore@math.tamu.edu

Guergana Petrova, Jonathan Siegel
Texas A&M University
gpetrova@math.tamu.edu, jwsiegel@tamu.edu

MS149

Learning Stochastic Closures Via Conditional Diffusion Model and Neural Operator

Closure models are widely used in simulating complex multiscale dynamical systems such as turbulence and the earth system, for which direct numerical simulation that resolves all scales is often too expensive. For those systems without a clear scale separation, deterministic and local closure models often lack enough generalization capability, which limits their performance in many real-world applications. In this work, we propose a data-driven modeling framework for constructing stochastic and non-local closure models via conditional diffusion model and neural operator. Specifically, the Fourier neural operator is incorporated into a score-based diffusion model, which serves as a data-driven stochastic closure model for complex dynamical systems governed by partial differential equations (PDEs). We also demonstrate how accelerated sampling methods can improve the efficiency of the data-driven stochastic closure model. The results show that the proposed methodology provides a systematic approach via generative machine learning techniques to construct data-driven stochastic closure models for multiscale dynamical systems with continuous spatiotemporal fields.

Jinlong Wu
University of Wisconsin-Madison
jinlong.wu@wisc.edu

MS149

Separable Operator Networks

Operator learning has become a powerful tool in machine learning for modeling complex physical systems governed by PDEs. Although DeepONet show promise, they require extensive data acquisition. Physics-informed DeepONets (PI-DeepONet) mitigate data scarcity but suffer from inefficient training processes. We introduce Separable Operator Networks (SepONet), a novel framework that significantly enhances the efficiency of physics-informed operator learning. SepONet uses independent trunk networks to learn basis functions separately for different coordinate axes, enabling faster and more memory-efficient training via forward-mode automatic differentiation. We provide a universal approximation theorem for SepONet proving the existence of a separable approximation to any nonlinear continuous operator. Then, we comprehensively benchmark its representational capacity and computational performance against PI-DeepONet. Our results demonstrate SepONet's superior performance across various nonlinear and inseparable PDEs, with SepONet's advantages increasing with problem complexity, dimension, and scale. For 1D time-dependent PDEs, SepONet achieves up to 112 faster training and 82 reduction in GPU memory usage compared to PI-DeepONet, while maintaining comparable accuracy. For the 2D time-dependent nonlinear diffusion equation, SepONet efficiently handles the complexity, achieving a 6.44% mean relative ℓ_2 test error, while PI-DeepONet fails due to memory constraints.

Xinling Yu
University of California, Santa Barbara
xyu644@ucsb.edu

Sean Hooten
Hewlett Packard Labs, Hewlett Packard Enterprise
sean.hooten@hpe.com

Ziyue Liu, Yequan Zhao
University of California, Santa Barbara
ziyueliu@ucsb.edu, yequan_zhao@ucsb.edu

Marco Fiorentino
Hewlett Packard Labs, Hewlett Packard Enterprise
marco.fiorentino@hpe.com

Thomas Vaerenbergh
Hewlett Packard Labs, HPE Belgium
thomas.van-vaerenbergh@hpe.com

Zheng Zhang
University of California, Santa Barbara
zhengzhang@ece.ucsb.edu

MS150

Fundamental Physics in the Era of AI and Machine Learning

The advent of large data sets has revolutionized various fields of physics, from high-energy physics to astronomy and cosmology. Modern experimental detectors such as the LHC, Gaia space mission, and LIGO, along with upcoming experiments like the Rubin Observatory and Square Kilometer Array (SKA), necessitate a redefinition of our approach to fully leverage these advancements. Concurrently, the rise of machine learning (ML) and artificial intelligence (AI) is transforming our ability to analyze and interpret vast amounts of data. The application of cutting-edge AI/ML techniques to Big Data is poised to unlock unprecedented insights into fundamental physics, addressing pivotal questions such as the search for physics beyond the Standard Model, the particle nature of dark matter, the cosmological formation of structures, and the history of the Universe. This paper explores the integration of AI/ML in fundamental physics, emphasizing collaborative efforts between AI/ML experts and physicists. By showcasing the potential of these interdisciplinary approaches, we aim to highlight the transformative impact on understanding and advancing fundamental physics.

Abdullah Dabdoub
University of Paris-Saclay
abdullahdabdoub@gmail.com

Mohamed Labadi
University of Chlef
Frontiers Labs
mohamed.labadi2@gmail.com

Samir Abdelmalek
University of Tipaza
s.abdelmalek@cu-tipaza.dz

Bujar Ajdari
Medical UniversityofGraz
Medical UniversityofGraz
bujar.ajdari@stud.medunigraz.at

MS150

Machine Learning Surrogates for Multi-Scale Modeling of Biofuels Pyrolysis

Simulation is a key component in understanding and optimizing the biofuel pyrolysis process; however, the multi-scale nature of this process and corresponding simulation limits makes it difficult to optimize reactor scale operation depending on microscale material property. Particle scale models are capable of incorporating microstructure dependent processes (such as asymmetric heat transfer), but do not directly connect to reactor control parameters (such as input flow speed). In contrast, macroscale reactor models directly simulate fluid and particle flow, but the larger length scales necessitate neglecting microstructure processes. Here we demonstrate a machine learning surrogate model for incorporating microstructure effects in macroscale two fluid Euler simulations via reparameterization of particle properties, trained on 3D FEM particle models. We show results for both a direct surrogate of an offline-optimized reparameterization and a differentiable model that optimizes and predicts the parameterization in a single step by leveraging a differentiable 0D initial value problem solver. Both these models demonstrate improved predictive properties at the macroscale compared to the baseline parameterization but differ in generalizability and volume of training data required. Finally, we discuss the prospects of adaptively training such models through goal-oriented sampling to minimize overall computational time required for down-stream engineering applications.

Hilary Egan
National Renewable Energy Laboratory
hilary.egan@nrel.gov

MS151

Quantum Finite Temperature Arnoldi's Method

We present a quantum algorithm for approximating thermal expectation values, called the quantum finite temperature Arnoldi's method, which is inspired by the classical algorithm called finite temperature Lanczos method. Our algorithm is based on constructing Krylov spaces on a quantum computer, then estimating expectation values of using a form of Gauss quadrature based on Arnoldi's method.

William Kirby
IBM Quantum
william.kirby@ibm.com

MS152

An Efficient Proximal-Based Approach for Solving Nonlocal Allen-Cahn Equations

In this talk, we present an efficient approach for the spatial and temporal discretization of the nonlocal Allen-Cahn equation, which incorporates various double-well potentials and an integrable kernel, with a particular focus on a non-smooth obstacle potential. While nonlocal models offer enhanced flexibility for complex phenomena, they often lead to increased computational costs and there is a need to design efficient spatial and temporal discretization schemes, especially in the non-smooth setting. To address this, we propose first- and second-order energy-stable time-stepping schemes combined with the Fourier collocation approach for spatial discretization. We provide energy stability estimates for the developed time-stepping schemes. A key aspect to our approach involves a representation of a solution via proximal operators. This together with the spatial and temporal discretizations enables direct evaluation of the

solution that can bypass the solution of nonlinear, non-smooth, and nonlocal system. This method significantly improves computational efficiency, especially in the case of non-smooth obstacle potentials, and facilitates rapid solution evaluations in both two and three dimensions. We present several numerical experiments to illustrate the effectiveness of our approach.

Olena Burkovska
Oak Ridge National Laboratory
burkovskao@ornl.gov

Ilyas Mustapha
Kansas State University
ilyas1@ksu.edu

MS152

The Use of Multifidelity Methods to Expedite Uncertainty Quantification on Nonlocal Problems

Nonlocal models are essential for accurately describing certain physical phenomena where traditional partial differential equations (PDEs) fall short. One major characteristic of nonlocal models is that they account for interactions between points separated by distances up to a specified length scale, commonly referred to as the horizon. However, the reduced sparsity caused by distant interactions makes nonlocal models computationally more intensive than their local counterparts. This higher computational demand can limit the feasibility of using nonlocal models in outer-loop applicationssuch as uncertainty quantification, optimization, inference, or controlthat require multiple model evaluations. To address this challenge, we propose a multifidelity approach for nonlocal models suited for uncertainty quantification. This method integrates surrogate models of varying fidelity with the high-fidelity model. The low-fidelity surrogate models are built by reducing mesh refinement, the horizon, or both. The multifidelity approach accelerates Monte Carlo simulations of the quantity of interest by conducting most evaluations with the cheaper surrogate models while reserving a limited number of high-fidelity evaluations for accuracy. Our results demonstrate that this approach achieves about two orders of magnitude in speed-up compared to using the high-fidelity model alone.

Parisa Khodabakhshi
Lehigh University
pak322@lehigh.edu

Olena Burkovska
Oak Ridge National Laboratory
burkovskao@ornl.gov

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

Max Gunzburger
Florida State University
Department of Scientific Computing
mgunzburger@fsu.edu

MS152

Dynamic Electromechanical Interactions in Piezoelectric Materials Through Peridynamic Framework

Peridynamics (PD) is a theoretical framework that employs integro-differential equations as an alternative to the traditional partial differential equations used in classical continuum mechanics. The theory of PD incorporates a nonlocal length parameter, hence classifying it as a non-local theory. This study investigates the dynamic behaviour of piezoelectric material in the framework of peridynamic theory. An analysis of the propagation of waves has been performed to determine the dispersion relation. This study considers Gaussian, exponential, and constant functions as nonlocality functions. Graphs are demonstrated to illustrate the variations in frequency, phase velocity, and group velocity with changes in the size of the horizon, the number of points inside the horizon, different nonlocality functions, and shape parameters of nonlocality functions. Also, the significant impact of electromechanical coupling on frequency, phase velocity, and group velocity is presented and studied using numerical computation and graphical demonstration. Understanding their behaviour and properties from a wave and vibrational perspective is essential for designing efficient and effective piezoelectric devices. Their ability to convert mechanical energy into electrical energy and vice versa makes them indispensable in sensors, actuators, transducers, energy harvesters, and medical imaging devices.

Subrata Mondal
National Institute of Technology Andhra Pradesh
Department of Mathematics, School of Sciences
pf122102@student.nitandhra.ac.in

Sudarshan Dhua
National Institute of Technology Andhra Pradesh
dhuasudarshan@gmail.com

MS152

CabanaPD: A Meshfree GPU-enabled Peridynamics Code for Exascale Fracture Simulations

Peridynamics is a nonlocal reformulation of classical continuum mechanics suitable for material failure and damage simulation, which has been successfully demonstrated as an effective tool for the simulation of complex fracture phenomena in many applications. However, the nonlocal nature of peridynamics makes it highly computationally expensive, compared to classical continuum mechanics, which often hinders large-scale fracture simulations. In this talk, we will present CabanaPD, a meshfree GPU-enabled peridynamics code for large-scale fracture simulations. CabanaPD is built on top of two main libraries: Kokkos and Cabana, both developed throughout the Exascale Computing Project (ECP). CabanaPD is performance-portable and exascale-capable, and it is designed to run on U.S. Department of Energys supercomputers, including Frontier, the first exascale machine and today's top supercomputer worldwide.

Pablo Seleson, Sam Reeve
Oak Ridge National Laboratory
selesonpd@ornl.gov, reevest@ornl.gov

MS152

Coarse Graining Kinetic Monte Carlo into a Nonlocal Continuum Model

Kinetic Monte Carlo (KMC) methods model the evolution of a system over time by associating a rate of change with

each degree of freedom. The rates are typically computed by subjecting each degree of freedom to a virtual fluctuation. For each such virtual fluctuation, the model computes the resulting change in the total energy of the system, and this energy change determines the rate. Among the fluctuations, one is accepted according to a stochastic algorithm that tends to favor those degrees of freedom with the largest rate, and the time step is advanced, leaving the other degrees of freedom unchanged. Depending on the energy landscape of the system, applications of KMC methods can vary widely in computational speed, motivating the development of a deterministic surrogate model. In this work, a nonlocal continuum model similar to peridynamics is proposed for accomplishing this. The nonlocal model adopts a continuum form of the total energy with a rate equation that reduces the total energy over time, without the use of random fluctuations. The parameters are calibrated using the KMC model. The proposed method is demonstrated with a KMC Potts model of a two-phase system. It is shown that the nonlocal model captures some essential features of grain growth in polycrystals, including the finite surface energy at grain boundaries. This feature causes these interfaces between grains to reduce their total area, leading to reduced curvature as the system evolves.

Stewart Silling
Sandia National Laboratories
sasilli@sandia.gov

MS153

AMG Convergence Theory for Nonsymmetric and Indefinite Problems

Algebraic multigrid (AMG) methods are known for their effectiveness in solving sparse symmetric positive definite (SPD) linear systems, where convergence theory is well-established through the A-norm. However, for nonsymmetric systems, the absence of a similar energy norm leaves the convergence theory of AMG less developed. The application of AMG to nonsymmetric and indefinite systems often relies on strategies adapted from SPD theory, but the incorporation of general relaxation schemes is still not well understood. In the SPD context, optimal AMG interpolation offers critical insights for determining the best two-grid convergence rate when combined with a symmetrized relaxation scheme. This study extends optimal AMG convergence theory to nonsymmetric systems by utilizing a matrix-induced orthogonality between the left and right eigenvectors in a generalized eigenvalue problem that connects the system matrix and relaxation operator. This extension allows us to derive a measure for the spectral radius of the two-grid error transfer operator that is mathematically equivalent to the derivation in the SPD case. Furthermore, this approach can be applied to symmetric indefinite problems, such as those arising from saddle point systems, to determine the precise convergence rate of two-grid methods based on optimal interpolation.

Ahsan Ali
University of New Mexico
ahsan@unm.edu

MS153

Sparse Approximate Inverse Smoothers for Multigrid Preconditioning in GPU-Accelerated Poisson Solvers

Two classes of preconditioners have proven effective for high-order finite element discretizations: p-multigrid with a Chebyshev-accelerated Jacobi smoother and the low-order refined preconditioner. However, both approaches require a substantial number of iterations when applied to highly deformed meshes. Sparse Approximate Inverse (SAI) and Factorized Sparse Approximate Inverse (FSAI) preconditioners have demonstrated their ability to enhance the convergence of iterative methods, with inherently parallel construction. In this work, we employ SAI and FSAI as smoothers within a p-multigrid preconditioning. The smoothers are constructed using the matrix-free operator, as the matrix is not explicitly stored in high-order finite element methods. To further reduce computational costs, we implement a thresholding strategy and a mixed-precision algorithm for smoother operations. We evaluate the performance of these smoothers on highly deformed meshes using GPUs. Our numerical results demonstrate that these smoothers are highly effective in reducing both iteration counts and runtime.

Yichen Guo, Tim Warburton, Eric De Sturler
Virginia Tech
ycguo@vt.edu, tcew@vt.edu, sturler@vt.edu

Noel A. Chalmers
AMD Research
noel.chalmers@gmail.com

MS153

Semi-Structured Algebraic Multigrid in Hypre

Due to their optimal linear complexity, multigrid methods are widely used for the solution of large-scale linear systems arising from PDE discretizations. Algebraic multigrid methods are particularly attractive due to their black-box approach, constructing a multigrid hierarchy based only on the linear system to be solved. This approach generally neglects the structure that is often present in the underlying PDE discretizations, however. This talk describes the Semi-structured Algebraic Multigrid (SSAMG) algorithm, which leverages the structure present in the linear system while maintaining a black-box algebraic approach. SSAMG is implemented in hypre, a library of high-performance linear solvers specializing in multigrid algorithms. This talk will give an overview of the approach and design of SSAMG and present the latest numerical results highlighting the benefits of SSAMG compared to fully unstructured AMG.

Wayne Mitchell, Wayne Mitchell, Robert D. Falgout, Victor Magri, Ulrike M. Yang
Lawrence Livermore National Laboratory
mitchell82@llnl.gov, mitchell82@llnl.gov, falgout2@llnl.gov, paludettomag1@llnl.gov, yang11@llnl.gov

MS153

Automated Design of Multigrid Methods Using Genetic Programming

Although multigrid is asymptotically optimal for many important PDE problems, its efficiency relies heavily on the careful selection of the individual algorithmic components. In contrast to recent approaches that learn optimal multigrid components using machine learning techniques, we adopt a complementary strategy, employing evolutionary algorithms to construct efficient multigrid cycles. Here, we will present its application to generate efficient BoomerAMG methods in hypre. BoomerAMG is extended to support flexible cycling, incorporating level-specific smoothing sequences and non-recursive cycling patterns; and these are

automatically generated using genetic programming (GP). Numerical experiments demonstrate the potential of these non-standard GP cycles to improve multigrid performance both as a solver and a preconditioner.

Dinesh Parthasarathy
University of Erlangen-Nuremberg, Germany
dinesh.parthasarathy@fau.de

Wayne Mitchell, Wayne Mitchell
Lawrence Livermore National Laboratory
mitchell82@llnl.gov, mitchell82@llnl.gov

Harald Koestler
University Erlangen-Nürnberg
harald.koestler@fau.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS153

An Energy Minimization Amg Method for H-Curl

A new AMG algorithm is presented for curl-curl electromagnetics problems that have been discretized with first order edge elements. The algorithm enforces a commuting relationship between interpolation and gradient operators using constraints within an energy minimization AMG approach. This commuting relationship guarantees that the kernel of the algebraically generated discrete curl-curl operators corresponds to discrete gradients. This is an essential condition for achieving good AMG convergence rates. The new AMG algorithm takes any already-computed nodal interpolation operator to then define an edge interpolation operator with the desired commuting relationship. In this way, it is possible to leverage traditional AMG schemes that generate nodal interpolation operators. We illustrate how the new algorithm can often produce ideal edge interpolation operators when it is given an ideal nodal interpolation operator (e.g., one based on linear interpolation for meshes with a certain regular structure). The computational cost of the energy minimization procedure is discussed. Numerical results are provided showing the efficacy of the AMG approach based on comparisons with other possible AMG strategies for curl-curl operators.

Raymond S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

MS154

Machine Learning with Diffractive Optical Neural Networks

In this talk, I will present our recent progress in diffractive optical neural network (DONN) systems for performing machine learning tasks efficiently and in high-throughput manners. I will discuss a fully reconfigurable DONN architecture and its deployment for scientific computing applications beyond machine vision.

Weilu Gao
University of Utah
weilu.gao@utah.edu

MS154

Ai-Enhanced Design Space Exploration for Negative Capacitance Field Effect Transistors

Ferroelectric materials are highly promising for energy-efficient microelectronics, with applications in ferroelectric field-effect transistors (FeFETs), non-volatile memory devices, and neuromorphic computing. Due to their low power consumption, fast polarization switching, and compatibility with CMOS technology, they are well-suited for next-generation high-density, low-power devices. However, while phase-field simulations are commonly used to provide physical insights into domain pattern evolution and switching properties, their high computational costs limit their applicability to realistic systems. In this talk, we will explore the parameter space defined by dielectric and ferroelectric thicknesses, Landau and gradient energy coefficients, multigrain shapes and sizes, and voltage using exascale-enabled phase-field simulations coupled with a machine learning framework. We will analyze input-output relationships using various ML models, including Gaussian Process Regression (GPR), Feedforward Neural Networks (FNN), and Convolutional Neural Networks (CNN). Additionally, we will evaluate the potential of a Graph Neural Network (GNN) framework to investigate mixed-phase ferroelectric oxides with an arbitrary number of grains for negative-capacitance field-effect transistors.

Prabhat Kumar
Lawrence Berkeley National Laboratory
prabhatkumar@lbl.gov

Jorge Munoz, Christian Fernandez-Soria
University of Texas at El Paso
jamunoz@utep.edu, cafernandez12@miners.utep.edu

Andrew J. Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

Camden Wilson
University of Texas at El Paso
cawilson4@miners.utep.edu

Yingheng Tang, Zhi Jackie Yao
Lawrence Berkeley National Laboratory
ytang4@lbl.gov, jackie_zhiyao@lbl.gov

MS154

Optical Neural Engine for Solving Scientific Partial Differential Equations

Solving partial differential equations (PDEs) is crucial for scientific research and development. Data-driven machine learning (ML) approaches are emerging to accelerate the time-consuming and computation-intensive numerical simulations of PDEs. Although optical systems offer high-throughput and energy-efficient ML hardware, their use in solving PDEs has yet to be demonstrated. Here, we present an Optical Neural Engine (ONE) architecture that combines diffractive optical neural networks for Fourier space processing and optical crossbar structures for real-space processing to solve two-dimensional time-dependent and time-independent PDEs across various disciplines, including Darcy flow, Poisson's equation of demagnetization, the Navier-Stokes equation in incompressible fluids,

Maxwells equations in nanophotonic metasurfaces, and coupled PDEs in multiphysics systems. We numerically and experimentally demonstrate the capability of the ONE architecture, which leverages high-performance dual-space processing to outperform traditional PDE solvers and rival state-of-the-art ML models. Moreover, it can be implemented using optical computing hardware, offering low-energy, highly parallel constant-time processing, and real-time reconfigurability for tackling multiple tasks with the same architecture. The demonstrated architecture provides a versatile and powerful platform for large-scale scientific and engineering computations.

Yingheng Tang
Lawrence Berkeley National Laboratory
ytang4@lbl.gov

Ruiyang Chen, Minhan Lou, Jichao Fan
The University of Utah
u1326032@utah.edu, minhan.lou@outlook.com, jichao.fan@utah.edu

Cunxi Yu
University of Maryland, College Park
cunxiyu@umd.edu

Andrew J. Nonaka, Zhi Jackie Yao
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, jackie_zhiyao@lbl.gov

Weilu Gao
University of Utah
weilu.gao@utah.edu

MS154

Ai-Driven Design and Optimization of Ferroic Based Electronics

Artificial intelligence (AI) is transforming many scientific areas, especially electronics where it reduces the time and cost of designing and optimizing devices. AI effectively manages spatial and temporal discrepancies among physical mechanisms in new electronic devices, which often require heterogeneous structures for functionality and miniaturization. Ferroic materials are invaluable for their energy-saving properties but pose significant challenges for numerical modeling. Traditional algorithms struggle with the vast differences in scale and complexity of the equations governing these interactions. Consequently, developers of post-CMOS technologies often resort to trial-and-error due to inadequate simulation tools. The use of neural networks for solving these complex equations has expanded, marking the growth of scientific machine learning (SciML). We are tapping into machine learning techniques, such as large multimodal models and Fourier neural operators, particularly for applications involving ferroelectric and ferromagnetic materials. These methods enable more efficient simulations and faster predictions of complete time sequences, significantly speeding up design timelines.

Zhi Jackie Yao, Andrew J. Nonaka, Prabhat Kumar
Lawrence Berkeley National Laboratory
jackie_zhiyao@lbl.gov, ajnonaka@lbl.gov, prabhatkumar@lbl.gov

Jorge Munoz
University of Texas at El Paso
jamunoz@utep.edu

Lavanya Ramakrishnan
Lawrence Berkeley National Laboratory
lramakrishnan@lbl.gov

Jialin Song
Simon Fraser University
jialin_song@berkeley.edu

Pu Ren
Northeastern University
ren.pu@northeastern.edu

N. Benjamin Erichson
International Computer Science Institute
erichson@icsi.berkeley.edu

MS155

Well-Posed Variational Formulations for Physics-Informed Learning

Physics-informed neural networks typically minimize a loss function – most commonly the mean-squared error of the interior and boundary residuals – to approximate the forward partial differential equation (PDE) problem. This loss may be alternatively as a least squares variational problem. In this talk, we demonstrate that many common loss functions lead to ill-posed variational problems due to incompatibility of the interior and boundary norms or the loss function being incompatible with solution regularity (e.g. singular solutions). We present a theoretical framework for developing well-posed loss functions for approximating PDEs using neural networks (NNs) and demonstrate its effectiveness on examples with singular solutions, such as Stokes flow.

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

Justin Dong
Lawrence Livermore National Laboratory
dong9@llnl.gov

MS155

Approximation Rates for Neural Operator Surrogates of Elliptic Pdes

We consider the problem of approximating the solution operator of a PDE (e.g., from coefficients to solution), viewed as a map between subsets of infinite dimensional spaces. In recent years, several techniques based on neural networks (NNs) have been developed to tackle this problem, which has many applications in computational science whenever an application requires multiple solutions of similar problems. In this talk, I will present some theoretical quantitative results on the approximation of solution operators of elliptic PDEs by NN-based surrogates. I will discuss the convergence rates of neural operators and how they depend on the smoothness of the coefficients in the input sets. I will consider PDEs in smooth and non-smooth domains, with bounded coefficients, and with coefficients that are log-normally distributed.

Carlo Marcati
University of Pavia

carlo.marcati@unipv.it

MS155

A New Paradigm for Machine Learning

The fundamental problem of machine learning is often formulated as the problem of function approximation. For example, we have data of the form $\{(x_j, y_j)\}$, where y_j is the class label for x_j, and we want to approximate the class label as a function of the input x. The standard way for this approximation is to minimize a loss functional, usually with some regularization. Surprisingly, even though the problem is posed as a problem of function approximation, approximation theory has played only a marginal role in this theory. We develop an idea of using approximation theory/harmonic analysis tools more directly, assuming the manifold hypothesis.

Hrushikesh Mhaskar
Claremont Graduate University
Dept. of Mathematics
hmhaska@gmail.com

MS155

A Multi-Level Neural Network Approach for Accurate Operator Approximations

The solution of boundary-value problems using deep learning approaches, such as the physics-informed neural networks or the deep operator networks, have been extensively investigated in recent years. However, achieving high accuracy in the approximations obtained from these methods often remains a significant challenge. A multi-level neural network approach was proposed in [1] that allows one to iteratively reduce the errors, sometimes within machine precision, when approximating a solution using PINNs. In this work, we extend the multi-level approach to approximate linear operators using the Green operator networks (GreenONets) as described in [2]. Starting with an initial approximation of the operator, we correct the solution by considering a different Green operator network involving higher frequencies. The method enables one to iteratively reduce the high-frequency contributions present in the residuals. Numerical examples will be presented to demonstrate the efficiency of the proposed multi-level approach. [1] Z. Aldirany, R. Cottereau, M. Laforest, and S. Prudhomme. Multi-level neural networks for accurate solutions of boundary-value problems, Computer Methods in Applied Mechanics and Engineering 419 (2024): 116666. [2] Z. Aldirany, R. Cottereau, M. Laforest, and S. Prudhomme. Operator approximation of the wave equation based on deep learning of Green's function, Computer & Mathematics with Applications, 159 (2024) 21–30.

Serge Prudhomme
École Polytechnique de Montréal
serge.prudhomme@polymtl.ca

Ziad Aldirany, Charlélie Bilodeau
Polytechnique Montréal
ziad.aldirany@polymtl.ca, charlelie.bilodeau@polymtl.ca

Régis Cottereau
Aix-Marseille Université, CNRS, Centrale Marseille, LMA UM
cottereau@lma.cnrs-mrs.fr

Marc Laforest
École Polytechnique de Montréal
marc.laforest@polymtl.ca

MS156

Invariant-domain-preserving and Conservative High-Order Accurate Discontinuous Galerkin Scheme for Compressible Flow Simulation

Research on fluid dynamic simulations has been very active and of interest to both mathematicians and engineering researchers in recent years. The equations raised in modeling compressible flow, such as the compressible Euler equations and the compressible Navier-Stokes equations, are fundamental in gas dynamics with various applications in aerospace engineering and numerous other important areas. Advanced optimization schemes provide powerful tools for finding minimizers when working with large datasets. In this talk, we combine the numerical methods for partial differential equations (PDEs) and large-scale non-smooth optimization techniques to construct an efficient high-order limiter. In high-resolution fluid dynamic simulations, the size of the data set that needs to be processed at each time step can be quite large. The generalized Douglas-Rachford algorithm has a provable asymptotic linear convergence rate. With optimal algorithm parameters, it is more efficient in terms of minimal memory requirements and low iteration costs. Our optimization method scales well for each iteration, with a complexity of $\mathcal{O}(N)$ where N is the total number of mesh cells. Our method is robust, efficient, and high-order accurate, which preserves the conservation and invariant domain with large time step size for compressible flow simulations and can be further extended to other types of PDEs.

Chen Liu
Department of Mathematics, Purdue University, West Lafayette
chenl@uark.edu

Xiangxiong Zhang
Purdue University
Department of Mathematics
zhan1966@purdue.edu

MS156

Numerical Investigation on the Biological Phase-Field Patterning

In this talk, we discuss a class of discontinuous Galerkin methods under the scalar auxiliary variable framework (SAV-DG) to solve a biological patterning model in the form of parabolic-elliptic partial differential equation system. In particular, mixed-type discontinuous Galerkin approximations are used for the spatial discretization, aiming to achieve a balance between the high resolution and computational cost. Second and third order backward differentiation formulas are considered under SAV framework for discrete energy stability. Numerical experiments are provided to show the effectiveness of the fully discrete schemes and the governing factors of patterning formation.

Yuan Liu
Wichita State University
yuan.liu@wichita.edu

MS156

Structure-Preserving Finite-Element Schemes for

the Euler-Maxwell and Euler-Poisson Equations

We discuss structure-preserving numerical discretization techniques for the Euler-Poisson and Euler-Maxwell equations that find applications in the modeling and simulation of fluid plasma, self gravitation, and nanoscale optical device modeling. A key feature of the methods presented is that they maintain a discrete energy law, as well as hyperbolic invariant domain properties, such as positivity of the density and a minimum principle of the specific entropy, on a fully discrete level. We first introduce and discuss the underlying algebraic discretization technique based on collocation, convex limiting, and a high-order IMEX splitting technique and then discuss how the method is applied to the coupled Euler-Poisson and Euler-Maxwell system. We demonstrate how a careful choice of continuous and discontinuous finite element spaces for the PDE subsystems combined with a block elimination of the source subsystem leads to an energy stable formulation with robust linear algebra. Most crucially, the scheme is able to cope with the inherent multiple time scales of the coupled PDE system. A detailed discussion of algorithmic details is given, as well as proofs of the claimed properties.

Jordan N. Hoffart, Matthias Maier
Texas A&M University
jordanhoffart@tamu.edu, maier@tamu.edu

Ignacio Tomas
Department of Mathematics & Statistics, Texas Tech Universit
igtomas@ttu.edu

MS156

Bound-Preserving Discontinuous Galerkin Solutions for Multiphase Flows in Porous Media

Multiphase flows in porous media are mathematically characterized by systems of coupled nonlinear partial differential equations. Suitable discretizations must be locally mass conservative to avoid numerical unbounded instabilities. Interior penalty discontinuous Galerkin methods have been successfully applied to modeling multiphase flows for heterogenous media with discontinuous permeability fields and discontinuous capillary pressures. However, while the numerical solutions are stable, they are not guaranteed to satisfy physical bounds. In addition, overshoot and undershoot phenomena are observed in the neighborhood of the saturation front. In this work, we present iterative flux limiters combined with slope limiters. The resulting saturations are bound-preserving. Robustness of the proposed limiting technique is tested on both incompressible and compressible two-phase flows with heterogeneity in the permeability field, with gravity and with anisotropy.

Beatrice Riviere
Rice University
Computational & Applied Mathematics Department
riviere@rice.edu

MS156

Energy-Consistent Time Discretization of Port-Hamiltonian Systems

Various ordinary and partial differential equations arising from physics can be written as port-Hamiltonian systems. Their Hamiltonian function represents an energy that is conserved or dissipated along solutions. Numerical schemes are energy-consistent, if the Hamiltonian is preserved or dissipated also along approximate solutions. This type of structure preservation property is not in general satisfied for standard continuous Petrov-Galerkin (cPG) time discretizations of arbitrary order. In this talk we present a modification of the cPG method of arbitrary order, which is energy consistent for a general class of port-Hamiltonian systems.

Jan Giesselmann
Technical University of Darmstadt
giesselmann@mathematik.tu-darmstadt.de

Attila Karsai
Technische Universität Berlin
karsai@math.tu-berlin.de

Tabea Tscherpel
Universitat Bielefeld
tscherpel@mathematik.tu-darmstadt.de

MS157

Dynamic Tomography Regularization with Optimal Space-Time Priors

We consider a sequence of sparse and dynamic tomography problems by regularizing in both spatial- and temporal domains using the cylindrical shearlet representation system and sparsity promoting regularization. This choice of sparsity promoting regularization is motivated by the ability to optimally approximate functions in the class of cartoon-like videos, and properties of the (quasi-)Banach decomposition spaces for $p > 0$. Using statistical inverse learning methods we obtain convergence rates for $p > 1$ in different noise conditions which are supported by numerical tests using both simulated and real dynamic tomography measurements. This is a joint work with T. A. Bubba, D. Labate and L. Ratti

Tommi Heikkilä
Lahti University of Technology LUT
tommi.heikkila@lut.fi

MS157

Efficient Methods for Dynamic Image Reconstruction with Motion Estimation

Large-scale dynamic inverse problems are typically ill-posed and suffer from complexity of the model constraints and large dimensionality of the parameters. A common approach to overcome ill-posedness is through regularization that aims to add constraints on the desired parameters in both space and temporal dimensions. In this work, we propose an efficient method that incorporates a model for the temporal dimension by estimating the motion of the objects alongside solving the regularized problems. In particular, we consider the optical flow model as part of the regularization that simultaneously estimates the motion and provides an approximation for the desired image. To overcome high computational cost when processing massive scale problems, we combine our approach with a generalised Krylov subspace method that efficiently solves the problem on relatively small subspaces. The effectiveness of the prescribed approach is illustrated through numerical experiments arising in dynamic computerized tomography and image deblurring applications.

Toluwani Okunola
Tufts University

toluwaokunola@gmail.com

MS157

A Joint Reconstruction and Model Selection AP Proach for Large-Scale Inverse Modeling

Inverse models can often incorporate predictor variables, similar to regression, to help estimate natural processes or parameters of interest from observed data. A core challenge is to identify a few predictor variables that are most informative of the model, given limited observations. This problem is typically referred to as model selection. In this work, we develop a one-step approach, where model selection and the inverse model are performed in tandem. We reformulate the problem so that the selection of a small number of relevant predictor variables is achieved via a sparsity-promoting prior. Then, we describe hybrid iterative projection methods based on flexible Krylov subspace methods for efficient optimization. These approaches are well-suited for large-scale problems with many candidate predictor variables. We evaluate our results against traditional, criteria-based approaches and demonstrate the applicability and potential benefits of our approach using examples from atmospheric inverse modeling.

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Malena Sabate Landman
Emory University
malena.sabatelandman@maths.ox.ac.uk

Julianne Chung
Department of Mathematics
Emory University
julianne.mei-lynn.chung@emory.edu

Jiahua Jiang
University of Birmingham
j.jiang.3@bham.ac.uk

Scot Miller
Johns Hopkins University
smill191@jhu.edu

MS158

Building Surrogate Models Using High Order Derivative Information

Efficient algorithms for solving inverse problems require repeated evaluation of parameter-to-output maps, and derivatives of these maps (e.g., gradients, Hessian vector products, and in some cases higher derivative actions), nested within outer loops. This is an expensive proposition because these maps are typically high dimensional, non-linear, and implicitly depend on the solution of a partial differential equation (PDE) or system of PDEs. Hence, it is desirable to construct cheap surrogate models for these mappings that can be used in place of expensive high-fidelity numerical simulation. Linearized models may be constructed using fast randomized methods, but often nonlinearities are important and should not be neglected. Here we present an efficient approach to construct nonlinear surrogate models based on higher order Taylor series. The derivative tensors in the Taylor series are far too large to be formed, so we fit them with tensor network models using random symmetric high order directional derivatives of the parameter-to-output map as training data. We show that our approach is able to form high order (e.g., 7th order) Taylor series surrogate models for high-dimensional PDE-based mappings. Although these surrogate models require an expensive "offline" construction phase, they are cheap to evaluate "online," and are accurate in terms of both function values and derivatives.

Nicholas Alger
The University of Texas at Austin
Oden Institute
nalger225@gmail.com

Blake Christierson
University of Texas at Austin
bechristierson@utexas.edu

Chen Peng
Georgia Institute of Technology
pchen402@gatech.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS158

Efficient Probabilistic Approaches for Solving Combinatorial Optimal Experimental Design Problems

In this talk, I present a fully probabilistic approach for solving binary optimization problems with black-box objective functions and with budget constraints. This problem is prevalent in optimal control applications such as optimal design of experiments including optimal sensor placement for parameter identification and Bayesian inverse problems. In the probabilistic approach, the optimization variable is viewed as a random variable and is associated with a parametric probability distribution. To inherently model hard (budget) constraints, we extend and employ a family of conditional Bernoulli models. The original optimization problem is replaced with an optimization over the expected value of the original objective, which is then optimized over the chosen conditional probability distribution parameters. This approach (a) is generally applicable to binary optimization problems with nonstochastic black-box objective functions and budget constraints; (b) accounts for budget constraints by employing conditional probabilities that sample only the feasible region and thus considerably reduces the computational cost compared with employing soft constraints; and (c) does not employ soft constraints and thus does not require tuning of a regularization parameter, for example to promote sparsity, which is challenging in sensor placement optimization problems.

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

MS158

The Information Geometric Regularization of the Euler Equation

This talk presents an inviscid regularization for mitigating shock formation in the barotropic Euler equations. Solutions of Euler's equations are paths on the manifold of diffeomorphisms. Shocks form when the deformation map reaches the boundary of this manifold. In this work, we

regularize the barotropic Euler equation by modifying the geometry of the diffeomorphism manifold. This modified geometry is motivated by semidefinite programming and the information geometry of the fluid density. In the modified geometry, geodesics do not cross the boundary of the manifold but instead approximate it asymptotically, preventing shock formation while preserving the long-time behavior of the solutions.

Ruijia Cao
Cornell University
rc948@cornell.edu

Florian Schaefer
Georgia Institute of Technology
fts@gatech.edu

MS159

High-order Finite Element Methods for Three-dimensional Multicomponent Convection-diffusion

We derive and analyze a broad class of finite element methods for simulating the stationary, low Reynolds number flow of concentrated mixtures of several distinct chemical species in a common thermodynamic phase. In particular, we discretize the Stokes–Onsager–Stefan–Maxwell (SOSM) equations, which model bulk momentum transport and multicomponent diffusion within ideal and non-ideal mixtures. Unlike previous approaches, the methods are straightforward to implement in two and three spatial dimensions, and allow for high-order finite element spaces to be employed. We accomplish this by reformulating the SOSM equations in terms of the species mass fluxes and chemical potentials, and discretize these unknown fields using stable $H(\mathrm{div})$-L^2 finite element pairs. We prove that the methods are convergent and yield a symmetric linear system for a Picard linearization of the SOSM equations. We also discuss how the proposed approach can be extended to the Newton linearization of the SOSM equations, which requires the simultaneous solution of mole fractions, chemical potentials, and other variables. Our theoretical results are supported by numerical experiments and we present an example of a physical application involving the microfluidic non-ideal mixing of hydrocarbons.

Aaron Baier-Reinio
University of Oxford
baierreinio@maths.ox.ac.uk

Patrick E. Farrell
Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS159

An Overview of High-Order Finite Element Methods

In this talk, we give a brief, high-level overview of high-order finite element methods (FEMs). We focus on p- (hp-version) FEMs, in which the polynomial degree is increased (and mesh is refined) to obtain convergence. We review classical theoretical results from Babuka, Suri, and Guo and recent extensions. We also discuss implementation aspects, such as the choice of basis functions and preconditioners.

Charles Parker
University of Oxford
charles.parker@maths.ox.ac.uk

MS159

A Posteriori Error Analysis of Hybrid High-Order Methods for the Elliptic Obstacle Problem

In this talk, *a posteriori* error analysis of the elliptic obstacle problem is addressed using hybrid high-order methods. The elliptic obstacle problem is a nonlinear model that describes the vertical movement of an object restricted to lie above a barrier while subjected to a vertical force. The method involve cell unknowns represented by degree-r polynomials and face unknowns represented by degree-s polynomials, where $r = 0$ and s is either 0 or 1. The discrete obstacle constraints are specifically applied to the cell unknowns. The analysis hinges on the construction of a suitable Lagrange multiplier, a residual functional and a linear averaging map. The reliability and the efficiency of the proposed *a posteriori* error estimator is discussed, and the study is concluded by numerical experiments supporting the theoretical results.

Ritesh Ritesh
Indian Institute of Technology
iritesh281@gmail.com

Kamana Porwal
Indian Institute of Technology Delhi
kamana@maths.iitd.ac.in

MS159

A High-Order Uniformly Bounds-Constrained Finite Element Method via Variational Inequalities

The solutions to partial differential equations frequently satisfy bounds constraints. When using finite element or finite difference methods, if one wishes to construct an approximate solution that satisfies these same bounds, great care is required. In a finite element context, one can replace a discrete variational problem with a discrete variational inequality. This allows for the selection of an approximate solution from a set of functions which satisfy the bounds constraints. Solving nonlinear optimization problems, though incurring a practical expense, bypasses known order barriers for linear problems and allows for the possibility of high-order and uniformly bounds-constrained finite element methods. It is difficult to work with the entire set of bounds-constrained polynomials. However, the polynomials whose coefficients, when represented in the Bernstein basis, satisfy the bounds constraints form a convenient subset with which to work. Selecting an approximation from this set via a variational inequality, one obtains an approximation which is uniformly bounds-constrained, independent of the mesh used. Recent work seeks to extend this approach to collocation-type Runge-Kutta methods. Using a stage-value formulation, the collocating polynomial can be cast in the Bernstein basis to enforce bounds constraints uniformly in time. Examples are shown in which optimal order accuracy is observed.

Robert C. Kirby
Baylor University
robert_kirby@baylor.edu

Daniel Shapero
Polar Science Center
Applied Physics Lab
shapero@uw.edu

John D. Stephens
Baylor University
john_stephens2@baylor.edu

MS159

Efficient High Order Finite Element Approximation of the Nonlinear Wave Equations in Three Dimensions

The solutions to nonlinear wave equations often exhibit complicated behaviour but the solution is nevertheless quite smooth, which makes high order methods attractive from an approximation theoretic point of view. However, high order methods can, if not implemented appropriately, be computationally expensive and have a CFL condition that degrades rapidly with the polynomial order. In addition, the associated linear systems that arise in the discretisation are ill-conditioned. In this talk we present an approach that addresses these issues. A key component is the provision of a uniform in h and p preconditioner for the mass matrix on unstructured tetrahedral elements that can be applied in a complexity $\mathcal{O}(p^4)$. In fact, all of the components needed for the implementation right from the computation of the residual through to graphical post-processing can be implemented in a complexity $\mathcal{O}(p^4)$. This means that we are able to take advantage of fully implicit time-stepping schemes in order to combat the effect of the CFL condition and carry out simulations with high order elements of challenging non-linear wave equations in three dimensions.

Mingyu Wang
Brown University
mingyu_wang@brown.edu

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

MS160

Genai Aided Synthetic Twin of a Population for Developing Accurate Multilayered Anomaly Detection for Early Outbreak Detection

The synthetic data generation power of Generative Artificial Intelligence (GenAI) in different domains (e.g., image, video, text, video, etc.) motivates us to create synthetic twins of real health data to overcome availability, biases, and privacy. This innovative technique addresses critical gaps in existing public datasets during the SARS-CoV-2 pandemic timeline by generating a data-rich dynamic synthetic twin of the population that maintains the statistical signature of real-world data. Using synthetic and real-world datasets, we apply a co-kurtosis-based projection to effectively detect anomalies within multivariate time-series data for different population resolutions. This approach enhances our early detection capability in infection timelines in individuals and locates outbreak hotspots across populations. We also conduct a thorough evaluation of uncertainties in both synthetic and real-world data, optimizing anomaly detection in physiological metrics. Our results demonstrate significant improvements in the early detection of anomalies across multivariate datasets. This research not only illuminates the intricacies of sounding alarms in response to irregularities in physiological measurements derived from wearable data but also hints at promising applications in healthcare and beyond.

Methun Kamruzzaman
Sandia National Laboratories
System Biology Department
mkamruz@sandia.gov

Jorge Salinas
Sandia National Laboratories
7011 East Ave, Livermore, CA 94551.
jsalin@sandia.gov

Hemanth Kolla, Kenneth Sale, Uma Balakrishnan, Kunal Poorey
Sandia National Laboratories
hnkolla@sandia.gov, klsale@sandia.gov, ubalakr@sandia.gov, kpoorey@sandia.gov

MS160

Towards Verifiable Cancer Digital Twins for Precision Medicine: Tissue Level Modeling Protocol for Prostate Cancer

Cancer exhibits substantial heterogeneity, manifesting as distinct morphological and molecular variations across tumors, which frequently undermines the efficacy of conventional oncological treatments. Intra-tumoral diversity necessitates the development of precision oncology therapies utilizing multiphysics, multiscale mathematical models for cancer. This study integrates agent-based modeling with cellular systems biology to predict tissue-level responses in castration-resistant prostate cancer. Utilizing patient-specific multi-omic data, our framework encodes physics-based cellular behaviors and rule-based interactions in an agent-based model to observe the resultant patient-specific collective tissue dynamics. We conduct local sensitivity analyses to delineate clustering tendencies under different androgen uptake conditions. We then employ machine learning to build surrogates for these models, enhancing computational efficiency, enabling extensive global sensitivity analysis, and exploring verification, validation, and uncertainty quantification methods. This protocol can facilitate the development of verifiable cancer digital twins, enabling the prediction of patient-specific responses to clinical therapies and advancing precision medicine in oncology. *We acknowledge funding from the National Cancer Institute through the Physical Sciences Oncology Network.*

Sharvari Kemkar
University of Pennsylvania
Department of Bioengineering,
sskemkar@seas.upenn.edu

Ravi Radhakrish
Department of Bioengineering
University of Pennsylvania
rradhak@seas.upenn.edu

Uma Balakrishnan
Sandia National Laboratories
ubalakr@sandia.gov

MS160

Applications of Agent-Based Modeling to Digital Twins in Cancer Immunology: Lessons Learned

and Future Directions

Tumors are adaptive ecosystems that are driven by single-cell behaviors, cell-cell interactions, and biophysical constraints. The complex dynamics of these cancer ecosystems make it difficult to tailor and optimize treatment planning to individual patients, to predict the outcomes of the treatments, and to predict treatment failure and choose an alternative plan before adverse events. Digital twinswhich create a virtual copy of a patient for dynamical and predictive modeling–have been posited as an approach to address these challenges [Hernandez-Boussard et al., Digital twins for predictive oncology will be a paradigm shift for precision cancer care, Nat. Med. 2021]. In this talk, we introduce agent-based modeling (ABM) in cancer biology using the open source PhysiCell framework, apply it to explore cancer-immune interactions, and use high-throughput model exploration to investigate feasibility and challenges of building patient-specific digital twins for cancer immunology, with new insights on the impact of stochasticity on patient forecasting [Rocha et al., A multiscale model of immune surveillance in micrometastases: towards cancer patient digital twins, bioRxiv (2023)]. We close by presenting recent advances to ease the construction of cancer immunology simulation models, and lay out our vision for integrating easily-constructed simulation models, high performance computing, and novel patient-derived experimental model systems to create cancer patient digital twins.

Paul Macklin
Indiana University
paul.t.macklin@gmail.com

Heber Rocha
Indiana University Bloomington
hlimadar@iu.edu

MS160

ELEQTRONeX: A Gpu-Accelerated Exascale Framework for Modeling Non-Equilibrium Quantum Transport in Nanomaterials

Non-equilibrium electronic quantum transport is essential for both current and future electronic, optoelectronic, and spintronic devices. Integrating atomistic to mesoscopic length scales within the same nonequilibrium device simulations has been challenging due to the high computational cost associated with the coupled multiphysics and multiscale requirements. In this work, we introduce **ELEQTRONeX** (**ELE**ctrostatic **Q**uantum **TR**ansport modeling **O**f **N**anomaterials at **eX**ascale), a massively parallel, GPU-accelerated framework designed for self-consistently solving the nonequilibrium Green's function formalism alongside electrostatics in complex device geometries. Through tailored algorithms optimized for GPU multithreading, we achieve substantial reductions in computational time and demonstrate excellent scalability across up to 512 GPUs and billions of spatial grid cells. We validate our framework by computing band structures, current-voltage characteristics, conductance, and drain-induced barrier lowering in various 3D configurations of carbon nanotube field-effect transistors. Additionally, we showcase its effectiveness for complex device/material geometries where periodic approaches are not applicable, such as arrays of misaligned carbon nanotubes and simulations with randomly dispersed trap charges that require fully 3D simulations.

Saurabh S. Sawant
Lawrence Berkeley National Lab
saurabhsawant@lbl.gov

MS161

A Compound-Fast Multirate Method for Coupled Free Flow with Porus Media Flow Problems

Multirate methods are well suited for solving coupled multi-physics flow problems because they allow for the utilization of subdomain specific time step sizes that are based on the level of flow activity. In our case,a large time step is used to integrate a slow changing porous media flow problem and a much smaller time step for the fast changing free flow problem. While computationally efficient, fully decoupled multirate schemes suffer small time step restrictions under small parameter regimes. In this talk, we consider a compound step multirate method that combines the stability of the monolithic approach and the computational efficiency of decoupling the flow problems. Our approach solves a few fully coupled monolithic problems on a coarse time mesh, then updates a free flow problem on a fine time mesh using the porous media flow solution from the monolithic step on the interface. We prove the stability and convergence of the scheme and demonstrate performance on manufactured solutions as well as realistic model problems.

Prince Chidyagwai
Loyola University Maryland
Department of Mathematics and Statistics
pchidyagwai@loyola.edu

MS161

Performance of High-Order, Adaptive, and Operator-Splitting Time Integration Approaches for HPC Simulations of Phase-Field Models and Additive Manufacturing

The choice of time integration scheme can substantially impact the computational cost of the numerical solution of systems of partial differential equations for large-scale multiphysics simulations. In this presentation we discuss time integration strategies for two types of multiphysics materials science simulations: phase-field simulations of solid-solid phase transformations in metal alloys and thermal simulations of metal additive manufacturing processes. For the phase-field simulations we present the results of a series of tests using the SUNDIALS time integration library in the MEUMAPPS C++ phase-field code. Using an implicit-explicit additive Runge-Kutta approach, we demonstrate improved performance for higher-order methods over the commonly used first-order implicit-explicit Euler scheme. We show that the implementation with SUNDIALS has no measurable additional overhead over a manually coded time integrator, maintaining a large GPU speedup (18x speedup on one full node of OLCF Frontier, comparing 8 GPUs to 64 CPU cores) and strong scaling to 1000s of GPUs. For the additive manufacturing application, we discuss results using the Adamantine and AdditiveFOAM part-scale thermal simulation codes. We demonstrate that proper handling of the volumetric heat source term and latent heat evolution is crucial for performant time integration of the otherwise diffusion-dominated system.

Stephen DeWitt
Oak Ridge National Laboratory
dewittsj@ornl.gov

David J. Gardner

Lawrence Livermore National Laboratory
gardner48@llnl.gov

John Coleman, Bruno Turcksin
Oak Ridge National Laboratory
colemanjs@ornl.gov, turcksinbr@ornl.gov

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

Balasubramaniam Radhakrishnan
Oak Ridge National Laboratory
radhakrishnb@ornl.gov

MS161

Operator-Splitting Approaches for Port-Hamiltonian Dae Systems in Multiphysics Simulation

In the simulation of differential-algebraic equations (DAEs), it is essential to employ numerical schemes that take into account the inherent structure and maintain explicit or hidden algebraic constraints without altering them. This paper focuses on operator-splitting techniques for coupled systems and aims at preserving the structure in the port-Hamiltonian framework. The talk explores two decomposition strategies: one considering the underlying coupled subsystem structure and the other addressing energy-associated properties such as conservation and dissipation. We show that for coupled index-1 DAEs with and without private index-2 variables, the splitting schemes on top of a dimension-reducing decomposition achieve the same convergence rate as in the case of ordinary differential equations. Additionally, we discuss an energy-associated decomposition for index-1 pH-DAEs and introduce generalized Cayley transforms to uphold energy conservation. The effectiveness of both strategies is evaluated using port-Hamiltonian benchmark examples from electric circuits.

Andreas Bartel
University of Wuppertal
bartel@math.uni-wuppertal.de

Malak Diab
Universitat Wuppertal
mdiab@uni-wuppertal.de

Andreas J. Frommer
Bergische Universitaet Wuppertal
Fachbereich Mathematik und Naturwissenschaften
frommer@uni-wuppertal.de

Michael Guenther
Bergische Universitat Wuppertal
guenther@uni-wuppertal.de

Nicole Marheineke
Trier University
marheineke@uni-trier.de

MS161

Stability and Efficiency Enhancements of Operator-Splitting Methods

Operator-splitting methods are widely used for the time integration of differential equations, especially those that arise from multi-scale or multi-physics models, because a monolithic approach may be inefficient or even infeasible. The most common operator-splitting methods are the first-order Lie–Trotter (or Godunov) and the second-order Strang (Strang–Marchuk) splitting methods. High-order splitting methods with real coefficients require backward-in-time integration in each operator and hence may be impacted by instability. However, besides the methods themselves, there are many other ancillary aspects to an overall operator-splitting method that are important in practice but often overlooked. For example, the order in which operators are integrated and the choice of sub-integration methods can significantly affect the performance of an operator-splitting method. In this paper, we design a new four-stage, third-order, 2-split operator-splitting method with seven sub-integrations and an optimized linear stability region. We then propose two general strategies to further improve its stability and efficiency for a specific problem, namely, to choose the ordering of operators to maximize linear stability and to choose low-order explicit sub-integrators for unstable sub-integrations. We demonstrate about a 40% improvement in the performance from the combined use of these strategies relative to standard implementations on a benchmark problem from cardiac electrophysiology.

Raymond J. Spiteri
University of Saskatchewan, Canada
Department of Computer Science
spiteri@cs.usask.ca

Siqi Wei
University of Saskatchewan
siqi.wei@usask.ca

Victoria Guenter
Department of Computer Science
University of Saskatchewan
v.guenter@usask.ca

MS161

Evaluation of IMEX, Multirate, and Operator Splitting Methods for Coupling Processes in a Cloud Microphysics Model

Cloud microphysical parameterizations play a crucial role in accurately simulating both regional weather and global climate. The parameterizations include numerous physical processes that are often coupled through some combination of parallel splitting, sequential splitting, and subcycling. Recent studies have shown that small modifications to the process coupling approach can substantially reduce the time discretization error, providing strong evidence that process splitting is a significant source of discretization error. To capitalize on this opportunity to achieve higher accuracy by improving process coupling, this work evaluates the performance of various implicit-explicit (IMEX), multirate, and operator splitting time integration approaches. The results presented are obtained on a recently developed cloud microphysics implementation (a subset of the P3 microphysics) that can rapidly leverage higher-order schemes by evaluating physical processes in a manner expected by time integrator software libraries (e.g., SUNDIALS). Overall, this work will demonstrate the benefits to be gained by the atmosphere community by utilizing higher-order time integration schemes in cloud microphysics parameterizations, as well as embracing implementations that can deploy those schemes. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-

07NA27344. LLNL-ABS-840173.

Justin Dong
Lawrence Livermore National Laboratory
dong9@llnl.gov

Sean P. Santos
Pacific Northwest National Laboratory
sean.santos@pnnl.gov

Steven B. Roberts, Carol S. Woodward
Lawrence Livermore National Laboratory
roberts115@llnl.gov, woodward6@llnl.gov

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

MS162

Diffuse Interface Methods for the StokesBiot Coupling

We consider the interaction between a poroelastic structure, described using the Biot model in primal form, and a free-flowing fluid, modelled with the time-dependent incompressible Stokes equations. We propose a diffuse interface model in which the governing equations in the Stokes and Biot regions are multiplied by an approximate indicator function of the respective subdomain, which smoothly transitions to zero in a diffuse region of width $\mathcal{O}(\epsilon)$ between them; this allows the weak forms to be integrated uniformly across the domain, and obviates tracking the subdomains or the interface between them. We prove convergence in weighted norms of a finite element discretisation of the diffuse interface model to the continuous diffuse model; here the weight is a power of the distance to the diffuse interface. We in turn prove convergence of the continuous diffuse model to the standard, sharp interface, model. Numerical examples verify the proven error estimates, and illustrate application of the method to fluid flow through a complex network, describing blood circulation in the circle of Willis.

Francis Aznaran, Martina Bukac
University of Notre Dame
faznaran@nd.edu, mbukac@nd.edu

Boris Muha
Department of Mathematics
Faculty of Science, University of Zagreb
borism@math.hr

Abner J. Salgado
Department of Mathematics
University of Tennessee
asalgad1@utk.edu

MS162

Stabilized High-Order Multirate Time-Integration for Multiphysics Problems Via Paired-Explicit Runge-Kutta Methods

We are going to discuss the application of multirate time-integration schemes to coupled multiphysics problems. Hyperbolic multiphysics problems are characterized by different characteristic speeds. Through the CFL condition, these varying characteristic speeds translate into (different) stability constraints. For classic explicit time-integration methods, the most restrictive (local) stability constraints determines the maximum (global) admissible timestep, which leaves room for efficiency gains. This motivates the use of stabilized multirate time-integration schemes to integrate such systems efficiently. We employ the paired-explicit Runge-Kutta schemes by Vermeire et al. which enable a seamless coupling of composing schemes. By construction, these schemes are conservative and consistent and are optimized for the discretized PDE. Recently, we extended these schemes to fourth order of accuracy. Another benefit of these schemes is that they are readily implemented in codes which employ a method-of-lines approach. In fact, by implementing these methods as partitioned Runge-Kutta schemes, one may leave the entire spatial discretization process untouched. We demonstrate significant savings in terms of flux evaluations resulting in speedup compared to a range of state-of-the-art optimized Runge-Kutta schemes while maintaining temporal accuracy. The presented applications cover coupled acoustic-inviscid flow simulations and Euler Equations with self-gravity.

Daniel Doehring
RWTH Aachen University
doehring@acom.rwth-aachen.de

MS162

Sedimentation in Suspensions of Rod-Like Particles: A Kinetic Fluid Model and Its Approximation by Hyperbolic Moment Equations

We focus on a model (Helzel & Tzavaras), which defines a sedimentation process in dilute suspension of rigid rod-like particles under the influence of gravity. This model is defined by a system of partial differential equations which couples a kinetic Smoluchowski equation to a macroscopic flow equation. The Smochluchowski equation is a drift-diffusion equation on the sphere, which needs to be solved at every point of the flow domain. Our goal is to construct an efficient numerical method for this high-dimensional problem (i.e., a 5-dimensional problem plus time). This will be done by replacing the high-dimensional Smoluchowski equation by a lower-dimensional system of moment equations. The moment equations form a hyperbolic system with a source term, which can be approximated using standard numerical methods such as LeVeque's wave propagation algorithm. A numerical solution for the coupled moment systems is computed using an operator splitting method. Our computational studies demonstrate that accurate results are achieved when a sufficiently large number of moment equations are used.

Bella My Phuong Quynh Duong
Heinrich-Heine-University
myduo100@uni-duesseldorf.de

MS162

High-fidelity simulations for continuum-rarefied gas flows past aerospike blunt body with R13 kinetic approach

High-speed vehicles face significant aerodynamic drag and aerothermal heating during atmospheric travel. Blunt body configuration is commonly used to reduce aerodynamic heating on the surface; however, it also causes a considerable increase in drag. While this design is effective for re-entry vehicles requiring high drag coefficients and minimal aerodynamic heating, it is less suited for other

high-speed applications. An aerospike positioned ahead of the blunt body is one strategy to reduce drag and aerodynamic heating in such vehicles. When the Knudsen number increases, the flow becomes more rarefied, eventually reaching the free molecular regime. Under these conditions, the Navier-Stokes-Fourier equations, without accounting for velocity slip and temperature jump boundary conditions, may not accurately predict drag and aerodynamic heating. To overcome these limitations, the gas kinetic models are often used. This study employs the regularized 13-moment (R13) model, which captures rarefaction effects with high accuracy. The R13 model, combined with Maxwell slip and Smoluchowski jump conditions, is used to simulate continuum-rarefied gas flow over an aerospike blunt body. High-fidelity simulations are conducted using an explicit mixed modal discontinuous Galerkin method on triangular meshes. The study explores the effects of varying the spike length to blunt body diameter ratio across a wide range of continuum-rarefied flows.

Satyvir Singh, Manuel Torrilho
RWTH Aachen University
singh@acom.rwth-aachen.de, mt@acom.rwth-aachen.de

MS162

On Nonlinear Closures for Moment Equations Based on Orthogonal Polynomials

In the present work, an approach to the moment closure problem on the basis of orthogonal polynomials derived from Gram matrices is proposed. Its properties are studied in the context of the moment closure problem arising in gas kinetic theory, for which the proposed approach is proven to have multiple attractive mathematical properties. Numerical studies are carried out for model gas particle distributions and the approach is compared to other moment closure methods, such as Grad's closure and the maximum-entropy method. The proposed "Gramian" closure is shown to provide very accurate results for a wide range of distribution functions.

Eda Yilmaz
Applied and Computational Mathematics,
RWTH AACHEN UNIVERSITY
yilmaz@acom.rwth-aachen.de

MS163

Data-driven Shock-capturing Indicator for High-order Methods Using Supervised Machine-learning Strategy

This paper presents a data-driven shock-capturing indicator for high-order methods using a supervised machine-learning strategy. High-order methods are susceptible to Gibbs-Wilbraham oscillations near shock waves. These challenges have driven the design of shock stabilization techniques, such as limiting and artificial viscosity methods. At the same time, detecting the shock location is a prerequisite for the shock stabilization. Existing shock detection methods are commonly based on stability criteria or non-smoothness features necessitating tunable parameters which entails a trade-off between accuracy and robustness. To achieve both accuracy and robustness without tunable parameters, we have developed a data-driven shock-capturing indicator using decision tree classifiers capable of handling unstructured-mixed meshes effectively. The training dataset is generated by projecting smooth and discontinuous analytic functions onto polynomial space, labeling elements intersecting discontinuities as troubled-cells. A gradient boosting algorithm is used to train decision tree ensemble models, resulting in accurate and efficient shock classifiers combined with proven shock stabilization techniques, such as the shock-capturing PID controller and the adaptive subcell limiting strategy. Through extensive numerical tests, we validated the proposed indicators performance in terms of accuracy, robustness, and cost-efficiency.

Dongseok Kim
Seoul National University, Seoul, Korea
dskim@snu.ac.kr

Jayeon Joo
Seoul National University
zoonature@snu.ac.kr

Hojun You
Sejong University
hojun.you@sejong.ac.kr

Chongam Kim
Seoul National University
chongam@snu.ac.kr

MS163

A Simple and Robust Adaptive Flux Reconstruction Shock-Capturing Method for the Nonlinearly Stable Flux Reconstruction High-Order Method

The flux reconstruction (FR) method has gained popularity in the research community. It recovers promising high-order methods through modally filtered correction fields, such as the Discontinuous Galerkin (DG) method, on unstructured grids over complex geometries. Under a class of energy stable flux reconstruction (ESFR) schemes, the FR method allows for larger timesteps than DG while ensuring stability for linear advection on linear elements. For nonlinear problems, split forms and entropy-conserving flux differencing approaches have become popular as they guarantee robustness for unsteady problems on coarse unstructured grids. Nonlinearly stable flux reconstruction (NSFR) combines the key properties of provable nonlinear stability and the increased timestep from ESFR. Using bound-preserving limiters, the NSFR method can be used for problems involving strong shocks. The limiters preserve positivity while maintaining a high-order of accuracy, but they do not mitigate nonphysical oscillations. This study aims to mitigate oscillations using an adaptive implementation of the FR correction parameter such that dissipation is added at locations where shocks are detected. The proposed method requires two components, a shock sensor and a consistent method of scaling the correction parameter. This study will investigate the ideal sensor and explore options for scaling the correction parameter. Results will include the 2D Double Mach Reflection and the 2D Astrophysical Mach Jet cases.

Shruthi Srinivasan, Siva Nadarajah
McGill University
sai.srinivasan@mail.mcgill.ca, siva.nadarajah@mcgill.ca

MS163

Entropy Filtering for High-Order Combustion

Discontinuous spectral element methods such as flux reconstruction (FR) combine the superior accuracy of spectral schemes with the geometric flexibility of finite volume schemes. As such they are an excellent candidate for scale

resolving flow simulations. However, on account of their use of polynomials to represent the solution state they encounter issues for flows which contain shocks. A particularly effective mitigation for this is the entropy filtering (EF) approach of Dzanic and Witherden which adaptively filters the solution to ensure positivity of density and pressure, and enforcement of a minimum entropy condition on the flow. In this talk I will outline recent work towards enhancing EF such that it can be applied to reacting flow problems. Specifically, we will discuss the importance of how entropy is computed. Moreover, we will also show how weighting the filter on a per-variable basis can lead to improved results within the context of multi-species flows.

Freddie D. Witherden
Texas A&M University
fdw@tamu.edu

MS164

Moment Neural Operator: a Digital Twin for Capturing Discontinuity

This paper introduces the Moment Neural Operator (MNO), a novel approach for learning mappings between discontinuous functions, addressing a significant limitation in existing Neural Operators (NO) that primarily target smooth functions. By leveraging moment-based discretization, which represents functions through their polynomial-weighted averages, MNO effectively captures sharp discontinuities without the unphysical oscillations observed in traditional methods such as the Fourier Neural Operator (FNO). The framework employs a neural network to approximate the eigen-decomposition of moment matrices, enabling the reconstruction of discontinuous target functions via the Christoffel-Darboux polynomial. We validate MNO on scalar conservation laws, including Burgers' equation and the LWR traffic flow model, demonstrating its capability to manage discontinuities and reduce high-frequency noise. Experimental results highlight MNO's competitive accuracy and superior handling of sharp transitions compared to FNO. This work underscores the potential of moment-based techniques for advancing neural operators and inspires future exploration of implicit representations in function approximation.

Qi Gao
Columbia University
qg2179@columbia.edu

Kuang Huang
Department of Mathematics
The Chinese University of Hong Kong
khuang@math.cuhk.edu.hk

Xuan Di
Department of Civil Engineering and Engineering Mechanics
Columbia University
sharon.di@columbia.edu

MS164

Digital Twins for Complex Systems: Challenges and Opportunities

A digital twin (DT) is a computational model of a physical system that continually updates its knowledge of the system by assimilating observational data, and in turn informs decisions and controls the system to achieve a desired goal, over a continually evolving time horizon. DTs have enormous potential to transform the role of models and data in decision-making for complex systems. At the same time, they present significant mathematical, statistical, and computational challenges. This stems from the complexity and scale of mathematical models describing many natural and engineered systems, the numerous uncertainties that underlie them, the complexity of observing systems and indirect and multimodal nature of the data they produce, and the critical societal impact of model-based decision making. In this talk I will discuss mathematical challenges and opportunities for realizing digital twins of complex systems.

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS164

Robust Digital Twin for Geological Carbon Storage

While recent work has shown that conditional neural networks can be used to effectuate prior-to-posterior mappings during ensemble Kalman filtering, their accuracy relies on the quality of the training ensemble, consisting of the simulated time-advanced nonlinear state and possibly nonlinear observations. Its precision also depends on the generative neural networks' ability to generalize. In this work, we will report progress on how to make our Digital Twin more robust with respect to distribution shifts and modeling errors, e.g. due to erroneous choices for the physics within the Forecast Step. To this end, we will follow a two-pronged approach consisting of augmenting the Forecast step, e.g. by including different models for the rock physics that map the Digital Twin's state to seismic properties, and by adding more physics during the prior-to-posterior mapping. In its current implementation, training of our Digital Twin's neural networks relies on relative low-fidelity amortization of the Forecast Ensemble, which is susceptible to distribution shifts. By conducting non-amortized inference informed by the physics and low-fidelity prior-to-posterior mapping, we propose an approach that is more robust than the current implementation of the Digital Twin.

Felix J. Herrmann
Georgia Institute of Technology
felix.herrmann@gatech.edu

Abhinav Gahlot
Georgia Tech
agahlot8@gatech.edu

Rafael Orozco, Ziyi Yin, Grant Bruer
Georgia Institute of Technology
rorozco@gatech.edu, ziyi.yin@gatech.edu, gbruer@gatech.edu

MS164

Greedy Sampling in High Dimensions via the Polytope Division Method

In the last two decades, model order reduction has been established as an important tool for the solution of high-dimensional parametrized partial differential equations. However, even with the development and success of new methods that exploit machine learning tools, the problem of offline training cost and, relatedly, offline sampling, remains. Most methods still rely on a random sampling of the parameter space, which especially in high-dimensional parameter spaces necessitates large amounts of full-order (expensive) training data. We explain the Polytope Divi-

sion Method (PDM), a greedy-type method, to determine where in the parameter domain to sample, therefore reducing the offline training cost. PDM splits the parameter space into polytopes and investigates the quality of the sample in the barycenter of each polytope.

Evie Nielen
Eindhoven University of Technology
e.t.e.nielen@tue.nl

Oliver Tse
Eindhoven University of Technology (Tu/e)
o.t.c.tse@tue.nl

Karen Veroy
Department of Mathematics and Computer Science
Eindhoven University of Technology
k.p.veroy@tue.nl

MS164

Optimal Power Procurement for Green Cellular Wireless Networks under Uncertainty

This work derives an optimised power procurement strategy for cellular networks powered by uncertain renewable energy sources and conventional fossil-fuel based electricity to minimise power procurement costs, subject to uncertainty in wireless channels. The approach is based on the use of stochastic optimal control with embedded Markovian dynamics for uncertain renewable energy and wireless channels. A data-driven, parametric stochastic differential equation (SDE) is used to model renewable energy, with parameters calibrated using German wind production data. An SDE is derived to model wireless channel fading according to the Nakagami fading model. Convex approximations are used to handle probabilistic constraints arising from the network quality of service threshold that needs to be satisfied. The solution procedure involves developing a computationally optimal numerical scheme to solve the associated non-linear Hamilton-Jacobi-Bellman equation, building an efficient stochastic subgradient method for the associated, non-smooth dual problem and constructing a fixed-point iteration scheme to simultaneously tackle probabilistic constraints. Lastly, we build a rolling horizon framework to continuously incorporate day-ahead forecasts.

Shyam Mohan Subbiah Pillai
RWTH Aachen
subbiah@uq.rwth-aachen.de

Nadhir Ben Rached
University of Leeds
n.benrached@leeds.ac.uk

Raul Tempone
KAUST
RWTH Aachen University
raul.tempone@kaust.edu.sa

MS165

Scientific ML for Streamflow Simulations: Evaluating ODE for Performance

Recent developments in neural network (NN) models, particularly Long Short-Term Memory (LSTM) networks, have achieved impressive accuracy in streamflow simulations. However, the interpretability of NNs remains a concern, often leading them to be labeled as "black boxes". Hybrid models, which offer a promising approach by combining the strengths of machine learning and physical modeling, have garnered increasing attention. In these models, NNs are embedded within ordinary differential equations (ODEs), replacing traditional descriptions of physical processes. This approach aims to preserve the interpretability of physical models while leveraging the enhanced performance of NNs. A key requirement for the success of such hybrid models is the use of gradient-based optimization techniques such as backpropagation, allowing for the seamless integration of physical equations with NN layers and facilitating the training of coupled models. The performance of hybrid models depends on four key components: the data, the NN architecture, the ODEs, and the ODE solver. The impact of the ODE solver is unclear. Our research is focused on assessing the impact of different ODE solvers on hybrid models, particularly in terms of computational accuracy and performance metrics such as time and memory usage. Understanding how these solvers influence model outcomes is crucial for advancing the applicability of hybrid models in hydrology.

Jesus P. Curbelo
Department of Computer Science
University of Saskatchewan
jesus.curbelo@usask.ca

Raymond J. Spiteri
University of Saskatchewan, Canada
Department of Computer Science
spiteri@cs.usask.ca

Martyn Clark
Department of Civil Engineering, Schulich School of Engineer
martyn.clark@ucalgary.ca

MS165

Manifold Neural ODEs: Scalable Training of Semi-Explicit Neural ODEs with Arbitrary Constraints

Physical systems often have known conservation quantities, such as conservation of momentum, charge, and energy. Despite the promise of scientific machine learning (SciML) in combining data-driven techniques with mechanistic modeling, existing approaches for incorporating conservation laws in neural differential equations (NDEs) face significant limitations. Scalability issues and poor numerical properties prevent these neural models from being used for modeling physical systems with complicated conservation laws. In this talk, we will review some classical methods on how conservation laws are satisfied in index-1 semi-explicit differential-algebraic equations (DAEs). We shall build upon this to introduce Manifold Neural ODEs, a novel method that ensures that learned NDEs obey fundamental system constraints (conservation laws). We will explore the SciML ecosystem we have built around neural networks (Lux.jl), nonlinear solvers (NonlinearSolve.jl) and differential equations (DifferentialEquations.jl) that form the foundation of our method. Using problems ranging from rigid-body dynamics to orbital dynamics we will show how our method improves upon past research in scaling and overall performance.

Avik Pal, Christopher V. Rackauckas
Massachusetts Institute of Technology

avikpal@mit.edu, accounts@chrisrackauckas.com

MS165

Linearly Stabilized Schemes for the Time Integration of Stiff Nonlinear PDEs

In many applications, the governing PDE to be solved numerically contains a stiff component. When this component is linear, the use of an implicit time stepping method that is unencumbered by stability restrictions is preferred. On the other hand, if the stiff component is nonlinear, the complexity and cost per step of using an implicit method is heightened, and the use of explicit methods may be preferred. In this talk, we consider new and existing linearly stabilized schemes for the purpose of integrating stiff nonlinear PDEs in time. These schemes compute the nonlinear term explicitly and, at the cost of solving a linear system with a matrix that is fixed throughout, are unconditionally stable, thus combining the advantages of explicit and implicit methods. Applications are presented to illustrate the use of these methods.

Steven Ruuth
Simon Fraser University
steven_ruuth@sfu.ca

MS165

Compositional Physics-Informed Neural Network

Physics-Informed Neural Networks (PINN) have recently been applied to solve many forward and inverse problems. Traditionally, solving time-dependent problems using PINN involves sampling from a residual loss function on a fixed space-time domain, which can result in poor generalizability beyond its training time domain. We introduce Compositional Physics-Informed Neural Network (CPINN) where forward flow maps are learned by a variant of PINN while preserving compositional structure of flow maps. We show the error of CPINN beyond the training time can be bounded by the training error on a fixed training time interval and sampling error of the residual. We verify our error bound and the efficacy of long term prediction of CPINN on a variety of examples.

Pratham Lalwani, Andy Wan
UC Merced
plalwani@ucmerced.edu, andywan@ucmerced.edu

MS166

Parameterized Wasserstein Gradient Flow with Application in Porous-Medium Equation

Wasserstein gradient flow (WGF) is a powerful tool for understanding and analyzing density evolution processes. In the seminal work by Jordan, Kinderlehrer, and Otto, it was shown that the Fokker-Planck equation (FPE) is the gradient flow of the relative entropy functional under the Wasserstein metric. Since then, WGFs have shown extensive applications in optimal transport theory, optimization problems and more. However, numerical computation of WGF remains a challenging problem, especially when the state space is of high dimension. Furthermore, it is often desirable to find a sampler that generates samples following the solution of WGF rather than the actual density function solving WGF in many real-world statistics and machine learning applications. We develop a fast and scalable numerical approach to solve Wasserstein gradient flows (WGFs), which is particularly suitable for high-dimensional cases. We demonstrate the efficiency of the algorithm with solving porous-medium equation and provide samples in high-dimensional space and provide error estimates in Wasserstein metric.

Yijie Jin
Georgia Institute of Technology
yijiejin@gatech.edu

MS166

Design and Scalability of Flow-based Samplers for Bayesian Inference

In Bayesian inverse problems, we use a prior distribution to encode our knowledge about a state variable of interest and transform that prior using data to obtain a posterior distribution. Sampling this posterior is a key task for quantifying uncertainty in estimates of and predictions based on the state variable. In this talk, we discuss the dynamic transport approach to sampling the posterior where samples from the prior are transformed into posterior samples by ODE dynamics. We design dynamic transport samplers that begins with an explicit path of distributions between the prior and posterior and identifies a gradient velocity field via the solution of Poisson equations, which will cause the distribution of the samples to follow this path. The gradient structure of the velocity field has an optimal transport interpretation, while particular choices of path can lend Fisher-Rao gradient flow structure to the sampler. We solve the weak form of the Poisson equations in reproducing kernel Hilbert space, leading to tractable, gradient-free, closed-form interacting particle systems for sampling. We discuss implementation strategies that allow these interacting particle systems to scale, including dimension reduction, random Fourier feature approximation, and adaptive time-stepping based on the Wasserstein metric derivative. We demonstrate the utility of these interacting particle system samplers for Bayesian inverse problems in real-world application domains.

Aimee Maurais
Massachusetts Institute of Technology
maurais@mit.edu

MS166

Efficient Computational Methods for Wasserstein Natural Gradient Descent

Natural gradient descent (NGD) is a well-known preconditioning technique that incorporates the geometry of the forward model space to accelerate gradient-based optimization techniques in inverse and learning problems. One such famous geometry is the optimal transportation (OT) or the Wasserstein geometry, which is useful when recovering or learning probability measures. One of the critical challenges in NGD is the preconditioning cost. If performed naively, this cost is particularly taxing for the OT geometry due to the high computational cost of OT distances. In this talk, I'll present an efficient way of performing large-scale NGD with a particular emphasis on OT geometry. This is a joint work with Yunan Yang (Cornell) and Wanzhou Lei (Brown Grad School).

Levon Nurbekyan
Department of Mathematics
Emory University
levon.nurbekyan@emory.edu

Wanzhou Lei

Harvard University
wanzhoulei@g.harvard.edu

Yunan Yang
Cornell University
yunan.yang@cornell.edu

MS166

Volumetric Extension Methods for Computing the Optimal Transport Mapping on Surfaces

We first introduce the Monge problem of Optimal Transport on compact and orientable surfaces $\Gamma \subset \mathbb{R}^3$. This computation yields a pushforward map m from a source probability measure μ to a target probability measure ν on Γ that minimizes a cost functional. We then show how to properly extend the problem to a tubular neighborhood T_ϵ of Γ, by carefully defining an extended cost functional and extended probability measures. The resulting formulation defines a new Optimal Transport problem on T_ϵ whose solution m_ϵ can be used to find the Monge map m for the Optimal Transport problem on Γ. We will then be able to devise a simple discretization of the Optimal Transport PDE using a Cartesian grid on T_ϵ. We conclude with examples showing the success of computations on real-world applications.

Axel Turnquist
University of Texas at Austin
agrt@utexas.edu

MS167

Efficient Measure Transport for Bayesian Inference via LazyDINOs

We introduce LazyDINO, a variational inference method for fast and scalable solutions to large-scale nonlinear Bayesian inverse problems with costly parameter-to-observable (PtO) maps. LazyDINO efficiently optimizes a low-dimensional, structure-exploiting transport map, known as a lazy map [Brennan et al., NeurIPS, (2020)], using a fast ridge function surrogate of the PtO map. This optimized map then generates approximate posterior samples based on observed data. The surrogate is constructed by minimizing error bounds in the lazy map optimization and posterior approximation. When neural networks parameterize the surrogate, it becomes a derivative-informed neural operator (DINO) [O'Leary-Roseberry et al., JCP, (2024)], trained using joint samples of the PtO map and its Jacobian. Our results show that LazyDINO consistently outperforms existing methods, including Laplace approximation, lazy maps, and neural posterior estimation, in solving complex infinite-dimensional PDE-constrained Bayesian inverse problems. Moreover, with the same training cost, LazyDINO is 1025 times more accurate in amortized posterior approximation compared to LazyNO, which uses conventional operator learning.

Joshua Chen
The University of Texas at Austin
joshua@oden.utexas.edu

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

Michael C. Brennan
MIT, CCSE
mcbrenn@mit.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS167

Derivative-Informed Neural Operators for Managing Risks Associated with Rare Events

Rare events often represent catastrophic failures of engineered systems that are to be avoided in its design and operation. However, model-based estimation of such events via conventional sampling methods may require a prohibitive number of model evaluations. In this talk, we discuss novel neural operator methodologies that aim to accelerate estimation and optimization of rare events involving PDE models with uncertain parameters. The methods are then illustrated over a range of physically relevant examples, including fluid flow and nonlinear structural mechanics.

Dingcheng Luo
Oden Institute for Computational Engineering and Sciences,
The University of Texas at Austin
dc.luo@utexas.edu

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

Joshua Chen
The University of Texas at Austin
joshua@oden.utexas.edu

Thomas O'Leary-Roseberry
The University of Texas at Austin
Oden Institute for Computational and Engineering Sciences
tom@oden.utexas.edu

Harsha Honnappa
Purdue University
honnappa@purdue.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS167

Scalability of the Second-Order Reliability Method for Extreme Event Estimation of Stochastic Differential Equations

Estimating the probability of extreme events, which are often associated with rare system failures or catastrophes, is an important problem in many scientific and engineering disciplines. Asymptotically, such probabilities can be estimated using a Laplace approximation, which is referred

to as the second-order reliability method in engineering, or precise large deviation theory in mathematics. The method involves (i) solving an optimization problem - finding the most likely realization of the random parameter that leads to a prescribed outcome - for the exponential leading-order asymptotics of the extreme event probability, and (ii) subsequently calculating a determinant to get a prefactor for an asymptotically sharp estimate. In this talk, I will discuss how to carry out both of these steps numerically in a scalable way for extreme events in stochastic differential equations with additive or multiplicative Brownian noise. This is an infinite-dimensional problem, with randomness induced by the Brownian motion. In particular, I will highlight the necessity to treat the determinant calculation in step (ii) correctly from an infinite-dimensional point of view to get scalability. This leads to either a Fredholm or Carleman-Fredholm determinant computation, depending on whether the second variation of the noise-to-event map is trace-class or only Hilbert-Schmidt. The talk is based on, and extends, the results presented in [Schorlepp et al, Stat. Comput. 33(6), 137 (2023)].

Timo Schorlepp
Ruhr University Bochum
timo.schorlepp@nyu.edu

MS167

Large Deviation Theory-Informed Importance Sampling for Rare Event Estimation and Control

Rare and extreme events like hurricanes, energy grid blackouts, dam breaks, earthquakes, and pandemics are infrequent but have severe consequences. Because estimating the probability of such events can inform strategies that mitigate their effects, scientists must develop methods to study the distribution tail of these occurrences. However, calculating small probabilities is hard, particularly when involving complex dynamics and high-dimensional random variables. In this talk, I will discuss our proposed method for the accurate estimation of rare event or failure probabilities for expensive-to-evaluate numerical models in high dimensions, and its application to rare event control. The proposed approach combines ideas from large deviation theory and adaptive importance sampling. The importance sampler uses a cross-entropy method to find an optimal Gaussian biasing distribution, and reuses all samples made throughout the process for both, the target probability estimation and for updating the biasing distributions. Large deviation theory is used to find a good initial biasing distribution through the solution of an optimization problem. Additionally, it is used to identify a low-dimensional subspace that is most informative of the rare event probability. We compare the method with a state-of-the-art cross-entropy-based importance sampling scheme using examples including a tsunami problem.

Shanyin Tong
Columbia University
st3503@columbia.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS168

High-Fidelity Simulation of Low-Pressure Turbines Using a Spectral Element Framework

Understanding the flow physics in low-pressure gas turbines associated with upstream wakes, end-wall effects, and free stream turbulence require high-fidelity numerical methods that can both capture the geometric complexities, but also incorporate the unsteadiness exhibited due to the rotor-stator interactions. To investigate these effects thoroughly we carry out large-scale simulations with our recent spectral element framework, enabling analysis of the complex flow phenomena within the turbine and the mechanisms driving laminar-turbulence transition. Due to the improved performance and workflow of our new framework we are able to explore a larger set of the parameters the turbine is subject to during operational conditions, giving larger insights into the secondary flows within the low-pressure turbine environment than before and more accurately predicting losses as well as the dynamic stress on the blades. We present an overarching picture of our simulations with applications in turbomachinery and especially highlight the workflow enabled through our spectral element framework for high-fidelity simulation of low-pressure turbines on recent large-scale HPC systems.

Martin Karp, Victor Baconnet, Ardeshir Hanifi
KTH Royal Institute of Technology
makarp@kth.se, baconnet@kth.se, hanifi@kth.se

Dan Henningson
Department of Mechanics,
Royal Institute of Technology, Stockholm
henning@mech.kth.se

MS168

Machine-Learning Approaches for Enhancing Convergence of Large-Scale Solution Strategies

We introduce a new class of hybrid preconditioners for solving parametric linear systems of equations. The proposed preconditioners are constructed by hybridizing the deep operator network, namely DeepONet, with standard iterative methods. Exploiting the spectral bias, DeepONet-based components are harnessed to address low-frequency error components, while conventional iterative methods are employed to mitigate high-frequency error components. Our preconditioning framework utilizes the basis functions extracted from pre-trained DeepONet to construct a map to a smaller subspace, in which the low-frequency component of the error can be effectively eliminated. Our numerical results demonstrate that the proposed approach enhances the convergence of Krylov methods by a large margin compared to standard non-hybrid preconditioning strategies. Moreover, the proposed hybrid preconditioners exhibit robustness across a wide range of model parameters and problem resolutions.

Alena Kopanicakova
Brown University
alena.kopanicakova@usi.ch

MS168

Non-Linear Preconditioning Techniques for Large Scale Applications

The phase-field method has become widely adopted in computational mechanics for fracture modeling due to its ability to accurately capture crack initiation, propagation, branching, and merging without the need for explicit criteria. This approach simplifies the compu-

tational process by avoiding the frequent remeshing required as cracks evolve. However, this approach introduces new difficulties, particularly related to the highly nonlinear, non-convex, and non-smooth nature of the underlying energy landscape. To address these issues, we propose a nonlinear additive/multiplicative field-split preconditioned Newton method. We propose a field-split-based additive/multiplicative Schwarz preconditioned Newton method to solve the fracture problem by employing a right preconditioner that can handle inequality constraints. The effectiveness and robustness of our proposed method will be validated through benchmarking against the traditional alternate minimization technique.

Hardik Kothari
Università della Svizzera italiana
hardik.kothari@usi.ch

MS168

A new higher-order convergent derivation of the Fick-Jacobs equation and applications to computational neuro-engineering

Modeling physical processes in network structures, e.g. pipe systems, cardiovascular networks or neurons can quickly become an extremely expensive computational problem when modeling such networks in three dimensions. To reduce this cost, we are interested in modeling networks as collections of one-dimensional cables by reducing the three-dimensional equations to one dimension. This dimension reduction must take into account the non-constant radius of the network branches. Jacobs proposed an intuitive and straight-forward derivation of this reduction in 1935 based on Fick's laws, but it wasn't until 1992 that Zwanzig further studied this Fick-Jacobs model and demonstrated that it was unstable for domains with a quickly changing radius. Since then, multiple corrections on the diffusion coefficient or temporal derivative have been proposed to stabilize the Fick-Jacobs model for steep domains, with limited success. We propose a new simple derivation of the Fick-Jacobs model utilizing a high-order expansion of the computed flux that, when applied in a discrete computational setting, produces second-order spatial convergence and provides a much more accurate numerical solution at a wide range of radial gradients than previous correction methods.

Gillian Queisser, Zachary Miksis
Temple University
gillian.queisser@temple.edu, miksis@temple.edu

MS168

Immersed Domain Approach for Fluid Structure Contact Interaction

Fluid-structure interaction (FSI) has gained significant traction in recent decades, with applications spanning various disciplines, including geophysics and biomedicine. In FSI, computational techniques are defined by the choice of discrete domain representation, falling into two main categories: "boundary-fitted" or "non-boundary-fitted" meshes. Boundary-fitted methods offer high accuracy, but their viability is limited in the presence of large solid deformations. Non-boundary-fitted methods maintain separate and non-matching fluid and structure meshes. Here, structure and fluid are described within a Lagrangian and Eulerian framework, respectively. However, higher mesh resolution is to maintain comparable accuracy, making parallel computing necessary. We present an immersed domain approach for the numerical solution of fluid-structure-contact-interaction (FSCI) problems. The fluid and structure are coupled within the whole overlapping volume, while the different structures in contact are coupled on their surfaces. These couplings are achieved with the method of dual Lagrange multipliers. The nonlinear solution procedure is achieved by solving a sequence of the statically condensed system where only the fluid variables are unknown. We show our general algorithmic framework and our primary parallel computing tools and present a fully coupled simulation of a liquid diaphragm pump with contact interaction between elastic valves displaced by the fluid and the valve seats.

Patrick Zulian
UniDistance Suisse
patrick.zulian@unidistance.ch

Maria Nestola
Institute of Computational Science, USI, Lugano, Switzerland
nestom@usi.ch

Rolf Krause, Rolf Krause
KAUST / UniDistance Switzerland
krause.r@gmx.net, krause.r@gmx.net

MS169

Randomized Countsketch Qr Factorization for Wide and Short Matrices

In this talk we introduce randomized countsketch QR on wide and short matrices. Our approach addresses the computational challenge of the QRCP, particularly in the costly pivot identification phase, by compressing the matrix while preserving its geometric properties. While previous work done on randomized QRCP has focused on reducing the number of rows, our approach complements these efforts by then reducing the number of columns using the countsketch matrix. Applying then QRCP on this sketched matrix allows us to identify the most influential columns in the original matrix. We show that this proposed method reveals the numerical rank of a matrix in an analogous way to the traditional QRCP. Through experimental validation on various matrix types, we evaluate the effectiveness of our chosen pivots.

Israa Fakih
Ecole polytechnique federale de Lausanne
israa.fakih@epfl.ch

Laura Grigori
EPFL and PSI, Switzerland
laura.grigori@epfl.ch

MS169

Sparse Cholesky Factorization Utilizing GPUs

The solution of sparse symmetric positive definite linear systems is an important computational kernel in large-scale scientific and engineering modeling and simulation. We will solve the linear systems using a direct method, in which a Cholesky factorization of the coefficient matrix is performed using a right looking approach and the resulting triangular factors are used to compute the solution. Sparse Cholesky factorization is compute intensive. In this work we investigate techniques for reducing the factorization time in sparse Cholesky factorization by offloading some of the dense matrix operations on a GPU. We will

describe the techniques we have considered. We achieved up to 4x speedup compared to the CPU-only version.

M. Ozan Karsavuran, Esmond G. Ng
Lawrence Berkeley National Laboratory
mokarsavuran@lbl.gov, egng@lbl.gov

BARRY W. Peyton
Dalton State College
bpeyton@daltonstate.edu

MS169

2D Sparse Communication Methods for Maximum Weight Matching Applications on GPUs

Large-scale distributed computations on graphs and matrices are often burdened by complex communication requirements and their associated overheads. One method to mitigate these overheads is to consider 2D distributions, where a given rank in the distributed computation only considers a small block of the (adjacency) matrix. Such a distribution enables effective load balance and the ability to utilize a hierarchical communication scheme with desirable scaling behavior. However, for irregular computations, which are often driven by a 'wave front' of updates, such communication schemes can be inefficient. In this work, we discuss our approach of addressing this problem through the development of 'sparse communication methods' for 2D distributions, specifically targeting GPU-based systems. We implement and optimize our methods, and we demonstrate their applicability to the maximum weight matching problem, commonly used in the context of multilevel coarsening within applications like graph partitioning. We further demonstrate that our method scales to extremely large meshes, strong scaling to 256 GPUs on matrices with tens of billions of nonzeros. Additionally, we analyze our matching algorithm from the lens of efficiency and quality in the context of graph partitioning applications.

Michael Mandulak
RPI
mandum@rpi.edu

George Slota
Rensselaer Polytechnic Institute
gmslota@gmail.com

MS169

What Makes "scidac" Scidac? The Future of Fastmath and Application Partnerships

Now in its fifth instantiation, the DOE Office of Science's Scientific Discovery through Advanced Computing (SciDAC) Program has proven to be an enormously successful motivator of progress in scientific computing. While earlier versions of the program varied somewhat in structure, since SciDAC3 there have been two components. Application Partnerships in specific domain science areas have been where the rubber meets the road in SciDAC application scientists and computational scientists, mathematicians, and computer scientists partnering to make major improvements in domain science capabilities. In parallel with the partnerships, the SciDAC Institutes work on crosscutting research that is broader in scope than the individual partnerships. There are currently two SciDAC institutes broadly speaking, the FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics) Institute focuses on applied mathematics and its expression in scalable and performant mathematical libraries, while the RAPIDS Institutes focus is on Computer Science, Data, and Artificial Intelligence. In this talk, we provide overall background and context for the FASTMath-themed talks in this minisymposium. We present a short background of the evolution of the Institutes in SciDAC, lessons learned from successful (and unsuccessful) engagement of the institutes with the partnerships, and discuss some potential future directions as SciDAC moves into its sixth installment in 2025.

Dan Martin, Esmond G. Ng
Lawrence Berkeley National Laboratory
dfmartin@lbl.gov, egng@lbl.gov

MS169

Mixed Precision Algorithms in hypre

The hypre software library provides parallel solvers and preconditioners for a variety of high- performance computing architectures. Recently, the use of mixed precision in algorithms has become of high interest since it provides reduced memory use and faster performance for lower precision operations. Development of hypre started more than twenty-five years ago, and generally focused on using double precision. While hypre can also be configured at single precision, it was not originally designed to allow the use of several different precisions in combination. Recently, the hypre team developed the capability of using mixed precision in hypre. This talk will describe some mixed precision algebraic multigrid methods developed in hypre and present some results.

Ulrike M. Yang, Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
yang11@llnl.gov, oseikuffuor1@llnl.gov

MS170

Randomized Time Integrators for Sequential-in-Time Training of Neural Network Parametrizations with Neural Galerkin Schemes

Sequential-in-time training of parametrizations such as neural networks with Neural Galerkin and related schemes can be challenging due to the tangent-space collapse phenomenon, which is a form of overfitting. In particular, applying standard time integrators that truncate information corresponding to small singular values over time can lead to non-negligible accumulation of errors that lead to instabilities and a loss of accuracy. In this work, we discuss time integration with random subspace embeddings and show that it is an appealing approach in terms of empirical efficiency and accuracy in sequential-in-time training.

Yijun Dong
New York University
yd1319@nyu.edu

Yifan Chen
Caltech and NYU
yifan.chen@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University

pehersto@cims.nyu.edu

MS170

Forward and Inverse Modeling Strategies with SciML for Real-time Tsunami Forecasting

Most problems in earth systems are characterized by partial observability and high dimensionality. This complicates the use of many SciML methods that operate best on systems with near-complete observability and struggle with both forward and inverse modeling for applications in earth systems. Furthermore, forecasting alone is insufficient for real-world disaster modeling, as it also needs robust uncertainty analysis and model stability to be usable. This talk presents novel approaches to disaster modeling with SciML in both forward and inverse modeling scenarios. We frame a problem with real-world data, historically consistent conditions, and accurate geometry for the problem of rapid tsunami forecasting after a seismic event. Our forward modeling method based on the Neural Galerkin projection can rapidly forecast tsunami wave height evolution and wave arrival time at various coastal locations of interest. Our inverse modeling approach based on the Senseiver can make high-resolution, full-field wave reconstructions given sparse measurements corresponding to actively deployed ocean buoys. Crucially, we demonstrate our ability to provide error bars and variability ranges for our forecasts to iterate over various scenarios and uncertainty in initial conditions. As our methods are not restricted to tsunami modeling, we frame our approach within the larger context of earth system modeling and the challenges and opportunities for AI in natural disaster management.

Arvind T. Mohan
Los Alamos National Lab
arvindm@lanl.gov

Edward McDugald
University of Arizona, United States
Los Alamos National Laboratory, United States
emcdugald@lanl.gov

Shane Coffing
Los Alamos National Laboratory
sxc@lanl.gov

John Tipton
Los Alamos National Laboratory, United States
jtipton@lanl.gov

Javier Santos
Los Alamos National Laboratory
jesantos@lanl.gov

Darren Engwirda
CSIRO, Australia
d.engwirda@gmail.com

MS170

When Big Neural Networks Are Not Enough: Physics, Multifidelity, Kernels

Modern machine learning has shown remarkable promise in multiple applications. However, brute force use of neural networks, even when they have huge numbers of trainable parameters, can fail to provide highly accurate predictions for problems in the physical sciences. We present a collection of ideas about how enforcing physics, exploiting multifidelity knowledge and the kernel representation of neural networks can lead to significant increase in efficiency and/or accuracy. Various examples are used to illustrate the ideas.

Amanda Howard, Wenqian Chen
Pacific Northwest National Laboratory
amanda.howard@pnnl.gov, wenqian.chen@pnnl.gov

Saad Qadeer
University of North Carolina Chapel Hill
saad.qadeer@pnnl.gov

Shady Ahmed, Panos Stinis
Pacific Northwest National Laboratory
shady.ahmed@pnnl.gov, panos.stinis@pnnl.gov

MS170

Generative Modeling of Complex Stochastic Spatiotemporal Dynamics

Modeling complex stochastic spatiotemporal dynamics, such as turbulent flows, is challenging due to their chaotic and random nature. Traditional numerical simulations, like eddy-resolved methods, offer detailed insights but are computationally expensive and limited in scalability. While deep learning-based surrogate models have emerged as alternatives, they often struggle to capture the inherent stochasticity of these dynamics due to their deterministic frameworks. This study introduces the Conditional Neural Field-Based Latent Diffusion (CoNFiLD) Model, a novel generative framework designed to effectively model complex spatiotemporal dynamics. CoNFiLD integrates a conditional neural field with a latent diffusion model, leveraging probabilistic diffusion processes for memory-efficient generation of diverse dynamics. The framework employs Bayesian conditional sampling, enabling the unconditional model to generate conditioned outputs across various scenarios without retraining, enhancing adaptability and versatility. Through rigorous numerical experiments, CoNFiLD demonstrates its effectiveness in replicating stochastic and chaotic dynamics, such as turbulence, overcoming the limitations of conventional deep learning models. The results underscore the models potential as a robust, efficient, and scalable approach for simulating complex spatiotemporal phenomena, setting a new standard in generative modeling for stochastic systems.

Meet Parikh, Xiantao Fan, Xin-Yang Liu, Pan Du, Jian-Xun Wang
University of Notre Dame
mparikh2@nd.edu, fanxiantao1@gmail.com, xliu28@nd.edu, pdu@nd.edu, jwang33@nd.edu

MS171

On Combining Neural Operator and Denoising Diffusion Model to Predict Material Response

Finite element analysis, a common approach for solving solid mechanics problems, is typically associated with high computational costs. Moreover, due to the heterogeneity and complexity of materials, measurement of response is often computationally challenging. To alleviate this, deep learning (DL) techniques are being utilized to predict the mechanical behavior of materials. However, many existing models struggle with generalization and demand extensive data preprocessing to optimize performance, particularly when the material response spans wide ranges. To address this, we propose a DL framework that utilizes ma-

terial geometry and loading to predict material responses across multiple geometries. The framework comprises two components: a U-Net based conditional denoising diffusion probabilistic model and a DeepONet. We leverage the generative power of the denoising diffusion model to estimate the normalized stress maps, while the DeepONet learns the scaling required to rescale the underlying stress maps. The outputs from these models are combined to generate accurate stress fields. We assess the performance of this framework through a series of experiments involving material geometries of varying complexity.

Purna Vindhya Kota, Meer Mehran Rashid, Lori Graham-Brady, Somdatta Goswami
Johns Hopkins University
pkota2@jh.edu, mrashi12@jhu.edu, lori@jhu.edu, sgoswam4@jhu.edu

MS171

Integrating Neural Operators with Diffusion Models Improves Spectral Representation in Turbulence Modeling

We integrate neural operators with diffusion models to address the spectral limitations of neural operators in surrogate modeling of turbulent flows. While neural operators offer computational efficiency, they exhibit deficiencies in capturing high-frequency flow dynamics, resulting in overly smooth approximations. To overcome this, we condition diffusion models on neural operators to enhance the resolution of turbulent structures. Our approach is validated for different neural operators on diverse datasets, including a high Reynolds number jet flow simulation and experimental Schlieren velocimetry. The proposed method significantly improves the alignment of predicted energy spectra with true distributions compared to neural operators alone. This enables the diffusion models to stabilize longer forecasts through diffusion-corrected autoregressive rollouts, as we demonstrate in this work. Additionally, proper orthogonal decomposition analysis demonstrates enhanced spectral fidelity in space-time. This work establishes a new paradigm for combining generative models with neural operators to advance surrogate modeling of turbulent systems, and it can be used in other scientific applications that involve microstructure and high-frequency content. See our project page: `https://vivekoommen.github.io/NO_DM/`

Vivek Oommen, Aniruddha Bora, Zhen Zhang
Brown University
vivek_oommen@brown.edu, aniruddha_bora@brown.edu, zhen_zhang2@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS172

Neural Universal Differential Equation Hypernetwork Surrogates for Agent-Based Disease Models

Agent-based models (ABMs) provide a rich framework for modeling outbreaks and interventions in epidemiology by explicitly accounting for diverse individual interactions and environments. However, these models are usually stochastic and highly parameterized, requiring precise calibration for accurate predictions. When considering realistic numbers of agents, this high dimensional calibration can be computationally prohibitive thereby requiring efficient surrogate models to replace expensive model evaluations. In this talk, we present ABM surrogate models using Neural Ordinary Differential Equations (NODEs), machine learning (ML) models based on differential equations that can be calibrated to data using a continuous analogue of backpropagation. The goal of the NODE surrogates is to reproduce the time - series outputs of the ABM as a function of several parameters. To achieve this, we take the approach of a hypernetwork using one ML model to predict the weights of another. An initial neural network takes as input the parameters of the ABM model and outputs weights defining a NODE that produces subsequent time series outputs. We compare this methodology to other surrogate modeling approaches including random forests.

Wyatt H. Bridgman, Connor Robertson
Sandia National Laboratories
whbridg@sandia.gov, cjrobe@sandia.gov

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

MS172

Latent Space Dynamics Identification of Path Integral Molecular Dynamics Data

Classical molecular dynamics (MD) is frequently based on (simple) pair-potentials. This simplification ignores the influence of the neighborhood composition, hence, has limited predictive power for multi-component systems. One possibility for enhanced accuracy is offered by Path-integral MD (PIMD). PIMD solves the quantum Hamiltonian with only a small number of approximations that hold at high temperature resulting in increased accuracy for multi-component systems. However, the practical applicability of PIMD is substantially limited by the required incorporation of electronic degrees of freedom and the associated considerable increase in overall computational effort. For example, for the specific problem of ablator material mixing into the deuterium-tritium (DT) fuel region manifesting in thermonuclear fuel compression, explicit PIMD simulations can only reasonably simulate a fraction of the total necessary to resolve meaningful mixing. To achieve sensible temporal and spatial scales for this application example, we propose to train an auto-encoder on PIMD data of the temporal evolution of mass density across the bimaterial interface. The resulting latent space representation forms the basis for inexpensive temporal forecasting done via Dynamic Mode decomposition or Sparse Identification of Nonlinear Dynamics.

Jannik Eisenlohr
Michigan State University
eisenl10@msu.edu

Youngsoo Choi
Lawrence Livermore National Laboratory
choi15@llnl.gov

Michael Murillo, Luciano Silvestri
Michigan State University
murillom@msu.edu, silves28@msu.edu

MS172

Dimensionality Reduction in Viscoelastic Fluid Dynamics: Insights from Machine Learning Models

In this study, we investigate nonlinear dimensionality re-

duction in viscoelastic fluids considering energy-compatible families of kernel functions tailored to specific viscoelastic stress models. Rooted mathematically by Reproducing Kernel Hilbert Space (RKHS), linear and nonlinear viscoelastic models are investigated in order to highlight the critical role of well-chosen metrics in advancing scientific research applications in complex fluid systems. Through numerical experiments on an unsteady viscoelastic lid-driven cavity flow, we demonstrate the effectiveness of energy-compatible kernels in identifying energetically-dominant coherent structures in viscoelastic flows across various Reynolds and Weissenberg numbers. Notably, features extracted using Kernel Principal Component Analysis (KPCA) with energy-compatible kernel functions yield more accurate reconstructions of mechanical energy than traditional methods, such as ordinary Principal Component Analysis (PCA) with naively defined state vectors or KPCA with ad-hoc kernel choices. Finally, we also present preliminary results considering autoencoders in complex viscoelastic fluid flows.

Cassio M. Oishi
Department of Mathematics, Statistics and Computation
UNESP - Univ Estadual Paulista
cassiooishi@gmail.com

Steven Brunton
University of Washington
sbrunton@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Samuel E. Otto
University of Washington
ottoncdr@uw.edu

MS173

Algorithms and Fine-grained Theory for Randomized Rank-revealing Factorizations

Randomized algorithms have gained increased prominence within numerical linear algebra and they play a key role in an ever-expanding range of problems driven by a breadth of scientific applications. In particular, randomized methods can be used to accelerate the computation of rank-revealing algorithms that play a key role in fast solvers and rank-structured matrices. In this talk we will discuss the randomized Golub-Klema-Stewart algorithm and provide accompanying theoretical analysis that demonstrates how the performance of randomized algorithms depends on matrix structures beyond singular values (such as coherence of singular subspaces.

Anil Damle
Cornell University
Department of Computer Science
damle@cornell.edu

MS173

Decomposition Strategies for Nonlinear Preconditioning

Nonlinear preconditioning refers to transforming a nonlinear algebraic system to a form for which Newton-type algorithms have improved success through quicker advance to the domain of quadratic convergence. We first place these methods in the context of a proliferation of variations distinguished by being left- or right-sided, multiplicative or additive, non-overlapping or overlapping, and partitioned by field, subdomain, or other criteria. We present the Nonlinear Elimination Preconditioned Inexact Newton(NEPIN), which is based on a heuristic bad/good heuristic splitting of equations and corresponding degrees of freedom. Various small nonlinearly stiff algebraic and discretized PDE problems are considered for insight and we illustrate performance advantage and scaling potential on 3D two-phase flow in porous media.

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS173

Direct Solver for a High Order Multidomain Spectral Collocation Scheme in 3D Using Gpus

We describe an efficient sparse direct solver for smooth variable-coefficient elliptic PDEs on 3D domains utilizing CPUs and GPUs. This solver is for the Hierarchical Poincar-Steklov (HPS) method, a multidomain spectral collocation scheme that enables approximation with high polynomial order and is well suited for solving oscillatory problems. In our method we construct a sparse system that solves for the boundaries of each subdomain ("leaf node") formed in the HPS discretization. After the leaf boundary values are obtained, we can solve for their interiors using numerical differential operators based on Chebyshev polynomials to get the total solution. We use batched linear algebra operations on GPUs to quickly solve for leaf interiors and construct the Dirichlet-to-Neumann (DtN) maps needed to form the sparse system. To factorize the sparse system we use multilevel direct solvers such as MUMPS. The factorized sparse system can be applied to multiple distinct problems with the same differential operator but different boundary conditions and body loads. Given a fixed leaf polynomial order, our method shows excellent linear scaling with spatial degrees of freedom in the DtN assemblies and batched leaf solves and quadratic scaling with factorization. With polynomial order up to 14 and 10 points per wavelength, it can solve oscillatory 3D problems with a relative error less than 10^{-6}.

Joseph L. Kump
University of Texas at Austin
josek97@utexas.edu

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

Anna Yesypenko
University of Texas at Austin
annayesy@utexas.edu

MS173

Non-overlapping Domain Decomposition Methods for Time Parallel Solution of PDE-constrained Optimization Problems

In this talk, we will explore non-overlapping domain decomposition methods and their application to parabolic PDE-constrained optimization problems. We will compare the difference between decomposing in space with decomposing in time. Then, we will discuss some properties of

time domain decomposition methods, such as Dirichlet-Neumann method, Neumann-Neumann method, based on the forward-backward structure of the optimality system. We will also comment on the classical Schwarz method, which fails to converge when applied to non-overlapping spatial subdomains. For each method, several variants can be identified, some of these are only good smoothers, while others could lead to efficient solvers.

Liu-Di Lu
University of Geneva
Geneva, Switzerland
liudi.lu@unige.ch

Martin J. Gander
University of Geneva
martin.gander@unige.ch

MS173

Randomized Algorithms for Compressing and Inverting Rank-Structured Matrices

This talk describes algorithms for computing data-sparse representations of rank-structured matrices, specifically H2 matrices. The algorithms are black-box, meaning that they only interact with the matrix to be compressed through its action on vectors, making them useful for tasks like forming Schur complements or matrix-matrix multiplication. A key highlight is the algorithm for "Randomized Strong Recursive Skeletonization" (RSRS), which simultaneously compresses and factorizes an H2-matrix with a strong admissibility criterion within the black-box framework. In this talk, we will showcase the application of RSRS in developing an efficient sparse direct solver for elliptic PDEs.

Anna Yesypenko
University of Texas at Austin
annayesy@utexas.edu

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

MS174

Structure-Preserving Kernel Approximation of Hamiltonian Systems Using Generalized Interpolation

We propose an intrinsically symplectic kernel approximation scheme for learning the dynamics of a Hamiltonian system from time series data. Many non-dissipative physical phenomena can be modeled as Hamiltonian systems, which result in energy-conserving models. This conservation is reflected in the flow map of a Hamiltonian system being symplectic. Instead of directly approximating the time series data (which would not preserve the underlying structure), a symplectic flow map is learned based on residuals from symplectic time integration schemes, similar to the approach described in [Horn, P., et al., *A Generalized Framework of Neural Networks for Hamiltonian Systems*, Available at SSRN 4555182023, 2023]. By doing so, the symplectic structure of the underlying system is preserved. To learn a symplectic flow map, we apply kernel methods using suitable generalized interpolation conditions in the framework of [Wendland, H., *Scattered Data Approximation*, Vol. 17, Cambridge University Press, 2004]. To achieve efficient evaluation of the flow map, we employ a kernel-based model with a small expansion size, which is attained by greedily selecting interpolation points. The time series data is then approximated by evaluating the learned symplectic flow map. We demonstrate the efficiency of our approach through numerical experiments.

Robin Herkert, Tobias Ehring, Bernard Haasdonk
University of Stuttgart
robin.herkert@ians.uni-stuttgart.de,
tobias.ehring@ians.uni-stuttgart.de,
haasdonk@mathematik.uni-stuttgart.de

MS174

Nonlinear Reduction of Partial Differential Operators Via Convolutional Architectures

Intrusive Reduced Order Methods (ROMs) based on a linear compression, such as Proper Orthogonal Decomposition (POD), are well-established tools for reducing the computational complexity of large-scale systems. However, linear techniques often face significant limitations when applied to transport-dominated problems, where such approaches may require hundreds of modes to reach an acceptable accuracy. To overcome these challenges, we propose a novel nonlinear ROM framework that employs neural networks to compress the differential operators, and we test such method in finite volume and finite element discretizations. Our approach explores various convolutional-type architectures, including continuous convolutional layers [Coscia, D., et al. "A continuous convolutional trainable filter for modelling unstructured data." Computational Mechanics, 2023, 72.2: 253-265], to effectively handle the sparsity of large operators while optimizing memory-usage. The final ROM consists of a reduced-order equation where the operators are compressed using these advanced convolutional techniques. The reduced solution is then backmapped to the original space through a nonlinear decoder, which can be implemented as either a convolutional network or a multilayer perceptron (MLP). This innovative framework provides a promising alternative to traditional ROMs, offering enhanced accuracy in complex, nonlinear scenarios.

Anna Ivagnes
SISSA
anna.ivagnes@sissa.it

Giovanni Stabile
Biorobotics Institute, Sant'Anna School of Advanced Studies
Pisa, Italy
giovanni.stabile@santannapisa.it

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

MS174

Local and Global Approximation Strategies for (Parametrized) Friedrichs' Systems

In this contribution we are concerned with parametrized linear PDE-operators exhibiting *Friedrichs' structure*, i.e. who allow for a first-order reformulation

$$A_\mu u := \sum_{i=1}^{d} A_\mu^i \partial_{x_i} u + C_\mu u, \quad A_\mu^i \in [L^\infty(\Omega)]_{sym}^{m\times m},$$

and fulfill a general positivity criterion. The associated class of PDE problems $A_\mu u = f$ includes, among others,

many well known linear problems whose solution sets behave very differently. Using the abstract Friedrichs' theory we deduce a criterion for exponential approximability which gives rigorous statements also for rarely discussed examples. In particular, we introduce a new interpretation of solution sets with solutions originating from different potentially parameter-dependent function spaces, motivated by certain graph-spaces which naturally occuring the Friedrichs' framework. In the case of multiscale problems, various spectral methods build upon local approximation spaces capturing fine-scale behavior. In order to compute these spaces locally, a localized training method has been proposed in [Buhr and Smetana,2018] which uses the compactness of a local transfer operator. We generalize this approach to Friedrichs' systems and deduce a criterion which together with a generalized Caccioppoli-inequality implies the required compactness. In both settings, numerical results for different application-driven examples will be shown.

Lukas Renelt
University of Münster
lukas.renelt@uni-muenster.de

Mario Ohlberger
University of Muenster
Applied Mathematics Muenster
mario.ohlberger@uni-muenster.de

Christian Engwer
INAM, University of Münster
christian.engwer@uni-muenster.de

MS174

Solving Elliptic PDEs with a Priori Error Bounds Using Kernel Methods and the Deep Ritz Approach

In this contribution, kernel approximations are applied as ansatz functions within the deep Ritz method to approximate weak solutions of elliptic partial differential equations with weak enforcement of boundary conditions using Nitsche's method. We prove an a priori error estimate for the approach, in which the error decays with the fill width of the centers used in the kernel approximation. Moreover, the rate of the decay is determined by the smoothness of the weak solution. We also describe how to apply the procedure in practice and give implementational details. By means of numerical examples, we investigate the performance of the proposed approach and present practical results that confirm the theoretical findings.

Tizian Wenzel
Universität Hamburg
wenzel@math.lmu.de

Hendrik Kleikamp
University of Münster
hendrik.kleikamp@uni-muenster.de

MS175

Time-Dependent Hamiltonian Simulation: Quantum Algorithm and Superconvergence

Simulation of quantum dynamics, emerging as the original motivation for quantum computers, is widely viewed as one of the most important applications of a quantum computer. The task becomes more challenging as the underlying unitary becomes more oscillatory. In such cases, an algorithm with commutator scaling and a weak dependence, such as logarithmic, on the derivatives of the Hamiltonian is desired. We introduce a new time-dependent Hamiltonian simulation algorithm based on the Magnus series expansion that exhibits both features. Importantly, when applied to unbounded Hamiltonian simulation in the interaction picture, we prove that the commutator in the second-order algorithm leads to a surprising fourth-order superconvergence, with an error preconstant independent of the number of spatial grids. The proof of superconvergence is based on operator calculus and semiclassical analysis that is of independent interest.

Di Fang
Duke University
di.fang@duke.edu

MS176

Inference of Kernel Parameters for Nonlocal Equations

Nonlocal equations are characterized by the choice of interaction kernel. Frequently, this kernel is given in terms of a priori unknown parameters such as the horizon or the fractional exponent. In this talk, we consider the problem of parameter inference for nonlocal kernel functions.

Christian Glusa
Sandia National Laboratories
caglusa@sandia.gov

MS176

Dynamic Brittle Fracture As a Well Posed Nonlocal Initial Value Problem

A nonlocal model for dynamic damage evolution consisting of two branches one elastic and the other inelastic is considered. Evolution from the elastic to the inelastic branch depends on material strength and is mediated through the constitutive law relating force to strain. The field theory is of peridynamic type and also involves a 2 point phase field that depends on the displacement. The energy for the model interpolates between elastic energy for small strains and surface energy for sufficiently large strains that fail the material. For three dimensional problems with a flat crack, power balance delivers the crack tip velocity in terms of the rate of work done by the load and the change in both the kinetic energy and elastic potential energy of the specimen. The fracture energy is the Griffith fracture energy. Subsequent passage to the limit of vanishing non-locality in a pre-cracked plate subjected to mode I loading delivers a sharp fracture evolution that recovers classic dynamic fracture mechanics. Several numerical examples are given that illustrate the method.

Robert P. Lipton
Department of Mathematics
Louisiana State University
lipton@lsu.edu

Debdeep Bhattacharya
University of Utah
d.bhattacharya@utah.edu

MS176

A Model-Based Approach for Continuous-Time Policy Evaluation with Unknown Lvy Process Dy-

namics

Reinforcement learning (RL) is active branch of machine learning focused on learning optimal policies to maximize cumulative rewards through interaction with the environment. While traditional RL research primarily deals with Markov decision processes in discrete time and space, we explore RL in a continuous-time framework, essential for high-frequency interactions such as stock trading and autonomous driving. Our research introduces a PDE-based framework for policy evaluation in continuous-time environments, where dynamics are modeled by Lvy processes. We also formulate the Hamilton-Jacobi-Bellman (HJB) equation for the corresponding stochastic optimal control problems governed by Lvy dynamics. Our approach includes two primary components: 1) Estimating parameters of Lvy processes from observed data, and 2) Evaluating policies by solving the associated integro-PDEs. In the first step, we use a fast solver for the fractional Fokker-Planck equation to accurately approximate transition probabilities. We demonstrate that combining this method with importance sampling techniques is vital for parameter recovery in heavy-tailed data distributions. In the second step, we offer a theoretical guarantee on the accuracy of policy evaluation considering modeling error. Our work establishes a foundation for continuous-time RL in environments characterized by complex, heavy-tailed dynamics.

Xiaochuan Tian
UC San Diego
xctian@ucsd.edu

Qihao Ye
University of California, San Diego
q8ye@ucsd.edu

MS176

Nonlocal Constitutive Operator: Discovering Hidden Physics from Data

Neural operators, which can act as implicit solution operators of hidden governing equations, have recently become popular tools for learning the responses of complex real-world physical systems. Nevertheless, most neural operator applications have thus far been data-driven and neglect the intrinsic preservation of fundamental physical laws and mathematical guarantees in models. In this work, we introduce a novel integral neural operator architecture called the Peridynamic Neural Operator (PNO) that learns a nonlocal constitutive law from data. This neural operator provides a forward model in the form of peridynamics, with objectivity and momentum balance laws automatically guaranteed. Additionally, we demonstrate the expressivity, efficacy, and guarantee of mathematical properties such as solution uniqueness and existence of our model. To verify the applicability of our approach, we apply the PNO in learning a material model and microstructure field from both synthetic and experimental data sets.

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

Siavash Jafarzadeh
Lehigh University
sij222@lehigh.edu

Xiaochuan Tian
UC San Diego
xctian@ucsd.edu

Zhongqiang Zhang
Worcester Polytechnic Institute
zzhang7@wpi.edu

Stewart Silling
Sandia National Laboratories
sasilli@sandia.gov

MS177

Multigrid Parallel-in-Time Methods for Nonlinear Hyperbolic PDE Systems

Sequential time-stepping using only spatial parallelism is becoming a computational bottleneck because the world's largest parallel computers now have millions of parallel processor cores due to stagnating processor speeds. In this context, parallelization in time can provide additional concurrency leading to further speedups. Parallel-in-time methods have been demonstrated to work well for parabolic PDEs, but remain a challenge for hyperbolic PDEs. In this talk, we present new developments for the multigrid reduction-in-time (MGRIT) parallel-in-time method that solve nonlinear hyperbolic PDEs and hyperbolic systems in a small number of iterations with convergence factor independent of mesh resolution. Crucial ingredients include modified semi-Lagrangian coarse-grid operators, careful linearization strategies, and characteristic block preconditioners for systems. Results are presented for linear advection and linear acoustic systems, and for the nonlinear Burgers equation, shallow water equations and compressible Euler equations with shocks.

Hans De Sterck
University of Waterloo, Canada
hdesterck@uwaterloo.ca

Oliver A. Krzysik
University of Waterloo
okrzysik@uwaterloo.ca

Robert D. Falgout
Lawrence Livermore National Laboratory
falgout2@llnl.gov

Jacob B. Schroder
Department of Mathematics and Statistics
University of New Mexico
jbschroder@unm.edu

MS177

Fast Multipole Attention for Transformer Neural Networks

Transformer-based machine learning models have achieved state-of-the-art performance in many areas. However, the quadratic complexity of the self-attention mechanism in Transformer models with respect to the input length hinders the applicability of Transformer-based models to long sequences. To address this, we present Fast Multipole Attention (FMA), a new attention mechanism that uses a divide-and-conquer strategy to reduce the time and memory complexity of attention for sequences of length n from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$ or $\mathcal{O}(n)$, while retaining a global receptive field. The hierarchical approach groups queries, keys, and values into $\mathcal{O}(\log n)$ levels of resolution, where groups at greater distances are increasingly larger in size

and the weights to compute group quantities are learned. As such, the interaction between tokens far from each other is considered in lower resolution in an efficient hierarchical manner. This multi-level divide-and-conquer strategy is inspired by fast summation methods from n-body physics and the Fast Multipole Method.

Yanming Kang
University of Waterloo
yanming.kang@uwaterloo.ca

Hans De Sterck
University of Waterloo, Canada
hdesterck@uwaterloo.ca

Giang Tran
University of Waterloo
giang.tran@uwaterloo.ca

MS177

Implicitly Extrapolated Geometric Multigrid for the Gyrokinetic Poisson Equation in Plasma Applications

We are interested in the solution of a Poisson-like equation that arises as part of a 5D coupled problem in the context of Tokamak plasma simulations. The gyrokinetic Poisson equation is solved on a large number of poloidal cross sections of the Tokamak geometry. While these cross sections can initially be treated as circular domains, it has been found that deformed geometries offer more realistic and advantageous results. Geometric multigrid methods tailored for curvilinear coordinates are however less commonly used. In this presentation, we introduce a specialized geometric multigrid algorithm using optimized line smoothers to enable parallel scalability. We use a finite difference discretization of the energy potential of the Poisson equation, which results in a symmetric matrix and allows for a matrix-free implementation with low memory usage. Additionally, we propose an implicit extrapolation technique that increases the order of convergence from linear to at least quadratic order. Finally, we will present scalability results for our OpenMP implementations of the solver.

Carola Kruse, Sai Aakash Dasari
Cerfacs
kruse@cerfacs.fr, dasari@cerfacs.fr

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS177

Near-Optimal Multigrid Preconditioners for Boundary Control of PDE-Constrained Optimization Problems with Reduced Measurements

In this talk we construct and examine multigrid preconditioners (PCMG) for numerically solving boundary control of elliptic equations when minimizing to measurements on a lower dimensional manifold. Application-driven examples include pointwise and (boundary) trace measurements. Using a control-to-state operator the reduced approach is employed, eliminating the state and adjoint variables to provide an unconstrained problem when solving for the Neumann control. We show that PCMG solves the optimization problem in a near optimal optimal manner with respect to the discretization, where normally in this reduced dimension measurement space we expect a lower order result. Importantly, this ensures PCMG increases in quality as the mesh is refined, resulting in decreasing CG iterations for solving the corresponding linear systems at higher resolutions. Numerical experiments confirm our results.

Michael A. Retzlaff
University of Maryland, Baltimore County
mretzla1@umbc.edu

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

MS178

Scientific Machine Learning in the New Era of Ai: Foundations, Visualization, and Reasoning

The rapid advancements in artificial intelligence (AI), propelled by data-centric scaling laws, have significantly transformed our understanding and generation of both vision and language. However, natural media, such as images, videos, and languages, represent only a fraction of the modalities we encounter, leaving much of the physical world underexplored. We propose that Scientific Machine Learning (SciML) offers a knowledge-driven framework that complements data-driven AI, enabling us to better understand, visualize, and interact with the diverse complexities of the physical world. In this talk, we will delve into the cutting-edge intersection of AI and SciML. First, we will discuss the automation of scientific analysis through multi-step reasoning grounded with formal languages, paving the way for more advanced control and interactions in scientific models. Second, we will demonstrate how SciML can streamline the visualization of intricate geometries, while also showing how spatial intelligence can be adapted for more robust SciML modeling. Finally, we will explore how scaling scientific data can train foundation models that integrate multiphysics knowledge, thereby enhancing traditional simulations with a deeper understanding of physical principles.

Wuyang Chen
Simon Fraser University
wuyang@sfu.ca

MS178

On Diffusion Models for Modeling Spatial-Temporal Data

TBD

N. Benjamin Erichson
International Computer Science Institute
erichson@icsi.berkeley.edu

MS178

WaveCastNet: An Ai-Enabled Wavefield Forecasting Framework for Earthquake Early Warning

Large earthquakes can be destructive and quickly wreak havoc on a landscape. To mitigate immediate threats, early warning systems have been developed to provide time to take precautions and prevent damage. The success of these systems relies on fast, accurate predictions of ground motion intensities, which is challenging due to the complex

physics of earthquakes, wave propagation, and their intricate spatial and temporal interactions. To improve early warning, we propose a novel framework, WaveCastNet, for forecasting ground motions from large earthquakes. WaveCastNet integrates a novel convolutional Long Expressive Memory (ConvLEM) model into a sequence to sequence (seq2seq) forecasting framework to model long-term dependencies and multi-scale patterns in both space and time. WaveCastNet, which shares weights across spatial and temporal dimensions, requires fewer parameters compared to more resource-intensive models like transformers and thus, in turn, reduces inference times. Importantly, WaveCastNet also generalizes better than transformer-based models to different seismic scenarios, including rare and critical situations with higher magnitude earthquakes. Importantly, our proposed approach does not require estimating earthquake magnitudes and epicenters, which are prone to errors using conventional approaches; nor does it require empirical ground motion models, which fail to capture strongly heterogeneous wave propagation effects.

Dongwei Lyu
Department of Mathematics, UC Berkeley
International Computer Science Institute
dwlyu@uchicago.edu

MS178

A Scalable Framework for Learning the Geometry-Dependent Solution Operators of Partial Differential Equations

Solving partial differential equations (PDEs) using numerical methods is a ubiquitous task in engineering and medicine. However, the computational costs can be prohibitively high when many-query evaluations of PDE solutions on multiple geometries are needed. Although artificial intelligence (AI) focused on learning PDE operators on a fixed geometry, generic and scalable frameworks that can alleviate the computational burdens on multiple geometries are yet to be developed. We aim to address the challenge by introducing a generic AI framework, DIffeomorphic Mapping Operator learNing (DIMON), which allows AI to learn geometry-dependent solution operators of different types of PDEs on various geometries. A subnetwork composed of a rotation-invariant transformer neural network is used to encode the domain shape with efficacy and efficiency. We present several examples to demonstrate the performance of the framework in learning both static and time-dependent PDEs on parameterized and non-parameterized domains; including solving a system of multiscale PDEs that characterize the electrical propagation on thousands of personalized heart digital twins. Accurate, efficient, and scalable, DIMON can reduce the computational costs of solution approximations on multiple geometries from hours to seconds with significantly less computational resources, thus ushering in fast prediction of PDE solutions with AI on multiple geometries, and advancing applying AI in engineering and medicine.

Minglang Yin
Brown University
myin16@jhu.edu

Jiwoo Noh, Ryan Brody
Johns Hopkins University
jnoh10@jh.edu, rbrody2@jhmi.edu

Nicolas Charon
University of Houston
ncharon@central.uh.edu

Lu Lu
Yale University
lu.lu@yale.edu

Mauro Maggioni
Johns Hopkins University
mauromaggionijhu@icloud.com

Natalia A. Trayanova
Johns Hopkins University
Institute for Computational Medicine
ntrayanova@jhu.edu

MS178

State Space Models for Forecasting Time Series

State-space models (SSMs) have recently emerged as a framework for learning and generating long-range sequences. They leverage linear, time-invariant (LTI) systems to achieve fast and numerically stable training and inference. Their effectiveness and stability, however, depend heavily on the initialization and parameterization of the models. In this talk, we analyze these factors and derive better model designs by studying the LTI systems from a frequency perspective. We show examples where the mathematical understanding of the LTI systems improves the performance of SSMs on time-series tasks.

Annan Yu
Cornell University
Center for Applied Mathematics
ay262@cornell.edu

MS179

Least-Squares Neural Network Method for Solving Scalar Nonlinear Hyperbolic Conservation Laws

Solutions of nonlinear scalar hyperbolic conservation laws (HCLs) are often discontinuous due to shock formation; moreover, locations of shocks are a priori unknown. This presents a great challenge for traditional numerical methods because most of them are based on continuous or discontinuous piecewise polynomials on fixed meshes. By using neural network (NN) as the class of approximating functions, recently we proposed a space-time least-squares neural network (LSNN) method. The method shows a great potential to sharply capture shock without oscillation, overshooting, or smearing. In this talk, we will give a brief introduction of NN as a class of approximating functions with moving meshes and use a simple example to show why the NN is superior to piecewise polynomials on fixed meshes when approximating discontinuous functions with unknown interface. We will then describe the LSNN method and discuss its pros and cons and related open problems.

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Junpyo Choi, Min Liu
Purdue University
choi508@purdue.edu, liu66@purdue.edu

MS179

Galerkin Neural Network Approximation of

Parameter-Dependent PDEs

We present preliminary results using Galerkin Neural Networks to approximate solutions to parameter-dependent PDEs and contrast our approach with the classical finite element reduced basis method. Our study proposes new parametric continuity results for second-order elliptic equations from which we obtain robust error estimators that inform an adaptive solution strategy. We numerically demonstrate robustness of these estimators and study convergence rates for problems with highly oscillatory parameters and other features that are challenging for traditional approaches.

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

Alexey Izmailov
Brown University
alizma@brown.edu

MS179

Efficient Kernel Surrogates for Neural Network-based Regression

Despite their promise in performing a variety of learning tasks, a theoretical understanding of the effectiveness of Deep Neural Networks (DNNs) has so far proven elusive, partly due to the difficulties inherent in studying their generalization properties on unseen datasets. Recent work has shown that randomly initialized DNNs in the infinite width limit converge to kernel machines relying on a Neural Tangent Kernel (NTK) with known closed forms. This suggests, and experiments corroborate, that empirical kernel machines can also act as surrogates for finite width DNNs. The computational cost of assembling the full NTK, however, makes this approach practically infeasible. In this talk, we will discuss the performance of the Conjugate Kernel (CK), a low-cost approximation to the NTK. For the regression problem of smooth functions and classification using logistic regression, we shall show that the CK performance is only marginally worse than that of the NTK and, in certain cases, much superior. In particular, we will present bounds for the relative test losses, verify them with numerical tests, and identify the regularity of the kernel as the key determinant of performance. In addition to providing a theoretical grounding for using CKs instead of NTKs, our framework suggests a recipe for improving DNN accuracy inexpensively. We present demonstrations of this on the foundation model GPT-2 and physics-informed operator networks.

Saad Qadeer
University of North Carolina Chapel Hill
saad.qadeer@pnnl.gov

Andrew Engel
PNNL
andrew.engel@pnnl.gov

Amanda Howard
Pacific Northwest National Laboratory
amanda.howard@pnnl.gov

Adam Tsou
Stony Brook University
adam.tsou@stonybrook.edu

Max Vargas, Panos Stinis, Tony Chiang
Pacific Northwest National Laboratory
max.vargas@pnnl.gov, panos.stinis@pnnl.gov, tony.chiang@pnnl.gov

MS179

Tensor Neural Networks for Steady Fokker-Planck Equations in High Dimensions

We solve high-dimensional steady-state Fokker-Planck equations on the whole space by using a sum of tensor products of one-dimensional feedforward networks or a linear combination of several selected radial basis functions. These networks allow us to efficiently exploit auto-differentiation in major Python packages while using radial basis functions can fully avoid auto-differentiation, which is rather expensive in high dimensions. We then use the physics-informed neural networks and stochastic gradient descent methods to learn the tensor networks. One essential step is to determine a proper numerical support for the Fokker-Planck equation. We demonstrate numerically that the tensor neural networks in physics-informed machine learning are efficient for steady-state Fokker-Planck equations from two to ten dimensions.

Taorui Wang
WPI
twang13@wpi.edu

Zheyuan Hu, Kenji kawaguchi
National University of Singapore
e0792494@u.nus.edu, kenji@nus.edu.sg

Zhongqiang Zhang
Worcester Polytechnic Institute
zzhang7@wpi.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS180

Variational Time-Implicit Schemes for Wasserstein Gradient Flow and Reaction-Diffusion Systems

We design and compute first-order implicit-in-time variational schemes with high-order spatial discretization for initial value gradient flows in generalized optimal transport metric spaces. We first review some examples of gradient flows in generalized optimal transport spaces from the Onsager principle. We then use a one-step time relaxation optimization problem for time-implicit schemes, namely generalized Jordan-Kinderlehrer-Otto schemes. Their minimizing systems satisfy implicit-in-time schemes for initial value gradient flows with first-order time accuracy. We adopt the first-order optimization scheme ALG2 (Augmented Lagrangian method) and high-order finite element methods in spatial discretization to compute the one-step optimization problem. This allows us to derive the implicit-in-time update of initial value gradient flows iteratively. The proposed method is unconditionally stable for convex cases. Numerical examples are presented to demonstrate the effectiveness of the methods in two-dimensional PDEs, including Wasserstein gradient flows, Fisher–Kolmogorov-Petrovskii-Piskunov equation, and two and four species reversible reaction-diffusion systems. This is a joint work with Stanley Osher from UCLA and Wuchen Li from U.

South Carolina.

Guosheng Fu
University of Notre Dame
gfu@nd.edu

Wuchen Li
University of South Carolina
wuchen@mailbox.sc.edu

Stanley J. Osher
UCLA
sjo@math.ucla.edu

MS180

A High-Order Bounds-Preserving Numerical Scheme for the Cahn-Hilliard-Navier-Stokes Equations with the Logarithmic Potential

A high order numerical method is developed for solving the Cahn-Hilliard-Navier-Stokes equations with the Flory-Huggins potential. The scheme is based on the Q_k finite element with mass lumping on rectangular grids, the second-order convex splitting method and the pressure correction method. The unique solvability, unconditional stability, and bound-preserving properties are rigorously established. The key for bound-preservation is the discrete L^1 estimate of the singular potential.

Daozhi Han
State University of New York at Buffalo
daozhiha@buffalo.edu

MS180

Efficient Numerical Schemes for Multi Component Mixtures of Fluids

n/a

Justin Swain
University of North Texas
justinswain2@my.unt.edu

MS180

Edge-Preserving Multilevel Methods

We present wavelet-based multilevel methods for recovering edge information in signal restoration. At each level, we solve a total variation regularized problem, which we solve by an iterative reweighting least square approach. We present numerical examples that show the effectiveness of these proposed methods.

Danyh Tolah
Arizona State Universtiy
dtolah@asu.edu

MS180

Analysis of Elliptic Transmission Problems with Sign-Changing Coefficients Using the Least Squares Finite Element Method

In the electromagnetic field, certain semiconductor materials or metals may exhibit negative dielectric constants within specific frequency ranges. This phenomenon alters the sign of the coefficients in the elliptic transmission problem. We examine elliptic transmission problems with potentially sign-changing coefficients. Traditional solution methods impose symmetrical restrictions on the grid; however, we employ the least squares finite element method to circumvent these grid restrictions and perform a priori error estimation. Numerical experiments confirm the accuracy of our analysis.

Yangwen Zhang
University of Louisiana at Lafayette
yangwen.zhang@louisiana.edu

MS181

Rough Feature Estimation and Heavy-Tailed Priors

We consider using Markovian alpha-stable and Student's t random field priors for Bayesian inversion. These are heavy-tailed prior models, which implicates that the posterior turns out to be high-dimensional, multimodal and heavy-tailed. The benefit of such prior models is that they promote Bayesian inversion with different smoothness properties, that is, you can construct simultaneously smooth and rough features. Sampling and getting estimators from the resulting posterior distributions requires care. We review some basic techniques based on MCMC, HMC, variational Bayes and optimisation literature for getting estimators with fast enough computation times. As numerical examples, we demonstrate the applicability of the proposed models and method for synthetic and X-ray tomography problems, and real-world imaging in industry.

Lassi Roininen
Lappeenranta-Lahti University of Technology
lassi.roininen@lut.fi

MS181

Sparsity Promoting Eeg Source Localization with Time and Spatial Regularization for Large-Scale Datasets

Source localization is crucial for interpreting electroencephalograms and requires solving large and extremely ill-posed inverse problems. As standard approaches are intractable for such large dynamic datasets at scale, we propose a source localization procedure utilizing sparsity promoting spatial and temporal regularization and use novel computationally efficient variable projected augmented Lagrangian methods for its solution. We compare performance of these methods to standard solvers and demonstrate the utility of our novel methods for extremely large dynamic source localization problems where standard methods cannot be used.

Jack Michael Solomon
Emory University
jack.michael.solomon@emory.edu

MS182

Accessible HPC Using Julia

Developing scientific code that makes use of modern HPC systems is crucial for addressing some of society's most pressing challenges (ranging from climate modeling, to analyzing novel therapeutics). However, writing code to make use of HPC architectures requires specialized knowledge, impeding rapid exploration, and preventing researchers not familiar with HPC from contributing their expertise. Yet, a common critique of high-level languages such as Python, R, and Matlab is that they often "get in the way" of computing experts who want low-level control over hardware.

Here we demonstrate how the Julia language can be used to develop HPC workflows without needing a low-level understanding of HPC hardware, while still allowing computing experts low-level control over the underlying hardware. Furthermore, we show how the Julia language is uniquely capable of allowing both HPC experts, and scientists to reuse each other's code – making HPC accessible to a broad community.

Johannes P. Blaschke
Lawrence Berkeley National Laboratory
jpblaschke@lbl.gov

MS182

JACC: A Julia Package for HPC Meta-Programming and Performance Portability

In this work we present Julia ACCelerated or JACC which is a package in the Julia language aimed at providing programming productivity and performance portability for users developing HPC based applications in Julia. JACC provides a unified frontend layer on top of different backends available in Julia like CPU and GPUs (CUDA, HIP and oneAPI). As a result, the user can use the same code to run on multiple backends. This improves the programming productivity along with providing performance portability. We test this near zero overhead model on multiple key kernels that are used in applications like lattice Boltzmann, the Gray-Scott model as well as the conjugate gradient method. Moreover, we will also demonstrate other capabilities in JACC like JACC.BLAS, JACC.shared and JACC.multi.

Het Y. Mankad
Carnegie Mellon University
mankadhy@ornl.gov

Het Y. Mankad, Pedro Valero Lara, William F. Godoy, Keita Teranishi, Jeffrey Vetter
Oak Ridge National Laboratory
mankadhy@ornl.gov, valerolarap@ornl.gov, godoywf@ornl.gov, teranishik@ornl.gov, vetter@ornl.gov

MS182

Automating Heterogeneous Parallelism in Julias Sciml Ecosystem

Julia's SciML is an ecosystem similar to SciPy or MATLABs built-in numerical solver libraries in that it provides the standard numerical solvers for the Julia ecosystem. Everything from ODE solvers, nonlinear solvers, optimization routines, and more are provided with one common interface. Something that makes the Julia ecosystem stand out is its direct compatibility with machine learning and its deep integration with GPUs. In this talk we will focus on the latter, showcasing how the SciML stack employs various GPU toolsets to automate the process of translating a complex CPU-based model to a GPU-based model. We will discuss the way that this differs from standard machine learning frameworks, why it achieves 20x-100x acceleration over PyTorch and Jax for ODE solvers, and some of the ongoing efforts being taken to accelerate small optimization problems which traditionally have not been able to use parallelism.

Utkarsh R
MIT
utkarsh5@mit.edu

MS182

Portable Differentiation using Enzyme

Enzyme is an automatic differentiation tool for differentiating the LLVM IR language. With Julia relying on LLVM as a compiler backend, Enzyme is a natural fit to serve as a highly optimized and flexible generator of derivative code for computational simulation kernels. We go over the current status of Enzymes Julia support through Enzyme.jl, future plans, and a demo of a simplified implementation of the Burgers equations.

Michel Schanen
Argonne National Laboratory
mschanen@anl.gov

William S. Moses
University of Illinois at Urbana-Champaign, U.S.
wsmoses@illinois.edu

Valentin Churavy
Massachusetts Institute of Technology
vchuravy@mit.edu

MS183

A Space-Time DPG Finite Element Formulation for Pulsed Laser Propagation

In this talk, we propose a modified non-linear schrödinger equation for modeling pulse propagation in optical waveguides. The proposed model bifurcates into a system of elliptic and hyperbolic equations depending on wave-guide parameters. The proposed model leads to a stable first-order system of equations. As a stable first-order system of equations, this model distinguishes itself from the canonical nonlinear schrödinger equation. We have employed the space-time discontinuous Petrov-Galerkin finite element method to discretize the modified system of equations. Here, we present stability analysis for both the elliptic and hyperbolic systems of equations. We demonstrate the stability of the proposed model using several numerical examples on space-time meshes.

Ankit Chakraborty
University of Texas at Austin
ankit.chakraborty@austin.utexas.edu

Leszek Demkowicz
Institute for Computational Engineering and Sciences (ICES)
The University of Texas
leszek@ices.utexas.edu

Judit Munoz-Matute
Basque Center for Applied Mathematics
judith.munozmatute@gmail.com

Jacob Grosek
Air Force Research Laboratory, Kirtland Air Force Base
jacob.grosek.1@us.af.mil

MS183

High-Order Biorthogonal Functions in H(div) and

H(curl)

It is well known that finite element basis functions based on orthogonal polynomials yield in better condition numbers. Not only is this possible for H^1, but also for the whole de-Rham complex, including $H(\mathrm{Div})$ and $H(\mathrm{Curl})$. In this talk, we will present a biorthogonal basis of each space in the de-Rham complex. Each basis leads to sparse projection operators, which enables us to project from high order finite element spaces to a different high order space optimally. Furthermore, we will present an example of this using a high order variant of the auxiliary space preconditioner by Hipmair and Xu (2007).

Tim Haubold
Universität Göttingen
t.haubold@math.uni-goettingen.de

Sven Beuchler
Leibniz University Hannover
beuchler@ifam.uni-hannover.de

Joachim Schöberl
Vienna University of Technology
Institute for Analysis and Scientific Computing
joachim.schoeberl@tuwien.ac.at

MS183

Performance and Scalability of Lower-Order Refined Preconditioning for Spectral/hp Element Methods for Complex, 3D Geometries

The iterative solution of higher-order spectral/*hp* element methods for simulating high Reynolds (Re) number incompressible flows around complex, 3D industrial geometries is limited due to the computational costs involved. Recently, an efficient preconditioning technique called the *Lower-Order Refined* (LOR) preconditioner is introduced within the incompressible Navier-Stokes (IncNS) equations solver of the open-source spectral/*hp* element method framework Nektar++. LOR uses a spectrally equivalent, lower-order (P=1) discretisation to precondition the high-order problem and is often combined with the algebraic multigrid (AMG) method to solve the resulting lower-order problem. The current work presents the application of the LOR preconditioner for large-scale 3D industrial problems in racecar aerodynamics. The algorithm's iterative performance and parallel scalability are analysed on CPU architectures by testing problems as large as 10^8 degrees of freedom on processor counts ranging to 10^4. The memory footprint of this preconditioner is evaluated, and the performance bottlenecks are identified by profiling the various sub-routines of the algorithm. The application regime for the AMG method is understood for solving the lower-order problem, and the best hyperparameters and coarse grid sizing are determined.

Parv Khurana
Imperial College London
McLaren Racing Limited
p.khurana22@imperial.ac.uk

Spencer Sherwin
Imperial College London
s.sherwin@imperial.ac.uk

David Moxey
Department of Engineering
King's College London
david.moxey@kcl.ac.uk

Hari Sundar
Tufts University
hari.sundar@tufts.edu

Julien Hoessler, Athanasios Chatzopoulos
McLaren Racing Limited
julien.hoessler@mclaren.com,
a.chatzopoulos@mclaren.com

MS183

A P-Adaptive Hermite Method for Electromagnetic Waves in Nonlinear Optical Media

High-order Hermite methods are especially well-suited for linear hyperbolic problems and rely on a Hermite interpolation procedure in space and a local time-stepping method to evolve the data in time. Dissipative and conservative Hermite methods of arbitrary order have been developed for Maxwells equations in linear dispersive media described by generalized Lorentz models. The dissipative Hermite method evolves the spatial derivatives of the electromagnetic fields through order m to achieve a $2m+1$ rate of convergence. Remarkably, the stability condition depends only on the maximum wave speed and is independent of the order. In this talk, we propose a dissipative Hermite method for nonlinear optics problems. We consider Kerr-type nonlinear media modeled by Maxwells equations with auxiliary differential equations describing the linear and nonlinear responses of the underlying media. The method relies on a recursion relation for the system of ordinary differential equations required to evolve the Hermite polynomial coefficients in time. The fifth-order Runge-Kutta method is used as the local time-stepping method. The approach is free of any nonlinear algebraic solver and requires solving small local linear systems of equations, where the dimension is independent of the order. Moreover, p-adaptivity is straightforward in this framework. Numerical examples in 1-D and 2-D are performed and the expected convergence order is observed for reasonable values of m.

Yann-Meing Law
Polytechnique Montreal
yann-meing.law@polymtl.ca

Zhichao Peng
The Hong Kong University of Science and Technology
pengzhic@ust.hk

Daniel Appelö
Virginia Tech
appelo@vt.edu

Thomas M. Hagstrom
Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

MS183

Subspace and auxiliary space solvers for high-order interior penalty discretizations in H(div)

Pressure-robust discretizations of the Stokes and incompressible Navier–Stokes equations can be obtained using $H(\mathrm{div})$-conforming finite element spaces for the velocity approximation, together with interior penalty discontinuous Galerkin discretizations of the Laplacian. We develop

auxiliary space and subspace correction preconditioners for these interior penalty discretizations that are robust with respect to discretization parameters, including mesh size and polynomial degree. The preconditioners reduce the problem to standard H^1-conforming problems, which can be treated using a large number of standard preconditioners, including highly scalable algebraic multigrid methods. The extension to coupled Stokes systems is performed by way of standard block preconditioners.

Will Pazner
Department of Mathematics
Portland State University
pazner@pdx.edu

MS184

A Low-Complexity Data-Driven Algorithm in Realizing True Time-Delay Beamformers

Beamforming enhances link capacity in millimeter-wave MIMO communication systems by leveraging spatial multiplexing and multiple propagation paths. A Butler matrix is a beamforming network designed to optimize phased arrays of antenna elements through FFT beams. However, the wideband FFT beams depend on frequency angles and are prone to the beam squint problem. Thus, to answer the challenges posed by the beam squint problem, we propose implementing a true-time delays-based wideband multibeam approach via the delay Vandermonde matrix (DVM). We introduce a structured neural network to realize multi-beam beamformers via the DVM structure. We learn weights within the neural network that adhere to the DVM structure. We also utilize sparse submatrices based on the DVM factorization to efficiently minimize space and computational complexities in the network. Finally, we show that our structure-imposed neural network achieves a low complexity and accurate performance in realizing multi-beam beamformers, surpassing the capabilities of conventional neural networks. This is a joint work with Hansaka Aluvihare, Arjuna Madanayake, and Xianqi Li. This work was supported by the National Science Foundation award numbers 2229473 and 2229471.

Sirani M. Perera
Embry-Riddle Aeronautical University
pereras2@erau.edu

MS184

Inner Product Free Krylov Subspace Methods for Large Scale Inverse Problems

We describe new Krylov subspace methods for solving large-scale linear inverse problems. The approach is a modification of the Hessenberg iterative algorithm that is based on an LU factorization and is therefore referred to as the least squares LU (LSLU) method. We also discuss how to incorporate Tikhonov regularization in an efficient hybrid manner. Theoretical findings and numerical results show that LSLU and the hybrid approach can be effective in solving large-scale inverse problems without requiring any inner products, and have comparable performance with existing iterative projection methods. This is a joint work with Ariana Brown, Julianne Chung, and Malena Sabate Landman.

James G. Nagy
Emory University
Department of Math and Computer Science
jnagy@emory.edu

MS184

Fast Converting Iterative Multiscale Mixed Methods

Linear solvers for subsurface flow problems based on multiscale mixed methods have been recently studied and are naturally parallelizable in multi-core computers. They can handle efficiently the solution of large problems in very heterogeneous formations of interest to the industry. Efficiency in the numerical solutions is dictated by the choice of selected interface spaces: the smaller the dimension of these spaces, the better, in the sense that fewer multiscale basis functions need to be computed and smaller interface linear systems need to be solved. Thus, in the solution of large computational problems, it is desirable to work with piecewise constant or linear polynomials. In these cases, it is well known that the flux accuracy, when computed in terms of fine grid solutions, is of the order of 10^{-1}. In this work, we focus on the development of a practical, efficient, and accurate solver for large problems. We consider subdomains with small overlapping regions, and we introduce the concept of a smoothing step, to handle small-scale errors in the multiscale solution. Additionally, we introduce novel informed spaces for calculating multiscale basis functions and develop an iterative method. Several numerical studies are presented to illustrate the very fast convergence of the new iterative solver. We consider a problem with an analytical solution followed by a study of two-dimensional solutions of several layers of the permeability field of the SPE 10 project.

Dilong Zhou
The University of Texas at Dallas
dxz200000@utdallas.edu

MS186

High Order Entropy Stable Methods for Blood Flow Simulations

The blood flow equations are a 1D nonlinear hyperbolic system which serves as a reduced model of arterial blood flow. We are interested in constructing entropy stable discontinuous Galerkin (DG) methods for this system, which enforce a semi-discrete cell entropy inequality while retaining high order accuracy. An entropy stable method for the blood flow equations was introduced by Bürger, Valbuena, and Vega [(2023). *Numer. Methods Partial Differ. Eq.*. 39, 24912509] using primitive variables. We propose an entropy stable DG method for the blood flow equations based on conservative variables, based on a new condition for entropy stability introduced in Chan, Shukla, et al. [(2024). *Journal of Computational Physics*, 112876] that eases the analysis of non-conservative terms in the 1D blood flow equations.

Raven Johnson, Jesse Chan
Rice University
rsj5@rice.edu, jesse.chan@rice.edu

MS186

Simulation-Guided Design of Congenital Aortic Valve Repair

Congenital heart defects affect approximately one in every hundred births and are the leading cause of infant mortality in the United States. Despite successes in surgical

treatment, suboptimal outcomes remain common. Surgical treatment of complex, rare congenital heart valve defects typically follows an empirical, retrospective, "guess and check' approach. Simulation-guided design tools provide a flexible, controllable and efficient means to address the clinical need to predict optimal surgical repairs. This talk will present new methods for fluid-structure interaction simulations of heart valves and their surgical treatment. In a novel, nearly first-principles method for model generation called design-based elasticity, a system of partial differential equations representing the mechanical equilibrium of the valve under pressure is derived. The solution of these equations, via tuning parameters and boundary conditions, is designed to represent the predicted loaded configuration of the valve. A full model is then constructed from the loaded configuration. In fluid-structure interaction simulations, these models are highly effective, producing realistic flows under physiological pressures over multiple cardiac cycles. I will then discuss simulation-guided design of surgical bicuspidization of the aortic valve, a repair technique for severe congenital aortic pathology, and preliminary results on in vitro validation and clinical translation to surgical practice.

Alexander D. Kaiser
Stanford University
Institute for Computational & Mathematical Engineering
alexdkaiser@stanford.edu

Perry Choi, Amit Sharir, Yoshikazu Ono, Masafumi Shibata, Yellappa Palagani
Stanford University
Cardiothoracic Surgery
pschoi@stanford.edu, asharir@stanford.edu, yoshi.ono@stanford.edu, shibatam@stanford.edu, yp275@stanford.edu

Alison Marsden
School of Medicine
Stanford University, USA
amarsden@stanford.edu

Michael Ma
Stanford University
Cardiothoracic Surgery
mma@stanford.edu

MS186

Modeling Supraventricular Tachycardia Using Dynamic Computer-Generated Left Atrium

Supraventricular Tachycardia (SVT) occurs when the heart's atria beat rapidly or irregularly compared to the ventricles. Although not immediately fatal, this disharmony contributes to strokes, heart attacks, and heart failure. Catheter ablation is the primary treatment, wherein an electrophysiologist creates a 3D heart map, guiding a catheter to burn aberrant tissue with RF energy. Despite advances, gaps persist in understanding SVT triggers and optimal ablation sites, especially in cases like atrial fibrillation (AF). To address these gaps, our team has created a model of the left atrium that beats in real time and is adjustable down to the level of individual muscles. Users can implement ablation strategies on our digital twin to quickly gain insights outside of the operating room. Patient data can be imported directly from a CT scan and electro-cardial mapping. This approach accelerates SVT comprehension without endangering lives. Our work holds life-saving potential that could revolutionize cardiac care.

Bryant Wyatt, Gavin McIntosh
Tarleton State University
wyatt@tarleton.edu, gavin.mcintosh@go.tarleton.edu

Avery Campbell
Oncor Electric Delivery
avery.grant.campbell@gmail.com

Melanie Little
MD Anderson School of Health Professions
melanie.little@go.tarleton.edu

Leah Rogers
Tarleton State University
leah.rogers@go.tarleton.edu

Brandon Wyatt
Biosense Webster
brandonwyatt@me.com

MS187

Bayesian Learning of Canonical Realizations in Linear Time-invariant Dynamical Systems

Understanding and predicting the behavior of dynamical systems given partial observations remains a significant challenge across various disciplines. We present a framework that integrates analytical approaches for dimensional reduction with a Bayesian approach for uncertainty quantification, specifically designed to exploit invariances in the representation of the dynamics and in the data. Our focus is on systems observed over time, where the data are noisy and modeled by an observation function that provides only incomplete information about the system state. This setup frequently produces sparse or redundant data, along with aspects of the dynamics that are not identifiable, complicating the inference process. To address these issues, our approach employs *reduction* techniques that simplify the representation of the data while preserving the essential, observable dynamics of the system. This is achieved by performing inference for canonical forms of the dynamics, which streamline the complexity of the system without losing critical information. By casting this inference problem in the Bayesian setting, we naturally quantify the (dependent) uncertainties of the noise model, transition operator, observation function, and other internal parameters of the dynamics that are not directly observable. We first demonstrate our approach in the setting of linear time-invariant systems, and then discuss extensions to more general classes of dynamics.

Andrey Bryutkin
Mathematics
MIT
bryutkin@mit.edu

Matthew Levine
Broad Institute
mattlevine22@gmail.com

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

MS187

A Dynamic Network Model for Thermally-Driven

Reactive Transport Near Chemical Equilibrium

Predicting the fluid, thermal, and solutal transport in an evolving complex network of pores requires a fundamental description of the transport processes and their coupling to the underlying reaction chemistry. To tackle the dynamics under various competing timescales (chemical, advective, thermal and mass diffusive) and solution-coupled boundary conditions, we perform a small-amplitude perturbation analysis on the leading-order equations derived from an existing first-principle model [Tilley et. al. 2021]. In network edges, we obtain a spectral decomposition of temperature and species transport near chemical equilibrium. By imposing flux conservation laws at network vertices via a weakly nonlinear analysis, we close the network model by describing the time evolution of temperature and species at interior vertices. This work introduces a general approach to pore network modeling with PDE dynamics near equilibrium and provides a firm analytical background for adaptation to nonlinear dynamics.

Binan Gu
Worcester Polytechnic Institute
bgu@wpi.edu

Burt S. Tilley
Worcester Polytechnic Institute
Mathematical Sciences Department
tilley@wpi.edu

Thomas Baumann
Institute of Hydrochemistry
Technische Universitaet Munich
tbaumann@tum.de

MS187

A Digital Twin for Real Time Bayesian Inference and Prediction of Tsunamis

Hessian-based algorithms for the solution to Bayesian inverse problems typically require many actions of the Hessian matrix on a vector. For problems with high-dimensional parameter fields or expensive-to-evaluate forward operators, a direct approach is often computationally intractable, especially in the context of real-time inversion. One way to overcome the computational bottleneck of Hessian matrix-vector multiplications in these large-scale inverse problems is to exploit the structure of the underlying operators. A particular class of operators for which we can exploit the structure very effectively are those representing autonomous systems. The evolution of such systems with respect to any given input may depend on the system's current state but does not explicitly depend on the independent variable (e.g., time). We present a scalable and computationally efficient approach for Bayesian inversion of problems involving autonomous systems. Our approach splits the computation into a precomputation ("offline") phase and a real-time inversion ("online") phase. Contrary to other methods, this approach does not employ a lower-fidelity approximation but instead uses the full discretization obtained from the PDE-based model. The method is applied to a real-time tsunami Bayesian inverse problem involving a time-invariant dynamical system. Scalability and efficiency of the implementation are demonstrated for state-of-the-art GPU-accelerated compute architectures.

Stefan Henneking
University of Texas at Austin
stefan@oden.utexas.edu

Sreeram R. Venkat
University of Texas Austin
201 E 24th St, Austin, TX 78712
srvenkat@utexas.edu

Milinda Fernando
Oden Institute, UT Austin
milinda@oden.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS187

A Goal-Oriented Quadratic Latent Dynamic Network Surrogate Model for Parameterized Systems

In this work, we develop a goal-oriented surrogate model for time-dependent systems with parameterized inputs using a latent dynamics network. Motivated by the success of quadratic operator inference, we employ quadratic operators for both the latent and reconstruction models. Compared to the traditional proper orthogonal decomposition-based operator inference approach, our latent dynamics model automatically discovers the hidden quadratic manifold that best approximates the evolution of the quantities of interest, as well as the underlying dynamical relationship between the parameterized inputs and outputs. We further impose a stability constraint on our model, motivated by a Lyapunov-function-based stability analysis of the underlying dynamics. Additionally, we introduce an eigenvalue penalization term in the optimization process to enhance stability, yielding a quadratic dynamical model with a generalized negative definite linear term. We demonstrate our approach on an unsteady incompressible Navier-Stokes flow problem: vortex shedding from a cylinder with a random inflow field. Numerical results confirm that our surrogate, which maps the random inflow field to several outputs including lift and drag, achieves both high approximation accuracy and stability.

Yuhang Li, Stefan Henneking
University of Texas at Austin
lyh122508@utexas.edu, stefan@oden.utexas.edu

Peng Chen
Georgia Institute of Technology
pchen402@gatech.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

MS187

Deep Learning Framework for History Matching CO2 Storage with 4D Seismic and Monitoring Well Data

Geological carbon storage entails the injection of megatonnes of CO_2 into subsurface formations. The properties of these formations are usually highly uncertain, which makes management of large-scale storage operations challenging. In this paper we introduce a history matching strategy that enables the calibration of formation properties based on early-time observations. Our framework involves two fit-for-purpose deep learning surrogate models that provide predictions for in-situ monitoring well data and interpreted 4D seismic saturation data. These two

types of data are at very different scales of resolution, so we construct separate, specialized deep learning networks for their prediction. This approach is more straightforward and efficient than training a single surrogate that provides global high-fidelity predictions. The deep learning models are integrated into a hierarchical Markov chain Monte Carlo (MCMC) history matching procedure. History matching is performed on a synthetic case with and without 4D seismic data, which allows us to quantify the impact of 4D seismic on uncertainty reduction. The use of both data types is shown to provide substantial uncertainty reduction in key geomodel parameters and to enable accurate predictions of CO_2 plume dynamics. The overall history matching framework developed in this study represents an efficient way to integrate multiple data types and to assess the impact of each on uncertainty reduction and performance predictions.

Nanzhe Wang
Stanford University
nzwang@stanford.edu

Louis J. Durlofsky
Stanford University
Dept of Energy Resources Engineering
lou@stanford.edu

MS188

The Challenges of Using Machine Learning in Weather Models

In this presentation, we will review existing and emerging machine learning practices used in atmospheric models. We will highlight the rationale behind the adoption of specific numerical strategies, examine the reasons why others have been set aside, and introduce new ideas for future research directions.

Stephane Gaudreault
Environment Canada and Climate Change
stephane,gaudreault@ec.gc.xa

MS188

Ensembling for Hierarchical Deep Learned Time-Steppers

Dynamical systems play a crucial role in scientific computing, providing valuable insights into the behavior of complex physical systems. However, many current methods rely on governing equations and exhibit numerical stiffness, especially in multiscale dynamics. One promising approach involves the use of a multiscale timestepper. This paper proposes expanding the timestepper to consider an ensemble of all possible paths. This method offers improved performance on datasets with limited data or noise. By leveraging the versatility of a multiscale timestepper and incorporating ensemble-based techniques, our approach offers a more robust and accurate framework for modeling dynamical systems in scientific computing. Additionally, it provides a means to estimate the uncertainty of predictions, enhancing the reliability of computational simulations.

Katherine Johnston
University of Washington
kjohnst@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS188

The Many Faces of Exponential Integration

In this talk we will discuss several ways in which the ideas of exponential integration can be used to construct efficient schemes for stiff systems of differential equations. We will present a new framework to develop and to analyze new class of schemes we call stiffness resilient methods. Previously proposed exponential integrators are typically derived using either classical or stiff order conditions. These order conditions are complex and difficult to solve to construct high order schemes. Classically derived methods can also suffer from the order reduction phenomenon. The new φ-order conditions we propose allow greatly simplify construction of exponential methods with favorable properties. The structure of the error of these methods is designed to prevent order reduction for many important stiff problems. At the same time stiffness resilient schemes are easy to derive using our proposed approach. In addition, we will discuss applying exponential integration to problems in fields such as plasma physics and weather prediction. We will discuss how special exponential-type methods can be constructed to take advantage of the structure of the problem to further improve efficiency of the time integration.

Mayya Tokman
University of California, Merced
School of Natural Sciences
mtokman@ucmerced.edu

MS188

Achieving Non-Linear Temporal Stability Using Projection Runge-Kutta Methods

Many physical systems yield measurable quantities that either remain constant or evolve monotonically over time. Examples of such behaviour include explicitly conserved quantities, such as mass, momentum, and energy, or implicitly conserved quantities such as entropy in the compressible Euler equations with smooth solutions. When solving these Partial Differential Equations (PDEs) numerically, an important indicator of the quality of the numerical solution is preserving physical conservation of these parameter. In the current work we will present a framework for Quasi-Orthogonal Runge-Kutta Projection Methods (QORK). These allow for conservation of auxiliary variables, such as general forms of energies or entropy, through a projection operation at the end of each time step. These projects are designed to preserve order of accuracy, efficiency, and non-linear stability. Validation will be performed for a range of ODE and PDE systems verifying analytical proofs of the above properties.

Brian Vermeire, Mohammad Najafian
Concordia University
brian.vermeire@concordia.ca, mohammadreza.najafianzadehnajafabadi@mail.concordi

MS189

Multi-Tensor AI/ML Uniquely Able to Discover and Validate Actionable and Mechanistically Interpretable Predictors from Noisy and Small-Cohort High-Dimensional Multi-Omic Clinical Data

Prediction in medicine remains limited. The entire multiome affects disease. And, unlike typical AI/ML, e.g., neu-

ral networks and deep learning, our multi-tensor AI/ML is uniquely able to discover and validate predictors from multi-omic clinical data. We demonstrated the quantum mechanics-based algorithms, i.e., the multi-tensor comparative spectral decompositions, in the discovery and validation of predictors in, e.g., brain, lung, nerve, ovarian, and uterine cancers [doi: 10.1063/1.5142559, 10.1063/1.5099268]. The algorithms identified the predictors repeatedly, in federated and imbalanced public datasets from as few as 50100 patients, showing that the algorithms are batch- and demographics-agnostic, and the predictors are actionable in the general population. The glioblastoma (GBM) brain cancer predictor, the first to encompass the whole genome, was additionally prospectively and retrospectively experimentally validated to be the most accurate and precise predictor of survival and response to treatment. All other attempts to associate a GBM tumors DNA copy numbers with the patients outcome failed, establishing that the algorithms find what all others miss, and the predictors outperform all others where they exist. Here, we introduce a unified framework for the algorithms [doi: 10.1145/3624062.3624078, 10.1073/pnas.0509033102, 10.1073/pnas.0530258100]. We also present recent experiments validating the mechanistic interpretability of the predictors.

Orly Alter
University of Utah and Prism AI Therapeutics, Inc.
orly@sci.utah.edu

Elizabeth Newman
Emory University
elizabeth.newman@emory.edu

David P. Perry, Sri Priya Ponnapalli
University of Utah
david.p.perry@sci.utah.edu, priya@sci.utah.edu

Jessica W. Tsai
University of Southern California
jtsai@chla.usc.edu

MS189

The Multilinear Generalized Singular Value Decomposition (ML-GSVD) and Its Applications in Wireless Communications

In this talk, we introduce a new Multilinear Generalized Singular Value Decomposition (ML-GSVD) for two or more matrices with one common dimension. The ML-GSVD extends the Generalized Singular Value decomposition (GSVD) of two matrices to higher orders. The proposed decomposition allows us to jointly factorize a set of matrices with one common dimension. In comparison with other approaches that extend the GSVD, the ML-GSVD preserves the essential properties of the original (matrix-based) GSVD, such as the orthogonality of the second-mode factor matrices as well as the subspace structure of the third-mode factor matrices. Moreover, we introduce an ALS-based algorithm to compute the ML-GSVD, which has been inspired by techniques to compute the PARAFAC2 decomposition. Furthermore, we present applications of the ML-GSVD in MIMO wireless communication systems. In particular, we will discuss multi-user MIMO broadcast systems with rate splitting at the transmitter. To this end, we show how the GSVD (for two users) and the ML-GSVD (for $K > 2$ users) can be used to define the number of common and private streams and how to adjust the message split. In this case, the use of the ML-GSVD simplifies the resource allocation and scheduling tasks significantly, since the common messages should be transmitted to selected groups of users where the selection of these groups depends on the current channel conditions.

Martin Haardt, Liana Khamidullina
Ilmenau University of Technology
martin.haardt@tu-ilmenau.de, liana.khamidullina@tu-ilmenau.de

MS189

Tubal Quasi Tensor Algebra In Hilbert Spaces

In the past two decades, the study matrix mimetic tensor algebra has rapidly expanded, with applications in signal processing, medical imaging, time-series analysis, and more. These methods consider a third order tensor as a matrix of tubes (real valued vectors oriented inwards), which is further equipped with a binary multiplication operation making it a ring. In this talk, I will define and consider tubal quasitensors, which are matrices whose entries are tubes that belong to an infinite dimensional vector space. Motivated by applications in signal processing and functional data analysis, where the input data are elements in some functions space, we explore the structure of tubal quasitensors over separable Hilbert spaces. Once a binary multiplication of tubes is defined, we construct an induced operator algebra, and demonstrate the construction of a C*-algebra and Hilbert C*-modules over a commutative *-ring of bounded linear operators on the Hilbert space. The main result is the construction of a matrix mimetic SVD for tubal quasitensors, analogous to the star-M SVD for finite dimensional tubal tensors. We discuss applications in multiway signal representation and time series annotation. Moreover, the construction provides an elucidating view of the compressibility of input data by means of their modulus of continuity, thus providing a new perspective on the choice of M-transform in both the finite and infinite dimensional cases.

Uria Mor
Tel Aviv University
Weizmann Institute of Science
uriamor@mail.tau.ac.il

Haim Avron
Tel Aviv University
haimav@tauex.tau.ac.il

MS189

Mapping the Progression to Malignancy in Ovarian Cancer Using Tensor Factorizations

Advances in factorization methods have allowed the systematic examination of data represented as a higher-order tensor. Studies applying these methods to biological data have shown them to capture features omitted by other methods while remaining robust to variations in data provenance and other experimental factors. Recently, these methods have been extended to model serial measurements (e.g. time series) by using factorizations with respect to the discrete cosine transform (DCT-II). We leverage the time progression modeling capability of TCAM, a novel factorization method, to explore the precancer and early cancer process in ovarian cancer proteomics datasets. Ovarian cancer remains one of the cancers with ineffective screening protocols, and the majority of the cases are still diagnosed at advanced stages of the disease. The biology of the early

stages of ovarian cancer is currently being established, and many of the mechanisms in the onset of malignancy remain unknown. Our analysis of proteomics datasets with healthy, precancer, and cancerous stages of disease in individuals found both established and novel genes implicated in the onset of malignancy. We discuss our methodology for examining single-cell datasets, the significance of our findings, and we give recommendations on using heterogeneous data sources. Finally, we explore the use of transforms other than discrete cosine and discuss experiments in transform selection.

Daniel Salinas Duron
Mount St. Marys University
d.salinasduron@msmary.edu

Malika Y. Dosso
Mount St. Mary's University
m.y.dosso@email.msmary.edu

MS190

Sample Efficient Algorithms for Stochastic Optimization

Risk-aware stochastic optimization problems arise in a wide range of applications, including engineering design, optimal control, and machine learning. The scale, computational cost, and complexity of these models often make classical optimization techniques impractical. To overcome these challenges, we have developed novel optimization methods that are both efficient and scalable, as well as well-suited for distributed computing implementations. Our methods utilize adaptive sampling strategies to progressively improve the accuracy of step computations, ensuring both computational efficiency and scalability. Furthermore, they incorporate second-order information by leveraging the inherent stochasticity of the problem. We establish global convergence rates and demonstrate that these algorithms achieve optimal worst-case iteration and sample complexity bounds. Finally, we showcase the effectiveness of our algorithms through their application to large-scale machine learning models.

Raghu Bollapragada
Operations Research and Industrial Engineering
The University of Texas at Austin
raghu.bollapragada@utexas.edu

MS190

Reduced Sample Complexity in Scenario-Based Control System Design via Constraint Scaling

The scenario approach is widely used in robust control system design and chance-constrained optimization, maintaining convexity without requiring assumptions about the probability distribution of uncertain parameters. However, the approach can demand large sample sizes, making it intractable for safety-critical applications that require very low levels of constraint violation. To address this challenge, we propose a novel yet simple constraint scaling method, inspired by large deviations theory. Under mild nonparametric conditions on the underlying probability distribution, we show that our method yields an exponential reduction in sample size requirements for bilinear constraints with low violation levels compared to the classical approach, thereby significantly improving computational tractability. Numerical experiments on robust pole assignment problems support our theoretical findings.

Jaeseok Choi
Pennsylvania State University
jxc6747@psu.edu

Anand Deo
Indian Institute of Management Bangalore
anand.deo@iimb.ac.in

Constantino Lagoa
Penn State University
Pennsylvania State University
lagoa@psu.edu

Anirudh Subramanyam
Pennsylvania State University
subramanyam@psu.edu

MS190

Rare Event Probability Estimation in Complex High-dimensional Spaces Through a Novel Gradient-free Inverse Importance Sampling Scheme

This work presents a novel gradient-free importance sampling-based framework for estimating rare event probabilities, building on our foundational Approximate Sampling Target with Post-processing Adjustment (ASTPA) approach. ASTPA uniquely constructs and directly samples an unnormalized target distribution, relaxing the optimal importance sampling distribution (ISD). The targets normalizing constant is then estimated using our inverse importance sampling (IIS) scheme, employing an ISD fitted based on the obtained samples. In this work, a gradient-free sampling method within ASTPA is developed through a guided dimension-robust preconditioned Crank-Nicolson (pCN) algorithm. To boost the sampling efficiency of pCN in our context, a computationally effective, general discovery stage for the rare event domain is devised, providing (multi-modal) rare event samples used in initializing the pCN chains. Considering the key significance of estimating normalizing constants in various fields, we demonstrate the broad applicability of our IIS scheme. We also show that an approximately fitted ISD is adequate for IIS, thus avoiding the scalability issues pertinent to density estimation methods, as showcased by examples with dimensions up to 500. The unbiasedness and coefficient of variation of the ASTPA estimator are analytically proven. Finally, diverse problems are presented, demonstrating the advantages of the suggested framework compared to several state-of-the-art sampling methods.

Elsayed Eshra, Kostas G. Papakonstantinou
Pennsylvania State University
eme5375@psu.edu, kpapakon@psu.edu

MS190

Extreme Event Prediction Using Sparse DEIM

Discrete empirical interpolation method (DEIM) estimates a function from its incomplete pointwise measurements. Unfortunately, DEIM suffers large interpolation errors when few measurements are available. Here, we introduce Sparse DEIM (S-DEIM) for accurately estimating a function even when very few measurements are available. To this end, S-DEIM leverages a kernel vector which has been neglected in previous DEIM-based methods. When the function is generated by a continuous-time dynamical

system, we propose a data assimilation algorithm which approximates the optimal kernel vector using sparse observational time series. We prove that, under certain conditions, data assimilated S-DEIM converges exponentially fast towards the true state. Finally, we demonstrate the application of S-DEIM for prediction of rogue waves from sparse observational data.

Mohammad Farazmand
NC State University
farazmand@ncsu.edu

MS191

Adaptive Mesh Refinement for Atmospheric and Oceanic Flows on GPUs

ERF and REMORA are new simulation codes for modeling atmospheric and oceanic flows, respectively. Both codes use adaptive mesh refinement (AMR) to efficiently achieve high resolution in the regions of most interest, which may change dynamically as a simulation evolves. They take advantage of the performance portability and support for AMR provided by the software framework, AMReX. In this talk, we will describe our general strategies for building new capability as well as present simulation results from both codes.

Ann S. Almgren, Aaron M. Lattanzi, Mahesh Natarajan
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov, amlattanzi@lbl.gov,
MaheshNatarajan@lbl.gov

Eliot Quon
NREL
eliot.quon@nrel.gov

Hannah Klion, Jean M. Sexton
Lawrence Berkeley National Laboratory
klion@lbl.gov, jmsexton@lbl.gov

MS191

FFTX: Speeding Up Kernels That Contain Fourier Transforms

FFTX is a performance-portable, open-source FFT software system for CPUs and GPUs analogous to FFTW for CPU systems. It is more than an FFT library, as it supports application-specific optimizations corresponding to integrating more of the algorithms into the analysis / code-generation process, based on the open-source SPIRAL tool chain for FFTs and tensor algebra algorithms developed at Carnegie-Mellon University and SpiralGen, Inc. Using SPIRAL, FFTX can represent kernels that include FFTs composed with algorithmic operations such as multiplication by a (possibly matrix-valued) symbol and batching. By combining substeps in an integrated algorithm, the amount of data traffic can be reduced by significant amounts, compared to implementations that call linear algebra and FFT library routines as black boxes that have been optimized separately. FFTX applies the size-specific analysis and automatic code generation techniques in SPIRAL to generate code customized to particular architectures. In this talk, we will demonstrate the results of applying the FFTX framework to improve the performance of kernels including free-space convolution, a pseudo-spectral Maxwell solver, and others, on supercomputers with NVIDIA and AMD GPU systems.

Peter McCorquodale
Lawrence Berkeley National Laboratory
PWMcCorquodale@lbl.gov

MS191

Multi-Level Is the New Annealing: Efficiently Learning Posterior Densities Without Gradients Via Transport Maps

Many crucial decision-making processes in Bayesian inference are described by a low-dimensional input of interest to models that are computationally expensive and functionally a black-box. We provide a method that uses nonlinear measure transport to approximate a sequence of unnormalized posteriors induced by different likelihood densities without any derivative information. By attempting to incorporate knowledge gleaned from each likelihood, we are able to incorporate likelihoods that are low-accuracy due to computational choices (e.g., "mesh sizing") as well as likelihoods that represent uncertainty differently (e.g., "tempering"). In this method, many model evaluations can be parallelized; further, we end up with an approximate posterior density of the highest fidelity model, which we can evaluate and exactly sample from without any model evaluations. We demonstrate the efficacy of this method on a few examples in PDE-based inference problems.

Daniel Sharp
Massachusetts Institute of Technology
dannys4@mit.edu

Bart Waanders
Sandia National Labs
bartv@sandia.gov

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

MS191

Solving the Real-Time Boltzmann Transport Equation with Adaptive and Multirate Time Integration Methods

Electron dynamics can be modeled by the electron real-time Boltzmann transport equation (rt-BTE) with first-principles electron-phonon (e-ph) collisions. Solving the lattice (phonon) rt-BTE with e-ph and phonon-phonon (ph-ph) collisions remains challenging due to the different timescales of e-ph and ph-ph interactions. This presentation will overview multirate time integration capabilities in the SUNDIALS library and then describe interfacing between the PERTURBO code and SUNDIALS to efficiently advance coupled electron and phonon rt-BTEs in time. We show results indicating a significant speed-up using adaptive step size and multirate infinitesimal (MRI) methods from SUNDIALS. Prepared by LLNL under Contract DE-AC52-07NA27344. LLNL-ABS-862642

Jia Yao, Ivan Maliyov
California Institute of Technology
jyyao@caltech.edu, imaliyov@caltech.edu

David J. Gardner, Carol S. Woodward
Lawrence Livermore National Laboratory
gardner48@llnl.gov, woodward6@llnl.gov

Marco Bernardi
Caltech

bmarco@caltech.edu

MS192

Energetic Variational Approaches (EnVarA) for Active Materials and Reactive Fluids

Active/reactive fluids convert and transduce energy from their surrounding into a motion and other mechanical activities. These systems are usually out of mechanical or even thermodynamic equilibrium. One can find such examples in almost all biological systems. In this talk I will present a general theory for active fluids which convert chemical energy into various type of mechanical energy. This is the extension of the classical energetic variational approaches for mechanical systems. The methods will cover a wide range of both chemical reaction kenetics and mechanical processes. This is a joint project with many collaborators, in particular, Bob Eisenberg, Yiwei Wang and Tengfei Zhang.

Chun Liu
Department of Applied Mathematics, Illinois Tech
Chicago, IL 60616
cliu124@iit.edu

MS192

Energetic Variational Neural Network Discretizations of Gradient Flows

Numerous applications in physics, material science, biology, and machine learning can be modeled as gradient flows. In this talk, we present a structure-preserving Eulerian algorithm for solving L2-gradient flows and a structure-preserving Lagrangian algorithm for solving generalized diffusions by employing neural networks as tools for spatial discretization. Unlike most existing methods that construct numerical discretizations based on the strong or weak form of the underlying PDE, the proposed schemes are constructed based on the energy-dissipation law directly. This guarantees the monotonic decay of the system's energy, which avoids unphysical states of solutions and is crucial for the long-term stability of numerical computations. To address challenges arising from nonlinear neural-network discretization, we adopt a temporal-then-spatial discretization approach on these variational systems. The proposed neural-network-based schemes are mesh-free, allowing us to solve gradient flows in high dimensions.

Yiwei Wang
UC Riverside
yiweiw@ucr.edu

Ziqing Hu
University of Notre Dame
zhu4@nd.edu

Chun Liu
Department of Applied Mathematics, Illinois Tech
Chicago, IL 60616
cliu124@iit.edu

Zhiliang Xu
University of Notre dame
zx2@nd.edu

MS192

A New Family of Thermodynamically Consistent Models for Yield-Stress Fluids.

In this presentation, we formulate a new class of rheological models for yield-stress fluids by introducing an internal dynamic variable based on a Gaussian chain assumption of the polymers and prescribing materials parameters as functions of the internal variable. The resulting models, leveraging the framework of the classical Oldroyd-B model and extending the model established by Kamani et al (2021), demonstrate energy dissipation properties. The internal variable dynamics capture the materials transient response to changes in deformation and applied loadings, through an effective relaxation time, elastic modulus, and polymeric viscosity, which are functions of the internal variable. To assess the models validity and range of applicability, we compare two representatives with the Kamani-Donley-Rogers (KDR) model in a set of material and rheometric functions, highlighting both divergences and parallels between the two types of models. Our numerical results on the material functions and rheological parameters illustrate the practical applicability and advantage of the new models over the KDR model. Specifically, the new family of models comply with the second law of thermodynamics and, in addition, can describe a broader range of transient rheological properties of yield-stress fluids.

Qi Wang
University of South Carolina
qwang@math.sc.edu

MS192

General Numerical Framework for Structure-Preserving Reduced Order Models of Thermodynamically Consistent Reversible-Irreversible PDEs

In this talk, I will present a newly developed numerical framework to derive structure-preserving reduced-order models for thermodynamically consistent PDEs. Our approach focuses on two key aspects: (a) a systematic method for generating reduced order models that respect the underlying thermodynamic principles of the original PDE systems, and (b) a strategy for constructing accurate, efficient, and structure-preserving numerical algorithms to solve these reduced order models. The frameworks generality allows it to be applied to a wide range of PDE systems governed by thermodynamic laws. We will demonstrate the effectiveness of this approach through several numerical examples.

Zengyan Zhang
Pennsylvania State University
zzz5527@psu.edu

Jia Zhao
Binghamton University
jzhao10@binghamton.edu

n/a TBA
TBA
n/a

MS193

Public Health, Mathematical Epidemiology, and Legal issues; the role of AI in interpreting the law in disease dynamics.

It is not hard to understand the necessity of having a broad yet deep understanding of public health problems in different societies. To fully grasp the reasons speeding up

a contagious disease, many factors varying from the dynamics of the disease to socio-political situations and constraints shape the dynamics of public health. In the meantime, we should pay special attention to the role of law and its various branches, particularly legal epidemiology, cutting across multiple disciplines, which affect the prospects of decision-making within public health. A retrospective analysis of different phases of pandemics lays bare the challenges arising from the interaction between law and public health decision-making. Therefore, shedding light on the dynamics and nuanced boundary between negative and positive rights, judicial opinions, and other dimensions of rules and statutes is necessary. In this research, by infusing mathematical models of epidemic control with the notion of law and its interpretation, we meticulously assess the role of the legal aspect of public health decision-making to show how unintended consequences can come about, which hamper our ability to curb diseases. Furthermore, using AI and its tools will enable us to apply legal interpretation to our model, leading to a deeper comprehension of disease and better solutions.

Rouzbeh Aghaieebeiklavasani
Faculty of Law and Political Science, University of Tehran
robert12t.robert@gmail.com

Gholam Reza Rokni Lamouki
School of Mathematics, Statistics and Computer Science
College of Science, University of Tehran
rokni@ut.ac.ir

MS193

Massive Agent-Based Simulation of Epidemic Dynamics: Assessing the Impact of Detailed Socio-Demographics and Heterogeneous Spatial-Temporal Complexity on Disease Spread

The COVID-19 pandemic highlighted the urgent need for advanced modeling approaches that can accurately capture disparities and heterogeneous impacts in disease dynamics, thereby enhancing the effectiveness of resource allocation and decision-making. In this presentation, I will discuss an agent-based model designed to simulate disease transmission across the United States, with a focus on varying spatial resolutions. This model replicates the contact patterns of over 320 million agents as they engage in daily activities at schools, workplaces, and within their communities. I will also discuss the impact of racial and ethnic characteristics and the integration of heterogeneous contact patterns in educational and occupational settings on disease spread. This model provides a robust platform for conducting "what if" scenario analysis and providing insights into potential strategies for mitigating the impacts of infectious diseases.

Prescott C. Alexander
Los Alamos National Laboratory
Los Alamos, NM
palexander@lanl.gov

Thomas harris
Los Alamos National Laboratory
University of Melbourne
thomash1@lanl.gov

Micaela Richter
Los Alamos National Laboratory,
The Ohio State University
mrichter@lanl.gov

Emily Shinkle
Los Alamos National Laboratory
eshinkle@lanl.gov

Joy Kitson
Los Alamos National Laboratory,
University of Maryland
jkitson@lanl.gov

Joe Tuccillo
Oak Ridge National Laboratory
tuccillojv@ornl.gov

Nidhi Parikh, Sara Del Valle, Timothy C. Germann
Los Alamos National Laboratory
nidhip@lanl.gov, sdelvall@lanl.gov, tcg@lanl.gov

MS193

Data-Driven Population Health Surveillance at Scale

High-quality, accurate, and real-time information about disease spread is critical for rapid response to a biothreat. We leverage the Department of Energys unique capabilities in computational science to provide state-of-the-art workflows for near-real-time population-wide health surveillance. The Silent Watcher is an integrated large language model (LLM) designed for autonomous analysis and interpretation of data for biothreat monitoring. The autonomous system can be used to monitor public health reports and identify emerging biothreats. When a threat is detected, it extracts information to enable agent-based models (ABMs) of scenarios that can be used to support public health decision making. FrESCO is a modular deep-learning natural language processing library that can be used to ingest unstructured clinical text and transform it into commonly used medical coding schemes during an emerging biothreat scenario. It also includes a prototype for differentially private federated learning allowing for collective training and inference across multiple health care institution without compromising the privacy and security of patients. ENABLE is a framework that facilitates population scale ABM simulations of disease spread, integrating information collected through the Silent Watcher and FrESCO frameworks. These interoperable workflows can be used to facilitate data-driven epidemiological modeling for rapid response to large-scale health crises.

Heidi Hanson, John Gounley
Oak Ridge National Laboratory
hansonha@ornl.gov, gounleyjp@ornl.gov

Nicholas Chia, Thomas Brettin, Rick L. Stevens
Argonne National Laboratory
chia@anl.gov, bretting@anl.gov, stevens@anl.gov

Adam Spannaus
The University of Tennessee, Knoxville
spannaus@utk.edu

Christopher Stanley
Oak Ridge National Laboratory
stanleycb@ornl.gov

Maksudul Alam
ORNL
alamm@ornl.gov

S.K. Seal, Sifat Moon
Oak Ridge National Laboratory
sealsk@ornl.gov, moons@ornl.gov

MS193

ExaEpi: An Exascale-Capable Agent-Based Model for Epidemiology

Agent-based models of diseases like COVID-19 have proven valuable in shaping the national response and preparedness. ABMs have three major advantages over other modeling techniques: ABMs can capture emergent phenomena; ABMs provide a fundamental and natural description of a system; and ABMs are quite flexible and adaptable. However, their use for forecasting and control has been limited due to difficulties in calibrating them to the multitude of data streams available during an outbreak and quantifying the uncertainties of the model. The goal of the exascale-ready ABM code ExaEpi is to tackle these challenges and expand its capabilities by leveraging the adaptive mesh refinement framework, AMReX, to simultaneously model both discrete agents and continuous fields. With this we will be able to create a generalized ABM for epidemiology and model a variety of time-evolving diseases. Coupling these efforts to novel compartmental modeling techniques, we will be able to calibrate these ABMs against a wide-variety of multi-scale data - with a full accounting of the uncertainties in the model. Given ExaEpis ability to run many ensembles quickly on exascale compute facilities, we will be able to integrate diverse data streams via continuous fields. Our end goal is to create new workflows that incorporate both reinforcement learning and surrogate models to optimally evaluate a variety of intervention scenarios and generate forecasts, with uncertainties, to guide policy decisions.

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

Peter Nugent
Lawrence Berkeley National Laboratory
penugent@LBL.gov

Arnav Singhal
Brown University
LBNL
arnav_singhal@brown.edu

Erin Acquesta
Sandia National Laboratories
eacques@sandia.gov

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA
ghosh5@llnl.gov

Nathan Keilbart
LLNL
keilbart1@llnl.gov

Tan Nguyen
LBNL
tannguyen@lbl.gov

MS193

Advances in Automation, Computation, and Algorithms for Decision Support in Public Health

The COVID-19 pandemic revealed fundamental gaps in data, analytics, infrastructure, and cross-sector collaborations needed to effectively monitor and respond to rapidly evolving health threats. Automated data ingestion and scalable epidemic analysis has the potential to narrow or eliminate these inefficiencies. Informed by our experiences in supporting public health decision makers, we will begin by presenting an open-source, event-based, automated analysis platform, AERO. AERO automates the execution of workflows on HPC and cloud-based resources and is implemented as a service that builds upon Globus and other cloud services to provide security, efficient data management, scalable workflow execution, and secure data access. The platform is decentralized, utilizing a bring your own storage and compute model to ensure scalability and adaptability to individual storage and compute needs and access. We will present AERO public health use cases where data is automatically monitored and specific events launch epidemic analyses on HPC resources, where the analysis results are made securely available to public health partners. We will also provide brief updates on calibration and data assimilation algorithms that we are developing that can make use of the always on nature of AERO for effective epidemic response. We will discuss future directions in this promising interdisciplinary area that we believe will facilitate increased and better cross-sector collaborations.

Jonathan Ozik
Argonne National Laboratory
jozik@anl.gov

Valerie Hayot-Sasson
University of Chicago
vhayot@uchicago.edu

Nicholson Collier, Abby Stevens, Arindam Fadikar
Argonne National Laboratory
ncollier@anl.gov, stevensa@anl.gov, afadikar@anl.gov

Mickael Binois
Inria
mickael.binois@inria.fr

Kyle Conroy, Olga Morozova, Kyle Chard
University of Chicago
kyleconroy@bsd.uchicago.edu,
omorozova@bsd.uchicago.edu, chard@uchicago.edu

Charles Macal
Argonne National Laboratory
macal@anl.gov

MS194

Reduced Order Electronic Wavefunctions in Density Functional Theory Calculations

The Density Functional Theory electronic structure problem can be solved very accurately using wavefunctions-based approaches, using plane-waves (pseudo-spectral) or real-space (finite difference or finite elements) discretizations. The accuracy however comes with a high computational cost associated with iteratively solving for these wavefunctions repeatedly when the parameters of the problem (atomic coordinates) change with time, for example in a first-principles molecular dynamic. In this talk, we will discuss possible ways of using reduced-order models to lower the cost of these computations while preserving their

accuracy.

Jean-Luc Fattebert
Oak Ridge National Laboratory
fattebertj@ornl.gov

Siu Wun Cheung, Youngsoo Choi, Seung Whan Chung, Jonas Kaufman, Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
cheung26@llnl.gov, choi15@llnl.gov, chung28@llnl.gov, kaufman22@llnl.gov, oseikuffuor1@llnl.gov

MS194

Learning Data-Driven Reduced-Order Models for Gyrokinetic Plasma Turbulence Simulations

The future of commercially viable fusion reactors hinges upon our ability to successfully model and predict turbulence-driven transport. The computational cost of high-fidelity gyrokinetic simulations, however, remains too severe a bottleneck for both design optimization and control tasks. Therefore, developing computationally cheap yet accurate reduced-order models (ROMs) is a promising avenue to bridge this gap. This presentation focuses on constructing data-driven, non-intrusive ROMs using limited linear simulation data generated by the gyrokinetic GENE code. Our goal is to accurately predict the frequency and growth rate of the most unstable eigenmode of the high-dimensional linear operator. To achieve this, we employ Dynamic Mode Decomposition (DMD) with time-delay embeddings. We begin by evaluating the effectiveness of this approach in predicting the system's state beyond the training horizon. Subsequently, we explore its generalizability across parametric variations, such as the binormal wavenumber and physical parameters characterizing the underlying species (e.g., density and temperature gradients). In the broader context of fusion research, this highlights the possibility of leveraging data-driven ROMs for turbulent transport to produce accurate and physically meaningful results at a fraction of the cost of high-fidelity simulations. This, in turn, can enable tasks such as the design and control of optimized fusion devices which would be infeasible otherwise.

Kevin Gill, Benjamin Faber
University of Wisconsin-Madison
kevin.gill@wisc.edu, bfaber@wisc.edu

Ionut-Gabriel Farcas
Virginia Tech
farcasi@vt.edu

Silke M. Glas
University of Twente
s.m.glas@utwente.nl

MS194

Assessing Robustness and Reliability of Machine Learning Interatomic Potentials

The use of machine learning (ML) techniques has drastically changed the field of interatomic potentials by allowing researchers to study complex systems at a significantly reduced computational footprint than brute-force quantum simulations. However, the large number of trainable parameters, the complexity of the architecture and sensitivity to the training dataset often obscure systematic biases in these ML potentials. ML potentials are assessed and ranked by calculating deviations in predicted forces and energies from ground truth values obtained from quantum simulations. But a second step of assessment happens when properties obtained from an interatomic potential are compared with quantum simulation results or experiments. This two-step verification assumes that any improvement in the predicted atomic forces and energies of configurations (in training or test sets) directly translates to improvement in property predictions. We demonstrate that traditional evaluation metrics are often not sufficient to provide a holistic picture of the correlation between model error and errors in physical properties calculated by using ML potentials and propose a new set of evaluation metrics to rank and assess interatomic potentials. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory (LLNL) under Contract DE-AC52-07NA27344 and was funded by LDRD with project tracking code 23-SI-006.

Amit Samanta
Lawrence Livermore National Laboratory
samanta1@llnl.gov

Vincenzo Lordi
Lawrence Livermore National Labs
lordi2@llnl.gov

Tianjiao Li
Physics Division, Lawrence Livermore National Laboratory
Cornell University
tl787@cornell.edu

Kaiyang Wang
Cornell University
wang147@llnl.gov

Collin Lewin
MIT
lewin1@llnl.gov

Joshua Vita
Lawrence Livermore National Laboratory
vita1@llnl.gov

MS194

Accelerating Kinetic Simulations of Electrostatic Plasmas with Local Reduced-Order Modeling

Despite the advancements in high-performance computing and modern numerical algorithms, the cost remains prohibitive for multi-query kinetic plasma simulations. In this work, we develop data-driven reduced-order models (ROM) for collisionless electrostatic plasma dynamics, based on the kinetic Vlasov-Poisson equation. Our ROM approach projects the equation onto a linear subspace defined by principal proper orthogonal decomposition (POD) modes. We introduce an efficient tensorial method to update the nonlinear term using a precomputed third-order tensor. We capture multiscale behavior with a minimal number of POD modes by decomposing the solution manifold into multiple time windows and creating temporally-local ROMs. We consider two strategies for decomposition: one based on physical time and the other on electric field energy. Applied to 1D-1V simulations, specifically the benchmark two-stream instability case, we demonstrate that the energy-windowing reduced-order mode (EW-ROM) is more efficient and accurate than the time-windowing reduced-order model (TR-ROM).

With the tensorial approach, EW-ROM solves the equation approximately 100 times faster than Eulerian simulations while maintaining a maximum relative error of 7% for the training data and 9% for the testing data.

Ping-Hsuan Tsai
Virginia Tech
pinghsuan@vt.edu

Seung Whan Chung
Lawrence Livermore National Laboratory
chung28@llnl.gov

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA
ghosh5@llnl.gov

John Loffeld, Youngsoo Choi, Jonathan Belof
Lawrence Livermore National Laboratory
loffeld1@llnl.gov, choi15@llnl.gov, belof1@llnl.gov

MS195

A High-order Domain Decomposition Preconditioner for PDEs on Moving Surfaces

We present an effective preconditioner for PDEs on moving surfaces based on a high-order accurate domain decomposition method. On static surfaces, recently developed fast direct solvers may be used to efficiently solve elliptic PDEs to high-order accuracy with multiple righthand sides, e.g. to accelerate implicit timestepping schemes. However, when the surface itself changes in time or deforms in response to the PDE solution, the direct solver must be rebuilt from scratch at every time step. In this talk, we show how a direct solver constructed for a previous surface PDE serves as an effective preconditioner for a "nearby" surface PDE, allowing the cost of constructing the direct solver to be amortized over many time steps. We demonstrate this approach on some time-dependent reaction-diffusion systems on moving surfaces.

Daniel Fortunato
Flatiron Institute
dfortunato@flatironinstitute.org

MS195

Efficient Solvers for the Sparse System That Arises from the Three-Dimensional HPS Discretization

Numerical results indicated that the Hierarchical Poincare-Steklov (HPS) discretization is effective at solving two-dimensional Helmholtz problems even in the high-frequency regime. The efficiency of the method is thanks to its nested dissection-inspired direct solver. Unfortunately, while it is straightforward to extend the discretization to three-dimensional problems, the direct solver's natural extension is inefficient. This talk presents alternative methods for solving the linear system that arises from the HPS discretization. The methods utilize both sparse direct solver and iterative solver techniques. This combination allows them to be easily parallelized. Preliminary numerical results will be presented.

Adrianna Gillman
University of Colorado at Boulder
Department of Applied Mathematics
adrianna.gillman@colorado.edu

MS195

A Sparse Hierarchical HP-finite Element Method on Disks, Annuli, and Cylinders

We present a sparse high-order hierarchical hp-finite element method tailored for disk and annulus domains. The mesh is composed of an innermost disk (excluded if the domain is an annulus) and a series of concentric annuli. The basis is formed by Zernike annular polynomials, allowing the use of recently developed quasi-optimal complexity techniques for efficient expansion and evaluation. For the 2D Helmholtz equation, the method decouples the Fourier modes, reducing the 2D problem into a sequence of independent 1D problems that can be solved in parallel with an optimal complexity factorization. Consequently, the overall 2D solution process, including the expansion of the right-hand side and the evaluation of the solution on a grid, achieves quasi-optimal complexity. We validate the effectiveness of this approach on PDEs, such as the time-dependent Schrdinger equation, in scenarios where the coefficients and data exhibit discontinuities in the radial direction. Furthermore, we demonstrate that the basis can be integrated with the Alternating Direction Implicit (ADI) method of Fortunato and Townsend, extending the methods applicability to 3D cylindrical domains. Reference: A sparse hierarchical hp-finite element method on disks and annuli, I. P. A. Papadopoulos, S. Olver, 2024 (https://arxiv.org/abs/2402.12831).

Ioannis Papadopoulos
Weierstrass Institute
papadopoulos@wias-berlin.de

Sheehan Olver
Imperial College London
s.olver@imperial.ac.uk

MS195

Multiscale Domain Decomposition Preconditioners for Heterogeneous Porous Media Flows

Multiscale domain decomposition methods have already shown great scalability in parallel high-performance simulations of heterogeneous subsurface flows [Jaramillo et al., Comput. Geosc. 26(3), 2022] at the cost of accuracy compared to the solution of the undecomposed discretizations. To take advantage of the excellent scalability properties of multiscale methods without losing accuracy, we introduce a family of preconditioners based on multiscale domain decomposition methods for solving heterogeneous subsurface flows. We build the preconditioners using the Multiscale Robin Coupled Method - MRCM [Guiraldello et al. J. Comput. Phys. 355, 2018], a domain-decomposition method based on the imposition of Robin-type boundary conditions on each subdomain. This method is combined with iterative linear solvers, such as the Preconditioned Conjugate Gradient method, to solve the linear systems arising from the (undecomposed) Finite Volume discretization of heterogeneous porous media flows. The results show that multiscale methods can be successfully used as preconditioners to improve efficiency when combined with different smoother methods, reducing the number of iterations. Results also demonstrate the parallel performance of the MRCM preconditioners, validating it as a viable alternative compared to other existing preconditioning techniques.

Fabricio Sousa

Institute of Mathematics and Computational Sciences
University of Sao Paulo
f.s.sousa@gmail.com

Pablo Carvalho
University of Sao Paulo
Brazil
pablogscarvalho@gmail.com

Franciane F. Rocha
The University of Texas at Dallas
fr.franciane@gmail.com

Rafael Guiraldello
Piri Technologies LLC
USA
rafaeltrevisanuto@gmail.com

Roberto Ausas
Instituto de Ciencias Matematicas e de Computacao
Universidade de Sao Paulo
rfausas@icmc.usp.br

Gustavo Buscaglia
Instituto de Ciencias Matematicas e Computacao
Brasil
gustavo.buscaglia@icmc.usp.br

Felipe Pereira
Department of Mathematical Sciences
University of Texas at Dallas, Dallas, Texas
felipepereira62@gmail.com

MS195

Achieving Optimal Parallel Scalability for Fast Direct Solvers for Elliptic PDEs

Elliptic PDEs are ubiquitous in science and engineering, and in many cases their solution is the bottleneck to overall performance and scalability. Nearly linear complexity and scalability can be obtained for solving elliptic PDEs using iterative methods like Multigrid. However, the convergence of iterative methods is sensitive to the conditioning of the operator. In that regard, direct solvers are more robust and attractive because of their predictable computational cost. Fast direct solvers have seen significant advancements in recent years achieving performance comparable to iterative methods when used for multiple solves. In this work, we discuss the challenges in achieving optimal parallel scalability and a study into the scale at which fast direct solvers can be competitive with iterative solvers, even when used for a small number of solves. We will also present preliminary scalability results for a distributed memory Hierarchical Poincare Steklov solver that demonstrates optimal parallel scalability.

Hari Sundar
Tufts University
hari.sundar@tufts.edu

MS196

Data-Driven Manifold Learning for Efficient Reduced-Order Modeling

Model order reduction techniques for partial differential equations are a subject of active research. A crucial component of any reduced-order model, intrusive or non-intrusive, is the employed dimensionality reduction procedure. Traditionally, proper orthogonal decomposition (POD) has been utilized in various reduced-order modeling approaches. However, POD is limited to providing linear compression, which is restrictive for problems with complex nonlinear phenomena, especially those exhibiting slow Kolmogorov N-width decay. To address this limitation, nonlinear dimensionality reduction techniques based on deep learning, such as various autoencoder architectures—fully connected, convolutional, or graph-based—have been employed, which can learn an accurate low-dimensional manifold. However, these methods are computationally intensive and require hyperparameter tuning. In this talk, we present our investigations into learning a manifold that can achieve the intrinsic dimension of the underlying solution manifold using kernel POD. Herein, we improve upon the state-of-the-art forward and backward mappings to enhance the reduction procedure, focusing on efficient and effective compression from both informational and computational standpoints. We demonstrate the applicability of our proposed framework over numerous challenging nonparametric and parametric convection-dominated dynamical systems.

Harshit Kapadia
Max Planck Institute Magdeburg
kapadia@mpi-magdeburg.mpg.de

Lihong Feng
Max Planck Institute, Magdeburg, Germany
feng@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS196

Two-Stage Model Reduction Approaches for the Efficient and Certified Solution of Parametrized Optimal Control Problems

In this talk we present an efficient reduced order model for solving parametrized linear-quadratic optimal control problems with linear time-varying state system. The fully reduced model combines reduced basis approximations of the system dynamics and of the manifold of optimal final time adjoint states to achieve a computational complexity independent of the original state space. Such a combination is particularly beneficial in the case where a deviation in a low-dimensional output is penalized. We propose different strategies for building the involved reduced order models, for instance by separate reduction of the dynamical systems and the final time adjoint states or via greedy procedures yielding a combined and fully reduced model. These algorithms are evaluated and compared for a two-dimensional heat equation problem. The numerical results show the desired accuracy of the reduced models and highlight the speedup obtained by the newly combined reduced order model in comparison to an exact computation of the optimal control or other reduction approaches.

Hendrik Kleikamp, Lukas Renelt
University of Münster
hendrik.kleikamp@uni-muenster.de, lukas.renelt@uni-muenster.de

MS196

LyapInf: Data-Driven Estimation of Stability

Guarantees for Nonlinear Dynamical Systems

Analyzing the stability of a nonlinear dynamical system is central to understanding system behavior and designing controllers, and we use a Lyapunov function for this purpose. It is possible to guarantee the stability of a system if one can find a Lyapunov function that is positive definite and decreasing over time along the orbit of the system, thus providing a sufficient condition for stability. A Lyapunov function also characterizes an estimate of the domain of attraction, which indicates the region under which the system states asymptotically converge to equilibrium. The construction of a Lyapunov function is done analytically and ad hoc for certain nonlinear systems. However, doing so for systems with high nonlinearities and different dimensions is a challenging task. To address this problem, we present a data-driven method for discovering Lyapunov functions, called Lyapunov function inference (LyapInf). This new method fits a quadratic Lyapunov function to the state trajectory data of the dynamics via optimization, where the process of inferring a Lyapunov function is based on the non-intrusive model reduction method of Operator Inference. This method learns one of many possible Lyapunov functions that ensures stability and estimates the domain of attraction without access to the system model. In this work, we demonstrate this new method on several numerical examples.

Tomoki Koike, Elizabeth Qian
Georgia Institute of Technology
tkoike3@gatech.edu, eqian@gatech.edu

MS196

Performance Bounds for Reduced Order Models

Classically, reduced order models produce a problem adapted reduced basis, by greedy methods or POD, to compute quantities of interest in real time or high query scenarios. To determine their performance, they are compared to the best possible performance of any linear method, given by the Kolmogorov n-width. While very successful in the elliptic regime, more recently problems where classical reduced order models show unsatisfactory performance came into focus, e.g. transport dominated problems with sharp gradients or shocks. Newer methods address these shortcomings by nonlinear techniques: transported subspaces, optimal transport or neural networks, to name a few. How well do these perform? Since the Kolmogorov n-width lower bounds linear methods only, it is no longer a suitable metric. We provide a new benchmark based on ideas from nonlinear width and metric entropy that provides lower performance bounds for all nonlinear method that are stable. Moreover, we compute the benchmark for a class of transport PDEs.

Gerrit Welper
University of Central Florida
gerrit.welper@ucf.edu

Donsub Rim
Washington University in St. Louis
rim@wustl.edu

MS197

ZFP: Balancing Storage and Accuracy using Compressed Arrays

The increasing cost of data movement has inspired much research into reduced-precision representations and algorithms in recent years, enabling smaller data volumes with little loss in solution accuracy. Toward that end, novel number formats like BFloat16, Posits, and several others have been proposed that attempt to make better trade-offs between accuracy and dynamic range for HPC and AI applications. There are, however, limitations to how much data can be reduced via such scalar formats. We describe ZFP, a compressed number format for multidimensional numerical arrays that dramatically increases accuracy per bit of storage over traditional number formats. ZFP partitions multidimensional arrays into small chunks that are independently compressed, decompressed, and cached. ZFP arrays provide a common array API with support for multidimensional indexing, iterators, proxy pointers, and multithreaded access, making them a drop-in replacement for conventional uncompressed arrays. ZFP arrays support fine-grained selection of storage footprint for each array, or even user-prescribed error tolerances, while retaining full random access to individual array elements. We demonstrate the superior accuracy achieved by ZFP over competing number formats in a number of HPC applications, such as PDE solvers, data analysis, and visualization, while delivering 5–100x reduction in memory footprint and data movement over double precision with acceptable accuracy.

Peter Lindstrom
Lawrence Livermore National Laboratory
pl@llnl.gov

MS197

Solving hyperbolic PDEs with mixed and variable precision

Exploiting low and mixed-precision can increase throughput in computation, i.e., reduce memory requirements, lower bandwidth requirements and increase instruction-level parallelism. On the other hand, lower precisions can affect the stability and results of an algorithm. This is particularly true for hyperbolic PDE problems, where solutions are characterised by wave propagation, such that solution features are propagated just as well as any introduced error. Looking at the high-order ADER discontinuous Galerkin method for solving hyperbolic PDEs, we present a comparative study on where and which lower numerical precision can be used in the Exascale Hyperbolic PDE Engine ExaHyPE2. We compare results and performance for benchmark problems in the context of seismic and fluid simulations, using mixed-precision kernels, variable precision over the domain, and adaptive precision using runtime criteria to evaluate the benefits and cost of each of these approaches. Finally, we examine how these can be combined to maximise benefits.

Marc Marot-Lassauzaie, Michael Bader
Technical University of Munich
marc.marot@tum.de, bader@in.tum.de

MS197

An Extension of C++ with Memory-Centric Specifications for Hpc to Support Beyond Ieee Floating Point Formats

The C++ programming language and similar languages are prone to inefficient memory usage of data structures. Compilers insert extra bits to ensure individual members align with bytes and cache lines, but they cannot optimize based on the ranges of integers, enums, or bitsets. Additionally, C++ lacks built-in support for data exchange via MPI and

for custom floating-point precisions. We propose an extension to C++ that allows developers to use attributes to guide the compiler on optimal memory layouts. For example, it allows the developer to specify if multiple booleans should be packed into a bit field, if floats should have fewer significant bits than the IEEE standard, or if a custom MPI datatype is needed for certain attributes. This extension allows but does not mandate the use of standard alignment and padding rules, does not depend on external libraries, and keeps the code compliant with standard C++. Our implementation based of Clang/LLVM demonstrates the potential to improve both runtime performance and memory efficiency, as shown through benchmarks using smoothed particle hydrodynamics (SPH). These results highlight the possible gains in performance and development productivity.

Pawel Radtke
Computer Science, Durham University
pawel.k.radtke@durham.ac.uk

MS198

Simulation of Spin Chains in Open Quantum Systems

We propose an algorithm that combines the inchworm method and the frozen Gaussian approximation to simulate the Caldeira-Leggett model in which a quantum particle is coupled with thermal harmonic baths. The frozen Gaussian approximation allows decomposition of the reduced density operator into Gaussian wave packets, so that the simulation can be done in parallel. For each wave packet, the evolution is modeled by the inchworm method, which is a bold-line approach for open quantum systems. Our numerical example of a model with the double-well potential shows that the interference pattern gradually disappears as the coupling between the system and the bath increases, demonstrating the effect of quantum decoherence.

Zhenning Cai
National University of Singapore
matcz@nus.edu.sg

Siyao Yang, Siyao Yang
University of Chicago
siyaoyang@uchicago.edu, siyaoyang@uchicago.edu

Geshuo Wang
National University of Singapore
geshuowang@u.nus.edu

MS198

On the Sign Problem in Auxiliary-Field Quantum Monte Carlo

In this talk, we will examine different types of Hubbard-Stratonovich transformation for a generic fermionic Hamiltonian. In particular, we will focus on the severity of the sign problem depending on the transformation and point out new insights we gained from numerical simulations. The talk will not assume any quantum chemistry prior knowledge and be made accessible to applied mathematicians.

Joonho Lee
Harvard University
joonholee@g.harvard.edu

MS198

Accurate Solution of the Many Body Problem Using Auxiliary Field Quantum Monte Carlo

Obtaining a solution to the many-electron Schrdinger equation stands as one of the grand challenges in chemistry and physics. Highly accurate solutions are achievable for materials containing elements from the first two rows of the periodic table, enabling us to predict the energies and properties, sometimes even achieving accuracy on par with experimental results. However, the predictive power diminishes significantly when attempting to study the properties of transition metal-containing clusters or transition states. In this presentation, I will describe a promising new method known as Auxiliary Field Quantum Monte Carlo (AFMQC), which has the potential to address many of the challenges that conventional quantum chemistry methods face. Nevertheless, this method encounters three significant challenges: (a) it is relatively expensive, (b) obtaining properties beyond energy values can be challenging, and (c) systematically improving its accuracy can also be costly. It's worth noting that some of these challenges are not unique to AFQMC but afflict many other quantum Monte Carlo methods. During this talk, I will introduce novel techniques that have allowed us to overcome these three challenges.

Sandeep Sharma
University of Colorado, Boulder
sandeep.sharma@colorado.edu

MS198

Optimal Control of Open Quantum Systems

The control of flying qubits carried by itinerant photons is ubiquitous in quantum communication networks. In addition to their logical states, the shape of flying qubits must also be tailored to match the remote receiver. In this paper, we introduce the quantum optimal control theory to the design of flying-qubit shaping protocols. A gradient-based algorithm is proposed for the generation of arbitrary-shape flying qubits with general non-ideal emitters. Simulations show that, as a joint control with the traditionally used tunable coupler, coherent driving fields can be applied to the shaping when the coupling strength is fixed or limited. The optimized control protocols can effectively suppress unwanted level leakage and multi-photon radiation. The method provides a systematic approach to high-fidelity control of flying qubits using realistic quantum devices.

Guofeng Zhang
The Hong Kong Polytechnic University
guofeng.zhang@polyu.edu.hk

MS199

Optimization Using Pathwise Algorithmic Derivatives of Electromagnetic Shower Simulations

In the context of high-energy physics, algorithmic differentiation may allow to systematically improve detector designs based on end-to-end simulations of particle detectors. However, such simulations are often stochastic programs, which pose additional mathematical challenges to AD beyond the technical complexities; e.g., the algorithmic derivative of a mean approximating an expected value may be completely different from the derivative of the ex-

pected value. As a first step to understand these challenges, we have differentiated a Monte-Carlo simulation of the energy deposition of electromagnetic showers in a simple sampling calorimeter. In this setup, we found that AD provides nearly unbiased and sufficiently low-variance derivatives once one particular physics process, multiple scattering, is turned off. In this talk, we will summarize our methodology and results, and outline further steps on the way to differentiable particle simulations.

Max Aehle
University of Kaiserslautern-Landau
max.aehle@scicomp.uni-kl.de

Nicolas R. Gauger, Max Sagebaum
TU Kaiserslautern
nicolas.gauger@scicomp.uni-kl.de,
max.sagebaum@scicomp.uni-kl.de

MS199

Online Optimisation of Machine Learning Collision Models to Accelerate Direct Molecular Simulation of Rarefied Gas Flows

We develop an online optimisation algorithm for in situ calibration of collision models in simulations of rarefied gas flows. The online optimised collision models are able to achieve similar accuracy to Direct Molecular Simulation (DMS) at significantly reduced computational cost for 1D normal shocks in argon across a wide range of temperatures and Mach numbers. DMS is a high fidelity method of simulating rarefied gases which numerically integrates the trajectories of colliding molecules. However, DMS is substantially more computationally expensive than the popular Direct Simulation Monte Carlo (DSMC) method, which uses simple phenomenological models of the collisions. We aim to accelerate DMS by replacing the computationally costly Classical Trajectory Calculations (CTC) with a neural network collision model. A key feature of our approach is that the machine learning (ML) collision model is optimised online during the simulation on a small dataset of CTC trajectories generated in situ during simulations.

Nick Daultry Ball
University of Oxford
Nicholas.DaultryBall@maths.ox.ac.uk

Jonathan F. MacArt
Princeton University
jmacart@nd.edu

Justin Sirignano
University of Oxford
justin.sirignano@maths.ox.ac.uk

MS199

A Taxonomy of Automatic Differentiation Pitfalls

Automatic differentiation is a popular technique for computing derivatives of computer programs. While automatic differentiation has been successfully used in countless engineering, science and machine learning applications, it can sometimes nevertheless produce surprising results. In this paper we categorize problematic usages of automatic differentiation, and illustrate each category with examples such as chaos, time-averages, discretizations, fixed-point loops, lookup tables, linear solvers, and probabilistic programs, in the hope that readers may more easily avoid or detect such pitfalls. We also review debugging techniques and their effectiveness in these situations.

Jan Hückelheim
Argonne National Laboratory
jhueckelheim@anl.gov

Harshitha Menon
Lawrence Livermore National Laboratory
gopalakrishn1@llnl.gov

William Moses
MIT CSAIL
wmoses@mit.edu

Bruce Christianson
University of Hertfordshire
b.christianson@herts.ac.uk

Paul Hovland
Argonne National Laboratory
hovland@mcs.anl.gov

Laurent Hascoet
Inria
hascoet@inria.fr

MS199

Adjoint Kinetic Simulation of Neutral Particles with Reversible Random Number Generators

PDE-constrained optimization problems appear in various areas of engineering and computational science. In such problems, one aims to find a design that minimizes some cost functional depending on the solution of a PDE model. Solving such problems requires both the ability to evaluate the model, as well as to compute the gradient of the objective function to the design parameters. We consider the case of fusion reactor design, where we use a Monte Carlo particle simulation to simulate neutral particles in the reactor edge. These simulations produce outputs and, consequently, a cost functional subject to stochastic noise. Care must then be taken when computing the gradient of the cost function to the design parameters, to ensure convergence. The discrete-adjoint method computes a consistent gradient through a simulation that retraces the original particle trajectories backwards in time. For complex simulations, however, storing these trajectories requires a large amount of memory. In this work, we tackle this memory issue by recomputing reversed trajectories from scratch. To this end, we use a reversible random number generator, i.e., one that can step both forwards and backwards through its sequence at identical cost. We demonstrate that this approach significantly reduces the memory required in a straightforward implementation of a simplified 1D coupled plasma-neutral model.

Emil Løvbak
Karlsruhe Institute of Technology
emil.loevbak@kit.edu

Giovanni Samaey
KU Leuven
giovanni.samaey@kuleuven.be

MS200

Incremental Hierarchical Tucker Decomposition

Efficient compression of data streams is vital in data-

driven applications, where large-scale data must be processed and stored under strict resource constraints. Traditional one-shot compression algorithms often require substantial memory, making them impractical for many real-world scenarios. Incremental methods, like those based on backpropagation, demand extensive datasets and multiple training passes, leading to prolonged training times. We introduce the Hierarchical Incremental Tucker (HIT) algorithm, which incrementally updates a low-rank tensor decomposition in the hierarchical Tucker format. HIT is the first algorithm capable of incrementally approximating tensor streams in this format while maintaining provable relative approximation error guarantees, even for infinite data streams. Leveraging the hierarchical Tucker format, HIT achieves significant computational efficiency and minimal memory usage, making it ideal for resource-limited hardware. We will present a comprehensive evaluation of HIT against state-of-the-art algorithms, including TT-ICE, using tensor streams from PDE-driven simulations and cyber-physical systems like video frames. Our results show reduced computation time and memory requirements, with improved generalization to unseen data, underscoring HIT's practical advantages in various applications.

Doruk Aksoy
University of Michigan
Department of Aerospace Engineering
doruk@umich.edu

MS200

Efficient Incremental Tucker Decomposition for Streaming Scientific Data

The Tucker decomposition has emerged as a popular format for compressing large datasets generated from high-fidelity scientific simulations. Several software packages (Tensor Toolbox, TuckerMPI) enable computing the Tucker decomposition of static data, but relatively fewer works address compressing a streaming tensor. In this work, we develop a streaming Tucker algorithm, tailored to scientific simulations where the data tensor grows along a single time-like dimension. At any point, we seek to update the existing factorization with a new tensor slice, without accessing the already incorporated tensor slices in their raw, uncompressed form. Throughout this process, we ensure that a user-specified relative error tolerance is met. We present an implementation within the TuckerMPI framework and apply it to both synthetic and simulated (combustion) datasets. By comparing against the standard (batch) algorithm, we show that our proposed approach provides significant improvements in terms of memory usage. If the Tucker ranks stop growing along the streaming tensor mode, our approach also incurs less wall-time compared to the batch version. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energys National Nuclear Security Administration under contract DE-NA0003525.

Saibal De, Hemanth Kolla
Sandia National Laboratories
sde@sandia.gov, hnkolla@sandia.gov

Eric Phipps
Sandia National Laboratories
Center for Computing Research
etphipp@sandia.gov

Zitong Li
University of California Irvine
zitongl5@uci.edu

Antoine Meyer, Francesco Rizzi
NexGen Analytics
antoine.meyer@ng-analytics.com, francesco.rizzi@ng-analytics.com

MS200

Incorporating Quantities of Interests into Low-Rank Tensor Decompositions

High-dimensional data is ubiquitous across scientific computing disciplines, including plasma physics, fluids, earth systems, and mechanics. Such data is naturally represented as a tensor, consisting of the value of each simulation variable at each point in space and time. Tensor decomposition methods, akin to matrix factorization methods for two-dimensional data, facilitate powerful analysis/reduction of such data, including data compression, surrogate modeling, pattern identification, and anomaly detection. However, existing tensor decomposition methods target simple statistical error metrics, most commonly least-squares loss, resulting in low-rank models of the data that fail to faithfully represent important physics quantities of interest or invariants arising from conservation principles. In this work, we explore new formulations of two common tensor decomposition methods, the Canonical Polyadic (CP) and Tucker decompositions, that attempt to better preserve these quantities by incorporating them directly in the optimization problems that define the resulting low-rank models. We then explore solving these extended optimization problems and investigate their ability to preserve these quantities compared to their overall reconstruction accuracy. Computational results of applying this approach to CP and Tucker decomposition of data arising from simulation of plasma physics and combustion will be presented.

Daniel M. Dunlavy
Sandia National Laboratories
dmdunla@sandia.gov

Eric Phipps
Sandia National Laboratories
Center for Computing Research
etphipp@sandia.gov

Hemanth Kolla
Sandia National Laboratories
hnkolla@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Edward Phillips
Sandia Labs
egphill86@gmail.com

MS201

Uncertainty-Aware Generative Molecular Design

Deep generative models have seen increased application in life and material sciences, particularly for small molecular drug design. However, training a generative model from scratch can be challenging for new tasks with a low amount

of available data for the specific molecular properties. On the other hand, the black-box nature of the task of interest e.g. experimental or in-silico measurement of the target property of molecules makes the task-specific fine-tuning of a pre-trained generative model computationally challenging due to its large number of parameters. This talk will focus on a model uncertainty-guided fine-tuning strategy for a pre-trained variational autoencoder (VAE) based generative model. Specifically, we leverage a low-dimensional active subspace around the pre-trained model to quantify the uncertainty of the high-dimensional VAE parameters which is computationally intractable otherwise. The resulting model uncertainty class is explored by black-box optimization to search for a diverse set of models that can generate molecules with improved properties compared to the pre-trained model.

A N M Nafiz Abeer
Texas A & M University
nafiz.abeer@tamu.edu

Sanket Jantre, Nathan Urban, Byung-Jun Yoon
Brookhaven National Laboratory
sjantre@bnl.gov, nurban@bnl.gov, byoon@bnl.gov

MS201

Investigating and Mitigating Memorization in Diffusion Based Generative Modeling of Turbulence Systems

Diffusion-based generative models have demonstrated great potential in modeling complex stochastic and chaotic physical phenomena such as turbulence. However, their tendency to memorize training data raises concerns about their generalization and capability to produce novel and diverse instantaneous flow realizations. This study systematically investigates the extent of memorization in diffusion models applied to turbulent flow synthesis, examining its impact on the quality and novelty of generated turbulence data. We employ correlation analysis and quality assessments to quantify memorization effects, exploring how factors such as training data diversity, sample size, and model capacity contribute to this phenomenon. The study further extends to conditional generative models, analyzing how memorization manifests within parametric datasets. To address these challenges, we propose a set of mitigation strategies designed to reduce memorization while maintaining the fidelity of the generated flow fields. Our findings offer insights into the trade-offs involved in using diffusion models for turbulence synthesis and present actionable approaches to enhance their robustness and reliability.

Meet Parikh, Xin-Yang Liu, Xiantao Fan, Jian-Xun Wang
University of Notre Dame
mparikh2@nd.edu, xliu28@nd.edu, fanxiantao1@gmail.com, jwang33@nd.edu

MS201

Normalizing Flows for Bayesian Inference: A Case Study with Linear Regression

Variational Inference (VI) is a scalable means of sampling from intractable posterior distributions arising in Bayesian inference. Recently, Normalizing Flows (NFs) have been used to aid VI for sampling from complex, multimodal posteriors; something out of reach for mean-field and structured VI approaches. Despite its potential, there are few theoretical studies on the approximate posterior obtained from Normalizing Flows aided VI (FAVI). The computational cost of FAVI depends heavily on the NF family, but there is no work quantifying the nature of the FAVI approximate posterior at a particular complexity of the NF. We provide a framework to answer these questions within the Bayesian linear regression context. We derive the variational posterior from FAVI with Inverse Auto-Regressive Flows that approximates the true posterior for linear regression with $p = 2$ predictors and a Gaussian prior. Since Bayesian inference facilitates uncertainty quantification, we also derive the loss in credible interval coverage resulting from using FAVI to approximate the true posterior, as a function of correlation ρ between the predictors. We find that given sufficient complexity of the NF there is no loss in coverage from FAVI regardless of ρ. In contrast, the coverage for mean-field VI increases in $|\rho|$. We then extend our results to $p > 2$. Our results are presented across complexity of the NFs and include the complexity required for FAVI to exactly recover the true posterior.

Sumegha Premchandar
Wells Fargo
premchan@msu.edu

Shrijita Bhattacharya, Tapabrata Maiti
Michigan State University
bhatta61@msu.edu, maiti@msu.edu

MS201

Latent Space Modeling for AI-Enabled Physics Research

Latent spaces, which provide compressed representations of high-dimensional data, play a pivotal role in data-driven scientific research techniques. In this talk, we provide an overview of representation learning in the context of generative modeling for astrophysical studies – from traditional emulation techniques to state-of-the-art foundation models. We highlight efficient anomaly identification and feature extractions, smoother data interpolations, faster likelihood estimations for Bayesian inference, and other research methodologies enhanced by robust latent space modeling. A key focus of the talk is on addressing the challenges posed by unsupervised learning in scientific contexts, where traditional latent space modeling can lead to entangled representations that obscure physical parameters. We introduce a novel latent space modeling approach that leverages auxiliary information to achieve disentanglement, aligning latent factors with known physical variables, thereby improving model interpretability and robustness against adversarial attacks.

Nesar Ramachandra
Argonne National Laboratory
nramachandra@anl.gov

Arkaprabha Ganguli
Argonne National Laboratory
Mathematics and Computer Science Division
aganguli@anl.gov

Julie Bessac
US National Renewable Energy Laboratory
julie.bessac@nrel.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division

emconsta@mcs.anl.gov

MS202

Robust OED for Large-Scale Nonlinear Bayesian Inverse Problems

This work presents a scalable approach for the robust optimal experimental design of nonlinear inverse problems. Typically, an optimal design is one which maximizes some utility quantifying the quality of the solution of an inverse problem. This optimal design, however, is dependent on elements of the inverse problem, such as the simulation model, the prior, or the measurement error model. Robust optimal experimental design aims to produce an optimal design that is aware of the additional uncertainties encoded in the inverse problem and remains optimal even after variations in them. In this work, we follow a worst-case scenario approach to develop a new framework for the robust optimal design of nonlinear Bayesian inverse problems. The proposed framework a) is designed for infinite-dimensional Bayesian inverse problems constrained by nonlinear partial differential equations; b) develops efficient approximations of the utility, namely, the expected information gain; c) employs eigenvalue sensitivity techniques to develop analytical forms and efficient evaluation methods of the gradient of the utility with respect to the uncertainties we wish to be robust against; d) employs a probabilistic optimization paradigm that properly defines and efficiently solves the resulting combinatorial max-min optimization problem. The proposed approach is numerically validated for an optimal sensor placement problem in a classical nonlinear inverse problem.

Abhijit Chowdhary
North Carolina State University
achowdh2@ncsu.edu

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

MS202

Optimal Experimental Design for High-Energy-Density Physics

High-energy-density (HED) physics is a field of research that studies the behavior of materials under extreme conditions of temperature and pressure. It plays an important role across a range of scientific disciplines, including condensed matter physics, astrophysics, plasma physics, and fusion energy research. At these extreme conditions, much of the material behavior is not well understood. Experiments are therefore crucial for developing and refining HED physics and computational models for simulation. HED experiments are extremely expensive and often performed at oversubscribed facilities. A careful design of experiments is thus critical for maximizing the value produced by each experiment and the data it generates. We achieve this by employing methods of optimal experimental design (OED). Using a Bayesian formulation, OED quantifies the expected information gain (or expected uncertainty reduction from prior to posterior) of the unknown model parameters provided by an experiment, and maximizes it across the design space. Numerically, we estimate the expected information gain with a nested Monte Carlo. We apply this OED method on a series of experiments to perform a material characterization under shocked conditions, present validation to the identified designs, and discuss insights revealed by the results.

Codie Fiedler Kawaguchi, Xun Huan, Eric Johnsen
University of Michigan
cfkawa@umich.edu, xhuan@umich.edu,
ejohnsen@umich.edu

Alex Rasmus, Elizabeth C. Merritt
Los Alamos National Laboratory
rasmus@lanl.gov, emerritt@lanl.gov

MS202

Goal-Oriented Optimal Design of Infinite-Dimensional Bayesian Inverse Problems

Our work focuses on goal-oriented optimal experimental design (gOED) within the framework of large-scale linear Bayesian inverse problems. Specifically, we address a nonlinear goal-functional defined on the inversion parameter, which is considered in an infinite-dimensional space. Traditional optimal experimental design aims to identify data acquisition locations (designs) that minimize uncertainty in the inversion parameter or maximize information gain. In contrast, gOED seeks designs that, when employed in the inversion process, reduce uncertainty in the goal itself. Our main contribution is the development of a gOED criteriona functional of sensor-weightingthat, when minimized, yields such designs. This criterion is derived via a second-order Taylor expansion of the goal functional, followed by an analytical computation of the expected variance of this expression. We propose three computational methods for evaluating the gOED criterion, each suitable for efficient application to large-scale problems. Our numerical experiments, conducted for both quadratic and nonlinear goal-functionals, demonstrate that our methods provide a more accurate characterization of the goal compared to the classical A-optimal approach and C-optimality (obtained via linearization of the goal-functional).

John N. Neuberger
North Carolina State University
jnneuber@ncsu.edu

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Bart Waanders
Sandia National Labs
bartv@sandia.gov

MS202

A Linearise-Then-Optimise Approach to Large-Scale Bayesian Optimal Experimental Design

We consider optimal experimental design (OED) within the context of large-scale non-linear Bayesian inverse problems. We propose a novel linearise-then-optimise (LtO) approach to avoid repeatedly computing local (Laplace) approximations to the posterior at the 'training' stage. The proposed LtO approach can also be applied to the problem of marginalised OED. To account for the significant modelling errors and uncertainties induced by the approach we apply the Bayesian approximation error (BAE) approach which is carried out off-line in a similar fashion to training

the experimental design. We demonstrate the effectiveness of the approach on several PDE-based problems motivated by several real-world optimal experimental design situations. This is based on ongoing work with K. Koval and A. Alexanderian.

Ruanui Nicholson
Department of Engineering Science
University of Auckland
ruanui.nicholson@auckland.ac.nz

MS202

Leveraging Machine Learning for Active Experimentation and Data Generation

A typical learning process of a neural network can be described by the interaction between the data, the weights and the model architecture. In a standard regime, this interaction involves updating the weights of the neural network model with respect to a series of data batches for a fixed architecture. As this data is generated a priori, there is an inherent uncertainty in the learning process due to the discrepancy between the data required to learn efficiently and the data available. In this talk, we will demonstrate that this uncertainty could be quantified and further leveraged to improve performance in a neural network training process through active experimentation. Towards this end, we will enunciate a novel mathematical formulation where the learning problem is formulated as a dynamical system and the dynamics are represented by a differential equation. Subsequently, we will elucidate, "how to quantify this uncertainty during learning?" and provide insights into the effect of the this uncertainty. We will end this talk with insights on "how to use uncertainty to correct the behavior of the neural network model while learning" In particular, we will show that way to experiment and generate data on the fly.

Krishnan Raghavan
Mathematics and Computer Science
Argonne National Laboratory
kraghavan@anl.gov

George Papadimitriou
Information Sciences Institute
georgpap@isi.edu

Hongwei Jin
Argonne National Laboratory
jinh@anl.gov

MS203

A Computational Framework for the Inverse Design of Anisotropy in Microstructures

Multiscale materials generally exhibit an anisotropic mechanical response which makes the inverse design a twofold problem. First, we must learn the type and orientation of anisotropy and then find the optimal design parameters to achieve a desired mechanical response. Given the stress-strain data for any material, we first solve a forward problem using partially Input Convex Neural Networks (pICNNs) to get a polyconvex representation of the strain energy with respect to the deformation gradient whereas it can have an arbitrary form with respect to the design parameters. While training the neural network, we also find the type of anisotropy, if any, along with the preferred directions. Once the model is trained, we solve the inverse problem using an evolution strategy to get the design parameters that give a certain stress-strain response. We test the framework against synthetic macro- and micro-scale data and it is able to not only learn the polyconvex potentials but also recover the correct parameters for the inverse design problem.

Asghar A. Jadoon
University of Texas at Austin
ajadoon@utexas.edu

Max Rosenkranz
TU Dresden
max.rosenkranz@tu-dresden.de

Manuel Rausch
The University of Texas at Austin
manuel.rausch@utexas.edu

Karl Kalina
TU Dresden
karl.kalina@tu-dresden.de

Reese Jones
Sandia National Laboratories
rjones@sandia.gov

Jan Fuhg
The University of Texas at Austin
jan.fuhg@utexas.edu

MS203

Latent Space Denoising Diffusion for Inverse Design of Elastoplastic Microstructures with Targeted Constitutive Models

Designing multi-phase, heterogeneous materials such as metal-ceramic composites, dual-phase steels, and polymer blends, which consist of elastoplastic phases with complex and varied property distributions, presents a high-dimensional design challenge. This work introduces a denoising diffusion algorithm to address this challenge by embedding the desired behaviors in a latent space. The algorithm utilizes denoising diffusion probabilistic models to refine microstructural designs iteratively, generating realistic samples with tailored mechanical responses. Reversing a Markov diffusion process, it efficiently manipulates the multi-phase topology and material parameters of microstructures, producing a diverse array of prototypes with targeted stiffness, yield surface responses, and hardening mechanisms. Neural network surrogates replace high-fidelity finite element simulations, enabling rapid identification of prototypes within the desired performance range. The results demonstrate the effectiveness of the denoising diffusion approach in generating microstructures with precisely tuned elastoplastic properties, all within the latent space informed by the training data. Numerical experiments validate the algorithm's capability in inverse design, providing insights into the intricate relationships between microstructural geometry, topology, and their resultant macroscopic elastoplastic behaviors.

Sara Karimi, Nick Vlassis
Rutgers, The State University of New Jersey
sk3192@scarletmail.rutgers.edu, nick.vlassis@rutgers.edu

MS203

Discovery Through Mechanical Testing and Dic in

3D-Printed Viscoelastic Polymers

Despite the surge in popularity of physics-augmented neural networks (PANN) that offer a data-driven approach to material modeling, most demonstrations of these methods rely on synthetic data. Training PANNs on experimental data requires special considerations because, unlike synthetic data that can readily cover all possible deformation and stress states, experimental data is often sparse and limited by practical constraints. Additionally, while synthetic data is typically generated from traditional models with relatively simple functional forms, experimental data present a challenge due to the potential complexity and unknown nature of the underlying physics. In this talk, we will explore training PANN constitutive models using experimental data from 3D-printed viscoelastic polymers. We focus on a multi-material printing process that adjusts properties through polymer blending, and we train PANN models to represent the composition of these blends. We will address the challenge of working with sparse data from simple experiments like tension and torsion. In addition, we will discuss using digital image correlation strain measurements (DIC) to generate more comprehensive training, testing, and validation datasets.

Steven J. Yang, Govinda Padmanabha
Cornell University
sjy32@cornell.edu, ga355@cornell.edu

Reese Jones
Sandia National Laboratories
rjones@sandia.gov

Daniel T. Seidl
Sandia National Lab
dtseidl@sandia.gov

Nikolaos Bouklas
Cornell University
nb589@cornell.edu

MS203

Characterizing Nonlinear Finite Deformation Viscoelastic Material Properties Using Digital Image Correlation and Deep Learning Deep-VM Algorithms

Accurately characterizing viscoelastic material properties at different strain rates is vital in various engineering disciplines. For instance, understanding their quasistatic behavior is crucial for precisely controlling soft robotics and flexible electronics precisely, while quantifying their ultra-high-rate behavior is essential for improving the efficiency and safety of ultrasound and laser surgeries. Traditional methods for extracting these material properties often rely on analytical or numerical models coupled with iterative comparisons to experimental data, which can be computationally expensive. Here we propose novel deep learning-based methods, called Deep-VM algorithms, that can directly extract nonlinear, finite deformation viscoelastic material properties from Digital Image Correlation (DIC) full-field experimental measurements. By assuming the Kelvin-Voigt viscoelastic material model, our approach uses DIC-measured spatiotemporal deformation data to independently learn the rate-independent and rate-dependent stress components through two parallel neural network channels. We demonstrate the effectiveness of our algorithm using both synthetic and experimental datasets of hydrogels subjected to various loading scenarios, including point-force loading and creep tests. Additionally, we explore the application of our algorithm in assessing viscoelastic properties at different rates through experiments of needle and pulsed-laser-induced cavitation in hydrogels.

Lehu Bu
The University of Texas at Austin
lhbu@utexas.edu

Congjie Wei, Chenglin Wu
Texas A&M University
congjiewei@tamu.edu, chenglinwu@tamu.edu

Jin Yang
University of Texas at Austin
jin.yang@austin.utexas.edu

MS204

Sunny.jl for Atomic-scale Magnetism

Sunny.jl is a package for simulating materials magnetism at atomic scales, where quantum effects become important. It is being developed by an interdisciplinary team that spans theoretical and experimental physics, as well as computational science. Julia has served as an excellent platform for this project. Code can be rapidly prototyped in natural math syntax, and then later tuned for very high computational efficiency. Julia's just-in-time compilation model enables an interactive development experience with no compromise to performance. This talk reports on some of the unique features that Sunny provides for simulating quantum spin dynamics, and how various Julia features have helped to streamline the development workflow.

Kipton Barros
Los Alamos National Laboratory
kbarros@lanl.gov

MS204

The ITensors.jl Library and Sparse Tensor Format

Tensor networks are a set of computational tools designed to produce efficient low-rank approximations of high dimensional data. ITensors.jl is a Julia library dedicated to tensor handling and tensor network computations. Originally motivated by quantum physics, it can also be used as a generic, algorithm-agnostic tensor library. It relies on intelligent indices to determine which axes must be contracted. This feature allows to write simple and efficient code including a large number of indices.In this talk, we first give an introduction to ITensor, its key features and review some use cases in quantum physics. In a second time, we focus on the problem of sparse storage format for a tensor. This is a challenging problem as the commonly used compressed sparse rows and compressed sparse columns matrix formats do not generalize easily to higher dimensions. We detail the block sparse tensor implementation used in ITensor and how to efficiently make use of symmetries and conserved quantities.

Olivier Gauthe, Karl Pierce
Flatiron Institute
ogauthe@flatironinstitute.org,
kpierce@flatironinstitute.org

Matthew Fishman
Flatiron Institute, New York
mfishman@flatironinstitute.org

Miles Stoudenmire
Flatiron Institute
mstoudenmire@flatironinstitute.org

Steven White
University of California Irvine
srwhite@uci.edu

MS204

Towards New Workflows in Computational Materials Science with Julia

Current solutions to key 21st-century challenges (such as climate change, food insecurity, healthcare, and communications) are fundamentally limited by the functional properties of known materials. Designing new materials increasingly relies on computational modeling with state-of-the-art workflows frequently interweaving first-principles with empirical modeling as well as data-driven approaches. To make advances, we often must establish novel connections across fields such as physics, chemistry, computer science, and applied math. The Julia language and user community are well-poised to address these challenges. However, with many well-established software solutions and their respective user communities already in existence, Julia tools cannot just duplicate existing functionality. Rather, it is crucial to integrate and add value to existing (often monolithic) codebases and software ecosystems. In this talk, I will discuss efforts in this area by the JuliaMolSim community, with a particular focus on our work in developing interfaces to facilitate interoperability and qualitatively new types of computational workflows that can enable new cutting-edge computational materials research.

Rachel Kurchin
Carnegie Mellon University
rkurchin@andrew.cmu.edu

MS204

LATTICEQCD.JL: Simulation of Quantum Chromo-Dynamics in 4 Dimensional Spacetime

We present our code (LatticeQCD.jl) for quantum chromodynamics (QCD), which describes microscopic world inside of nucleons. QCD calculation has been implemented by Fortran and C++ on supercomputers or GPU clusters because it requires huge numerical resource, i.g. Monte-Carlo with inversions of matrices with 10^{16} x 10^{16}, and has been succeeded to calculate crucial numbers used in experiments. We implemented a code for QCD in Julia, which achieves compatible speed with a Fortran code. As one of the applications in high-performance computing with Julia, we will show benchmarks with parallel computations with Julia and MPI.

Yuki Nagai
University of Tokyo
nagai.yuki@mail.u-tokyo.ac.jp

Akio Tomiya
Tokyo Woman's Christian University
akio@yukawa.kyoto-u.ac.jp

MS205

Third-Order A-Stable Alternating Implicit Runge-Kutta Schemes

We design pairs of six-stage, third-order, alternating implicit Runge–Kutta (RK) schemes that can be used to integrate in time two stiff operators by an operator-split technique. We also design for each pair a companion explicit RK scheme to be used for a third, nonstiff operator in an IMEX fashion. The main application we have in mind are (non)linear parabolic problems, where the two stiff operators represent diffusion processes (for instance, in two spatial directions) and the nonstiff operator represents (non)linear transport. We identify necessary conditions for linear A(α)-stability by considering a scalar ODE with two (complex) eigenvalues lying in some fixed cone of the half-complex plane with nonpositive real part. We show numerically that it is possible to achieve A(0)-stability when combining two operators with negative eigenvalues, irrespective of their relative magnitude. Finally, we show by numerical examples including two-dimensional nonlinear transport problems discretized in space using finite elements that the proposed schemes behave well.

Alexandre Ern
CERMICS
Inria Paris
alexandre.ern@enpc.fr

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

MS205

Formulation and Discretization of Fully Hyperbolic Approximations of Multiphysics Problems

We present a framework for constructing a first-order hyperbolic system whose solution approximates that of a desired higher-order evolution equation. Constructions of this kind have received increasing interest in recent years, and are potentially useful as either analytical or computational tools for understanding the corresponding higher-order equation. We perform a systematic analysis of a family of linear model equations and show that for each member of this family there is a stable hyperbolic approximation whose solution converges to that of the model equation in a certain limit. We then show through several examples that this approach can be applied successfully to a very wide range of nonlinear PDEs of practical interest.

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Abhijit Biswas
King Abdullah University of Science and Technology
abhijit.biswas@kaust.edu.sa

MS205

A Time Adaptive Partitioned Solver for Thermal and Mechanical Fluid-Structure Interaction

We consider partitioned solvers for mechanical and thermal ?uid-structure interaction where different sub-solvers are used for the different domains. We are in particular interested in improving the computational efficiency through the use of different and adaptive time steps in the sub solvers. This can be achieved by the use of waveform iterations, which has recently been implemented in the open source coupling library PreCICE. An alternative to relaxation are Quasi-Newton methods and these have recently

been combined with waveform iterations for fixed time steps. In this talk, we further extend the Quasi-Newton waveform iterations to the time adaptive case, where both of the sub-solvers use an adaptive time stepping scheme. This is achieved by sampling the waveforms to a fixed time grid in the QN-method. Furthermore, we also discuss the properties of this extension as well as its implementation in the open source coupling library PreCICE. Lastly, we also show that using a time adaptive solver results in faster run times for a thermal fluid structure interaction test case, were a piece of steel is cooled with air, as well as a simple mechanical fluid structure interaction test case.

Niklas Kotarsky
Centre for Mathematical Sciences at Lund University
niklas.kotarsky@math.lu.se

Philipp Birken
Lund University
philipp.birken@math.lu.se

MS205

Adaptive Multi-Rate Time Integration in the Arkode Library

Multiphysics models couple two or more physical processes together in a single simulation. These combinations may include systems of differential equations with different type (parabolic, hyperbolic, etc.), with different degrees of nonlinearity, and that evolve on disparate time scales. As a result, such simulations prove challenging for "monolithic" time integration methods that treat all processes using a single approach. As the first talk in this mini-symposium, we will introduce the challenges that these types of problems impose on time integration methods, laying a common foundation for all the speakers in this session. Following that introduction, we will focus specifically on algorithms for multirate systems of ordinary differential equations. These applications are characterized by having different components of the problem that evolve at distinct time scales; multirate methods therefore allow evolution of these components using different time step sizes. Among such methods, multirate infinitesimal (MRI) variants allow near-complete freedom in how "fast" time scales are evolved. With this freedom, however, comes the challenge in determining optimal time steps to use at each scale. In this talk, we discuss our recent work to develop efficient methods to adapt time step sizes in MRI methods, and their implementation in the ARKODE solver library (a component of the SUNDIALS suite).

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Vu Luan
Texas Tech University
Department of Mathematics and Statistics
vu.luan@ttu.edu

MS205

Multirate Nonlinearly Partitioned RungeKutta Methods

We present Multirate Nonlinearly Partitioned Runge-Kutta Methods (MNPRK), which are a multirate extension of our recent work on nonlinearly partitioned Runge-Kutta methods. These MNPRK methods can be viewed as a nonlinear generalization of subcycling and additive multirate methods. We will discuss order conditions and stability properties of these methods related to directional nonlinear coupling between slow and fast scales. Some examples, such as implicitly wrapped strong stability preserving methods, will be presented.

Brian Tran
Los Alamos National Laboratory
btran@lanl.gov

Tommaso Buvoli
Tulane University
tbuvoli@tulane.edu

Ben Southworth
Los Alamos National Laboratory
southworth@lanl.gov

MS206

Super-Time-Stepping Methods for Diffusion Within Gyrokinetic Simulations

To foster the solution efforts of diffusion within gyrokinetic simulations of fusion plasmas, we have added new functionality to the ARKODE library in SUNDIALS to support a variety of Super Time Stepping (STS) methods. These methods have been designed for systems of ordinary differential equations where the right-hand side contains mildly stiff terms, but where one still wishes to treat those terms using an explicit time integration method. In order to circumvent the traditional linear stability limitations of explicit methods, STS methods utilize a large number (s) of internal stages, but instead of using these to maximize the order of accuracy STS methods maximize the linear stability. Of particular importance for the *practicality* of these methods is that they may be implemented using significantly fewer than s storage vectors, and thus these are among the class of "low-storage Runge-Kutta (LSRK) methods. In particular, the new ARKODE module supports the RungeKuttaChebyshev, RungeKuttaLegendre, and RungeKuttaGegenbauer families of STS methods. The module additionally supports a variety of strong stability preserving (SSP) LSRK methods. In this talk, we describe each of the above families of methods, their implementation in ARKODE, and demonstrate results that use these methods for gyrokinetic simulations of fusion plasmas.

Mustafa Aggul
Southern Methodist University
Department of Mathematics
maggul@smu.edu

MS206

Asymptotic-preserving and Energy-conserving Numerical Methods for Nonlinear Dispersive Wave Equations

We explore the development and application of IMEX time integration methods for nonlinear dispersive wave equations. Our primary focus is on first-order hyperbolic approximations, which offer the advantage of circumventing the need to discretize higher-order spatial and temporal derivatives. Additionally, we investigate asymptotic-preserving schemes, which preserve energy or entropy conservation laws on a discrete level.

Sebastian Bleecke

Johannes Gutenberg University Mainz
sbleecke@uni-mainz.de

MS206

Unconditional Strong Stability Preserving Additive Methods

In this talk we present a class of high order unconditionally strong stability preserving (SSP) implicit two-derivative RungeKutta schemes up to fourth order, SSP implicit-explicit (IMEX) multi-derivative RungeKutta schemes up to third order where the time-step restriction is independent of the stiff term, and unconditionally SSP additive downwind and multi-derivative RungeKutta schemes up to second order. The unconditional SSP property for a method of order p ¿ 2 is unique among SSP methods, and depends on a backward-in-time assumption on the derivative of the operator. We show that this backward derivative condition is satisfied in many relevant cases where SSP IMEX schemes are desired. We devise unconditionally SSP implicit RungeKutta schemes of order up to p = 4, and IMEX Runge Kutta schemes of order up to p = 3.; we also devise unconditionally additive methods of order 2 where one component is based on downwinding. For the multi-derivative IMEX schemes, we also derive and present the order conditions. These methods are positivity and asymptotic preserving when applied to a range of problems, including a hyperbolic relaxation system, the Broadwell model, and the Bhatnagar-Gross-Krook (BGK) kinetic equation. This is a joint work with AndrewChristlieb, Sining Gong, Zachary J. Grant, Jingwei Hu, and Ruiwen Shu.

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

MS206

New Operator Splitting Capabilities in the Suite of Nonlinear and Differential/Algebraic Equation Solvers (SUNDIALS)

For time-dependent simulations that couple several physical processes, operator splitting techniques are ubiquitous. Coupling is simple, and they allow each process to be evolved independently, possibly using a different integrator and time step tailored to the unique characteristics of that process. In this talk, we will discuss a new implementation of operator splitting methods in the SUNDIALS library. In addition to standard, low order Lie-Trotter and Strang splittings, it includes many high order schemes as well as support for custom coefficients. It builds on an extensive SUNDIALS infrastructure of integrators, (non)linear solvers, and vectors with support for parallel and GPU computing. This talk will cover implementation experiences, examples of how to use the methods, and numerical tests.

Steven B. Roberts, Carol S. Woodward
Lawrence Livermore National Laboratory
roberts115@llnl.gov, woodward6@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Cody J. Balos, David J. Gardner
Lawrence Livermore National Laboratory
balos1@llnl.gov, gardner48@llnl.gov

MS207

Accelerating Microstructure Evolution Predictions Via Fully Convolutional Network-Based Spatiotemporal Models

Advances in spatiotemporal predictive learning have opened new pathways for accelerating simulations in scientific domains, yet many methods remain computationally expensive due to complex architectures and training pipelines. In this work, we develop a framework based on SimVPv2, a fully convolutional neural network model, for accelerating microstructure evolution predictions. Building on the simplicity of SimVPv2, the proposed approach eliminates the need for heavy U-Net-like architectures, leveraging streamlined stacks of convolutional layers combined with an efficient Gated Spatiotemporal Attention mechanism. This design achieves state-of-the-art performance while significantly reducing computational complexity. Our approach demonstrates superior predictive accuracy and efficiency on benchmark microstructure evolution datasets, achieving faster training and inference with fewer FLOPs. Extensive evaluations across diverse microstructural processes, such as grain growth and Spinodal decomposition, highlight the model's robustness and versatility. By offering a powerful yet straightforward solution, our framework establishes a solid baseline for spatiotemporal learning in materials science, enabling rapid and reliable predictions to advance material design and discovery workflows. This is a joint work with Michael Trimboli, Ke-gang Wang, and Sirani M. Perera

Xianqi Li
Department of Mathematics and Systems Engineering
Florida Institute of Technology
xli@fit.edu

MS207

Quadratic B-Splines and the Geometry of Non-Planar 3D Fracture Growth

Physics based fracture growth modelling is computationally expensive because of the need to solve coupled partial differential equations at different scales and its associated re-triangulation requirements. This justifies working with analytical representation of fractures such as NURBS. Instead of the latter, we propose a simplified representation using periodic B-splines. Then fracture growth reduces to adding a periodic 3D curve to a preexisting sequence curves, namely the fracture representation in the previous iteration. The main advantages over the NURBS representation are the ability to estimate the stress intensity factors at any point of the fracture using simple interpolation techniques and the possibility of polyhedral bounding of the whole fracture. This representation shares with the NURBS representation the possibility of fracture refinement to any resolution.

Marco Paluszny
Universidad Nacional de Colombia
Cl. 59a #63-20, Medellín, Antioquia
mpalusznykl@unal.edu.co

Adriana Paluszny
Imperial College
apaluszn@imperial.ac.uk

Ellya Kanimova
Imperial College, London
ellya.kanimova21@imperial.ac.uk

MS208

A Geometric Objective for Graph Coarsening with Machine Learning Applications

In this talk, I will delve into a geometric perspective of graph coarsening, which is a technique for solving large-scale graph problems by working on a smaller version of the original graph. It has a long history in scientific computing and has recently gained popularity in machine learning, particularly in methods that preserve the graph spectrum. Compared to the previous spectral perspective, the proposed geometric perspective is especially useful when working with a collection of graphs, such as in graph classification and regression. Specifically, we consider a graph as an element on a metric space equipped with the Gromov–Wasserstein (GW) distance, and bound the difference between the distance of two graphs and their coarsened versions. Minimizing this difference can be done using the popular weighted kernel K-means method, which strengthens existing spectrum-preserving methods.

Yifan Chen
Hong Kong Baptist University
yifanc@hkbu.edu.hk

Rentian Yao
University of Illinois at Urbana-Champaign
University of British Columbia
rentian2@illinois.edu

Yun Yang
UIUC
USA
yy84@illinois.edu

Jie Chen
MIT-IBM Watson AI Lab, IBM Research
chenjie@us.ibm.com

MS208

Compressing Deep Learning Models with Interpolative Decompositions

After training complex deep learning models, a common task is to compress the model to reduce compute and storage demands. When compressing, it is desirable to preserve the original models per-example decisions (e.g., to go beyond top-1 accuracy or preserve robustness), maintain the networks structure, automatically determine per-layer compression levels, and eliminate the need for fine tuning. We introduce a principled approach that satisfies these criteria. Our technique leverages techniques from (randomized) numerical linear algebra to effectively compress networks and is adaptable to many types of network layers (including attention mechanisms in transformers). We will discuss the central mathematical ideas that enable our methodology to work and scale to very large models, and provide extensive numerical experiments that illustrate its efficacy.

Megan Flynn
Cornell University
mr2268@cornell.edu

Anil Damle
Cornell University
Department of Computer Science
damle@cornell.edu

MS208

Sympgnns: Symplectic Graph Neural Networks for Identifying High-Dimensional Hamiltonian Systems and Node Classification

Existing neural network models to learn Hamiltonian systems, such as SympNets, although accurate in low-dimensions, struggle to learn the correct dynamics for high-dimensional many-body systems. Herein, we introduce Symplectic Graph Neural Networks (SympGNNs) that can effectively handle system identification in high-dimensional Hamiltonian systems, as well as node classification. SympGNNs combines symplectic maps with permutation equivariance, a property of graph neural networks. Specifically, we propose two variants of SympGNNs: i) G-SympGNN and ii) LA-SympGNN, arising from different parameterizations of the kinetic and potential energy. We demonstrate the capabilities of SympGNN on two physical examples: a 40-particle coupled Harmonic oscillator, and a 2000-particle molecular dynamics simulation in a two-dimensional Lennard-Jones potential. Furthermore, we demonstrate the performance of SympGNN in the node classification task, achieving accuracy comparable to the state-of-the-art. We also empirically show that SympGNN can overcome the oversmoothing and heterophily problems, two key challenges in the field of graph neural networks.

Alan John Varghese
School of Engineering, Brown University
alan_john_varghese@brown.edu

Zhen Zhang
Brown University, U.S.
zhen_zhang2@brown.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS208

Efficient Visualization of Implicit Neural Networks Via Weight Matrix Analysis

An implicit neural representation (INR) is a neural network that approximates a function over space and possibly time. Memory-intensive visualization tasks, including modern 4D CT scanning methods, represent data natively as INRs. While such INRs are prized for being more memory-efficient than traditional data on a lattice, discretization to a regular grid is still required for many visualization tasks. In this talk, we present an algorithm to store high-resolution voxel data only for regions with significant detail, reducing memory requirements. To identify these high-detail areas, we use an interpolative decomposition pruning method on the weight matrices of the INR. The information from pruning is used to guide adaptive mesh refinement, allowing automatic mesh generation, tailored to the underlying resolution of the function.

Jennifer E. Zvonek
Cornell University
jez34@cornell.edu

Andrew Gillette
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
gillette7@llnl.gov

MS209

Co-Evolving Networks for Opinion and Social Dynamics in Agent-Based Models

Digital media has fundamentally transformed social interactions and reshaped the ways in which individual opinions influence and are influenced by these interactions. Thus, investigating the interplay between opinion and social dynamics has become crucial for understanding the formation of social structures and collective outcomes. We extended state-of-the-art models by introducing a stochastic agent-based model that captures the co-evolution of opinion and social dynamics [Djurdjevac Conrad, Kppl, Djurdjevac; Feedback Loops in Opinion Dynamics of Agent-Based Models with Multiplicative Noise, Entropy24(10), 2022]. Our model considers agents moving in a social space influenced by both positions and opinions of other. Agents with similar opinions of the same stance exhibit social closeness, while opinion dissimilarity reinforces social distancing between agents. Opinion dynamics is driven by agents spatial proximity and their opinion similarity. By analyzing the underlying social and opinion interaction networks, we explore the mechanisms influencing the appearance of emerging phenomena, like echo chambers and opinion consensus [Djurdjevac Conrad, Vu, Nagel; Co-evolving networks for opinion and social dynamics in agent-based models, Accepted in Chaos:An interdisciplinary Journal of Nonlinear Science, 2024]. We apply our model to survey data and show that it can capture the co-evolution of political identity and individual opinions regarding governmental issues.

Natasa Djurdjevac Conrad, Sören Nagel
Zuse Institute Berlin
natasa.conrad@zib.de, nagel@zib.de

Nhu Quang Vu
Freie Universität Berlin
vunh99@zedat.fu-berlin.de

Jonas Köppl
Weierstrass Institute for Applied Analysis and Stochastics
koeppl@wias-berlin.de

Ana Djurdjevac
Freie Universität Berlin
adjurdjevac@zedat.fu-berlin.de

MS209

How to Win An Election in a Dynamically Evolving Opinion Space

We model dynamically changing candidate positions in the face of a dynamic electorate. We use Hegselmann-Krause opinion dynamics to model the changes in the electorate. Candidates move along the gradient and are maximizing the votes. We use the combined candidate-voter model to demonstrate the possibility of discontinuous jumps in candidate behavior as parameters of the model are varied.

Natasa Dragovic
University of Saint Thomas
natasa.dragovic@stthomas.edu

Christoph Borgers
Tufts University
christoph.borgers@tufts.edu

Arkadz Kirshtein
Texas A&M -Corpus Christi, Texas
arkadz.kirshtein@tamucc.edu

MS209

Inverse Adaptive Learning Design on Networks

Learning the rules that govern self-organization in complex systems is crucial for advancing our ability to model and design such systems. Generalized Kuramoto models have been widely employed for this purpose, effectively capturing the synchronization of systems of oscillators. However, these models require prior knowledge of the governing equations, which limits their flexibility and applicability. To overcome this challenge, we benchmark modern machine learning approaches, including Graph Neural Networks (GNNs) and Koopman Augmented Networks (KANs), to determine their effectiveness in inferring and predicting emergent dynamics without relying on explicit equation forms. Our benchmarks reveal the strengths and limitations of these tools in capturing synchronization and swarming behaviors across various network topologies and conditions, providing insights into their potential for advancing the study of self-organization in complex systems. This work contributes to the growing field of machine learning-driven modeling, offering a pathway toward more flexible and scalable approaches in adaptive system design.

Bianca Dumitrascu
Columbia University, New York
bmd2151@columbia.edu

MS209

A New Dangling Centrality Metric and Its Comparison with Existing Centrality Measures for Real World Datasets

In this lecture, I will introduce the Dangling Centrality metric which is a new vision to social network measures and analyze its relationship with established centrality metrics across diverse datasets, including Amazon product networks, a Protein-Protein Interaction (PPI) network, and a Bitcoin network. Our findings highlight the importance of considering node absence in network dynamics, with Dangling Centrality showing strong correlations with traditional metrics. Through analysis of both complex and simple networks, I will demonstrate the metric's ability to identify key products, proteins, and individuals, with results supported by regression plots and Pearson's correlation coefficients.

Ubaida Fatima
NED University of Engineering and Technology Karachi
NED University of Engineering and Technology Karachi
ubaida@neduet.edu.pk

MS209

Modeling Candidate-Voter Dynamics

We consider a candidate-voter dynamics model based on the Hegselmann-Krause opinion dynamics reformulated from the probability density perspective. We analyze properties of the resulting PDE and look at their effect on the

original particle model.

Arkadz Kirshtein
Texas A&M -Corpus Christi, Texas
arkadz.kirshtein@tamucc.edu

Natasa Dragovic
University of Saint Thomas
natasa.dragovic@stthomas.edu

MS210

Efficient Numerical Methods for PDE-constrained Optimization Problems in Diffeomorphic Image Registration

We present our work on efficient algorithms for diffeomorphic image registration, formulating the problem as an optimization task governed by transport equations for image intensities and a geodesic equation on the group of diffeomorphisms. This inverse problem is inherently ill-posed and infinite-dimensional in the continuum, leading to high-dimensional, ill-conditioned inversion operators after discretization. To address these challenges, we develop efficient numerical methods focused on optimization, including strategies for numerical time integration and preconditioning. We validate the performance of our methods through tests on both synthetic and real-world data, demonstrating their effectiveness in overcoming the mathematical and computational complexities involved.

Jannatul Ferdous Chhoa
University of Houston
jchhoa@cougarnet.uh.edu

Yunhui He
Department of Mathematics
University of Houston
yhe43@central.uh.edu

Andreas Mang
University of Houston
Department of Mathematics
amang@Central.UH.EDU

MS210

Towards a Unified Electro-Fluid-Mechanical Framework for Cardiac Modeling

The electrical activity of the heart is responsible for initiating the contraction of cardiac muscle, and the specialized cardiac conduction system (CCS) is responsible for coordinating the activation of the heart's chambers. Consequently, cardiac rhythm disorders can impair or even totally disrupt the pumping function of the heart. This talk details progress on work to incorporate a comprehensive model of cardiac electrophysiology within a detailed model of cardiac fluid dynamics and fluid-structure interaction (FSI). We address FSI using the immersed finite element-difference method, an extension of the immersed boundary method for finite element-based structural mechanics models. Our integrated EP-FSI model includes tissue-scale models of electrical impulse propagation within the cardiac chambers based on the monodomain or bidomain equations along with a CCS model that includes the sinoatrial node, atrioventricular node, His bundles, and Purkinje fibers. This talk will describe key numerical approaches, including finite element models that integrate volumetric and cable-like models of cardiac electrophysiology, and will detail initial simulation results obtained using this integrated model.

Laryssa Abdala
University of North Carolina
Department of Mathematics
laryssa@live.unc.edu

Simone Rossi
University of North Carolina
srossi@email.unc.edu

David Wells
University of North Carolina at Chapel Hill
drwells@email.unc.edu

Boyce E. Griffith
University of North Carolina - Chapel Hill
boyceg@email.unc.edu

MS210

Extraction of Coronary Vessel Networks from Computed Tomography and Corresponding Flouroscopy Images

Vessel network models extracted strictly from an MRA have proven to be lacking in completeness and consistency. Flouroscopy images have different, but complementary properties from their corresponding MRA images. If corresponding flouroscopy images are available then they can be used, in combination with the MRA image, to generate enhanced vessel network models thus alleviating some of the gaps and errors present in the strictly MRA-based models. A necessary prerequisite to the combination of these images is a solution to a multimodal and multidimensional registration problem. Results of a centerline-based approach to registration is presented for a cohort of 12 sets of patient data each consisting of a single MRA image and several corresponding flouroscopy images are presented.

Dan Lior
Baylor College of Medicine
dan.lior@bcm.edu

MS210

Closed Loop Solute Transport in Blood Vessels and Organs

Hypoplastic left heart syndrome (HLHS) is a congenital heart disease that accounts for 2-3% of congenital heart diseases in the United States and 40% of all neonatal cardiac deaths. HLHS causes oxygenated blood to mix with deoxygenated blood, resulting in death. This raises a critical need to accurately model the transport of oxygen in blood vessels and organs throughout the human body to improve outcomes in patients with HLHS. Previously, numerical reduced models have been created to solve for blood flow and concentration of one solute. These models reduce the dimensions of the vessels and organs to improve computational efficiency. This work extends the models from open network of blood vessels to closed loops and includes organs such as the heart. Appropriate transmissibility conditions at each vessel junction and organ bed are constructed, that are based on balance laws. The class of interior penalty discontinuous Galerkin methods is used for the discretization of the models.

Arshia Singhal
Rice University

as327@rice.edu

MS211

Variable Projection Methods for Large-scale Separable Nonlinear Inverse Problems

Variable projection methods are among the classical and efficient methods for solving separable nonlinear least squares problems. This talk will introduce variable projection methods for solving large-scale inverse problems and a local convergence analysis when inexact inner solvers are used. Numerical experiments will be presented, demonstrating the effectiveness of these methods in the context of semi-blind deblurring problems.

Malena I. Espanol
Arizona State University
malena.espanol@asu.edu

MS211

Generalized Sparse Bayesian Learning Techniques for Inverse Problems

Recovering sparse generative models from limited and noisy measurements presents a significant and complex challenge. Given that the available data is frequently inadequate and affected by noise, assessing the resulting uncertainty in the relevant parameters is crucial. Notably, this parameter uncertainty directly impacts the reliability of predictions and decision-making processes. In this talk, we explore the generalized sparse Bayesian learning framework. This advanced statistical modeling technique facilitates the quantification of uncertainty in parameter estimates by treating involved quantities as random variables. The strategy employs distinct conditionally Gaussian priors for each parameter vector alongside generalized gamma-distributed hyper-parameters to impose sparsity. By amalgamating these sparsity-promoting priors with established Bayesian inference techniques, a new family of algorithms is developed.

Jan Glaubitz
Massachusetts Institute of Technology
Glaubitz@mit.edu

Anne Gelb
Dartmouth College
Department of Mathematics
annegelb@math.dartmouth.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

MS211

Carleman Estimates for Mean Field Games

The theory of Mean Field Games is a powerful apparatus, which can model almost any societal phenomenon via a system of two nonlinear parabolic partial differential equations. However, the price to pay is a huge complexity of this system. To address the complexity issue, Carleman estimates were introduced recently in this field by Klibanov. As a result, stability and uniqueness theorems for the first time were obtained for both forward and inverse problems of Mean Field Games. Furthermore, Carleman estimates enable one to construct various versions of the effective convexification method for these problems. Both theoretical and numerical results will be presented.

Michael Klibanov
University of North Carolina at Charlotte
mklibanv@uncc.edu

MS211

Deep-learning-assisted Real-time Algorithms for Inverse Problems Governed by Partial Differential Equations

Inverse problems are pervasive in sciences, engineering, and mathematics. Inverse problems are challenging due to their ill-posed nature. Inverse problems governed by nonlinear PDEs are even more challenging due to prohibited computational requirements. Yet many practical problems, including data assimilation, parameter calibration, and imaging, require real-time inverse solutions. We propose an approach to overcome the aforementioned challenge by synergizing rigorous computational inverse approaches with the speed of deep learning. We present several approaches from purely data-driven to rigor-driven deep learning methods to solve large-scale inverse problems. Several results are presented for inverse problems governed by heat equations, Burgers' equations, and Naver-Stokes equations.

Tan Bui-Thanh, Hai Nguyen
The University of Texas at Austin
tanbui@oden.utexas.edu, hainguyen@utexas.edu

Russell Philley
University of Texas at Austin, U.S.
rsphilley@utexas.edu

MS211

On the Problem of Active Control of Fields Through the Lenses of Integral Operators, Inverse Sources and Optimization

In this talk we will present first a theoretical introduction for the general problem of active control of fields and then focus on the particular case of electromagnetic fields. We will then show how by employing techniques from inverse source analysis and through suitable layer potential operators one can prove the existence of a class of active control sources achieving the desired effects. We then present how these controls can be further studied through the equivalent formulation of the underlying problem as a constrained minimization problem and in particular focus on the implementation of a practical strategy for far field control of electromagnetic fields by using a swarm of flying drones with applications to communications. Numerical support will be presented for several interesting applied scenarios.

Daniel Onofrei
Department of Mathematics
University of Houston
dtonofre@central.uh.edu

MS212

Galerkin RBF Operator Learning and Surrogate Models

Data-driven physics models and operators require efficient representation of solutions; one such efficient basis is radial basis functions. We build a Gaussian radial basis function approximation scheme, where the shape parameters of the Gaussians are informed by data. With these RBFs,

we solve a Galerkin formulation of the desired operator or PDE. Our approach to do so rests fundamentally on the observation that products of Gaussian RBFs and their gradients are integrable as polynomial moments of a product distribution with a closed-form expressions to assemble mass and stiffness matrices without quadrature error. This straightforward observation enables a meshfree discretization requiring no numerical quadrature, making it suitable for high-dimensional problems. We prove that variational problems posed in this Gaussian basis converge as the number of Gaussian RBFs is increased, and we show results for a handful of PDE data recovery problems, highlighting the convergence and flexibility of our approach.

Jonas Actor
Sandia National Laboratories
jaactor@sandia.gov

MS212

Neural Networks in Numerical Pdes

In this talk, we will first give a brief introduction of NNs from numerical analysis perspective and use a simple example to show why NNs are superior to piecewise polynomials on fixed meshes when approximating discontinuous functions with unknown interface. We will then describe a space-time least-squares neural network (LSNN) method. The method shows a great potential to sharply capture shock without oscillation, overshooting, or smearing. The exceptional approximation powers of NN come with a price: the procedure for determining the values of the nonlinear parameters of NN entails solving a high-dimensional nonconvex optimization problem. We will describe our newly developed training algorithm for shallow ReLU NN.

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Min Liu
Purdue University
liu66@purdue.edu

MS212

Efficient and Accurate Training of Physics-Informed Deep Operator Networks with the Conjugate Kernel

Recent work has shown that the empirical Neural Tangent Kernel (NTK) can significantly improve the training of physics-informed Deep Operator Network (DeepONets). The NTK, however, is costly to calculate, greatly increasing the cost to train such systems. In this paper, we study the performance of the empirical Conjugate Kernel (CK) for physics-informed DeepONets, an efficient approximation to the NTK that has been observed to yield similar results. For physics-informed DeepONets, we show that the CK performance is comparable to the NTK, while significantly reducing time complexity for training DeepONets.

Sarah Murphy
University of North Carolina, Charlotte
shelfert@charlotte.edu

Amanda Howard, Tony Chiang, Panos Stinis
Pacific Northwest National Laboratory
amanda.howard@pnnl.gov, tony.chiang@pnnl.gov, panagiotis.stinis@pnnl.gov

MS212

MITONet: A Novel Data-Driven, Physics-Guided Neural Emulator for Predicting Coastal Hydrodynamics

We introduce the Multiple Input Temporal Operator Network (MITONet), a novel approach that leverages recent advancements in operator learning, specifically within the Deep Operator Network framework. MITONet is designed to efficiently emulate high-dimensional numerical solvers for complex, non-linear problems governed by time-dependent, parametric partial differential equations. By utilizing a latent space representation and incorporating temporal bundling, MITONet provides a powerful and flexible solution for capturing intricate dynamics over time. We demonstrate its effectiveness in forecasting hydrodynamic processes, such as ocean currents and wave propagation, emphasizing its capability to generate accurate predictions for complex coastal dynamics. Our results highlight MITONets accuracy, even in scenarios involving out-of-distribution data, demonstrating its potential for real-world applications.

Peter Rivera-Casillas
US Army Corps of Engineers ERDC
Vicksburg, MS
peter.g.rivera-casillas@erdc.dren.mil

Sourav Dutta
The University of Texas at Austin, Oden Institute
sourav.dutta@austin.utexas.edu

Shukai Cai
University of Texas at Austin
Oden Institute for Computational Engineering & Sciences
shukai.cai@utexas.edu

Mark Loveland
US Army Corps of Engineers ERDC
Information Technology Laboratory
mark.d.loveland@erdc.dren.mil

Corey Trahan
US Army Engineer Research & Development Center
corey.j.trahan@erdc.dren.mil

Matthew W. Farthing
US Army Engineer Research and Development Center
matthew.w.farthing@usace.army.mil

Jonghyun Lee
University of Hawaii at Manoa
jonghyun.harry.lee@hawaii.edu

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

MS213

A Truncated Generalized Singular Value Decomposition for Comparative Analysis of Noisy Data

Among the most important and powerful matrix factorizations is the singular value decomposition (SVD), which

provides dimensionality reduction and denoising benefits through its provably optimal low-rank representations. In this presentation, we will bring the optimality of the truncated SVD to comparative data analysis via the generalized SVD (GSVD). Evidenced in applications ranging from genomics to inverse problems, the GSVD reveals features that are either common between two matrices, or exclusive to either one of them. However, for cases in which the data are noisy, the patterns identified by the GSVD can inherit noisy properties as well. To reveal non-noisy features, we propose an algorithm for a truncated GSVD by first optimally reducing noise through a truncated SVD followed by a statistically-motivated GSVD procedure on rank-deficient matrices. We provide insights into the theoretical guarantees of our new algorithm, and empirically demonstrate performance through illustrative examples.

Elizabeth Newman, Yizhou Chen
Emory University
elizabeth.newman@emory.edu, yizhou.chen@emory.edu

Orly Alter
University of Utah and Prism AI Therapeutics, Inc.
orly@sci.utah.edu

MS213

Joint Analysis of X-ray Ptychography and X-ray Fluorescence

X-ray fluorescence can be used for revealing the internal elemental composition of a sample, while x-ray ptychography can produce high-resolution and high-contrast transmission. In this work, we exploit the complementarity between X-ray fluorescence and ptychography, integrate both data modalities and formulate a nonlinear optimization-based approach for reconstruction of the elemental composition of a given object with improved resolution.

Elle Buser
Emory University
ebuser@emory.edu

Yuanzhe Xi
Department of Mathematics
Emory University
yuanzhe.xi@emory.edu

Zichao Di
Argonne National Lab
wendydi@mcs.anl.gov

MS213

Structured Column Subset Selection with An Application to Optimal Experimental Design

This paper addresses a structured column subset selection problem (CSSP) with applications to optimal experimental design for sensor placement. Traditional CSSP methods select the most informative columns from a matrix. We extend these methods to structured column selections, focusing on scenarios where columns represent observations under different experimental conditions and are grouped by experimental variables or sensor configurations. We develop two tensor-based algorithms for structured CSSP: TenSVD and SVDTen. Both use the Golub-Klema-Stewart method and data structure to enhance computational efficiency. We analyze their computational costs and performance guarantees under the D-optimality criterion, demonstrating their effectiveness through model problems in Bayesian inverse settings.

Hugo Diaz-Norambuena
North Carolina State University
hsdiazno@ncsu.edu

Srinivas Eswar
Mathematics and Computer Science Division
Argonne National Laboratory
seswar@anl.gov

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS213

Empowering Scientific Research with a Scalable, Hardware-Agnostic Tiled Linear Algebra Framework in Julia

High-Performance Computing (HPC) is essential for scientific research, enabling complex simulations and analyses across fields. However, the specialized knowledge required to utilize HPC effectively can be a barrier for many scientists. This work leverages state-of-the-art scalable linear algebra and abstract array interfaces to provide a hardware-agnostic, mixed-precision, scalable tiled abstraction layer to enhance accessibility and usability. It takes advantage of multiple-dispatch to build upon earlier work in abstract interfaces providing hardware-agnosticism and mixed-precision. It also integrates tiled out-of-core linear algebra approaches. The new abstraction layer that implements both aspects facilitates the extensibility of the resulting tiled abstraction towards any tiled operation. It provides end-users with a single function call for all use cases, and developers with the applicability of a single implementation to any hardware and precision. Developers also benefit from a data dependency-driven scheduling implementation system. We demonstrate the practical benefits by implementing the Singular Value Decomposition (SVD), showing how the approach streamlines the development process and accelerates scientific discovery. This work aims to democratize HPC resources by bridging the gap between advanced computational capabilities and user accessibility, empowering a broader range of scientists to fully leverage modern computing technologies.

Evelyne Ringoot, Valentin Churavy, Julian Samaroo
Massachusetts Institute of Technology
eringoot@mit.edu, vchuravy@mit.edu,
jpsamaroo@gmail.com

Rabab Alomairy
KAUST
rabab.omairy@kaust.edu.sa

Alan Edelman
Department of Mathematics
Massachusetts Institute of Technology and Julia Computing
edelman@mit.edu

MS214

Space-Time Non-Local Multi-Continua Multiscale

Method for Channelized-Media Parabolic Equations

Many applied problems feature coupled spatial and temporal heterogeneities that complicate homogenization or upscaling. Non-local multi-continua are proposed to address complex spatial heterogeneities in various papers. We extend this approach to a parabolic problem with time-dependent heterogeneous coefficients, developing the space-time Non-local multi-continua (space-time NLMC). This method efficiently constructs multiscale basis functions through local energy minimization within oversampled space-time regions. These functions demonstrate exponential decay outside their domain and offer computational advantages over classical methods, providing a systematic and flexible way to achieve accurate solutions while reducing computational overhead.

Jiuhua Hu
Department of Mathematics
University of Wisconsin - Madison
jiuhuahu@vt.edu

Wing Tat Leung
City University of Hong Kong
wtleung27@cityu.edu.hk

Eric Chung
The Chinese University of Hong Kong
Department of Mathematics
eric.t.chung@cuhk.edu.hk

Yalchin Efendiev
Dept of Mathematics
Texas A&M University
efendiev@math.tamu.edu

MS214

Partially Explicit Generalized Multiscale Finite Element Methods for Poroelasticity Problem

We develop a partially explicit time discretization based on the framework of constraint energy minimizing generalized multiscale finite element method (CEM-GMsFEM) for the problem of linear poroelasticity with high contrast. Firstly, dominant basis functions generated by the CEMGMsFEM approach are used to capture important degrees of freedom and it is known to give contrastindependent convergence that scales with the mesh size. In typical situation, one has very few degrees of freedom in dominant basis functions. This part is treated implicitly. Secondly, we design and introduce an additional space in the complement space and these degrees are treated explicitly. We also investigate the CFL-type stability restriction for this problem, and the restriction for the time step is contrast independent.

Xiu Su
University of California Santa Barbara
xsu2@ucsb.edu

Wing Tat Leung
City University of Hong Kong
wtleung27@cityu.edu.hk

Wenyuan Li, Sai-Mang Pun
Texas A&M University
wenyuanli@tamu.edu, smpun@math.tamu.edu

MS214

The Interplay Between Deep Learning and Model Reduction

The integration of Reduced Order Modeling (ROM) and Deep Learning (DL) presents a promising avenue for enhancing computational efficiency and predictive accuracy. ROM techniques reduce complex systems into low-dimensional representations, preserving essential dynamics, while DL methodologies excel at learning intricate patterns from raw data. By com- bining ROM with DL, we aim to develop approaches that inherit merits from both disciplines. On one hand, we explore leveraging ROM as a preprocessing step to train DNNs with limited labeled data, addressing data scarcity. On the other hand, we intend to utilize DL to expedite ROM construction and learn ROMs directly from observational data. In this talk, we will elaborate on specific efforts made along this line.

Min Wang
University of Houston
mwang55@cougarnet.uh.edu

MS214

Multimodality Scientific Foundation Models for Multiscale Multiphysics Problems

Multi-operator learning uses an operator embedding structure to train a single neural network on data from multiple operators. Thus, multi-operator learning is capable of predicting a range of operators within one model. We propose pretraining and fine-tuning strategies for solving PDEs using multi-operator learning. One key aspect is that by increasing the number of families of operators used in pretraining, a PDE foundation model can be fine-tuned to downstream tasks involving new PDEs with a limited number of samples, thus outperforming single operator neural networks. Specifically, a multi-operator learning model pre-trained with data from diverse PDE families can predict unseen operators after fine-tuning with only a limited number of operators from the new family, enabling them to serve as a data-free PDE solver. We also show that the proposed training and fine-tuning method is able to predict new operators in zero-shot prediction without samples. Additionally, we introduce a PDE-agnostic meta-learning algorithm to improve the adaptability of the model to various PDEs by providing a better parameter initialization process. To address the needs of applications with limited computing resources, we explore low-rank adaptation methods that reduce computational costs while enhancing solver accuracy. Lastly, by examining the scaling law with respect to the number of operator families, we establish and highlight its potential for broad adaptation in PDE-solving tasks.

Zecheng Zhang
Florida State University
zz23e@fsu.edu

Jingmin Sun
Department of Mathematical Science
Carnegie Mellon University
jingmins@andrew.cmu.edu

Hayden Schaeffer
Department of Mathematics
University of California, Los Angeles, U.S.

hayden@math.ucla.edu

MS215

Satisfying the Fluctuation Dissipation Balance with Adaptive Mesh Refinement

In computational thermodynamics, an additional term is added to the fluid equations to account for microscopic fluctuations acting as a reverse action to the energy dissipation. Choosing a fluctuation term proportional to the square root of the viscous operator ensures that the fluctuations will neither overwhelm nor be overwhelmed by the dissipation in the system. However, when adaptive mesh refinement is introduced to the computational domain, the natural viscous operator is no long positive definite and satisfying the fluctuation dissipation balance becomes nontrivial. In this poster we 1) outline a second order finite difference scheme for a MAC grid that yields a symmetric laplace operator L and 2) prove this operator is also positive definite by deriving the corresponding divergence operator D such that L = DDT and D satisfies the fluctuation dissipation balance.

Bryn Barker
University of North Carolina at Chapel Hill
brynbarker@unc.edu

MS215

LibMobility: a Fast, GPU Library for Simulating Fluctuating Stokesian Suspensions

We introduce a fast and modular GPU library, libMobility, that solves the hydrodynamic mobility problem for Stokesian particle suspensions in open, periodic, and confined geometries. In particular, we've implemented fast routines to apply mobility matrices and their principle square root, ensuring that the library is immediately available for use in Brownian dynamics applications. In addition to being a high-performance library, libMobility is designed to have a streamlined installation and be easy to integrate into existing codes. We present examples with fluctuating membranes, slender filaments, and active particles that demonstrates both the flexibility and the computational advantage of the library.

Ryker Fish
Colorado School of Mines
rfish@mines.edu

MS215

A Posteriori Error Estimator for Poroelasticity with Uncertain Inputs

This presentation provides an overview of recent advances in stochastic Galerkin mixed finite element methods (SG-MFEMs) applied to parameter-dependent linear elasticity equations. We begin with a unique three-field partial differential equation (PDE) model, where Youngs modulus is represented as an affine function of a finite or countable set of parameters. The presentation then examines the SG-MFEM approximation. Furthermore, we introduce a novel a posteriori error estimation scheme that assesses errors in the natural weighted norm crucial for the stability of the weak formulation. The reliability and efficiency of the proposed error estimator are demonstrated, with emphasis on the independence of the associated constants from both the Poisson ratio and the SG-MFEM discretization parameters, ensuring robustness in the incompressible limit. The first part concludes by exploring error reduction proxies related to potential enrichments of SG-MFEM spaces, proposing their use in developing an adaptive algorithm that halts when the estimated error meets a user-defined tolerance.

Arbaz Khan
IIT Roorkee
arbazkha@gmail.com

MS215

Dressed Diffusion By Hydrodynamic Coupling to a Soft Mode

Interfaces introduce non-trivial effects on the motion of a particle diffusing nearby. The effect of hard surfaces is well documented, but the influence of soft boundaries on Brownian motion is yet to be understood. Here, we study the motion of a diffusing particle interacting with a soft mode through a hydrodynamic-like interaction. We propose a toy model consisting of a point-like particle diffusing in a sinusoidal trap to which a second harmonically-trapped point-like particle is coupled through a hydrodynamic-like interaction. We derive analytical formulas for the long-time diffusivity. We show that the diffusivity of the particle is enhanced by the elastrohydrodynamic coupling. Our results are supported by Langevin numerical simulations.

Juliette Lacherez
University of Bordeaux
juliette.lacherez@u-bordeaux.fr

MS215

Applications of Stochastic Methods in Simulating Mesoscale Flows

In this minisymposium we examine some specific numerical implementations of stochastic fluid dynamics algorithms, with examples including Fluctuating Hydrodynamics (FHD), and Browian, Stokesean, and Langevin dynamics. In each case this will involve coverage of material including Greens function and grid based approaches to modelling systems such as suspensions and thin films, with related topics such as Peskin Kernels and grid refinement also discussed.

Daniel R. Ladiges
Center for Computational Science and Engineering
Lawrence Berkeley National Laboratory
drladiges@lbl.gov

MS216

Hassle-Free Prediction of Excitation Wave Quenching Using Actors

Cardiac arrhythmias can be viewed as undesirable excitations in myocardial tissue. From this perspective, the role of a defibrillator is to quench these excitations and return the heart to its normal function. We are specifically interested in the quenching of stable (travelling-wave) excitations, which is a more insidious yet less well-studied problem. Methods for finding protocols and thresholds for such quenching are manually as well as computationally intensive. We use the actor model of concurrent programming to accelerate threshold prediction for excitation wave quenching. Our novel approach employs any number of actors to carry out the simulations of potential thresholds in parallel with a supervisor actor making decisions about the ongoing simulations. In particular, when an actor's

simulation finishes, the supervisor can use the new information to reason about all other ongoing simulations. The supervisor can then, for example, start a new simulation with an updated version of the threshold. It can also take information from completed simulations to stop ones that are deemed unnecessary and re-allocate resources to new ones. This approach accelerates time to solution, albeit at an increased computational cost overall, because some computation is necessarily wasted when trying to anticipate future simulations. However, a major advantage is the hassle-free nature of the process, which essentially runs from start to finish with no further intervention from the user.

Kyle Klenk
Department of Computer Science
University of Saskatchewan
kyle.klenk@usask.ca

MS216

Multi-Fidelity Approaches to Cardiac Reconstruction with Ensemble Kalman Filters

In cardiac experiments optical mapping is used to measure the spatiotemporal dynamics of excitation across the surface(s) of the heart. The reconstruction of excitation patterns through the unobserved depth of the tissue is essential to understanding arrhythmogenic patterns and realizing the potential of computational models in cardiac medicine. We have interpolated experimental recordings of periodic cardiac stimulation on the epicardial and endocardial surfaces of a section of canine ventricle, which are treated as observations with a prescribed uncertainty in an ensemble data assimilation scheme. We model the physical system underlying these observations with a hierarchy of approximations, from low-dimensional phenomenological models to comparatively high-dimensional ionic models. This approach combines efficient exploration of generative processes with low-fidelity models, while maintaining the interpretability of high-fidelity models. Results from a variety of spatial and temporal regimes are reported.

Christopher Marcotte
Computer Science
Durham University
christopher.marcotte@durham.ac.uk

MS216

Adaptive Design of Computer Experiments for Multi-Fidelity Emulation

Many biological systems are represented by complex mathematical equations, such as ODEs or PDEs, implemented as computer codes. Examples include models designed to simulate protein structures or the electrical activities of neurons. Often, a hierarchy of computer models is available, offering varying levels of fidelity in describing the system under study. While higher-fidelity simulators provide more accurate representations, they also come with higher computational costs compared to lower-fidelity models. To mitigate these costs, expensive models are often approximated by statistical surrogates, which are significantly faster to run. Gaussian process (GP) emulators are a key class of surrogate models. Training a GP requires a set of training runs, a.k.a., design points, typically selected by a space-filling design that ensures uniform distribution across the input space. However, recent interest has shifted towards adaptive designs, where the selection of the next sample depends on previously observed model outputs. To address this, we propose a novel adaptive design method called Variance of Improvement for Global Fit (VIGF). This method is extended to multi-fidelity emulation, where evaluations from both low- and high-fidelity models are integrated to effectively leverage all available information. We assess the applicability of VIGF on a set of test functions, with results suggesting that it outperforms other methods in predicting these functions in most cases.

Hossein Mohammadi
University of Exeter
h.mohammadi@exeter.ac.uk

MS217

Mass-Conservative Emulation of a Compatible Finite Element Model with a Graph Neural Network

The computational burden associated with running ensemble simulations of glacier change - and the resulting restriction on rigorous uncertainty quantification - is among the principal scientific bottlenecks for actionable predictions of land ice change. In an effort to make models faster, we present a graph neural network-based emulator of the monolayer higher-order approximation to the Stokes' equations as implemented with a compatible velocity-thickness finite element pair. This emulator fundamentally encodes rotation and reflection-invariance, and we couple it to a classic DG0 finite element method for ice transport ensuring that the hybrid model retains the essential property of near-exact mass conservation. We use as training data FEM-model-derived synthetic glaciers grown atop deglaciated topography from western North America - as a test case we then demonstrate that the hybrid NN-FEM model can accurately and efficiently simulate glacial cycles in novel mountainous regions.

Douglas Brinkerhoff, Jake Downs
University of Montana
douglas1.brinkerhoff@umontana.edu,
iostream333@gmail.com

MS217

Non-Intrusive Reduced-Order Modeling for Dynamical Systems with Spatially Localized Features

In this talk we present a non-intrusive reduced-order modeling framework for dynamical systems with spatially localized features characterized by slow singular value decay. The proposed approach builds upon two existing methodologies for reduced and full-order non-intrusive modeling, namely Operator Inference (OpInf) and sparse Full-Order Model (sFOM) inference. We decompose the domain into two complementary subdomains which exhibit fast and slow singular value decay, accordingly. The dynamics on the subdomain exhibiting slow singular value decay are learned using sFOM while the dynamics with intrinsically low dimensionality on the complementary subdomain are learned using OpInf. The resulting, coupled OpInf-sFOM formulation leverages the computational efficiency of OpInf and the high spatial resolution of sFOM and thus enables fast non-intrusive predictions for localized features with slow singular value decay, such as transport-dominated dynamics. Furthermore, we propose a novel, stability-promoting regularization technique with a closed-form solution based on the Gershgorin disk theorem for both sFOM and OpInf models and evaluate the efficiency of the coupled approach in terms of offline and online speedup. Finally, we demonstrate the capabilities of the

coupled OpInf-sFOM formulation for testcases such as the one-dimensional, viscous Burgers equation and a two-dimensional parametric model for the Pine Island Glacier ice thickness dynamics.

Leonidas K. Gkimisis
Max Planck Institute
Computational Methods in Systems and Control Theory
gkimisis@mpi-magdeburg.mpg.de

Marco Tezzele
Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
marco.tezzele@emory.edu

Marco Tezzele
Emory University
marco.tezzele@emory.edu

Nicole Aretz
The University of Texas at Austin
nicole.aretz@austin.utexas.edu

Thomas Richter
University of Magdeburg
Germany
thomas.richter@ovgu.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

Karen Willcox
The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences
kwillcox@utexas.edu

MS217

Assessing Machine Learning's Impact on Cryospheric Sciences

Recently, Machine Learning (ML) has emerged as a powerful tool to observe, model, and understand Earth's frozen regions. ML use cases for the cryosphere have surged in the past decade, ranging from learning sea ice dynamics directly from satellite data to accelerating numerical ice sheet models with machine learning emulators. While ML has served as a tool to accelerate and optimize scientific processes, recent discussions question whether ML can serve as a tool to drive scientific discovery for the cryosphere. In other fields, such as molecular biology and computational linguistics, ML has proven it can help solve previously unanswered scientific questions at scale; however, it remains unclear if ML can fulfill such hopes for the cryospheric sciences. Here, we identify key challenges and scientific questions across cryospheric subfields and analyze if and to what degree ML has aided in addressing those. Furthermore, we identify areas and tasks ML could impact and transform in the coming years. We invite ML practitioners and domain experts to join forces, ensuring that impactful cryospheric problems are solved with the appropriate methods.

Julia Kaltenborn
McGill University
julia.kaltenborn@mail.mcgill.ca

Kim Bente
School of Computer Science
The University of Sidney
kim.bente@sydney.edu.au

Andrew McDonald
University of Cambridge
British Antarctic Survey
arm99@cam.ac.uk

Chris Derksen
Environment and Climate Change Canada
chris.derksen@ec.gc.ca

David Rolnick
McGill University
Mila - Quebec AI Institute
drolnick@cs.mcgill.ca

MS217

A Data-Driven Coupled Neural Operator Approach to Model Ice-Sheet Dynamics

Traditionally, ice sheet dynamics have been modeled using classical discretization methods such as the Finite Element Method (FEM). FEM-based models tend to be computationally expensive and therefore not ideal in the context of uncertainty quantification where a large number of evaluations are required for accurate predictions. We propose machine-learning-based surrogates as a more efficient alternative to these models. We introduce a set of fully-coupled neural operators, where the first operator is a vanilla Deep Operator Network that maps the ice thickness field and basal friction field (one of the most uncertain parameters) to the ice speed; whereas the second operator is a flow map that evolves the dynamics of the ice sheet by mapping the ice velocity and thickness at a time to the ice thickness at a next step. In particular, two approaches are considered for the flow map: i) a vanilla architecture, where the operator is a conventional self-composed Deep Operator Network and ii) a SVD-based operator, where we precompute the operator basis functions using the data. The operators are first trained independently and then coupled together to conduct a probabilistic analysis of the ice sheet dynamics due to uncertainty in the basal friction field. Preliminary results demonstrate that the proposed framework is much faster than the FEM model while providing accurate predictions.

Dario Rodriguez
University of Illinois at Urbana-Champaign
darior2@illinois.edu

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

MS217

Identifying Energy Balance Drivers and Feedbacks of Greenland Ice Sheet Surface Melt Using Causal Inference

The mass loss from the Greenland Ice Sheet (GrIS) has accelerated during the past decades, with surface mass balance (SMB) decrease becoming the dominant contributor due to enhanced surface melt. SMB directly results from the complex interactions and feedbacks between the ice sheet and the atmosphere across different time scales.

The relevant processes could be nonlinear, and their exact causal dependencies cannot be detected in conventional correlative approaches. Causal inference can extract underlying physically plausible networks from these interactions and allow us to focus on fewer highly relevant dependencies, thus delivering better interpretability. In this study, we investigate causal links among the key processes that contribute to the summer surface melt of the GrIS with causal inference. First, the causal discovery method PCMCI+ is applied to the outputs of historical CESM2 simulations using a fully coupled configuration with a dynamic GrIS. Significant energy balance drivers and related feedbacks to surface melt will be identified with quantified strength and time lag. The resulting causal graph is evaluated by comparing to the graphs derived from high-resolution regional climate model outputs. Then we apply the same method to identify the causal dependencies in the outputs of CESM2 simulations under the preindustrial and a strong warming scenario (4CO2), by comparing which the potential changes of the melt regime under global warming are estimated.

Ziqi Yin, Aneesh Subramanian, Rajashree Datta
University of Colorado Boulder
ziqi.yin@colorado.edu, ansu6268@colorado.edu,
rajashree.datta@colorado.edu

Adam Herrington
National Center for Atmospheric Researc
aherring@ucar.edu

Danni Du
University of Colorado Boulder
danni.du@colorado.edu

Sahara Ali, Omar Faruque
University of Maryland, Baltimore County
sali9@umbc.edu, omarf1@umbc.edu

Jianwu Wang
Department of Information Systems
University of Maryland, Baltimore County
jianwu@umbc.edu

MS218

Modeling and Computing Two-Phase Dielectric Fluid Flows: A Consistent Phase Field Formulation and Efficient Numerical Algorithm

We develop a technique for modeling and simulating two-phase dielectric flows and their interactions with external electric fields in two and three dimensions. We first present a thermodynamically consistent and reduction-consistent phase field formulation for two-phase dielectric fluids, which honors the mass/momentum conservations and thermodynamic principles, and is endowed with the property that, when only one fluid component is present, the two-phase formulation will exactly reduce to that of the corresponding single-phase system. In particular, this model accommodates an equilibrium solution that is compatible with the physical requirement of zero velocity. This property provides a simpler method for computing the equilibrium state of two-phase dielectric systems. We further present an efficient algorithm, with a spectral-element discretization for 2D and combined Fourier-spectral/spectral-element discretization for 3D, for simulating this class of problems. This algorithm computes different dynamic variables successively in a decoupled fashion. It involves only coefficient matrices that are constant/time-independent in the linear systems upon discretization, even when the physical properties of the dielectric fluids are different. This property enables the use of Fourier-spectral discretization and FFT for 3D simulations. Several simulations are provided to demonstrate the performance and to compare with theoretical models and experimental data.

Jielin Yang, Suchuan Dong
Purdue University
yang1659@purdue.edu, sdong@purdue.edu

MS218

Dynamically Regularized Lagrange Multiplier Schemes with Energy Dissipation for the Incompressible Navier-Stokes Equations

In this paper, we present efficient numerical schemes based on the Lagrange multiplier approach for the Navier-Stokes equations. By introducing a dynamic equation (involving the kinetic energy, the Lagrange multiplier, and a regularization parameter), we form a new system which incorporates the energy evolution process but is still equivalent to the original equations. Such nonlinear system is then discretized in time based on the backward differentiation formulas, resulting in a dynamically regularized Lagrange multiplier (DRLM) method. First- and second-order DRLM schemes are derived and shown to be unconditionally energy stable with respect to the original variables. The proposed schemes require only the solutions of two linear Stokes systems and a scalar quadratic equation at each time step. Moreover, with the introduction of the regularization parameter, the Lagrange multiplier can be uniquely determined from the quadratic equation, even with large time step sizes, without affecting accuracy and stability of the numerical solutions. Fully discrete energy stability is also proved with the Marker-and-Cell (MAC) discretization in space. Various numerical experiments in two and three dimensions verify the convergence and energy dissipation as well as demonstrate the accuracy and robustness of the proposed DRLM schemes.

Kha Doan, Thi Thao Phuong Hoang
Auburn University
kcd0030@auburn.edu, tzh0059@auburn.edu

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Rihui Lan
Ocean University of China
lanrihui@ouc.edu.cn

MS218

A New Class of Higher-order Stiffly Stable Schemes with Application to the Navier-Stokes Equations

How to construct stable second- and higher-order fully decoupled schemes for the incompressible Navier-Stokes equations has been a long standing open problem. A main issue is that stability regions of usual multistep time discretization decrease as their order of accuracy increase, so they do not possess enough stability to control the higher-order explicit treatment of the pressure in a fully decoupled scheme. We shall construct a new class of IMEX schemes, by using Taylor expansion at $t_{n+\beta}$ (with $\beta \geq 1$ as a parameter) for updating the solution at t_{n+1}, whose stability region increases with β, thus allowing us to choose β

according to the stability and accuracy requirement. In particular, by choosing suitable β, we are able to construct higher-order unconditionally stable (in H^1 norm), fully decoupled consistent splitting schemes for the Navier-Stokes equations, and derive uniform optimal error estimates. We shall also present ample numerical results to show the computational advantages of these schemes for some nonlinear parabolic systems, including in particular Navier-Stokes equations. This talk is based on the joint work with Dr. Fukeng Huang.

Jie Shen
Purdue University
Department of Mathematics
shen7@purdue.edu

MS219

Model-form Error Correction using Universal Differential Equations for an Agent-Based Model of Infectious Disease

Agent-based models (ABMs) have emerged as a powerful tool for studying infectious diseases. By introducing heterogeneity and stochasticity through individual agent behaviors, ABMs allow for the emergence of complex patterns at the population level. However, scaling up and calibrating ABMs to provide further information and additional validation and verification remains challenging. On the other hand, ordinary differential equations (ODEs) models are often employed to model the average behavior of populations over time. ODE models adopt a deterministic approach, making them easier to calibrate. Due to the strengths and limitations of both modeling frameworks, developing ODE models as surrogates for ABMs is an active research area within scientific communities. In recent years, the emerging field of scientific machine learning (SciML) has aimed to integrate traditional mathematical modeling with advancements in machine learning to address these challenges. These advancements are laying the foundation for SciML to provide novel surrogates that decompose the global and local behaviors inherent to ABMs. This presentation demonstrates universal differential equations (UDEs) as an approach to bridge the gap between ODE models and ABM models. Using UDE models as surrogates for ABMs allows us to preserve the foundational ODE that represents global disease dynamics while coupling it with a neural network model to approximate functions for the local behaviors of the ABM.

Kyle Nguyen
Sandia National Laboratories
Livermore, CA
kcnguye@sandia.gov

Erin Acquesta, Ray Jaideep
Sandia National Laboratories
eacques@sandia.gov, jairay@somnet.sandia.gov

MS219

Scalable Algorithm for Strongly Connected Component Updates in Large Dynamic Graphs

Real-world networks, spanning diverse domains such as social interactions, biological systems, and autonomous agent collaborations, are inherently dynamic. As these networks evolve over time, their structural changes lead to changes in fundamental properties, including shortest paths, centrality measures, vertex coloring, and strongly connected components (SCCs). In the context of SCCs, the addition of a new edge between nodes in different SCCs can merge them into a single SCC, while the deletion of edges within an SCC can cause it to split into multiple smaller SCCs. As network sizes increase, the complexity of updating these properties grows significantly. To address this challenge, we design and implement a novel algorithm for efficiently updating SCCs in dynamic graphs. Our approach leverages meta-graph representation and a hub reference technique to streamline the SCC update process as the network evolves. We have developed implementations of our algorithm for both shared memory systems and GPU architectures, enhancing its applicability across various computational environments. Our work contributes to the field of network analysis by providing a scalable solution for maintaining up-to-date information on strongly connected components in large, dynamic networks.

Aashish Pandey
University of North Texas
aashishpandey@my.unt.edu

S.M. Shovan, Sajal K. Das
Missouri University of Science and Technology
sskg8@mst.edu, sdas@mst.edu

Boyana Norris
University of Oregon
norris@cs.uoregon.edu

Sanjukta Bhowmick
University of North Texas
sanjukta.bhowmick@unt.edu

MS219

The Tricks Required for Scientific Machine Learning to Work on Real Data

The combination of scientific models into deep learning structures, commonly referred to as scientific machine learning (SciML), has made great strides in the last few years in incorporating models such as ODEs and PDEs into deep learning through differentiable simulation. Such SciML methods have been gaining steam due to accelerating the development of high-fidelity models for improving industrial simulation and design. However, many of the methods from the machine learning world lack the robustness required for scaling to industrial tasks. What needs to change about AI in order to allow for methods which can guarantee accuracy and quantify uncertainty? In this talk we will go through the details of how one can enable robustness in building and training SciML models. Numerical robustness of algorithms for handling neural networks with stiff dynamics, continuous machine learning methods with certifiably globally-optimal training, alternative loss functions to mitigating local minima, integration of Bayesian estimation with model discovery, and tools for validating the correctness of surrogate models will be discussed to demonstrate a next generation of SciML methods for industrial use. Demonstrations of these methods in applications such as two-phase flow HVAC systems, modeling of sensors in Formula One cars, and lithium-ion battery packs will be used to showcase the improved robustness of these approaches over standard (scientific) machine learning.

Chris Rackauckas
Massachusetts Institute of Technology
Boston, MA

chris.rackauckas@juliahub.com

MS219
Network-Based Epidemic Control

In response to the rapid global transmission of infectious diseases, it is imperative to develop optimal strategies that control the spread of epidemics while minimizing the societal impact of interventions. To address this challenge, we propose two distinct frameworks. The first framework introduces a lockdown policy that strategically controls travel rates between population nodes within a network, aiming to minimize infections while ensuring the continuity of societal functions. We provide a convergence analysis to demonstrate the effectiveness of the solution approach. The second framework focuses on pandemic control through optimal isolation measures, designed to balance economic costs and the reduction of infection rates. By extending the traditional SIR model, we incorporate isolation strategies, transforming the problem of finding optimal isolation measures into a matrix balancing problem. This allows for efficient computation of isolation strategies, even in large networks, and ensures timely intervention for both asymptomatic and symptomatic cases. We validate the robustness of these frameworks by applying them to publicly available data on inter-county travel frequencies, specifically analyzing the spread of infections across the 14 counties of Massachusetts. Our results underscore the practical utility of these strategies in managing the spread of infectious diseases while mitigating the economic and social costs associated with such interventions.

Mahtab Talaei
Boston University
mtalaei@bu.edu

Apostolos Rikos
Hong Kong University of Science and Technology (Guangzhou)
apostolosr@hkust-gz.edu.cn

Alex Olshevsky, Laura White, Ioannis Paschalidis
Boston University
alexols@bu.edu, lfwhite@bu.edu, yannisp@bu.edu

MS219
Bayesian Decision-Making for Efficient Epidemiological Model Calibration

Computational epidemiological models are widely used for simulating and forecasting epidemic processes. In order for these models to effectively capture the dynamics and inherent characteristics of the epidemic processes, efficient model calibration to observed data is critical. In this talk, we focus on efficient calibration methods via Bayesian decision-making for the family of compartmental epidemiological models, explicitly considering the potential computational challenges across various scenarios. Specifically, we introduce model calibration methods based on a gray box Bayesian optimization (BO) scheme, which leverages the functional structure of the compartmental epidemiological model to improve calibration performance. Furthermore, we present model calibration methods via a decoupled decision-making strategy for BO, which takes advantage of the decomposable nature of the functional structure. We will present performance assessment results, comparing the calibration efficiencies of the various schemes and demonstrating how gray box variants of BO can further enhance calibration performance. We expect that the proposed Bayesian decision-making scheme for efficient model calibration can be extended to rapidly calibrate complex epidemiological models, such as agent-based models, which are computationally taxing.

Byung-Jun Yoon
Brookhaven National Laboratory, U.S.
bjyoon@tamu.edu

Xiaoning Qian
Texas A&M University
Brookhaven National Laboratory
xqian@tamu.edu

Puhua Niu
Texas A&M University
niupuhua.123@tamu.edu

MS220
A Transformer-Based Methodology to Correct 7-Day Forcecasts of STOFS-2D-Global: Q4 2024 Performance Evaluation

Predicting total water levels globally is a task replete with uncertainty. To address short- and medium-term forecast errors, one approach is to exercise machine learning capabilities that enhance the output of global circulation models. This is an enrichment strategy, one that can render a rapid enriched forecast with minimized model bias. Previously, we demonstrated this sort of methodology with the Temporal Fusion Transformer (TFT). We trained and evaluated the TFT on a 3-year hindcast of NOAAs ADCIRC-based STOFS-2D-Global and demonstrated excellent bias-minimization skill. Of course, with its reanalysis winds, this hindcast certainly does not mimic operational forecasting; STOFS-2D-Globals meteorological forecasts from GFS are subject to nonnegligible errors especially beyond the short-term period. To what extent the degrading meteorological forecast skill in the forecast horizon compromises TFT skill has yet to be quantified. Recently, we have begun exercising the TFT to minimize the bias of STOFS-2D-Global forecasts at NOAAs CO-OPS stations along US coastlines. For each midnight STOFS-2D-Global forecast cycle from September through December 2024, we evaluate TFTs performance. We characterize skill at different intervals within the horizon to address the degradation. Additionally, we compare performance across stations to ascertain if TFT performance is spatially correlated or contingent on other factors such as sensitivity of water levels to winds.

Albert Cerrone
University of Notre Dam
acerrone@nd.edu

Leendert Westerink
Dept. of Civil & Environmental Engineering & Earth Sciences
lwesterink@nd.edu

Joannes Westerink
Dept. of Civil & Environmental Engineering & Earth Sciences
University of Notre Dame
jjw@nd.edu

MS220
On the use of NASAs Cyclone Global Navigation

Satellite System (CYGNSS) wind data in storm surge modeling

Tropical cyclones (TCs) routinely cause flooding that results in loss of life and economic disruption. The primary driving force behind TC flooding arises from (shear) stresses at the ocean surface due to winds, which are estimated by a drag formulation that is a function of the surface-level (10-m) winds. In turn, these winds are typically estimated by simple parametric TC wind models, which provide the so-called gradient balanced winds (V_{gr}) that occur well above the surface-level winds (V_{10}), thereby requiring an adjustment based on an empirical surface wind reduction factor, i.e., $V_{10} = \mathrm{SWRF} \times V_{\mathrm{gr}}$. In practice, the SWRF is typically taken as a fixed constant value; however, previous studies have shown that this simple approach can lead to inaccurate estimates of TC winds and that the "ability of a gradient balance model to match the observed surface winds could be improved by using a SWRF that varies with radius [of the storm]" [1]. Therefore, using surface-level wind speed data retrieved from NASAs Cyclone Global Navigation Satellite System (CYGNSS) mission, a radial-dependent SWRF function is derived and implemented into the ADCIRC (Advanced Circulation) model, with the goal of obtaining improved results over the "standard" technique of using a fixed SWRF. [1] P.J. Vickery, et al, A Hurricane Boundary Layer and Wind Field Model for Use in Engineering Applications, J. Appl. Meteor. Climatol., 48, 381405, 2009.

Ethan Kubatko
Department of Civil, Environmental and Geodetic Engineering
The Ohio State University
kubatko.3@osu.edu

Aaron Sines, Mohammad Al-Khaldi, Joel Johnson
The Ohio State University
sines.49@osu.edu, al-khaldi.2@osu.edu, johnson.1374@osu.edu

MS220

Above Expectations: Sensitivity Analysis for High Impact, Low Probability Storm Surge Events

Computational storm surge models are critical to both emergency and long-term coastal flood risk analysis. These models are dependent on many physical and numerical parameters, which often are related to each other, have time-varying effects, and contain large uncertainties, producing sometimes unanticipated impacts on storm surges and their dynamics. While a number of studies explore how such parameters influence expected, i.e. mean, values of storm surges, it is important to understand the implications for extremes. In this work, we explore and quantify the impact of parameters on low probability, high impact storm surge levels using approaches developed from structural engineering. The methodology allows us to gain insight into the effects of uncertainties on extreme events, including those that have developed as a consequence of climate change.

Talea Mayo
Emory University
talea.mayo@emory.edu

Kevin Mackie
University of Central Florida
kevin.mackie@ucf.edu

MS220

Numerical Simulations of Sediment Transport Over Erodible Beds Using Extended Shallow Water Models

Sediment transport contributes significantly to the hydromorphodynamic changes of free-surface flows in shallow water environments. This is transport modeled based on the shallow water assumption without considering the variation of the water velocity along the vertical direction. However, this variation of the velocity is crucial for a better approximation of the near-bed velocity for the fluids. Vertical variations are accounted for using a moment model, which considers a polynomial expansion of the horizontal velocity. In this work, we present both a model and simulations considering both vertical variation of velocity and erosion-deposition effects for sediment transport, providing a computationally efficient framework for predicting sediment dynamics. The extended model consists of 4 parts: (1) the standard shallow water equations, (2) the so-called moment equations for the evolution of the basis coefficients, (3) an evolution equation for the sediment concentration, and (4) a transport equation for the bed. This allows for bedload and suspended load transport in one coupled model. The complex coupled model is given in explicit form and allows mathematical analysis. Using hyperbolic regularization, the computation of realistic eigenvalues and stability of the model can be proven. Several numerical test cases are also presented. Finally, we show the validation of the coupled model against a set of experimental data for the evolution of free surface and bed deformations.

Afroja Parvin
KU Leuven, Belgium and Peking University, China
afroja.parvin@kuleuven.be

Julian Koellermeier
University of Groningen
j.koellermeier@rug.nl

Giovanni Samaey
KU Leuven
giovanni.samaey@kuleuven.be

MS221

Mesh Adaptation for Ice Sheet Modeling

Modeling ice-sheet dynamics presents several challenges, one of which is accurately tracking the "grounding line" (GL) - the boundary where an ice sheet transitions from being grounded to floating in the ocean. The quality of the mesh near the GL significantly impacts the accuracy of the solution and our ability to make reliable predictions. As ocean temperatures rise near the GL, the GL retreats, necessitating adaptive meshing to maintain computational efficiency during multi-decadal simulations. In this work, we integrate the mesh adaptivity library Omega_h into the Finite Element implicit solver library Albany, which is used for simulating ice-sheet dynamics within the Energy Exascale Earth System Model (E3SM). We present preliminary application and performance results from a time-dependent problem on a 2D unstructured mesh running on NVIDIA GPUs on the Perlmutter cluster at NERSC.

Luca Bertagna
Sandia National Laboratories
lbertag@sandia.gov

Cameron W. Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS221

Exascale In Situ Machine Learning for Turbulent Flows

A stabilized finite element flow solver has been developed using PETSc and libCEED for application to scale resolving simulations of turbulent flows on peta and exascale computers. As the data generation rate of exascale computers far exceeds the IO capacity, the solver has also integrated in situ machine learning so that subgrid scale and other turbulence model closures can be learned at a very low tax to the ongoing simulation. In some cases of interest, the tax is near or even below the tax of writing the data to disk for offline training as is typically done. In this talk we will provide a brief overview of the solver, the added components necessary to extract training data in situ, the storage of that extracted data into a database, and the concurrent ML based model training and its interaction with that database. Efficient inference, whether it be part of the training or in subsequent model application, together with scalability of all aspects of the workflow will be discussed.

Riccardo Balin
Argonne National Labs
rbalin@anl.gov

James Wright, Aviral Prakash
University of Colorado Boulder
james.wrightiii@colorado.edu,
aviral.prakash@colorado.edu

John Evans
Institute for Computational Engineering and Sciences
The University of Texas at Austin
john.a.evans@colorado.edu

Jed Brown
University of Colorado Boulder
jed@jedbrown.org

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

MS221

Multiscale Global Impurity Transport in Fusion Devices

In fusion reactors, background plasma species escape the confinement of the magnetic field and reach a special region in the tokamak region called divertor, which is equipped to handle the bombardment of such particles. This sputters impurity particles from the divertor which may escape into the core of the plasma or strike back the divertor surface, causing reflection and further sputtering of new impurities. Simulation tools required to model such a physical process encounter a wide range of temporal and spatial scales. For example, such tools must carefully account for complex geometric features at the device scale (order of metres), while also being able to account for the sheath region (order of millimetres) near the plasma facing surfaces to model the particle-surface interaction. The impurity transport and surface interaction problem thus requires an integrated approach involving multiple high fidelity models which can interact with one another. In this talk we will discuss a global impurity transport GPU-accelerated code, GITRm, developed to be a part of such an integrated approach. The utility and effectiveness of GITRm in simulating impurity transport in realistic tokamak geometries will be demonstrated through multiple cases of interest including the ITER and DIII-D tokamaks. These simulations demonstrate GITRm's ability to capture localized behavior in the collector probes far from the divertor surfaces and effectively handle multi-species impurity transport.

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@rpi.edu

Dhyanjyoti Nath
Scientific Computation Research Center
Rensselaer Polytechnic Institute
nathd@rpi.edu

Timothy Younkin
Oak Ridge National Laboratory
younkintr@ornl.gov

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS221

Material Point Methods Implementable in Unstructured Meshes and Their Consistency and Accuracy Improvements

Material point method (MPM) uses Eulerian mesh and Lagrangian particles. The original MPM can be easily used in unstructured mesh as the finite element method. Because of the discontinuity of the gradients of the shape functions across cell boundaries, as particles move across them, numerical noises on nodal forces are generated, often leading to the failure of a calculation. To address this issue, the generalized interpolation material point (GIMP) and the convective particle domain interpolation (CPDI) methods are developed. In these methods a particle is not a point but a finite domain, which can occupy multiple cells. Nonlocal operations are required to map between the particle and nodal quantities making the method difficult to implement in unstructured meshes. To address the cell-crossing noise, the dual domain material point (DDMP) method maps part of the stress on particles to nodes. The nodal force is then calculated with two contributions, from the nodal stress and from the remaining stress on particles. As a particle approaches a cell boundary, the remaining stress is reduced to zero to eliminate the discontinuity on the nodal force. In DDMP, particles remain as geometric points and mappings between particles and nodes are local. This method can be implemented with unstructured meshes. Recently, DDMP is further improved with the lo-

cal stress difference (LSD) scheme, significantly increasing numerical accuracy and consistency while reducing mesh dependency.

Duan Z. Zhang
Los Alamos National Laboratory
dzhang@lanl.gov

Kyle Perez
Los Alamos National Laboratory
Los Alamos National Laboratory
kperez@lanl.gov

Jiajia Waters
Los Alamos National Lab
jwaters@lanl.gov

Paul Barclay
Los Alamos National Laboratory
barclay@lanl.gov

MS222

Tensor Train Space Time Spectral Collocation Method for Solving 3D Maxwells Equations

We introduce a tensor train space-time spectral collocation method for solving three dimensional Maxwells equations. We discretize both time and space using spectral collocation points on a staggered grid, which maintains the divergence free property of the magnetic field. Numerical experiments show that the tensor train format of the method achieves exponential convergence, while improving the full tensor solver performance by order of magnitude in term of CPU time and memory usage.

Dibyendu Adak, Duc Truong, Korobkin Oleg, Mario Ivan I. Ortega, Kim Rasmussen
Los Alamos National Laboratory
dadak@lanl.gov, dptruong@lanl.gov, korobkin@lanl.gov, mortega@lanl.gov, kor@lanl.gov

Boian Alexandrov
Los Alamos National Laboratory
T-1: Physics and Chemistry of Materials
boian@lanl.gov

MS222

A New Dynamic Low Rank Method for Thermal Radiative Transfer

The Thermal Radiative Transfer (TRT) equations — which play a critical role in modeling hohlraums for Interial Confinement Fusion — model the energy exchange between electromagnetic radiation and a background material. They comprise a set of high-dimensional, stiff, nonlinear PDEs that couple the distribution function for the energy density of the photon field — which describes the radiation energy density at a given point in space, passing through a given direction, and with a given frequency — with the material energy. We present a new Dynamic Low Rank (DLR) method for TRT, which offers an attractive tradeoff between computational cost and accuracy. The basic idea of DLR methods is to represent the photon distribution function as an SVD-like product of dynamically evolving, quasi-optimal spatial and angular basis functions. Although there has been promising recent work on leveraging DLR for linear transport problems, we have developed a version of DLR that is tailored for TRT. In particular, the method is implicit-in-time (without the typical DLR time-splitting), preserves the diffusion limit, is formulated with positive-definite operators, and allows efficient AMG solvers. We demonstrate the efficiency and robustness of our method on a standard challenging benchmark problem.

Terry Haut
Lawrence Livermore National Laboratory
haut3@llnl.gov

MS222

Leveraging Tensor Networks for Solving PDEs

Tensor Networks have emerged as a powerful tool for tensor approximation, enabling the solution of large-scale linear and multilinear algebra problems that are intractable with classical methods. This talk will introduce the application of tensor networks to accelerate numerical partial differential equation solvers. We will focus on our recent results in two areas: (1) solving time independent neutron transport equations and (2) implementing space-time spectral collocation methods for linear and nonlinear equations. Numerical examples will demonstrate that tensor network-based methods achieve significant compression and speedup compared to full tensor approaches, potentially revolutionizing computational efficiency in scientific computing.

Duc Truong, Dibyendu Adak, Mario Ivan I. Ortega, Kim Rasmussen
Los Alamos National Laboratory
dptruong@lanl.gov, dadak@lanl.gov, mortega@lanl.gov, kor@lanl.gov

Boian Alexandrov
Los Alamos National Laboratory
T-1: Physics and Chemistry of Materials
boian@lanl.gov

MS222

Low-Rank Tucker Tensor Methods for Stochastic Dynamic Optimization

The numerical solution of dynamic optimization problems is often limited by the memory required to store the state trajectory, which is used to evaluate the objective function and its derivatives. Recently, [R. Muthukumar et al., SIAM Journal on Optimization 31(2), pp. 12421275 (2021)] introduced a trust-region method for dynamic optimization that employs randomized sketching to compress the state trajectory, resulting in inexact derivative computations. We extend this approach to stochastic optimization, where we use a low-rank Tucker sketch to compress both space and sample dimensions. By adaptively learning the sketch rank, the trust-region algorithm achieves rigorous convergence guarantees. Due to the randomness introduced by the sketch, the traditional secant update formulae can produce poor Hessian approximations. In particular, the difference of two gradients, computed from two different sketches, may be inconsistent. To overcome this, we employ a sketched approximation of the Hessian application, in lieu of computing the gradient difference. We numerically demonstrate the improved stability of this approach on an example from PDE-constrained optimization.

Radoslav G. Vuchkov
Sandia National Laboratories
rgvuchk@sandia.gov

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS223

When Reduced Precision is not an Option

Prevalent scientific applications are largely designed assuming availability of double-precision arithmetic on demand with the use of lower precision arithmetic motivated by performance improvement. With the arrival of hardware that is designed for lower precision, and expectation of further decline in the availability of higher precision hardware, we are faced with the possibility of having to rely on software based solutions. In this presentation I will describe scenarios where low and/or mixed precision arithmetic is not feasible, and what it might mean for the design of scientific software architecture.

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

MS223

Probabilistic Rounding Uncertainty Analysis for Floating-Point Statistical Models

Advancements in computer hardware now allow low- and mixed-precision arithmetic to improve efficiency, especially on new architectures. It is thus critical that the rounding uncertainty be rigorously quantified alongside traditional sources of uncertainty including those from observations, sampling, and numerical discretization. Traditional deterministic rounding uncertainty analysis (DBEA) assumes that the absolute rounding errors equal the unit roundoff u, considering the worst-case scenario. This work presents a novel probabilistic rounding uncertainty analysis called VIBEA. By treating rounding errors as i.i.d. random variables and leveraging concentration inequalities, VIBEA provides high-confidence estimates for rounding uncertainty using higher-order rounding error statistics. The presented framework is valid for all problem sizes n, unlike DBEA, which requires $nu < 1$. Further, it can account for the potential cancellation of rounding errors, resulting in rounding uncertainty estimates that grow slowly with n. We demonstrate that quantifying rounding uncertainty alongside traditional sources allows for a more efficient allocation of computational resources, balancing efficiency with accuracy. This study takes a step towards a comprehensive mixed-precision approach to enhance model reliability and optimize resource allocation in predictive modeling. The talk will conclude with a vision for end-to-end, formally verified numerics for scientific computing.

Sahil Bhola
University of Michigan, Ann Arbor
sbhola@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS223

Efficient Mixed-Precision Memory-Bound BLAS based on Memory Accessors with Applications to Sparse Direct Solvers

Mixed-precision algorithms can take advantage of the lower precision of operands to enhance performance. In various contexts, it is beneficial to decouple the storage precision from the compute precision: the data is stored and accessed in low precision, but the computations are kept in high precision. This "memory accessor' approach benefits from reduced data accesses and improved accuracy, and can simplify the programming of mixed precision software packages. In this work, we develop such a memory accessor and investigate how it can accelerate sparse linear solvers. In particular, we assess its impact on custom floating-point datatypes unsupported by hardware and on structures such as Block Low-Rank formats. When considering BLAS-2 memory-bound operations like `trsv` we observe that the storage cost adequately matches the performance of the operation, in multiple parallel settings, provided that the conversion from storage to compute precision is efficient. For custom datatypes, we take advantage of the recent AVX512-VBMI2 instruction set to reach the required efficiency. We will present preliminary performance experiments using the sparse solver MUMPS, both as a direct method and as a preconditioner for Krylov methods.

Patrick R. Amestoy
Mumps Technologies
patrick.amestoy@mumps-tech.com

Antoine Jego
LIP6, Unité Mixte de Recherche de Sorbonne Université
antoine.jego@lip6.fr

Jean-Yves L'Excellent
Mumps Technologies
jean-yves.l.excellent@mumps-tech.com

Théo Mary
CNRS
Sorbonne Université
theo.mary@lip6.fr

Grégoire Pichon
Université Claude Bernard
Inria
gregoire.pichon@inria.fr

MS223

Solving Big Problems with Little Numbers

The future of simulations lies in leveraging hardware features designed for the AI market, particularly in low-precision computations. Modern NVIDIA GPUs exemplify this trend, offering significant performance gains through low-precision computations, resulting in reduced elapsed time, smaller memory footprints, and energy savings. We harness these capabilities to develop fast mixed-precision linear algebra algorithms. Our adaptive precision conversion strategy dynamically adjusts computation accuracy, maintaining high precision only where necessary within the matrix operator, while still meeting application-specific precision requirements. This approach revolutionizes computational efficiency for geospatial statisticians, bioinformaticians, and geophysicists, having significant implications for environmental computational statistics, genome-wide association studies in computational biology, and seismic imaging for CO2 sequestration.

Hatem Ltaief

King Abdullah University of Science & Technology (KAUST)
hatem.ltaief@kaust.edu.sa

MS223

Solver and Benchmarking Perspectives on Mixed Precision

Both floating-point and fixed-point representations are now part of the computational science toolbox. These include limited precision, quantized integers, and modular precision ecosystems, and the mix thereof such as graded and block formats. New solvers are adapting to this new landscape of representation formats to maintain the accuracy that can be relied upon while delivering some of the performance gains that the new hardware promises. The recent advances of mixed-precision solvers will be presented to highlight both their benefits and design constraints.

Piotr Luszczek
Massachusetts Institute of Technology
luszczek@icl.utk.edu

MS224

Automatic Second Quantization Algebra with Differentiation

The second quantization is a widely used method for many-body theories such as electronic structure theories and circuit QED. This method compactly parameterizes away the spin statistic constraints, but the translation from a second quantized model to a numerical implementation remains challenging even for experts. This talk presents our recent work on modeling and automating this process with primitive concepts from programming language theories. We then discuss how the resulting model can be symbolically differentiated and compiler optimized to yield viable code for practical theories such as couple cluster.

Kangbo Li
Cornell University
kl935@cornell.edu

MS224

PDE Perspective on Many-body Problems in Quantum Optics

Quantum optics is the quantum theory of the interaction of light and matter. In this talk, I will describe a real-space formulation of quantum electrodynamics. The goal is to understand the propagation of nonclassical states of light in systems consisting of many atoms. In this setting, there is a close relation to kinetic equations for nonlocal PDEs with random coefficients.

John C. Schotland
Yale University
john.schotland@yale.edu

MS224

Re-anchoring Quantum Monte Carlo with Tensor-Train Sketching

We propose a novel algorithm for calculating the ground-state energy of quantum many-body systems by combining auxiliary-field quantum Monte Carlo (AFQMC) with tensor-train sketching. In AFQMC, having a good trial wavefunction to guide the random walk is crucial for avoiding sign problems. Typically, this trial wavefunction is fixed throughout the simulation. Our proposed method iterates between determining a new trial wavefunction in the form of a tensor train, derived from the current walkers, and using this updated trial wavefunction to anchor the next phase of AFQMC. Numerical results demonstrate that our algorithm is highly accurate for large spin systems, achieving a relative error of 10^{-5} in estimating ground-state energies. Additionally, the overlap between our estimated trial wavefunction and the ground-state wavefunction achieves a high-fidelity. We provide a convergence proof, highlighting how an effective trial wavefunction can reduce the variance in the AFQMC energy estimate.

Ziang Yu
University of Chicago
ziangyu@uchicago.edu

MS225

Strong Convergence of Path Sensitivities

It is well known that the Euler-Maruyama discretisation of an autonomous SDE using a uniform timestep h has a strong convergence error which is $O(h^{1/2})$ when the drift and diffusion are both globally Lipschitz. In this presentation we prove that the same is true for the approximation of the path sensitivity to changes in a parameter affecting the drift and diffusion, assuming the appropriate number of derivatives exist and are bounded. This appears to fill a gap in the existing stochastic numerical analysis literature and is particularly important when considering the use of Multilevel Monte Carlo methods for estimating option sensitivities in mathematical finance.

Mike Giles
University of Oxford
mike.giles@maths.ox.ac.uk

MS225

Sensitivity Calculation for Monte Carlo Particle Simulations in Nuclear Fusion Reactors

Sensitivity-based optimization and uncertainty quantification tools are increasingly adopted in the nuclear fusion community for the design of future reactors and the validation of existing models, see, e.g., Refs. [1,2]. Most of the previous work in the nuclear fusion community focuses on simulations with deterministic continuum models. However, this continuum approximation is often violated for the neutral particles (atoms and molecules). Hence, a (partially) kinetic treatment is required for the neutrals, where individual particles are followed with a Monte Carlo (MC) particle-tracing code. The statistical noise from the MC part makes sensitivity calculation extremely challenging. The correlation between the particle trajectories of the original and perturbed simulation is easily lost for Finite Difference (FD) calculations [3]. Ref. [4] shows that Algorithmic Differentiation (AD) maintains correlation and leads to a statistical error reduction of the sensitivities with up to a factor 10^5 compared to FD for low-collisional conditions. However, there are issues with diverging sensitivities originating from long-lived particles in high-collisional conditions. We demonstrate how the AD error is affected by MC estimator choice. [1] M. Baelmans et al., Nuclear Fusion 57 (2017)
[2] S. Carli et al., Journal of Comp. Phys. 491 (2023)
[3] W. Dekeyser et al., Contrib. Plasma Phys. 58 (2018)

[4] N. Horsten et al., Contrib. Plasma Phys. (2024)

Niels Horsten, Stefano Carli, Wouter Dekeyser
KU Leuven
niels.horsten@kuleuven.be, stefano.carli@kuleuven.be, wouter.dekeyser@kuleuven.be

MS225

Using the Scalable Parallel Random Number Generators (sprng) Library for Reproducibility and Sensitivity Analysis in Monte Carlo Computations

The Scalable Parallel Random Number Generators (SPRNG) library is a widely used tool for stochastic simulation. It was designed to permit rapid and completely reproducible computations on a wide variety of parallel hardware. Here we focus on the ways that reproducibility is created in SPRNG. We first look at the use of parameterized random number generators (RNGs) as the basis for creating a reproducible set of parallel RNGs organized with a single seed to control the streams and numbers to be used in a simulation. Then we consider that SPRNG was also designed to permit the splitting of sequences in a reproducible fashion. We motivate the need for splitting RNGs and then show how a global design feature of SPRNG permits total reproducibility even in the presence of splitting. We finish with a status report on SPRNG and comment on the ongoing work to continue SPRNG's development and deployment.

Michael Mascagni
Computer Science/Computational Science
Florida State University, Tallahassee
mascagni@fsu.edu

MS225

Differentiable Monte Carlo Rendering and PDE Solvers

Physics-based Monte Carlo rendering algorithms generate photorealistic images by simulating the flow of light through a detailed mathematical representation of a virtual scene. In contrast, physics-based differentiable rendering algorithms focus on computing derivatives of images with respect to arbitrary scene parameters such as camera pose, object geometry, and material properties. Recently, Monte Carlo methods have also been applied to solve elliptical PDEs such as the Poisson equation. Compared with conventional finite-element and boundary-element methods, these techniques are "grid-free' and, thus, capable of handling extremely complex domains. On the other hand, differential variants of these PDE solvers estimate derivatives of PDE solutions with respect to arbitrary parameters including boundary conditions and domain shapes. Recent advances in differentiable Monte Carlo rendering and PDE solvers have enabled solving many challenging inverse problems—the search of scene/PDE parameters optimizing user-specified objective functions—using gradient-based methods such as stochastic gradient descent (SGD). In this talk, we present some of our recent works on these topics and discuss remaining challenges and open problems.

Shuang Zhao
University of California, Irvine
shz@ics.uci.edu

MS226

Structured Random Sketching for Tensor Decompositions

In this talk I will address randomized methods for computing tensor decompositions including Tucker and Tensor Train. The kernel computation is computing a sketch of an unfolding/matricization of the tensor, and we impose structure on the random matrix in order to exploit structure in the unfolding and reduce computational cost. I will discuss theoretical results on the accuracy of these approaches, their accuracy in practice, and the performance improvement they achieve over deterministic methods.

Grey Ballard
Wake Forest University
ballard@wfu.edu

MS226

Parametric Kernel Low-Rank Approximations Using Tensor Train Decomposition

Kernel matrices arise from kernel functions corresponding with integral equations, Gaussian processes, etc. However, kernel matrices arising from practical applications are often dense, large, require expensive kernel evaluations to compute, and depend on specific hyper-parameters; thus, there is a need for methods to compute and store low-rank approximations of parametric kernel matrices efficiently. We propose a method divided into two phases, an offline phase and an online phase, where the offline phase dominates the computational cost, and the online phase is relatively inexpensive. During the offline phase, we apply multi-variate polynomial approximation over a tensorized grid of Chebyshev nodes on the parametric kernel function, and then apply the Tensor Train (TT) decomposition on the tensorized grid of Chebyshev nodes, efficiently computed via TT cross approximation. During the online phase, we instantiate a kernel matrix for a particular parameter and compute and store its low-rank approximation using only tensor contractions and matrix multiplication. Our method has linear complexity in terms of the size of the parametric kernel matrix, and its utility and efficiency are demonstrated by applying it to various kernels arising from different application areas, varying spatial configurations of the source and target points, and varying properties of the kernel matrix (symmetry and positive semidefinite).

Abraham W. Khan
North Carolina State University, Mathematics Department
awkhan3@ncsu.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS226

Convergence Implications of Asynchronous Generalized Canonical Polyadic Tensor Decomposition

Recent advances for computing canonical polyadic (CP) tensor decompositions have leveraged randomized algorithms, in the form of generalized CP (GCP) with Adam optimization, in tandem with parallel and distributed codes to address the challenge of ever-growing datasets. Devel-

opments in related fields, such as federated learning, propose asynchronous algorithmic solutions to technical issues present in distributed data problems, e.g., deadlocks and underutilization. Lewis and Phipps demonstrated how an asynchronous SGD algorithm can achieve strong scaling but the impact on solution time and accuracy remains unclear. In this work we study the algorithm effectiveness of asynchronous SGD applied to distributed GCP vis-a-vis computational cost and convergence on synthetic and publicly-available real-world datasets of varying sizes, dimensions, and sparsity patterns using several loss functions, including Gaussian and Poisson loss.

Jeremy Myers
Sandia National Laboratories
jermyer@sandia.gov

Eric Phipps
Sandia National Laboratories
Center for Computing Research
etphipp@sandia.gov

MS227

Deep Learning Approach for Weather Prediction and Climate

Currently, weather forecasting is undergoing a paradigm shift moving from the physics-based traditional numerical weather prediction (NWP) to machine learning (ML) based models. Advancing in numerical modeling and data assimilation have made slow but steady progress over the last 4 decades, leading to a quiet revolution in weather forecasting. Recently, advances in machine learning, availability of high-quality data, and advances in hardware (e.g. GPUs/TPUs) have set the stage for deep learning to tackle problems for weather and climate. In the last 2 years, several deep learning-based models for weather forecasting have been demonstrated with skill approaching or exceeding the best available NWP weather forecasts. These models include Graphcast, Fourcastnet, FuXi, Pangu-weather, ClimaX, Fengwu, and Stormer, each with vastly different training methods, machine learning architectures, and variables predicted. Here, we present 2 state-of-the-art deep-learning models for weather and climate. 1.) Stormer, a simple transformer-based model that achieves state-of-the-art performance on medium-range weather forecasting with minimal changes to the standard transformer backbone. 2.) LUCIE a neural operator-based atmospheric emulator that remains stable during autoregressive inference for 100 years with little climate drift.

Troy Arcomano
Texas A&M University
troyarcomano@tamu.edu

MS227

Physics-Informed Graph Neural Networks for Extreme Flood Modeling

Recently, the growing risk of extreme precipitation has highlighted the need for high resolution, accurate flood maps. Traditional numerical methods for flood modeling face challenges such as intense computational burdens, particularly when simulating large domains at high resolution. These shortcomings create an opportunity for employing novel solutions that leverage artificial intelligence. This project employs physics-informed neural operators via the Shallow Water Equation Graph Neural Network (SWE-GNN), a deep learning model that can be used for two-dimensional hydrodynamic simulation over unseen topographies. Randomly sampled topographic domains surrounding hydrologic flowlines in Texas are used to generate a training dataset of numerically driven simulations. Extreme precipitation events are used as forcing inputs to capture regional extreme flood events for training. Simulations are then encoded as mesh-based graphs with nodes containing hydraulic information at each timestep including water discharge, surface height, and precipitation. A SWE-GNN is trained on the simulation data with neighbor-sampled wetted subgraphs and evaluated on a watershed withheld from the training dataset near Houston, TX. The SWE-GNN demonstrates promising capabilities for advancing hydrological systems predictability by reducing computation time by orders of magnitude and enabling any hydrodynamic model to use GPU via surrogate modeling.

Jeremy Feinstein, Ziang He, Sofia Sahagun, Eugene Yan
Argonne National Laboratory
jfeinstein@anl.gov, hez@anl.gov, ssahaguncovarrubias@anl.gov, eyan@anl.gov

MS227

Multimodal Generative Pretraining for Astrophysics

Deep Learning has seen a recent shift in paradigm, from training specialized models on dedicated datasets, so-called Foundation Models, trained in a self-supervised manner on vast amounts of data and then adapted to solve specific tasks with state-of-the-art performance. This new paradigm has been exceptionally successful not only for large language models (LLMs) but in other domains such as vision models. However applications of this new approach in astrophysics are still very scarce, for reasons ranging from new architectures to the (surprising) lack of availability of suitable large scale datasets. In this talk, I will discuss our recent work on deploying such a Foundation Model approach in the context of astronomical data. Specifically, I will present our efforts in compiling the first large scale ML-ready multimodal dataset in the field of astrophysics, covering a wide variety of scientific data (e.g. images, optical spectra, time series), and representing over 70TB of observational data from many different instrument. Based on this dataset, I will present our efforts to build a large scale multimodal generative model, capable of solving multiple downstream tasks.

Francois Lanusse
University of California, Berkeley
U.S.
flanusse@berkeley.edu

MS227

Neural Conditional Simulation for Complex Spatial Processes

In spatial statistics, classical statistical inference such as parameter estimation or spatial interpolation is computationally intensive or intractable for many spatial processes. This intractability often comes from the joint likelihood or the conditional distribution which are rarely available in closed form. A key objective in spatial statistics is to simulate from the conditional distributionthe distribution of the spatial process at unobserved locations given the observed locationsto enable spatial interpolation. In this talk, we propose using neural diffusion models for conditional simulation of complex spatial processes. Using a masking approach, we train a score-based diffusion model

within a stochastic differential equation (SDE) framework to learn the conditional reverse process–a process which reverse-diffuses Gaussian noise into samples from conditional distributions. The masking approach involves modifying the diffusion model so that the partially observed field and mask indicating observed and unobserved locations are inputs to the neural network approximating the score. As a result, the diffusion model only requires unconditional samples from the spatial process during training and is amortized with respect to the mask, provided the mask pattern is similar to those used during training. Finally, we discuss methods for validating that the generated samples do indeed come from accurate approximations of the true conditional distributions.

Julia Walchessen, Mikael Kuusela
Carnegie Mellon University
jwalches@andrew.cmu.edu, mkuusela@andrew.cmu.edu

MS228

Hybrid Neural and Differential Dynamical Models

We present an approach for learning subgrid-scale model effects in the simulation of partial differential equations (PDEs) using the method of lines, focusing on their representation in the PDE formulation through neural ordinary differential equations (NODEs). Our method leverages data assimilation, where fine-scale simulation data is utilized, with hybrid modeling that blends traditional PDE solvers and machine learning directly at the equation level. This approach addresses the computational challenges of fine temporal and spatial grid scales, enhancing accuracy and efficiency. We capture subgrid-scale dynamics and approximate coupling operators by employing NODEs in conjunction with partial system knowledge. The resulting method improves the efficiency of low-order solvers while effectively parameterizing subgrid scales. We demonstrate the efficacy of this approach through numerical experiments on the two-scale Lorenz 96 equation, the convection-diffusion equation, and the Navier-Stokes equations.

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

Shinhoo Kang
Department of Aerospace Engineering
The University of Texas at Austin
shinkang@korea.ac.kr

MS228

Operationalization of AI Models with Quantified Uncertainties Monitoring Out of Distribution Shifts

Out-of-Distribution (OOD) detection is crucial for maintaining the reliability and effectiveness of deployed AI models in safety-critical applications, as they adapt to real-world data that deviates from their training distribution. This deviation, known as OOD drift, can significantly degrade model performance. Real-world environments produce varied and unpredictable OOD data, requiring advanced detection mechanisms. For classification tasks, the network confidence score has been widely used as an indicator of epistemic uncertainty, and has shown some success in OOD detection with specific model architectures and datasets. However, it is well-known that neural networks can produce arbitrarily high confidence for inputs far away from the training data, making such uncertainty metrics unreliable indicators of OOD inputs. Although OOD detection has been extensively studied, challenges remain in developing efficient algorithms for large-scale AI models suitable for online deployment. This talk will focus on a fast and scalable topologically-inspired approach of quantifying AI model uncertainties, with applications in large-scale image classification and language modeling tasks. By extracting features from intermediate layers of an AI model, our approach naturally defines an uncertainty metric which can be generalizable to operationalizing different types of AI models in practice.

Sudeepta Mondal
RTX
sudeepta.mondal2@rtx.com

Ganesh Sundaramoorthi
RTX Technology Research Center
ganesh.sundaramoorthi@rtx.com

MS228

Analysis of Multiscale Sampling Markov Chain Monte Carlo Methods

In Uncertainty Quantification of Subsurface Flows, a fundamental problem is how to obtain the permeability field of a porous medium given some measurements of the flow. The approach of Bayesian Analysis translates it into the sampling of a complicated posterior probability distribution, thus requiring Markov Chain Monte Carlo (MCMC) methods. Since each iteration of the Markov Chain must solve an expensive partial differential equation, the simulation of large, realistic problems is still a very challenging task. Multiscale Sampling methods are a new derivative-free MCMC approach that takes inspiration from Domain Decomposition methods: Sampling is performed locally in disjoint subdomains, and different techniques are possible to couple those local fields into a global one. They have shown promising experimental results in the literature, with a significant improvement in convergence. However, their theoretical properties have not yet been studied. In this talk we provide new results that show that the Multiscale Sampling methods are well-defined, in the sense that their Markov Chain has a limiting distribution. Moreover, by carefully keeping track of each step of the construction of those algorithms, we are able to describe their corresponding instrumental distribution. It is hoped that this detailed analysis can be used as a framework to study further improvements on this class of MCMC methods.

Lucas Seiffert, Felipe Pereira
The University of Texas at Dallas
lxs190035@utdallas.edu, luisfelipe.pereira@utdallas.edu

MS228

Deep and Reinforcement Learning for Optimal Decision Making

While reinforcement learning (RL) has shown promising performance, its sample complexity continues to be a substantial hurdle, restricting its broader application across a variety of domains. Imitation learning (IL) utilizes oracles to improve sample efficiency, yet it is often constrained by the quality of the oracles deployed. which actively interleaves between IL and RL based on an online estimate of their performance. Robust policy improvement (RPI) draws on the strengths of IL, using oracle queries to fa-

cilitate explorationan aspect that is notably challenging in sparse-reward RL particularly during the early stages of learning. As learning unfolds, RPI gradually transitions to RL, effectively treating the learned policy as an improved oracle. This algorithm is capable of learning from and improving upon a diverse set of black-box oracles. In this talk, I will present our latest developments and applications of RPI algorithmic approaches.

Matthew Walter
Toyota Technological Institute at Chicago
mwalter@ttic.eu

Xuefeng Liu
Department of Computer Science
University of Chicago
xuefeng@uchicago.edu

MS229

Variational Autoencoders in Inverse Problems and Optimal Experiment Design

Stochastic inverse problems and optimal experiment design (OED) are crucial in science and engineering, relying on prior models (priors). Traditionally, priors are simplified to enable analytical solutions, but these limit model applicability, especially in high-dimensional data scenarios. For example, inverse problems with many unknowns are often underdetermined due to physical limits in sensors, requiring effective priors to address ill-posedness. Recent advances in generative machine learning models, such as GANs, VAEs, and diffusion models, have enabled the construction of priors that leverage legacy data and bypass common assumptions. These models provide a more flexible way to represent prior knowledge, improving the accuracy and consistency of solutions to inverse problems and the effectiveness of experimental designs. This work addresses optimal sensor placement in high-dimensional spaces, a key challenge in engineering applications with limited sensor budgets. We propose using VAEs to formulate an OED problem in the latent space of the VAE, providing a more efficient solution path. Our approach integrates Bayesian and frequentist perspectives, showing that under certain conditions, these reduce to the same optimization problem. We demonstrate our method through optimal placement of sensors in prototypical examples across different physics. Our approach offers a robust and scalable solution, opening new possibilities for applying machine learning models in OED.

Wilkins Aquino
Duke University
wilkins.aquino@duke.edu

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

Chandler Smith, Timothy Walsh
Sandia National Laboratories
chasmit@sandia.gov, tfwalsh@sandia.gov

MS229

Recent Advances in Weak Form Data-Driven Modeling in Mechanics

Recent advances in data-driven modeling approaches have proven highly successful in a wide range of fields in science and engineering. In particular, learning governing equations via mimizing an equation error criteria, offers a powerful and explainable scientific machine-learning framework. In this talk, I will present our weak form approach and briefly discuss how it addresses several ubiquitous challenges within conventional model development, discretization, parameter inference, and model refinement. In particular, I will describe our equation learning (WSINDy) and parameter estimation (WENDy) algorithms as well as how they can be used for different Eulerian and Lagrangian frameworks in continuum mechanics. Our approach has exhibited surprising performance, accuracy, and robustness properties. I will demonstrate these performance properties via applications to several benchmark continuum mechanics and other engineering problems in ordinary, partial, and stochastic differential equations as well as coarse-graining and reduced order modeling.

David M. Bortz
University of Colorado
Department of Applied Mathematics
dmbortz@colorado.edu

Daniel Messenger
Department of Applied Mathematics
University of Colorado - Boulder
daniel.messenger@colorado.edu

April Tran
University of Colorado at Boulder
chi.tran@colorado.edu

MS229

Equivariant Graph Convolutional Neural Networks for the Representation of Homogenized Anisotropic Microstructural Mechanical Response

Composite materials with different microstructural material symmetries are common in engineering applications where grain structure, alloying and particle/fiber packing are optimized via controlled manufacturing. In fact these microstructural tunings can be done throughout a part to achieve functional gradation and optimization at a structural level. To predict the performance of particular microstructural configuration and thereby overall performance, constitutive models of materials with microstructure are needed. In this work we provide neural network architectures that provide effective homogenization models of materials with anisotropic components. These models satisfy equivariance and material symmetry principles inherently through a combination of equivariant and tensor basis operations. We demonstrate them on datasets of stochastic volume elements with different textures and phases where the material undergoes elastic and plastic deformation, and show that the these network architectures provide significant performance improvements.

Ravi Patel
Sandia National Laboratories
rgpatel@sandia.gov

MS229

Micropolar Deep Material Network

This work builds on the recent developments of the deep material network and extends the paradigm to micropolar materials in order to efficiently model size effects in composite solids at the continuum level. Micropolar continua naturally incorporate the size effect using local equations,

at the cost of extra degrees of freedom. This cost manifests as an actual computational cost when the equations are numerically solved, and it is quite steep when simulations are scaled to practical problems. Deep material network is presented as a solution to this issue because simulation times are decreased by 100 times in many cases, and retraining is not an issue, as the deep material network is able to extrapolate constitutive behaviors it was not trained on.

Fatemeh Pourahmadian
University of Colorado Boulder
fatemeh.pourahmadian@colorado.edu

Noah Francis
University of Colorado, Boulder
noah.francis@colorado.edu

Dongil Shin
Pohang University of Science and Technology
dongilshin@postech.ac.kr

Ricardo Lebensohn
Los Alamos National Laboratory
lebenso@lanl.gov

Remi Dingreville
Sandia National Laboratories
rdingre@sandia.gov

MS230

Low rank methods for linear and nonlinear PDE

In this talk, we present our recent work on low-rank methods for linear and nonlinear PDE. These techniques are particularly effective in reducing computational complexity while maintaining the essential features of the original solution. Focus will be on techniques for preserving complete positivity and trace in open quantum systems, and on a new CrossDEIM approximation that allow us to efficiently handle nonlinear problems. We demonstrate the efficacy of our approach through various examples, including the Bratu problem, the Allen-Cahn equation, and the Lindblad equation. Our results highlight the potential of low-rank, structure-preserving methods in efficiently solving complex problems in scientific computing.

Daniel Appelo
Virginia Tech
appelo@vt.edu

Yingda Cheng
Virginia Polytechnic Institute and State University
yingda@vt.edu

MS230

Approximation of Radiation Boundary Kernels for Spheroidal Surfaces

We consider numerical evaluation and approximation of the frequency-domain Dirichlet-to-Neumann kernel appropriate for a spheroidal boundary and outgoing waves.

Stephen Lau
Mathematics
The University of New Mexico
lau@math.unm.edu

MS230

Optimization of Complete Radiation Boundary Conditions for a GPU Cluster

In this work, I discuss the implementation and computational efficiency of complete radiation boundary conditions (CRBC) for the 3D acoustic wave equations. CRBC is proposed as a more complete alternative to the widely used perfectly matched layer (PML). The goal is to integrate this method into an existing PDE-solving codebase optimizing for scalability and clarity of implemention. The target hardware for this work is a 20-node NVIDIA DGX Superpod equipped with 160 A100 GPUs. Future work will explore hyperbolic PDEs possed for topological Metamaterials which support surface modes along a material interface.

Mason A. Mccallum
Southern Methodist University
mmccallum@smu.edu

MS230

A Novel PML-type Technique Based on Real Coordinate Transformation

The Perfectly Matched Layer (PML) is a widely utilized artificial absorbing layer for truncating computational domains in wave scattering problems. While the conventional PML based on complex coordinate transformations is well-established, we propose a novel PML-inspired technique employing real coordinate transformations (RCT). This approach significantly simplifies the PML equations by converting the slow, algebraically decaying factor of the outgoing wave into one that decays exponentially. Notably, unlike traditional methods, the RCT-based layer is not artificialthe computed field within the layer accurately reconstructs the outgoing wave of the original scattering problem in the unbounded domain. Furthermore, by extracting the infinite oscillatory pattern within the real compressed layer (RCL), finite element methods (FEM) and other domain discretization techniques can achieve the same level of accuracy as the classic PML.

Bo Wang
School of Mathematics and Statistics
Hunan Normal University
bowang@hunnu.edu.cn

Li-Lian Wang
School of Physical & Mathematical Sciences
Nanyang Technological University
lilian@ntu.edu.sg

MS230

A Null Infinity Layer for Wave Scattering

Null infinity is the asymptotic region where outgoing waves or radiation propagate far from a source or scatterer, providing a natural framework for analyzing wave behavior on unbounded domains. Building on this concept, I present a geometric approach for solving time-harmonic wave scattering problems without resorting to domain truncation. By mapping the unbounded domain to a bounded one and scaling the oscillatory decay toward null infinity, this method enables efficient and accurate computations of the far-field pattern on a finite numerical grid. Inspired by de-

velopments in numerical relativity for time-domain wave equations, the approach integrates mathematical ideas from relativity, hyperbolic geometry, and conformal compactification. A carefully designed transformation layer confines these mappings to an annular domain. I demonstrate the effectiveness of the method with one- and two-dimensional numerical examples.

Anil Zenginoglu
University of Maryland
anil@umd.edu

MS232

A Physical and Generative Framework for Data-driven Modeling with Uncertainty Quantification

In the era of big data, advancements in sensor technologies have enabled frequent access to rich spatio-temporal data. This development has paved the way for a new paradigm, where predictive dynamical models are automatically extracted from data streams, minimizing the need for extensive prior knowledge and ad-hoc tuning. We present a generative framework for distilling physical models from raw data. Our approach involves discovering low-dimensional latent variables and identifying a latent dynamical model in this new set of variables. This is achieved by coupling variational autoencoders for dimensionality reduction with Variational Identification of Nonlinear Dynamics (VINDy) to learn a probabilistic dynamical model from a set of candidate features. Once trained, the model is employed in an online generative phase to compute full-time solutions for new control parameters or initial setups. The probabilistic framework naturally enables uncertainty quantification, providing uncertainty-aware estimates. Additionally, data assimilation techniques using Kalman filters are integrated into the framework to track the evolution of the physical system and adapt the digital model accordingly. The identified model can account for external inputs, enabling control actions on the system, thus closing the digital-physical loop within a digital twin framework. We demonstrate the effectiveness of our method on a wide range of high-dimensional, nonlinear dynamical systems.

Paolo Conti
Politecnico di Milano
paolo.conti@polimi.it

Jonas Kneifl
Institute of Engineering and Computational Mechanics
University of Stuttgart
jonas.kneifl@itm.uni-stuttgart.de

Luca Rosafalco, Andrea Manzoni
Politecnico di Milano
luca.rosafalco@polimi.it, andrea1.manzoni@polimi.it

Attilio Frangi
Politecnico di Milano
Dipartimento di Ingegneria Civile e Ambientale, DICA
attilio.frangi@polimi.it

Joerg Fehr
Institute of Engineering and Computational Mechanics
University of Stuttgart
joerg.fehr@itm.uni-stuttgart.de

Steve Brunton
University of Washington
sbrunton@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS232

From Ageing-Aware Battery Discharge Prediction to Optimal Health-Aware Operation

In Prognostics and Health Management (PHM), the integration of machine learning has paved the way for advanced predictive models that ensure the reliable operation of complex industrial assets. However, the challenges posed by sparse, noisy, and incomplete data require the integration of prior knowledge and inductive bias to improve generalization, interpretability, and robustness. Inductive bias, which refers to the set of assumptions embedded in machine learning models, plays a crucial role in guiding these models to generalize effectively from limited training data to real-world scenarios. Particularly within the context of PHM, where physical laws and domain knowledge are integral, the use of inductive bias can significantly enhance a models ability to predict system behavior under diverse operating conditions. By embedding physical principles into learning algorithms, inductive bias reduces the dependence on large datasets, ensures that model predictions are physically consistent, and enhances the generalizability and interpretability of the models. This talk will explore the various forms of inductive bias applied in PHM, with a special emphasis on how incorporating structural inductive biases into learning architecturessuch as through physics-informed graph neural networkscan effectively address the aforementioned challenges by capturing the complexities of system dynamics and data quality.

Olga Fink
École Polytechnique Fédérale de Lausanne
olga.fink@epfl.ch

MS232

Model Order Reduction and Scientific Machine Learning As Enablers for Digital Twins

We focus on some perspectives about Reduced Order Methods (ROM) for parametric Partial Differential Equations (PDEs) with a special interest in parametric problems arising in offline-online Computational Fluid Dynamics (CFD) with a special attention to the enhancement provided by Scientific Machine Learning (SML) and to digital twin(s). Efficient parametrisations (random inputs, geometry, physics) are very important to be able to properly address an offline-online decoupling of the computational procedures and to allow competitive computational performances. Current ROM developments in CFD include: (i) a better use of stable high fidelity methods, to enhance the quality of the reduced model too, also in presence of bifurcations and loss of uniqueness of the solution itself, (ii) capability to incorporate turbulence models and to increase the Reynolds number; (iii) more efficient sampling techniques to reduce the number of the basis functions, retained as snapshots, as well as the dimension of online systems. All the previous aspects are quite relevant – and often challenging – in complex real-world CFD problems to focus, for example, on real time simulations for complex parametric industrial, environmental and biomedical flow problems, or even in a control flow setting with data assimilation and uncertainty quantification. Some model problems will be illustrated by focusing on few benchmark

study cases, applied to problems of interest.

Gianluigi Rozza
SISSA, International School for Advanced Studies, Trieste, Italy
grozza@sissa.it

MS232

InVAErt Networks for Amortized Inference and Identifiability Analysis of Lumped Parameter Hemodynamic Models

Estimation of cardiovascular model parameters from electronic health records (EHR) poses a significant challenge primarily due to lack of identifiability. Structural non-identifiability arises when a manifold in the space of parameters is mapped to a common output, while practical non-identifiability can result due to limited data, model misspecification, or noise corruption. To address the resulting ill-posed inverse problem, optimization-based or Bayesian inference approaches typically use regularization, thereby limiting the possibility of discovering multiple solutions. In this study, we use inVAErt networks, a neural network-based, data-driven framework for enhanced digital twin analysis of stiff dynamical systems. We demonstrate the flexibility and effectiveness of inVAErt networks in the context of physiological inversion of a six-compartment lumped parameter hemodynamic model from synthetic data to real data with missing components.

Guoxiang Tong, Daniele E. Schiavazzi
University of Notre Dame
gtong@nd.edu, dschiavazzi@nd.edu

Carlos Sing Long
Institute for Mathematical and Computational Engineering
Pontificia Universidad Catolica de Chile
casinglo@uc.cl

MS233

Space-Time Parallel Scaling of Parareal with a Fourier Neural Operator As Coarse Propagator

Iterative parallel-in-time algorithms like Parareal can extend scaling beyond the saturation of purely spatial parallelization when solving initial value problems. However, they require the user to build coarse models to handle the inevitable serial transport of information in time. This is a time-consuming and difficult process since there is still limited theoretical insight into what constitutes a good and efficient coarse model. Novel approaches from machine learning to solve differential equations could provide a more generic way to find coarse-level models for parallel-in-time algorithms. This paper demonstrates that a physics-informed Fourier Neural Operator (PINO) is an effective coarse model for the parallelization in time of the two-asset Black-Scholes equation using Parareal. We demonstrate that PINO-Parareal converges as fast as a bespoke numerical coarse model and that, in combination with spatial parallelization by domain decomposition, it provides better overall speedup than both purely spatial parallelization and space-time parallelization with a numerical coarse propagator.

Abdul Qadir Ibrahim
Hamburg Univeristy of Technology
abdul.ibrahim@tuhh.de

MS233

MultigridDeferred Correction Time Integrators

Parallel time-integration is a promising approach to accelerate numerical simulations of time-dependent PDEs, harnessing the computational capability of modern supercomputers. There are two relatively modern parallel time integration techniques that non-intrusively take a user-specified serial time-integrator to construct a parallel time-integrator. Revisionist Integral Deferred Correction (RIDC) utilizes the deferred correction framework to increase the order of accuracy of the specified serial integrator, in approximately the same wall-time as the sequential time integrator. Multigrid Reduction in Time (MGRIT) utilizes a hierarchy of time discretizations to provide acceleration over the specified sequential time integrator. This talk is concerned with combining these two philosophies, taking a user-specified sequential time integrator and designing a parallel time integrator that computes a more accurate solution in a shorter wall time.

Benjamin W. Ong
Michigan Technological University
ongbw@mtu.edu

Robert D. Falgout
Lawrence Livermore National Laboratory
falgout2@llnl.gov

MS234

Micro-Macro Multi-Level Spectral Deferred Correction Method

Spectral Deferred Correction (SDC) methods are an iterative technique to solve initial value problems numerically. SDC methods can be seen as applying a suitable preconditioner to a Picard iteration, resulting in faster and more reliable convergence to a collocation solution. A notable variation of SDC is multi-level SDC (MLSDC), which is inspired by non-linear multigrid methods and involves performing correction sweeps across a hierarchy of levels. In this talk, we introduce a new method called Micro-Macro Multi-Level SDC (M3LSDC) for second-order ordinary differential equations (ODEs). M3LSDC builds on the MLSDC approach but employs a reduced-order model on the coarse level, rather than a coarser discretization. We will demonstrate the advantages of this method using numerical examples, such as the Duffing and Lorentz equations.

Ikrom Akramov, Daniel Ruprecht
Institute of Mathematics
Hamburg University of Technology
ikrom.akramov@tuhh.de, ruprecht@tuhh.de

Sebastian Götschel
Hamburg University of Technology
sebastian.goetschel@tuhh.de

Ignace Bossuyt, Giovanni Samaey
KU Leuven
ignace.bossuyt1@kuleuven.be, giovanni.samaey@kuleuven.be

MS234

Space-Time Parallel Simulations on GPUs with

pySDC

Spectral deferred correction (SDC) methods offer various opportunities for concurrency in the time direction. Recent developments in diagonal preconditioners have enabled small scale parallelism with particularly high parallel efficiency by solving for all stages simultaneously. We combine this with spectral methods in space, which can be parallelized easily via distributed Fourier transforms to obtain massively parallel schemes. By using GPUs, we arrive at implementations that efficiently cater to modern HPC systems at scale. Our implementations form part of pySDC, a Python code that enables rapid prototyping of a wide range of SDC and parallel-in-time related aspects of time integration. We demonstrate excellent strong and weak scaling for multiple PDEs, showcasing the practical capabilities of both the method and the implementation.

Thomas Baumann
Forschungszentrum Juelich
Technische Universitaet Hamburg Harburg
t.baumann@fz-juelich.de

MS234

Polynomial Integrators and SDC Iterations

Polynomial block methods are multivalued general linear methods constructed using continuous polynomials that approximate the ODE solution in time. These methods naturally allow for small-scale parallelism in the form of simultaneous right-hand-side evaluations and output computations. In this talk, I will describe how to combine SDC iterations with polynomial block methods to unlock parallelization both across the method and across the steps. I will discuss the stability and convergence of these methods, and present numerical experiments that highlight certain advantages of polynomial integrators over the well-known Parareal method.

Tommaso Buvoli
Tulane University
tbuvoli@tulane.edu

MS234

Robust Semi-Implicit Multilevel SDC Methods for Time-Integration of Conservation Laws

Semi-implicit spectral deferred correction (SDC) methods provide a systematic approach to construct time integration methods of arbitrarily high order for nonlinear evolution equations including conservation laws. They converge towards A- or even L-stable collocation methods, but are often not sufficiently robust themselves. In this paper, a family of SDC methods inspired by an implicit formulation of the Lax-Wendroff method is developed. Compared to fully implicit approaches, the methods have the advantage that they only require the solution of positive definite or semi-definite linear systems. Numerical evidence suggests that the proposed semi-implicit SDC methods with Radau points are L-stable up to order 11 and require very little diffusion for orders 13 and 15. The excellent stability and accuracy of these methods is confirmed by numerical experiments with 1D conservation problems, including the convection-diffusion, Burgers, Euler and Navier-Stokes equations. For enhancing the computational efficiency we consider multilevel methods and Krylov acceleration. Depending on the progress made, preliminary results on space-time adaptivity and application to multidimensional flow solvers will be shown.

Jörg Stiller
Dresden University of Technology
joerg.stiller@tu-dresden.de

Erik Pfister
TU Dresden
erik.pfister@tu-dresden.de

MS235

Approximating Riemannian Gradient Flows on Quantum Computers for Ground State Problems

Adaptive quantum algorithms have been proposed to overcome challenges of existing hybrid quantum-classical algorithms related to ansatz selection and the optimization landscape structure. Instead of fixing an ansatz, in these approaches a quantum circuit is successively grown informed by measurement data. In this talk, I describe adaptive quantum algorithms as approximations of Riemannian gradient flows on the unitary group; an optimization framework where one directly optimizes over quantum circuits rather than gate parameters. I show that despite the existence of saddle points, the full gradient flow converges to the ground state for almost all initial states. While the full gradient flow is in general not efficiently implementable on quantum computers, I go on to discuss several approximation schemes that exhibit similar convergence behavior but require only a polynomial overhead when it comes to device implementations.

Arenz Christian
Arizona State University
christian.arenz@asu.edu

MS235

Optimal Control of Open Quantum Systems

The optimal control problem for open quantum systems can be formulated as a time- dependent Lindbladian that is parameterized by a number of time-dependent control variables. We present algorithms for solving this optimal control problem efficiently, i.e., having a poly-logarithmic dependency on the system dimension, which is exponentially faster than best-known classical algorithms. Our algorithms are hybrid, consisting of both quantum and classical components. The quantum procedure simulates time-dependent Lindblad evolution that drives the initial state to the final state, and it also provides access to the gradients of the objective function via quantum gradient estimation. The classical procedure uses the gradient information to update the control variables.

Xiantao Li
Pennsylvania State University
xxl12@psu.edu

MS235

Quantum Digital Twins

We present Quantum Digital Twins (QDTs), which are digital clones of existing quantum computers. Our QDTs are built within a robust generalized Bayesian framework, enriched by and integrated with the optimal transportation theory, realizing bidirectional interactions between quantum computers and virtual models on classical systems. We present the application of a simple QDT to a transmon

quantum device using Ramsey quantum measurements.

Mohammad Motamed
University of New Mexico
motamed@unm.edu

MS235

Incorporating Multi-Qubit Effects in Maxwell-Schrdinger Methods for Efficiently Modeling Superconducting Circuit Quantum Computers

To achieve quantum error correction on superconducting circuit quantum processors, multi-qubit gate performance needs to be significantly improved. This is challenging due to the complex multi-qubit interactions, which include unintended classical and quantum crosstalk and leakage out of the computational basis. To successfully design optimal control pulses, it is critical that the underlying modeling method used in the optimization process captures these complex effects. Maxwell-Schrdinger methods that model the self-consistent interactions between classical electromagnetic fields and qubit dynamics have shown promise for these control optimization purposes, but have never been demonstrated for multi-qubit systems. Here, we show how multi-qubit interactions that occur in superconducting circuit quantum processors can be rigorously incorporated into Maxwell-Schrdinger methods. We discuss the formulation of the coupled equations of motion for this system and how the multi-qubit exchange coupling rate due to quantum electromagnetic interactions can be characterized in a pre-processing step through classical electromagnetic simulations. Details on a numerical method for solving this coupled Maxwell-Schrdinger system are also discussed. Simulations of cross-resonance gates, a common multi-qubit gate in this hardware platform, are presented to demonstrate the need for this modeling approach, as well as its efficiency compared to standard modeling approaches.

Ghazi Khan, Thomas Roth
Purdue University
khang@purdue.edu, rothte@purdue.edu

MS236

Variable Projection Method for Mixed Gaussian Models

The variable projection method has been developed as a powerful tool for solving separable nonlinear least squares problems. It has proven effective in cases where the underlying model consists of a linear combination of nonlinear functions, such as exponential functions. A modified version of the variable projection method to address a challenging semi-blind deconvolution problem involving mixed Gaussian kernels is employed. The aim is to recover the original signal accurately while estimating the mixed Gaussian kernel utilized during the convolution process. The numerical results obtained through the implementation of the proposed algorithm are presented.

Jordan T. Dworaczyk
Arizona State University
jtdworac@asu.edu

MS236

Inverse Source Problems for Fractional Parabolic Equations: A Regularization Method and Reconstruction Algorithm

Fractional Differential Equations (FDEs) arise from many real-world applications in diverse areas, such as physics, Mechanics and dynamic systems, and signal and image processing. The forward and inverse problems for FDEs have been investigated intensively in recent decades. As usual, inverse problems for FDEs are generally ill-posed, which increases the challenges in studying them. This talk will discuss a regularization method for inverse problems of finding a factor of the source term for fractional parabolic equations. This method guarantees the Holder-type error estimate with the optimal order. This ends with an algorithm for recovering the source and numerical illustrations. This is a joint work with Nguyen Van Duc (Vinh University).

Thi-Phong Nguyen
New Jersey Institute of Technology
thiphong.nguyen@njit.edu

MS236

ComputationalComputational Boundary Control Methods for Acoustic Inverse Boundary Value Problems

We consider inverse boundary value problems for the acoustic wave equation from near field data represented by the Neumann-to-Dirichlet map. We develop linearized boundary control methods to analyze the stability and reconstruction. The analysis leads to reconstructive algorithms that are numerically validated. This is joint work with Lauri Oksanen (University of Helsinki) and Tianyu Yang (Michigan State University).

Yang Yang
Michigan State University
yangy5@msu.edu

MS237

Layer Parallel Training of Transformer Networks

Transformers utilize self-attention and feed-forward mechanisms to achieve state-of-the-art performance on a variety of language modeling tasks. However, as these networks become increasingly larger and deeper, additional sources of parallelism are needed to accelerate their training. By leveraging the residual structure inherent in transformers, we interpret the network as an Euler time step and apply parallel-in-time techniques to achieve speed-up. We demonstrate this speed-up on a variety of problems, including BERT and GPT-style networks.

Shuai Jiang
Sandia National Laboratories, U.S.
sjiang@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

Alena Kopanicakova
Brown University
alena_kopanicakova@brown.edu

Marc Salvadó
Università della Svizzera italiana
marc.salvado@usi.ch

Jacob B. Schroder
Department of Mathematics and Statistics

University of New Mexico
jbschroder@unm.edu

MS237

Machine Learning-based Surrogate Models for an Efficient Homogenization of Open-porous Materials

Open-porous materials like aerogels typically feature a complex nanostructure which presents a significant challenge for simulating their mechanical properties. In the case of aerogels sizes and shapes of the nanopores can vary a lot depending on the type of aerogel and the conditions during its synthesis. To account for the nanostructured morphology of the material, multiscale numerical homogenization approaches are commonly employed. Based on the well-established FE^2-method [1] a homogenization approach can be applied which considers a beam frame model for the representation of a fibrillar nanostructure [2]. To reduce the computational load of solving each microscopic problem the beam frame solver can be replaced by a machine learning-based surrogate model which is trained to predict the resulting stresses of the beam frame model for a given deformation. The machine learning approach is expected to yield a high potential of reducing the total computational load of the multiscale method while maintaining flexibility and good convergence. [1] Feyel, F. *Multiscale FE^2 elastoviscoplastic analysis of composite structures* Comp. Mat. Sci., 1999. [2] Klawonn, A., Lanser, M., Mager, L., and Rege, A. *Computational homogenization for aerogel-like polydisperse open-porous materials using neural network-based surrogate models on the microscale.* Accepted for publication, in revised form, in Comp. Mech., Springer, December 2024.

Axel Klawonn
Department of Mathematics and Computer Science and Center
for Data Centand Simulation Science, University of Cologne
axel.klawonn@uni-koeln.de

Martin Lanser
Universitaet zu Koeln
Mathematisches Institut
martin.lanser@uni-koeln.de

Lucas Mager
Department of Mathematics and Computer Science
Universitaet zu Koeln
lucas.mager@uni-koeln.de

Janine Weber
Universität zu Köln
Department of Mathematics and Computer Science
janine.weber@uni-koeln.de

MS237

A Nonoverlapping Domain Decomposition Method for Extreme Learning Machines: Elliptic Problems

Extreme learning machine (ELM) is a methodology for solving partial differential equations (PDEs) using a single hidden layer feed-forward neural network. It presets the weight/bias coefficients in the hidden layer with random values, which remain fixed throughout the computation, and uses a linear least squares method for training the parameters of the output layer of the neural network. It is known to be much faster than Physics informed neural networks. However, classical ELM is still computationally expensive when a high level of representation is desired in the solution as this requires solving a large least squares system. In this talk, we propose a nonoverlapping domain decomposition method (DDM) for ELMs that not only reduces the training time of ELMs, but is also suitable for parallel computation. We introduce local neural networks, which are valid only at corresponding subdomains, and an auxiliary variable at the interface. We construct a system on the variable and the parameters of local neural networks. A Schur complement system on the interface can be derived by eliminating the parameters of the output layer. The auxiliary variable is then directly obtained by solving the reduced system after which the parameters for each local neural network are solved in parallel. A method for initializing the hidden layer parameters suitable for high approximation quality in large systems is also proposed.

Chang-Ock Lee
Department of Mathematical Sciences
KAIST
colee@kaist.edu

MS238

Advancing Multi-Scale Kinetic Modeling in Strongly Magnetized Relativistic Plasmas: An Analytical Particle Pusher Approach

Originally developed by F.H. Harlow at LANL for hydrodynamic simulations, Particle-in-Cell (PIC) methods are now widely used to model kinetic plasmas in many applications. However, simulating multi-scale physics in strongly magnetized relativistic plasmas, especially for particle acceleration, remains challenging for conventional PIC methods. This study leverages a known analytical solution for particle motion in constant electromagnetic fields to develop approximate solutions for non-uniform fields. The approach improves simulation accuracy and scale for strong magnetic fields, surpassing standard methods. We demonstrate the benefits of this analytical particle pusher through benchmarks and comparisons with traditional PIC methods. Additionally, we apply the new algorithm to realistic scenarios, such as relativistic shocks, spanning weak to strong magnetizations. This study aims to shed light on the enhanced capabilities of the proposed analytical particle pusher and its potential impact on advancing our understanding of plasma dynamics in multi-scale astrophysical environments.

Guangye Chen, Nicole Lloyd-Ronning, Chengkun Huang, Fan Guo, Lucian Sahd
Los Alamos National Laboratory
gchen@lanl.gov, lloyd-ronning@lanl.gov,
huangck@lanl.gov, guofan@lanl.gov, lsahd@lanl.gov

Patrick Kilian
Space Science Institute
pkilian@spacescience.org

MS238

Overcoming (some) Challenges in Simulating Charged Particle Transport Through Tailored Numerical Methods

Proton Beam Therapy (PBT) is a type of radiotherapy used for cancer treatment. Due to the sharply peaked dose-depth curve characteristic of protons, and the fact that protons stop at a finite depth inside the tissue, PBT

is especially useful when treating tumours situated close to vital organs, which need to be spared from radiation damage. To produce a treatment plan suited to a specific patient, an accurate forward model for proton radiation is required. Many evaluations of this forward model are needed to produce an optimal treatment plan, meaning that in addition to being accurate, the forward model also needs to be quick to evaluate, either numerically or analytically. This talk will cover some models commonly used for modelling proton radiation, including the Boltzmann Radiation Transport Equation and simpler PDE models that can be derived as approximations of it. We will discuss some of the complexities of accurately approximating the solutions to these equations through numerical methods, and present some ways in which these complexities can be tackled.

Veronika Chronholm
University of Bath
mvc34@bath.ac.uk

MS238

A Micro-macro Decomposition for Implicit time-stepping of the BGK Model

There has been recent interest in the acceleration of solvers to nonlinear systems arising from implicit temporal discretizations of the space-homogeneous BGK equation. One such method, called High-Order Low-Order (HOLO), uses a fluid model to accelerate standard source iteration schemes used to solve the implicit system. In this talk, we present a Micro-Macro (MM) method for solving this system. Micro-Macro methods represent the kinetic distribution as a sum of a local Maxwellian equilibrium and a micro perturbation. In highly-collisional regimes, the micro component is small and can be compressed reducing the overall storage cost of the distribution. However, study of the MM method has been primarily focused on explicit and implicit-explicit temporal discretizations where structural conditions on the micro component can be easily preserved. We build a MM method for fully-implicit timestepping that offers the similar acceleration benefits as the HOLO method, while also allowing compression of the storage cost of the distribution. We compare the accuracy and performance of the MM and HOLO methods. Additionally, we show the compression benefits of the MM method as a function of the collision scale.

Cory Hauck
Oak Ridge National Laboratory and University of Tennessee
hauckc@ornl.gov

Paul Laiu
Oak Ridge National Laboratory
laiump@ornl.gov

Stefan Schnake
Oak Ridge National Laboratory, U.S.
schnakesr@ornl.gov

MS238

Entropic Fokker-Planck Method for Polyatomic Gases

When gases are far from thermal equilibrium, conventional continuum models fail to accurately describe their behavior. Instead, a kinetic model that describes gases statistically is necessary so that molecular interactions can be taken into account. The Boltzmann equation is the most widely used kinetic model and is typically solved by particle Monte-Carlo methods (DSMC). These methods offer high physical accuracy but are computationally expensive, especially for near-equilibrium flows, as collisions must be calculated explicitly. An alternative kinetic model is given by the Fokker-Planck (FP) equation, which approximates the Boltzmann equation by modeling the effect of binary collisions as a drift-diffusion process. This approach allows for a more efficient particle method via the underlying Langevin equation, eliminating the need for explicit collision calculations. In this talk, we introduce a novel FP equation for polyatomic gases, derived to ensure correct relaxation rates of key moments and adherence to the H-theorem. Numerical experiments demonstrate that our FP method achieves good agreement with DSMC results while significantly reducing computational time, particularly in near-equilibrium regimes.

Lukas Netterdon
Applied and Computational Mathematics
RWTH Aachen University
netterdon@acom.rwth-aachen.de

Veronica Montanaro
Laboratory for Computational Engineering
Empa, Switzerland
veronica.montanaro@empa.ch

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

Hossein Gorji
Laboratory for Computational Engineering
Empa, Switzerland
mohammadhossein.gorji@empa.ch

MS238

Challenges in Simulation of Charged Particle Transport

Proton Beam Therapy (PBT) is a type of radiotherapy used for cancer treatment. Due to the sharply peaked dose-depth curve characteristic of protons, and the fact that protons stop at a finite depth inside the tissue, PBT is especially useful when treating tumours situated close to vital organs, which need to be spared from radiation damage. To produce a treatment plan suited to a specific patient, an accurate forward model for proton radiation is required. Many evaluations of this forward model are needed to produce an optimal treatment plan, meaning that in addition to being accurate, the forward model also needs to be quick to evaluate, either numerically or analytically. This talk will cover some models commonly used for modelling proton radiation, including the Boltzmann Radiation Transport Equation and simpler PDE models that can be derived as approximations of it. We will discuss some of the complexities of accurately approximating the solutions to these equations through numerical methods, and present some ways in which these complexities can be tackled.

Tristan Pryer
University of Bath

tmp38@bath.ac.uk

MS239

Non-Intrusive Model Reduction of Coupled Aero-Thermal Simulations

This presentation develops a non-intrusive reduced-order model (ROM) applied to a coupled aero-thermal simulation problem. Aero-thermal simulations are crucial for modeling and designing ablative thermal protection systems for atmospheric reentry vehicles but are typically incredibly computationally expensive and rely on large-scale production codes, which are not straightforward to include in design and analysis workflows. Non-intrusive model reduction allows one to construct a much more computationally inexpensive ROM without requiring access to the full-order model (FOM) source code and while maintaining high accuracy. We compare proper orthogonal decomposition with recent nonlinear dimension reduction techniques and learn the dynamics in the reduced space using a radial basis function interpolant trained using FOM trajectory data and carefully constructed synthetic data. The proposed ROM is demonstrated on a simulation of a reentry flight vehicle with a paraboloid geometry.

Alejandro N. Diaz
Rice University
andiaz@sandia.gov

Patrick J. Blonigan, John Tencer, Kanad Khanna, Jonathan Murray
Sandia National Laboratories
pblonig@sandia.gov, jtencer@sandia.gov, kkhanna@sandia.gov, jmurray@sandia.gov

MS239

Multicontinuum Homogenization and Its Application

In this talk, I will discuss an approach for multicontinuum homogenization. The main idea of this approach is to identify multiple continua at each macroscopic point. We will discuss the idea and the method derived from it. Further, we will discuss its relation to previously developed multiscale methods. Numerical examples of interface problems will also be provided in this talk.

Wing Tat Leung
City University of Hong Kong
wtleung27@cityu.edu.hk

MS239

A Data-Driven Subspace Optimized Method for Electrical Impedance Tomography

Electrical Impedance Tomography (EIT) is a non-invasive imaging technique that reconstructs the internal conductivity distribution of a subject based on electrical measurements taken from its boundary. The reconstruction process is often ill-posed and computationally expensive, which challenges the accuracy and efficiency of EIT. This paper presents a novel data-driven subspace optimization method designed to enhance the performance of EIT. The proposed approach is inspired by direct imaging methods and optimizes the subspaces in which the reconstruction and induced currents occur, reducing computational complexity while maintaining accuracy. The method leverages a large dataset of precomputed simulations to identify a low-dimensional subspace that captures the most significant features of the conductivity distribution. Experimental results demonstrate that the proposed method not only accelerates the reconstruction process but also improves the quality of the resulting images, outperforming conventional techniques in both synthetic and real-world scenarios.

Ying Liang
Purdue University
ying.liang@duke.edu

Hongkai Zhao
Duke University
zhao@math.duke.edu

MS239

Multicontinuum Splitting Scheme for Multiscale Flow Problems

In this work, we propose a multicontinuum splitting scheme for multiscale problems, illustrated through a parabolic equation with a high-contrast coefficient. Utilizing the framework of multicontinuum homogenization, we introduce smooth macroscopic variables and decompose the solution space into two components to effectively separate the dynamics at different speeds (or the effects of contrast in high-contrast cases). By treating the fast component (depending on the contrast) implicitly and the other one explicitly, we construct partially explicit time discretization schemes, which can reduce computational cost. Contrast-independent stability conditions are derived. Additionally, we discuss possible methods to obtain an optimal decomposition of the solution space, which relaxes the stability conditions while enhancing computational efficiency. A Rayleigh quotient problem in tensor form is formulated, and simplifications are achieved under certain assumptions. Finally, we present numerical results for various coefficient fields with different continua to validate our proposed approach. It can be observed that the multicontinuum splitting schemes enjoy high accuracy and efficiency.

Yalchin Efendiev
Dept of Mathematics
Texas A&M University
efendiev@math.tamu.edu

Wing Tat Leung
City University of Hong Kong
wtleung27@cityu.edu.hk

Buzheng Shan
Texas A&M University
shanbzh@tamu.edu

Min Wang
University of Houston
mwang55@cougarnet.uh.edu

MS239

Decoupling Schemes for Multiscale Multicontinuum Problems

Multiscale multicontinuum models are used to describe problems with high-contrast properties at a macro level. Upscaling or multiscale methods are utilized to construct a macro-scale model by identifying the macroscale continuum and connection inside and between continua. In up-

scaling methods, macro-scale variables are represented as average coarse-scale characteristics and should be defined based on prior knowledge about connectivity. In a generalized multiscale method, local spectral problems can be used to identify the macro-scale continuum. Both methods, with accurate identification of the continuums and connections between them, lead to an accurate representation of the problem at the macro level. In this work, we introduce an effective decoupling method that separates the equations for each continuum. The approach is built on an implicit-explicit approximation by time with an explicit treatment of the coupling term. We have developed and analyzed continuum decoupling techniques for multiscale multicontinuum models with application to flow in fractured and high-contrast domains. Numerical results demonstrate that the proposed scheme is stable, accurate, and computationally efficient.

Maria Vasilyeva
Texas A&M UniversityCorpus Christi
maria.vasilyeva@tamucc.edu

MS240

Mathematical Modeling for Enhanced Power Absorption in Hybrid Wave Energy Systems Under Ocean Currents

The increasing demand for renewable energy has sparked interest in hybrid wave energy technologies, which integrate multiple energy capture mechanisms for greater efficiency. In this context, the present study focuses on a hybrid wave energy converter that integrates a piezoelectric plate with a U-shaped oscillating water column device strategically placed over a uniform seabed. A mathematical model is developed to assess the impact of ocean wave-current interactions on this system, using classical small amplitude potential flow theory to analyze hydrodynamic factors and a coupled boundary element-finite difference method to manage complex fluid-structure interactions, enabling the detailed evaluation of power extraction under various structural and oceanic scenarios. The study demonstrates that ocean currents significantly enhance power generation, with the Doppler shift in apparent angular frequency notably enhancing energy capture in intermediate and long wave regimes. The results also reveal that the hybrid system offers superior efficiency and more stable power output compared to a stand-alone oscillating water column wave energy converter device. This model serves as a crucial tool for optimizing system design, addressing challenges in wave energy utilization, and advancing future research and industrial applications in renewable energy.

Santanu K. Dash
Birla Institute of Technology & Science, Pilani
Hyderabad Campus, Hyderabad- 500078, Telangana, India
p20210443@hyderabad.bits-pilani.ac.in

MS240

Computational Framework for Automated Analysis of Dendritic Spine Geometry

Recent advances in connectomics have enabled the reconstruction of neural tissues at nanometer resolution using electron microscopy (EM), providing new insights into brain structures, particularly dendritic spines. These dynamic protrusions are crucial for synaptic plasticity and play key roles in various neurological processes, including learning, memory, and disease states. Traditional methods of spine analysis, while effective for isolated spines, often require human expert intervention to fully capture the complexity of dendritic spines embedded within intricate networks. This paper proposes an advanced, automated computational framework that combines discrete differential geometry, machine learning, and 3D image processing techniques to analyze dendritic spines in complex environments. The approach emphasizes the continuous nature of spine morphology, moving beyond fixed classification schemes.

Abdel Kader A. Geraldo, Thomas Fai
Brandeis University
ageraldo@brandeis.edu, tfai@brandeis.edu

MS240

Algorithms for Computing Membrane Curvature

Computation of membrane bending forces involves the numerical approximation of the Laplace-Beltrami operator on a meshed surface. The popular cotangent scheme is known to converge pointwise to the Laplace-Beltrami operator for certain restricted classes of meshes, but in general it fails to converge. Several promising discretizations of the Laplace-Beltrami operator based on approximations of the fundamental solution to the heat equation have been proposed. These heat kernel-based discretization schemes guarantee pointwise convergence for a broad class of meshes at the cost of increased computational complexity. A straightforward implementation of a heat kernel-based method results in an algorithm whose computational cost is quadratic in the number of vertices in the mesh. We will present numerical experiments that show the computational cost of heat kernel-based algorithms can be significantly improved using an adaptive quadrature scheme.

Wilson Lough
University of Colorado Boulder
wilson.lough@colorado.edu

MS241

Estimating the Unobserved: Adjoint Based State Estimation and Sensitivity Studies Using the Ice-Sheet and Sea Level System Model (issm) in West Antarctica

Improving estimates of ice/ocean parameters and states reduces sea level change uncertainty. However, improving estimates under the Antarctic Ice Sheet (AIS) is non-trivial, as direct observations are typically not feasible. Data assimilation using models of physical processes can constrain otherwise unobservable states and parameters of the AIS. Capabilities developed within the Ice-sheet and Sea-Level System Model (ISSM) allow for the simultaneous inversion of such unobservables. More specifically, newly improved automatic differentiation/adjoint model capabilities in ISSM allow for the simultaneous inversions of basal friction, ice rheology, sub-ice shelf melt, grounding line positions, and other parameters. These estimates are validated against observations of velocity and thickness changes, then referenced against other estimates within the area of study. This study evaluates this methodology over Ronne Ice Shelf and surrounding areas in West Antarctica from 1992-2023. By providing a physically-consistent estimate of unobservable states and parameters, this method provides insights into the evolution of the AIS over the past 30 years, facilitates coupling to other state estimates such as the ECCO ocean circulation model, and helps constrain future sea

level change. This work is performed at the California Institute of Technologys Jet Propulsion Laboratory under a contract with the National Aeronautics and Space Administration.

Daniel Cheng, Ian Fenty, Eric Larour
Jet Propulsion Laboratory
daniel.l.cheng@jpl.nasa.gov, ian.fenty@jpl.nasa.gov, eric.larour@jpl.nasa.gov

Nicole Schlegel
NOAA Geophysical Fluid Dynamics Laboratory
nicole.schlegel@noaa.gov

MS241

A Python Library for Solving Ice Sheet Modeling Problems Using a Unified Framework: Physics-Informed Neural Networks for Ice and Climate (pinnicle)

Accurately predicting the future contributions of ice sheets to sea level rise is hindered by incomplete knowledge of glaciological processes (e.g., basal friction, ice rheology) and the sparse availability of critical observational data (e.g., subglacial topography). Traditional models address these challenges through data assimilation techniques that solve inverse problems governed by conservation laws. However, these methods often lack flexibility and require extensive data, limiting their effectiveness in data-sparse regions.To overcome these limitations, we introduce PINNICLE (Physics-Informed Neural Networks for Ice and CLimatE), an open-source Python library designed for ice sheet modeling. PINNICLE integrates physics-informed neural networks with observational data to solve forward and inverse problems within a unified framework. It currently supports various models, including the Shelfy-Stream Approximation (SSA) and Mono-Layer Higher-Order (MOLHO) models, for both steady-state and transient simulations. PINNICLE is user-friendly, requiring minimal hyperparameter tuning, and is built on DeepXDE, compatible with machine learning frameworks such as TensorFlow, PyTorch, and JAX. We demonstrate PINNICLE's capabilities through case studies on the Greenland and Antarctic ice sheets. This framework advances ice sheet modeling by improving predictive accuracy and reducing uncertainties in ice dynamics.

Gong Cheng, Mansa Krishna, Mathieu Morlighem
Dartmouth College
Gong.Cheng@dartmouth.edu, mansa.krishna.gr@dartmouth.edu, Mathieu.Morlighem@dartmouth.edu

MS241

Physics-Informed Deep Learning for Data Assimilation of Ice Shelves (DIFFICE.jax)

In this talk, I will introduce DIFFICE.jax: A DIFFerentiable neural-network solver for data assimilation of ICE shelves, written in JAX. I will discuss how PDE-constrained neural networks, trained with real-world data from Antarctica, can help uncover the viscosity model governing ice-shelf dynamics. Despite its crucial role in ice-sheet flow, the flow law of glacial iceits rheologycannot be directly measured at the ice-sheet scale in the field. Various geophysical-scale phenomena may cause the rheology derived in the lab to deviate from actual field behavior. By using physics-informed deep learning, we leverage differentiable modeling and backpropagation to solve inverse problems, constrained by real-world satellite observations and governing equations. We infer glacial rheology that differs from the assumptions commonly used in climate simulations, highlighting the need to reassess the rheology of geophysical complex fluids beyond laboratory settings. Finally, I will discuss some challenges and promises of physics-informed neural networks for data assimilation.

Ching-Yao Lai, Yongji Wang
Stanford University
cyaolai@stanford.edu, yongjiw@stanford.edu

MS241

Estimating Uncertainties in Subglacial Topography Constrained by Mass Conservation Law through Stochastic Simulations

Despite its critical role in ice sheet models, subglacial topography in Antarctica generally remains sparsely measured by airborne radar with 10s to 100s kilometers gaps in-between radar lines. Physics of mass conservation could relate well-observed surface velocity with bed elevation. However, numerically solving topography from mass conservation creates one topography map that does not ensure realistic roughness and could not reveal the inherent uncertainties from sparse elevation measurements. In this study, we propose a novel approach to generate an ensemble of realistically rough and mass-conserving subglacial topography through Monte Carlo Markov Chain sampling. We iteratively perturb the topography with geostatistics methods that reconstruct the topography roughness observed in radar data. The perturbed topography is accepted with a probability derived from the mass conservation law. The method tested on the Denman and Totten Glacier has successfully generated multiple different topography constrained by radar measurements, the mass conservation equation, and the topographic roughness observed in the radar data. This method provides a means to incorporate realistically rough topography into ice sheet models while avoiding artifacts caused by the violation of mass conservation. Furthermore, multiple subglacial topography realizations allow the propagation of uncertainties caused by sparsity of radar measurement to the result of downstream ice sheet models.

Niya Shao
niya shao
niyashao@ufl.edu

Michael Field, Emma MacKie
Department of Geological Science, University of Florida
michael.field@ufl.edu, emackie@ufl.edu

Felicity McCormack
Securing Antarctica's Environmental Future, School of Earth,
Atmosphere and Environment, Monash University, Australia
felicity.mccormack@monash.edu

MS242

The Exascale Computing Project: Teams of Teams (Multi-Team Systems) Study

Collaboration and team science are emerging areas of interest in research and software production. Historically, multi-institutional research collaborations are difficult to initiate and maintain, negatively impacting communica-

tion, negotiation, and dialogue between industry, government, and academic researchers. The Exascale Computing Project (ECP) was created to facilitate broader research collaboration under a shared funding structure and extended timeline to support scientific discovery. We conducted interviews with ECP teams, representing a variety of domain specialities, research institutions, programming backgrounds, and application areas. Using grounded theory as our analytical framework, we assessed how ECPs structure created an environment of increased trust between projects and how software shared between teams facilitated sustained collaboration. Based on our findings within ECP projects, we share key takeaways and recommendations for sustainable multi-institutional collaboration and best shared software practices.

Hana Frluckaj
University of Texas at Austin
hanafrla@utexas.edu

Elaine M. Raybourn
Sandia National Laboratories
emraybo@sandia.gov

Benjamin Sims
LANL
bsims@lanl.gov

Reed Milewicz
Sandia National Laboratory
rmilewi@sandia.gov

MS242

Teaming Up with AI: Using Algorithms and Data Analysis to Enhance Team Formation

This presentation explores the potential of online computational systems to facilitate the formation of more diverse and connected teams. I will begin by examining various conceptualizations of group diversity, synthesizing them into a well-defined optimization problem. Next, I will present the implementation of these algorithms designed to address this optimization challenge, along with their performance across multiple datasets. The effectiveness of these algorithms is then evaluated through a controlled laboratory experiment, providing insights into their real-world applicability. I conclude by discussing the broader implications of this research and outlining promising avenues for future investigation. This work contributes to our understanding of how technology can enhance team composition, potentially leading to more innovative and effective collaborations in various fields.

Diego Gomez-Zara
University of Notre Dame
dgomezara@nd.edu

MS242

Fostering Team Science for All

The field of computational science and engineering is well established as the third pillar of scientific inquiry. Many significant advances in basic science, national security, and U.S. innovation have been made possible over the past several decades through computational science and engineering. The interleaving of high-performance computing, and more recently artificial intelligence, has created additional capabilities as well as challenges, necessitating highly skilled labor ready to tackle bigger and bigger and more complex problems through inter- and multi-disciplinary teams. However, several U.S. government agencies and others have described critical workforce development and recruitment challenges in advanced scientific computing. This session explores why we face critical workforce development challenges and potential solutions to recruitment, retention, and development. We will review the importance of cultivating a workforce that supports people from all walks of life and backgrounds and explore approaches to facilitating workplaces where everyone thrives and contributes to increase innovation.

Mary Ann E. Leung
Sustainable Horizons Institute
mleung@shinstitute.org

MS242

Team Learning for Better Scientific Software

Given the importance of team learning for the effectiveness of interdisciplinary, collaborative research funded by the US NSF and US DOE, a workshop was organized to initiate the development of a set of theoretically-grounded, empirically-driven best practices for learning in scientific software teams. In this talk, I will focus on sharing insights, outcomes, and lessons learned over the course of the workshop. The workshop, funded under the Better Scientific Software Fellowship, was attended by experts from the computational and social sciences, ensuring a broad set of perspectives was represented in the development of guidelines for team learning in scientific software projects. In addition to reviewing the knowledge products of the workshop, I will provide information about how to access and contribute to the maintenance of these resources in this talk. Finally, I will discuss future opportunities to continue the team learning for better scientific software workshop as a series.

Olivia B. Newton
University of Montana
olivia.newton@umt.edu

MS243

Algorithmic Differentiation for Julia: Overview and Source Transforms Ad

In this talk, we review the basics of AD overall and the specificities of Source Transform AD for generating adjoint code specifically. The accent is put on the use of the Tapenade tool for which we describe the novelties in order to include Julia as input language. We will present our first results on Julia code differentiated using Tapenade and detail the various design choices made to accommodate for this new language.

Jean-Luc Bouchot
RWTH Aachen
jean-luc.bouchot@inria.fr

Jean-Luc Bouchot
INRIA Sophia Antipolis
jean-luc.bouchot@inria.fr

MS243

Enhancing the Stability of Quantized Neural Networks: Error Propagation and Gradient Correction

Quantization of neural networks offers significant compu-

tational and memory efficiency gains, making it essential for deploying deep learning models on resource-constrained hardware. However, low-bit precision in weights and activations introduces numerical errors that accumulate during training, particularly in backpropagation, where non-differentiability and quantization noise disrupt gradient flow. This work investigates error propagation in quantized training, analyzing the limitations of the straight-through estimator (STE) and the feasibility of gradient correction terms. We propose an adaptive gradient scaling approach under stochastic rounding to counteract quantization-induced distortions, improving convergence while preserving computational efficiency. Empirical validation on benchmark datasets demonstrates a reduction in gradient mismatch and enhanced training robustness. This work aims to contribute to the reliability of quantization-aware training and expand its applicability in real-world AI use cases.

Oana Marin
PeraCompute Technologies
oanam@peracompute.org

MS243

On the Julia Interface of Adol-C

Recently, an interface to the Julia environment was added to ADOL-C. This newly developed package ADOLC.jl combines the various functionalities of ADOL-C with the convenient usability of Julia. It allows to use all drivers available within ADOL-C for the suitable calculation of derivative information also of very high order. We report on the design of ADOLC.jl and its usage. Furthermore, we analyse the performance of the derivative calculation in Julia using ADOL-C and show corresponding numerical results.

Tim Siebert, Sri Harshitha Tadinada, Andrea Walther
Humboldt-Universität zu Berlin
tim.siebert@hu-berlin.de, sri.tadinada@hu-berlin.de, andrea.walther@hu-berlin.de

MS244

Performance Portability on CPUs, GPUs, and FPGAs for a SYCL Implementation of a Discontinuous Galerkin Shallow Water Model

An increasingly diverse landscape of high performance computing hardware makes the development of simulation codes which, in addition to portability, also demonstrate performance portability a critical issue in many areas of computational science. The goal of performance-portable software is to achieve a sizeable fraction of the practically achievable performance on a given hardware while maintaining, at the same time, a unified codebase. We propose a performance-portable SYCL-based implementation of the 2D shallow water equations which uses a discontinuous Galerkin discretization on unstructured triangular meshes. Using just a few compile-time defined algorithm and data structure specializations our model code demonstrates good performance on a range of CPUs and GPUs of different types as well as on Intel Stratix FPGAs.

Markus Büttner, Vadym Aizinger
University of Bayreuth
markus.buettner@uni-bayreuth.de, vadym.aizinger@uni-bayreuth.de

MS244

Differentiable Physics Simulators for AI-Accelerated Hazard Modeling

Engineering design for landslide hazards requires accurate simulation and parameter identification. Yet, traditional methods for computing derivatives through forward simulations are computationally intensive and unstable. Current simulation frameworks also lack integration capabilities with machine learning models due to their inability to compute gradients of input parameters. To address these challenges, we propose a Differential Programming simulator that combines automatic reverse differentiation with second-order gradient-based optimization, creating a fully differentiable Material Point Method (MPM) simulator. Our approach identifies material properties by iteratively updating parameters to minimize the difference between simulated and observed behavior using L-BGFS. This method achieves higher accuracy by solving exact PDEs than model-agnostic reinforcement learning approaches. We extend this framework with physics-embedded Differentiable Graph Network Simulators (GNS), achieving 165x speedup for granular flow prediction compared to CPU simulations. A novel hybrid GNS/MPM approach provides 24x acceleration while preserving physical conservation laws. Additionally, we develop methods to derive material parameters from video data using Neural Radiance Fields, allowing iterative optimization of material properties and loading conditions through differentiable simulation.

Krishna Kumar
University of Texas at Austin
krishnak@utexas.edu

MS244

SWEMniCS: A Software Toolbox for Modeling Coastal Ocean Circulation, Storm Surges, Inland and Compound Flooding

Flooding from storm surges, rainfall runoff, and their interaction into compounding events are major natural hazards in coastal regions. To assess risks of damages to life and properties alike, numerical models are needed to guide emergency responses and future assessments. Numerical models, such as ADCIRC have over many decades shown their usefulness in such assessments. However, these models have a high threshold in terms of new user engagement as development and compilation is not trivial for users trained in compiled programming languages. Here, we develop a new open-source finite element solver for the numerical simulation of flooding. The numerical solution of the underlying PDEs is developed using the finite element framework FEniCSx. The goal is a framework where new methods can be rapidly tested before time-consuming development into codes like ADCIRC. We validate the framework on several test cases, including large-scale computations in the Gulf of Mexico for Hurricane Ike (2008).

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

Mark Loveland
US Army Engineer Research and Development Center
markloveland@utexas.edu

Benjamin Pachev
Oden Institute - The University of Texas at Austin
benjaminpachev@utexas.edu

Jennifer Proft
University of Texas at Austin
jennifer.proft@austin.utexas.edu

Eirik Valseth
Oden Institute, UT Austin
eirik@utexas.edu

MS244

Tsunami Modeling with Hyperbolic-Dispersive Systems in GeoClaw

The nonlinear shallow water equations are one of the most popular depth-averaged water wave systems for tsunami modeling. Consequently, great efforts have been made to develop large-scale, robust, and efficient numerical schemes for this system. GeoClaw is a collection of algorithms for the discretization of shallow flows over complex topographies with a wave-propagation formulation of Godunov-type methods and adaptive mesh refinement. The shallow water equations neglect dispersive effects, which may be important in modeling tsunamis that involve shorter-wavelength perturbations. On the other hand, dispersive water wave models lack an inherent wave-breaking mechanism and are often formulated as systems of PDEs with high-order and mixed time-space derivatives, which requires the implicit inversion of differential operators. In recent years, hyperbolic approximations of such systems have been proposed as potentially advantageous, mainly due to the availability of explicit shock-capturing numerical schemes for such first-order systems. This work is concerned with the implementation, in GeoClaw, of a hybrid solver that transitions from hyperbolic-dispersive systems to the shallow water equations as required by the conditions of the problem. We evaluate the performance of this solver with standard benchmarking tests and real tsunami data, and run simulations of hypothetical scenarios.

Carlos Muñoz
King Abdullah University of Science and Technology (KAUST)
carlos.munozmoncayo@kaust.edu.sa

Marsha Berger
Flatiron Institute
mberger@flatironinstitute.org

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

Kyle Mandli
Center for Computational Mathematics, Flatiron Institute
kyle.mandli@gmail.com

MS244

Static and Moving Cut Mesh Methods for Flood Modeling

In the context of flood modeling, cut meshes can be used for a variety of applications. At their simplest, cut meshes allow complex domain geometries to be represented well on relatively coarse background meshes. When paired with discontinuous Galerkin or finite volume methods, cut meshes also allow less intrusive treatment of sharp features in bathymetry and the ability to capture moving wet-dry interfaces and moving domains. However, such uses for cut meshes present their own challenges, e.g. the small cell problem compounded by interface tracking. Here we present our recent work developing advanced cut mesh methods for use with the shallow water equations.

Christina Taylor
University of Texas at Austin
christina.taylor@austin.utexas.edu

MS245

Jet: A New Graph Partitioner for GPUs

We present a new graph partitioner, Jet, that runs on both CPU and GPU. Jet is a high-quality multilevel graph partitioner. Most current graph partitioners (such as Metis) only run on CPU. It is quite challenging to port partitioners to GPU due to the highly irregular memory access and need for fine-grained operations. Jet uses novel, efficient GPU algorithms for both coarsening and refinement. We show that Jet also provides superior quality partitioning on a wide variety of graphs and meshes, compared to current serial and parallel partitioners. Finally we discuss ongoing work to extend Jet to distributed memory for multi-GPU partitioning.

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Michael Gilbert
Pennsylvania State University, U.S.
msg5334@psu.edu

Kamesh Madduri
Pennsylvania State University
madduri@psu.edu

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

MS245

Pcms: A Geometry and Discretization Aware Coupling Tool for Fusion Devices

Coupling of fusion device codes present unique challenges in the scale of computations, range of coordinate systems, high dimensionality of phase space, and geometric complexity. These challenges require new approaches that enable efficient coupling on exascale supercomputers and afford adherence to physical constraints such as conservation that are required for stable coupling schemes. In this talk, we present developments in PCMS to support geometry informed coupling with geometry that is derived from parameterized CAD models and magnetic geometry such as geqdsk, VMEC, and DESC. We also present progress towards five-dimensional mesh-based conservative field transfer methods that will enable distribution-function based

coupling of fusion-plasma codes. Lastly, we indicate a set of future PCMS developments towards integration of physics and engineering systems. This research was supported by the U.S. Department of Energy Office of Science FES and ASCR through four SciDAC-5 Partnership Centers (1) StellFoundry: High-fidelity Digital Models for Fusion Pilot Plant Design (DE-AC02-09CH11466), (2) HiFiStell: High-Fidelity Simulations for Stellarators (DE-SC0024548), (3) Computational Evaluation and Design of Actuators for Core-Edge Integration (CEDA) (DE-AC02-09CH11466), (4) Center for Advanced Simulation of RF Plasma Material Interactions, and the FastMath SciDAC institute (DE-SC0021285).

Jacob S. Merson, Abhiyan Paudel, Fuad Hasan, Matthew McCall
Rensselaer Polytechnic Institute
mersoj2@rpi.edu, paudea@rpi.edu, hasanm4@rpi.edu, mccalm@rpi.edu

Cameron W. Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Angel Castillo-Crooke
Rensselaer Polytechnic Institute
castia5@rpi.edu

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS245

Reconstruction of CAD Models into Single Trimmed Surfaces

Classical methods by which computer-aided design (CAD) geometries are represented for both design and analysis involve meshing. This poses problems, as the meshing process takes a significant amount of time and labor, results in an approximate representation of the geometry of interest, and typically generates models capable only of low-order analyses. This research focuses on accurately and efficiently rebuilding given CAD surfaces or meshes, such as shell-like components of aeronautical structures or automotive vehicles, into single trimmed surfaces (rather than a piecewise-linear approximation), making them suitable for use in isogeometric analysis. Our proposed method simply and accurately represents these structures by converting a CAD geometry into a feature-aware triangulation, which is then flattened using theory from conformal or differential geometry to define the parametric domain and trimming curves of the intended spline. Subsequently, the bijection between this flattened geometry and its original spatial representation is then used to inform a mapping of a B-spline back into the spatial domain. Once the spatial surface is achieved, it is trimmed using the boundary of the original geometry, thereby reconstructing the intended geometry as a single trimmed B-spline surface. Resulting spline geometries are ultimately used for structural engineering analysis to demonstrate their suitability for analysis.

Rebecca Rasmussen, Kendrick Shepherd
Brigham Young University
becca_ras@byu.edu, kendrick_shepherd@byu.edu

MS245

High-Order Mesh Adaptivity Using Goal-Oriented Error Estimation

We explore a new approach to goal-oriented mesh optimization in PDE-driven computational simulations. Targeting high-order mesh adaptivity, the approach combines geometric quality optimization with control of the PDE solution error. By leveraging the Target-Matrix Optimization Paradigm (TMOP) in conjunction with adjoint sensitivity analysis, the method employs a goal-oriented error estimation framework to optimize the mesh adaptively. We demonstrate that the optimized mesh simultaneously achieves good geometric quality and reduces the PDE solution error in regions critical to a predefined computational objective.

Vladimir Z. Tomov, Ketan Mittal, Mathias Schmidt
Lawrence Livermore National Laboratory
tomov2@llnl.gov, mittal3@llnl.gov, schmidt43@llnl.gov

MS245

Developing Particle Simulation Code in Fusion Plasma Physics and Gas Dynamics Using Unstructured Mesh Infrastructure and PUMIPic

In this presentation, we report our work developing two particle simulation code, XGCm and Comet. Both XGCm and Comet work on either CPU or GPU accelerators. XGCm is a gyrokinetic particle-in-cell code for fusion plasma physics, modeling the turbulent transport in tokamak device. Comet is a direct simulation Monte Carlo code for modeling and simulating hypersonic and rarefied gas flows. Both of them use some type of particle method rooted in the same kinetic modeling framework. Central to them are the unstructured mesh infrastructure Omega_h, and mesh-centric particle data structure PUMIPic. Omega_h provides the functionalities to perform distributed mesh operations on GPUs, while PUMIPic ensures that particle operations can be performed efficiently on GPUs. Several additional open-source libraries are employed in the code to interact with GPUs, including Kokkos, Cabana, and PETSc. Firstly, numerical methods and algorithms to perform different particle operations on GPUs are discussed in connection with PUMIPic particle data structure. Secondly, validation studies are discussed. The circular geometry cyclone case is used in XGCm, where the calculated turbulence growth rate shows excellent agreement with existing simulation results. Comet simulation results of hypersonic rarefied gas flows are compared with existing experimental data. Finally, the parallel performances of the two code are shown, including strong and weak scaling results.

Chonglin Zhang
University of North Dakota
chonglin.zhang@und.edu

Gerrett Diamond
MathWorks
diamog320@gmail.com

Cameron W. Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS246

Tensor Networks for Turbulent Combustion Models

Computational fluid dynamics (CFD) simulations using finite-rate chemistry (FRC) and detailed chemical mechanisms provide high-fidelity representations of complex reactions but are computationally intensive, especially for mechanisms involving thousands of reactions and species. The stiffness of chemical source terms, due to a wide range of reaction time scales, further complicates this by leading to numerical instability. While flamelet-based combustion models offer a cost-effective alternative, they require higher-dimensional manifolds to accurately capture the physics of high-speed propulsion systems, which increases computational complexity exponentially, a problem known as the curse of dimensionality. In our previous work, we introduced tensor trains (TTs) as a low-rank, efficient alternative to artificial neural networks (ANNs) for representing high-dimensional flamelet manifolds, particularly in hypersonic propulsion. The TT approach decomposes multi-dimensional flamelet data into low-rank tensors, significantly reducing computational complexity from exponential to linear scaling. Implementing the TT framework within the JENRE Multiphysics Framework led to a two-order-of-magnitude reduction in memory requirements for a 5D flamelet table, while maintaining accuracy in 2D CFD simulations. In this study, we extend the TT framework to 3D reacting CFD simulations, demonstrating its potential to advance hypersonic propulsion technologies.

Sinan Demir
Argonne National Laboratory
sdemir@anl.gov

Pierson T. Guthrey
Lawrence Livermore National Laboratory
guthrey1@llnl.gov

Ryan Johnson, Brian Bojko
US Naval Research Laboratory
ryan.f.johnson4.civ@us.navy.mil,
brian.t.bojko.civ@us.navy.mil

MS246

Novel Performance-Portable Tensor Network Algorithms

Tensor network algorithms, which are pivotal in quantum many-body physics, machine learning, data compression, and pde solvers have been exploding in popularity. Of particular interest is the use of tensor network methods to solve high-dimensional time-dependent pdes, such as the Vlasov equations, which describe the evolution of the distribution function of a plasma in phase space. We discuss and benchmark various algorithmic building blocks for solving the Vlasov equations via tensor networks on GPUs. Our goal is to use these kernels to address the challenges of high-dimensional phase space discretizations and the need for efficient linear solver algorithms. Further, we introduce a framework where this can be done in a performance-portable way, ensuring that these algorithms are performant on all architectures, from laptops to next-generation architectures.

Pierson T. Guthrey
Lawrence Livermore National Laboratory
guthrey1@llnl.gov

MS246

Tensor Network Algorithms for High-Order Correlations in Statistical Mechanics

Modelling the energy transport and equation of state of confined liquids and shocked plasmas in the presence of large density gradients requires an accurate understanding of high dimensional particle correlation functions. Typical solution schemes attack this problem through an assumption that the liquid is slowly varying, which significantly reduces the dimensionality of the solution space. Tensor network algorithms offer an alternative paradigm, where a low rank tensor train representation of the high dimensional functions allows for an efficient computation of the fully heterogeneous density correlation functions. We investigate the efficacy and efficiency of such an approach, showing results from the solution of highly heterogeneous plasmas.

Zach Johnson
Michigan State University
john8248@msu.edu

MS246

First Efforts at Tensor-Train for Space-Time DPG Problems

Plasma physics problems are typically simulated using either hydrodynamics models or particle-in-cell (PIC) models. In some physics regimes of interest for pulsed-power simulation, the former introduces unacceptable levels of model-form error, while the latter requires an infeasible number of macro-particles for accurate simulation. For such regimes, direct discretization of the Vlasov equations is appropriate. However, the full Vlasov model requires three spatial and three velocity dimensions, plus a time dimension, for a total of seven discretization dimensions. The curse of dimensionality thus leads to exhaustion of computational resources: the number of degrees of freedom needed to resolve the scale of the problem often exceeds what is feasible on the largest supercomputers. We seek to mitigate the curse of dimensionality in Vlasov models by combining two discretization technologies: the discontinuous Petrov-Galerkin (DPG) finite element methodology and Tensor-Train compression techniques. DPG is a *minimum-residual* method, with mechanisms to control the norm in which the residual minimized, and natural mechanisms for estimating error in that norm. Tensor-Train methods, meanwhile, allow compression of PDE operators and solution representations in a *tunable* fashion. This work is part of a new collaboration among Sandia, Lawrence Livermore, and Los Alamos National Laboratories, and in early stages of development. We present our progress to date.

Nathan V. Roberts
Sandia National Laboratories
nvrober@sandia.gov

MS247

Enhancing X-Ray Tomography Through Sequential Experimental Design and Uncertainty Quantifica-

tion

This study introduces an advanced methodology for enhancing X-ray tomography by integrating Bayesian modeling and sequential experimental design. A Bayesian model is constructed to quantify uncertainty associated with measurement processes, incorporating factors such as beam shifts and other sources of variability. An adaptive algorithm is utilized to guide the selection of subsequent measurement actions, aiming to optimize information acquisition. This approach involves forecasting the outcomes of the subsequent N measurements to refine the experimental design in real-time. By leveraging uncertainty quantification within a sequential framework, the methodology seeks to improve imaging precision and reliability in X-ray tomography.

San Dinh
Argonne National Laboratory
sdinh@anl.gov

Zichao Wendy Di
Mathematics and Computer Science Division
Argonne National Laboratory
wendydi@mcs.anl.gov

Matt Menickelly
Argonne National Laboratory
mmenickelly@anl.gov

MS247

Next-Generation Gaussian Processes for Function Approximation, Uncertainty Quantification, and Decision-Making

Gaussian processes (GPs) and Gaussian-related stochastic processes are powerful tools for function approximation, uncertainty quantification, global optimization, and autonomous data acquisition due to their robustness, analytical tractability, and natural inclusion of Bayesian uncertainty estimates. Even so, Gaussian processes are often criticized for poor approximation performance and neck-breaking computational costs in real-life applications. The reason for this gap, however, is not the methodology itself but rather a user-caused lack of flexibility and domain awareness of the underlying prior probability distribution, the likelihood, poor training, and insufficiently expressive utilization of the posterior for decision-making. In this talk, I want to address many challenges of GPs that might inhibit optimal performance, guiding the audience to Next-Generation Gaussian Processes. We will begin our journey by looking at a few motivating examples and some theoretical preliminaries before diving into commonly encountered challenges and what can be done to address them. We will discover hands-on solutions and tools that make implementation easy. The key takeaway for the audience will be a better understanding of Gaussian processes and ways to customize them for optimal performance.

Marcus M. Noack
Lawrence Berkeley National Laboratory
MarcusNoack@lbl.gov

MS247

Accounting for Probabilistic Uncertainty in Continuous Lyophilization via Polynomial Chaos Expansion

Lyophilization, also known as freeze drying, is a process that is commonly used to increase the stability of various pharmaceutical products in pharmaceutical manufacturing, e.g., mRNA vaccines, allowing for higher storage temperature. While the current trends in the pharmaceutical industry are moving towards continuous manufacturing, the majority of industrial lyophilization processes are still being operated in a batch mode. Typical lyophilization is subject to several uncertainties, such as product quality, operating conditions, and stochasticity in the process (e.g., ice nucleation). Therefore, the development and operation of continuous lyophilization should be able to account for these uncertainties efficiently, ideally in real time. This work presents a framework to account for the probabilistic uncertainty in continuous lyophilization via polynomial chaos expansion (PCE), a systematic computationally efficient approach to represent random variables with a set of orthogonal polynomials. First, a mechanistic model is developed for the continuous lyophilization of suspended vials, the state-of-the-art continuous lyophilization technology. Then, the non-intrusive PCE technique is applied to the model with uncertain parameters. Finally, case studies associated with uncertainty quantification and model-based optimization and control using our PCE-based model are discussed, with comparison to the traditional Monte Carlo approach.

Prakitr Srisuma, George Barbastathis
Massachusetts Institute of Technology
prakitrs@mit.edu, gbarb@mit.edu

Richard D. Braatz
Massachusetts Institute of Technology
Department of Chemical Engineering
braatz@mit.edu

MS247

Posterior Covariance Structures in Gaussian Processes

In this talk, we present a comprehensive analysis of the posterior covariance field in Gaussian processes, with applications to the posterior covariance matrix. Our geometric analysis reveals how the Gaussian kernel's bandwidth parameter and the spatial distribution of the observations influence the posterior covariance as well as the corresponding covariance matrix, enabling straightforward identification of areas with high or low covariance in magnitude. Drawing inspiration from the a posteriori error estimation techniques in adaptive finite element methods, we also propose several estimators to efficiently measure the absolute posterior covariance field, which can be used for efficient covariance matrix approximation and preconditioning.

Yuanzhe Xi
Emory University
yxi26@emory.edu

Difeng Cai
Southern Methodist University
ddcai@smu.edu

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS248

Predictive Adaptive Mesh Refinement Using Gaus-

sian Process Regression

Flow and transport PDEs can elicit solutions with strong gradients or wave-like behavior that propagate in time. Signals that initially occur in the interior or at the boundary are not immediately experienced by all points in the domain. AMR is an efficient numerical technique for dynamically increasing resolution in areas of the domain where it is needed, while leaving the rest of the domain at a coarser resolution. AMR has been traditionally based on a deterministic approach to initiate the refinement process, making use of thresholds on solution gradient data such as vorticity or a concentration front. It would be an advantage to replace the AMR criterion with a predictive approach based on AI/ML methods to make use of stochastic approximation of real time data in a simulation, requiring less user interaction and knowledge. We propose two approaches in this study: function thresholding and truncation error thresholding. In function thresholding, grid points with dependent values or derivatives that exceed a user defined threshold are refined. In truncation error thresholding, grid points that possess a truncation error estimate that exceeds a threshold are refined. We have trained a Gaussian Process Regression model on 2D flows to learn the coordinate origin, width, and height of refinement regions. Results using non-stationary and stationary kernels are presented, showing inferred and deterministic refinement regions, and Intersection-over-Union to characterize accuracy.

Christopher Paolini, Jana Gonzalez
San Diego State University
paolini@engineering.sdsu.edu, jgonzalez2264@sdsu.edu

David Trebotich
Lawrence Berkeley National Laboratory
dptrebotich@lbl.gov

MS248

Tensorial Acousto-elastodynamics for Seismic Applications in Complex Media: A Mimetic Finite Difference Approach

Seismic imaging and inversion scenarios are often characterized by complex geometries that require non-Cartesian computational solution domains. Most existing finite-difference approaches used in seismology are developed for Cartesian meshes, which are ineffective at representing curvilinear interfaces and accurately modeling the full elastic wavefield interactions with the irregular free-surface domain boundary. Developing efficient and stable numerical approaches that accurately mimics the underlying physical and mathematical properties of acousto- and elastodynamics (e.g., conservation laws, symmetry, duality) and its associated boundary conditions is a challenging task. In this work we demonstrate that these issues largely can be avoided in earth imaging problems by computing solutions of elastodynamics in a generalized coordinate system that is conformal to interfaces of interest. We develop a mimetic finite-difference time-domain (MFDTD) tensorial elastodynamics solver that uses coordinate mappings to transform an irregular physical domain to a regular computational grid.This numerical solver represents the key computational kernel for several challenging full-wavefield propagation simulations including anisotropic elastic wave simulation from irregular topography surfaces and modeling acoustic wave simulation for a moving hydroacoustic source in the presence of a time-varying free surface.

Jeffrey Shragge
Colorado School of Mines
jshragge@mines.edu

MS248

Data-Driven Mimetic Framework for Turbulence Modeling

It is well recognized that the modeling of turbulent fluid flows poses several numerical challenges. The non-linear convective term of the Navier Stokes equation requires special care when performing the spatial discretization. Spatial operators that satisfy the discrete version of the continuous integration by parts property result in accurate discretizations for the convective term. Mimetic difference schemes are based on obtaining discrete high order operators that are equivalent to the tensor calculus quantities of divergence and gradient. They possess a discrete equivalent of a global conservation law (i.e., the extended Gauss divergence theorem), thereby leading to accurate discretizations of the nonlinear convective terms. In this work, we investigate the use of the high order mimetic operators in developing ground truth data as applied to the Burgers equation. This ground truth is then supplied as inputs to a neural network as velocity and gradients of velocity, and the network is trained to obtain a machine-learned subgrid scale (SGS) turbulence model. The performance of this machine-learned mimetic based model is compared with the prevalent SGS models such as the static dynamic Smagorinsky models.

Anand Srinivasan
University of California, Irvine
anands3@uci.edu

MS248

Isomorphic Structures and Operator Analysis in Mimetic Discretizations

This study presents a comprehensive examination of the structural and operatorial foundations within mimetic discretizations, with a focus on bridging the gap between discrete and continuous function spaces. By scrutinizing the mimetic gradient and divergence operatorscentral to the discretization of the NAVIER-STOKES equationswe study their kernel and image spaces, establishing their isomorphisms through rigorous mathematical proofs. Our methodology leverages discrete scalar and vector function spaces, delineated by grid spacing, to define linear mappings that unveil the subspace relationships and quotient space structures integral to understanding these operators' roles in computational fluid dynamics. Central to our findings is the application of the first isomorphism theorem, which facilitates a deeper insight into how mimetic discretizations reflect the continuous properties of differential operators within a discrete framework. This allows for an exploration into the algebraic and topological implications of such discretizations, notably in the context of the NAVIER-STOKES equations. Furthermore, we extend our investigation to encompass subalgebras, ideals, their quotients, and the formulation of short exact sequences that mirror the continuous interplay between gradient, divergence, and LAPLACIAN operators.

J. de Curto
BARCELONA Supercomputing Center
jdecurto@icai.comillas.edu

Irene de Zarza
Universidad Francisco de Vitoria /
Universitat Oberta de Catalunya

dezarza@uoc.edu

MS249

Learning Symmetry-Preserving Closure Models for Large Eddy Simulation

Symmetries are among the fundamental properties of physics and partial differential equations [Frisch 1995]. The Navier-Stokes equations are invariant to symmetry groups such as translations, rotations, Galilean transformations, and scaling. These fundamental symmetries are further related to physical properties such as energy-conservation. It is therefore desirable that simplified models and numerical discretizations of the Navier-Stokes equations also respect these symmetries. Symmetry-preserving discretizations have successfully been used for the incompressible Navier-Stokes equations [Verstappen 2003]. Large eddy simulation (LES) aims to resolve large scales of turbulent flows only, to reduce the computational cost. This requires choosing a *closure model*. Recently, deep learning has been used to learn closure models. This requires training data, which is discrete by nature. We showed that discretizing the equations first, before filtering and applying a closure model, removes model-data inconsistencies and improves stability [Agdestein 2024]. In this work, we extend this framework with symmetry-preserving neural network closure models. We use group-equivariant steerable layers [Weiler 2023] to enforce symmetries, making the entire discrete LES model physically consistent and stable.

Syver D. Agdestein
Centrum Wiskunde & Informatica
sda@cwi.nl

MS249

A Learning-Based Multiscale Model for Underground Fluid Transport Processes

Transport of water through permeable geological formations couples various phenomena. As the water flows through the permeable medium, water reacts with the medium changing the morphology, mechanical properties, and permeability of the medium; this in turn affects the flow and chemistry. These varied phenomena occur at various length and time scales, which presents a complex multiscale challenge. Traditional multiscale modeling approaches have been developed for particular processes to pass the information across scales. However, the actual implementation of this technique is prohibitively expensive due to the repeated solution of partial differential equations at the finer scales. We present a methodology to overcome this challenge by creating a high-fidelity, computationally efficient surrogate of the lower scale behavior that can directly be used at the upper scale to prevent repeatedly solving equations. This framework uses one-time off-line data generated from the pore-scale simulations to train the solution to the partial differential equations over a neural network and obtain a learned surrogate. The surrogate is an inexpensive, approximate solution to the lower-scale problem, which can be used to solve the macroscopic problem without further modeling. We focus on the transport of flow through the porous medium in the presence of chemical reactions, where the convection-diffusion equations are complicated due to fluid-rock interactions and evolving microstructure.

Mina Karimi
California Institute of Technology
mkarimi@caltech.edu

Kaushik Bhattacharya
Howell N. Tyson Sr. Professor of Mechanics
California Institute of Technology
bhatta@caltech.edu

MS249

Bayesian Inference for Patient-Specific Mechanical Models of Tumor Growth

Equipping predictive models with quantitative uncertainties is critical for establishing trust in computational models and enabling robust decision making for personalized medicine. In particular, mechanistic modeling of high-grade gliomas proves challenging due to intra-tumoral heterogeneity at the cellular, molecular, and dynamic scales. The difficulty in resolving heterogeneity in the tumor microenvironment translates to challenges in the diagnosis and management of the disease. We present a methodology that utilizes non-invasive quantitative magnetic resonance imaging (MRI) data to calibrate the spatially varying parameters of a reaction-diffusion partial differential equation (PDE) describing tumor invasion and proliferation. We generate a three-dimensional patient-specific computational domain from the first MRI scan in the time course and estimate time courses of the tumor extent and cellular density. Further, we couple the reaction-diffusion PDE to a linear elastic model to account for the effect of tissue stress response on tumor infiltration. A Bayesian formulation of the inverse problem is developed to account for uncertainty arising from the infrequently collected and indirectly informative MRI data. We analyze the ill-posed nature of the problem and discuss model error and selection of appropriate priors. We demonstrate the scalability of the approach with a large-scale example incorporating publicly available MRI data and a realistic three-dimensional geometry.

Graham T. Pash
University of Texas at Austin, US
gtpash@utexas.edu

Umberto Villa
The University of Texas at Austin
uvilla@austin.utexas.edu

David A. Hormuth II
Oden Institute for Computational Engineering and Sciences
Livestrong Cancer Institutes
david.hormuth@austin.utexas.edu

Thomas E. Yankeelov
University of Texas at Austin
Livestrong Cancer Institutes
thomas.yankeelov@utexas.edu

Karen E. Willcox
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
kwillcox@oden.utexas.edu

MS250

Hybrid Neural Differentiable Modeling for Spatiotemporal Wall-Bounded Turbulence

Efficient simulation of complex turbulence is crucial for numerous engineering applications, particularly in accurately reconstructing near-wall flow features in wall-bounded tur-

bulence. Traditional approaches often rely on subgrid-scale (SGS) closure and wall models, which approximate small-scale features but frequently compromise accuracy in complex flow scenarios. This study introduces a novel neural differentiable modeling framework to enhance both the predictability and efficiency of wall-bounded turbulence simulations. Our approach seamlessly integrates deep neural networks (DNNs) with physics-based numerical solvers within a differentiable programming framework, leveraging the flexibility of DNNs and the generalizability of physics-based models. The proposed neural solver incorporates joint SGS-wall models, accurately representing equilibrium/non-equilibrium effects while addressing discrepancies in numerical discretizations. Validation across multiple scenarios, including channel flow and turbulent boundary layers, demonstrates that the framework substantially enhances the fidelity and computational performance of turbulence simulations. This work sets a new benchmark in CFD by highlighting the transformative potential of differentiable programming in advancing the simulation of complex turbulent flows.

Xiantao Fan, Xin-Yang Liu, Meet Parikh, Jian-Xun Wang
University of Notre Dame
fanxiantao1@gmail.com, xliu28@nd.edu, mparikh2@nd.edu, jwang33@nd.edu

MS250

Robust and Interpretable Data-Driven Turbulence Models with a Focus on Hypersonic Aerodynamics

We present a data-driven turbulence modeling approach for compressible flows. Our approach leverages tensor-basis neural networks and a neural-network-based variable turbulent Prandtl number to improve closure relationships. To improve the robustness of the formulation, we employ a calibrated deep ensemble formulation and Lipschitz continuous neural networks. We train a global data-driven turbulence model, which builds upon a nominal k-epsilon model, and present results across a suite of low-speed and hypersonic test cases. For hypersonic boundary layers, the inclusion of the wall-normal Reynolds stress has an appreciable impact on the wall-normal momentum balance and wall quantities. For shock boundary layer interactions and compression ramps, the use of a variable Prandtl number reduces the over-prediction in wall heating at and downstream of the reattachment point.

Eric Parish
Sandia National Lab
eric.parish9@gmail.com; ejparis@sandia.gov

Matthew Barone
Sandia National Laboratories
mbarone@sandia.gov

David Ching, Cyrus Jordan, Nathan Miller, Steven Beresh
Sandia National Lab
dching@sandia.gov, cjorda@sandia.gov, nmille1@sandia.gov, sjberes@sandia.gov

MS250

A Holistic View of How UQ Fits into the Credibility of Scientific Computing Including Machine Learning

In any investigation in scientific computing uncertainty quantification (UQ) is an essential activity. UQ is applied to classical applications and recently to Scientific Machine Learning (SciML). A broader challenge for scientific computing is credibility and trust for results. The set of activities seen as UQ is only a piece of a broader tapestry of actions providing confidence in results. This confidence is primarily driven by comparison with objective reality seen in experiments and observations. This comparison is known as validation determining if models describe that reality and how well. A process known as verification tells us if a model is properly represented computationally. Error and uncertainty is the medium of proof in both activities. Together these processes coupled with UQ and a handful of other steps provide an evidence basis for credibility and trust in results. This rubric is well developed and accepted for classical scientific computing. For SciML the rubric must change to match the manner and quality steps in those methodologies. Developing trust in SciML is an essential step to integrate it with the broader community.

William J. Rider
Sandia National Laboratories
wjrider@sandia.gov

MS251

Towards Dynamic Amr on Mixed-Element Unstructured Meshes With Conforming and Non-Conforming Interfaces

Adaptive Mesh Refinement's role in accurately simulating unsteady multi-scale fluid dynamics phenomena is rapidly expanding due to its ability to position mesh resources in locations of critical flow regions. While AMR on unstructured, mixed element-type meshes has presented significant challenges in terms of algorithmic and software infrastructure, as well as flow-solver interfacing, we will discuss our approach to coupling legacy Computational Fluid Dynamics solvers with the t8code framework in the non-conforming mesh setting and the development of an automatic conforming face interface.

Andrew Kirby
University of Wyoming
akirby@uwyo.edu

Dimitri Mavriplis
Scientific Simulations LLC, U.S.
mavripl@scientific-sims.com

MS251

A GPU-Accelerated AMR Framework for Compressible Combustion/Detonation Simulation

Our aim is to perform three-dimensional large-scale compressible combustion/detonation computation. We describe our experiences in developing a flexible framework for high-efficiency and -fidelity simulation of complex non-reactive/reactive phenomena by leveraging the AMReX platform. A variety of testing examples are performed to verify the accuracy and efficiency of our framework. We focus especially on the algorithms of grid-wise operations, portability of the codes from CPU to CPU/GPU heterogeneous system and the implementation of a GPU-accelerated chemistry solver.

Yuqi Wang
National University of Defense Technology, China
wangyuqi@nudt.edu.cn

Yadong Zeng
University of Minnesota

University of Minnesota
zengx372@umn.edu

Jianhan Liang
Science and Technology on Scramjet Laboratory, National Univ
jhleon@vip.sina.cn

MS252

Div-conforming Unfitted Finite Element Methods

In this work, we will present unfitted finite element formulations for the numerical approximation of problems in $H(\mathrm{div})$, e.g., the Darcy equations for flow in porous media. In particular, we will consider a mixed finite element discretisations (e.g., Raviart-Thomas and BDM elements) on a background mesh and properly chosen stabilisation terms to make the approximation stable and have a well-posed system matrix. The stabilisation terms will rely on bulk ghost-penalty penalties on cell aggregates, including (at least) one interior cell. The scheme will be designed so that the unfitted method inherits the exact divergence-free property of this finite element approximation in the body-fitted case. The proposed formulation's numerical analysis and numerical experimentation will also be presented.

Santiago Badia
Monash University
santiago.badia@monash.edu

Anne M. Boschman
University of Nottingham
anne.boschman@nottingham.ac.uk

Alberto Martin
Australian National University
alberto.martin@monash.edu

Ricardo Ruiz-Baier
Monash University
ricardo.ruizbaier@monash.edu

MS252

Advances on the Shifted Boundary Method: Neumann and Contact Boundary Conditions, and Applications

Scientific computing is routinely assisting the design of systems, components, or materials with complex shapes, but it is often underestimated how the time and labor cost of mesh generation affects the overall analysis cycle. The situation is even more critical for high order discretizations. Methods that could ease these limitations could more effectively interface with meta-algorithms from Optimization, Uncertainty Quantification, Reduced Order Modeling, Machine Learning, and Artificial Neural Networks. Embedded/immersed/unfitted boundary methods seem good candidates for this purpose, because they obviate the need for body-fitted meshing, but may be difficult to implement, due to complex cell cutting operations at boundaries, which also negatively affect numerical stability and matrix conditioning. The Shifted Boundary Method is a simple, stable, and accurate embedded method, which eliminates the need of cut cells. Boundary conditions are imposed on a surrogate discrete boundary, lying in proximity of the true boundary interface, and composed of full edges/faces of the grid. Appropriate field extension operators are then constructed by way of Taylor expansions, with the purpose of imposing accurate boundary conditions on the surrogate boundary that mimic the effect of the true boundary conditions. The performance of the SBM will be demonstrated on large-scale problems with particular emphasis on the treatment of Neumann and contact boundary conditions.

Jason Haydel Collins
Duke University
jason.collins@duke.edu

Kangan Li
Pennsylvania State University
kbl5610@psu.edu

Alexei Lozinski
Laboratoire de Mathématiques de Besançon
CNRS et Université de Franche-Comté
alexei.lozinski@univ-fcomte.fr

Guglielmo Scovazzi
Duke University
guglielmo.scovazzi@duke.edu

MS252

Matrix-Free Evaluation of Cutfem Operators

In CutFEM, the grid does not need to conform to the domain, allowing for the use of computationally convenient Cartesian grids. On these grids, the tensor product structure of the operators can be exploited, significantly improving computational efficiency. Unlike traditional FEM, where sparse matrix computations are typically memory-bound due to storage and access patterns, matrix-free approaches are compute-bound, resulting in faster performance. The evaluation of the finite element operator in CutFEM can be divided into three parts. First, the operator on uncut cells is handled similarly to standard FEM, where the use of tensor product structure allows for substantial speedup in cell-wise operations. Second, the contributions of cut cells must be evaluated. As the number of these cells decreases with mesh refinement, it becomes feasible to store the matrices for them. Finally, special measures are necessary to ensure the well-posedness of the formulation regardless of the cut location. This is achieved through a ghost penalty method, where the jumps of all derivatives across the faces of cut cells are summed. By exploiting the tensor product structure, this stabilization term can be evaluated efficiently, requiring only two 1D matrices for the entire mesh while preserving the computational benefits of tensor product operators.

Michal Wichrowski
Heidelberg University
wichrowski@uni-heidelberg.de

MS253

Eliminating Artificial Boundary Conditions in Time-Dependent Density Functional Theory Using Fourier Contour Deformation

I will introduce a method for solving the time-dependent Schrdinger equation in free space with a compactly-supported time-dependent potential. The method avoids artificial boundary conditions using a Fourier pseudospectral-type method, but with the spatial Fourier variable deformed to a complex contour. This approach yields the free space solution restricted to a computational subdomain, with controlled accuracy. The scheme is high-order accurate in space and time, has quasi-optimal com-

putational complexity, operates on Cartesian spatial grids in any dimension, permits spatially-uniform applied fields, and allows solutions to leave the computational domain and return later. I will present results from simulations of absorption and photoemission spectroscopy using time-dependent density functional theory, demonstrating recovery of accurate spectra from simulations on small computational domains.

Jason Kaye, Alex H. Barnett
Flatiron Institute, Simons Foundation
jkaye@flatironinstitute.org, abarnett@flatironinstitute.org

Leslie Greengard
Simons Foundation and Courant Institute
New York University
lgreengard@flatironinstitute.org

Umberto de Giovannini
Università degli studi di Palermo
umberto.de-giovannini@mpsd.mpg.de

Angel Rubio
Flatiron Institute, Simons Foundation
arubio@flatironinstitute.org

MS253

Virtual Beamline++: Comprehensive Nonlinear Optical Modeling for High Energy Lasers

The field of high energy, high power, and high average repetition rate lasers requires sophisticated modern modeling tools to predict and understand beam performance. Virtual Beamline++ (VBL++) has been developed to aid in the design and operation of future laser systems. This model couples numerical wave propagation to numerical models of nonlinear phenomena through various split and multi-step algorithms. This talk will show how turn-key models can be coupled to cutting-edge multi-step numerical algorithms to guide the operation and design of large-scale laser systems. This talk will discuss both how the model is utilized in the current environment and plans for its development in the future. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-ABS-868715.

Samuel Mclaren
Lawrence Livermore National Laboratory
mclaren3@llnl.gov

MS253

Boundary Conditions for the Einstein Field Equations

A brief overview regarding the initial-boundary value problem for Einstein's field equations with radiative-type boundary conditions will be provided.

Olivier Sarbach
Universidad Michoacana de San Nicolás de Hidalgo, Mexico
olivier.sarbach@umich.mx

MS254

Parametric Model Reduction of Mean-field and Stochastic Systems on the Space of Probability Densities

The aim of this work is to learn models of population dynamics of physical systems featuring stochastic and mean-field effects and a dependence on physics parameters. The learned models can act as surrogates of classical numerical models to efficiently predict the system behavior. By population dynamics we refer to the evolution of the ensemble of samples that represent the system as opposed to the trajectories of individual samples. Crucially, different sample dynamics can collectively give rise to the same population dynamics. We aim to exploit this redundancy to achieve a reduction of complexity in the model. Building on the Benamou-Brenier formula from optimal transport, we use a variational problem to infer parameter- and time-dependent gradient fields that represent approximations of the population dynamics. The inferred gradient fields can then be used to rapidly generate sample trajectories that mimic the dynamics of the physical system on a population level over varying physics parameters. We demonstrate on several mean-field-type, high-dimensional, and chaotic particle systems that our approach accurately predicts population dynamics over a wide range of parameters and outperforms state-of-the-art diffusion-based and flow-based models that simply condition on time and physics parameters.

Jules Berman
NYU
jmb1174@nyu.edu

Tobias Blickhan, Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
tobias.blickhan@nyu.edu, pehersto@cims.nyu.edu

MS254

Latent-EnSF: A Latent Ensemble Score Filter for High-Dimensional Data Assimilation with Sparse Observation Data

Accurate modeling and prediction of complex physical systems often rely on data assimilation techniques to correct errors inherent in model simulations. Traditional methods like the Ensemble Kalman Filter (EnKF) and its variants as well as the recently developed Ensemble Score Filters (EnSF) face significant challenges when dealing with high-dimensional and nonlinear Bayesian filtering problems with sparse observations, which are ubiquitous in real-world applications. In this talk, we present a novel data assimilation method, Latent-EnSF, which leverages EnSF with efficient and consistent latent representations of the full states and sparse observations to address the joint challenges of high dimensionlity in states and high sparsity in observations for nonlinear Bayesian filtering. We introduce a coupled Variational Autoencoder (VAE) with two encoders to encode the full states and sparse observations in a consistent way guaranteed by a latent distribution matching and regularization as well as a consistent state reconstruction. With comparison to several methods, we demonstrate the higher accuracy, faster convergence, and higher efficiency of Latent-EnSF for two challenging applications with complex models in shallow water wave propagation and medium-range weather forecasting, for highly sparse observations in both space and time. See our paper at https://arxiv.org/abs/2409.00127

Phillip Si, Peng Chen
Georgia Institute of Technology

psi6@gatech.edu, pchen402@gatech.edu

MS254

Diffusion Models for Inverse Problems Done Right

Pretrained diffusion models (DMs) have recently been popularly used in solving inverse problems (IPs). The existing methods mostly interleave iterative steps in the reverse diffusion process and iterative steps to bring the iterates closer to satisfying the measurement constraint. However, such interleaving methods struggle to produce final results that look like natural objects of interest (i.e., manifold feasibility) and fit the measurement (i.e., measurement feasibility), especially for nonlinear IPs. Moreover, their capabilities to deal with noisy IPs with unknown types and levels of measurement noise are unknown. In this talk, we advocate viewing the reverse process in DMs as a function and propose a novel plug-in method for solving IPs using pretrained DMs, dubbed DMPlug. DMPlug addresses the issues of manifold feasibility and measurement feasibility in a principled manner, and also shows great potential for being robust to unknown types and levels of noise. Through extensive experiments across various IP tasks, including two linear and three nonlinear IPs, we demonstrate that DMPlug consistently outperforms state-of-the-art methods, often by large margins especially for nonlinear IPs.

Hengkang wang
University of Minnesota
wang9881@umn.edu

Xu Zhang
Amazon
spongezhang@gmail.com

Haihui Li, Xuxiang Wan
University of Minnesota
lixx5027@umn.edu, wan01530@umn.edu

Tiancong Chen
University of Minnesota, Twin Cities
Department of Computer Science and Engineering
chen6271@umn.edu

Ju Sun
Department of Computer Science and Engineering
University of Minnesota
jusun@umn.edu

MS254

Nonlinear Filtering in Stochastic Dynamical Systems

We present a nonlinear filtering approach for solving inverse problems sequentially in stochastic dynamical systems. By utilizing a Bayesian framework, we construct a generative modeling method to identify the state trajectories from noisy data in a robust manner.

Nicole Yang
Emory University
tianjiao.yang@emory.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

MS255

On the Performance of a Parallel and Randomized Single Value Decomposition Algorithm

The single value decomposition (SVD) is a core operation for modal decomposition techniques, which are widely used to reduce the complexity of high-dimensional mathematical models. pyLOM [Eiximeno et al., A HPC open source reduced order model suite for fluid dynamics applications, 2024] has recently unlocked the possibility of performing the SVD at a high-performance computing level. The parallel SVD from pyLOM is tailored for tall and skinny matrices as those obtained from computational fluid dynamics simulations. It has shown ideal scalability on a matrix of size $(5.83 \times 10^7, 56)$ when using up to 2700 Intel Sapphire Rapids. However, pyLOM currently struggles with the augment of columns in the matrix, due to the increase in the amount of memory needed. This work aims to upgrade the present pyLOM capabilities. To do so, the possibility of running a randomised SVD [Erichson et al., Randomized Matrix Decompositions using R, 2018] on top of the current parallel algorithm will be added. The first step of the new algorithm is performing a parallel randomization of the data. This reduces the number of columns of the matrix to the desired number of vectors of the new basis. Then, the parallel SVD from pyLOM can be used to compute the modes of the system on the randomized data matrix. The presentation will include a discussion on the implementation of the new algorithm as well as a detailed description of its parallel performance and the gains over the current version.

Benet Eiximeno
Barcelona Supercomputing Center
benet.eiximeno@bsc.es

Arnau Miró
Barcelona Supercomputing Center
Universitat Politècnica de Catalunya
arnau.miro@bsc.es

Nathan Kutz
University of Washington
kutz@uw.edu

Ivette Rodríguez
Universitat Politècnica de Catalunya
ivette.rodriguez@upc.edu

Oriol Lehmkuhl
BSC
oriol.lehmkuhl@bsc.es

MS255

Randomized Orthogonalization Techniques and Eigenvalue Solvers

In this talk we discuss recent progress in using randomization for solving eigenvalue problems. We discuss first numerically stable algorithms for orthogonalizing a set of vectors and their usage in the Arnoldi iteration. We then present associated Krylov subspace methods for solving eigenvalue problems as the implicitly restarted Arnoldi Method and discuss experimental results.

Laura Grigori

EPFL and PSI, Switzerland
laura.grigori@epfl.ch

Jean-Guillaume de Damas
Inria
jean-guillaume.de-damas@inria.fr

MS255

Randomized Preconditioned Cholesky-Qr

We analyze "rpCholesky-QR", a randomized preconditioned Cholesky-QR algorithm for computing the thin QR factorization of real mxn matrices with rank n. The perturbation analysis is transparent and identifies clearly all factors that contribute to error amplification; it requires only a minimal amount of assumptions and produces interpretable bounds, rather than first-order estimates. The numerical experiments demonstrate the accuracy of rpCholesky-QR for matrices that are highly ill-conditioned, have worst-case coherence, and can have a large number of columns, up to n=m/3. A sampling amount of 3n is sufficient, even under these worst-case conditions. The two-norm deviation from orthonormality of rpCholesky-QR increases with the condition number of the preconditioned matrix, as is corroborated by the perturbation theory. This is joint work with James Garrison.

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

James Garrison
North Carolina State University
jegarri3@ncsu.edu

MS255

Novelty Sampling for Fast, Effective Data Reduction

We present a new algorithm for reducing a large data set to a small number of landmark data points. The landmarks are randomly selected, yet they account for nearly all the "novelty" in the data. To generate landmarks, we randomly propose data points and accept/reject with probabilities depending on the previous selections. After the generation step, the landmarks can be used to quickly make predictions and find clusters in the data. Landmark-based learning has a memory footprint which is independent of the data size, so the approach is suitable for distilling large data sets with $N >= 10^9$ data points.

Robert Webber
California Institute of Technology
rwebber@ucsd.edu

Joel A. Tropp
Computing and Mathematical Sciences
California Institute of Technology
jtropp@cms.caltech.edu

Ethan N. Epperly
California Institute of Technology
eepperly@caltech.edu

MS256

Digital Twins for Risk-Aware Personalized Medicine

This talk discusses the mathematical and computational ingredients needed to build trustworthy digital twins for personalized medicine. We consider a patient-specific digital twin supporting a human decision-maker for risk-aware clinical decision-making in brain cancer. We build predictive digital twins as probabilistic graphical models for encoding uncertainty and using Bayesian methods to execute the digital twin's calibration, prediction, and decision tasks. When deployed, a digital twin should be able to (i) provide interpretable decisions with explainability of underlying models being key, (ii) quantify the effects of multiple sources of uncertainty and account for underlying risk in safety-critical decisions, and (iii) be computable on actionable timescales. We achieve explainability and interpretability through a model-centric view of the digital twin that places physics-based models at the core. The digital twins are made patient-specific through continuous calibration using the patient's data that leads to predictions with quantified uncertainty under different control actions. We then solve the optimal control under uncertainty problem balancing treatment efficacy and toxicity that leads to a suite of risk-informed decisions for the patient. For these decisions to be useful for the clinicians, the digital twin needs to be executed within acceptable timescales. We achieve fast computations through explainable surrogate models and the use of multifidelity methods.

Anirban Chaudhuri
MIT
anirbanc@oden.utexas.edu

Anirban Chaudhuri
University of Texas Austin
anirbanc@oden.utexas.edu

Karen E. Willcox
Oden Institute for Computational Engineering and Sciences
University of Texas at Austin
kwillcox@oden.utexas.edu

MS256

Digital Twins of the Earth's Weather with a Focus on Extreme Events

Earth's climate is changing rapidly under the effect of global warming, leading to more frequent and severe extreme weather events [1,2]. These weather extremes, in turn, are exacting heavy socioeconomic and environmental tolls [3], prompting an urgent need for better understanding and predicting them. In this talk, we present some recent results obtained for the tropical Indo-Pacific region, using methods arising in dynamical system theory. In particular, we show that changes in weather patterns are leading to more weather extremes, namely heatwaves and extreme precipitation. These extremes can only be partially explained by El Nio-Southern Oscillation-driven variability. We then present the use of explainable AI tools to investigate the onset and precursors of these extremes. More specifically, we try to bridge existing human knowledge and AI knowledge to better understand their behaviour and

predictability.

Gianmarco Mengaldo
National University of Singapore
College of Design and Engineering
mpegim@nus.edu.sg

MS256

Using Data Consistent Inversion to Build Population-Informed Priors for Bayesian Inference

Bayesian inference is an essential tool for building predictive digital twins as it incorporates information from observational data to reduce uncertainties in the underlying computational models. However, data from an individual physical asset may be sparse, requiring highly-informative priors for Bayesian inference; in practice, the physical knowledge needed to construct such priors may not be available. Consequently, this work presents a novel approach for leveraging data from a population of related assets to construct informative Bayesian priors. Specifically, we use data-consistent inversion to estimate population-informed priors that improve inference on an individual asset. We demonstrate this approach on a computational mechanics application. Numerical examples show that utilizing population-informed priors significantly increases the KullbackLeibler divergence, i.e. the information gain from the posterior to the prior, in comparison to standard prior specification. These results are complemented with theory for linear-Gaussian inference that establishes the conditions under which using our approach is guaranteed to improve posterior estimates of uncertainty.

Rebekah White
Sandia National Laboratories
rebwhit@sandia.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Tim Wildey
Optimization and Uncertainty Quantification Dept.
Sandia National Laboratories
tmwilde@sandia.gov

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS256

Multi-Fidelity Delayed Acceptance for PDE Inverse Problems with Progressive Neural Network Surrogates

Performing uncertainty quantification (UQ) for inverse problems requires repeated multi-scenario evaluations. When complex partial differential equations are involved, the high computational cost of full-order models (FOM), such as Finite Elements, makes traditional methods like Markov Chain Monte Carlo (MCMC) impractical. Additionally, relying on inexpensive data-driven surrogate models is challenging due to limited high-fidelity data and the high accuracy required for inverse problems. Conversely, low-fidelity data can be obtained more efficiently, simplifying regression tasks. To address these challenges, we propose a Multi-Fidelity Delayed Acceptance scheme. Inspired by the Multi-Level Delayed Acceptance scheme, our method introduces a flexible framework combining solvers with different fidelities, not strictly related to simulation refinement levels. By leveraging progressive neural network surrogate models, our approach uses the FOM only during the training phase, allowing for coarser simulations at identification time. This model sequentially performs multi-fidelity regression from multiple sources, enabling on-the-fly selection of fidelity levels. The strategy is tested on a set of benchmark problems, including isotropic groundwater flows and MEMS accelerometer calibration. Results show that our method consistently reduces the computational cost associated with MCMC sampling, offering a promising approach for efficient UQ in a wide range of engineering problems.

Filippo Zacchei, Paolo Conti, Matteo Torzoni
Politecnico di Milano
filippo.zacchei@polimi.it, paolo.conti@polimi.it,
matteo.torzoni@polimi.it

Attilio Frangi
Politecnico di Milano
Dipartimento di Ingegneria Civile e Ambientale, DICA
attilio.frangi@polimi.it

Andrea Manzoni
Politecnico di Milano
andrea1.manzoni@polimi.it

MS257

Instability and Self-Propulsion of Autophoretic Filaments: Theory and Numerics

Chemically active filaments are autophoretic particles that can swim spontaneously in a viscous fluid owing to the slip flow induced by local chemical gradients produced at their surface. These gradients arise from spatial variation in the activity, which represents the concentration flux and may vary across the filament. While autophoretic filaments take advantage of both, their geometric anisotropy and their surface activity to break the symmetry in order to self-propel; their flexibility may allow for different conformations and modes of transport (e.g., tumbling, pumping, translating, or rotating). By means of analytical predictions and numerical simulations, we investigate the stability of a flexible autophoretic filament, with a uniform axisymmetric activity, against planar perturbations and the resulting emergence of self-propulsion. In the highly flexible regime, we find that the filament is susceptible to a buckling instability. The latter results from the competition between the stress induced by phoretic effects and elastic restoring forces. Our findings can be exploited to achieve precision transport of artificial microswimmers through complex environments.

Ursy Makanga, Akhil Varma
Max Planck Institute for the Physics of Complex Systems
ursy@pks.mpg.de, varma@pks.mpg.de

MS257

A Simulation Platform for Slender, Semiflexible, and Inextensible Fibers with Brownian Hydrodynamics and Steric Repulsion

Aside from their ubiquity in biology, physics, and engineering, filaments present unique challenges from an applied-mathematical point of view. Their slenderness, inextensibility, semiflexibility, and meso-scale nature all require numerical methods that can handle multiple lengthscales in the presence of constraints. Accounting for Brownian mo-

tion while keeping the dynamics in detailed balance is difficult, as is including a background solvent, which couples the dynamics of multiple filaments together. In this talk, I present a simulation platform for deterministic and Brownian inextensible filament dynamics which includes nonlocal fluid dynamics and steric repulsion. In the former case, I define the mobility using line integrals of Rotne-Prager-Yamakawa regularized singularities, and numerically preserve the symmetric positive definite property by using a thicker regularization width for the nonlocal integrals than for the self term. For steric repulsion, I introduce a soft local repulsive potential defined as a double-integral over two filaments, then present a scheme to evaluate the nonzero components of the integrand. I demonstrate that Langevin dynamics sample from the equilibrium distribution of free filament shapes, and that the modeling error in using the thicker regularization is small. I conclude with two examples, sedimenting filaments and cross-linked fiber networks, in which nonlocal hydrodynamics does and does not generate long-range flow fields, respectively.

Ondrej Maxian
University of Chicago
ondrej@uchicago.edu

MS257

An Overview of Brownian Dynamics with Constraints

This talk will serve as a short and general introduction to numerical methods for simulating Brownian dynamics with constraints. Ill discuss why we need to use methods that ensure that both hydrodynamic interactions and thermal noise are captured, and why these methods must carefully account for geometric and mechanical constraints present in the system. Ill discus specific examples which include: colloidal particles suspended in domains with complex boundaries, inextensible filaments and membranes, as well as rigid particles with arbitrary shapes. Ill end the talk with some discussion on recent advances and open problems in the field.

Brennan Sprinkle, Ryker Fish
Colorado School of Mines
bsprinkl@mines.edu, rfish@mines.edu

MS258

Using Integrators As Constraints in Quantum Collocation

Piccolo.jl solves constrained direct optimization problems for quantum optimal control (QOC). Framing QOC as a constrained direct optimization problem allows *Piccolo.jl* users to maximize the utility of their devices, designing minimum time controls, robust controls, interpolatable controls, and more. Quantum dynamics enter direct optimization problems as integrator constraints between adjacent knot points. Efficient and accurate integrators are critical for practical numerical optimization of QOC problems. Bespoke integrator designs are often the key for successful control across the quantum technology landscape. In this talk, we show how *Piccolo.jl* allows users to take advantage of this integrator design space. We discuss how integrators for quantum sensing have been adapted for shaken lattice interferometers. We show how numerically-efficient Padé integrators attain accuracy through symmetry preservation. We also compare how the demands of open quantum system integrators manifest for indirect and direct methods, and how to manage the common case of time-dependent control carriers. Each example is demonstrated using the *Piccolo.jl* open source ecosystem.

Andrew Goldschmidt
University of Chicago
andygold@uchicago.edu

Aaron Trowbridge, Jack Champagne
Carnegie Mellon University
atrowbri@andrew.cmu.edu, jchampag@andrew.cmu.edu

MS258

Optimal Quantum Control for Multi-qubit Operations via Time-parallel Multiple-shooting

Quantum optimal control is essential in quantum computing, serving as the pulse-level interface between the quantum compiler and hardware. When applied to larger multi-qubit operations, it holds the potential to yield higher fidelity and shorter durations for quantum algorithms compared to standard compilation using single and two-qubit gates. However, the significant computational demands of multi-qubit optimal control require the use of large-scale High-Performance Computing (HPC) platforms to harness greater computational concurrency. This presentation introduces a multiple-shooting approach for quantum optimal control that enables concurrency along the time domain. This approach partitions the time domain into multiple windows, with the intermediate states at window boundaries treated as additional optimization variables. Continuity of state is enforced through equality constraints. This structure facilitates parallel-in-time computation of state evolution across different time windows, leading to a substantial acceleration in the evaluation of the objective function and its adjoint-based gradient. We employ a quadratic penalty method to solve the time-parallel constrained optimization problem and demonstrate the effectiveness and efficiency of this method through numerical experiments on quantum Fourier transform gates in systems with 2, 3, and 4 qubits, where we observe an 80x speedup in gradient evaluation for the 4-qubit case.

Stefanie Guenther
Lawrence Livermore National Laboratory
guenther5@llnl.gov

MS258

Universally Robust Quantum Control

We study the robustness of the evolution of a quantum system against small uncontrolled variations in parameters in the Hamiltonian. We show that the fidelity susceptibility, which quantifies the perturbative error to leading order, can be expressed in superoperator form and use this to derive control pulses which are robust to any class of systematic unknown errors. The proposed optimal control protocol is equivalent to searching for a sequence of unitaries that mimics the first-order moments of the Haar distribution, i.e. it constitutes a 1-design. We highlight the power of our results for error resistant single- and two-qubit gates. Reference: Phys. Rev. Lett. 132, 193801 (2024)

Anthony Kiely
University College Dublin

anthony.kiely@ucd.ie

MS258

Determining the Control Mechanisms for Creating Quantum Gates

It has become common place to create one and two qubit quantum gates, both in simulation as well as in the laboratory using a variety of platforms. However, the nature of the transformation from an initial unit matrix to the final unitary gate is a matter of quantum control mechanism about which little is understood. A typical gate has non-zero and zero elements in its matrix, and it is equally important to understand the mechanism by which both types of elements are "created" by the control. The presentation will provide gate mechanism analysis through an interpretation of the interfering quantum pathway amplitudes associated with the significant contributing terms to the Dyson expansion for each particular gate. Illustrations will be shown for creating the CNOT and SWAP gates as closed systems. In the case of the CNOT gate, two distinct control fields will be utilized that both individually yield the same gate while producing distinct mechanisms. It will also be pointed out that the zero elements in a gate are typically created by massive destructive interference of the quantum pathway amplitudes.

Herschel Rabitz
Department of Chemistry
Princeton University
hrabitz@Princeton.EDU

MS259

Data-Driven Inversion of Full-Field Deformation Data with Neural Ordinary Differential Equation Fields

Accurately capturing the nonlinear mechanical behavior of soft heterogeneous materials remains a signifi-cant challenge across various fields, including biomechanics and soft robotics. While recent advancements have introduced data-driven strain energy functions, these models have primarily been trained on homogenized stress-strain data. However, in practical applications, stress-strain relationships exhibit spatial variability, ne-cessitating the identification of material parameters from full-field deformations under measurable boundary conditions. To address this, we propose a physics-guided framework for parameter identification in data-driven material models using full-field deformation data. While the approach is broadly applicable to any data-driven model, we demonstrate its effectiveness with our previously introduced neural ordinary differen-tial equation (NODE)-based model. The method is based on minimizing the strong form of the balance of linear momentum, ensuring physically consistent parameter estimation. To validate our approach, we per-form finite element simulations using the identified parameters and compare the results against experimental data. We assess the methods performance using both synthetic and experimental strain fields.

Adrian Buganza Tepole
Purdue University
abuganza@purdue.edu

MS259

Fully Automated Framework for 3D Segmentation and Velocity Reconstruction in 4D Flow Mri for Hemodynamics Analysis

Accurately reconstructing hemodynamic flows from 4D Flow MRI data is crucial for advancing cardiovascular diagnostics. Yet, the technique faces challenges such as low spatial resolution, lengthy scan times, and susceptibility to artifacts, which can degrade the quality of hemodynamic evaluations, particularly in areas of rapid velocity change. We introduce scalable hemodynamic reconstruction for 4D flow MRI (SHR-4D), which leverages a forward measurement model that integrates geometry and velocity fields to enhance the accuracy and resolution of hemodynamic assessments. This integration of geometry and velocity fields allows SHR-4D to work with either magnitude and velocity data or just the velocity data, significantly enhancing its applicability. By employing maximum likelihood estimation and subsampling techniques, SHR-4D efficiently processes high-dimensional image data, significantly reducing noise and improving the precision of velocity field reconstructions and geometry segmentation. Through testing on synthetic datasets and in vitro experiments, we demonstrate superior performance in flow segmentation and velocity reconstruction. The method's adaptability, scalability, and effectiveness in handling different imaging conditions and data types make it a valuable tool for improving cardiovascular disease diagnosis and research.

Atharva Hans
Purdue University
hans1@purdue.edu

MS259

Advancing Systems Biology Through Computational Models and Ml/AI Techniques

The recent success of bioinformatics, driven by advances in molecular biology and molecular medicine, has led to an exponential increase in the volume and diversity of biological data. This includes nucleotide and protein sequences, annotations, high-throughput experimental data, and biomedical literature. In parallel, systems biology has emerged as a critical field, shifting away from the reductionist approach that previously dominated biological research. This shift necessitates the coordinated efforts of biologists, data analysts, mathematical modelers, and computer scientists. The accumulation of large-scale biological databases calls for novel computational technologies and research strategies. In this context, computational models and tools, particularly those utilizing machine learning and artificial intelligence (ML/AI) techniques, have proven highly successful. Methods such as clustering and classification for gene expression data play a pivotal role in knowledge discovery, modeling, and optimization tasks. These methods enable the development of computational models capable of predicting the responses of complex biological systems to various perturbations. This paper explores the integration of ML/AI techniques and computational methods in systems biology, highlighting their impact on understanding and modeling biological complexity.

Mohamed Labadi
University of Chlef
Frontiers Labs
mohamed.labadi2@gmail.com

Abdullah Dabdoub
University of Paris-Saclay
abdullahdabdoub@gmail.com

Meriem Labadi

University of Souk-Ahras and Frontiers Labs
m.labadi@frontierslabs.com

Pauline Kors
Medical University of Graz
pauline.kors@stud.medunigraz.at

Samir Abdelmalek
University of Tipaza
s.abdelmalek@cu-tipaza.dz

MS259

Advancing Computational Cardiology: Neural Network Surrogates for Patient-Specific Electrophysiology Models

Neural network surrogate models are emerging as powerful tools for patient-specific cardiac electrophysiology modeling, bridging the gap between computational cardiology and clinical application. Traditional approaches often depend on models that require extensive retraining to adjust for variations in individual patients and pathological conditions. Our novel framework integrates physics-based modeling with data-driven methods to create a versatile neural network surrogate capable of handling various cardiac anatomies. We employ Branched Latent Neural Maps (BLNMs) to efficiently encode complex space-time fields into a neural network, enabling accurate and robust patient-specific parameter estimation. By leveraging statistical shape modeling, our approach augments a cohort of Tetralogy of Fallot patients by incorporating anatomical variability, improving the model's generalizability and predictive accuracy across diverse geometries. Furthermore, the framework enables the prediction of activation times in the cardiac cycle, providing critical insights into patient-specific electrical activity. This approach allows for the seamless adaptation of the surrogate model across patient cohorts, eliminating the need for geometry-specific retraining. Our pipeline contributes to the clinical translation of computational cardiology by offering fast, reliable, and adaptable tools for personalized cardiac care.

Elena S. Martinez
Stanford University
elenasm@stanford.edu

Beatrice Moscoloni
Delft University of Technology
b.moscoloni@tudelft.nl

Matteo Salvador
Stanford University
msalvad@stanford.edu

Mathias Peirlinck
Delft University of Technology
m.peirlinck@tudelft.nl

Alison L. Marsden, Fanwei Kong
Stanford University
amarsden@stanford.edu, fwkong@stanford.edu

MS259

Sciml for Aortic Hemodynamics and Wall Mechanics: Lessons from Surrogate Learning for Uncertainty Quantification

We will present a selection of real problems with surrogate learning for computational biomechanics with a focus on uncertainty quantification. Examples are Gaussian process (GP) and multi-fidelity surrogates for CFD simulations of non-Newtonian aortic haemodynamics [Ranftl et al 2021], as well as Bayesian neural networks (NN) for aortic wall mechanics [Ranftl et al 2022]. In that regard, Bayesian inference offered a straight-forward handling of missing data and outliers, a frequent challenge in biological and medical data. We will discuss not just what worked well but also what didnt or remains an open question. E.g. multi-fidelity GPs or physics-informed NNs proofed challenging in dynamic or strongly non-linear systems with large deformations, respectively. If time permits, we will touch upon the inconvenient issue with the definition of the boundary conditions and the boundary itself based on e.g. MRI, which emerges typically in patient-specific biomechanics models, and its consequences for the Machine Learning (ML) part. In essence, it is ill-advised to consider the learning task isolated from the input space specification or study objectives. The talk will focus neither on computational biomechanics nor on ML, but on the specific challenges and opportunities at the interface thereof encountered in the examples. [1] Ranftl et al. IJNMBE 39.4 (2021) e3576. [2] Ranftl et al. CMAME 401 (2022): 115594.

Sascha Ranftl
Graz University of Technology
Institute of Theoretical Physics-Computational Physics
sr7843@nyu.edu

MS260

Approximate Bayesian Computation Using Distribution Random Forests

We introduce an Approximate Bayesian Computation (ABC) framework for estimating the posterior distribution and the maximum likelihood estimate (MLE) of the parameters of models defined by intractable likelihood functions. This framework can describe the possibly skewed and high dimensional posterior distribution by a novel multivariate copula-based distribution, based on univariate marginal posterior distributions which can account for skewness and be accurately estimated by Distribution Random Forests (DRF) while performing automatic summary statistics (covariates) selection, and based on robustly estimated copula dependence parameters. The framework employs a novel multivariate mode estimator to perform for MLE and posterior mode estimation, and an optional step to perform model selection. The posterior distribution estimation accuracy of the ABC frameworks is illustrated through simulation studies involving models with analytically computable posterior distributions; and involving exponential random graph and mechanistic network models, each defined by an intractable likelihood from which it is costly to simulate large networks. The ABC frameworks are also illustrated through analyses of massive complex networks. Finally, we also consider (copula-free) generalized Bayes ABC, applicable to high-dimensional parametric models, using component-wise Gibbs sampling of the posterior distribution based on robust loss functions for the likelihood estimated via DRF.

George Karabatsos
University of Illinois Chicago
Department of Mathematics, Statistics & Computer Sciences

georgek@uic.edu

MS261

Structure-Preserving Dynamical Low-Rank Methods for Kinetic Equations

Dynamical low-rank methods have gained interest in recent years as a viable solution to the curse of dimensionality in the numerical solution of kinetic equations. Depending on the regime and nature of the solution, a direct application of these methods often results in high ranks or comes at the cost of losing physical properties such as conservation. Using the Vlasov-Fokker- Planck equation as a model problem, which is the governing equation for plasma dynamics, we demonstrate that when tuned to the problem's structure, the dynamical low-rank method can provide an efficient and robust approximation compared to the full tensor method. This talk is based on the joint work with Jack Coughlin and Uri Shumlak.

Jingwei Hu
Department of Applied Mathematics
University of Washington
hujw@uw.edu

MS261

A Galerkin Approach for Trajectorial Average in Maximally Superintegrable Systems with Application in Plasma Physics

Kinetic equations are usually in the form of a Liouville equation plus interaction terms (for example collisional operators). That structure often results in a multiscale problem, since the time scale of the Hamiltonian dynamics is in general different from those of interactions. When the advection in phase space corresponding to the Poisson bracket is much faster than the interaction processes, numerical simulations can be expensive due to the CFL condition. To avoid that, one would expect to perform average along the particle trajectories. In this work, we propose a novel structure-preserving method by investigating the relation of function spaces. Consequently, it can be categorized as a Galerkin approach. We will also show how it was used to solve the particle-wave system in plasma physics.

Kun Huang
University of Texas at Austin
k_huang@utexas.edu

Irene M. Gamba
Department of Mathematics and ICES
University of Texas
gamba@math.utexas.edu

MS261

A Reduced Order Model Enhanced Iterative Solver for Parametric Radiative Transfer Equation

Radiative transfer equation (RTE) models particles propagating through and interacting with a background medium. Applications, such as uncertainty quantification, medical imaging, and shape optimization, require solving RTE many times for various parameters. Source Iteration (SI) with Diffusion Synthetic Acceleration (DSA) is a popular iterative solver for RTE. SI-DSA can be seen as a preconditioned fixed point iteration. DSA serves as the preconditioning step to accelerate the convergence. DSA is based on the diffusion limit of a kinetic correction equation. However, when the underlying problem is far from its diffusion limit, DSA may become less effective. Furthermore, DSA does not exploit low-rank structures of the solution manifold concerning parameters of parametric problems. To address these issues, we enhance SI with data-driven reduced-order models. These data-driven ROMs still build on the original kinetic description of the correction equation and leverage low rank structures concerning parameters of parametric problems. A new preconditioner, which exploits the advantages of both data-driven ROMs and the classical DSA method, is developed.

Zhichao Peng
The Hong Kong University of Science and Technology
pengzhic@ust.hk

MS261

High-Order Micro-Macro Decomposition Schemes for Boltzmann-BGK

The kinetic Boltzmann equation with the Bhatnagar-Gross-Krook (BGK) collision operator allows for the simulation of gas dynamics over a wide range of Knudsen numbers with a simplified collision operator. Efficient numerical methods for Boltzmann-BGK should be asymptotic-preserving, which allows the numerical method to be stable at fixed mesh parameters for any value of the Knudsen number, including in the fluid (very small Knudsen numbers), slip flow (small Knudsen numbers), transition (moderate Knudsen numbers), and free molecular flow (large Knudsen numbers) regimes. In this work, we develop a novel micro-macro decomposition scheme for solving the Boltzmann-BGK and Boltzmann-ES-BGK systems. Kinetic flux vector splitting techniques are used to develop accurate fluxes and to enforce boundary conditions. High order is achieved using discontinuous Galerkin finite element methods. Several numerical examples are shown to demonstrate the effectiveness of the proposed numerical scheme.

James A. Rossmanith, Preeti Sar
Iowa State University
Department of Mathematics
rossmani@iastate.edu, sar@iastate.edu

MS262

Particle Flows and Diffusions for High-Dimensional Nonlinear Bayesian Inference

Recently, many new fully nonlinear data-assimilation methods have been developed in the data-assimilation, mathematical, statistical, and machine learning communities that can be applied to high-dimensional systems. Of particular interest are so-called particle flow filters and diffusion-based filters because of their efficiency, making them computationally compatible with existing methodologies for e.g. weather forecasting. Particle Flow filters minimize the distance between a set of particles and the posterior pdf in an iterative way. They rely on accurately estimating the gradient of the logarithm of the prior, which is difficult from a set of prior particles. Diffusion-based methods are becoming mainstream in machine learning as methods to generate samples from highly non-Gaussian densities using Langevin dynamics, and, with a small addition, turn out to have great potential in data assimilation. Of special importance are so-called stochastic interpolants, that can be show to converge in a finite number of iteration steps. Interestingly, these diffusion-based methods do not need the gradient of the logarithm of the prior for their

performance. In this presentation I will explain how these methods are connected and how we can make them useful for high-dimensional problems using examples, and how developments in traditional data assimilation can considerably accelerate some of these methods.

Peter Jan van Leeuwen
Colorado State University
peter.vanleeuwen@colostate.edu

MS262

A High-Dimensional Deterministic Convergent Particle Filter

Mixture-model particle filters such as the ensemble Gaussian mixture filter require a stochastic resampling procedure in order to converge to exact Bayesian inference. Gaussian-mixture kernel density estimation is inefficient in high-dimensions. We combine the recent advancement of the Ensemble Epanechnikov Mixture Filter and the optimal transport deterministic resampling strategy for a high-dimensional particle filter that provably converges to exact Bayesian inference.

Andrey A. Popov
University of Hawai'i at Manoa, U.S.
apopov@hawaii.edu

MS262

Efficient Randomization Techniques for Data Assimilation

Large-scale inverse problems involve fusing incomplete and noisy information from multiple sources, such as model simulations, measurements from sensors, and physical experiments, to obtain a consistent description of the state of the underlying physical system. Solving the Bayesian formulation of these problems enables quantifying the uncertainties associated with the solution. However, solving Bayesian problems presents a major challenge: Solving Bayesian inverse problems is computationally demanding, often requiring hundreds to thousands of expensive simulations to accurately estimate the parameters and their uncertainties. Randomized algorithms provide an attractive means to reduce the computational cost. In this work, we will explore efficient randomization techniques as a means to develop scalable solvers and pre-conditioners to mitigate the computational costs associated with solving Bayesian inverse problems.

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Amit Subrahmanya
Virginia Tech
amitns@vt.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS263

Mesoscale-microscale Coupling for Wind Energy Simulations

Wind farm simulations require data from mesoscale atmospheric simulations as initial and boundary conditions for the microscale turbine environments. A file-based coupling of mesoscale and microscale solvers is computationally inefficient and restricted in the amount of data that can be exchanged. We will present a runtime, or in-memory, coupling of the Energy Research and Forecasting (ERF) atmospheric flow solver with the AMR-Wind turbine flow solver. Whereas ERF solves the compressible flow equations, the AMR-Wind code solves for anelastic or incompressible flow. As both codes are built upon the AMReX framework, the multiblock feature in AMReX enables the data transfer by defining a mapping between the data structures of the two simulation classes. We implement a mass balance correction in AMR-Wind to enforce the incompressibility constraint for the boundary velocity field received from ERF. We will demonstrate a one-way coupling of ERF to AMR-Wind for a typical atmospheric boundary layer simulation, which also allows for different temporal and spatial resolutions at the interface and hence more efficient use of computational resources. We will discuss other performance related challenges such as load balancing. Finally, we will present our progress on a two-way coupling, which enables the microscale solver to feed back high resolution data to the mesoscale model for simulating interactions between distant wind farms or studying their effect on the regional environment.

Mukul Dave
ASU Research Computing
Arizona State University
mukul92@gmail.com

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

Marc Henry de Frahan
National Renewable Energy Laboratory
marc.henrydefrahan@nrel.gov

MS263

Comparing Design and Performance of Two Versions of the E3SM Coupler

Coupled simulations in the Energy Exascale Earth System Model (E3SM) integrate complex models of the atmosphere, ocean, rivers, land surface, and sea ice into a single system. The software to transfer data between components and perform interpolation between different meshes makes up most of the coupler of E3SM. To improve computational efficiency and scientific productivity, efforts have been made to transition from the current coupler based on the Model Coupling Toolking (MCT) to a topology-aware library, the Mesh Oriented datABase (MOAB). MOAB provides a comprehensive description of the topology for each submodel in E3SM. This approach supports arbitrary grids and adaptive decompositions, ensuring seamless operations like mesh intersections, conservative projection operator computation, and remapping weight application. Integrating the MOAB mesh database within E3SM, with interfaces to TempestRemap, allows for the generation of remapping weights during simulations, mostly eliminating the need for offline pre-processing. We present detailed algorithmic and performance results of the new MOAB coupler and compare it with the existing E3SM-MCT coupler. Potential enhancements to the coupler, including high-order projection methods for property preservation, are also discussed. This transition aims to optimize computational efficiency and scientific output by leveraging MOABs advanced capabilities in mesh topology and data

management.

Robert Jacob, Iulian Grindeanu, Vijay S. Mahadevan
Argonne National Laboratory
jacob@anl.gov, iulian@anl.gov, mahadevan@anl.gov

MS263

Enhanced Sub-Mesoscale Feature Resolution in Multiscale Ocean Simulations

Submesoscale eddies and fronts are critical components of the ocean's circulation, driving mixing and restratification in the upper ocean, controlling the advection and transport of temperature, salinity, and suspended tracers, and promoting subduction and upwelling pathways that ventilate the ocean's interior. However, fully resolving these spatiotemporal scales adaptively across the entire ocean domain using only a global model incurs significant computational costs. In this talk, we present a hierarchically nested modeling approach to better capture the effects of coastal and open-ocean dynamics in climate models. By integrating a high-resolution Regional Ocean Modeling System (ROMS) with the global Model for Prediction Across Scales-Ocean (MPAS-O) for specific regions, better resolution of the sub-mesoscale processes can be achieved with minimal computational impact. We discuss the numerical coupling algorithms implemented with the MOAB unstructured mesh library infrastructure for seamless one-way and two-way offline and online coupling, aimed at improving the fidelity with which MPAS-O resolves nonlinear ocean dynamics. Additionally, we demonstrate the flexibility of this multiscale modeling approach through open-ocean cases with refinements in the Gulf Stream and North Atlantic regions, as well as coastal refinements in the Bay of Bengal and Chesapeake Bay regions.

Vijay S. Mahadevan
Argonne National Laboratory
mahadevan@anl.gov

Kyle Hinson
Pacific Northwest National Laboratory
kyle.hinson@pnnl.gov

Iulian Grindeanu
Argonne National Laboratory
iulian@anl.gov

Robert Hetland
Pacific Northwest National Laboratory
robert.hetland@pnnl.gov

MS263

Coupling Large-Scale Wave (WaveWatch III) and Atmospheric (Energy Research Forecasting) Simulations

WaveWatch III (WW3) is a community wave modeling framework which solves the spectral action density balance equation to describe the oceans wave heights, frequencies, and directions. In this talk, we describe how we couple the Energy Research Forecasting (ERF) code, which solves the compressible Navier-Stokes equations for large-scale weather modeling, with WW3. Using MPI and AMReXs MPMD capability, we send the significant wave heights and mean wavelengths from WW3 to compute a surface roughness parameter within ERF through a fixed point iteration. The average wind velocities and wind directions from ERF are sent to WW3 to compute an atmospheric source term, generating a more accurate simulation of wind and waves across the atmospheric surface layer. This two-way coupling between large-scale simulations allows us to visualize a more detailed picture of the effects of wind flow across the oceans surface, as well as compare our model to previously uncoupled or one-way coupled versions. Paying particular attention to load balancing as a function of the time-steps of each simulation, we are then able to discern more efficient coupling routines for different physical domains and weather conditions.

Jack Pham
University of California, Merced
tpham82@ucmerced.edu

Jean M. Sexton, Aaron M. Lattanzi
Lawrence Berkeley National Laboratory
jmsexton@lbl.gov, amlattanzi@lbl.gov

Francois Blanchette
University of California, Merced
fblanchette@ucmerced.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

MS263

Language Interoperable Native Runtime Coupling of AMReX and pyAMReX : Multiscale Flow Modeling Demonstration

Continuum modeling of rheologically complex fluids with unknown constitutive relationships, such as that between local flow state and viscosity, requires a multiscale methodology wherein particle-based simulations can either be conducted across the entire rheological space to pre-compute the constitutive response, or in-situ to inform local rheology at every fluid cell at each time step of a fluid simulation. However, both these approaches can quickly become computationally intractable due to a priori unknown and vast rheological space, whereas the latter involves unavoidable dependency on time-consuming particle-based methods. This need for extensive particle-based simulations can be significantly reduced by employing an uncertainty quantified approximation of the complex fluids constitutive response. This method adaptively focuses only on the under-sampled regions of the rheological space with high-variance, thereby improving efficiency by minimizing the number of particle simulations. This study leverages an ensemble deep learning framework designed to obtain the underlying problem-specific constitutive relationship. The framework is implemented by integrating an AMReX-based scalable incompressible flow solver (incflo) with PyLAMMPS, facilitated by pyAMReX and PyTorch. The efficacy and accuracy of this approach will be demonstrated through continuum modeling of granular material and complex-rheology suspension dynamics and compared against discrete element simulations.

Bhargav Sriram Siddani
Lawrence Berkeley National Laboratory
bsiddani@lbl.gov

Axel Huebl
Lawrence Berkeley National Laboratory (LBNL), USA
axelhuebl@lbl.gov

Weiqun Zhang

Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
WeiqunZhang@lbl.gov

Andrew J. Nonaka, Ishan Srivastava
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, isriva@lbl.gov

MS264

Benchmark Dataset and Framework Accelerate Seismic AI Research

Revolutionary artificial intelligence (AI) research requires a robust benchmark dataset and workflow. Benchmark, like the widely known ImageNet, aims to create a general-use dataset for scientists to compete, share research updates, and have a quantitative metric. Either active seismic imaging for CO2, hydrogen, or geothermal uses or passive seismic imaging for hazard management relies on large-scale benchmark models/datasets for an upgrade in monitoring capability. Benchmark can be a viable avenue for researchers active in active/passive monitoring to excel in AI research. Two examples, i.e., distributed acoustic sensing (DAS) data processing and earthquake location imaging, will be used to demonstrate how benchmarks help accelerate seismic AI research and enable the best-ever performance.

Yangkang Chen
The University of Texas at Austin
yangkang.chen@beg.utexas.edu

MS264

Gpu Acceleration of Acoustic Sub-Seabed Imaging

We present a GPU-based acceleration of an acoustic imaging process used as part of a sub-seabed surveying technology. Imaging involves a Kirchhoff pre-stack time migration-based process that back-propagates received acoustic energy using two-way travel time calculations, followed by summation over all source and receiver pairs for each imaging point. The original implementation was CPU-based and parallelized using standard C++ threads to distribute computational work across CPU cores. As the operations performed at each imaging point are independent, there is a high degree of data parallelism and arithmetic intensity, making the problem suited for the many-core GPU architecture. We describe porting of the code to NVIDIA GPUs using CUDA, and address details such as floating-point aspects, memory transfers and domain decomposition. To remove noise spikes, the original CPU algorithm uses an alpha-trimming technique that sorts the contributing values at each image point, and trims a percentage of high and low values before doing the summation. Implementation on the GPU uses a modified alpha-trim procedure that eliminates the need for an expensive sort operation. We show comparison of images from the CPU and GPU codes, and provide performance results demonstrating significant speedup with the GPU code.

Clyde Clements, Richard Charron, Ryan Laidley, Cevdet Kilic
Kraken Robotics
cclements@krakenrobotics.com, rcharron@krakenrobotics.com, rlaidley@krakenrobotics.com, ckilic@krakenrobotics.com

MS264

Machine Learning Tools For Seismic Low Frequency Extrapolation

Low frequency seismic data and accurate initial models are essential for circumventing the inherent cycle-skipping problem in full waveform inversion (FWI). However, acquiring low frequency seismic data is challenging due to receiver limitation and noise contamination. Seismic low frequency extrapolation refers to the process of predicting the unrecorded low frequency components from the available high frequency seismic data. Deep learning algorithms, such as convolutional neural networks (CNNs) and generative adversarial networks (GANs), have shown promise in performing low frequency extrapolation. However, these methods typically require powerful hardware and can be cumbersome to train effectively. In this talk, we explore the possibility of using machine learning algorithms, such as Random Forest, Gradient Boosting, and Gaussian Process Regression, as viable alternatives to deep learning algorithms for seismic low frequency extrapolation. The training and test data are obtained from the forward solve of the Marmousi velocity model. In order to reduce the input/output data dimensions, we train the machine learning models to extrapolate the amplitude and phase spectra of the time traces. The machine learning models demonstrate low frequency extrapolation capabilities comparable to their deep learning counterparts at a reduced computational cost.

Rohin Harikumar
The University of Texas at Dallas
rxh210003@utdallas.edu

Sue Minkoff
University of Texas - Dallas
sminkoff@utdallas.edu

MS264

Stationary Phase Analysis for Scattered Waves in Ambient Noise Cross-correlations

Stacked cross-correlation functions have become ubiquitous in the ambient seismic imaging and monitoring community as approximations to the Green's function between two receivers. While the quantification of this approximation to the ballistic arrivals is well established, the equivalent analysis for scattered waves is inadequate compared to the exponential growth of its applications. To provide a clearer understanding, we derive analytical stationary phase solutions for ambient noise cross-correlations focusing on the scattered waves, which show up after the ballistic arrivals. Under idealistic conditions, when 1) noise sources are uniformly distributed around the receivers and 2) each source is recorded separately, we show that while the non-ballistic arrivals in the stacked cross-correlation functions accurately reconstruct the scattered waves, the stationary phase zones for scattered waves drastically differ from those for the ballistic arrivals. The theoretical results provide a solid foundation for data-driven ambient noise imaging and inversion.

Yunyue Elita Li, Feng Zhu
Purdue University
elitali@purdue.edu, zhu1240@purdue.edu

Jizhong Yang

Tongji University
10yangjizhong@tongji.edu.cn

MS264

Industry-Scale Uncertainty-Aware Full Waveform Inference with Generative Models

In this talk, we address the challenge of Bayesian sampling for ill-posed inverse problems involving expensive PDE forward operators. To tackle this, we propose the use of generative models and discuss three key challenges along with our solutions. First, we address the constraints of GPU memory by introducing patch-based training and memory-efficient invertible networks, enabling the handling of large-scale data. Second, we demonstrate that these models can be effectively trained using only prior data available from acoustic observations and borehole well data, making the approach practical for real-world applications. Third, we explore various strategies to enhance the reliability of posterior samples by incorporating physics-based constraints through the wave operator. We conclude by presenting experiments that showcase the capabilities of these methods in field-data Full Waveform Inversion (FWI) workflows, scaling up to industry-scale problems with large 2D (512x7000 degrees of freedom) and 3D models.

Rafael Orozco, Tuna Erdinc, Mathias Louboutin, Felix Herrmann
Georgia Institute of Technology
rorozco@gatech.edu, herdinc3@gatech.edu, mathias.louboutin@gmail.com, felix.herrmann@gatech.edu

MS266

Algorithmic Differentiation on Operator Graphs and Applications

SoNICS is a new generation CFD sotfware which models its computation as an operator graph. Such a graph contains two kinds of nodes: operators correspond to the different tasks that are performed during SoNICS's execution; and data correspond to the values exchanged between these operators. A tool allowing to automatically generate such a graph is illustrated in [Michael Lienhardt, Maurice H. ter Beek, Ferruccio Damiani: Product lines of dataflows. J. Syst. Softw. 210: 111928 (2024)]. In this presentation, we discuss how we implemented forward and backward Algorithmic Differentiation in SoNICS, using a new process suited for operator graphs. This process is structured in three parts. First, the operator graph is analysed to identify which derivatives of which operators are needed to compute the expected derivative. Second, the identified differentiated operators are generated as new operators using standard source-transformation technics (such as the one implemented in tapenade [Laurent Hascot, Valrie Pascual: The Tapenade automatic differentiation tool: Principles, model, and specification. ACM Trans. Math. Softw. 39(3): 20:1-20:43 (2013)]). And finally, these new operators are combined with existing ones to obtain a new operator graph that performs the computation of the expected derivative. We illustrate our approach on few representative CFD examples.

Sébastien Bourasseau
ONERA, The French air and space agency
sebastien.bourasseau@onera.fr

Betrand Michel
Onera, the French Aerospace Lab
bertrand.michel@onera.fr

Michaël Lienhardt, Jorge Nunez Ramirez
ONERA
michael.lienhardt@onera.fr, jorge.nunez@onera.fr

Bruno Maugars, Cédric Content
Onera, the French Aerospace Lab
bruno.maugars@onera.fr, cedric.content@onera.fr

MS266

Enzyme-MLIR: A Case Study of High-level Optimizations on Automatic Differentiation

Automatic differentiation is key to numerous workloads in scientific computing and machine learning. Programs in these domains frequently contain high-level structure, like solvers, linear algebra, energy conservation, and more. Even without automatic differentiation, exploiting this structure is key to good performance, both numerically and computationally. This is especially true when one needs to determine how to parallelize and partition problems onto multiple cores or devices. Historically, the only mechanisms to teach automatic differentiation tools about such structure has been through the application of custom differentiation rules. While these permit the users to exploit structure in specific subroutines, they are both cumbersome to write for an entire codebase, and fail to account for optimization opportunities between functions – for example, exploiting the input to a function is symmetric because it is the result of a given operation. This talk will describe ongoing work in the Enzyme-MLIR project when enables Enzyme's compiler-based automatic differentiation to be applied to arbitrary low and high-level operations. We discuss the impact of implementing novel optimizations for algebraic simplification, and enabling users to write custom optimizations rules in the transform dialect. We show several case studies applying EnzymeMLIR to both existing JaX and Julia codes.

William S. Moses
University of Illinois at Urbana-Champaign, U.S.
wsmoses@illinois.edu

Oleksandr Zinenko
Stealth
zinenko@zinenko.fr

MS266

Relaxed Adjoints

Algorithmic Differentiation [Griewank and Walther, Evaluating Derivatives, SIAM, 2008], [Naumann, The Art of Differentiating Computer Programs, SIAM, 2012] yields (tangents and) adjoints for implementations of multivariate vector functions $F : \mathbb{R}^n \to \mathbb{R}^m$ as computer programs $y = F(x)$ with machine accuracy at the given x. This sounds like good news. Is it? Accuracy and performance (adjoint, in particular) are certainly highly desirable. However, restriction to individual points turns out to be a rather hard constraint. Derivative information is local. Global behavior (e.g., over compact subdomains) is not captured. Convex relaxations of (tangents and) adjoints yield guaranteed enclosures over compact subdomains. Interval expansions and McCormick relaxations will be considered. The resulting "relaxed adjoints" will be discussed in the context of approximate computing (e.g., for energy efficiency) [Vassiliadis et al., Towards automatic significance analysis for approximate computing, Proceedings of the 2016 International Symposium on Code Generation and Optimization], pruning of neural networks [Kich-

ler et al., Towards Sobolev Pruning, Proceedings of the 2024 Platform for Advanced Scientific Computing Conference] and deterministic global optimization [Deussen and Naumann, Subdomain separability in global optimization, Journal of Global Optimization, 2023].

Uwe Naumann, Neil Kichler
RWTH Aachen University, Germany
naumann@stce.rwth-aachen.de, kichler@stce.rwth-aachen.de

MS267

Osqp with GPUs & FPGAs: Accelerating Quadratic Programming on Heterogeneous Systems

In this talk, we present our recent work in the OSQP (Operator Splitting Quadratic Programming) solver that implements end-to-end acceleration of the solver on both GPU and FPGA platforms. At the core of this acceleration is a linear algebra abstraction layer that decouples the core optimization steps from numerical operations, enabling a unified OSQP solver API regardless of the computational backend. This layer allows for seamless backend selection at build-time, including CPU backends with sparse operations or Intel MKL, a GPU backend using cuSPARSE, and an FPGA backend using RSQP. OSQP libraries with these backends are packaged with high-level interfaces, such as the OSQP Julia and Python packages, enabling users to easily solve quadratic optimization problems on GPUs. Additionally, we will outline our current work on a computational backend using CUDA graphs, which aims to improve the performance of the GPU accelerated OSQP solver by reducing the amount of time spent in kernel launches and other CPU-based control tasks. We will present the overall architecture the graph-based solver uses, and discuss our experiences modifying the existing OSQP abstraction layer to support a graph-based compute backend.

Ian McInerney
The University of Manchester
i.mcinerney17@imperial.ac.uk

Bartolomeo Stellato
Princeton University
bstellato@princeton.edu

MS267

Krylov Methods for Portable and High-Performance Computing

Krylov.jl is a Julia package that leverages Krylov methods for high-performance computing. This presentation will explore its adaptability to heterogeneous GPU architectures from Intel, AMD, and NVIDIA, including supercomputers like Aurora, Frontier, and Venado. We will focus on its performance, memory efficiency, GPU compatibility, and support for various floating-point systems.

Alexis Montoison
GERAD and Polytechnique Montréal
Canada
alexis.montoison@polymtl.ca

Dominique Orban
GERAD and Dept. Mathematics and Industrial Engineering
Ecole Polytechnique de Montreal
dominique.orban@gerad.ca

MS267

Structure Exploitation and Preconditioning for Real-Time GPU-Accelerated Optimization at the Edge

GPU acceleration is rapidly becoming a crucial paradigm for enhancing the performance of computationally demanding numerical optimization tasks, particularly in the context of real-time optimal control at the edge for robotics. However, to achieve true real-time performance, it is essential to maximize computational efficiency by exploiting the inherent structure and sparsity of optimization problems. As such, due to their parallel-friendly structure, Krylov-based methods have emerged as effective computational patterns for GPUs. However, their success critically depends on the development of robust preconditioners. In this talk, we will discuss our recent work both developing and extending new preconditioners and on improving the underlying structure of Kyrlov solvers for improved edge GPU acceleration. We will demonstrate the impact of these advancements through practical applications on edge robots, showcasing how these techniques enable significant performance gains in real-time optimization scenarios.

Brian Plancher
Barnard College, Columbia University
bplancher@barnard.edu

MS267

GPU Implementation of Algorithm NCL

For smooth constrained optimization problems, Algorithm NCL is equivalent to the augmented Lagrangian algorithm of LANCELOT. It is immune to LICQ difficulties, and has only about 10 subproblems to be solved (using interior methods and second derivatives). The AMPL implementation of NCL used IPOPT or KNITRO. To utilize GPUs we use MadNLP.jl (a nonlinear interior method working on GPU) and CUDSS.jl (a Julia interface to the NVIDIA library cuDSS). The NCL transformation allows the KKT systems to be reduced to smaller systems that can be solved by Cholesky-type factorization. We describe the reduction to smaller systems and present numerical experiments.

Michael A. Saunders
Stanford University
saunders@stanford.edu

Alexis Montoison
GERAD and Polytechnique Montréal
Canada
alexis.montoison@polymtl.ca

Dominique Orban
GERAD and Dept. Mathematics and Industrial Engineering
Ecole Polytechnique de Montreal
dominique.orban@gerad.ca

François Pacaud
École des Ponts ParisTech
francoispacaud8@gmail.com

Sungho Shin
MIT

sushin@mit.edu

MS267

Nonlinear Optimization of Energy Systems on the Next Generation of Supercomputers

In this talk, we discuss how recent developments in algebraic modeling techniques combine with advances in nonlinear programming methods to close gaps that have emerged between such modern computing infrastructure and nonlinear optimization practice. We discuss ExaSim.jl, an algebraic modeling tool built around the concept of SIMD abstractions. By offering a natural abstraction for repetitive model structures, ExaSim.jl facilitates the automatic generation of GPU-compatible (and vendor-agnostic) sparse gradient and Hessian oracles for the solution of structured nonlinear programs as arise naturally in the context of many PDE-constrained or graphical optimization models. Complemented by recent methodological advances in GPU acceleration of (condensed-space) interior point methods, this modeling paradigm enables the solution of large-scale structured nonlinear programs on diverse GPU-based computing platforms with minimal intrusion. We demonstrate a prototypical workflow and the resultant computational advantages with two examples from the realm of energy systems engineering: AC optimal power flow (ACOPF) and model-based control of battery energy storage systems. Notably, the presented approach enables a 10-fold speed-up on a wide range of ACOPF instances and real-time capable control of battery systems based on a detailed pseudo-two-dimensional PDE model previously considered intractable for control purposes.

Flemming Holtorf
Department of Chemical Engineering
Massachusetts Institute of Technology
holtorf@mit.edu

Sungho Shin
MIT
sushin@mit.edu

MS268

Efficient Batch Multiobjective Bayesian Optimization

n/a

Kade Carlson
Penn State
kkc5441@psu.edu

MS268

Control Co-design of Morphing Aerial Autonomous Systems by Multi-objective Bayesian optimization

n/a

Ashwin Renganathan, Daning Huang
Pennsylvania State University
ashwin.renganathan@psu.edu, daning@psu.edu

MS268

Surrogate-assisted Optimization for Aerospace Applications

n/a

Michael Joly
Raytheon
michael.joly@rtx.com

MS268

Intelligent Quantification of Uncertainty using Active Learning.

n/a

Murali Krishnan Rajasekaran Pillai
GE Aerospace
muralikrishnan.rajasekharanpillai@ge.com

MS269

Ground State Energy Estimation from Noisy Quantum Observables

Within the many-body physics and chemistry communities, ground state energy estimation is one of the most promising applications of quantum computing. Current quantum computers, however, are still not fully realized and are prone to producing noisy outputs. In this talk, we will discuss a novel hybrid (part classical and part quantum) algorithm that obtains real-time measurements from a quantum computer and classically processes them via a denoising procedure and dynamic mode decomposition (DMD) to obtain estimates to the ground state energy. We benchmark our method against other leading hybrid algorithms and demonstrate rapid convergence across a variety of molecules and perturbative noise

Hardeep Bassi
University of California Merced
hbassi2@ucmerced.edu

MS269

Towards Provably Efficient Quantum Algorithms for Large-Scale Machine Learning Models

Large machine learning models are revolutionary technologies of artificial intelligence whose bottlenecks include huge computational expenses, power, and time used both in the pre-training and fine-tuning process. In this work, we show that fault-tolerant quantum computing could possibly provide provably efficient resolutions for generic (stochastic) gradient descent algorithms, scaling as $O(T^2 \text{ polylog}(n))$, where n is the size of the models and T is the number of iterations in the training, as long as the models are both sufficiently dissipative and sparse, with small learning rates. Based on earlier efficient quantum algorithms for dissipative differential equations, we find and prove that similar algorithms work for (stochastic) gradient descent, the primary algorithm for machine learning. In practice, we benchmark instances of large machine learning models from 7 million to 103 million parameters. We find that, in the context of sparse training, a quantum enhancement is possible at the early stage of learning after model pruning, motivating a sparse parameter download and re-upload scheme. Our work shows solidly that fault-tolerant quantum algorithms could potentially contribute to most state-of-the-art, large-scale machine-learning problems.

Jin-Peng Liu
University of California, Berkeley

jliu1219@terpmail.umd.edu

MS269

Operator Learning Formulation of Numerical Renormalization Group

Variational algorithms have been designed by optimizing the energy or integration error as the loss function. However, these loss functions do not directly describe the target output and thus introduce bias, often leading to the local minimum not being relevant to target prediction. From a machine learning perspective, we generalize the traditional interpretation of numerical renormalization group formulation from Wilson and White to directly optimize the target observable error. We explore such generalization and propose the operator learning renormalization group (OLRG). OLRG is a generalization of the NRG and DMRG frameworks for simulating the quantum many-body system. It recursively builds a simulatable system to approximate a target system of the same number of sites via operator maps. OLRG uses a loss function to minimize the error of a target property directly by learning the operator map in lieu of a state ansatz. Using operator maps opens the opportunity to optimize more flexible deep learning, tensor network, and parameterized quantum circuit models in the DMRG and generative learning style. This loss function is designed by a scaling consistency condition providing a provable bound for real-time evolution. We explored this alternative variational framework numerically with both classical and quantum settings and demonstrated the correctness of our theory.

Xiuzhe (Roger) Luo
University of Waterloo
rogerluo.rl18@gmail.com

MS269

Learning Quantum Dynamics Through Signal Processing

In this talk I will introduce two applications of signal processing methods to learning quantum dynamics. Conservation laws provide important information about the symmetry, ergodicity, and transport properties of quantum systems. In the first application, we apply the singular value decomposition to data generated from the classical shadow procedure in order to learn the conservation laws in an unknown quantum dynamics. The dynamics of a closed quantum system is governed by its Hamiltonian, learning which enables us to predict the evolution of the quantum system. In the second application we will apply the compressive sensing technique to learning the Hamiltonian of a quantum system without geometric locality. These examples showcase the potential of using signal processing tools to extract useful information from quantum systems.

Yu Tong
Duke University
yutong9410@gmail.com

MS270

Generative Models for Uncertainty Quantification in Inverse Problems

Inverse problems are ubiquitous, and solving them is challenging due to their ill-posed nature. The Bayesian paradigm lends itself naturally to the probabilistic treatment of inverse problems: the posterior distribution characterizes the relative plausibility of solutions, which also helps quantify uncertainties. Bayesian inference is computationally challenging: high-dimensionality, nonlinear processes, and intractable likelihoods are some examples. We propose a likelihood-free simulation-based inference framework for Bayesian inference using conditional score-based diffusion models. These models use neural networks to approximate the score function of the posterior distribution for different realizations of the measurement. Notably, training requires simulations of the forward model, meaning gradient computations or likelihood evaluations are not necessary, making the framework suitable for problems with black-box forward models. After training the score network network, Langevin dynamics-based Markov chain Monte Carlo enables efficient posterior sampling when a measurement vector is available. There is no need for retraining when new measurements are available. We demonstrate the efficacy of the proposed approach on an array of inverse elasticity problems with applications in biomechanics, which concern the inference of spatially varying material constitutive parameters of specimens from noisy response (displacement/strain) measurements.

Agnimitra Dasgupta
Department of Civil and Environmental Engineering
University of Southern California
adasgupt@usc.edu

Assad Oberai
University of Southern California
aoberai

MS270

Proximal Method for Diffusion Model Adaptation

We address the challenge of adapting diffusion models across domains with distribution shift, focusing on scenarios where a small labeled dataset (e.g., from a clinical trial) is available alongside larger, relevant but unlabeled datasets. Our goal is to train a diffusion model capable of generating accurate labels for target data. Our approach leverages proximal causal learning techniques to facilitate adaptation without explicitly modeling latent confounders that may cause distribution shifts between domains. We consider two key settings: (1) Concept Bottleneck, where an observed "concept" variable mediates between covariates and labels, and (2) Multi-domain, where training data from various source domains exhibit different distributions of latent confounders. We develop a two-stage kernel estimation method to adapt diffusion models to complex distribution shifts in both settings. This method allows us to harness information from unlabeled data to enhance the performance of our model on the target domain, effectively bridging the gap between the limited labeled data and the broader unlabeled datasets. Our experiments demonstrate that this proximal approach outperforms existing methods, particularly those attempting to explicitly recover latent confounders. The results suggest that our method offers a promising direction for transfer learning in diffusion models, especially in scenarios with limited labeled data and significant domain shifts.

Huan He
University of Pennsylvania
huan.he@pennmedicine.upenn.edu

MS270

A Unified Framework for Annealing-Based Diffu-

sion Inverse Problem Solvers

Over the last five years, diffusion models have become the state-of-the-art approach for many foundational generative modeling tasks, including image generation. More recently, several approaches have been developed that use diffusion models as an image prior for severely ill-posed image reconstruction problems. These approaches aim to sample from the posterior distribution of a given Bayesian inverse problem by first training a diffusion model to learn the prior distribution and then incorporating the known likelihood function of the inverse problem into the iterative, stochastic differential equation (SDE) based sampling process of the diffusion model. In this talk, we consider two of these approaches, both of which have achieved state-of-the-art performance on various image reconstruction tasks, and show that they can be integrated into a single unified framework that yields a new family of algorithms through combinations of the design choices of the two approaches. The algorithms in this family are systematically evaluated on a set of canonical image restoration tasks (inpainting, deblurring, and phase retrieval), as well as on a model problem whose posterior distribution is available analytically in closed form.

Evan Scope Crafts, Umberto Villa
The University of Texas at Austin
escopec@utexas.edu, uvilla@austin.utexas.edu

MS270

Hybrid Diffusion Models for Sampling High-Dimensional Multi-Modal Probability Distributions

In this talk, we propose a hybrid diffusion model for efficient sampling of high-dimensional, multi-modal probability distributions. Traditional MCMC methods, such as the MetropolisHastings and Langevin Monte Carlo sampling methods, are effective for sampling from single-mode distributions in high-dimensional spaces. However, these methods struggle to produce samples with the correct ratio for each mode in multi-modal distributions. To address the challenges posed by multi-modality, we adopt a divide-and-conquer strategy. We start by minimizing the energy function with initial guesses uniformly distributed within the prior domain to identify all the modes of the energy function. Then, we train a classifier to segment the domain corresponding to each mode. After the domain decomposition, we train a diffusion-model-assisted generative model for each identified mode within its support. Once each mode is characterized, we employ bridge sampling to estimate the normalizing constant, allowing us to directly adjust the ratios between the modes. Our numerical examples demonstrate that the proposed framework can effectively handle multi-modal distributions with varying mode shapes in up to 200 dimensions.

Neil Zhang, Guannan Zhang
Oak Ridge National Laboratory
zhangz2@ornl.gov, zhangg@ornl.gov

MS271

Parallel Implementaion of Navier Stokes Equation with Mimetic Operators

This study presents a parallel computing framework for solving the Navier-Stokes equations using mimetic methods. Mimetic methods, with their ability to conserve key properties of the differential operators, offer a computationally efficient and simple-to-implement way of accurately simulating fluid flow phenomena. The scalability and performance of the parallel implementation are evaluated for large problem sizes and various computational resources. The computational problems are implemented in C++ using the Armadillo and SuperLU sparse-solving libraries. Parallelization is through MPI domain decomposition, and strong scaling up to 1666 cores is demonstrated. Additionally, we investigate the complexity of domain decomposition concerning mimetic operators and strategies to minimize the communication overhead on the overall performance of the solver. Results demonstrate the effectiveness of the parallel mimetic solver in accurately capturing fluid flow behavior while achieving significant speedup on parallel architectures.

Jared Brzenski
San Diego State University
San Diego, CA 92182, United States
jbrzenski@sdsu.edu

MS271

MOLE: Mimetics Operators Library Enhanced

MOLE is an open-source library that implements high-order mimetic operators. It provides discrete analogs of the most common vector calculus operators: divergence, gradient, curl, and Laplacian. These operators act on functions discretized over staggered grids (uniform, nonuniform, and curvilinear), and they satisfy local and global conservation laws (Dumett and Castillo, 2022a, 2023a). MOLEs operators can be utilized to develop code for solving partial differential equations (PDEs). I will describe how to solve partial differential equations using MOLE.

Jose E. Castillo
Computational Science Research Center
San Diego State University
jcastillo@sdsu.edu

MS271

High-order Accurate Discretizations for Overset Grids

High-order accurate approximations often provide more efficient numerical solutions to Partial Differential Equations (PDEs) than corresponding low-order accurate ones. This talk will discuss high-order accurate conservative and non-conservative finite difference approximations for overset grids with application to incompressible fluid flows, electromagnetics, multigrid solvers, and the solution of indefinite Helmholtz problems. Overset grids enable the use of efficient finite difference approximations on complex geometry. An overset grid typically consists one or more background Cartesian grids together will boundary fitted curvilinear grids. The component grids overlap and the discrete solutions are connected through interpolation. Overset grids are particularly useful for problems with moving geometry such as with fluid-structure interactions.

William Henshaw
Rensselaer Polytechnic Institute
wdhenshaw@gmail.com

MS272

Embedded Operator Learning: Quantifying Model Error in Learning Constitutive Operators

Neural Operators (NO), a family of scientific machine

learning tools that learn the governing physical laws from the data, has been popularized over the last decade in modeling complex physic systems. In heterogeneous material modeling where predictive errors can potentially lead to catastrophic results, significant attention needs to be paid to assessing reliability of NO models and establishing quantified estimates of uncertainty in their predictions. However, only a few works had focused on extending the uncertainty quantification (UQ) to the application of NO. In this work, we introduced an algorithm combining Bayesian inference and NO, which learns the physics and enables uncertainty prediction simultaneously. Built based on Bayesian Neural Network, our proposed method provided statistical calibration of NO with respect to various datasets, enabling representation and quantification of the modeling error and data noise, together with the model parameters. For illustration, we encoded this UQ framework into several current popular and mature NO's such as Fourier Neural Operator (FNO) and Graph Neural Operator (GNO). As demonstrations, we consider a wide variety of scientific applications, where our method outperforms by giving accurate results with reasonable confidence region and ensemble predictions.

Yiming Fan
Lehigh University
yif319@lehigh.edu

MS272

Memory Modeling in Multiscale Viscoelastic Materials

The simulation of multiscale viscoelastic materials poses a significant challenge in computational materials science, requiring expensive numerical solvers that can resolve dynamics of material deformations at the microscopic scale. The theory of homogenization offers an alternative approach to modeling, by locally averaging the strains and stresses of multiscale materials and hence eliminating their smaller scales. Although the constitutive law relating the local stress to strain in a material is instantaneous, homogenization introduces history dependence into the constitutive model. Furthermore, this history dependence of the averaged stress on the averaged strain is only known implicitly and must be learned from numerical or experimental data. In the one dimensional setting, we give the first full characterization of the memory dependent homogenized constitutive law of multiscale viscoelastic materials. In one and higher dimensions, we develop a data-driven method, that across a wide range of periodic multiscale materials, accurately predicts their homogenized constitutive laws, thus enabling us to simulate their deformations under forcing. We use the approximation theory of neural operators to provide guarantees on the generalization of our approach.

George Stepaniants
California Institute of Technology
gstepan@caltech.edu

MS272

Solving Partial Differential Equations with Machine Learning Approximants and B-Splines Based on a Dual Variational Principle

In recent work (A. Acharya, Variational principles for nonlinear PDE systems via duality, QAM, 2023), an approach to devise a variational principle for linear and nonlinear PDEs such as Navier–Stokes equations, nonlinear elasticity, and time-dependent heat equation was proposed. In Acharya (2023), the main idea is to treat the primal PDE as a constraint and to invoke an arbitrarily chosen (lends flexibility) auxiliary potential with strong convexity properties to be optimized. This leads to requiring a convex dual functional to be minimized subject to Dirichlet boundary conditions on dual variables, with the guarantee that even PDEs that do not possess a variational structure in primal form can be solved via a variational principle. In this presentation, we adopt ideas from Acharya (2023) to tackle linear and nonlinear problems using machine learning approximants and B-splines. We leverage the dual variational principle, which provides greater flexibility—since the Dirichlet boundary conditions on dual variables can be arbitrarily chosen. The parametric dependence of the dual solution on the choice of Dirichlet boundary conditions and the auxiliary potential can offer practical advantages in training and optimization of neural networks. We apply the dual variational principle to a shallow neural network with RePU activation function and also use B-splines to solve convection-diffusion and transient heat conduction problems. This is joint work with Amit Acharya (CMU).

N. Sukumar
University of California, Davis
nsukumar@ucdavis.edu

MS273

Verification and Validation of Turbulence Modeling for Rans Closures

Scientific Machine Learning (SciML) is revolutionizing critical scientific applications essential for national security by improving efficiency and utilizing available data effectively. Ensuring the credibility of SciML in these critical applications is therefore paramount. The Reynolds Averaged NavierStokes (RANS) equations, fundamental for simulating compressible fluid flows, suffer from model-form errors that limit their applicability. Addressing this, Parish et al [AIAA 2023-2126] introduced a data-driven turbulence modeling strategy to improve RANS models. Utilizing multi-step training on 8 diverse datasets, including channel flows at different Reynolds numbers, duct flow, periodic hill, and hypersonic boundary layers, the study demonstrated the models efficacy. The research focuses on predicting Reynolds stress term discrepancies, emphasizing hyperparameter sensitivity and out-of-distribution dataset performance. Extensive validation efforts ensure the reliability of the ML models in capturing elusive model-form errors. Future phases will further test dataset and hyperparameter sensitivity to ensure credibility and applicability. This study significantly enhances RANS simulations accuracy and reliability, impacting broader fluid dynamics and ML fields. The systematic investigation advances the assessment of credibility in computational simulations involving SciML, offering valuable insights for turbulence modeling and contributing to the refinement of RANS simulations.

Uma Balakrishnan, William J. Rider, Matthew Barone
Sandia National Laboratories
ubalakr@sandia.gov, wjrider@sandia.gov,
mbarone@sandia.gov

Eric Parish
Sandia National Lab

eric.parish9@gmail.com; ejparis@sandia.gov

MS273

Enforcing Learnability and Model Consistency in Data Driven Turbulence Modeling

This work discusses advances made in feature-based data-driven modeling that have carefully considered some of the issues affecting its use for robust, generalizable, predictive modeling of turbulent flows. It details how augmentations are chosen and introduced to existing computational models to improve their deficiencies while also enforcing physical realizability. These augmentations are embedded into the solver as part of the learning process, allowing for the model parameters to be directly learned in terms of the low-fidelity states and enforcing model consistency. Furthermore, it describes an approach for designing features that extract the most salient aspects of the flow and allow for interpretable and learnable augmentations. Specifically, it describes the careful design of the feature-space representation and the idea of localized learning, which are key to providing a balance between robustness and generalizability to unseen configurations. These methodologies are then applied to hypersonic shock-boundary layer interactions and low-speed separated flows to improve heat transfer and separation location predictions, respectively, and demonstrate their capabilities for turbulence modeling.

Niloy Gupta
Computational AeroSciences Lab
University of Michigan
niloy@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS273

Reduced Data-Driven Turbulence Closure for Capturing Long-Term Statistics

The wide range of spatial and temporal scales present in turbulent flow problems forms a computational bottleneck, which large eddy simulations (LES) circumvent by a coarse graining procedure. The effects of the unresolved fluid motions enter the coarse-grained equations as an unclosed subgrid scale term. We present a simple, stochastic, turbulence closure model based on a reduced subgrid scale term. This subgrid scale term is tailor-made to capture the statistics of a small set of spatially-integrated quantities of interest (QoIs), with only one unresolved scalar time series per QoI. In contrast to other data-driven surrogates the dimension of the "learning problem" is reduced from an evolving field to one scalar time series per QoI. We use an *a-posteriori*, nudging approach to find the distribution of the scalar series over time. This approach has the advantage of taking the interaction between the solver and the surrogate into account. A stochastic surrogate parametrization is obtained by random sampling from the found distribution for the scalar time series. We compare the new method to an *a-priori* trained convolutional neural network (CNN). Evaluating the new method is computationally much cheaper and gives similar long-term statistics.

Rik Hoekstra
CWI
Amsterdam
rik.hoekstra@cwi.nl

Daan Crommelin
CWI Amsterdam and University of Amsterdam
Daan.Crommelin@cwi.nl

Wouter Edeling
CWI
wouter.edeling@cwi.nl

MS273

Constrained Machine Learning for Turbulence Modeling

The constants and functions in Reynolds-averaged Navier-Stokes (RANS) turbulence models are coupled. Consequently, modifications of a RANS model often negatively impact its basic calibrations, which is why machine-learned augmentations are often detrimental outside the training dataset. A solution to this is to identify the degrees of freedom that do not affect the basic calibrations and only modify these identified degrees of freedom when recalibrating the baseline model to accommodate a specific application. This approach is colloquially known as the rubber-band approach, which we formally call constrained model recalibration. This talk will identify the degrees of freedom in one-equation and two-equation models that do not affect the log law calibration. By subsequently interfacing data-based methods with these degrees of freedom, we train models to solve historically challenging flow scenarios, including the round-jet/plane-jet anomaly, airfoil stall, secondary flow separation, and recovery after separation. In addition to good performance inside the training dataset, the trained models yield similar or slightly better performance as the baseline model outside the training dataset.

Xiang Yang
Penn State University
xzy48@psu.edu

MS274

Scalable Distributed Mesh Overset Algorithms in p4est

Modern simulations venture increasingly into the realms of computational geometry and object and data location. While practical tasks such as intersecting a distributed set of source/sink or observation objects with the mesh are near trivial on uniform meshes, executing them efficiently on arbitrarily partitioned adaptive meshes is far from obvious and only recently being tackled and scaled, making for current and exciting research. In this talk, we discuss the mathematical design concept of the p4est software for scalable adaptive mesh management, with a special focus on new algorithms for generic data search and location. These features direcly correspond to application needs, and we outline how to benefit from interfacing to the library. p4est is actively used in many projects, including geoscientific applications, and we close with selected examples.

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

Hannes Brandt
University of Bonn
brandt@ins.uni-bonn.de

Tim Griesbach
Universität Bonn, Germany

Institute for Numerical Simulation
tim.griesbach@uni-bonn.de

MS274

Binarized Octree Generation for Deep Levels of Adaptive Mesh Refinement

Cartesian Adaptive Mesh Refinement (AMR), built upon octree structures, is a proven technique for multiscale simulations. While explicit octrees have traditionally been employed in AMR implementations, their memory-intensive nature and inefficiency in traversal hinder performance for extreme-scale simulations requiring deep dynamic adaptations. To address these challenges, we propose a binarized-octree generation method based on an implicit pointerless approach. This method employs Morton encoding for node indexing and a red-black tree for dynamic insertion and deletion of elements. The red-black tree ensures strict adherence to the Z-order curve, resolving data locality issues caused by hash table collisions in hashed linear-octree methods. Additionally, its inherent sorted structure simplifies domain partitioning for parallel computations. By leveraging bitwise representations throughout, our method surpasses hardware constraints on octree depth, enhancing runtime performance with only a slight increase in memory usage.

Inanc Senocak
Mechanical Engineering and Materials Science
University of Pittsburgh
senocak@pitt.edu

MS275

A Fictitious Domain Approach for the Finite Element Discretization of Fsi

In this talk I present some recent advances on a fictitious domain approach for the discretization of fluid-structure interaction problems proposed in [1]. This formulation is based on the introduction of a distributed Lagrange multiplier so that the problem fits in the framework of saddle point systems and possible stable choices of the finite element spaces which have been investigated. Our formulation allows for solving the Navier-Stokes equation and the elasticity equation on meshes independent of each other, at the price of computing a coupling term which involves test funtions defined on both meshes. A discussion on how to deal carefully with such term will also be presented, see [2] and [3]. We shall show also that the formulation is robust with respect small cells. The presented results have been obtained in collaboration with Daniele Boffi, Fabio Credali and Najwa Alshehri. [1] D. Boffi and L. Gastaldi, *A fictitious domain approach with Lagrange multiplier for fluid-structure interactions*, Numer. Math. 135 (2017), 711-732 [2] D. Boffi, F. Credali, and L. Gastaldi, *On the interface matrix for fluid-structure interaction problems with fictitious domain approach*, CMAME 401 (2022) [3] D. Boffi, F. Credali, and L. Gastaldi, *Quadrature error estimates on non-matching grids in a fictitious domain framework for fluid-structure interaction problems*, arxiv.org/abs/2406.03981

Lucia Gastaldi
Università di Brescia
Dipartimento di Matematica
lucia.gastaldi@ing.unibs.it

MS275

Polynomial Interpolation on the Square and Cube

Polynomial interpolation is a ubiquitous tool in numerical analysis. We present a method for deriving collocation points that attain faster convergence in integration and interpolation errors than standard tensor-product Chebyshev points in 2 and 3 dimensions, with ready generalization to higher dimensions. The natural polynomial basis associated to the points is efficient in "Euclidean" degree, which is the relevant measure for interpolating or integrating a generic analytic function. In addition to this efficiency advantage, there are natural algorithms based on the fast Fourier transform for working with the points; the conversion of point values to polynomial basis coefficients (and vice versa) can be performed in O(n log n) time. We compare the results of interpolating and integrating with these points to other "reduced" sets of points.

Richard C. Greene
New Jersey Institute of Technology
rcg4@njit.edu

MS275

Median Filters for Anisotropic Wetting / Dewetting Problems

We present new level set methods for multiphase, anisotropic (weighted) motion by mean curvature of networks, focusing on wetting-dewetting problems where one out of three phases is stationary – a good testbed for checking whether complicated junction conditions are correctly enforced. The new schemes are vectorial median filters: The level set values at the next time step are determined by a sorting procedure performed on the most recent level set values. Detailed numerical convergence studies are presented, showing that the correct angle conditions at triple junctions (which include torque terms due to anisotropy) are indeed indirectly and automatically attained. Other standard benefits of level set methods, such as sub-grid accuracy on uniform grids via interpolation and seamless treatment of topological changes, remain intact.

Jiajia Guo
Department of Mathematics
University of Michigan
jiajiag@umich.edu

Selim Esedoglu
University of Michigan
esedoglu@umich.edu

MS275

A Narrow Band Finite Element Method for the Level Set Equation

A finite element method is introduced to track interface evolution governed by the level set equation. The method solves for the level set indicator function in a narrow band around the interface. An extension procedure, which is essential for a narrow band level set method, is introduced based on a finite element projection combined with the ghost penalty method. This procedure is formulated as a linear variational problem in a narrow band around the surface, making it computationally efficient and suitable for rigorous error analysis. The extension method is com-

bined with a discontinuous Galerkin space discretization and a BDF time-stepping scheme. We discuss the stability and accuracy of the extension procedure and evaluate the performance of the resulting narrow band finite element method for the level set equation through numerical experiments.

Maxim A. Olshanskii
Department of Mathematics
University of Houston
molshan@math.uh.edu

Arnold Reusken
Numerical Mathematics
RWTH Aachen University, Aachen, Germany
reusken@igpm.rwth-aachen.de

Paul SCHWERING
RWTH Aachen
schwering@igpm.rwth-aachen.de

MS276

Preserving Structure in Model Reduction on Manifolds

Capturing and preserving physical properties, e.g., system energy, stability and passivity, using data-driven methods is currently a highly-researched topic in surrogate modeling. To ensure that the desired physical properties are retained, structure-preserving projection techniques are used in the field in model reduction (MOR). In this talk, we present structure-preserving MOR with nonlinear projections, which are needed for problems with slowly decaying Kolmogorov-n-widths. To precisely define and highlight the quantities that we would like to retain, we start with a formulation of initial value problems on manifolds, which we consider as the full-order model (FOM). Already at this level, we define what we mean by adding structure to the FOM and how this can be detailed geometrically. This formalism allows to introduce a novel projection technique, the generalized manifold Galerkin (GMG). Once that we have derived the geometric formulation, we focus on data-driven ansatzes to realize the presented reduction methods. In this part of the talk, we will connect several existing techniques for data-driven realizations with GMG projections

Patrick Buchfink
University of Twente/University of Stuttgart
p.buchfink@utwente.nl

Silke M. Glas
University of Twente
s.m.glas@utwente.nl

Bernard Haasdonk
University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de

Benjamin Unger
University of Stuttgart, Germany
benjamin.unger@simtech.uni-stuttgart.de

MS276

Structure-Preserving Model Reduction of Conservative Pdes Via Lifting Transformations

Existing model reduction techniques for high-dimensional full-order models of conservative PDEs encounter computational bottlenecks when dealing with systems featuring non-polynomial nonlinearities. This talk will present a nonlinear model reduction method that employs lifting variable transformations to derive structure-preserving reduced-order models for conservative PDEs with general nonlinearities. The advantages of the proposed structure-preserving lifting approach will be illustrated by various examples, which include the sine-Gordon equation, the nonlinear wave equation with exponential nonlinearity, and the Klein-Gordon-Zakharov equation.

Harsh Sharma
University of California San Diego
hasharma@ucsd.edu

Juan Diego Draxl
UC San Diego
jdraxl@ucsd.edu

Boris Kramer
University of California San Diego
bmkramer@ucsd.edu

MS276

Tensor Parametric Operator Inference with Hamiltonian Structure

This work presents a tensor-based approach to constructing data-driven reduced-order models corresponding to semi-discrete partial differential equations. By expressing parameter-varying operators with affine dependence as contractions of a generalized parameter vector against a constant tensor, this method leverages the operator inference framework to capture parametric dependence in the reduced-order model via the solution to a convex, least-squares optimization problem. This leads to a simple and straightforward implementation which directly extends to learning parametric operators with symmetry constraints, a key feature required for constructing accurate surrogates of systems with a Hamiltonian structure. The method is demonstrated on a scalar heat equation with variable diffusion coefficient and a wave equation with variable wave speed.

Arjun Vijaywargiya
University of Notre Dame
arjunveejay@nd.edu

MS276

Gradient Preserving Operator Inference for (Port-)Hamiltonian Systems

The proper orthogonal decomposition reduced-order model (POD-ROM) has been widely used as a computationally efficient surrogate model in large-scale numerical simulations of complex systems. However, when applied to a (port-)Hamiltonian system, a direct application of the POD method can destroy the underlying structure and produce unphysical results. In this talk, we will explore the use of operator inference to learn ROMs in a data-driven setting. These ROMs could preserve the (port-)Hamiltonian structure, leading to robust and efficient simulations. Numerical error estimations will be derived, and the effectiveness of the ROMs will be demonstrated through several numerical examples.

Yuwei Geng, Jasdeep Singh
University of South Carolina

ygeng@email.sc.edu, jasdeep@email.sc.edu

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Boris Kramer
University of California San Diego
bmkramer@ucsd.edu

Zhu Wang
Department of Mathematics
University of South Carolina
wangzhu@math.sc.edu

MS277

Small Singular Values Can Increase in Lower Precision

Motivated by the popularity of stochastic rounding in the context of machine learning and the training of large-scale deep neural network models, we consider stochastic nearness rounding of real matrices A with many more rows than columns. We provide novel theoretical evidence, supported by extensive experimental evaluation that, with high probability, the smallest singular value of a stochastically rounded matrix is well bounded away from zero – regardless of how close A is to being rank deficient and even if A is rank-deficient. In other words, stochastic rounding *implicitly regularizes* tall and skinny matrices A so that the rounded version has full column rank. Our proofs leverage powerful results in random matrix theory, and the idea that stochastic rounding errors do not concentrate in low-dimensional column spaces.

Christos Boutsikas, Gregory Dexter, Linkai Ma
Purdue University
boutsikas8@gmail.com, gdexter@purdue.edu,
ma856@purdue.edu

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

Petros Drineas
Purdue University
Computer Science
pdrineas@purdue.edu

MS277

Near-optimal Hierarchical Matrix Approximation from Matrix-vector Products

We describe a randomized algorithm for producing a near-optimal hierarchical off-diagonal low-rank (HODLR) approximation to an $n \times n$ matrix $\mathbf{A}$, accessible only though matrix-vector products with $\mathbf{A}$ and $\mathbf{A}^\mathsf{T}$. Our algorithm can be viewed as a robust version of widely used "peeling" methods for recovering HODLR matrices and is, to the best of our knowledge, the first matrix-vector query algorithm to enjoy theoretical worst-case guarantees for approximation by any hierarchical matrix class. To control the propagation of error between levels of hierarchical approximation, we introduce a new perturbation bound for low-rank approximation, which shows that the widely used Generalized Nystrm method enjoys inherent stability when implemented with noisy matrix-vector products. We also introduced a novel randomly perforated matrix sketching method to further control the error in the peeling algorithm.

Tyler Chen
New York University
tyler.chen@nyu.edu

Feyza Duman Keles
NYU
fd2135@nyu.edu

Diana Halikias
Cornell University
dh736@cornell.edu

Cameron Musco
University of Massachusetts Amherst
cmusco@cs.umass.edu

Christopher Musco
New York University
cmusco@nyu.edu

David Persson
EPFL
david.persson@epfl.ch

MS277

Optimal Experimental Designs Via Column Subset Selection

We tackle optimal sensor placement for Bayesian linear inverse problems by establishing connections to the Column Subset Selection Problem (CSSP). We build on the Golub-Klema-Stewart (GKS) approach which involves computing the truncated Singular Value Decomposition (SVD) followed by a pivoted QR factorization on the right singular vectors. We study the effects of using the Federov exchange rule, greedily swapping sensors while improving the objective, after a GKS-style initialization. Theoretical guarantees on the number of swaps are established. Numerical experiments on model inverse problems demonstrate the performance of our approaches.

Srinivas Eswar
Mathematics and Computer Science Division
Argonne National Laboratory
seswar@anl.gov

Amit Subrahmanya
Virginia Tech
amitns@vt.edu

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS277

Randomized Sampling by Tagging for Rank-Structured Matrix Compression

Many applications in scientific computing involve matrices

that are dense but "data-sparse,' having some exploitable low-rank structure that enables compression without significant loss of information. Rank-structured matrices are data-sparse matrices that can be tessellated into subblocks, which are either small enough to apply deterministic algorithms or else well-approximated by matrices of low rank. For large problem sizes, randomized sampling of the input matrix by a structured random test matrix, or randomized sketching, has proven effective for rank-structured matrix compression. We introduce a new sampling technique called tagging that is suitable for large black-box applications in which the entries of the input matrix are not accessed directly, but rather through fast matrix-vector or matrix-matrix multiplication. We discuss the improved efficiency of randomized sampling with tagging, as well as the accuracy of the overall compression scheme, for strongly admissible block low-rank (BLR) matrices with shared bases across admissible blocks within the same block-row or block-column, termed uniform BLR. We also describe how tagging can be used for efficient matrix reconstruction by leveraging previously computed sample matrices.

Katherine J. Pearce, James Levitt
University of Texas at Austin
Oden Institute
katherine.pearce@austin.utexas.edu,
jlevitt@oden.utexas.edu

Anna Yesypenko
University of Texas at Austin
annayesy@utexas.edu

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

MS277

Randomized Strong Rank-Revealing QR

Finding the rank-revealing QR(RRQR) factorization of a $m \times n$ matrix M is to determine a permutation matrix P and an integer k, such that the factorization

$$MP = Q \begin{pmatrix} R_{11} & R_{12} \\ O & R_{22} \end{pmatrix},$$

reveals the numerical rank of M: the $k \times k$ submatrix R_{11} is well-conditioned and $\|R_{22}\|_2$ is small. This problem is of interest in many applications. To compute such a factorization, the cost of the existing algorithm, strong RRQR algorithm, can be one and a half of that of QRCP, which is substantially more expensive than unpivoted QR. Recently, random sketching becomes a powerful tool for speeding up many numerical linear algebra algorithms. In this study, we take advantage of random sketching and develop a randomized algorithm to compute a strong RRQR factorization, whose cost is slightly larger than that of unpivoted QR. In our algorithm, the permutation matrix is obtained by doing deterministic strong RRQR on the sketched matrix $M^{sk} = \Omega M$, where Ω is a sketching matrix. We proved that the strong rank-revealing property holds for both the factorization of M and that of M^{sk}. Based on that, we also propose a low rank approximation method. It provides a rank-ℓ approximation with L_2 norm errors similar as that of randomized SVD, while additional providing approximations for the ℓ leading singular values.

Zhipeng Xue
EPFL
zhipeng.xue@epfl.ch

Laura Grigori
EPFL and PSI, Switzerland
laura.grigori@epfl.ch

MS278

Bi-fidelity Variational Auto-encoder: A Bi-fidelity Deep Generative Model to Understand Uncertainty Propagation

Understanding uncertainty propagation through physical emulators is one of the key problems in computational mechanics. Our work proposes a variational auto-encoder (VAE) based deep generative model to transfer the information of low-fidelity data for high-fidelity uncertainty reconstruction. An effective algorithm is provided to train the bi-fidelity VAE structure with extremely unbalanced bi-fidelity data. We also show the theoretical support for the bi-fidelity VAE model, with objective function being a lower bound of the high-fidelity log likelihood. Besides, a theory of bi-fidelity variational information bottleneck is presented, for the purpose of bi-fidelity transfer learning. We prove that maximizing our model's objective function is maximizing a lower bound of the bi-fidelity information bottleneck result. In the end, the implementation of bi-fidelity VAE shows a promising future of our method.

Nuojin Cheng, Osman Asif Malik, Subhayan De
University of Colorado, Boulder
nuojin.cheng@colorado.edu, osman.malik@colorado.edu,
subhayan.de@colorado.edu

Stephen Becker
University of Colorado at Boulder
stephen.becker@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS278

Dictionary Learning and Inpainting for Data Assimilation with Partial Observation

Generative artificial intelligence extends beyond its success in image/text synthesis, proving itself a powerful uncertainty quantification (UQ) technique through its capability to sample from complex high-dimensional probability distributions. However, existing methods often require a complicated training process, which greatly hinders their applications to real-world UQ problems, especially in dynamic UQ tasks where the target probability distribution evolves rapidly with time. To alleviate this challenge, we have developed a scalable, training-free score-based diffusion model for high-dimensional sampling. We incorporate a parallel-in-time method into our diffusion model to use a large number of GPUs to solve the backward stochastic differential equation and generate new samples of the target distribution. Moreover, we also distribute the computation of the large matrix subtraction used by the training-free score estimator onto multiple GPUs available across all nodes. Compared to existing methods, our approach completely avoids training the score function, making it capable of adapting to rapid changes in the target probability distribution. We showcase the remarkable strong and weak scaling capabilities of the proposed method on

the Frontier supercomputer, as well as its uncertainty reduction capability in hurricane predictions when coupled with AI-based foundation models.

Siming Liang
Florida State University
sliang@fsu.edu

Hoang A. Tran
Oak Ridge National Laboratory
Computer Science and Mathematics Division
tranha@ornl.gov

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

Hristo G. Chipilski
Florida State University, Department of Scientific Computing
hchipilski@fsu.edu

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS278

A Training-Free Conditional Diffusion Model for Learning Stochastic Dynamical Systems

This study introduces a training-free conditional diffusion model for learning unknown stochastic differential equations (SDEs) using data. The proposed approach addresses key challenges in computational efficiency and accuracy for modeling SDEs by utilizing a score-based diffusion model to approximate their stochastic flow map. Unlike the existing methods, this technique is based on an analytically derived closed-form exact score function, which can be efficiently estimated by Monte Carlo method using the trajectory data, and eliminates the need for neural network training to learn the score function. By generating labeled data through solving the corresponding reverse ordinary differential equation, the approach enables supervised learning of the flow map. Extensive numerical experiments across various SDE types, including linear, nonlinear, and multi-dimensional systems, demonstrate the versatility and effectiveness of the method. The learned models exhibit significant improvements in predicting both short-term and long-term behaviors of unknown stochastic systems, often surpassing baseline methods like GANs in estimating drift and diffusion coefficients.

Yanfang Liu
Middle Tennessee State University
yanfang.liu@mtsu.edu

Yuan Chen
SIAM
chen.11050@osu.edu

Dongbin Xiu
Ohio State University
xiu.16@osu.edu

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS278

Error Estimate of a Training-free Diffusion Model for Density Estimation

We present convergence analysis and error estimates of a training-free diffusion model for density estimation. Unlike existing diffusion models that train neural networks to learn the score function, we developed a training-free score estimation method that mini-batch-based Monte Carlo estimators to directly approximate the score function at any spatial-temporal location an ordinary differential equation (ODE), corresponding to the reverse-time stochastic differential equation (SDE). In this way, the numerical error of the diffusion model can be analyzed by exploiting the error bounds of the Monte Carlo error. We observe that the total error based on Monte Carlo score estimation depends on the dimensionality of the problem. To improve the error estimates, we assume the target distribution can be approximated by a Gaussian mixture model. Under this assumption, the score function can be computed exactly without any discretization error, and the dimension dependence of the total error is included in the error between the Gaussian mixture model and the exact target distribution. We present numerical examples to verify the obtained error bounds.

Pengjun Wang
Auburn University
pzw0030@auburn.edu

Yanzhao Cao
Department of Mathematics & Statistics
Auburn University
yzc0009@auburn.edu

Toan Huynh
Florida State University
th24ba@fsu.edu

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

Feng Bao
Florida State University
bao@math.fsu.edu

Zezhong Zhang, Minglei Yang
Oak Ridge National Laboratory
zhangz2@ornl.gov, yangm@ornl.gov

MS280

Kernel Methods and Pinns for Solving Rough Partial Differential Equations

Following the promising success of kernel methods in solving non-linear partial differential equations (PDEs), we investigate the application of Gaussian process methods to solve PDEs with rough forcing terms. We introduce an optimal recovery scheme defined by a Reproducing Kernel Hilbert Space (RKHS) of functions of greater regularity than that of the PDEs solution. We present the theoretical framework and prove convergence guarantees for the recovery of solutions to the PDE. We also show how this framework can be applied to Physics Informed Neural Networks (PINNs). We illustrate its application numerically

to problems arising in nonlinear time-dependent stochastic partial differential equations.

Matthieu Darcy
Caltech
mdarcy@caltech.edu

MS280

Calibration and Uncertainty Quantification of Carbon Cycle Models Using Bayesian Emulation and History Matching

Simulators are a fundamental tool to assess ecosystem dynamics. Coupled with observational data, these models can infer unobserved ecosystem properties. An example of this is the carbon cycle model DALEC Crop, which models the effect of Nitrogen fertilisation on wheat growth in a field over a crop-growing season. DALEC Crops parameters are currently calibrated pixel-by-pixel using a computationally expensive model-data fusion framework, CARDAMOM. In addition to high computational costs, the current framework does not allow for exploring the effects of structural discrepancy and parametric uncertainty. This limits our understanding of the crop fields physical processes.
To facilitate quick model calibration and a thorough analysis of the uncertainties involved with DALEC Crop, this work presents history matching combined with emulation as an alternative to the current calibration framework. We use observations of LeafArea Index when history matching to find the set of input combinations for which the simulator gives acceptable matches to observed data. We then use the input space found with a DALEC Crop forward run to produce plausible yield predictions, skipping DALECs current computationally costly model calibration to demonstrate the potential to reduce computational costs by using suitable statistical methodology. Finally, we explore the differences between the outputs of the CARDAMOM calibration and the history matching.

Nina A. Fischer
University of Edinburgh
School of Mathematics
N.Fischer@ed.ac.uk

MS280

Polynomial-Model-Based Optimization: Polynomial Surrogates for Guiding Black-Box Optimization Processes

Black-box optimization is a challenging task that occurs in a wide range of applications and across all disciplines. Common tasks include hyper-parameter selection in machine learning algorithms, computer simulations or even laboratory experiments. As these problems are typically costly to evaluate in terms of time and computational resources, black-box optimization algorithms aim to minimize this cost by reducing the number of necessary evaluations. Polynomial-model-based optimization (PMBO) is a novel surrogate-based method that alternates polynomial regression with Bayesian optimization steps. Besides finding the optimum, PMBO delivers an analytic expression of the approximation of the objective function. The simple nature of polynomials opens the opportunity for interpretation and analysis of the inferred surrogate model, providing a macroscopic perspective on the landscape of the objective function. The methodology of PMBO will be illustrated and compared with state-of-the-art optimization algorithms on the benchmark suite "BBOB - Black-Box Optimization Benchmark" in the deterministic and noisy setting.

Janina Schreiber
CASUS - Center for Advanced Systems Understanding, Helmholtz
j.schreiber@hzdr.de

MS280

Conformal Prediction under Lvy-Prokhorov Distribution Shifts: Robustness to Local and Global Perturbations

Conformal prediction provides a powerful framework for constructing prediction intervals with finite-sample guarantees, yet its robustness under distribution shifts remains a significant challenge. We address this limitation by modeling distribution shifts using Lvy-Prokhorov (LP) ambiguity sets, which capture both local and global perturbations. We provide a self-contained overview of LP ambiguity sets and their connections to popular metrics such as Wasserstein and Total Variation. We show that the link between conformal prediction and LP ambiguity sets is a natural one: by propagating the LP ambiguity set through the scoring function, we reduce complex high-dimensional distribution shifts to manageable one-dimensional distribution shifts, enabling exact quantification of worst-case quantiles and coverage. Building on this analysis, we construct robust conformal prediction intervals that remain valid under distribution shifts, explicitly linking LP parameters to interval width and confidence levels. Experimental results on real-world datasets demonstrate the effectiveness of the proposed approach.

Julie Zhu
MIT
qianyu_z@mit.edu

Oliver Wang, Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
olivrw@mit.edu, ymarz@mit.edu, ymarz@mit.edu

Liviu Aolaritei, Michael I. Jordan
University of California, Berkeley
liviu.aolaritei@berkeley.edu, jordan@cs.berkeley.edu

MS281

Bayesian Optimal Design of Experiments for Learning Markovian Homogenized Models in Viscoplasticity

A material constitutive model maps the materials deformation history to its stress. The constitutive parameters are often not measured directly; instead, one needs to invert for their values using experimental measurements of the material response, such as the history of material deformation under prescribed loading conditions. To learn the constitutive law of a viscoplastic material, the experimental data must be rich enough to capture a range of complex material behavior, such as elasticity, plastic flow, and history-dependency. To this end, we seek to optimize experimental designs for learning a homogenized Markovian constitutive model for viscoplasticity. The studied model uses internal state variables (ISVs) to incorporate relevant microscale information, such as microstructure and dislocation, into the macroscale model, leading to memory effects in the material response. We use variational inference methods to optimize the design of dynamic compression experiments under various uncertainties. This gives us valuable insights into the effectiveness of experiments, including trade-offs in

information content for learning different aspects of material behaviors.

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

Kaushik Bhattacharya
Howell N. Tyson Sr. Professor of Mechanics
California Institute of Technology
bhatta@caltech.edu

MS281

An Energy-Stable Machine-Learning Model of Non-Newtonian Hydrodynamics with Molecular Fidelity

We introduce a machine-learning-based approach for constructing a continuum non-Newtonian fluid dynamics model directly from a micro-scale description. To faithfully retain molecular fidelity, we establish a micro-macro correspondence via a set of encoders for the micro-scale polymer configurations and their macro-scale counterparts, a set of nonlinear conformation tensors. The dynamics of these conformation tensors can be derived from a generalized extendable energy functional structure, and be directly learned from the micro-scale model with clear physical interpretation. The final model, named the deep non-Newtonian model (DeePN^2), takes the form of conventional non-Newtonian fluid dynamics models and ensures energy stability. Numerical results demonstrate the accuracy and robustness of DeePN^2.

Pei Ge, Huan Lei
Michigan State University
gepei@msu.edu, leihuan@msu.edu

MS281

Inferring a Rheology for Sea Ice from Data Generated with a Discrete Element Method

Sea ice is a fundamental component of the climate system that is generally treated as a continuum fluid in general climate models. A key component in continuum models for sea ice is its rheology, which establishes a relationship between strain-rate, the Cauchy stress tensor, and other variables such as ice concentration and thickness. Traditionally, these models have been derived analytically based on speculative principles. Here, we present a general framework for inferring rheologies from data based on a representation of constitutive laws in terms of scalar functions that depend on the principal invariants of the strain-rate tensor. These scalar functions can be represented with neural networks that are trained with data. Then, by coupling PyTorch and the finite-element module Firedrake, we can solve these PDE models involving a neural network. We will present some initial findings involving data generated with the discrete element method SubZero used for modelling sea-ice.

Gonzalo Gonzalez de Diego
New York University (NYU)
gg2924@nyu.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS281

Particle-Continuum Multiscale Modeling of Sea Ice Floes

Sea ice profoundly influences the polar environment and the global climate. Traditionally, sea ice has been modeled as a continuum under Eulerian coordinates to describe its large-scale features, using, for instance, viscous-plastic rheology. Recently, Lagrangian particle models, also known as the discrete element method models, have been utilized for characterizing the motion of individual sea ice fragments (called floes) at scales of 10 km and smaller, especially in marginal ice zones. This paper develops a multiscale model that couples the particle and the continuum systems to facilitate an effective representation of the dynamical and statistical features of sea ice across different scales. The multiscale model exploits a Boltzmann-type system that links the particle movement with the continuum equations. Notably, the particle model characterizing the sea ice floes is localized and fully parallelized in a framework that is sometimes called superparameterization, which significantly improves computational efficiency. Numerical examples demonstrate the effective performance of the multiscale model. Additionally, the study demonstrates that the multiscale model has a linear-order approximation to the truth model.

Quanling Deng
Australian National University
quanling.deng@anu.edu.au

Samuel N. Stechmann, Nan Chen
University of Wisconsin-Madison
stechmann@wisc.edu, chennan@math.wisc.edu

MS282

Coupling Multi-Physics Applications Using the Parallel, Adaptive Library ForestClaw

We present several new features in the ForestClaw library for parallel, adaptive simulations on a tree-based hierarchy of logically Cartesian grids. Major new features are full support for 3d octree refinement, memory saving refinement strategies that aim to prevent excessive intermediate memory use during refinement, restart features, and new coupling capabilities that allow simulations on distinct meshes to communicate at runtime. Other new features include addition of numerical gauges to match observation data with numerical output, and refinement based on regions. Results demonstrating these capabilities in coupling of acoustic gravity waves between the upper atmosphere and the ionosphere will be demonstrated.

Donna Calhoun, Scott Aiton
Boise State University
donnacalhoun@boisestate.edu, scottaiton@boisestate.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS282

Space-Time Adaptive Finite Elements for

Maxwell's Equations

We provide space-time adaptive discretization of Maxwell's equations in two dimensions using finite element exterior calculus on simplicial meshes. Our main contribution is in providing this adaptive h refinement in conjunction with space-time finite elements specified using a de Rham sequence of Whitney basis forms for Maxwell's equations. We formulate Maxwell's equations as a (p, E, H) system where p is an electrical variable akin to pressure, E is the electrical field and H is the magnetic flux density. p is introduced for enabling the discrete divergence of E to be exactly specified. For this problem, we provide some a priori error analysis and a discussion of our preliminary work with a posteriori residual estimators. We derive our estimators by adapting the classical ones for the Hodge Laplacian problem. We shall also present some analysis on the optimality of our estimators, and validate the theory through computational examples for space-time adaptivity for model problems in two spatial dimensions.

Ishani Choudhary
PhD student at IIIT-D
ishanic@iiitd.ac.in

MS282

Stopping Power and the Quantum Boltzmann Model

On December 5th, 2022, the national ignition facility achieved better than break-even fusion, putting in 2MJ of energy at the surface of tritium and deuterium-filled capsule, and they got out 3MJ. This was the first time humanity demonstrated that we could reproduce the processes happening in our sun, providing a possible path to a cheap, clean energy future. Motivated by the challenges with NIF, various interesting problems have emerged where innovation regarding numerical methods could greatly enhance our ability to model key processes that happen on NIF. One example is the study of stopping power for the alpha particles, the energetic helium particles produced by deuterium and tritium fusing. Stopping power refers to the ability of a material to slow down or absorb charged particles, such as electrons or protons, as they pass through it. It is often measured as energy loss per unit distance. The energetic alpha particles primarily interact with the hot electrons, and slowing the alpha particle imparts energy to the electrons. Because of the unique energy landscape on NIF, the electrons are best described by the Quantum Boltzmann Equation (QBE). The QBEs nonlocality and the models high dimensional nature make this a challenging system to solve efficiently. In this talk, we discuss the need for this model, the origins of this model, and introduce a simple method for solving the 1D 1V collisionless QBE based on Strang Splitting, a WENO Semi-Lagrangian update, an FFT-based approach to reduce the non-local convolution to a non-local ODE, followed by a second SL step. We explore the nature of what happens to the solutions of the model as the non-dimensional Planck constant approaches the strongly quantum realm, and we study what happens to the rank of the solution. We use this to motivate our ongoing work on low-rank solvers for the QBE.

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

William A. Sands, Stephen White
Michigan State University
wsands@udel.edu, whites73@msu.edu

MS282

Adaptive Discontinuous Galerkin Method for Relativistic Vlasov-Maxwell System

We present an noval adaptive discontinuous Galerkin (DG) method and its GPU-accelerated implementations for the relativistic Vlasov-Maxwell system (RVMS). RVMS is a kinetic model of collionless plasma. It describes the evolution of plasma distribution functions and interacting electromagnetic fields. One of the major challenges for conventional mesh-based methods is that the timestep size for stable numerical methods is limited to a small number when system approaches the relativistic regime. Our method follows operator splitting approach proposed in [Rossmanith and Seal, *A positivity-preserving high-order semi-Lagrangian discontinuous Galerkin scheme for the Vlasov-Poisson equations* (2011)] to update the free streaming branch with a semi-Lagrangian DG method that is unconditionally stable for any timestep size, and update the acceleration branch with a one-step local spacetime DG method. Both branches are capable of adaptive mesh refinement and structure-preserving corrections. We implement this noval DG method in a distributed shared memory code, which is robust and scalable on production-ready high performance computing clusters equipped with GPU acceleration. We provide a variety of benchmark problems to veriy the numerical method and its implementations, and offer profiling statistics and insights for future development.

Yifan Hu, James A. Rossmanith
Iowa State University
Department of Mathematics
yifanh@iastate.edu, rossmani@iastate.edu

MS282

High Fidelity Simulation of Multi-Species, Kinetic Plasma Dynamics with the Numerical Flow Iteration

A fundamental understanding of the evolution of kinetic instabilities and turbulence in plasma is crucial to design, e.g., nuclear fusion devices. Plasma evolution is modeled through the Vlasov equation living in the up to six-dimensional phase-space. For each species present an additional Vlasov equation has to be solved. The different distribution functions for each species couple to each other through the self-induced electro-magnetic field as well as collisions. Even modern large-scale supercomputers struggle resolving his problem due to its high-dimensionality. Additionally solutions are develop of fine, multi-scale structures over time, which are relevant to the plasma dynamics, and the velocity support may change over time. Previously the authors introduced a novel approach, the Numerical Flow Iteration (NuFI), to solve the Vlasov equation in the electro-static limit. This approach uses the characteristic method to directly evaluate the distribution function via an iterative-in-time scheme for the flow map. This avoids making additional discretization errors, allows for better structure-preservation while reducing the memory-requirement of the computation by several orders of magnitude at the price of a higher computational load. In this work we present approaches to extend NuFI to handle multiple species systems and discuss the implications of this extensions on the computational performance of this

approach.

Paul Wilhelm
RWTH Aachen University
paul_wilhelm@gmx.de

Fabian Orland
IT Center
RWTH Aachen University
orland@itc.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS283

A Mixed Formulation for the Linear Regularized 13-Moment Equations: Analysis and Numerical Solution

This talk presents a mixed finite element formulation for the linear regularized 13-moment equations of non-equilibrium gas dynamics. As an extension of the classical fluid equations, moment models are robust and have been frequently used in recent years, yet they are challenging to analyze due to their additional equations. By effectively grouping variables, we identify a 2-by-2 block structure, allowing the analysis of the continuous well-posedness within the abstract LBB framework. Due to the unique tensorial structure, we derive new theoretical tools, such as a generalized Korn inequality for symmetric and trace-free 2-tensors. Additionally, we showcase the numerical implementation within an open-source solver, addressing both stabilization and the use of generalized Taylor-Hood elements. The proposed method is validated using test cases with exact solutions. Furthermore, if time allows, we will explore how the moment hierarchy can be extended to derive higher-order tensorial equations, requiring more sophisticated analytical tools than those used in traditional elasticity or Stokes theory.

Lambert Theisen
Applied and Computational Mathematics
RWTH Aachen University
lambert.theisen@rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS284

Hyper-Differential Sensitivity Analysis with Respect to Model Discrepancy

Computational models have made great process toward simulating complex physical phenomena. With this has come a significant increase in the computational cost and complexity of the model. This makes high-fidelity models impractical in many query analysis such as optimization. To circumvent this challenge, a low-fidelity model may be constructed for which optimization is more amenable. Examples of low-fidelity models include averaging small length scales, omission of physical processes, projection into lower dimensions, and data-driven approximations. However, model error propagates through the optimization resulting in suboptimal solutions. We consider the use of hyper-differential sensitivity analysis to learn improved optimization solutions using limited offline evaluations of the high-fidelity model. We introduce a representation of model discrepancy, the difference between high and low-fidelity models, and calibrate it in a Bayesian framework. The model discrepancy posterior is propagated through the optimization problem using post-optimality sensitivity analysis to produce a posterior distribution for the optimization solution. By leveraging tools from large-scale optimization, in a Bayesian formulation, we present an interpretable and efficient approach to learn higher-fidelity optimization solutions using limited evaluations of the high-fidelity model. Thanks to our Bayesian framing, we also enable optimal experimental design to guide the high-fidelity model runs.

Joseph Hart, Bart G. Van Bloemen Waanders
Sandia National Laboratories
joshart@sandia.gov, bartv@sandia.gov

MS284

Scientific Machine Learning in Enhancing Model Predictability Through Model Form Error Quantification

In computational simulations, the development of surrogate models is critical for efficiently predicting experimental outcomes. Previous research has validated the effectiveness of Gaussian processes in modeling discrepancies between simulation predictions and experimental data. Nonetheless, there is an increasing interest in leveraging machine learning (ML) and artificial intelligence (AI) techniques to enhance prediction accuracy and computational efficiency. This study aims to investigate the application of ML algorithms, specifically Multi-Layer Perceptron (MLP), Recurrent Neural Network (RNN), and Convolutional Neural Network (CNN), in constructing surrogate models and identifying the optimal combination of model parameters (density, velocity, and location) to accurately predict experimental data. This approach aimed to identify the parameter combinations and demystify the underlying physics to provide more accurate predictions. The introduction of the discrepancy term should further enhance the model's accuracy. The ML/AI-based discrepancy modeling is expected to demonstrate promising improvements over traditional Gaussian process methods, suggesting a novel pathway for refining surrogate model predictions. *SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Kwesi Ohene-Obeng
The University of Texas at El Paso
kaoheneoben@miners.utep.edu

Kathryn Maupin
Sandia National Laboratories
kmaupin@sandia.gov

MS285

A Computational Exploration of a Hierarchical Noise Model for Active Transport.

A closed-form solution of differential equations (deterministic or stochastic) only exists for special cases. As such, computational studies are carried out to analyze differential equation models of various biological and chemical processes. Active transport is one such vital process whose mathematical models are usually analyzed computationally. Following suit, we used the Euler-Maruyama method to simulate 10^9 trajectories of a stochastic differential equation model of active transport in a high-performance cluster and calculated ensemble average statistics. Here,

we will discuss how this ensemble average statistics of the study mimics three key statistical features of active transport-transient super diffusion, non-Gaussian probability density function, and non-monotonic evolution of the non-Gaussian parameter. We further extend the computational study of the model by changing some parameter values to include diffusing diffusivities and show that the model can accommodate other statistical signatures like transient sub-diffusion, which is commonly observed during transport in complex environments.

Dipesh Baral
Washington State University, Pullman, WA
dipesh.baral@wsu.edu

Nikolaos Voulgarakis
Washington State University
n.voulgarakis@wsu.edu

MS285

Polynomial Approximation for Nonlinear Model Reduction by Moment Matching

Many physical applications involve modeling a complex system as a high-dimensional system of ordinary differential equations. The high dimensionality of these models introduces significant computational challenges in trying to analyse the system's behaviour, which is often known as the "curse of dimensionality'. One way to alleviate the complexity of such problems is by using reduced-order modeling to create a low-dimensional system that preserves important characteristics. This talk focuses on nonlinear model reduction using moment matching for nonlinear dynamical systems whose input is generated by a signal generator system. Under certain assumptions, there exists an invariant manifold associated with the solution of a "Sylvester-like' partial differential equation (PDE) which relates to the steady-state response of the system [Isidori, Nonlinear Control Systems, 1995]. We introduce a procedure to numerically approximate the solutions to these invariance equations using the Galerkin spectral method, which involves expanding the solution through a basis of globally supported functions and using a Newton iteration on the coefficients. These approximate solutions to the invariance equation can then be used to construct reduced-order models. We demonstrate this method's ability to achieve moment matching in interconnected systems involving implicitly defined nonlinear signal generators, and in nonlinear problems of dimension up to 1000.

Carlos Doebeli
Imperial College London
c.doebeli22@imperial.ac.uk

MS285

Efficient Trajectory Inference in Wasserstein Space Using Consecutive Averaging

Capturing data from dynamic processes through cross-sectional measurements is seen in many fields such as computational biology. Trajectory inference deals with the challenge of reconstructing continuous processes from such observations. In this work, we propose methods for B-spline approximation and interpolation of point clouds through consecutive averaging that is instrinsic to the Wasserstein space. Combining subdivision schemes with optimal transport-based geodesic, our methods carry out trajectory inference at a chosen level of precision and smoothness, and can automatically handle scenarios where particles undergo division over time. We rigorously evaluate our method by providing convergence guarantees and testing it on simulated cell data characterized by bifurcations and merges, comparing its performance against state-of-the-art trajectory inference and interpolation methods. The results not only underscore the effectiveness of our method in inferring trajectories, but also highlight the benefit of performing interpolation and approximation that respect the inherent geometric properties of the data.

Amartya Banerjee
University of North Carolina at Chapel Hill
amartya1@cs.unc.edu

Harlin Lee
University of North Carolina at Chapel Hill
School of Data Science and Society
harlin@unc.edu

Caroline Moosmueller
University of North Carolina at Chapel Hill
cmoosm@unc.edu

Nir Sharon
Tel Aviv University
nsharon@tauex.tau.ac.il

MS285

High-Dimensional Ergodic Control Via Interacting Particle Systems and Generative Modeling Tools

Ergodic stochastic control problems are intrinsically computationally challenging, as sample trajectories typically need to be simulated over a long time horizon. They are even more challenging in high dimensions, when numerous sample trajectories must be simulated. Recent advances have shown that a class of ergodic control problems can be related to quasi-stationary distributions (QSDs); in particular, the log-transform of a QSD is exactly equal to the optimal ergodic control. QSDs describe behavior of a diffusion process in bounded domains over long time intervals, even when a stationary distribution does not exist. Moreover, it is known that a Fleming-Viot interacting particle system can produce samples from QSDs. We use recent advances in score-based generative modeling to approximate solutions to high-dimensional ergodic control problems using samples from the Fleming-Viot system.

Benjamin J. Zhang
Brown University
benjamin_zhang@brown.edu

Paul Dupuis
Division of Applied Mathematics
Brown University
dupuis@dam.brown.edu

MS286

A New Re-Redistribution Scheme for Weighted State Redistribution with Adaptive Mesh Refinement

State redistribution (SRD) is a recently developed technique for stabilizing cut cells that result from finite-volume embedded boundary methods. SRD has been successfully applied to a variety of compressible and incompressible flow problems. When used in conjunction with adaptive mesh refinement (AMR), additional steps are needed to

preserve the accuracy and conservation properties of the solution if the embedded boundary is not restricted to a single level of the mesh hierarchy. In this work, we extend the weighted state redistribution algorithm to cases where cut cells live at or near a coarse-fine interface within the domain. The resulting algorithm maintains conservation and is demonstrated on several two- and three-dimensional example problems.

Isabel Barrio Sanchez
Department of Mathematics
University of Pittsburgh
isb42@pitt.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
Lawrence Berkeley National Laboratory, U.S.
jbbell@lbl.gov

Marc T. Henry de Frahan
National Renewable Energy Laboratory, U.S.
Marc.HenrydeFrahan@nrel.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
WeiqunZhang@lbl.gov

MS286

A Moving Embedded Boundary Approach for the Compressible Navier-Stokes Equations with Adaptive Mesh Refinement

A numerical method has been developed to perform compressible flow simulations involving moving boundaries using an embedded boundary approach within the block-structured adaptive mesh refinement (SAMR) framework of AMReX. We leverage the SAMR capability to obtain quantitatively accurate results whilst using robust, second-order finite volume schemes. A conservative, unsplit, cut-cell approach is utilized and a ghost-cell approach is developed for computing the flux on the moving, embedded boundary faces. A third-order least-squares formulation has been developed to compute the wall velocity gradients, and was found to significantly improve the performance of the solver in terms of the quantitative comparison of surface quantities such as the skin friction coefficient. Various test cases are performed to validate the method, and compared with analytical, experimental, and other numerical results in literature. Inviscid and viscous test cases are performed that span a wide regime of flow speeds acoustic, smooth flows and flows with shocks. A closed system with moving boundaries an oscillating piston in a cylinder, showed that the error in mass within the system decreases with refinement, demonstrating that the numerical scheme is conservative with grid refinement, but is not discretely conservative. Viscous test cases involve that of a horizontally moving cylinder, an inline oscillating cylinder, and a transversely oscillating cylinder.

Mahesh Natarajan
Lawrence Berkeley National Laboratory
MaheshNatarajan@lbl.gov

Ray W. Grout
National Renewable Energy Laboratory
ray.grout@nrel.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
WeiqunZhang@lbl.gov

Marc Day
National Renewable Energy Laboratory
marcus.day@nrel.gov

MS286

Energy Stable Embedded Boundary Approach for High-Order Non-Dissipative Schemes

High-order non-dissipative schemes are preferred for accurate turbulent flow computations because of their superior resolution and discrete conservation properties. However, handling practical geometries with non-dissipative schemes is challenging because of the meshing complexities associated with those geometries and the boundary instabilities that develop over non-trivial meshes. An embedded boundary (EB) simplifies mesh generation by using a Cartesian fluid grid and simply cutting out the solid body from the grid. However, an EB creates non-uniform cells, which often causes small-cell issues and boundary instabilities; the instabilities are particularly severe for high-fidelity calculations that avoid artificial/numerical dissipation. As a result, existing non-dissipative EB methods are mostly restricted to second-order accuracy with ad hoc approaches (e.g. cell merging, flux redistribution, etc.) to address the small-cell problem. This study uses energy stability concepts to derive provably time-stable EB discretizations for incompressible/compressible flows that are globally up to fourth-order accurate (sixth-order accurate in the interior) and address the small-cell issue by design. Moreover, the proposed dimensionally-split framework avoids geometry/solution reconstructions, improving the computational efficiency and simplifying implementation in an existing solver. The proposed method is evaluated by performing a series of inviscid and viscous flow simulations.

Nek Sharan, Uday Shanker
Department of Aerospace Engineering
Auburn University
nsharan@auburn.edu, nzs0089@auburn.edu

Mahesh Natarajan
Lawrence Berkeley National Laboratory
MaheshNatarajan@lbl.gov

MS286

An Overview of the Embedded Boundary Capability in Amrex

AMReX is a software framework for building massively parallel block-structured adaptive mesh refinement (AMR) applications that run on both CPU and GPU systems. It supports complex geometries with a multilevel embedded boundary (EB) approach. AMReX provides an implicit function approach for generating the geometry information. Additionally, it also supports geometry generation using STL files. The EB information is stored in distributed data containers that are accessible on AMR meshes. AMReX provides supports for common operations in EB algorithms, including stability fixes for small cut cell, refluxing in solvers for hyperbolic conservation laws, and sparse linear system solvers for implicit methods. In this talk, we

will give an overview of the EB capabilities in AMReX and showcase them with results from a number of AMReX based applications.

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
WeiqunZhang@lbl.gov

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
Lawrence Berkeley National Laboratory, U.S.
jbbell@lbl.gov

MS287

Preconditioning the Steady State Oseen Equations in An Rbf-Fd Setting

Radial basis function finite difference (RBF-FD) discretization has recently emerged as an alternative to classical finite difference or finite element discretization of (systems) of partial differential equations. In this talk, we briefly describe how to discretize the steady state Oseen equations in an RBF-FD setting. Particularly, we show how to deal with the pressure constraint. Afterwards we focus on the construction of preconditioners for the iterative solution of the resulting linear systems of equations. In RBF-FD, a higher discretization accuracy may be obtained by increasing the stencil size. This, however, leads to a less sparse and often also worse conditioned stiffness matrix which are both challenges for subsequent iterative solvers, which creates the need for good preconditioners. In our numerical results, we focus on RBF-FD discretizations based on polyharmonic splines (PHS) with polynomial augmentation. We illustrate the convergence of the method and the performance of the preconditioners in numerical experiments for the three-dimensional Oseen equations.

Michael Koch
Hamburg University of Technology
michael.koch@tuhh.de

Sabine Le Borne
Hamburg University of Technology
Department of Mathematics
leborne@tuhh.de

MS287

Arbitrary Domain Generation for Steady Convection Via Radial Basis Functions

A radial basis function scheme is developed for the computation of steady convection over arbitrary two-dimensional domains. The Boussinesq flow under consideration is internally forced with a Gaussian spatial distribution to represent heat deposition from a propagating laser. Scaling results are obtained via numerical solutions for the steady temperature and flow velocity with respect to shape parameters describing the domain geometry and the spatial distribution of internal forcing. Applications of the numerical scheme are then discussed as they pertain to the simulation of physical geometric configurations present in experiments involving laser propagation.

Jeremiah Lane
Air Force Institute of Technology
jeremiah.lane@afit.edu

MS287

Operator Learning with Kernels: Property-preservation, Uncertainty Quantification, and Spatial Adaptivity

Operator learning, the task of recovering maps between infinite-dimensional function spaces, has seen many recent successes in both scientific and engineering applications. Neural operators in particular leverage neural network architectures for the task of operator learning. In this talk, we discuss how kernels can be used to enhance both the approximation abilities and the interpretability of neural operators. We further show how kernel-based approximation can be used in conjunction with two major neural operators to obtain uncertainty quantification. We demonstrate the application of our methods to the recovery of solution operators to partial differential equations (PDEs).

Varun Shankar, Ramansh Sharma
Kahlert School of Computing
University of Utah
shankar@cs.utah.edu, ramansh@cs.utah.edu

MS287

A New Framework for Numerical Integration of Scattered Data Over Geometrically Complex Domains

The dominant approach to generating quadrature (or cubature) formulas is to pick a "nice" vector space of functions for which the formulas are exact, such as algebraic or trigonometric polynomials. For numerical integration over intervals, this approach gives rise to Newton-Cotes and Gaussian quadrature rules. However, for geometrically complex domains in higher dimensions, this exactness approach can be challenging, if not impossible since it requires being able to exactly integrate basis functions for the vector space over the domains (or some collection of subdomains). Another challenge with determining good quadrature formulas arises when the integrand is not given everywhere over the domain, but only as samples at predefined, possibly "scattered' points (i.e., a point cloud), which is common in applications involving experimental measurements or when quadrature is a secondary operation to some larger computational effort. In this talk we introduce a new framework for generating quadrature formulas that bypasses these challenges. The framework only relies on numerical approximations of certain Laplace operators and on linear algebra. We show how several classic univariate quadrature formulas can arise from this framework and demonstrate its applicability to generating accurate quadrature formulas for geometrically complex domains (including surfaces) discretized with point clouds.

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS288

Real Time Simulations of Nuclear Cloud Rise

Operational models of nuclear cloud rise are required to have short simulation times to facilitate timely response by decision makers and emergency personnel. Because of

this, current models are low dimensional approximations and rely on parameterized variables tuned to historical U.S. nuclear test results. This study examines the use of the Energy Research and Forecasting model (ERF) as an alternative. By simulating the fully three dimensional and time varying atmosphere dynamics, significantly fewer modeling assumptions are required. Four historical U.S. nuclear tests with varying heights of burst were simulated in ERF and simulations agree favorably with both historical test data and previous simulation data from the Weather Research and Forecasting model (WRF). Furthermore, it was demonstrated that cloud rise simulations could be performed in real time, such that the measured wall time is approximately equal to the amount of time simulated. This represents a significant improvement in capability for nuclear cloud rise simulations and may enable its use as an operational tool in emergency situations. Future work will focus on including additional physics to the code such as debris microphysics and urban terrain modeling. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-868852.

Brandon C. Blakeley
LLNL
blakeley1@llnl.gov

Katie Lundquist
Lawrence Livermore National Laboratory
lundquist3@llnl.gov

MS288

Super Droplet Method for Lagrangian Microphysics Simulations in ERF

Modelling cloud microphysics is essential for simulating atmospheric flows. Eulerian approaches use empirical models for cloud/rain-droplets dynamics to evolve moisture quantities. Spectral approaches are accurate but expensive since they evolve the droplet density distribution. The particle-based super-droplet method (SDM) simulates droplet dynamics without empirical assumptions. Super-droplets are characterized by their multiplicity (the number of physical droplets they represent) and physical attributes (size, position, and velocity). Physical processes (advection, sedimentation, condensation/evaporation, and coalescence) determine the evolution of the super-droplets. The accuracy of derived quantities (cloud/rain-water content and rain accumulation) depend on the number of super-droplets per grid cell. The computational cost, while higher than bulk parameterization models, is acceptable on modern HPC platforms. We discuss the implementation of an SDM-based microphysics model in ERF, an exascale code for atmospheric flow simulations. We verify our implementation by simulating cloud dynamics and precipitation. We present scalability studies on modern HPC platforms, specifically with respect to the number of super-droplets. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344 and funded by the LDRD Program at LLNL under project tracking code 24-SI-001.

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA
ghosh5@llnl.gov

Hassan Beydoun, Dana McGuffin, David J. Gardner, Katie Lundquist
Lawrence Livermore National Laboratory
beydoun1@llnl.gov, mcguffin1@llnl.gov, gardner48@llnl.gov, lundquist3@llnl.gov

MS288

An Introduction to ERF: Algorithms and Applications

The Energy Research and Forecasting (ERF) modeling code is being developed as part of a core capability within the U.S. Department of Energy Wind Energy Technologies Office (WETO) portfolio. ERF will target high-fidelity representations of atmospheric flows at the mesoscale, multi-scale coupling, and characterization of energy available for harvesting via wind turbines. ERF is built upon AMReX, a block-structured adaptive mesh refinement (AMR) software framework which provides the underlying infrastructure for block structured mesh operations and hierarchical parallelism i.e., an MPI+X model, where X may be OpenMP on multicore CPU-only systems, or CUDA, HIP, or SYCL on GPU-accelerated systems. In this talk, we will discuss details regarding ERFs numerics, applications for the software, benchmark results obtained with the code, and future outlook.

Aaron M. Lattanzi, Ann S. Almgren, Mahesh Natarajan
Lawrence Berkeley National Laboratory
amlattanzi@lbl.gov, asalmgren@lbl.gov, MaheshNatarajan@lbl.gov

Eliot Quon
NREL
eliot.quon@nrel.gov

MS288

Turbulent Inflow Generation with the Energy Research and Forecasting (ERF) Code

Atmospheric flows exhibit a broad range of spatial and temporal scales. Consequently, numerical simulations of weather systems often employ adaptive mesh refinement (AMR) techniques to locally refine in regions of interest. Coarse-fine interfaces pose a fundamental challenge to computational frameworks since the resolved and modeled turbulence may be dramatically different in each grid e.g., a transition from a coarse Reynolds averaged (RANS) to a fine large eddy (LES) description. As a first step towards modeling such phenomena in ERF, we implement a turbulent inflow generation technique at domain boundaries. The selected method triggers grid-level appropriate turbulence at inflow boundaries but is currently restricted to a single resolution. Future work with the method will involve general application to coarse-fine interfaces and the partitioning of transported turbulent kinetic energy (TKE) by the RANS model into resolved turbulence by the LES model.

Ting-Hsuan Ma
University of Pittsburgh
dustinma324@gmail.com

Aaron M. Lattanzi
Lawrence Berkeley National Laboratory
amlattanzi@lbl.gov

Jeff Mirocha
Lawrence Livermore National Laboratory

jmirocha@llnl.gov

MS289

Density-Equalizing Map with Applications

We present surface and volumetric mapping methods based on a natural principle of density diffusion. Specifically, we start with a prescribed density distribution in a surface or volumetric domain and then create shape deformations with different regions enlarged or shrunk based on the density gradient. Using the proposed methods, we can easily achieve different mapping effects with controllable area change. Applications to shape registration, morphing, remeshing, medical shape analysis, and data visualization will be presented.

Gary Choi
The Chinese University of Hong Kong
ptchoi@cuhk.edu.hk

MS289

Discontinuous Galerkin Methods for Embedded Geometries Described Via B-Rep

In this talk, a class of novel high-order accurate discontinuous Galerkin (DG) methods for solving systems of partial differential equations over geometries with industrially relevant complexity is presented. Geometries are described using the oriented boundary representation (B-Rep) of the volume, an industry standard for data transmission. DG methods are based on the use of discontinuous basis functions and suitably defined boundary terms to weakly enforce boundary and interface conditions; with respect to other techniques, such as finite element or finite volume methods, DG methods feature high-order accuracy with various types of meshes (e.g., simplicial, polyhedral, embedded-boundary), local conservation properties, local *hp* adaptivity, block-structured mass matrices, and high parallel efficiency. Numerical applications involve the solution of different sets of PDEs including the Laplace, elasticity, and Stokes problems and demonstrate the capabilities of the proposed methodology. The implementation of multigrid methods for large-scale applications is also discussed.

Vincenzo Gulizzi
LBNL
vincenzo.gulizzi@unipa.it

MS289

Combined Parameter and Shape Optimization of Electric Machines with Isogeometric Analysis Using Second Derivatives

In design optimization problem, both parameter and shape optimization play critical roles in improving model performance. However, traditional optimization approaches often treat these aspects separately. This work presents a framework that integrates parameter and free form optimization, where the domain and the Partial Differential Equation modeling the physical system are discretized using Isogeometric Analysis. Using the adjoint method, we compute the gradients of an objective function with respect to both, the parameters and the control points that define the geometry. Furthermore, we incorporate second derivatives with respect to the control point and the parameters into the framework to reduce the computational cost of the optimization process. The proposed method is demonstrated by its application to a nonlinear 2D magnetostatic simulations of a state-of-the-art Permanent Magnet Synchronous Motor reconstructed by the commercial software JMAG in the open-source framework GeoPde, but it can be applied to any structural optimization problem with an elliptic constraint. The results are compared with designs optimized by parameter or shape optimization alone. Our results show that the simultaneous adjustment of rotor parameters and shape can significantly reduce the permanent magnet mass and nearly eliminate the torque ripple and by using second derivatives we can reduce the computational cost significantly.

Theodor Komann
Technical University of Darmstadt
komann@mathematik.tu-darmstadt.de

Stefan Ulbrich
Technische Universitaet Darmstadt
Fachbereich Mathematik
ulbrich@mathematik.tu-darmstadt.de

MS289

Controlling Shape in Leaves and Batteries

Living systems grow robustly into a remarkable array of forms, often in the face of external perturbations. Elucidating how this happens is a key challenge in biology, with massive potential applications in living and bio-inspired systems, and will require weaving together tools from mathematics, physics, and biology. Here using bacterial cell envelopes and plant tree leaves as examples of growing structures that can be abstracted as thin elastic shells, I will describe how models when informed by experiments can be essential in discovering the regulatory mechanisms involved in controlling shape. Lastly, I will discuss ongoing work connecting data and model-driven approaches and how understanding shape control in living organisms may help us design better solid-state batteries.

Salem H. Mosleh
University of Maryland Eastern Shore
salmosleh@umes.edu

MS289

A Framework for Inverse Problems: Manifold Harmonics for Parameterizing the Differential Growth of Plates and Shells

Shape-shifting sheets are attracting growing interest across diverse fields, including soft robotics, 4D printing, and kinetic architecture. When external stimuli, such as temperature or light, act on an initially flat thin sheet they induce local residual stresses. This causes the sheet to undergo in-plane stretching and out-of-plane bending as it seeks a nearly stress-free configuration that can be embedded in 3D space. By controlling local differential growth, flat sheets can be programmed to morph into a given target shape. Finding the exact growth field that morphs an initial configuration into a given target one is generally a hard task and numerically prohibitive for large problems. However, parameterizing the growth field on a global basis can significantly reduce the size of the inverse problem. In this study, we use the manifold harmonics (MH) to parameterize the growth field onto the initial configuration of a given target shape. We first compute the manifold harmonics using the discrete Laplace-Beltrami operator, then express the growth field as a combination of linear weights associated with each basis. By ranking these harmonics

based on their eigenvalues, we can achieve a hierarchical spatiotemporal progression from the initial configuration to the target shape. We also investigate various mapping techniques to flatten target shapes and their corresponding harmonics.

Mahmoud S. Shaqfa
Massachusetts Institute of Technology
mshaqfa@mit.edu

MS290

Robust Adaptive Meshing for PDEs with Nonsmooth Solutions and Distribution Transformer

In computational science and engineering, adaptive mesh generation is necessary to resolve the singularity in the PDE solution in order to achieve desired accuracy with as little computational overhead as possible. An indispensable component is the a posteriori error estimation. For challenging problems, the robustness of a posteriori error estimates with respect to high contrast PDE parameters is critical for detecting singular layers and/or avoiding over-refinement. We will present robust a posteriori error estimators for PDEs with high-contrast coefficients, including elliptic interface problems and convection-reaction-diffusion equations. The robustness is demonstrated via extensive numerical examples. In the second part of the talk, we present a new deep neural network-based distribution transformer for generating structured distributions with applications to machine learning.

Difeng Cai
Southern Methodist University
ddcai@smu.edu

MS290

Adaptive Finite Difference Solution for Fractional Boundary Value Problems

It is known that the solution to the Dirichlet boundary value problem of the fractional Laplacian operator on a bounded domain exhibits large gradient near the domain boundary. Finite difference and finite element methods show slow convergence as the mesh is refined if a quasi-uniform mesh is used. In this talk we will present a study on how an adaptive moving mesh method can be used for the underlying problem. A newly developed grid-overlay finite difference method is used to discretize the fractional Laplacian. Numerical results will be presented to verify the second-order convergence of the method on adaptive meshes.

Weizhang Huang
University of Kansas
whuang@ku.edu

MS290

Recent Advances in the Target-Matrix Optimization Paradigm for High-Order Mesh Optimization

The Target-Matrix Optimization Paradigm (TMOP) is a technique for adapting high-order meshes. In this method, optimized mesh nodal positions are determined by minimizing a functional that depends on each element's current and target geometric parameters: size, aspect-ratio, skew, and orientation. In previous work, we have demonstrated use of TMOP for solution-driven adaptivity in two- and three-dimensional problems. Recently we have augmented the TMOP-based functional with a penalization term that depends on a subset of mesh nodes aligning with a target surface prescribed as the zero of a level-set function. Minimizing this functional for a given mesh and a level-set thus results in a body-fitted mesh with good element quality. This technique has proven to be effective for automating meshing in topology optimization featuring complex curvilinear geometries. We augment the mesh-fitting approach with *hp*-adaptivity to produce mixed-order meshes that are more computationally efficient than uniform meshes. We also incorporate Automatic Differentiation for computation of gradients of the mesh quality metrics required for the minimization of the TMOP function using the Newton's method. The proposed high-order mesh optimization approach is purely algebraic, and extends to different element types (quadrilaterals/triangles/tetrahedron/hexahedra) in two- and three-dimensions.

Ketan Mittal, Veselin Dobrev, Tzanio Kolev
Lawrence Livermore National Laboratory
mittal3@llnl.gov, dobrev1@llnl.gov, kolev1@llnl.gov

Patrick Knupp
Dihedral LLC
knupp.patrick@gmail.com

Mathias Schmidt, Vladimir Z. Tomov
Lawrence Livermore National Laboratory
schmidt43@llnl.gov, tomov2@llnl.gov

MS290

Adaptive Mesh Strategies for Wave Propagation Problems

Adaptive mesh refinement (AMR) is crucial for achieving efficient numerical simulations of large time-dependent problems. However, incorporating explicit-in-time methods into a spacetime AMR framework presents significant challenges, particularly due to the stringent stability constraints imposed by time-stepping algorithms. In this presentation, we introduce a novel spacetime adaptive mesh refinement technique that addresses these challenges through advanced refinement strategies. Numerical results for the linear transport equation and acoustic wave equation will be presented.

Paulina Sepúlveda
Pontificia Universidad Católica de Valparaíso, Chile
paulina.sepulveda@pucv.cl

MS291

Learning Reduced Order Models with a Parametrized Rational Structure Using Low-Rank Barycentric Forms

Rational approximation is a powerful tool for capturing the behavior of complex physical systems. The AAA algorithm [Nakatsukasa et al., The AAA Algorithm for Rational Approximation, 2018] computes such approximations in a data-driven setting by combining interpolation and least-squares optimization in an iterative framework. Many systems of practical interest have an underlying parametrized structure which should be preserved by rational approximation algorithms. The p-AAA algorithm [Carracedo Rodriguez et al., The p-AAA Algorithm for Data-Driven Modeling of Parametric Dynamical Systems, 2023] does this by extending the AAA algorithm to the multivariate rational approximation setting. In this framework rational approximants are represented in

terms of a multivariate barycentric form $H(s,\ldots,p) = \left(\sum_{i=1}^{k}\cdots\sum_{j=1}^{q}\frac{\beta_{i\ldots j}}{(s-\sigma_i)\cdots(p-\pi_j)}\right)\Big/\left(\sum_{i=1}^{k}\cdots\sum_{j=1}^{q}\frac{\alpha_{i\ldots j}}{(s-\sigma_i)\cdots(p-\pi_j)}\right)$. Especially, when moving to many variables $s,\ldots,p$, evaluating H and determining the barycentric coefficient tensor α becomes computationally demanding. Motivated by this fact, we introduce new barycentric forms which are based on low-rank decompositions of the tensor α. We discuss the computational benefits of our derived representation and demonstrate its effectiveness in practice.

Linus Balicki
Virginia Tech
Blacksburg, USA
balicki@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS291

Colora: Continuous Low-Rank Adaptation for Reduced Implicit Neural Modeling of Parameterized Partial Differential Equations

This work introduces reduced models based on Continuous Low Rank Adaptation (CoLoRA) that pre-train neural networks for a given partial differential equation and then continuously adapt low-rank weights in time to rapidly predict the evolution of solution fields at new physics parameters and new initial conditions. The adaptation can be either purely data-driven or via an equation-driven variational approach that provides Galerkin-optimal approximations. Because CoLoRA approximates solution fields locally in time, the rank of the weights can be kept small, which means that only few training trajectories are required offline so that CoLoRA is well suited for data-scarce regimes. Predictions with CoLoRA are orders of magnitude faster than with classical methods and their accuracy and parameter efficiency is higher compared to other neural network approaches.

Jules Berman
NYU
jmb1174@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS291

Structure-Preserving Model Reduction Method for Lure Network Systems

As the dimension and complexity of network systems continue to increase, their analysis and control become more and more challenging. To address this, different model reduction methods have been developed to simplify the large-scale systems. Conventional methods such as balanced truncation, moment matching and Hankel-norm approximation often overlook the underlying network structure of a system. In constrast, clustering-based model reduction offers a structure-preserving solution for network systems. However, the existing research has primarily focused on reducing linear network systems. In this work, we extend the clustering-based model reduction framework to simplify thenetwork of a class of nonlinear systems, namely Lur'e network systems. The presented method not only preserves the network structure in the reduced systems but also has the potential to retain some important properties including stability, controllability and observability. Furthermore, an explicit bound on the approximation error will also be discussed.

Yangming Dou
University of Groningen
y.dou@rug.nl

Xiaodong Cheng
Wageningen University & Research
xiaodong.cheng@wur.nl

Jacquelien M. Scherpen
Rijksuniversiteit Groningen
j.m.a.scherpen@rug.nl

MS291

Model Reduction of Network Systems and Its Applications in System Simulation Software

Structure-preserving approximation is still an active research area. By preserving or mimicking relevant geometric structures such as, e.g., conservation laws or symplecticities, unphysical solution behavior and numerical instabilities can be avoided in many cases. Our model reduction approach is analyzed using energy-based modeling concepts, such as the port-Hamiltonian formalism. A particular focus of the talk also lies on the realization of the snapshot-based model order- and complexity-reduction. While beneficial for the robustness and performance of the reduced models, the compatibility conditions pose a challenge in the training phase. Appropriate adaptions of the conventional model reduction methods will be presented.

Bjoern Liljegren-Sailer
Trier University
bjoern.sailer@uni-trier.de

MS292

Robust Discretization for Incompressible Fluids with Variable Density and Applications to Turbulent Thermal Convection

We present two time discretization of the incompressible Navier-Stokes equations with variable density and dynamical viscosity. The first method uses the couple momentum, equal to the density times the velocity, and pressure as primary unknowns. The diffusion term is rewritten appropriately so the resulting stiffness matrix is time-independent. The method yields a robust and conditionally stable scheme that is suitable for high order finite element and spectral methods. Then, we introduce a method that couples the momentum, velocity and pressure as primary unknown where the relation between the momentum and velocity is enforced weakly (i.e. it now becomes a constraints). The resulting scheme yields a better stability and energy preserving properties than the first one. Theoretical results on the convergence properties of both scheme will also be presented. Both methods are first validated numerically using manufactured solutions before being applied to two-layered fluids problems with turbulent thermal convection.

Loic Cappanera, Salvatore Giordano, An Vu
University of Houston
lmcappan@central.uh.edu, sgiordan@central.uh.edu,

atvu26@central.uh.edu

MS292

Finite Element Numerical Schemes for the Navier-Stokes-Cahn-Hilliard System with Degenerate Mobility

Many scientific, engineering, and industrial applications related with hydrodynamics and materials science applications are based on the understanding of the evolution of the interface between two or more immiscible fluids. A very effective approach for representing interface problems is the diffuse interface/phase field approach, which describes the interfaces by layers of small thickness and whose structure is determined by a balance of molecular forces, in such a way that the tendencies for mixing and de-mixing compete through a non-local mixing energy. In particular, the Navier-Stokes-Cahn-Hilliard system was introduced to model the behavior of a mixture of two incompressible, viscous Newtonian fluids with the same constant density. During this talk I will present two new numerical schemes to approximate the Navier-Stokes-Cahn-Hilliard system with degenerate mobility using finite differences in time and finite elements in space. The proposed schemes are conservative, energy-stable and preserve the maximum principle approximately (the amount of the phase variable being outside of the interval $[0, 1]$ goes to zero in terms of a truncation parameter). Additionally, I will present several numerical results to illustrate the accuracy and the well behavior of the proposed schemes, as well as a comparison with the behavior of the Navier-Stokes-Cahn-Hilliard model with constant mobility.

Giordano Tierra
University of North Texas
gtierra@unt.edu

MS292

Structure Preserving Methods for Ideal Mhd

In this talk, we present our recent developments in the numerical approximation of the MHD system. The equations are approximated using the continuous finite element method in space, with high-order time integration methods, such as Runge-Kutta and Crank-Nicholson schemes, applied for time evolution. Due to the convective nature of the MHD system, finite element approximations can be unstable, so we introduce viscous regularization as a numerical stabilization term. We also investigate various thermodynamic properties of the system, including invariant domains: positivity of density and internal energy, and entropy production of the specific entropy. At the PDE level, we prove that these properties hold. At the discrete level, we apply the Marchuk-Strang splitting technique to separate the MHD equations into convection and source components and prove that the scheme preserves invariant domains. Additionally, the divergence-free finite element space is used to approximate the magnetic field, ensuring the exact preservation of the involution constraints.

Murtazo Nazarov
Uppsala University
Sweden
murtazo.nazarov@it.uu.se

Ignacio Tomas
Department of Mathematics & Statistics, Texas Tech Universit
igtomas@ttu.edu

Tuan Anh Dao
Royal Institute of Technology
Sweden
tbd

MS292

A Conforming Interface Approach for Rarefied Gas and Solid Phase Transitions

In moving boundary problems, the physical field equations are coupled with the changing domain in a non-linear way. A prominent example of this is the Stefan problem, the classical phase-change problem. We aim to develop a method that describes the transition between a rarefied gas and a solid. The underlying equations are complex and produce strong boundary layer effects [Weniger et al., Int. J. Heat Mass Transf., 2023]. The most flexible methods for general phase transition problems are based on a fixed computational mesh that is not interface-conforming, e.g., the Level-Set method, Phase-field method or Volume-of-Fluids method [Elgeti and Sauerland, Arch. Comput. Methods Eng., 2015]. Interface conditions are then imposed by interpolation. In the context of rarefied gas phase transitions, the interface and its discontinuity conditions need to be enforced as accurately as possible, as the produced boundary layer effects shape the bulk solution. We therefore present a conforming interface method based on standard Finite Elements and the Level-Set method with exact remeshing by reconstructing the interface. This way, the interface conditions can be imposed directly at the interface, perfect mesh refinement is possible, and the flexibility of arbitrary topology changes of the Level-Set method is inherited. The method is designed to allow the Finite Element solver and the meshing tool to be interchangeable and adaptable, ensuring maximum compatibility and versatility.

Donat Weniger
RWTH Aachen University
weniger@aices.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen
mt@acom.rwth-aachen.de

MS293

Efficient Parallel Multigrid Methods with the DEAL.II Library

This talk presents the implementation of efficient multigrid methods for high-order finite element problems. We concentrate on hybrid h- and p-multigrid ingredients discussed in [N. Fehn et al., Hybrid multigrid methods for high-order discontinuous Galerkin discretizations, JCP 2020] and [P. Munch et. al., Efficient distributed matrix-free multigrid methods on locally refined meshes for FEM computations, ACM TOPC 2023] for the level transfer and matrix-free evaluation for the level operators. We will present comprehensive performance comparison of Chebyshev-accelerated smoothers of point-Jacobi, block-Jacobi and additive Schwarz type in terms of the compute intensity, memory access and communication in a massively parallel setting for both SIMD-accelerated multi-core CPUs and GPU systems.

Shubham K. Goswami
Department of Computational and Data Sciences
Indian Institute of Science, Bengaluru
shubham.goswami@rub.de

Martin Kronbichler
Ruhr University Bochum
martin.kronbichler@rub.de

MS293

Communication-Aware Multigrid-Based Solvers with Matrix-Free Finite Element Operators

We consider large-scale implicit solvers for the numerical solution of partial differential equations (PDEs). These solvers require the high-bandwith networks of an HPC system for a fast time to solution. However, increasing variability in performance of the HPC systems, most likely caused by variable communication latencies and network congestion, makes the execution time of solver algorithms unpredictable and hard to measure. In particular, latency variability makes the reliable comparison of different algorithms and implementations difficult or impossible on HPC. The purpose of this talk is twofold: first, we present statistical techniques for separating HPC performance data for individual iterations of our multigrid-preconditioned Krylov solver into different levels of system latency, allowing us to, for example, identify and remove time periods when extremely high system latencies throttle application performance and distort performance measurements. Second, we present new communication-hiding methods for the matrix-free matrix-vector products occurring during multigrid smoothing iterations and outer Krylov solver orthogonalization. We implement our communication hiding methods for the matvec routine in the large-scale mantle convection code Rhea, and provide performance analysis of the Stokes solver on the Anvil and Frontera supercomputers using our new analysis techniques.

Max P. Heldman
Boston University
maxh@vt.edu

Johann Rudi
Virginia Tech
jrudi@vt.edu

Julie Bessac
US National Renewable Energy Laboratory
julie.bessac@nrel.gov

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Jiaqi Fang, Michael Gurnis
Seismological Laboratory
California Institute of Technology
jfang@caltech.edu, gurnis@gps.caltech.edu

MS293

Cache-Optimized Implementation of Multigrid Methods

In order to generate the high computational intensity needed on modern hardware, we consider block smoothers based on domain decomposition. These are known to yield fast converging methods and can be implemented in several ways, additive and multiplicative or as a preconditioner and combined with the residual computation. Using a simple performance model, we will discuss how these choices affect the performance of such methods on hardware with memory bottlenecks. We then discuss a cache-aware CPU implementation and present results. On GPU, the performance is also limited by the bottleneck to main memory, but access to cached data in on-device memory can pose a second one. We show ways to alleviate this slow-down and obtain a second roofline model. The presentation closes with performance results for Poisson and Stokes problems exhibiting the feasibility of our approach.

Guido Kanschat
Universität Heidelberg, Germany
kanschat@uni-heidelberg.de

Cu Cui
University of Heidelberg
cu.cui@iwr.uni-heidelberg.de

Michal Wichrowski
Heidelberg University
wichrowski@uni-heidelberg.de

MS293

Matrix-Free Multigrid Block-Preconditioners for Higher Order Discontinuous Galerkin Discretizations

Efficient and suitably preconditioned iterative solvers for elliptic partial differential equations (PDEs) of the convection-diffusion type are used in all fields of science and engineering. To achieve optimal performance, solvers have to exploit the high arithmetic intensity of modern many-core CPUs. The computationally most expensive components of the solver are the repeated application of the linear operator and the preconditioner solve. For discretisations based on higher-order Discontinuous Galerkin methods, sum-factorisation techniques result in a dramatic reduction of the computational complexity for the matrix-free operator application. However, an algorithmically optimal *hp*-multigrid preconditioner also requires the repeated inversion of dense block-matrices in the DG smoother. The explicit assembly and direct solution of those block-matrices with a bandwidth-bound LU- or Cholesky-factorisation counteracts any gains from the efficient operator application. Here we present an alternative, fully matrix-free implementation of DG block-smoothers. By inverting the block-matrices iteratively, it is possible to harness the full computational power of the CPU. We implemented a hybrid multigrid algorithm for high order DG discretisations in the EXADUNE framework. The effectiveness of this approach is demonstrated by solving a set of representative elliptic PDEs of increasing complexity.

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

Eike H. Mueller
University of Bath
e.mueller@bath.ac.uk

Steffen Muething
Universitaet Heidelberg
steffen.muething@iwr.uni-heidelberg.de

Marian Piatkowski
IWR, Heidelberg University

marian.piatkowski@uni-heidelberg.de

MS294

Interface Design and Performance Evaluation of a Batched Sparse Direct Solver

Over the course of interactions with various application teams, the need for batched sparse linear algebra functions has emerged in order to make more efficient use of the GPUs for many small and sparse linear algebra problems. In this talk, we present our recent work on a batched sparse direct solver for GPUs. The sparse LU factorization is computed by the levels of the elimination tree, leveraging the batched dense operations at each level and a new batched Scatter GPU kernel. The sparse triangular solve is computed by the level sets of the directed acyclic graph (DAG) of the triangular matrix. Batched operations overcome the large overhead associated with launching many small kernels. For medium sized matrix batches with not-so-small bandwidth, using an NVIDIA A100 GPU, our new batched sparse direct solver is orders of magnitude faster than a batched banded solver and uses less than one-tenth of the memory.

Wajih Boukaram
King Abdullah University of Science and Technology
wajih.boukaram@gmail.com

MS294

Graph Neural Preconditioners for Iterative Solutions of Sparse Linear Systems

Graph coarsening is a technique for solving large-scale graph problems by working on a smaller version of the original graph, and possibly interpolating the results back to the original graph. Popularized by algebraic multigrid methods applied to solving linear systems of equations, graph coarsening finds a new chapter in machine learning, particularly graph-based learning models. However, it is challenging to naively apply existing coarsening methods, because it is unclear how the multigrid intuition matches the machine learning problem at hand. We develop an objective-driven approach by explicitly defining the coarsening objective, which admits a geometric interpretation-maintaining the pairwise distance of graphs. We derive the objective function by bounding the change of the distance and show its relationship with weighted kernel k-means clustering, which subsequently defines the coarsening method. We demonstrate its effective use in graph regression and classification tasks.

Jie Chen
MIT-IBM Watson AI Lab, IBM Research
chenjie@us.ibm.com

MS294

Matrix Completion Algorithm in Butterfly Format

Low-rank completion algorithms have been widely studied and applied in both matrix and tensor settings. These algorithms typically assume that only a fixed subset of entries of the matrix/tensor (permitting structured decompositions) have been provided and leverage iterative methods to construct an approximate low-rank representation. For matrices representing discretized oscillatory operators, a single low-rank representation can yield very high numerical ranks and require an impractical number of known entries for the completion algorithm. In these cases, butterfly decomposition can be leveraged to attain more efficient representations. In this work, we propose a matrix completion algorithm to reconstruct a butterfly decomposition, requiring only $O(n\log^a n)$ number of known entries. The proposed algorithm first applies low-rank completion to construct a coarse initial guess, then uses alternating least square algorithms to iteratively reconstruct the butterfly decomposition. Numerical results demonstrating the rapid convergence and computational efficiency will be presented.

Yang Liu
Lawrence Berkeley National Laboratory
liuyangzhuan@lbl.gov

Navjot Singh
University of Illinois Urbana-Champaign
navjot2@illinois.edu

Edgar Solomonik
University of Illinois at Urbana-Champaign
solomon2@illinois.edu

Xiaoye S. Li
Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS294

Fast Direct Solvers for Neural Network Least Squares Approximations

Neural network is a useful tool for function approximations. However, fast solvers for relevant dense linear systems are rarely studied. This work gives a comprehensive characterization of the ill conditioning of some dense linear systems arising from shallow neural network function approximations in one dimensions, and further show some intrinsic structures that make it feasible to design fast direct solvers.

Jianlin Xia
Purdue University
xiaj@purdue.edu

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

Tong Ding, Min Liu, Xinyu Liu
Purdue University
ding158@purdue.edu, liu66@purdue.edu,
liu1957@purdue.edu

MS295

Mgard: Compression and Refactoring of Scientific Data with Variable-Order Polynomial Predictors

MGARD (MultiGrid Adaptive Reduction of Data) is an algorithm for compressing and refactoring scientific data based on the theory of multigrid methods. The core algorithm is built around stable multilevel decompositions of conforming piecewise linear finite element spaces, enabling accurate error control in various norms and derived quantities of interest. We extend this construction to arbitrary order Lagrange finite elements and propose a reformulation of the algorithm as a lifting scheme with polynomial pre-

dictors of arbitrary order. Additionally, a new formulation using a compactly supported wavelet basis is discussed, and an explicit construction of the proposed wavelet transform for uniform dyadic grids is described.

Viktor Reshniak, Qian Gong
Oak Ridge National Laboratory
reshniakv@ornl.gov, gongq@ornl.gov

Scott Klasky
Oak Ridge National Laboratory
Oak Ridge, TN
klasky@ornl.gov

Rick Archibald
Data Analytics and Machine Learning Group
Oak Ridge National Labortory
archibaldrk@ornl.gov

MS295

Overview of Fastmath

The Scientific Discovery Through Advanced Computing (SciDAC) Program is a computational science program at DOE that aims to accelerate progress in scientific research using high performance computing. To achieve this goal, DOE funds multidisciplinary teams, each of which involves close collaboration among domain scientists, applied mathematicians, and computer scientists. The applied mathematicians and computer scientists in the multidisciplinary teams, or partnerships, largely come from two SciDAC Institutes, FASTMath and RAPIDS, which are responsible for providing expertise in applied mathematics, computer science, and machine learning, and developing state-of-the-art, scalable algorithms and tools for solving scientific problems. This talk will provide an overview of the FASTMath Institute and describe the organization of the institute into topical areas and the capabilities provided.

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

MS295

Advances in Fastmath Unstructured Mesh Components in Support of Doe Fusion Simulation Codes.

Accurate simulations of fusion energy system physics require consideration of the complex geometries of the system components. This presentation will present FASTMath efforts on (i) analysis geometry construction for general 3D configurations, (ii) fully automatic generation of well controlled adaptive meshes to ensure simulation fidelity and (iii) tools to support the coupling of individual physics analysis codes across multiple scales. The geometry construction tools support combining CAD and physics geometry to create the desired analysis geometry. Automatic mesh generators support creation of graded anisotropic meshes on arbitrary geometric domains, and specialized mesh generation tools support the specific needs of selected tokamak and stellarator simulation codes. Massively parallel unstructured mesh based particle simulations are supported by a distributed mesh infrastructure that scales for both the particles and mesh. Code coupling, from one-way mesh field coupling to in-memory two-way coupling, of large-scale parallel simulation codes is supported by a set of data coordination structures and mesh field transfer operators. The integration and use of these tools into fusion energy system simulation codes is demonstrated for several problems of interest.

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

Jacob S. Merson, Onkar Sahni
Rensselaer Polytechnic Institute
mersoj2@rpi.edu, sahni@rpi.edu

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

Cameron W. Smith, Dhyanjyoti Nath
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu, nathd@rpi.edu

Usman Riaz, Fuad Hasan
Rensselaer Polytechnic Institute
riazu2@rpi.edu, hasanm4@rpi.edu

Saurabh Tendulkar
Simmetrix Inc.
saurabh@simmetrix.com

MS295

Optimization of Fusion Simulations with Some Available Derivatives

In the taxonomy of optimization problems, derivative-based methods can typically be used only when the full gradient of the objective function is available. However, for some objectives, the full gradient may not be available, but some components are. In this talk, we describe an extension to derivative-free optimization algorithms that takes advantage of the limited available derivative information in this scenario. We present a trust-region algorithm with objective function models based on Hermite interpolation. The method reduces to the classical derivative-free method with models based on Lagrange interpolation when no derivatives are available. We analyze the convergence properties of this algorithm, and demonstrate numerical performance both on synthetic problems and on applications within the magnetic confinement fusion community.

Evan Toler
Argonne National Laboratory
etoler@anl.gov

MS296

Differentiable Simulations for Mechanical Property Measurement: From Articulated Robots to Non-Newtonian Fluids

Traditional methods for measuring mechanical properties often require specialized equipment that is costly, complex to set up, and limited to assessing a single property at a time. In this talk, we present an innovative approach using differentiable simulations to inversely determine mechanical properties from straightforward experimental observations. This method is validated across diverse systems, including articulated robots and non-Newtonian fluids like ketchup, demonstrating the versatility and efficiency of this simulation-based technique in various robotics and mate-

rial contexts.

Peter Yichen Chen
MIT CSAIL
pyc@csail.mit.edu

MS296

Comparison of Ad-Driven Finite Element Updating and Virtual Fields Method for Material Model Calibration from Full-Field Data

The two primary approaches in the experimental mechanics' community for calibration from full-field DIC data are known as finite element model updating (FEMU) and the virtual fields method (VFM). In the VFM, the objective function is a squared mismatch between internal and external virtual work or power instead of measured and predicted displacement or strain subject to PDE constraints In FEMU, an objective function that quantifies the weighted mismatch between model predictions and corresponding experimentally-measured quantities of interest is minimized by iteratively updating the parameters of an FE model using an optimization algorithm. While FEMU is seen as more flexible, practitioners generally prefer VFM over FEMU because the latter is more computationally demanding. However, these comparisons are usually made by approximating gradients with finite difference schemes. However, for both cases, forward and adjoint sensitivities can be used to compute the gradient at considerably less cost than its calculation from finite difference approximations. Hence, in this talk, we rigorously compare VFM and FEMU in the context of automatic differentiation which lowers the computational cost of FEMU significantly. To this end, the two methods are tested on different elastic and inelastic constitutive behaviors and the influence of noise on the accuracies of both methods is studied.

Jan Fuhg
The University of Texas at Austin
jan.fuhg@utexas.edu

Sanjeev Kumar
UT Austin
sanjeev.kumar@austin.utexas.edu

Brian Granzow
Sandia National Laboratories
bngranz@sandia.gov

Daniel T. Seidl
Sandia National Lab
dtseidl@sandia.gov

MS296

Calibration of Elastoplastic Constitutive Model Parameters from Full-field Data with Automatic Differentiation-based Sensitivities

The two most widely used inverse methods in the experimental mechanics community are finite element updating (FEMU) and the virtual fields method (VFM). Both approaches take full-field deformation measurements provided by digital image correlation (DIC) techniques and load data as inputs and yield estimates of constitutive model parameters. Techniques from local sensitivity analysis have been applied in various fields to produce numerically-exact sensitivities though the formulation and solution of auxiliary linear partial differential equations (PDEs). They require an invasive modification of standard FE formulations, but they also reduce the computational cost of inversion relative to standard approaches that rely on finite difference approximations. There are three challenges with the application of such methods to plasticity models. The first is the proper treatment of the coupled nature of the internal variables evolution equations and the equilibrium PDE in the derivation of the sensitivity computations. The second is the history-dependence inherent in these models, which causes the sensitivities to vary in time. The final obstacle is the complex and often tedious derivation of the sensitivities for a given constitutive model. We use automatic differentiation (AD) to solve this problem. We present results from the application of our AD-based approach to calibration from experimental measurements of full-field displacement data obtained via stereo DIC.

Daniel T. Seidl
Sandia National Lab
dtseidl@sandia.gov

Brian Granzow
Sandia National Laboratories
bngranz@sandia.gov

MS296

Sensitivity Analysis and Design Optimization of Lagrangian Shock Hydrodynamics

OptimiSM is a library for posing and solving problems in solid mechanics using the finite element method. The central theme of this project is exploring how to get better performance and robustness by taking advantages of the tools of variational calculus. OptimiSM uses Lagrangian field theory to pose hard nonlinear solid mechanics problems as optimization problems, and then uses powerful optimization methods to solve them efficiently and reliably.

Brandon Talamini
Lawrence Livermore National Laboratory
talamini1@llnl.gov

MS297

Particle-in-Cell Particle Grid Adaptivity and Resampling

Particle methods for advective flows require periodic resampling or remapping of distribution function because the particle grid becomes distorted from its original, presumable high quality, configuration from the dynamics of the physics and no longer be an effective grid, resulting in numerical problems such as "particle noise". This process is akin to adaptive mesh refinement in a continuum method and is a fundamental method required for modern discretizations of PDEs. This talk presents a new framework that allow for remapping a distribution to almost any new particle distribution with a simple and well posed algorithm that can support any adaptivity strategy, such as "quite start" in magnetized plasmas, or any application specific grid optimizations, as well as traditional grid efficiency metrics of mesh adaptivity and application specific metrics like minimizing entropy generation. This algorithm also conserves an arbitrary number of moments of the distribution exactly. The approach will be demonstrated with the two-stream instability electrostatic plasma problem in one spatial and one velocity dimension using PETSc.

Mark F. Adams
Lawrence Berkeley National Laboratory

mfadams@lbl.gov

MS297

Using TRIXI.JL for Adaptively Coupling Multiphysics Problems

To couple arbitrary multiphysics problems we make use of the numerical simulation framework Trixi.jl. In our implementation the coupling happens through the exchange of corresponding boundary values. The user is free to define any type of coupling function, which allows for different kind of physics. In particular, our method is capable of coupling a hierarchy of systems. To capture dynamical change we also show adaptively coupled simulations done with our implementation. There, the domain boundaries can chance dynamically with the change criterion left up to the user to define. One application is the propagation of magnetic fields in space where we solve the magnetohydrodynamic equations only for the part of the domain with a significant magnetic field.

Simon Candelaresi, Michael Schlottke-Lakemper
University of Augsburg
simon.candelaresi@uni-a.de, michael.schlottke-lakemper@uni-a.de

MS297

Simulating Magnetic Reconnection Using the Petsc Particle-In-Cell Framework

Numerical solutions to the Vlasov-Maxwell equations have important applications in solar and magnetospheric physics. At the microscale, electrons and ions capture energy released from powerful magnetic events, such as reconnection, and in-turn play a role in driving the macroscopic systems. Advancements in plasma modelling provide important insights used to design the next generation of solar imagers and other instruments. The goal of this research is to extend the existing electrostatic, Vlasov-Poisson, PETSc Particle-in-Cell (PETSc-PIC) framework to self-consistent electromagnetic systems. In the Vlasov-Maxwell PETSc-PIC framework, Maxwells equations are solved using the finite element method while the Vlasov equation is solved with conservative, explicit basic symplectic time integrators. Collisions may be added to the formulation by means of a particle-basis Landau collision operator. With an electromagnetic, Vlasov-Maxwell, PETSc-PIC model, we may study the kinetic-scale transfer of magnetic energy to particles in reconnection systems, such as Harris current sheets.

Daniel Finn
University at Buffalo
dsfinn@buffalo.edu

MS297

Entropy Minimization in Resampling

We look at the particle remapping problem for PIC discretizations of kinetic plasma physics problems. In addition to conservation and projection properties, we look at the behavior of the entropy of the distribution. Since we have the entropy gradient from evaluation of the collision operator, we explore its use in remapping particles.

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

MS298

Distributed Zeroth-Order Gradient Tracking Methods for Heterogenous Systems

We consider two classes of stochastic mathematical programs with equilibrium constraints (MPECs) over networks. These include single-stage and two-stage distributed stochastic MPECs. Of these, the first problem class is motivated by addressing heterogeneity in distributed learning, while the second problem class is motivated by transportation system design under uncertainty. We develop provably convergent gradient tracking algorithms for solving stochastic MPECs over networks, where computing agents can communicate with their neighbors over an undirected connected network. To this end, we propose two novel methods: single-stage distributed implicit zeroth-order stochastic gradient tracking (DiZS-GT-1s) and its two-stage variant (DiZS-GT-2s). We establish convergence properties and derive new iteration and sample complexity bounds. We further present preliminary numerical experiments to compare the performance of the proposed methods across various network configurations with the centralized counterparts.

Mohammadjavad Ebrahimi
Rutgers University
me586@scarletmail.rutgers.edu)

MS298

Convergence Analysis for Randomized Linear Gradient Compression with Applications to Private Training

In distributed optimization, the communication of model updates can be a performance bottleneck. Consequently, gradient compression has been proposed as a means of increasing optimization throughput. This may be particularly useful in distributed private training scenarios, where the participating entities may be geographically far apart and subject to higher communication costs. Recent work investigated how the iteration penalty depends on the interaction between compression and problem structure. It was shown that for several distributions of compression matrices, among them random orthogonal matrices, the impact of compression on convergence can be quantified in terms of the norm of the Hessian of the objective, using a norm defined by the compression scheme. The analysis reveals that in certain cases, compression can benefit from low-rank structure in the problem. In this work we investigate the effects of compression in the private training scenarios. We observe that for training model with differential privacy, including generative models such as differentially private diffusion models, our bounds predict that the penalty introduced by compression is significantly reduced compared to worst-case bounds that only consider the compression level, ignoring problem data.

Thomas Flynn
Brookhaven National Laboratory
thomasflynn918@gmail.com

MS298

Privacy-Preserving Federated Learning across DOE Supercomputers

Privacy-preserving federated learning (PPFL) enables se-

cure collaboration on foundation models (FMs) across distributed datasets. This presentation highlights APPFL, an open-source framework addressing key challenges in PPFL, including scalability, data heterogeneity, and privacy preservation. We showcase results from FL experiments conducted across DOE supercomputers (Frontier, Polaris, and Perlmutter), leveraging Globus Compute and Proxy Store for resource management. These experiments demonstrate APPFL's capability to train models of various sizes efficiently while maintaining privacy.

Kibaek Kim
Argonne National Laboratory
kimk@anl.gov

MS299

Multimodal Deep Learning for Medical Diagnosis

Acute respiratory distress syndrome (ARDS), a severe form of hypoxemic respiratory failure with high in-hospital mortality, poses diagnostic challenges that can delay crucial interventions. Existing predictive models have made use of individual data sources that include ventilator waveform data (VWD), chest x-rays, and electronic health records to screen patients for ARDS, but no model has incorporated all three. This talk will explore some of the deep learning-based approaches we have explored to combine the three modalities into a single predictive model to improve the accuracy and speed of ARDS detection. The approaches could serve as a framework for use in other healthcare settings, where modalities with complementary information could be combined to improve upon the predictive performance of single-modality models.

Stefan Broecker
University of California, Davis
sabroecker@ucdavis.edu

MS299

Bayesian Discovery on Governing Equations

Spatio-temporal modeling of real-world data presents significant challenges due to high-dimensionality, noisy measurements, and limited data. In this talk, we introduce two frameworks that jointly solve the problems of sparse identification of governing equations and latent space reconstruction: the Bayesian SINDy autoencoder and SINDy-SHRED. The Bayesian SINDy autoencoder leverages a spike-and-slab prior to enable robust discovery of governing equations and latent coordinate systems, providing uncertainty estimates in low-data, high-noise settings. In parallel, SINDy-SHRED (sparse identification of nonlinear dynamics with shallow recurrent decoder network) integrates Gated Recurrent Units (GRUs) with a shallow decoder network to model temporal sequences and reconstruct full spatio-temporal fields using only a few sensors. Our proposed methods utilize a SINDy-based regularization. Beginning with an arbitrary latent state space, the dynamics of the latent space progressively converges to a SINDy-class functional. We conduct a systematic experimental study including synthetic PDE data, real-world sensor measurements for sea surface temperature, and direct video data. Particularly, SINDy-SHRED demonstrates robust generalization in a variety of applications with minimal to no hyperparameter adjustments. The interpretable SINDy model of latent state dynamics enables accurate long-term video predictions, achieving state-of-the-art performance among all baseline methods considered.

Mars L. Gao
Universty of Washington
Paul G. Allen School of Computer Science & Engineering
marsgao@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS299

Coincident Learning for Beam-Based Rf Station Fault Identification Using Phase Information at Slac Linac Coherent Light Source

The vast amount of data generated by accelerators makes manual monitoring impractical due to its labor-intensive nature. Existing machine learning solutions often rely on labeled data, manual inspection, and hyperparameter tuning, which limits their scalability. To address these challenges, we leverage coincidence learningan unsupervised technique designed for multi-modal tasksto automatically detect anomalies by identifying coincident patterns of behavior across two distinct segments of the feature space. Specifically, we focus on anomaly detection for radio-frequency (RF) stations at the SLAC Linac Coherent Light Source (LCLS). By analyzing shot-to-shot data from the beam position monitoring system alongside data from RF stations, we can identify the source of changes in the accelerator's status. Previous studies on RF stations produced reasonable results using time-asynchronous amplitude data, but ignored the richer information from time-synchronous phase data due to its complexity. We find that using neural networks to analyze the phase data enables the detection of anomalies that amplitude-based detection missed. Additionally, the rich information contained in the phase data facilitates clustering of anomalies into distinct categories, each with unique signatures. This categorization brings us closer to identifying the root causes of issues within the RF stations.

Jia Liang
Stanford University
jialiang@stanford.edu

MS299

Improving Particle Filters Using Deep Probabilistic Models

Particle filters (PFs) are powerful tools for state estimation in non-linear and non-Gaussian systems. However, traditional PFs can be computationally expensive, especially in high-dimensional problems. These challenges arise from the reliance on importance sampling, typically using the transitional density as the importance density, which neglects the most recent observations and often leads to particle degeneracy. The Unscented Particle Filter (UPF) attempts to address this issue using a Gaussian importance density derived from the Unscented Kalman Filter (UKF). However, UPFs require UKF updates for every particle at every step, making them computationally more demanding. This work proposes a novel approach to improve the efficiency of PFs by using deep probabilistic models to learn a suitable importance density. This learned importance density is then used to generate particles more consistent with the true state, leading to more efficient estimation. The effectiveness of the proposed method is demonstrated

through numerical examples using the Lorenz96 model. Results show that the proposed method achieves lower estimation errors while using significantly fewer particles compared to traditional PFs and UPFs. This improved efficiency makes the proposed method particularly well-suited for high-dimensional, non-linear, and non-Gaussian systems.

Apoorv Srivastava
Stanford University
apoorv1@stanford.edu

MS300

Tensor Methods in Deep Learning and Ai4science

Tensor methods have a long and rich history across diverse scientific fields. Despite their ability to enable significant parameter savings, computational efficiency, and enhanced performance in various applications, their adoption in deep learning remains limited. As deep learning increasingly drives scientific innovation (AI4Science), we stand at a crossroads with new opportunities for broader integration of tensor-based methods. This presentation will explore the current landscape of tensor methods in deep learning, examining the barriers to their wider adoption and the opportunities they present for modern machine learning.

Jean Kossaifi
NVIDIA
jkossaifi@nvidia.com

MS300

Tensor Network Solvers for Neutron Transport

Tensor networks techniques, known for their low-rank approximation ability of large tensors and matrices, have recently been successfully applied to simple discrete ordinates, multi-group, neutron transport eigenvalue problems. Recent research has focused on solving realistic problems using tensor networks such as those solved when designing nuclear reactors. We discuss the advantages and limitations of tensor network methods when solving the discretized neutron transport equation. We discuss the impacts of physical geometric features and nuclear data on tensor rank, compression, and solution times. We also discuss a generalization of previous work to time-dependent neutron transport. As tensor networks become another tool in solving PDE's, the high-dimensional nature of the hyperbolic discretized neutron transport equation presents a rigorous challenge to these methods.

Mario Ivan I. Ortega, Duc Truong, Ismael Boureima, Kim Rasmussen
Los Alamos National Laboratory
mortega@lanl.gov, dptruong@lanl.gov,
iboureima@lanl.gov, kor@lanl.gov

Boian Alexandrov
Los Alamos National Laboratory
T-1: Physics and Chemistry of Materials
boian@lanl.gov

MS300

A Semi-Lagrangian Adaptive-Rank Method (slar) for Linear and Nonlinear Transport Pdes

High order semi-Lagrangian methods for kinetic equations has been well-developed in the past few decades. In this work, we propose a semi-Lagrangian finite difference method that explore the adaptive-rank structure of the Vlasov-Poisson solution to further improve computational efficiency. Besides using extra large time stepping sizes via the semi-Lagrangian setting, the proposed method explores the low rank structure of the Vlasov solution by the cross approximation of matrices, which is also known as the CUR decomposition or pseudo-skeleton approximation. Such approximation could be obtained by selecting the columns and rows that best represent the solution matrix via a randomized pivoting strategy. Following the semi-Lagrangian update of the Vlasov solution via cross approximation, we apply a singular value truncation, as well as a mass conservative projection, of the Vlasov solution, for numerical stability and local mass conservation. The computational complexity scales linearly with respect to the mesh size N per dimension, in contrast to a N^2 for traditional full rank schemes, in each time step. A wide range of benchmark tests are performed, to demonstrate the efficiency and effectiveness of the proposed scheme.

Nanyi Zheng, Daniel Hayes
University of Delaware
nyzheng@udel.edu, dphayes@udel.edu

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

Jingmei Qiu
University of Delaware
jingqiu@udel.edu

MS301

Dimensionality Reduction in Modeling Atmospheric Impacts on Extreme Weather

Accurately modeling atmospheric phenomena like the Quasi-Biennial Oscillation, El Nio-Southern Oscillation (ENSO), stratospheric warming events, and large high- and low-pressure systems is essential for predicting weather patterns, hurricanes, and extreme events like storm surges. However, existing computational models face challenges, particularly due to the high cost of fine-tuning physical parameters. Dimensionality reduction techniques offer a way to efficiently estimate and understand these parameters by extracting latent features that capture the core characteristics of complex geophysical data. By exploring these approaches, we gain insights into the influence of hyper-parameters, paving the way for more precise and computationally efficient atmospheric modeling.

Riti Bahl
Emory University
riti.bahl@emory.edu

MS301

Optimization for Data Driven Reduced-Order Modeling

The optimal state-feedback gain for the Linear Quadratic Regulator (LQR) and other control problems is computationally costly to compute for high-order systems. Reduced-order models (ROMs) can be used to compute feedback gains with reduced computational cost. However, the performance of this common practice is not fully understood. We discuss theoretical and empirical results showing when POD is optimal and propose optimization methods

to improve on POD when possible.

Laura Balzano, Rachel Newton, Peter Seiler
Electrical Engineering and Computer Science
University of Michigan
girasole@umich.edu, rhoadesr@umich.edu,
pseiler@umich.edu

MS301

Algorithms for Reduced-Order Modeling and Data Assimilation.

We consider algorithms for the reduced-order modeling (ROM) of complex turbulent flows that exhibit propagation of coherent structures and shock waves, such as found in rocket nozzles. These are multi-scale, multi-physics dynamical systems with complex geometries. Our work investigates the coupling of spatially localized ROMs, the required interface conditions and the assimilation of sparse data. We explore conditions for stability and accuracy of the methods, the use of DEIM and system transformations (lifting) and consider a different method based on generalized Gaussian quadrature. Additionally, we use a simpler chaotic system as a proxy for turbulence to gain deeper insights into the performance of these ROMs.

Marc De Vernon
Southern Methodist University
marcdevernon@gmail.com

MS301

Learning Physics-Based Reduced-Order Models from Data Using Nonlinear Manifolds

Despite the remarkable rise of available computing resources, the need for model order reduction to cope with complex problems is an ever-present reality. Reduced-order models are imperative in making computationally tractable outer-loop applications that call for repeated simulation tasks. They require that one numerically solves the differential equations describing the physical system of interest in low-dimensional reduced spaces, in contrast to the original full-order models. However, traditional model reduction techniques often fail to identify a low-dimensional linear subspace for approximating the solution to many physics-based simulations. In this talk I will present a novel method for learning projection-based reduced-order models of physics-based dynamical systems using nonlinear manifolds. First, we learn the manifold by identifying nonlinear structure in the data through a general representation learning problem. The proposed approach is driven by embeddings of low-order polynomial form. The algebraic structure of the system that governs the problem of interest in the reduced space is revealed by means of a projection onto the nonlinear manifold. The matrix operators of the reduced-order model are then approximated, in a least-squares sense, using data-driven operator inference. Numerical experiments on a number of nonlinear problems demonstrate generalizability and the increase in accuracy that can be achieved.

Rudy Geelen
Texas A&M University
rgeelen@tamu.edu

Laura Balzano
Electrical Engineering and Computer Science
University of Michigan
girasole@umich.edu

Stephen Wright
University of Wisconsin-Madison
Dept. of Computer Sciences
swright@cs.wisc.edu

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

MS301

Koopman Theory and Wavelets: A Framework for Wavelet-based Dynamic Mode Decomposition

In this work, we present an in-depth analysis of the Koopman operator using wavelet theory and introduce the wavelet-based observables. We establish that for the Morlet wavelet, the action of the Koopman operator on these observables takes a simple closed form within a weighted L2 space. To approximate the Koopman operator numerically, we employ the Extended Dynamic Mode Decomposition method combined with the proposed wavelet-based observables. The observables are computed using the Synchrosqueezing Transform. We demonstrate the effectiveness of this approach on the Lorenz and FitzHugh-Nagumo systems, showing that it performs robustly even in the presence of moderate noise.

Cankat Tilki, Serkan Gugercin
Virginia Tech
Department of Mathematics
cankat@vt.edu, gugercin@vt.edu

MS302

Robot Dynamics Learning and Control Using Port-Hamiltonian Neural ODE Networks

Accurate models of robot dynamics are critical for safe and stable control of autonomous mobile robots. This talk will present machine learning techniques for robot dynamics identification, and discuss the use of learned dynamics models for safe autonomous robot navigation. Physical systems have configurations that evolve on manifolds and satisfy energy conservation principles. This motivates the use of neural ordinary differential equations with (port-)Hamiltonian structure on the system's configuration manifold to learn models of system dynamics that respect kinematics and energy constraints. Such models may be learned from trajectory data or sensor observations. Their Hamiltonian structure provides a natural Lyapunov function candidate to aid in the synthesis of control laws for stabilization, trajectory tracking, and safety constraint satisfaction.

Nikolay Atanasov
UC San Diego
natanasov@ucsd.edu

Thai Duong
University of California San Diego
tduong@eng.ucsd.edu

MS302

Bringing Numerical Methods and Deep Learning with Physics-Constrained Differentiable Solvers

Machine learning (ML) is increasingly playing a pivotal role in spatiotemporal modeling. A number of open questions remain on the best learning strategies to maximize the util-

ity of machine learning while ensuring the validity of such predictions, particularly in limited data scenarios. This talk will focus on exploring machine learning strategies for neural PDE solvers, with an emphasis on broad learning strategies that are applicable across a wide variety of systems and neural network architectures. Some topics I will discuss include: using self-supervised learning to change the basis of learning with spectral methods to solve fluid dynamics and transport PDE problems, and simulation-in-the-loop approaches via incorporating PDE-constrained optimization as a layer in neural networks. In each of these settings, I will discuss how ML methods can be used with numerical methods through fully differentiable settings.

Aditi Krishnapriyan
Chemical Engineering and EECS, UC Berkeley
aditik1@berkeley.edu

Nithin Chalapathi
University of California Berkeley
nithinc@berkeley.edu

MS302

Certified Learning of Lagrangian Dynamics from Data with Uncertainty Quantification

I will show how to use Gaussian Process regression to learn variational ODEs or PDEs from data with guaranteed convergence. From a statistical framework, uncertainty quantification for observables such as the Euler-Lagrange operator and Hamiltonians can be derived. The method is based on learning local computational stencils that can recover solutions. This strategy is shown to have surprising extrapolation properties to unseen data regimes.

Christian Offen
Paderborn University
christian.offen@uni-paderborn.de

MS303

Modeling and Simulating Multiphysics Systems with Discrete Exterior Calculus

Solving multiphysics systems requires the application of specialized techniques to the formulation of complex problems to derive new implementations for each combination of physical equations in a system. We present the Decapodes.jl software package that automatically assembles such multiphysics implementations from specifications in the discrete exterior calculus using a synthetic differential geometry approach and using de Rahm complexes as a discrete model of differential geometry. Attention is paid to recent developments necessary for the solution of fluid mechanical models such as pressure projection for Euler's and vorticity formulations of incompressible Navier-Stokes equations.

James P. Fairbanks, Luke Morris
University of Florida
fairbanksj@ufl.edu, luke.morris@ufl.edu

MS303

An Application of the DivDiv Complex in Numerical Relativity

The numerical simulation of the evolution of spacetime poses many challenges. As a first step, in this talk we are interested in the simulation of the linearized equations around the Minkowski metric. The resulting problem requires to finding a matrix-valued field satisfying important but nontrivial constraints. The goal is to show how the system can be reformulated on the DivDiv complex, in a way that naturally translates the constraints to the complex. This new formulation can be seen as a special case of a time-dependent Hodge-Dirac problem. We will then show how this problem can be efficiently approached using the framework of exterior calculus.

Marien Hanot
Université de Montpellier
marien-lorenzo.hanot@univ-lille.fr

Kaibo Hu
University of Edinburgh
kaibo.hu@ed.ac.uk

MS303

Discretely De Rham Conforming Hierarchical B-Splines for Stable Fluid Flow

In this presentation, new locally-sufficient rules will be presented and computationally validated demonstrating accuracy and numerical stability of the spline-based hierarchical B-spline complex of discrete differential forms applied to fluid analyses. Numerous partial differential equations, including those important for fluid flow and electromagnetic phenomena, must discretely obey conservation laws for numerical stability. In many instances, these conservation laws can be best represented by using function spaces for analysis that discretely adhere to the de Rham complex of differential topology. However, naive approaches for hierarchical refinement of these spline spaces will typically yield solution spaces that are not inf-sup stable. Herein, rules that are necessary to guarantee stability are presented for hierarchical B-spline domains of arbitrary dimension. The rules are coupled with numerical results numerically indicating stability of the proposed methods for the Stokes complex. Finally, preliminary results applied to the Navier-Stokes equations are presented. In sum, the proposed methods fill an important need for high-order, geometrically-accurate, locally-refineable and stable finite element spaces for fluid and electromagnetic analysis in arbitrary dimensions.

Kendrick Shepherd, Daira Sofia Velasco Vega
Brigham Young University
kendrick_shepherd@byu.edu, dairasv@byu.edu

Caleb B. Goates
Brigham Young University
Department of Civil and Construction Engineering
calebgoates@gmail.com

MS304

Space-time Optimal Control of Electromagnetic Systems

Maxwells equations are an important modeling component in a variety of modern engineering applications, ranging from nanoscale optical devices to magnetic confinement fusion. Despite the significant advances in the mathematical analysis and numerical methods for the solution of Maxwells equations, little progress has been made in our understanding of optimal control and optimal design problems governed by Maxwells equations. The goal of this talk is to discuss the mathematical and computational challenges associated with such optimization problems. A

theoretical framework for the time-discrete forward problem is developed. Well-posedness of the optimization problem is then established by the direct method, and the first order necessary and sufficient optimality conditions are derived. A physics-compatible finite element pair for electric and magnetic field is used, and the implementation is carried out in an advanced electromagnetic software MrHyDE (Sandia National Labs). Capabilities of the Rapid Optimization Library (ROL) are used to carry out large scale optimization.

Yaw Owusu-Agyemang
George Mason University
yowusuag@gmu.edu

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

Harbir Antil
George Mason University
hantil@gmu.edu

Jimmie Adriazola
Arizona State University
jimmie.adriazola@asu.edu

Rohit Khandelwal
George Mason University
rkhandel@gmu.edu

MS304

Time-parallel Optimal Control of Nonlinear Dynamical Systems

A key computational challenge in solving dynamic optimization problems is due to the serial nature of conventional forward and reverse (adjoint) time stepping methods. While some computational gains can be achieved by parallelizing in time the forward and adjoint solvers, greater improvements are realized by parallelizing the solution of coupled optimality systems. In this talk, we present recent advances toward scalable and efficient time-parallel optimal control, which are grounded in transformations of optimality systems that lead to multigrid-in-time schemes. Our contributions include novel complexity estimates that rely on scalings inspired by the function-space characteristics of optimal control problems.

Denis Ridzal, Radoslav G. Vuchkov, Aurya Javeed
Sandia National Laboratories
dridzal@sandia.gov, rgvuchk@sandia.gov,
asjavee@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

MS304

Neural Network Approaches for High-Dimensional Optimal Control and Transport Problems

We present neural network approaches for high dimensional optimal control and optimal transport problems. The Neural-HJB method approximates optimal policies, enabling real-time control while alleviating the computational burden of traditional approaches. Similarly, the COT-Flow method applies DNNs to solve dynamic conditional optimal transport problems, which is critical for applications like sampling and density estimation in Bayesian inference. Empirical results highlight the effectiveness of both methods, showcasing their performance compared to state-of-the-art alternatives.

Deepanshu Verma
Clemson University
dverma@clemson.edu

Lars Ruthotto
Departments of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Bart Waanders
Sandia National Labs
bartv@sandia.gov

Nick Winovich
Sandia National Laboratories
nwinovi@sandia.gov

Ricardo Baptista
Caltech
rsb@caltech.edu

Youssef M. Marzouk, Youssef M. Marzouk, Oliver Wang
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu, olivrw@mit.edu

MS304

Implicit Shock Tracking - A PDE-Constrained Optimization Approach to Accurate Resolution of High-Speed Viscous Flows

Shock tracking or shock fitting, where the computational mesh is moved to align mesh element faces with discontinuities, represents non-smooth flow features with the inter-element jump in the solution basis without requiring additional stabilization. In our previous work, we introduced an implicit shock tracking framework that discretizes conservation laws on a mesh without knowledge of discontinuities formulated as a PDE-constrained optimization problem over the discrete solution variables and nodal coordinates of the mesh. The optimization problem is solved using a full space sequential quadratic programming method that simultaneously converges the mesh and flow solution to their optimal values. In this talk, we extend the implicit shock tracking framework to viscous problems where flow features (e.g., shocks, boundary layers) are thin features with steep gradients.

Matthew J. Zahr
University of Notre Dame
mzahr@nd.edu

MS305

A Benchmark Dataset for Evaluating the Privacy and Utility Risks in Synthetic Health Data

There has been a notable increase in interest and implementation of tabular synthetic data generation (SDG) models. SDG is typically achieved by modeling a joint distribution from real data and then generating a synthetic dataset (SD). The resulting SD should mimic the statistical properties of the real data while preserving privacy. The

privacy and utility risks vary among and within SDG techniques. Research shows that the utility and privacy risks of even the most effective SDG models, such as Generative Adversarial Networks (GANs), can differ depending on the dataset used. However, which dataset characteristics make certain SDG models perform better remains unclear. Comparisons across studies are complicated by the use of different datasets, making general conclusions difficult. A standardized approach is needed to evaluate the privacy and utility of SDG. Such an approach would allow developers, users, and regulators to compare SDG techniques objectively and set performance benchmarks for specific purposes. In this presentation, we discuss progress in developing optimized benchmark datasets to differentiate SDG models in terms of privacy and utility. The discussion covers (i) the modeling (parametrization) of the dataset space, (ii) various generative models employed, (iii) privacy metrics applied for optimization, (iv) the optimization algorithm used, and (v) the results achieved.

Maryam Khani
University of Ottawa
mkhani@ehealthinformation.ca

Fida Dankar
New York University Abu Dhabi
fdankar@ehealthinformation.ca

Khaled El Emam
University of Ottawa
kelemam@uottawa.ca

MS305

Evaluation Metrics for Crack Detection with Machine Learning

As part of the US Department of Energys 3013 surveillance program, a machine learning model has been developed to detect cracks in images taken of 3013 canisters used to store surplus plutonium material. An hourglass neural network is trained on efficiently labeled training data generated by a laser confocal microscope or scanning electron microscope. The highly imbalanced, feature based nature of this data creates a challenge for evaluating the performance of the model to optimize its performance.

Stephanie Gamble
Savannah River National Laboratory
stephanie.gamble@srnl.doe.gov

MS305

Challenges and Solutions for Developing Machine Learning Benchmarks on Mission-Driven Scientific Applications

Algorithms developed in fundamental research settings are often vastly different from algorithms deployed in real-world scientific operations. While machine learning (ML) benchmarks are frequently designed to serve as the foundation for novel state-of-the-art research, benchmarks can also be designed to serve as a bridge between novel research and operational ML. In this introductory presentation, we use ML benchmarks for anomaly detection as a lens to examine key challenges for creating ML benchmarks that are both accessible and realistic. Using examples from Sandia National Laboratoriesas well as the broader literaturewe survey differences between the benchmark data used in research and the benchmark data used in operations. We also explore the performance metrics used in fundamental research compared to full decision pipelines used in operations. Finally, we provide a brief overview of current and future solutions for these important problems. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Amelia Henriksen
Sandia National Laboratories
aahenri@sandia.gov

MS306

Distributed Memory Tensor Train Cross for Large Data

When working with large data sets, the curse of dimensionality places tight restrictions on storage requirements as well as performing practical computations. In recent years, significant work has been made towards dealing with these issues in various forms of tensor decompositions. In this work, we propose an algorithm for a subtensor parallel Tensor Train Cross (TT - Cross) decomposition which is fit for a distributed memory setting. We will show that this algorithm maintains a low storage requirement, as well as a small communication cost throughout its stages. This facilitates for the effective use of large computing systems, which will utilize local tensor information to construct a global TT - Cross approximation. Numerical tests will be presented to display the scaling results as well as storage requirements in a synthetic benchmark, as well as a tensor that arises in PDE simulations.

Daniel Hayes
University of Delaware
dphayes@udel.edu

MS306

Towards the Efficient Approximation of Higher-Order, Tensor-Network Contractions Via a Low-Rank, Matrix-Free Tensor Decomposition

Tensors (multi-dimensional arrays) are among the most basic tools for computational modeling. Unfortunately, tensors are plagued by the so called "curse of dimensionality" which refers to the exponential complexity associated with accessing and manipulating tensors. To combat the curse of dimensionality, mathematicians utilize (approximate) tensor decomposition formats to recast large and expensive tensors into sets of smaller and more manageable ones. These decompositions replace exact tensors in tensor-networks that bottleneck the performance of modeling algorithms. However, it has been shown[Pierce, Rishi and Valeev, Robust approximation of tensor networks, 2021] that errors associated with the approximation of tensors in a network of tensor contractions can have significant impacts on the overall accuracy of the network. This is because small perturbations introduced by single tensor approximations are propagated and amplified through a tensor-network. We investigate a general means to alleviate this error propagation by introducing a matrix-free tensor decomposition to approximate entire exact tensor-networks in quantum physics and chemistry applications. This method does not rely on either the tensor-network or decomposition format. However, in this study, we focus our investigation on the application of the canonical polyadic decomposition because of its favorable structure and relation to previous studies such as the tensor hyper-contraction method.

Karl Pierce, Joseph Tindall, Miles Stoudenmire

Flatiron Institute
kpierce@flatironinstitute.org, jtindall@flatironinstitute.org, mstoudenmire@flatironinstitute.org

MS306

Parallel Higher-Order Orthogonal Iteration for Tucker Decomposition with Rank Adaptivity

HOOI (Higher Order Orthogonal Iteration) also known as Tucker-ALS (Tucker Alternating Least Squares) is an iterative optimization algorithm that uses block coordinate descent optimization to compress a multidimensional data tensor into a Tucker tensor (a smaller core tensor and factor matrices for each dimension). Classic HOOI is rank specified, which means we can solve for the factor matrices an using an SVD algorithm to set the factor matrices to the first r_k columns of the left singular vectors. Different works have attempted to adapt rank specified HOOI into an error specified variant. But unlike the HOSVD and ST-HOSVD algorithms, there is no commonly accepted error-specified version of HOOI. We introduce HOOI-Adapt, a novel error-specified algorithm that uses a core analysis technique to select the ranks of the Tucker tensor. Additionally, sequential and parallel variants of the proposed algorithms have been implemented in TuckerMPI. Furthermore, we present performance results on the NERSC Perlmutter cluster and perform cost analysis of the proposed parallel algorithms using the Hockney (alpha-beta-gamma) performance model.

João Pinheiro
Wake Forest University, U.S.
deolj19@wfu.edu

MS306

Parallel Sylvester Tensor Equation Solvers in the Tensor-Train Format

Sylvester tensor equations are a class of linear tensor equations that appear frequently in computational mathematics and applications. In particular, discretized PDEs with Laplace-like operators generate these tensor equations with sparse given matrices. In this talk, we focus on designing a distributed parallel algorithm for Sylvester tensor equations in tensor-train (TT) format. The cornerstone of our solver is the factored alternating direction implicit (fADI) method to solve Sylvester matrix equations via a sequence of shifted linear solvers. We use a 1D process partitioning over the physical dimensions to distribute the TT cores. We also design a two-directional variation of the solver for better load balance, and an iterative refinement stage to improve accuracy.

Tianyi Shi
Lawrence Berkeley National Laboratory
tianyishi@lbl.gov

MS307

Mathematical Modelling of Cancerous Tumour Growth

The dynamics of cancer cells and their interactions with immune system has been a subject of interest in research over the years due to the complexity in the interactions between the tumour cells and the immune system. In this paper, a model on Cancerous tumour growth which which captures the effects of lymphocytes in the growth and treatment of cancer is presented. A numerical analysis of the model indicates that a person with low white blood cells count is at a high risk when affected by tumour growth. The equilibrium state of the models is used to determine conditions for tumour free equilibrium and to verify the effect of white blood cells. The simulation results reveals that the model employed is a robust way of studying the dynamics of tumour cells and shows the interactions between the tumour cells, immune system, drug response and give and an insight into some factors to be considered in the treatment of cancers.

Mary O. Durojaye
UNIVERSITY OF ABUJA, ABUJA
FCT, ABUJA, NIGERIA
mayojaye@yahoo.com

Peter Odey
University of Abuja
Nigeria
peterachi48@gmail.com

MS308

Structure Preserving Hybrid Finite Volume Finite Element Method for Compressible MHD

In this talk we present a novel and efficient numerical method for the compressible viscous and resistive MHD equations suitable for all Mach number regimes. The time-integration strategy is a semi-implicit splitting, combined with a hybrid finite-volume (FV) and finite-element (FE) discretization in space. The non-linear convection is solved by a robust explicit FV scheme, while the magneto-acoustic terms are treated implicitly in time. As a direct consequence, the resulting CFL stability condition is based only on the fluid velocity. The magneto-acoustic terms are discretized by compatible FE based on a continuous and a discrete deRham complexes designed using Finite Element Exterior Calculus (FEEC). Thanks to the use of FEEC, energy stability, magnetic-helicity conservation and the divergence-free conditions can be preserved also at the discrete level. A very efficient splitting approach is used to separate the acoustic and the Alfvénic modes in such a fashion that the original symmetries of the PDE governing equations are preserved. The final algorithm involves only linear, symmetric, and positive definite algebraic systems, handled by the simple matrix-free conjugate-gradient method. While the formulation of the method is very general, numerical results for a second-order accurate FV-FEEC scheme will be presented. Higher order accurate implementation of the presented method is currently under development and, if ready, some preliminary results will be shown.

Francesco Fambri
Max-Planck Institute of Plasma Physics
Garching
francesco.fambri@ipp.mpg.de

Eric Sonnendrücker
Max-Planck-Institut für Plasmaphysik, NMPP division, and
Technische Universität München, Department of Mathematics
eric.sonnendruecker@ipp.mpg.de

MS308

Particle Methods for the Lenard-Bernstein Colli-

sion Operator

An important part of modeling the behavior of plasmas is the diffusive action of collisions. There are several types of collision operators that maintain conservative properties of a system and model these particle interactions. One such operator is the Lenard-Bernstein (LB) operator. This work investigates using a deterministic particle discretization of this operator to model collisions.

Joanna Held
Iowa State University
Department of Mathematics
joheld@iastate.edu

MS308

A Recovery-Based Discontinuous Galerkin Scheme for the Vlasov-Fokker-Planck Equations

When investigating plasma regimes in which collisional physics are relevant, the choice of collision operator naturally determines what physical collisional mechanisms are included in the model. While approximate collision models are typically adequate when considering the larger-scale impact of collisions on the bulk plasma velocity distribution, a more accurate model is required when proper treatment of the high-energy tails of the distribution is relevant, such as in fusion plasmas. The nonlinear Rosenbluth/Fokker-Planck collision operator (FPO) is a highly accurate treatment of the velocity space advection and diffusion on a plasma population due to small-angle Coulomb collisions and includes a velocity-dependent effective collision frequency for accurate modeling of the high-energy tails, but this operator is notoriously nontrivial to efficiently implement numerically. We present results from a recovery-based discontinuous Galerkin implementation of the full nonlinear FPO in the plasma simulation framework `Gkeyll` that is highly accurate and properly captures the cross derivatives in the diffusion term. The scheme is conservative in mass, momentum, and energy through corrections to the drag and diffusion coefficients. Integrals are precomputed and written to computational kernels that are called at runtime, dramatically reducing the computational expense of this complex model.

John Rodman
Virginia Tech
Department of Aerospace and Ocean Engineering
jrodman@vt.edu

Jimmy Juno
Princeton Plasma Physics Laboratory
jjuno@pppl.gov

Bhuvana Srinivasan
University of Washington
srinbhu@uw.edu

MS308

Well-Balanced Schemes for Shallow Water Mhd

Originally introduced to describe a transition region in stars, the shallow water magnetohydrodynamics (SWMHD) model is now used throughout a number of solar physics and geophysical applications. In these applications, it is common to see phenomena that result from just a small perturbation of a steady-state solution. However, if using a standard method to try and capture these numerically, one may miss these small phenomena entirely unless the grid is refined significantly. This refinement may prove quite costly, and even completely unreasonable on large 3-dimensional simulations. Well-Balanced (WB) schemes provide one alternative solution to this issue. Such methods preserve (non-trivial) steady-states of the system to order machine precision, in turn allowing one to capture small perturbations of these steady-states on *coarse* meshes. In this talk, I share our proposed WB finite volume method for both the 1-D and 2-D SWMHD system. These methods also properly treat the divergence-free condition of the magnetic field on a discrete level. The WB and locally divergence-free properties of the proposed schemes are provable, and the proposed method has been successfully tested on several examples.

Michael Redle
RWTH Aachen University
redle@acom.rwth-aachen.de

MS308

A Robust Entropy-Stable Discontinuous Galerkin Scheme for the Multi-Ion Mhd System

We present an entropy-stable discontinuous Galerkin (DG) scheme for the multi-ion magnetohydrodynamics (MHD) equations [Toth et al., Multi-ion magnetohydrodynamics], along with subcell limiting strategies designed to enhance its robustness in challenging plasma simulations. Our approach begins with a continuous entropy analysis of the multi-ion MHD system. We then propose an algebraic transformation that ensures entropy consistency is carried over from the continuous framework to its discrete approximation. Additionally, we incorporate a generalized Lagrange multiplier (GLM) technique to enforce the divergence-free condition on the magnetic field. We develop robust, entropy-conservative, and entropy-stable high-order DG discretizations using collocated Legendre–Gauss–Lobatto summation-by-parts (SBP) operators. These discretizations are consistent with existing entropy-conservative (EC) and entropy-stable (ES) schemes for the single-fluid GLM-MHD equations found in the literature. Our schemes ensure compliance with the second law of thermodynamics at the semi-discrete level while preserving local node-wise conservation properties. To further enhance robustness, we apply subcell limiting strategies to the high-order DG discretization. We validate our scheme through numerical experiments, demonstrating its high-order convergence, entropic properties, and robustness in complex multi-species MHD simulations.

Andrés Rueda-Ramírez, Aleksey Sikstel
University of Cologne
aruedara@uni-koeln.de, asikstel@uni-koeln.de

Gregor Gassner
Mathematical Institute
University of Cologne
ggassner@uni-koeln.de

MS309

Towards Realistic QBO Representation Through Surrogate-Accelerated Multi-Objective Optimization

We develop an end-to-end uncertainty quantification workflow to calibrate the convectively generated gravity waves in the Energy Exascale Earth System Model and obtain a more realistic representation of the quasi-biennial

oscillation (QBO). We introduce a domain knowledge-informed compressed representation of high-dimensional, spatio-temporal dense wind fields through a parsimonious statistical model that learns the fundamental frequency of the stochastic process from noisy observables. This model is used to estimate a small set of interpretable, physically meaningful quantities of interest that capture the QBO's most salient attributes, such as oscillation amplitude and period. We subsequently train a probabilistic surrogate that approximates these fundamental characteristics as a function of key physics parameters driving the convectively generated gravity waves, effectively facilitating inference at a fraction of the cost of a full climate simulation. Lastly, we conduct inference for our inverse problem via surrogate-accelerated multi-objective optimization to improve the QBO. We encounter a tension between the oscillation amplitude and period, which constrains the QBO representation and impedes finding a single solution that satisfactorily matches both objectives. Therefore, we quantify the bicriteria trade-off and generate a representative set of Pareto optimal physics parameter values that balance the conflicting objectives.

Luis Damiano
Sandia National Laboratories
ladamia@sandia.gov

Walter Hannah
LLNL
hannah6@llnl.gov

James Benedict
Los Alamos National Laboratory
benedict@lanl.gov

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS309

Portable Implementation of Large-Scale Simulation-Constrained Optimization Approaches for Bayesian Inference of Ice-Sheet Models

The mass loss from the Greenland and Antarctic ice sheets significantly contributes to global sea level rise. Accurate projections of ice sheet mass loss require sophisticated modeling of ice-sheet dynamics and evolution, while accounting for uncertainties in observational data and computational models. In this presentation, we explore state-of-the-art methods for calibrating Greenland and Antarctic ice sheet models by inverting high-dimensional model parameters. These methods leverage large-scale PDE (Partial Differential Equation)-constrained optimization techniques and employ Bayesian inference to efficiently approximate the posterior distribution of inferred parameters. We detail some computational aspects of our scalable and portable implementation, which is based on algorithmic differentiation (AD) and adjoint methods. Finally, we present results from uncertainty quantification studies enabled by efficient sampling of the parameters' posterior distribution.

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Max Carlson
Sandia National Laboratories
maxcarl@sandia.gov

Trevor Hillebrand
Los Alamos National Laboratory
trhille@lanl.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Kim Liegeois, Jerry Watkins
Sandia National Laboratories
knliege@sandia.gov, jwatkin@sandia.gov

MS309

Multifidelity Bayesian Optimization for Steady-State Predictions using Gyrokinetic Simulations of Plasma Turbulence

The flux matching problem is an important consideration in turbulent gyrokinetic simulations, which are used to study the dynamics of a plasma in fusion devices like tokamaks. By reformulating the flux matching problem (required to predict steady state in plasma systems) as an optimization problem, the problem reduces to finding the optimal background plasma gradients that minimize the discrepancy between the turbulent fluxes computed by the gyrokinetic simulation and the transport model, which considers volumetric heating and losses. We present a multifidelity Bayesian optimization approach for solving the flux matching problem. Our multifidelity approach uses information from 3 different sources: an analytic model of transport (low-fidelity), a reduced model of turbulence (medium-fidelity) and a full gyrokinetic turbulence code (high-fidelity). We investigate how information from these different sources can be combined to reduce the overall computational cost of the optimization procedure.

Pieterjan Robbe
Sandia National Laboratories
pmrobbe@sandia.gov

Pablo Rodriguez-Fernandez, Aaron Ho, Nathan Howard
Massachusetts Institute of Technology
pablorf@mit.edu, aaronkho@mit.edu,
nthoward@psfc.mit.edu

Chris Holland
University of California San Diego
chholland@ucsd.edu

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

MS310

Using Arkouda/Arachne for Understanding Brain Connectome Graphs

Connectomics is a new field that tries to understand how

neurons connect at the synapse level, helping scientists learn about brain functions and how they go wrong in diseases. Recently, new datasets were published that allow for comparing connectomes, giving new views on brain development and differences between populations. Arachne, an advanced tool built on the Arkouda framework and based on the Chapel programming language, helps to solve these problems by offering scalable and parallel graph analytics for big connectome datasets. It makes it easier for researchers to use high-performance computing (HPC) resources directly from their laptops. In this talk, I will explain how Arachne uses data structure and parallel algorithms (Servers) like Subgraph Monomorphism to make exploring brain connectomes, such as those from the Hemibrain and Cerebellum, much more efficient, with speeds up to 100 times faster than current tools. Arachnes integration with Python makes it easy to use while still giving the power of HPC, making it very valuable for neuroscientists. We plan to expand Arachne support workflows for connectome analysis on a petabyte scale seamlessly. This is joint work with Mohammad Dindoost.

David A. Bader
New Jersey Institute of Technology
bader@njit.edu

Mohammad Dindoost
New Jersey Institute of Technology
U.S.
tbd

MS310

Optimizing Distributed Data Structures to Scale Walk and Path Sampling

Random walk and path sampling plays a crucial role in numerous graph algorithms, including tasks like graph embedding and link prediction. As large-scale graphs become increasingly common across various fields, the need for scalable walk sampling solutions has grown. Traditional methods often focus on shared memory systems, where the scalability is constrained by available compute and memory resources. This work presents the Random Neighbor Traversal Graph (RaNT-Graph), a distributed system designed for sampling billions of walks on large, scale-free graphs. RaNT-Graph's partitioning strategy effectively addresses the imbalances caused by high-degree vertices. Additionally, the incorporation of asynchronous walk messages, rejection sampling, and a distributed alias method allows RaNT-Graph to efficiently sample various types of algorithm-specific walks. Our approach demonstrates improved scalability compared to state-of-the-art distributed system across various walk sampling tasks. Notably, we show RaNT-Graph can sample 100 billion walks in about 15 minutes on 128 compute nodes, highlighting its capability to handle massive walk sampling workloads.

Lance Fletcher
Texas A&M University
lance.fletcher@tamu.edu

MS310

Pathways to Large Graph Foundational Model

Graph-based machine learning has been widely deployed in critical applications, such as molecular property prediction, drug discovery, material design, and cyber threat detection. Though with the successful applications, existing graph-based machine learning approaches face a major challenge of low generability. That is, most graph learning approaches are tailored to train from scratch for a single task on a particular graph, which requires data collection and deployment for each individual graph and task. Inspired by the success of existing foundation models in natural language processing and computer vision, the graph foundation model has been proposed recently with the goal of developing a graph model capable of generalizing across different graphs and tasks. In this talk, I will briefly discuss the pathway towards graph foundational model.

Yuede Ji
University of Texas at Arlington
yuede.ji@unt.edu

MS310

Scalable Motif Counting on Large-Scale Dynamic Graphs

Motifs, small subgraphs of k vertices such as triangles and K-truss, are used to characterize and align different networks, and are important to multiple domains such as biology and social networks. A computationally difficult problem, recent advances have seen success in counting small motifs for large graphs exceeding a billion edges on static networks. However, real world data is often changing and can often be better modeled as dynamic networks. Some individual motifs are well studied for both static and dynamic networks but scalable motif counting on dynamic networks remains lacking. We present a method for counting motifs of k=3,4 for batched fully-dynamic networks. Instead of recomputing motif counts for the changed network from scratch, we update the frequencies of the batched update for some motifs and use combinatorial arguments to compute the rest. Our method can leverage massively parallel architecture, both multi-core CPUs and GPUs, to scale to large networks. We evaluate our method on real-world networks and show that we outperform existing parallel motif counting methods on static networks.

Ali Y. Khan
University of North Texas
alikhan@my.unt.edu

Nigel Tan, Michela Taufer
University of Tennessee Knoxville
ntan1@vols.utk.edu, mtaufer@utk.edu

Sanjukta Bhowmick
University of North Texas
sanjukta.bhowmick@unt.edu

MS310

Efficient Techniques in Analyzing Large Dynamic Networks

Dynamic graphs, characterized by their evolving topologies, require continuous updates to graph properties. Traditional static graph algorithms, which re-compute these properties after each set of topological modifications, often become inefficient in such dynamic environments. In this talk, we first present a generic framework that helps in designing parallel algorithms to update graph properties across large networks subject to various changes. Using this framework, we design parallel algorithms to update properties, including shortest paths, vertex coloring, and strongly connected components. These update algorithms identify affected subgraphs by investigating the topological changes and then update properties only within these ar-

eas. While these update algorithms typically outperform static algorithms, we observe that factors such as batch size, type of changes, and their location within the graph can influence execution time. This observation points to a new research direction: developing adaptive graph algorithms capable of switching between full recomputation using traditional static algorithms and update algorithms, depending on the nature of changes in large dynamic networks.

Arindam Khanda
Missouri University of Science and Technology
arindam.khanda@gmail.com

MS311

Discovery of Constitutive Models Through the Use of a Differentiable Finite Element Solver

The predictive capabilities of numerical simulations in computational solid mechanics are significantly influenced by the accuracy of the phenomenological material models they utilize, and multiple strategies to increase their effectiveness have been studied through the years. Additionally, the discovery and calibration of these models still follows trial-and-error approaches. Recent advances in digital image correlation (DIC) provide high-dimensional spatial deformation data for the study of data-driven numerical simulations in computational solid mechanics. Conventional inverse problem approaches oftentimes first calibrate a chosen constitutive model to a set of spatial –and often uncertain– material parameter fields then computing forward predictions utilizing Bayesian methods to account for uncertainty. The fixed form of the chosen phenomenological constitutive law is oftentimes the overlooked bottleneck of the predictive capability of the data-driven model. In this talk, we propose an inverse problem framework based on differentiable FEM and machine learning-enabled constitutive models that is able to discover constitutive models based on limited data sets with experimental DIC data in mind. The idea is to rapidly discover expressive, interpretable, and physically constrained data-driven constitutive models that capture the rich deformation information embedded in the DIC spatial data.

Jingye Tan
Cornell University
jt894@cornell.edu

Jan Fuhg
The University of Texas at Austin
jan.fuhg@utexas.edu

Nikolaos Bouklas
Cornell University
nb589@cornell.edu

MS311

Efficient and Agile Automatic Differentiation for Finite Element Methods

Sensitivity analysis in high-fidelity simulations is crucial for topology and shape design optimization workflows. Topology optimization requires gradients of spatially varying parameters, such as density fields, while shape optimization can benefit from direct optimization of nodal positions in computational meshes. Both methods necessitate derivative analysis schemes with respect to finite element fields and nodal coordinates, often involving complex manual derivations or inaccurate finite difference methods. This talk introduces a new framework for calculating these derivatives efficiently. We propose parameterization fields for topology optimization and a shape displacement field that adjusts reference coordinates for shape optimization, discretized using finite element basis functions and interpolated at quadrature points. By exploiting the finite element methods structure, we compute sensitivities of the finite element residual via forward-mode automatic differentiation at each quadrature point, enabling rapid development of novel nonlinear optimization problems within HPC-ready code. This methodology is implemented in Serac, an open-source thermomechanical simulation engine built on the MFEM finite element framework, with applications including design optimization of nonlinear responsive liquid crystal elastomer structures and porous electrodes.

Jamie Bramwell
Lawrence Livermore National Laboratory
bramwell1@llnl.gov

MS311

Differentiable Physics for Generalizable Closure Modeling of Separated Flows

The computational study of turbulence in fluids is challenging due to the requirement for resolving all spatial and temporal scales. Recently, efforts have been directed towards developing machine learning (ML) methods to mitigate these challenges by constructing turbulence closures where unresolved quantities are modeled as a function of resolved quantities. However, in several ML-based turbulence closure models, prediction is severely limited when faced with varying flow geometries. This talk will present results from a differentiable programming framework to learn generalizable turbulence closure models. In particular, our differentiable framework includes the training of a graph-neural network (GNN) model for subgrid-scale stress tensor, which is integrated into a finite-element (FEM) solver. For this, the gradients computed by automatic differentiation during the training of GNNs are coupled with the discrete adjoint of the FEM solver. By chaining these two components, the learning of subgrid stresses can be conducted using sparse true flow-field information as targets. Moreover, GNN models for the subgrid stress are deployed for different types of geometries with different physics due to its mesh-invariant nature. Our formulation allows the development of a single GNN-based subgrid closure model that generalizes across different geometries and separation physics. It also can support the idea that generalizable ML closures can be achieved using differentiable physics.

Hojin Kim
Penn State University
hojin.kim@psu.edu

Romit Maulik
Pennsylvania State University
rmm7011@psu.edu

MS312

Dmswarmrt: A Ray Trace Package Developed in the Petsc Particle-In-Cell Framework

Ray tracing has applications in a variety of fields ranging from computational geophysics to the simulation of inertially confined fusion plasma, each requiring their own nuance to how the ray tracing is handled. This can be in terms of geometry, solver selection, etc. I will present

in this talk a general purpose ray tracing framework built within the Portable Extensible Toolkit for Scientific Computing (PETSc) using the Particle in Cell portion of the library with a focus on applications to ICF. This provides a framework with support for parallel unstructured meshes and access to PETSc's suite of scalable, composable solvers with its associated GPU compatibility.

Joseph Pusztay
University at Buffalo
Department of Computer Science and Engineering
josephpu@buffalo.edu

MS312

Embedding Multi-Species Probes in MHD Simulations of the Suns Corona

The dynamics of the Sun's outer atmosphere, or corona, are dominated by interactions between the plasma and the solar magnetic field, such as magnetic reconnection. In order to ground simulations of these phenomena in reality, it is important to synthesize corresponding observational data, both in-situ and remote-sensing. Doing so accurately for the solar atmosphere requires detailed information on dozens of chemical elements occupying hundreds of ionization states. Unfortunately, this can be computationally impractical to include in large-scale magnetohydrodynamic (MHD) simulations, which are a cornerstone of coronal modeling. To circumvent this problem, we describe a method to localize the multi-species component of the analysis to specific regions of interest in the plasma. We begin by performing numerical experiments of reconnection scenarios using a 2.5D resistive MHD model in our code SPRUCE. Then, using Lagrangian tracer particles embedded in these simulations as probes, we perform multi-species thermal-nonequilibrium calculations to predict the local behavior of the individual electron and ion populations, which can then be converted into synthetic in-situ spacecraft observations or synthetic remote-sensing spectra. Ultimately, this approach will help us understand the degree to which particular reconnection scenarios contribute to coronal energization and solar wind formation by improving our analyses of observational data.

Gregory Szypko, Stephen Bradshaw
Rice University
gms12@rice.edu, sb28@rice.edu

Grant Gorman
Sandia National Lab
gmgorman22@gmail.com

MS313

Spatio-Temporal Transport-Based Inference for Imaging Spectroscopy

Spatiotemporal analysis methods often rely on Gaussian assumptions to simplify the modeling of complex fields. However, many natural phenomena are inherently non-Gaussian, which require more general approaches to characterize the underlying distributions. We discuss a Bayesian framework for modeling multivariate spatiotemporal fields involving hundreds of correlated quantities of interest. We adapt recent ideas from measure transport to develop methods for efficient inference that is not limited to Gaussian characterizations. Given the high dimensionality, it is crucial to exploit the underlying problem structure to ensure computational feasibility. We leverage this structure through the identification of conditional independence relations between variables, which is directly related to sparsity in the corresponding transport maps. We apply this approach to imaging spectroscopy retrievals in Earth remote sensing.

Kelvin M. Leung
Massachusetts Institute of Technology
kmleung@mit.edu

MS313

A Structurally Informed Data Assimilation Approach for Discontinuous State Variables

Ensemble-based Kalman filtering data assimilation is a scientific process that combines available observations with numerical simulations to obtain statistically accurate and reliable state representations in dynamical systems. However, it is well known that the commonly used Gaussian distribution assumption introduces biases for state variables that admit discontinuous profiles, which are prevalent in nonlinear partial differential equations. In this talk, we focus on the design of a new structurally informed prior that exploits statistical information from the simulated state variables. In particular, based on the second moment information of the state variable gradient, we construct a new weighting matrix for the numerical simulation contribution in the data assimilation objective function. This replaces the typical prior covariance matrix used for this purpose. We further adapt our weighting matrix to include information in discontinuity regions via a clustering technique. Our numerical experiments demonstrate that this new approach yields more accurate estimates than those obtained using standard ensemble-based Kalman filtering on shallow water equations.

Tongtong Li
University of Maryland, Baltimore County
tongtong.li@umbc.edu

Yoonsang Lee
Dartmouth College
Department of Mathematics
yoonsang.lee@dartmouth.edu

anne Gelb
Dartmouth College
anne.gelb@math.dartmouth.edu

MS313

Preconditioning Techniques for Large-Scale Sparsity-Promoting Inverse Problems

Hybrid projection methods have proven to be a powerful technique for the solution of large-scale linear inverse problems with ℓ_2 regularization, enabling efficient regularization parameter selection via a reduced basis while also mitigating the number of matrix-vector products with the forward model $\mathbf{A}$. However, a major obstacle arises when applying these methods to solving sparsity-promoting inverse problems: the prescribed reduced basis must be very large to appropriately represent the sparsity in the solution, which greatly hinders the computational efficacy of the hybrid projection method. To remedy this, here we introduce a new preconditioning technique that seeks a reduced basis in a transformed space which drastically reduces the number of basis vectors required to appropriately capture the sparsity. We illustrate the efficacy of our technique using several numerical tests and compare our approach to alternatives such as recycling methods based on iterative basis

expansion and compression.

Jonathan Lindbloom
Dartmouth College
jonathan.t.lindbloom.gr@dartmouth.edu

Mirjeta Pasha
Tufts University
mirjeta.pasha@tufts.edu

Jan Glaubitz
Massachusetts Institute of Technology
Glaubitz@mit.edu

MS313

Sparsity Promoting Hierarchical Bayesian Model for Eit

The aim of Electrical Impedance Tomography (EIT) is to determine the electrical conductivity distribution inside a domain by applying currents and measuring voltages on its boundary. Mathematically, the EIT reconstruction task can be formulated as a non-linear inverse problem. The Bayesian inverse problems framework has been applied expensively to solutions of the EIT inverse problem, in particular in the cases when the unknown conductivity is believed to be blocky. In this talk, we demonstrate that by exploiting linear algebraic considerations it is possible to organize the calculation for the Bayesian solution of the nonlinear EIT inverse problem via finite element methods with sparsity promoting priors in a computationally efficient manner. The proposed approach uses the Iterative Alternating Sequential (IAS) algorithm for the solution of the linearized problems. Within the IAS algorithm, a substantial reduction in computational complexity is attained by exploiting the low dimensionality of the data. Numerical tests on synthetic and real data illustrate the computational efficiency of the proposed algorithm.

Monica Pragliola
University of Bologna
Department of Mathematics
monica.pragliola@unina.it

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS314

A Bayesian Network-based Framework for Uncertainty Management in Multimodal Scientific Simulations

Recent advancements in exascale multiscale and multiphysics simulations have significantly improved the modeling and prediction of complex systems, from manufacturing processes to climate change. These simulations integrate multiple modules, each representing distinct physical phenomena, with a range of sub-models varying in fidelity and complexity. Efficient development, rigorous verification, and reliable application require identifying critical modules that impact the accuracy and reliability of quantities of interest (QoIs) and prioritizing their enhancement. This paper presents a systematic framework for uncertainty quantification (UQ) to guide decision-making in large-scale scientific simulations. Using Bayesian networks, the framework models the information flow between modules, aggregating uncertainties across interconnected components. Global sensitivity analyses rank each module's contribution to the overall uncertainty in the target QoI, while Bayesian model selection identifies the most plausible sub-model for integration into subsequent simulation iterations. This iterative process continues until the simulation meets predefined accuracy and reliability standards. The framework is applied to the Ablative Boundary Layers At The Exascale (ABLATE), which simulates turbulent combustion and solid fuel ablation in hybrid rocket motors. Results demonstrate that the framework effectively manages resources, controlling both accuracy and confidence levels in QoI.

Danial Faghihi
University at Buffalo
danialfa@buffalo.edu

Georgios Georgialis, Alejandro Becerra, Abani Patra
Tufts University
georgios.georgalis@tufts.edu, alejandro.becerra@tufts.edu, abani.patra@tufts.edu

MS314

Hybrid Climate Simulations Using Machine Learning

The under-representation of cloud formation is a long-standing bias associated with climate simulations. Parameterisation schemes are required to capture cloud processes within current climate models but have known biases. We overcome these biases by embedding a Multi-Output Gaussian Process (MOGP) trained on high resolution Unified Model simulations to represent the variability of temperature and specific humidity within a climate model. A trained MOGP model is coupled in-situ with a simplified Atmospheric General Circulation Model named SPEEDY. The temperature and specific humidity profiles of SPEEDY are perturbed at fixed intervals according to the variability predicted from the MOGP. Ten-year predictions are generated for both control and ML-hybrid models. The hybrid model reduces the global precipitation bias by 18% and over the tropics by 22%. To further understand the drivers of these improvements, physical quantities of interest are explored, such as the distribution of lifted index values and the alteration of the Hadley cell. The control and hybrid set-ups are also run in a plus 4K sea-surface temperature experiment to explore the effects of the approach on patterns relating to cloud cover and precipitation in a warmed climate setting. We also present preliminary work including towards including wave breaking data-driven models into climate simulations.

Daniel Giles, James Briant
University College London
d.giles@ucl.ac.uk, james.briant.21@ucl.ac.uk

Cyril Morcrette
Met Office
cyril.morcrette@metoffice.gov.uk

Serge Guillas
University College London

s.guillas@ucl.ac.uk

MS315

Mixed Types of Uncertainty Quantification Using Random Fuzzy Sets

In this talk, we deal with mixed types of uncertainty using random variables and random fuzzy sets. Specifically, the stochastic uncertainty in the system is represented with random variables, while various epistemic uncertainties represented using non-probabilistic uncertain variables with random sets and fuzzy sets. The uncertainty in the statistics of output will then be quantified with random fuzzy sets. To reduce the computational cost, generalized polynomial chaos expansion will be adopted to approximate the full simulation. The proposed method will be applied to the biological systems (such as the system of odor capture or valveless pumping system) for mixed types of uncertainty analysis.

Isai Chavarri, Yanyan He, Minyi Chen
University of North Texas
isaichavarri@my.unt.edu, yanyan.he@unt.edu, minyichen@my.unt.edu

John Cain
Harvard University
jcain2@math.harvard.edu

MS315

Energy Dissipation Rate-Based Sampling: An Adaptive Refinement Strategy for Solving Thermodynamically Consistent Models Using Pinns

We introduce Energy Dissipation Rate-based Sampling (EDRS), an innovative adaptive refinement technique for Physics-Informed Neural Networks (PINNs) aimed at significantly enhancing their efficacy in tackling thermodynamically consistent models. The core is the strategic utilization of energy dissipation rate density as a key metric for selectively resampling critical collocation points, thereby refining the model's accuracy. EDRS notably outperforms the traditional residual-based adaptive refinement approach, demonstrating a sixfold improvement in relative mean square error for the Allen-Cahn equation. Leveraging the inherent mesh-free nature of PINNs, we deploy neural networks to adeptly approximate the solutions of thermodynamically consistent phase field models with dynamic boundary conditions in arbitrary domains. Our primary aim is to meticulously investigate how dynamic boundary conditions influence the overall dynamics of the bulk material. We conduct thorough simulations across both disk and ellipse-shaped domains, analyzing how static versus dynamic boundary conditions distinctly affect the system's behavior. This study not only underscores the potential of EDRS in enhancing the computational performance of PINNs but also enriches our comprehension of the pivotal role dynamic boundary conditions play in influencing the dynamic behavior of thermodynamic systems, offering valuable insights for future studies in computational physics and engineering.

Chunyan Li
The Pennsylvania State University
cvl6109@psu.edu

Wenkai Yu
Beijing Computational Science Research Center
yuwenkai@csrc.ac.cn

Qi Wang
University of South Carolina
qwang@math.sc.edu

MS315

Explicit-Solute Implicit-Solvent Molecular Simulation with Binary Level-Set, Adaptive-Mobility, and GPU

Coarse-grained modeling and efficient computer simulations are critical to the study of complex molecular processes with many degrees of freedom and multiple spatiotemporal scales. Variational implicit-solvent model (VISM) for biomolecular solvation is such a modeling framework, and its initial success has been demonstrated consistently. In VISM, an effective free-energy functional of solute-solvent interfaces is minimized, and the surface energy is a key component of the free energy. In this work, we extend VISM to include the solute mechanical interactions, and develop fast algorithms and GPU implementation for the extended variational explicit-solute implicit-solvent (VESIS) molecular simulations to determine the underlying molecular equilibrium conformations. We employ a fast binary level-set method for minimizing the solvation free energy of solute-solvent interfaces and construct an adaptivemobility gradient descent method for solute atomic optimization. We also implement our methods on the integrated GPU. Numerical tests and applications to several molecular systems verify the accuracy, stability, and efficiency of our methods and algorithms. Our fast computational techniques may enable us to simulate very large systems such as protein- protein interactions and membrane dynamics for which explicit-solvent all-atom molecular dynamics simulations can be very expensive.

Shuang Liu
University of North Texas
Shuang.Liu@unt.edu

Zirui Zhang
University of California, Irvine
ziruz16@uci.edu

Hsiao-Bing Cheng
UCSD
hscheng@math.ucsd.edu

Li-Tien Cheng, Bo Li
University of California, San Diego
l3cheng@ucsd.edu, bli@ucsd.edu

MS315

Investigation of a Fuzzy Fractional Diabetes Model Involving GlucoseInsulin Alliance Scheme with a Fuzzy Double Parametric Approach

This study aims to develop and analyze a comprehensive regulatory framework for managing glucose and insulin levels in the presence of diabetes mellitus. An innovative mathematical model of diabetes is explored using fractional calculus, incorporating the ABC fractional derivative. The model is further enhanced by introducing a fuzzy double parametric approach to account for uncertainties in biological parameters. The framework is solved using a semi-analytical technique, specifically the q-Homotopy Analysis Generalized Integral Transform Technique. The effectiveness of this q-Homotopy Analysis Generalized Integral Transform Technique is demonstrated by comparing results

with classical methods such as ADLTM and HPTM. The existence and stability of the solution are established using the Banach Fixed Point Theorem. Numerical simulations and graphical representations are generated using Maple software to illustrate the behaviour of glucose and insulin dynamics in this fuzzy fractional diabetes model for crisp and uncertain cases. The detailed investigation also explores the impact of glucose and insulin levels on the disease's progression, emphasizing the role of the fuzzy double parametric approach in enhancing the model's robustness.

Parthkumar P. Sartanpara
Department of Mathematics,
Sardar Vallabhbhai National Institute of Technology, Surat.
parthsartanpara@gmail.com

Ramakanta Meher
Sardar Vallabhbhai National Institute of Technology
Surat-395007, Gujarat, India
meher_ramakanta@yahoo.com

MS316

Multilevel Projection Algorithms for Handling Heterogenous Differential Privacy

Differential Privacy (DP) has emerged as a cornerstone for safeguarding sensitive information in data analysis and machine learning. However, the heterogeneity of real-world data and varying privacy requirements across clients pose significant challenges to traditional DP mechanisms. This work introduces Multilevel Projection Algorithms as a novel approach to address these challenges in federated learning scenarios. Our proposed algorithm dynamically project noisy gradients and model updates onto refined low-dimensional subspaces, accommodating diverse privacy budgets across heterogeneous clients. By leveraging hierarchical optimization techniques, we achieve an optimal balance between privacy guarantees and model performance. The multilevel projections enable efficient data handling, improve convergence rates, and reduce communication overhead while maintaining rigorous DP guarantees. We demonstrate the initial results through real-world dataset, highlighting its scalability, adaptability, and robustness of the approach.

Olivera Kotevska
Oak Ridge National Laboratory
kotevskao@ornl.gov

MS316

Federated Adaptive Global Pruning under Model Heterogeneity

Federated Learning (FL) has gained significant interest in training large AI models in a distributed computing environment benefiting from its capability to maintain the privacy of sensitive data of the participating parties. However, challenges remain in effectively handling of participating parties with heterogeneous computational power, such as edge devices. In this work, we propose a federated framework, called FedSpaLLM, that involves an adaptive global pruning scheme to enable collaborative training of large models, such as LLMs, on parties with heterogeneous computational power.

Yijiang Li
Argonne National Laboratory
yijiang.li@anl.gov

MS316

Federated Image Reconstruction for Multimodal Data

We propose a federated algorithm for reconstructing images using multimodal tomographic data sourced from dispersed locations, addressing the challenges of traditional unimodal approaches that are prone to noise and reduced image quality. Our approach formulates a joint inverse optimization problem incorporating multimodality constraints and solves it in a federated framework through local gradient computations complemented by lightweight central operations, ensuring data decentralization. Leveraging the connection between our federated algorithm and the quadratic penalty method, we introduce an adaptive step-size rule with guaranteed sublinear convergence and further suggest its extension to augmented Lagrangian framework. Numerical results demonstrate its superior computational efficiency and improved image reconstruction quality.

Minseok Ryu
Arizona State University
minseok.ryu@asu.edu

MS317

Paired Autoencoder for Inverse Problems

Nonlinear inverse problems, often arising from discretized PDEs, are challenging due to the computational cost of data misfit estimation. Likelihood-free approaches offer an alternative to traditional inversion techniques but often suffer from generalization and accuracy issues. We discuss a paired autoencoder framework as a novel likelihood-free estimator. This approach enables efficient solution construction, quality assessment, and iterative improvement. We demonstrate the effectiveness of our method through applications in full waveform inversion and inverse electromagnetic imaging.

Matthias Chung
Emory University, U.S.
matthias.chung@emory.edu

MS317

Variable Spiking Deep Operator For Energy Efficient Scientific Machine Learning

Spiking neural networks and spiking neurons are expected to drive the next generation of AI, as they offer energy-efficient, sustainable alternatives to the current deep learning-based AI algorithms. Their compatibility with neuromorphic hardware also makes them suitable for edge computing and applications where energy resources are limited. Having said this, at this stage, their performance in regression tasks, which are often encountered in computational mechanics and scientific machine learning, is lacking and requires further exploration. This work proposes a variable spiking deep operator that utilizes Variable Spiking Neurons (VSNs) within its architecture to promote energy efficiency. VSNs were developed in response to the challenge faced by vanilla spiking neurons in dealing with regression tasks. They communicate using non-binary spikes while promoting sparse communication. This gives them the advantage of carrying information-rich spikes while keeping the overall computations low. Com-

munication with non-binary spikes is also being explored in neuromorphic chipsets; for example, Intels Loihi 2 chipset supports graded spike communication. To test the developed variable spiking deep operator, we train it on a dataset generated from a time-dependent, two-dimensional partial differential equation. The results produced show that the developed framework gives a good approximation of the ground truth and compares well against the vanilla deep neural operator.

Shailesh Garg
Indian Institute of Technology Delhi
Indian Institute of Technology Delhi
amz218308@am.iitd.ac.in

Souvik Chakraborty
Department of Applied Mechanics
Indian Institute of Technology Delhi
souvik@am.iitd.ac.in

MS317

More of a Good Thing: Combining Multifidelity, Domain Decomposition, and New Architectures for Improved Physics-Informed Training

Physics-informed neural networks and operator networks have shown promise for effectively solving equations modeling physical systems. However, these networks can be difficult or impossible to train accurately for some systems of equations. One way to improve training is through the use of a small amount of data, however, such data is expensive to produce. We will introduce our novel multifidelity framework for stacking physics-informed neural networks and operator networks that facilitates training by progressively reducing the errors in our predictions for when no data is available. In stacking networks, we successively build a chain of networks, where the output at one step can act as a low-fidelity input for training the next step, gradually increasing the expressivity of the learned model. We will finally discuss the extension to domain decomposition using the finite basis method, including applications to newly-developed Kolmogorov-Arnold Networks.

Amanda Howard, Panos Stinis
Pacific Northwest National Laboratory
amanda.howard@pnnl.gov, panos.stinis@pnnl.gov

MS317

Geometric Neural-Operators for Manifold Point-Cloud Representations: PDE Solvers and Meshless Methods

Geometric Neural Operators (GNPs) are introduced for data-driven deep learning of operators for tasks in non-euclidean settings. The approaches allow for handling manifolds of general shape, including with point-cloud representations. We show how GNPs can be used (i) to estimate geometric properties, such as the metric and curvatures, (ii) to develop solvers for Partial Differential Equations (PDEs) on manifolds, and (iii) to solve Bayesian inverse problems for identifying manifold shapes. Further, we show how GNPs can learn representations for local patches of manifolds, which can be useful in meshless methods or other downstream tasks.

Blaine Quackenbush, Paul J. Atzberger
University of California-Santa Barbara
bquackenbush@umail.ucsb.edu, atzberg@ucsb.edu

MS317

Spectral Localization and Lipschitz Continuity of Neural Networks

A spectral localization technique is proposed to enforce the Lipschitz continuity of neural networks with learnable Lipschitz constants. The method is used to construct locally attractive stable manifolds of residual flows and is validated on various tasks including reinforcement learning, density estimation and certified robustness.

Viktor Reshniak
Oak Ridge National Laboratory
reshniakv@ornl.gov

MS318

Iterative Methods for T-Product Linear Systems

In applications involving inverse problems, large-scale data is a common challenge. In this presentation, we introduce an iterative method for approximating the solution of large-scale multi-linear systems, represented in the form A*X=B under the tensor t-product. Unlike previously proposed randomized iterative strategies, such as the tensor randomized Kaczmarz method (row slice sketching) or the tensor Gauss-Seidel method (column slice sketching), which are natural extensions of their matrix counterparts, our approach delves into a distinct scenario utilizing frontal slice sketching. In particular, we explore a context where frontal slices, such as video frames, arrive sequentially over time, and access to only one frontal slice at any given moment is available. This talk will present our novel approach, shedding light on its applicability and potential benefits in approximating solutions to large-scale multi-linear systems. This is joint work with Hengrui Luo (Rice University).

Anna Ma
UC Irvine
anna.ma@uci.edu

Hengrui Luo
Rice University
hrluo@rice.edu

MS318

Algorithms and Software for Large Scale Tensor Completion

Tensor completion is a technique used to recover missing or unobserved entries in a tensor by leveraging the structure and correlations within the data. It is widely applied in fields like data science, computer vision, and signal processing, where high-dimensional data often contains incomplete or sparse information. The correlations in data can be represented in the form of various tensor decompositions such as CP, Tucker, Tensor train etc. Also, depending on the noise and data distribution, a suitable objective function is minimized. Numerical optimization algorithms used in the literature to optimize for the tensor decompositions are mainly based on first order derivative information due to the nonlinear structure of the problem. We formulate algorithms based on second derivatives of the objective functions, and develop software infrastructure for the same to show that in many cases these algorithms outperform the gradient-based algorithms in terms of accuracy of predic-

tion while maintaining scalability.

Navjot Singh
University of Illinois Urbana-Champaign
navjot2@illinois.edu

Edgar Solomonik
University of Illinois at Urbana-Champaign
solomon2@illinois.edu

MS319

Nonlinear Balanced Truncation Model Reduction in High Dimensions

Nonlinear balanced truncation is a model reduction technique that reduces the dimension of nonlinear systems on manifolds and preserves either open- or closed-loop observability and controllability aspects of the nonlinear system. Three computational challenges so far prevented its deployment on large-scale systems: (a) the computation of Hamilton-Jacobi-(Bellman) equations that are needed to characterize controllability and observability, (b) a scalable strategy to compute the nonlinear coordinate transformation to balance the system, and (c) efficient reduced-order model (ROM) construction on the resulting nonlinear balanced manifolds. We present a novel unifying and scalable computational approach to balanced truncation for medium/large-scale control-affine nonlinear systems. The Taylor-series based approach solves a class of parametrized Hamilton-Jacobi-Bellman equations, and subsequently computes the polynomial manifold that achieves output diagonalization of the energy functions. This specific tensor structure for the coefficients of the Taylor series (tensors themselves) allows for scalability up to thousands of states. We also present initial results on reduced-order modeling. The talk will illustrate the strength and scalability of each of the three steps of the procedure on several semi-discretized nonlinear partial differential equations.

Boris Kramer, Nicholas Corbin
University of California San Diego
bmkramer@ucsd.edu, ncorbin@ucsd.edu

MS319

Numerical Investigation of the Fractal-fractional Order Model for Diabetes Mellitus Considering of Media-driven Awareness Program

Diabetes is rapidly emerging as a global epidemic, posing a significant threat to public health. Modeling the spread and management of diabetes is crucial for monitoring its growing prevalence and developing cost-effective strategies to mitigate its incidence and complications. This paper presents a fractal-fractional order nonlinear model for diabetes mellitus that incorporates the cumulative effect of media-driven diabetes awareness and education programs. The model is analyzed using a two-step Newtonian polynomial approach with the Caputo derivative. Key aspects such as equilibrium points, stability, and the existence and uniqueness of solutions are examined to ensure the robustness of the model. The obtained results are validated against previously published findings. Graphical and numerical results are obtained for different values of the fractional order. This research provides valuable insights for predicting disease trends and planning effective clinical management for diabetes patients.

Vishalkumar J. Prajapati
Sardar Vallabhbhai National Institute of Technology, Surat
prajapativishalkumar15@gmail.com

Lalchand Verma
Panipat Institute of Engineering & Technology
Panipat-132102, Haryana, India
lalchandverma81@gmail.com

Ramakanta Meher
Sardar Vallabhbhai National Institute of Technology
Surat-395007, Gujarat, India
meher_ramakanta@yahoo.com

MS319

Greedy Construction of Quadratic Manifolds for Nonlinear Dimensionality Reduction and Nonlinear Model Reduction

Dimensionality reduction with quadratic manifolds augments linear approximations in subspaces with quadratic correction terms. While previous works rely on linear approximations given by projections onto the first few leading principal components of the training data, we instead construct subspaces so that the corresponding linear approximations can be corrected most efficiently with quadratic terms. We present a greedy method for the subspace construction that selects basis vectors from leading as well as later principal components. The greedy selection allows us to determine a basis that can leverage the quadratic corrections most efficiently. This is in contrast to choosing as basis the leading principle components, which results in the best linear approximation but is not necessarily most informative for the quadratic correction terms. Properties of the greedily constructed manifold allow applying linear algebra reformulations so that the greedy method scales to data points with millions of dimensions. Numerical experiments demonstrate that an orders of magnitude higher accuracy is achieved with the greedily constructed quadratic manifolds compared to manifolds that are based on the leading principal components.

Paul Schwerdtner
New York University
ps5136@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS319

Shallow Recurrent Decoders for Reduced Order Models

Data-driven models are critically enabling in many application areas where the underlying dynamics are unknown or only partially known, or where high-fidelity simulations are computationally expensive to generate. The ability to produce accurate, low-rank, proxy models enable dynamic models to transform the representation and characterization of such systems. We develop a model reduction scheme based upon a Shallow REcurrent Decoder (SHRED) architecture. The scheme uses a neural network for encoding limited sensor measurements in time (sequence-to-sequence encoding) to full state-space reconstructions via a decoder network. Based upon the theory of separation of variables, the SHRED architecture is capable of (i) reconstructing full

spatio-temporal fields with as little as three point sensors, and (ii) forecasting the future state of the system using neural network roll-outs from the trained time encoding model. The SHRED reduced order model architecture is demonstrated on a number of spatio-temporal dynamics.

Jan P. Williams
Department of Mechanical Engineering
University of Washington
jmpw1@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS320

Uncovering Temporal Latent Structure Using Alternators

This talk discusses alternators, a novel family of non-Markovian dynamical models for sequences. An alternator features two neural networks: the observation trajectory network (OTN) and the feature trajectory network (FTN). The OTN and the FTN work in conjunction, alternating between outputting samples in the observation space and some feature space, respectively, over a cycle. The parameters of the OTN and the FTN are not time-dependent and are learned via a minimum cross-entropy criterion over the trajectories. Alternators are versatile. They can be used as dynamical latent-variable generative models or as sequence-to-sequence predictors. When alternators are used as generative models, the FTN produces interpretable low-dimensional latent variables that capture the dynamics governing the observations. When alternators are used as sequence-to-sequence predictors, the FTN learns to predict the observed features. In both cases, the OTN learns to produce sequences that match the data. Alternators can uncover the latent dynamics underlying complex sequential data, accurately forecast and impute missing data, and sample new trajectories. We showcase them in various applications, including neuroscience, chaos modeling, and climate science.

Adji Bousso Dieng
Princeton University
adji@princeton.edu

MS320

Diffusion Generative Modeling in Non-Euclidean Spaces

Machine learning in non-Euclidean spaces have been rapidly attracting attention in recent years, and this talk will give an example by describing a sequence of developments that eventually leads to generative modeling of data in Lie groups, which are useful for molecule/material synthesis, quantum sciences, robotics, etc. More precisely, I will begin with dynamics that performs optimization on Lie group, and turn it into an algorithm that samples from probability distributions on Lie groups. The performance of this sampler will also be quantified, without log-concavity condition or its common relaxations. Then I will describe how this sampler can lead to a structurally-pleasant diffusion generative model that allows users to, given training data that follow any latent statistical distribution on a Lie group, generate more data exactly on the same manifold that follow the same distribution. If time permits, some of the aforementioned applications will be demostrated.

Molei Tao
School of Mathematics, Georgia Institute of Technology
mtao@gatech.edu

MS320

Designing for Equivariance to Perceptual Variation in Images

Group equivariant convolutional neural networks have been designed for a variety of geometric transformations from 2D and 3D rotation groups, to semi-groups such as scale. Despite the improved interpretability, accuracy and generalizability afforded by these architectures, group equivariant networks have seen limited application in the context of perceptual quantities such as hue and saturation, even though their variation can lead to significant reductions in classification performance. In this paper, we introduce convolutional neural networks equivariant to variations in hue and saturation by design. To achieve this, we leverage the observation that hue and saturation transformations can be identified with the 2D rotation and 1D translation groups respectively. Our hue-, saturation-, and fully color-equivariant networks achieve equivariance to these perceptual transformations without an increase in network parameters. We demonstrate the utility of our networks on synthetic and real world datasets where color and lighting variations are commonplace.

Felix O'Mahony
University of Oxford
felixomahony@gmail.com

Yulong (Arthur) Yang, Christine Allen-Blanchette
Princeton University
yy0555@princeton.edu, ca15@princeton.edu

MS321

Finite Element Spaces for Double Forms

In 2018, Li developed finite element spaces for symmetric bilinear forms with tangential-tangential continuity, with applications to elasticity and numerical relativity. Although Li constructed these spaces for arbitrary order, his construction fundamentally relied on the existence of zeroth order (piecewise constant) spaces; these are the classical spaces of Regge calculus (1961). The existence of zeroth order spaces is somewhat surprising: for instance, if we consider antisymmetric bilinear forms instead of symmetric ones, then zeroth order finite element spaces do not exist in any dimension greater than two. Viewing bilinear forms as 1-forms tensor 1-forms, we generalize by considering k-forms tensor l-forms, also known as double forms. In the same way that bilinear forms have natural subspaces of symmetric bilinear forms and antisymmetric bilinear forms, double forms have natural subspaces as well. In any dimension, we determine which of these subspaces admit zeroth order finite element spaces, and we construct them if so. As a result, we obtain a unified framework that, among other things, encompasses the symmetric matrices with tangential-tangential continuity of Regge, the symmetric matrices with normal-normal continuity of Sinwel (2009), and the trace-free matrices with normal-tangential continuity of Gopalakrishnan, Lederer, and Schberl (2019).

Yakov Berchenko-Kogan
Florida Institute of Technology
yberchenkokogan@fit.edu

Evan S. Gawlik
Santa Clara University
egawlik@scu.edu

MS321

Canonical Finite Elements for Form-Valued Differential Forms

We provide a finite element discretization of ℓ-form-valued k-forms on triangulation in $\mathbb{R}^n$ for general k, ℓ and n. The construction generalizes finite element Whitney forms and their higher order versions for the de Rham complex, the Regge finite elements for the elasticity complex, and the HHJ elements for solving the biharmonic equations. Applications of the construction include discretization of strain and stress tensors in continuum mechanics and metric and curvature tensors in differential geometry in any dimension.

Kaibo Hu
University of Edinburgh
kaibo.hu@ed.ac.uk

Ting Lin
Peking University
lintingsms@pku.edu.cn

Qian Zhang
Michigan Technological University
qzhang15@mtu.edu

MS321

Bounded Commuting Projections in Feec Using Discrete Local Problems on Alfeld Splits

We construct projections onto the classical finite element spaces based on the Lagrange, Nédélec, Raviart-Thomas, and discontinuous elements on shape-regular simplicial meshes. Our projections are defined locally, are bounded in the L^2 norm, and commute with the corresponding differential operators. The cornerstone of the construction are local weight functions which are discrete (piecewise polynomials) built on mesh stars using the Alfeld split of the original simplicial mesh. We will discuss the construction of these projections and how the noted properties are satisfied. As an example, the L^2-stability of the projections is established by invoking discrete Poincaré inequalities on stars of simplices, for which we provide constructive proofs. Finally, we will discuss the modification of the construction to preserve homogeneous boundary conditions.

Pratyush Potu
Brown University
pratyush_potu@brown.edu

Alexandre Ern
Universite Paris-Est
CERMICS, Ecole des Ponts
ern@cermics.enpc.fr

Johnny Guzman
Brown University
johnny_guzman@brown.edu

Martin Vohlralik
INRIA
martin.vohralik@inria.fr

MS321

Finite Element Exterior Calculus for Hamiltonian PDEs

We consider the application of finite element exterior calculus (FEEC) methods to a class of canonical Hamiltonian PDE systems involving differential forms. Solutions to these systems satisfy a local *multisymplectic conservation law*, which generalizes the more familiar symplectic conservation law for Hamiltonian systems of ODEs, and which is connected with physically-important reciprocity phenomena, such as Lorentz reciprocity in electromagnetics. We characterize hybrid FEEC methods whose numerical traces satisfy a version of the multisymplectic conservation law, and we apply this characterization to several specific classes of FEEC methods, including conforming ArnoldFalkWinther-type methods and various hybridizable discontinuous Galerkin (HDG) methods. Interestingly, the HDG-type and other nonconforming methods are shown, in general, to be multisymplectic in a stronger sense than the conforming FEEC methods. This substantially generalizes previous work of McLachlan and Stern [Found. Comput. Math., 20 (2020), pp. 3569] on the more restricted class of canonical Hamiltonian PDEs in the de DonderWeyl graddiv form.

Ari Stern
Washington University in St. Louis
stern@wustl.edu

Enrico Zampa
University of Trento
enrico.zampa@unitn.it

MS322

Comparative Analysis of Neural Networks and Kernel Methods for Estimating Statistical Divergences

Statistical divergences measure the dissimilarity between probability distributions and are fundamental in Bayesian optimal experimental design. Estimating these divergences accurately from finite samples is challenging, especially in high-dimensional spaces. Recent advances have addressed these challenges by leveraging variational bounds, where the estimation involves optimization over a functional space typically implemented using either kernel methods or neural networks. While considerable progress has been made on the study of the convergence of these estimators, key questions remain unresolved, including the influence of the chosen functional space and the optimization algorithm. In this work, we conduct a comparative analysis of neural networks and kernel methods for estimating statistical divergences, with an emphasis on the convergence of the underlying optimization algorithms.In particular, we compare kernel methods to shallow neural networks and discuss the convergence of gradient descentfortraining two-layer neural networks in the mean-field limit.

Ayoub Belhadji
Massachusetts Institute of Technology
abelhadj@mit.edu

Julie Zhu
MIT
qianyu_z@mit.edu

Youssef M. Marzouk, Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu, ymarz@mit.edu

MS322

Optimal-Stopping Control for Bayesian Sequential Experimental Design

In sequential optimal experiment design (sOED), the number of experiments being planned is generally known and fixed. When the experiment campaign is allowed to terminate early, an important question arises: When should one stop doing more experiments? Simple stopping strategies based on reaching a threshold of accumulated information or reward is generally not optimal. We thus seek the *optimal stopping policy* for sOED. We formulate sOED in a Bayesian manner, where the reward (i.e., utility) terms are based on the information gain of model parameters or model predictions. We show that, in a risk-neutral setting, optimal stopping takes place when the terminal reward exceeds the expected future reward sum. We propose two computational approaches to solve for the optimal stopping policy: approximate dynamic programming and policy gradient. Both methods are demonstrated in a linear-Guassian benchmark where the posteriors can be obtained analytically, and in a sensor movement problem for detecting contaminant source.

Chen Cheng, Xun Huan
University of Michigan
chech@umich.edu, xhuan@umich.edu

MS322

Backwards Sequential Monte Carlo for Bayesian Optimal Experimental Design

The expected information gain (EIG) is a crucial quantity in Bayesian optimal experimental design (OED). However, evaluating the EIG can be computationally expensive since it requires the posterior normalizing constant, leading to a doubly-intractable integral. A rich literature exists for estimation of this normalizing constant, with sequential Monte Carlo (SMC) approaches being one of the gold standards. In this work, we leverage the idiosyncrasies of OED to develop a novel EIG-specific SMC estimator that starts with a sample from the posterior and tempers backwards towards the prior. The method arises from the observation that, in many cases, the Monte Carlo variance of standard SMC estimators for the normalizing constant of a single dataset are significantly lower than the variance of the normalizing constants across datasets; the latter thus contributes the majority of the variance for EIG estimates. This suggests the potential to slightly increase variance while drastically decreasing computation time by reducing the SMC population size and, taking this idea to the extreme, opens the door to unique estimators. We demonstrate our method on a coupled spring-mass system where we observe significant performance improvements.

Andrew Chin
Johns Hopkins University
achin23@jhu.edu

Tommie Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS322

Bayesian Optimal Design Accelerates Discovery of Material Properties from Bubble Dynamics

An optimal sequential experimental design approach is developed to characterize soft material properties at the high strain rates associated with bubble cavitation. The approach involves optimal design and model inference. The optimal design strategy maximizes the expected information gain in a Bayesian statistical setting to design experiments that provide the most informative cavitation data about unknown material properties. We infer constitutive models by characterizing the associated viscoelastic properties from measurements via a hybrid ensemble-based 4D-Var method (En4D-Var). The inertial microcavitation-based rheometry (IMR) method (Estrada et al. 2018, JMPS) is used to simulate bubble dynamics under high strain-rate laser-induced cavitation (LIC) and low-rate needle-induced cavitation (NIC). A library of constitutive models demonstrates the method. We generate synthetic experimental data representing the viscoelastic behavior of stiff and soft polyacrylamide hydrogels and collect measurements from LIC and NIC experiments. Accurate and efficient characterizations of the underlying models are presented.

Tianyi Chu
Georgia Institute of Technology
tchu72@gatech.edu

Joseph Beckett, Bachir Abeid, Jonathan Estrada
University of Michigan
jbecktt@umich.edu, babeid@umich.edu,
jbestrad@umich.edu

Spencer H. Bryngelson
Georgia Institute of Technology
shb@gatech.edu

MS322

Fast and Flexible Estimation of Mutual Information

We consider methods for estimating mutual information (MI) and differential entropy via variational approximations, addressing significant challenges in uncertainty quantification and MI computation. We present a novel approach to variational MI approximations by utilizing moment matching operations, thereby replacing the need for costly nonconvex optimization. This approach is applicable to implicit models that lack closed-form likelihood functions, providing substantial computational speedups. We further extend our variational MI estimators by incorporating Normalizing Flows, enhancing the flexibility of the variational distribution beyond the commonly used Gaussian assumptions. These new flow-based estimators are validated on large MI problems and diverse benchmarking tests, often outperforming traditional critic-based estimators. Their effectiveness is also demonstrated in Bayesian Optimal Experimental Design for online sequential decision making.

Jason Pacheco, Caleb Dahlke
University of Arizona

pachecoj@cs.arizona.edu, calebdahlke@arizona.edu

MS323

A Programming Model for Embedding ML Surrogates in Scientific Application

Recent advancements in Machine Learning (ML) have substantially improved its predictive and computational abilities, offering promising opportunities for surrogate modeling in scientific applications. By accurately approximating complex functions with low computational cost, ML-based surrogates can accelerate scientific applications by replacing computationally intensive components with faster model inference. However, integrating ML models into these applications remains a significant challenge, hindering the widespread adoption of ML surrogates as an approximation technique in modern scientific computing. We propose an easy-to-use directive-based programming model that enables developers to seamlessly describe the use of ML models in scientific applications. The runtime support, as instructed by the programming model, performs data assimilation using the original algorithm and can replace the algorithm with model inference. Our evaluation across five benchmarks, testing over 5000 ML models, shows up to 83.6x speed improvements with minimal accuracy loss (as low as 0.01 RMSE).

Zane Fink
University of Illinois
zanef2@illinois.edu

MS323

Efficient Digital Twin Training using Uncertainty-Guided Data Generation

Machine learning proxy models offer a powerful approach to accelerating and even replacing computationally expensive models. However, constructing these digital twins presents a unique challenge in efficiently generating training data. A naive uniform sampling of the input space can lead to a non-uniform sampling of the output space, resulting in gaps in training data coverage and potentially compromising accuracy. While massive datasets could eventually fill these gaps, the computational burden of full-scale simulations can make this impractical. In this talk, we introduce a framework for adaptive data generation that leverages uncertainty estimation to identify regions requiring additional training data and re-triggering simulations to fill the identified gaps. Essentially, this approach steers large-scale simulations towards generating the necessary data for training digital twins iteratively and thus reduces the data needed to train accurate digital twins. We will demonstrate the challenges at training at scale and the effectiveness of such methods on both a simple one-dimensional function and a complex multidimensional physics model.

Ana Gainaru
Oakridge
gainarua@ornl.gov

MS323

Small Dense Models of Large Sparse Graph Partitioning Problems

Graph Partitioning is a critical problem in numerous scientific and engineering domains including social network analysis, VLSI design and many more. Spectral methods are known to produce quality partitions while minimizing edge cuts for a wide range of problems. However, the computational cost associated with the calculation of the fiedler vector, an eigenvector associated with the second smallest eigenvalue of the graph Laplacian, remains a significant bottleneck. In this paper, we present an accelerated approach to spectral bisection partitioning by replacing the traditional eigenvalue calculation with a simple artificial neural network model to approximate the fiedler vector. We demonstrate that our approach achieves partitioning quality comparable to spectral bisection while significantly reducing the computational overhead, making it more scalable and efficient for large-scale problems.

Vishvam Patel
University of Alabama in Huntsville
vhp0001@uah.edu

MS324

On Quadratic Knapsack Limiting and Enforcing a Cell Entropy Inequality

Entropy stable nodal discontinuous Galerkin methods satisfy a cell entropy inequality. One strategy to enforce such an inequality is through knapsack limiting, which solves a knapsack problem to determine an optimal set of limiting coefficients which result in a semi-discrete cell entropy inequality while preserving nodal bounds. In this work, we provide a slight modification of this approach, where we utilize a quadratic knapsack problem instead of a standard linear knapsack problem. We prove that this quadratic knapsack problem can be reduced to efficient scalar root-finding. Numerical results demonstrate that the proposed quadratic knapsack limiting strategy is efficient, results in better time accuracy than the linear knapsack limiting approach, and behaves better under adaptive time-stepping.

Brian Christner, Jesse Chan
Rice University
bc71@rice.edu, jesse.chan@rice.edu

MS324

High Performance Computing Aspects of Invariant Domain Preserving Schemes

My talk will present performance analysis results for a class of invariant-domain preserving schemes with complex limiting for the compressible Euler equations. Guided by roofline performance models, bottlenecks in memory access, instruction latency and arithmetic work load are identified and improved by reformulations of the underlying numerical model. My talk will compare the achieved performance of a second order finite-element scheme against higher-order discontinuous Galerkin methods with and without limiters, highlighting the price to be paid for mathematically provable robustness.

Martin Kronbichler
Ruhr University Bochum
martin.kronbichler@rub.de

MS324

A Structure-Preserving High-Order Matrix-Free Finite Element Method

We present a three-dimensional matrix-free finite element method (MF-FEM) which provides an explicit and arbitrary high order approximation of the smooth solutions of

the advection-diffusion partial differential equations both in space and time. The scheme allows for an efficient diagonalization of the mass matrix without any loss of accuracy. This is achieved by coupling the MF-FEM formulation with a Deferred Correction (DeC) type method for the discretization in time. We will discuss the structure-preserving (SP) properties of our scheme, such as conservation, positivity, and local bounds preservation. SP is achieved by combining a low-order scheme with provable SP properties, and a high-order, possibly SP violating scheme, in a convex combination with positive weights. We express both low-order and high-order schemes in terms of element-based distribution coefficients and perform the convex limiting in the element-wise residual distribution form. The resulting scheme is verified using a series of challenging benchmark problems for the Euler equations of gas dynamics. Next, we will present a high-order MF-FEM scheme for the shallow water equations and show that our scheme is well-balanced with a proper choice of numerical quadratures and stabilization operators.

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

MS324

Realizability-Preserving Discontinuous Galerkin Method for Spectral Two-Moment Radiation Transport in Special Relativity

We present a realizability-preserving numerical method for solving a spectral two-moment model to simulate the transport of massless, neutral particles interacting with a background material moving with relativistic velocities. The model is obtained as the special relativistic limit of a four-momentum-conservative general relativistic two-moment model. Using a maximum-entropy closure, we solve for the Eulerian-frame energy and momentum. The proposed numerical method, with discontinuous Galerkin (DG) spatial discretization and implicit-explicit (IMEX) time stepping, is designed to preserve moment realizability, which corresponds to moments defined by a nonnegative phase-space density. With a series of numerical tests, we demonstrate the accuracy and robustness of this DG-IMEX method.

Yulong Xing
The Ohio State University
xing.205@osu.edu

MS325

Fast Multigrid Solvers Using Sparsity-promoting Bases for the High-order Simplicial de Rham Complex

We present a scalable multigrid solver for the Riesz maps of the L2 de Rham complex at high-order on triangular and tetrahedral meshes. The multigrid relaxation combines static condensation and parallel space decomposition methods in order to avoid the elevated memory and computational costs by exploiting local orthogonal bases. We introduce new finite elements discretizing the usual H(d)-conforming (d = grad, curl, or div) polynomial spaces, but with different basis functions. The new elements build upon the degrees of freedom from (Demkowicz et al. 2000), and consist of integral moments on a symmetric reference simplex against a numerically computed polynomial basis that is orthogonal in both the L2- and H(d)-inner products. On the reference simplex, the resulting stiffness matrix has diagonal interior block, and fully decouples the interior degrees of freedom. Thus, the Schur complement resulting from elimination of interior degrees of freedom is simply the interface block itself. This sparsity is not preserved on arbitrary simplices. Thus, we devise a preconditioning strategy that neglects the coupling of the interior degrees of freedom. We precondition the Schur complement with the interface block, and apply point-Jacobi for the interior block. The interface block is further preconditioned with a space decomposition method on small subdomains constructed around vertices and edges. We show iteration counts that are robust with respect to the polynomial degree.

Pablo Brubeck
University of Oxford
brubeckmarti@maths.ox.ac.uk

Patrick E. Farrell
Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

Robert C. Kirby
Baylor University
robert_kirby@baylor.edu

MS325

Performance-Portable Multigrid for Adaptive FEM for the Stokes Problem

Performance-Portable Multigrid for Adaptive FEM for the Stokes Problem We present recent advances in the Finite Element library deal.II to use performance portable linear solvers for the efficient solution of finite element problems on CPUs and GPUs. Here, I am presenting the current state of different options to solve variable viscosity Stokes problems that appear in Mantle Convection problems. We consider various options from matrix-free geometric multigrid directly implemented in deal.II using Kokkos, algebraic multigrid solvers through PETSc+Kokkos, and CPU based approaches.

Timo Heister
Clemson University
Mathematical Sciences
heister@clemson.edu

MS325

Advanced Finite Element Discretizations and Solvers in the MFEM Library

Efficient exploitation of exascale architectures requires rethinking of the numerical algorithms used in PDE-based simulations to expose fine-grain parallelism and maximize arithmetic intensity. In this talk we present an overview of MFEM (https://mfem.org), a scalable library for high-order finite element discretization of PDEs on general unstructured grids, which powers HPC applications in a wide variety of fields. We review some the recent research and development activities in the project, including high-order mesh adaptivity, differentiable simulations, matrix-free preconditioning, scalability on GPU-accelerated exascale supercomputers, and demonstrate their impact in several large-scale applications from the US Department of Energy.

Tzanio Kolev
Lawrence Livermore National Laboratory

kolev1@llnl.gov

MS325

A Space-Time Multigrid Method for Space-Time Finite-Element Discretizations

Space-time FEM, stage-parallel implicit Runge Kutta-methods, and, in general, parallel-in-time methods offer a promising solution to the limited scalability of algorithms on distributed systems when parallelism is only exploited in the spatial domain. These algorithms cannot run faster than a certain threshold even if more hardware resources are added. Identifying additional parallelism might be a solution. [1] presented the implementation of a stage-parallel preconditioner for implicit Runge-Kutta methods using multigrid and demonstrated that the scaling limit can indeed be shifted. This talk expands on the concept of time parallelism by introducing a space-time multigrid method based on tensor-product space-time finite-element discretizations. The method is facilitated by the matrix-free capabilities of deal.II. It addresses both high-order continuous and discontinuous variational time discretizations. We demonstrate its effectiveness for the heat and acoustic wave equations [2]. Furthermore, we present our progress in the extension of our solver to Stokes and Navier-Stokes problems and discuss implementational analogies to stage-parallel implicit Runge-Kutta methods [1]. [1] Munch, P., et al., 2024. Stage-parallel fully implicit RungeKutta implementations with optimal multilevel preconditioners at the scaling limit. [2] Margenberg, N. and Munch, P, 2024. A space-time multigrid method for space-time finite element discretizations of parabolic and hyperbolic PDEs

Peter Munch
Uppsala University
peterrmuench@gmail.com

Nils Margenberg
Helmut Schmidt Universität Hamburg
margenbn@hsu-hh.de

MT1

How to Give Good Talks

The workshop aims to enhance the participants' presentation skills. Participants will have the opportunity to reflect on what makes a good talk and prepare their own short talk, applying the tips and tricks relayed in the minitutorial. We will discuss the various types of talks researchers may give and how they can create great slides. We will also present different presentation styles for communicating research to the public, such as TEDx talks, 3-minute talks (e.g. FameLab) and elevator pitches. Participants should bring their laptops with them.

Katerina Kaouri
School of Mathematics
Cardiff University, UK
kaourik@cardiff.ac.uk

MT2

How to Give Good Talks

The workshop aims to enhance the participants' presentation skills. Participants will have the opportunity to reflect on what makes a good talk and prepare their own short talk, applying the tips and tricks relayed in the minitutorial. We will discuss the various types of talks researchers may give and how they can create great slides. We will also present different presentation styles for communicating research to the public, such as TEDx talks, 3-minute talks (e.g. FameLab) and elevator pitches. Participants should bring their laptops with them.

Katerina Kaouri
School of Mathematics
Cardiff University, UK
kaourik@cardiff.ac.uk

MT3

Understanding Generative AI: the Core Concepts

The minitutorial "Understanding Generative AI" offers a comprehensive yet concise exploration of generative models, catering to participants eager to grasp both theoretical and practical aspects within a 90-minute session. The tutorial begins with an introduction to generative AI, establishing its definition, historical context, and distinguishing characteristics from discriminative models. Core generative models such as Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and Transformer models are then explored, with a focus on their mathematical foundations and operational principles. Participants will gain insight into the diverse applications of generative AI across industries like healthcare, finance, and creative arts, illustrated through impactful case studies. The tutorial also addresses the ethical implications and societal impacts of generative AI, offering guidelines for responsible usage to ensure technology's positive influence. A live coding session forms the practical component, where attendees will observe the step-by-step implementation and deployment of a simple generative model, enhancing their understanding through hands-on demonstration. The session concludes with a summary of key takeaways and an interactive QA, ensuring participants leave with a solid understanding of generative AI, practical knowledge for implementation, and a critical awareness of ethical considerations. This tutorial is designed to provide a well-rounded introduction to generative AI, suitable for professionals and enthusiasts alike.

Xi Chen
UCSC
billchenxi@gmail.com

MT4

Fast Direct Solvers for Elliptic PDEs

The tutorial will describe a class of methods known as "fast direct solvers" that have attracted much attention in the linear algebra and numerical PDE communities in the past two decades. These algorithms address the problem of solving a linear equation Ax=b arising from the discretization of either an elliptic PDE or of an associated integral equation. The matrix A will be sparse when the PDE is discretized directly, and dense when an integral equation formulation is used. In either case, industry practice for large scale problems has been to use an iterative solver such as, e.g., multigrid or GMRES. A direct solver, in contrast, builds an approximation to the inverse of A (or alternatively, an easily invertible factorization such as LU or Cholesky). A major development in the last couple of decades has been the emergence of algorithms for performing this inversion in linear or close to linear time. Such methods must necessarily exploit that the inverse of the matrix A is "data sparse", typically in the sense that it can be tessellated into blocks that have low numerical

rank. The tutorial will be accessible to graduate students and postdocs with a foundation in numerical linear algebra, and the basics of methods for solving PDEs numerically. Participants will be given the opportunity to directly test out the concepts introduced through tutorial codes designed to support the presentation.

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

Adrianna Gillman
University of Colorado at Boulder
Department of Applied Mathematics
adrianna.gillman@colorado.edu

MT5

Fast Direct Solvers for Elliptic PDES Part II of II

The tutorial will describe a class of methods known as "fast direct solvers" that have attracted much attention in the linear algebra and numerical PDE communities in the past two decades. These algorithms address the problem of solving a linear equation Ax=b arising from the discretization of either an elliptic PDE or of an associated integral equation. The matrix A will be sparse when the PDE is discretized directly, and dense when an integral equation formulation is used. In either case, industry practice for large scale problems has been to use an iterative solver such as, e.g., multigrid or GMRES. A direct solver, in contrast, builds an approximation to the inverse of A (or alternatively, an easily invertible factorization such as LU or Cholesky). A major development in the last couple of decades has been the emergence of algorithms for performing this inversion in linear or close to linear time. Such methods must necessarily exploit that the inverse of the matrix A is "data sparse", typically in the sense that it can be tessellated into blocks that have low numerical rank. The tutorial will be accessible to graduate students and postdocs with a foundation in numerical linear algebra, and the basics of methods for solving PDEs numerically. Participants will be given the opportunity to directly test out the concepts introduced through tutorial codes designed to support the presentation.

Per-Gunnar Martinsson
University of Texas at Austin, U.S.
pgm@oden.utexas.edu

Adrianna Gillman
University of Colorado at Boulder
Department of Applied Mathematics
adrianna.gillman@colorado.edu

MT6

Data-Driven Reduced Modeling in the Time and Frequency Domains: Fundamentals, Best Practices, and Implementation

Recent advances in computational science and engineering enable detailed and accurate numerical simulations at scale. However, making quantitative predictions over long-time intervals or performing many-query tasks such as design optimization and uncertainty quantification remains challenging even on large supercomputers. Data-driven model reduction is emerging as an increasingly useful tool to address these challenges by constructing computationally efficient surrogates of physical processes using simulation and/or experimental data. This minitutorial will cover fundamentals of data-driven model reduction from two complementary perspectives: data in the time domain and data in the frequency domain. We will begin by delving into the theoretical underpinnings of key methods within these domains, followed by practical demonstrations of their application. The implementation of these methods often involves a multifaceted approach, encompassing data processing, linear algebra, and numerical analysis. To provide comprehensive examples, we will showcase real-world applications, highlighting common pitfalls and best practices for data preparation, method selection, and model evaluation. A code repository with hands-on Jupyter notebook demonstrations will be provided.

Ionut-Gabriel Farcas
Virginia Tech
farcasi@vt.edu

Shane A. McQuarrie
Sandia National Laboratories, U.S.
smcquar@sandia.gov

Steffen Werner
Virginia Tech
U.S.
steffen.werner@vt.edu

MT7

Data-Driven Reduced Modeling in the Time and Frequency Domains: Fundamentals, Best Practices, and Implementation Part II

Recent advances in computational science and engineering enable detailed and accurate numerical simulations at scale. However, making quantitative predictions over long-time intervals or performing many-query tasks such as design optimization and uncertainty quantification remains challenging even on large supercomputers. Data-driven model reduction is emerging as an increasingly useful tool to address these challenges by constructing computationally efficient surrogates of physical processes using simulation and/or experimental data. This minitutorial will cover fundamentals of data-driven model reduction from two complementary perspectives: data in the time domain and data in the frequency domain. We will begin by delving into the theoretical underpinnings of key methods within these domains, followed by practical demonstrations of their application. The implementation of these methods often involves a multifaceted approach, encompassing data processing, linear algebra, and numerical analysis. To provide comprehensive examples, we will showcase real-world applications, highlighting common pitfalls and best practices for data preparation, method selection, and model evaluation. A code repository with hands-on Jupyter notebook demonstrations will be provided.

Ionut-Gabriel Farcas
Virginia Tech
farcasi@vt.edu

Shane A. McQuarrie
Sandia National Laboratories, U.S.
smcquar@sandia.gov

Steffen Werner
Virginia Tech
U.S.

steffen.werner@vt.edu

MT8

PETSc the Portable Extensible Toolkit for Scientific Computations

PETSc, the portable extensible toolkit for scientific computations, is used as an algebraic backend in many scientific libraries around the world. It has been deployed on various architectures, from laptops to large exascale systems. In the tutorial, we will cover the basics of PETSc, first by describing its elementary classes, and then by providing examples showcasing how it can be used to solve (non)linear systems, using simple or more advanced preconditioners, with an emphasis on their connection with computational science and engineering. We will also present some other advanced capabilities for example using its sister library SLEPc, the scalable library for eigenvalue problem computations, to solve eigenvalue problems or compute singular value decomposition. Most of these exemples will be implemented live in either C, Fortran, or Python depending on the attendees preference.

Pierre Jolivet
CNRS
pierre.jolivet@cnrs.fr

Junchao Zhang
Argonne National Laboratory
jczhang@anl.gov

Tobin Isaac
Argonne National Lab
tisaac@anl.gov

MT9

BE: Accessible High-Performance Computing Using the Julia Language

Julia offers the flexibility of a high-productivity language while providing control, performance, and compatibility with high-performance computing (HPC) hardware. This mini-tutorial demonstrates how Julia makes modern HPC accessible. It covers resource configuration, distributed computing, code optimization for CPUs and GPUs, and versatile workflows. Participants will have the opportunity to experience it firsthand with a hands-on session on a GPU-powered supercomputer. Tutorial materials, including hands-on exercises and instructions in the use of the Perlmutter Supercomputer at NERSC will be provided here: https://github.com/JuliaParallel/julia-hpc-tutorial-siam-cse25 The first hour covers technical backgrounds and examples in contemporary uses of HPC systems: resource configuration and management; distributed computing; optimizing code of CPUs and GPUs; sophisticated workflows that can run (almost) anywhere! Participants will be provided with examples and training accounts on Perlmutter. The remaining 40 mins will give users the opportunity to work on these examples under the guidance of the instructor. You will need to bring a laptop capable of connecting jupyter.nersc.gov to the in-person session.

Johannes P. Blaschke
Lawrence Berkeley National Laboratory
jpblaschke@lbl.gov

MP2

Constrained Gradient Projection for Sparse Signal Recovery

Signal processing often faces challenges when determining sparse solutions for under-determined linear systems of equations. A common approach for solving these types of problems is Iterative Soft-Thresholding (IST). One of the main challenges with IST is that it may be slow to converge since IST places bounds on the Hessian. In a previous study by Figueiredo et al., gradient projection algorithms were used to address the bound-constrained quadratic programming (BCQP) formulation of these problems. This method does not require bounds on the Hessian and requires only the matrix-vector products, rather than explicit access to a large system. We propose an approach to improve the solutions to these under-determined systems by adding meaningful constraints. Common methods for solving constrained optimization problems include interior point methods, active set methods, and projected gradient methods. Although interior point and active set methods are highly accurate, both require a solution to a linear system at each iteration, which can become computationally expensive as the dimension of the problems increases. Since we are interested in approximating a solution quickly and efficiently, we focus on projected gradient methods, which are more computationally inexpensive compared to the methods previously mentioned. Moreover, we formulate the constrained subproblem in the dual space to further simplify the problem.

Jacqueline Alvarez, Roummel Marcia, Roummel Marcia
University of California, Merced
jalvarez94@ucmerced.edu, rmarcia@ucmerced.edu, rmarcia@ucmerced.edu

MP2

Leveraging Clinical BERT for Comorbidity Prediction in Obstructive Sleep Apnea Patients

Obstructive Sleep Apnea (OSA) is a prevalent sleep disorder characterized by repeated interruptions in breathing during sleep, leading to various significant health issues, including cardiovascular diseases, diabetes, and cognitive impairments. Given the high prevalence of OSA and its associated comorbidities, accurately classifying and predicting these comorbidities is crucial for assessing the overall severity of OSA. This study leverages Clinical BERT, a transformer-based model pre-trained on clinical text data, to predict the presence of comorbidities from clinical text data and assess the severity of OSA based on the Charlson Comorbidity Index (CCI). By analyzing clinical text data, the model aims to provide a comprehensive risk profile for OSA patients, ultimately aiding in better clinical decision-making and patient management.

Neila Bennamane
Florida International University, U.S.
nbenn017@fiu.edu

Destinee S. Morrow
Hood College
Lawrence Berkeley National Laboratory
dmorrow@lbl.gov

MP2

Novel Spectral Graph Clustering Algorithm on Signed Graphs

Graph theory has important applications in the field data science, as graphs can often be used to organize large

datasets and analyze how the data is interconnected. In many applications, a primary goal is to identify the clusters that occur within the graph so that data can be appropriately categorized. However, identifying the appropriate partition for the set of nodes is not always clear. Here, spectral graph theory and optimized graph cuts are used to generate a near optimal clustering of the vertices in the graph. Nonetheless, even an optimal clustering may not reflect the true labels of all the nodes. To this end, we employ an active learning algorithm, wherein vertices near the boundaries of the clusters have their labels sampled and utilized to measure the accuracy of the clustering. Sampled vertices will be used to generate negative edges between vertices corresponding to opposite clusters. While positive weights between nodes indicate how strongly connected two nodes are, negative weights can establish which nodes they repel. These negated edges generate a new graph from which a new clustering can be formulated. In this active learning approach, the process will repeat iteratively until a threshold of accuracy is reached. Utilizing active learning on graphs, we explore a new way of spectral clustering on signed graphs that has applications in dozens of fields where accurate clustering of data is relevant.

Paola Campos
California State University, Stanislaus
pcampos2@csutan.edu

Alex Cloninger
University of California, San Diego
acloninger@ucsd.edu

Sawyer Robertson
UC San Diego
s5robert@ucsd.edu

MP2

Evaluating Performance of Neural Network Quantization for Conjunctival Pallor Anemia Detection

Anemia is a widespread global health issue, particularly among young children in low-resource settings. Traditional methods for anemia detection often require expensive equipment and expert knowledge, creating barriers to early and accurate diagnosis. To address these challenges, we explore the use of deep learning models for detecting anemia through conjunctival pallor, focusing on the CP-AnemiC dataset, which includes 710 images from children aged 659 months. The dataset is annotated with hemoglobin levels, gender, age, and other demographic data, enabling the development of machine learning models for accurate anemia detection. We fine-tuned the MobileNet architecture, known for its efficiency in mobile and embedded vision applications, using data augmentation techniques and a cross-validation strategy. Our fine-tuned model achieved an accuracy of 0.9313, a precision of 0.9374, and an F1 score of 0.9773 demonstrating strong performance on the dataset. To optimize the model for deployment on edge devices, we performed post-training quantization, evaluating the impact of different bit-widths (FP32, FP16, INT8, and INT4) on model performance. Preliminary results demonstrate that while FP16 quantization maintains high accuracy (0.9250), precision (0.9370) and F1 score (0.9377), more aggressive quantization (INT8 and INT4) leads to significant performance degradation. Overall, our study supports further exploration of quantization schemes and hardware optimizations to assess trade-offs between model size, inference time, and diagnostic accuracy in mobile healthcare applications.

Sebastián Cruz Romero, Wilfred Lugo Beauchamp
University of Puerto Rico at Mayagüez
sebastian.cruz6@upr.edu, wilfredo.lugo1@upr.edu

MP2

Dynamic Inverse Problems: Efficient Methods for Dynamic Image Reconstruction with Motion Estimation

Large-scale dynamic inverse problems are typically ill-posed and suffer from complexity of the model constraints and large dimensionality of the parameters. A common approach to overcome ill-posedness is through regularization that aims to add constraints on the desired parameters in both space and temporal dimensions. In this work, we propose an efficient method that incorporates a model for the temporal dimension by estimating the motion of the objects alongside solving the regularized problems. In particular, we consider the optical flow model as part of the regularization that simultaneously estimates the motion and provides an approximation for the desired image. To overcome high computational cost when processing massive scale problems, we combine our approach with a generalized Krylov subspace method that efficiently solves the problem on relatively small subspaces. The effectiveness of the prescribed approach is illustrated through numerical experiments arising in dynamic computerized tomography and image deblurring applications.

Toluwani Deborah Okunola
Tufts University, U.S.
toluwani.okunola@tufts.edu

MP2

Automated Pipeline Supporting AMReX Base Codes for Device Parameters Inspections and Machine Learning Applications

Utilities to Execute Pipelines (UTEP) is being developed and adapted to the AMReX-based codes: FerroX, MagneX, and ARTEMIS to efficiently organize and extract the outputs from simulations that are being systematically sent to the high-performance supercomputer Perlmutter used for device parameters inspection. Additionally, the data is ready for machine learning applications, enhancing the scalability of the trained models to handle more diverse scenarios.

Christian Fernandez
The University of Texas at El Paso
cafernandez12@miners.utep.edu

Jorge Munoz
University of Texas at El Paso
jamunoz@utep.edu

Cesar Diaz
Stanford University
cesardc1@stanford.edu

Yingheng Tang, Prabhat Kumar, Andrew J. Nonaka, Zhi Jackie Yao
Lawrence Berkeley National Laboratory
ytang4@lbl.gov, prabhatkumar@lbl.gov, ajnon-

aka@lbl.gov, jackie_zhiyao@lbl.gov

MP2

Utilizing Latent Space Representation for Patient Phenotyping and Risk Stratification

Obstructive sleep apnea (OSA) is a common sleep-related disorder characterized by intermittent breathing pauses during sleep, which can significantly increase the risk of cardiovascular and metabolic diseases. The often undiagnosed nature of OSA and difficulties in identifying patients most at risk for associated comorbidities has led to sub-optimal personalized patient care. In addition, the complexity and inconsistency of clinical data in electronic health records (EHR) pose challenges in deriving reliable results in healthcare studies. We extracted and compared learned latent spaces– a deep learning technique that compresses data to uncover hidden patterns– to filter out the noise and irrelevant details from the EHR data. We then deep phenotyped OSA patients through unsupervised clustering using the latent representation, identified patient subgroups to uncover potential risk factors that drive subgroup differentiation, and developed a clinical tool to predict patient group assignment via supervised learning. These findings enhance the understanding of OSA deep phenotyping and improve patient comorbidity risk assessment.

Sophia Isabel Perez Ferrer
Hood College
sophiaisabelpf@gmail.com

MP2

Machine Learning the QCD Equation of State

Microseconds after the big bang, matter is theorized to have existed under conditions of extreme temperature and density. To better understand matter in the early universe, and the transition from deconfined quark matter to hadronic matter, nuclear physics has long pushed for a complete phase diagram. In my research I apply machine learning with a kNN algorithm to interpolate thermodynamic behavior. This approach saves valuable computing power, allowing us to sort unphysical equations of state for matter from valid equations without heavy computation. These equations of state are then applied in simulations and in comparison with data from high-energy collision experiments.

Justin Laberge
University of Texas at El Paso
jlaberge@miners.utep.edu

MP2

Coarse-Grained Modeling of Bacterial Adhesion Pili

Enterotoxigenic Escherichia coli (ETEC) is a major global health concern and a leading cause of infectious diarrheal diseases, affecting 200 million people each year and causing 50,000 deaths. The bacterias primary colonization factors, CS20 and CFA/I pili, are essential to its pathogenicity. These pili act like "arms" and are made of flexible, elongated filaments that enable the bacteria to attach to host cells and withstand external forces, such as fluid flow. Investigating pilus unwinding at the molecular level through experiments alone is difficult because of the structural complexity of these filaments, necessitating computational approaches. However, current all-atom simulations are limited to short length and time scales due to its high computational cost. Here, we introduce a novel coarse-grained approach using the GoMartini model to examine pilus unwinding mechanics at a fraction of the computational cost of all-atom methods. This coarse-grained method groups atomic interactions to create computationally efficient simulations that capture pilus unwinding dynamics within hours, as opposed to the days required for all-atom simulations. Our results highlight the potential of the GoMartini approach for efficiently studying large biomolecular complexes over extended time and length scales, advancing computational insights into bacterial adhesion and underlying molecular mechanisms.

Rocky Lu
The College of New Jersey
lujrocky@gmail.com

MP2

Efficient Detection of Protein-protein Interactions Using Sparse Signal Reconstruction

Proteins rarely act in isolation within cells but interact with specific partners to form functional complexes. Because such protein-protein interactions (PPIs) are crucial for understanding cellular information processing, numerous methods have been developed for mapping PPIs. One such method is immunopurification followed by mass spectrometry (IP-MS), which identifies all proteins that interact with a bait protein of interest. While IP-MS is used widely, mapping large PPI networks is expensive since each experiment only profiles one bait. To overcome this limitation, we developed a pooled method that leverages ideas from compressed sensing to profile many baits using a limited number of IP-MS runs. This method first measures multiple linear combination of the PPIs by modifying the immunopurification step to use combinatorial antibody pools to pull down multiple bait proteins. The resulting signals are then decoded using a sparse signal reconstruction algorithm that estimates the PPIs for each bait. We here demonstrate the theoretical validity of the method using simulated experiments and test the practical feasibility using data from an initial pilot experiment. Overall, pooled IP-MS will allow researchers to map larger interaction networks at reduced cost, thus democratizing the study of PPIs.

Moitrish Majumdar, Tomas Rube
University of California, Merced
mmajumdar@ucmerced.edu, trube@ucmerced.edu

Marko Jovanovic, Lena Street
Columbia University
mj2794@columbia.edu, las2256@columbia.edu

MP2

Impact of Deep Neural Network Hyperparameter on Solution Accuracy of Minimization Problems

This project applies the Deep Ritz Method, a deep learning-based approach, to approximate solutions for a one-dimensional (1D) minimization problem. The study examines how neural network hyperparameters, including depth, width, and learning rate, influence the accuracy of the solution. The algorithm's approximations are compared to the analytical solution of the minimization problem, with the accuracy assessed through the UNDEFINED

LATEX error.

Zaire Meachem, Ensela Mema, Natalie Olenkovsky, Keisha Vu
Kean University
zaire.meachem12@gmail.com, emema@kean.edu, olenkovn@kean.edu, vuk@kean.edu

MP2

Predictive Models for Assessing GHG Emissions in Aquaculture Systems

Aquaculture is the worlds dominant source of fish production, having recently surpassed capture fisheries. Most aquaculture operations occur in freshwater environments, utilizing excavated ponds. Recent studies have demonstrated that freshwater aquaculture systems are significant sources of greenhouse gases (GHG) such as carbon dioxide (CO2), and methane (CH4). These emissions are temporally variable, measurements are costly and time-consuming and there is a notable scarcity of predictive models tailored to the aquaculture sector. Drawing on temporally and spatially resolved data from aquaculture systems in Brazil, we use machine learning techniques to model the temporal variability in CO2, and CH4 fluxes from excavated ponds systems. Through our research, we aim to identify the factors and patterns driving temporal variability in GHG emissions from these distinct yet widespread aquaculture production systems. This could help correct datasets lacking temporally resolved measurements.

Maria Camila Mejia Garcia
University of Texas Rio Grande Valley
mariacmejiag80@gmail.com

Hansapani S. Rodrigo
School of Mathematical and Statistical Sciences
University of Texas Rio Grande Valley
hansapani.rodrigo@utrgv.edu

Rafael Almeida
Univeristy of Texas, Rio Grand Valley
rafael.almeida@utrgv.edu

MP2

ParticleFlow: Immersive Mixed Reality Simulations with Graph Neural Networks

Disasters from landslides and avalanches to floods and erosion cause significant damage and loss of life each year, affecting millions and highlighting the need for better understanding and prediction. Traditional simulation methods often fall short in providing real-time, accurate representations of these particulate systems due to their complexity and computational demands.The Graph Network-based Simulator (GNS is a machine learning-based tool using Graph Neural Networks (GNNs) that, addresses this by modeling granular flow dynamics. Our project integrates GNS into a mixed reality (MR) environment, incorporating Augmented Reality (AR).

Cristian Moran
Texas A&M University, San Antonio
cmora0305@jaguar.tamu.edu

Krishna Kumar
University of Texas at Austin
krishnak@utexas.edu

MP2

Compatible Energy Preserving Discretizations for Nonlinear Optical Wave Propagation: The Maxwell-Duffing Approach

We explore the modeling and numerical discretization of Maxwell's equations in nonlinear optical media, specifically focusing on the Maxwell-Duffing model. We present the constitutive laws governing electromagnetic wave propagation in non-magnetic, non-conductive media, describing the material's response using a nonlinear cubic Duffing model coupled with Maxwell's equations. The presentation includes the derivation of energy relations for the one-dimensional nonlinear Maxwell model. We introduce a high-order spatial discretization method based on fully discrete leap-frog finite-difference time-domain (FDTD) methods, designed for the accurate and stable simulation of nonlinear wave propagation. Numerical simulations highlight the effectiveness of these methods in capturing the complex dynamics of electromagnetic waves in nonlinear media. Special attention is given to the implementation of the Second order in time and higher order in space leap-frog scheme and its application to traveling wave solutions. We prove Energy Stability of the Higher Order Yee FDTD Schemes for the cubic Maxwell-Duffing Model and demonstrate these results through Numerical Simulations. This work provides critical insights into the mathematical and computational challenges of modeling nonlinear optical materials, offering robust techniques for advancing research in nonlinear photonics.

Victory Obieke
Oregon State University
victoryadoaora@gmail.com

MP2

Investigation of Oceanic Wave Solutions to a Modified (2+1)-Dimensional Coupled Nonlinear Schrodinger System

Oceanic wave characteristics can be investigated by a modified integrable generalized (2+1)-dimensional nonlinear Schrodinger (NLS) system of equations with variable coefficients. Variable coefficients enhance the modeling capability of the NLS equation, making it a more powerful tool for understanding and predicting wave behaviors in complex oceanic and other physical systems. Here, two newly modified methods, specifically the improved nonlinear Riccati equation method and the improved sub-equation method, have been proposed to investigate the aforementioned nonlinear system. Through the utilization of these methods, we successfully obtain traveling and solitary wave solutions for this nonlinear system. We emphasize several constraint conditions that serve to guarantee the existence of these solutions. In addition, using the Mathematica software package, we provide visual representations, including three-dimensional plots and their corresponding contour plots, for the solutions obtained. This work illustrates that the two proposed approaches provide straightforward and efficient means of acquiring various types of solitons, rational, trigonometric, hyperbolic, and exponential solutions. Moreover, they present a more potent mathematical tool for addressing a variety of other nonlinear partial differential equations that hold significance in the field of applied science and engineering.

Kayode Oluwasegun

Drexel University
kto32@drexel.edu

MP2

Molecular Dynamics Simulations of the Structural Effects of Oncogenic Mutations in the Nucleosome

The information available to the cell to initiate, control, or stall metabolic processes is housed in the genome, whose monomeric unit is the nucleosome core particle (NCP). The nucleosome core particle (NCP), in conjunction with DNA-binding proteins like chromatin remodelers and transcription factors, forms the fundamental cellular apparatus essential for transcription, DNA repair, and DNA replication. However, oncogenic mutations within the histone core can disrupt processes, including histone exchange and nucleosome sliding in different cancer types. Mutations in the core histone proteins have been shown to destabilize the H2B-H4 protein interface, impacting histone octamer stability. Here, we ran 36 ???? all-atomistic simulations of the wild-type (WT) and single point oncogenic mutated systems of the NCP on Anton 2[3]. We show structural and dynamic destabilization of the H2B-H4 interface in the histone core in the mutant systems. We use principal component analysis to characterize the shift in conformation in the histone core, comparing WT and oncogenic systems. Thermal Stability Assay (TSA) and nano-DSC confirm the shift in the NCP's melting temperature. This study will shed light on structural changes in the NCP core upon mutation and suggest how they might affect transcription.

Augustine Onyema, Christopher DiForte, Rutika Patel
Graduate Center, City University of New York
aonyema@gradcenter.cuny.edu,
cdiforte@gradchenter.cuny.edu,
rpatel1@gradcenter.cuny.edu

Sebastien Poget, Sharon Loverde
City University of New York
sebastien.poget@csi.cuny.edu,
sharon.loverde@csi.cuny.edu

MP2

Image Reconstruction Through Polar Coordinate Transformation for Ultrasound Imaging Applications

Image registration plays a crucial role in medical image processing by aligning datasets from different modalities and standardizing training images for downstream deep learning tasks. Traditional registration methods often face challenges when working with ultrasound images due to variations in acquisition between portable and cart-based systems. In this work, we propose a novel approach that first transforms ultrasound images into their native polar coordinates to improve registration accuracy. Our method intends to solve a constrained optimization problem to achieve simultaneous spatio-temporal image registration. By leveraging the geometric properties of polar coordinates, we anticipate that our approach will enhance alignment accuracy and reduce computational complexity, thereby providing a robust preprocessing framework for deep learning models. Through planned experiments on clinical data, we expect to demonstrate improvements in registration accuracy and computational performance compared to existing methods. If successful, this project could significantly advance ultrasound image translation tasks, ultimately enhancing the effectiveness of AI-driven diagnostic imaging and improving patient care outcomes.

Jocelyn Ornelas-Munoz, Erica M. Rutter, Roummel Marcia, Roummel Marcia
University of California, Merced
jornelasmunoz@ucmerced.edu, erutter2@ucmerced.edu,
rmarcia@ucmerced.edu, rmarcia@ucmerced.edu

David Messika-Zeitoun
University of Ottawa Heart Institute
dmessika-zeitoun@ottawaheart.ca

MP2

Colors of Cultures Mars Volcano Images

Colors of Cultures: Culture is defined by the Oxford dictionary as the arts and other manifestations of human intellectual achievement regarded collectively. Can we then use machine learning to determine with an image what culture(s) an image is likely derived from? There are several things to consider, such as what types of cultures do we want to explore and train the model on? Mars Volcano Images: Given images from space agencies as data, can we identify the locations of volcanoes in the images?

Joshua Peoples
Arizona State University, U.S.
jpeople1@asu.edu

MP2

In Silico Analysis of the Flagella-Propelled Cell Swimming in Viscous Fluids

Flagella are intricate cellular structures that serve many functions from cell propulsion to biomixing and transport. The goal of this project is to make a computational fluid dynamics model of a green algae cell with two flagella swimming through viscous fluid. This project employs the AMReX-based fluctuating hydrodynamics solver containing custom-developed algorithms for simulating the flagellum as the immersed chain of force particles with prescribed motions. In addition, we have updated the previously developed yt-centric JupyterLab notebook based on pyAMReX which serves as a native python API for AMReX codes. This implementation has produced better streamlined data importation, visualization, and analysis. We were able to compile and run code on Perlmutter that led to visualizing both fluid flow field and forces on the discretized flagellar particles.

Lauren Sowards, Quinton Fister, Gang Xu
University of Central Oklahoma, U.S.
lsowards@uco.edu, qfister@uco.edu, gxu@uco.edu

Daniel R. Ladiges
Center for Computational Science and Engineering
Lawrence Berkeley National Laboratory
drladiges@lbl.gov

Johannes P. Blaschke
Lawrence Berkeley National Laboratory
jpblaschke@lbl.gov

MP2

Verification of LPDDR4 Memory Controller

DRAM controllers are essential for managing the communication between a processor and DRAM memory, managing

complex operations such as refresh cycles, row and column accesses, and bank switching. Given the complexity of timing requirements and workload management, comprehensive verification is essential to detect potential errors such as data corruption, access violations, or performance degradation. This poster presents a methodology for verifying a memory controller through RTL digital simulation. We focus on timing correctness, data integrity, and performance. The results of this verification will be used to tape out the only existing (as of Jan 2025) open-source DDR controller model, designed to be flexible and parameterized using the Chisel language.

Ailsa Sun
University of California, Berkeley
23suna18@berkeley.edu

Ella Schwarz, Borovije Nickolic
University of California Berkeley
schwarzem@berkeley.edu, bora@berkeley.edu

MP2

Software Engineering Web-scraping

This project presents a flexible web scraping application built in Eclipse using Java. The program allows users to input any website URL and specify up to three CSS selectors to target and extract specific elements from the web page. The extracted data can be saved in multiple formats, including Excel and Notepad, providing users with versatile options for data management and reporting. To ensure efficient data storage and retrieval, the application integrates with an SQL database, enabling users to maintain structured records of their web scraping results. The solution is designed to be user-friendly while demonstrating key concepts of web automation, data extraction, and database integration.

Adriana Tirpak
Slippery Rock University
art1021@sru.edu

MP2

Accelerating PINNs for Heterogeneous Computing Systems

We aim to introduce one of the higher collective communication framework for high performance computing and deep learning workloads, Horovod. Horovod aims to make distributed deep learning workloads fast and scalable. As deep learning workloads become more demanding and computationally expensive we seek ways to cut back on the computational costs associated in the collective communication calls. To best profile such operations we use the built-in, native horovod timeline for minimal and no external overhead. We close out with an analysis of performance gains on a Physics Informed Neural Network Example, conclusion followed by a future work.

Alex Villa, Yuke Li, Xiaoyi Lu
University of California, Merced
avilla49@ucmerced.edu, yli304@ucmerced.edu, xiaoyi.lu@ucmerced.edu

MP2

Developing a Tutorial to Explore Adaptive Mesh Refinement in High Performance Computing Simulations

This project adapts a high-performance computing (HPC) application into a beginner-friendly tutorial on Adaptive Mesh Refinement (AMR) and HPC. Using Frontier (the worlds fastest supercomputer) and AthenaPK, we showcase AMRs role in refining simulations of phenomena, like blast waves. Compiling AthenaPK, running jobs with Slurm, to visualizing results on Andes with VisIt. We model a blast wave shaping an oak leaf, emphasizing AMRs impact and broadening participation in computing.

Wiktoria Zielinska
Oak Ridge National Laboratory
wiktoria41412@gmail.com

Suzanne Parete-Koon
Oak Ridge National Laboratory, U.S.
paretekoonst@ornl.gov

John Holmen
Oak Ridge National Laboratory
holmenjk@ornl.gov

PP1

Constrained Gradient Projection for Sparse Signal Recovery

Signal processing often faces challenges when determining sparse solutions for under-determined linear systems of equations. A common approach for solving these types of problems is Iterative Soft-Thresholding (IST). One of the main challenges with IST is that it may be slow to converge since IST places bounds on the Hessian, $\mathbf{A}^T\mathbf{A}$. In a previous study by Figueiredo et al., gradient projection algorithms were used to address the bound-constrained quadratic programming (BCQP) formulation of these problems. This method does not require bounds on the Hessian and requires only the matrix-vector products involving $\mathbf{A}$ and $\mathbf{A}^T$, rather than explicit access to $\mathbf{A}$. We propose an approach to improve the solutions to these under-determined systems by adding meaningful constraints. Common methods for solving constrained optimization problems include interior point methods, active set methods, and projected gradient methods. Although interior point and active set methods are highly accurate, both require a solution to a linear system at each iteration, which can become computationally expensive as the dimension of the problems increases. Since we are interested in approximating a solution quickly and efficiently, we focus on projected gradient methods, which are more computationally inexpensive compared to the methods previously mentioned. Moreover, we formulate the constrained subproblem in the dual space to further simplify the problem.

Jacqueline Alvarez
University of California, Merced
jalvarez94@ucmerced.edu

PP1

Optimizing Permeability Domain for Hydrogen Fuel Cells: A Numerical Approach Using Freefem++

Hydrogen fuel cells (HFCs) are becoming an important technology for cleaner energy. They produce electricity through chemical reactions without harmful emissions, making them useful in many areas like powering cars, buses, and providing energy in remote locations. As the

world moves towards more sustainable energy, improving the efficiency of fuel cells is essential. In this presentation, we introduce the effect of changing the permeability function inside the Gas Diffusion Layer (GDL), which is in the middle of a HFC, and how it affects the flow of fluids inside the cell. Permeability controls how gases move through the fuel cell, impacting many variables. Understanding these effects is important for making fuel cells work better. Using FreeFEM, we numerically solve the nonlinear system of partial differential equations derived from the convection-diffusion equation and the Darcy law, using the finite element method, considering three essential variables: oxygen concentration, pressure, and velocity. We will also present an optimum permeability function that helps maximize efficiency.

Afrah Alzahrani
King Fahd University of Petroleum and Minerals
g202213580@kfupm.edu.sa

PP1

Applying Acceleration to Krylov Subspace Eigenvalue Solvers

Many modern-day problems including dynamic analysis in structural engineering and electronic structure calculations involve solving a generalized eigenvalue problem with significantly large matrices. In 2002, Golub and Ye introduced a competitive inverse-free Kyrlov subspace iterative solver for this problem that focuses on large, sparse, symmetric matrices. In 2010, Quillen and Ye introduced a block generalization of the method that can handle clustered eigenvalues. Recently, momentum-type acceleration has been shown effective at a low cost to reduce the number of iterations in many iterative methods, including methods for eigenvalue problems. One such method is Nesterov accelerated gradient descent, a widely used extrapolation modification accelerating steepest descent. Here, we will explore applying Nesterov-like acceleration to the Krylov subspace method for the symmetric generalized eigenvalue problem. We will present numerical results demonstrating the effect of fixed, adaptive, and safeguarded adaptive choice of the momentum parameter.

Michelle Baker
University of Florida
mbaker4@ufl.edu

PP1

Structmats: A Matlab Package for Efficiently Computing with Matrices Having Displacement Structure

Matrices with displacement structure (e.g. Toeplitz, Vandermonde, Cauchy) are amenable to cheap representation and fast computations. The algebraic advantages of displacement structure are well-understood, and recent developments have uncovered connections to low-rank properties as well. These newer findings enable new computational methods, such as the use of rank-structured methods in concert with displacement structure to create fast and storage-lean inversion algorithms. Displacement structure arises naturally in applications (e.g. signal processing), and MATLAB is a popular tool in many of these circles. We introduce STRUCTMATS, an object oriented software package for MATLAB that efficiently implements both the established theory and our new findings on displacement structure. It is host to a collection of useful algorithms, including fast matrix-vector products, fast inversion routines, preconditioning tools, low rank and rank-structured approximation methods, and transforms that allow one to shift between popular matrix families via useful algebraic equivalencies. The package is designed so that the optimizations are handled behind the scenes; the only responsibility of the user is to instantiate the right class of matrix.

Levent Batakci
University of Washington
levent@uw.edu

Heather D. Wilber
University of Washington, U.S.
hdw27@uw.edu

PP1

A Data-Driven Reduced Order Model for Two-Layer Quasi-Geostrophic Oceans

The simulation of ocean flows is extremely challenging for several reasons. The first is of course the scale of the problem: the area of an ocean basin is of the order of millions of Km^2. The second reason is the nature of the flow itself, which requires very fine computational meshes to resolve all the eddy scales. The third reason is connected to the first two: high resolution meshes over very large domains lead to a prohibitive computational cost with nowadays computational resources. In this talk, we present a data-driven reduced order model (ROM) for the two-layer quasi-geostrophic equations, a simplified, yet nontrivial, model for ocean dynamics. The main building blocks of our ROM are Proper Orthogonal Decomposition (POD) and a Long Short Term Memory (LSTM) architecture. Using an extension of the classical double-gyre wind forcing test, we assess the accuracy of the POD-LSTM ROM both in the reconstruction and prediction of the flow and quantify the drastic reduction in the computational time allowed by the method.

Lander Besabe
Department of Mathematics
University of Houston
lybesabe@cougarnet.uh.edu

Michele Girfoglio
SISSA, Mathematics Area, mathLab
International School for Advanced Studies
mgirfogl@sissa.it

Annalisa Quaini
Department of Mathematics, University of Houston
quaini@math.uh.edu

Gianluigi Rozza
SISSA, International School for Advanced Studies,
Trieste, Italy
grozza@sissa.it

PP1

Pod-Based Reduced Order Modeling for the Generalized Kuramoto-Sivashinsky Equation

In this study, we present a comprehensive analysis of POD-based Reduced Order Models (ROMs), specifically Galerkin POD and POD-DEIM (Discrete Empirical Interpolation Method), applied to the one-dimensional generalized Kuramoto-Sivashinsky (gKS) equation. The dynamics

of the gKS equation ranges from a spatio-temporal chaos to a non-chaotic, quasi-periodic behaviour with a transition governed by a critical parameter γ. We developed ROMs using three different approaches, which can be classified into single-value ROM and multi-value ROM categories. In the single-value ROM category, two distinct approaches were employed, one for $\gamma = 0$, which corresponds to the classical Kuramoto-Sivashinsky equation exhibiting chaotic behaviour, and another for $\gamma \neq 0$. To construct the ROM for $\gamma = 0$, snapshots were collected from the system's attractors. For $\gamma \neq 0$, we adopted a different approach by running simulations over shorter periods across multiple trajectories to gather snapshots. To build the ROM using multiple γ values, we conducted simulations with various γ values and merged the resulting snapshots to create the data matrix. In all cases, the total simulation time and the number of snapshots collected were consistent. The ROM was tested across a range of γ, incorporating randomness in the initial conditions. Relative error and a reconstructed power spectrum analyses showed that the ROM performed well.

Md Rezwan Bin Mizan
University of Houston
mbinmiza@cougarnet.uh.edu

Maxim Olshanskiy
University Of Houston
Texas,USA
maolshan@cougarnet.uh.edu

Ilya Timofeyev
University of Houston
Texas, USA
itimofey@cougarnet.uh.edu

PP1

Nonlinear Eigenvalue Solver for Periodic Potential Problems

Asymptotically periodic potentials arise in many physical situations, such as the tight binding approximation in quantum mechanics. The associated partial differential equation on the potential can be converted into a spatial periodic eigenvalue problem, possibly with a defect in the potential. For these problems, we propose a novel solver that converts the unbounded domain problem into a nonlinear eigenvalue problem on a finite domain and uses a projector to enforce boundary conditions. We analyze the associated discretization error and provide perturbation bounds for the corresponding eigenvalues. We also demonstrate its performance on an example problem.

Alkiviades Boukas
Department of Mathematics, Cornell University, Ithaca, NY
ab2453@cornell.edu

Maximilian Ruth, David Bindel
Cornell University
maximilian.ruth@austin.utexas.edu, bindel@cornell.edu

PP1

Scalable Overset Computation Between a Forest-of-Octrees and An Arbitrary Distributed Parallel Mesh

Many physical systems cannot be modeled properly by a single mesh, e.g. because the domain contains holes, or several physics layers interact. A known natural solution to this is the combination of several overlapping meshes in a mesh overset (or Chimera grid) approach. To obtain meaningful results, the meshes have to yield a consistent solution wherever they overlap. This necessitates an efficient exchange of solution data between several meshes, which may be distributed in parallel. We present an algorithm that performs a one-directional mesh overset for a parallel distributed forest of octrees with another generic mesh of unrelated partition. The generic mesh is discretized by parallel distributed sets of query points, e.g. stemming from a quadrature rule applied to each cell. A parallel distribution of the forest following the Morton order can be encoded efficiently and enables communication-free searching in the forest's exact partition, which suffices to determine the communication pattern. Query points are then sent between the two meshes by non-blocking point-to-point communication, which may be overlapped with computation. Our algorithm is formulated in a modular manner. Thus, it can be generalized to tune the load balancing of the mesh overset or to guide adaptive refinement of the meshes around their intersection area. We present the algorithm's design and its application to these different tasks, including numerical results showcasing its scalability.

Hannes Brandt
University of Bonn
brandt@ins.uni-bonn.de

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

PP1

Accuracy and Estimation of Hierarchical Structures of Wave Equation Responses Based on Our Background Seismic Wave-Field

Distributed Acoustic Sensing (DAS), the repurposed use of dark telecommunication cables for seismological investigation, is proving to be of increased interest in the scientific community. Though extremely promising, DAS produces large amounts of data, some studies having multiple terabytes a day. Our research aims to find new ways of making this large amount of data easier to comb through so geophysicist and other interested scientists may have an easier time of finding moments of intrigue. In particular, were using a well known mathematical technique, Singular Value Decomposition (SVD), to compress and reconstruct cross-correlations of band passed signals received in the frequency domain. A simulation of 40 sensors with sources at different locations was used to test which frequencies and locations of source were the most lucrative in an SVD reconstruction. Data on how many singular values needed to reproduce a desired percentage of the cross-correlations was recorded. Finding ways of tracking changes in these cross correlations is now being performed. As a precursor to accurate estimation of hierarchical structures of wave equation responses based on our background seismic Wavefield, compression through SVD is a promising tool for improving the efficacy of DAS data.

Georgia A. Brooks
Colorado School of Mines
georgia_brooks@mines.edu

PP1

Accurate Temporal Integration Schemes for Non-

linear Adsorption Problems

We consider a nonlinear transport problem to model the chromatography process of high-capacity multimodal membranes. The absorption model for this process is nonlinear and can be represented by either explicitly or implicitly defined adsorption isotherms. A variety of time-integration schemes are developed through the use of Rothe's method. Contrasted with the traditional method of lines, this approach offers a flexibility in the manipulation of the temporally nonlinear adsorption term. We generate numerical solutions utilizing high-level software such as the FreeFEM, FEniCSx, and deal.ii environments. Further, nonlinear solution strategies are implemented in tandem with these time-integration schemes, incorporating the PETSc SNES and SUNDIALS KINSOL libraries. Lastly, a parameter sweep is performed, analyzing the performance of these methods for the space of feasible adsorption parameters.

Evan Butterworth
Clemson University
edbutte@clemson.edu

Lea Jenkins
clemson university
lea@clemson.edu

PP1

GPU Acceleration with the Fortran Standard Language

GPUs have become widespread in high performance computing (HPC) applications as they can often provide higher FLOP/s, memory bandwidth, and power efficiency compared to CPUs. However, developing HPC codes (new and legacy) to run effectively on GPUs has been challenging, as the standard programming languages have not directly supported GPU offload. This has led to the development of several external/extended language APIs with varying levels of portability and learning curves. In an attempt to improve portability and productivity for domain scientists, compiler/GPU vendors have recently started supporting GPU offload using preexisting standard language constructs in Fortran, C++, and Python. In Fortran, the 'do concurrent' loop syntax introduced in the Fortran 2008 standard has now been leveraged to support GPU-acceleration on NVIDIA, Intel, and AMD GPUs. In this presentation, we demonstrate this support using two production multi-GPU applications in heliophysics. We show the performance and portability of 'do concurrent' for GPU computing across vendors, noting when and where limited use of external APIs (e.g. OpenMP/OpenACC) are still needed/desired for compatibility and/or optimal performance.

Ronald M. Caplan, Miko Stulajter
Predictive Science Inc.
caplanr@predsci.com, miko@predsci.com

Jon Linker
Pedictive Science Inc.
linkerj@predsci.com

PP1

Solution Methodology for Nonlinear Eigenvalue Problems from Plasma Instabilities

Stability analysis of electron shear-flow in a crossed-field diode yields an eigenvalue problem, where the perturbation electric field amplitude and frequency represent the systems eigenvector and eigenvalue, respectively [Phys. Fluids 27, 2332-2345 (1984)]. This eigenproblem is expressed as a homogeneous two-point boundary value problem that is solved using shooting method to determine the dispersion relation for the plasma instability [Phys. Plasmas 18, 033108 (2011)]. In this work, we propose a reformulation that transforms the original eigenvalue equation into a generalized nonlinear eigenvalue problem to quantify electron shear-flow instability when the boundary value problem is singular. A polynomial representation of the resulting eigenproblem for the electrostatic and electromagnetic assumptions is shown to be of degree 9 and 15, respectively, which makes the use of direct methods computationally expensive. The present works investigates the use of iterative techniques based on the Krylov and rational Krylov methods [SIAM J. Sci. Comput. 36, A2842-A2864 (2014)] for solving these nonlinear eigenproblems with the aim of identifying the most advantageous computational formulation and solution methodology in terms of computational efficiency and favorable comparison with direct methods. This work is supported by the DOE NNSA LRGF under cooperative agreement DE-NA0003960. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Andres Castillo
Stanford University
amc58@stanford.edu

Anubhav Dwivedi
Stanford University
Stanford, CA
anubhavd@stanford.edu

Matthew Hopkins
Sandia National Laboratories
Albuquerque, NM
mmhopki@sandia.go

Kentaro Hara
Stanford University
kenhara@stanford.edu

PP1

Inverse Parameter and Shape Problem for an Isotropic Scatterer with Conductivity Coefficients

We investigate the direct and inverse problems for isotropic scatterers with two types of conductive boundary conditions. The study establishes the uniqueness of recovering coefficients from known far-field data at a fixed incident direction across multiple frequencies. Additionally, we consider the inverse shape problem, demonstrating the recovery of the scatterer's shape from far-field measurements at a single fixed frequency. We explore the direct sampling method for scatterer recovery, focusing on the factorization of the far-field operator. Our findings highlight the method's stability in the presence of noisy data and its applicability to two-dimensional partial aperture data. The theoretical results are verified with numerical examples, illustrating the performance of the direct sampling method.

Rafael Ceja Ayala
Arizona State University
rcejaaya@asu.edu

Isaac Harris

Purdue University
harri814@purdue.edu

Andreas Kleefeld
Forschungszentrum Jülich GmbH
Jülich Supercomputing Centre
a.kleefeld@fz-juelich.de

PP1

Adaptive Safeguarded Newton-Anderson at Singular and Nonsingular Points

We present an adaptive safeguarding scheme with a tunable parameter, which we call adaptive γ-safeguarding, that one can use in tandem with Anderson acceleration to improve the performance of Newtons method when solving problems at or near singular points. Newton-Anderson with adaptive γ-safeguarding converges locally for singular problems; it detects nonsingular problems automatically and responds by scaling the iterates towards standard Newton asymptotically. The result is a flexible algorithm that performs well for singular and nonsingular problems and can recover convergence from both standard Newton and Newton-Anderson with the right parameter choice. We will discuss three strategies one can use when implementing Newton-Anderson and γ-safeguarded Newton-Anderson to solve parameter-dependent problems near singular points. These strategies are demonstrated with two incompressible flow problems near bifurcation points.

Matt Dallas
University of Dallas
mdallas@udallas.edu

Sara Pollock
University of Florida
s.pollock@ufl.edu

Leo Rebholz
Clemson University
Department of Mathematical Sciences
rebholz@clemson.edu

PP1

Advancements in GPU-accelerated PIC Simulations with iPIC3D

Particle-in-cell simulations are essential for studying kinetic-scale plasma behaviour in laboratory settings, such as Tokamaks and stellarators, and astrophysical environments where phenomena like magnetic reconnection and shocks can accelerate particles to very high energies. We present the recent developments in the iPIC3D code, focusing on porting its two most computationally intensive modules, the particle mover and moment gatherer, to GPU architectures using CUDA. Initial comparisons with the CPU version of the code show a factor of 40 improvement in computational speed. Furthermore, we incorporate the exact energy-conserving semi-implicit method and the relativistic semi-implicit method, that enable energy conservation up to machine precision. In certain physical scenarios, this may be crucial to avoid artificial growth of energy of the system over long time scales and obtain physically viable results. We show excellent strong and weak scaling of iPIC3D on up to 1000 GPUs. Ongoing work is focused on enabling simulations across several thousands of GPUs, paving the way for exascale plasma simulations.

Pranab J. Deka, Fabio Bacchini
KU Leuven
pranab.deka@kuleuven.be, fabio.bacchini@kuleuven.be

PP1

Dynamically Regularized Lagrange Multiplier Methods for the Cahn-Hilliard-Navier-Stokes System

We present efficient and accurate numerical schemes for the Cahn-Hilliard-Navier-Stokes phase field model of binary immiscible fluids. By introducing two Lagrange multipliers and their dynamic equations, we reformulate the original model problem into an equivalent system that incorporates the energy evolution process. First- and second-order dynamically regularized Lagrange multiplier (DRLM) schemes are derived based on backward differentiation formulas, in which all nonlinear terms are treated explicitly and no stabilization term is imposed. These methods are mass conserving, unconditionally energy stable with respect to the original variables, and fully decoupled (requiring the solution of two biharmonic-type equations and two generalized Stokes systems, together with two nonlinear scalar equations at each time step). Both theoretical results and numerical performance of the DRLM schemes will be presented.

Kha Doan, Thi Thao Phuong Hoang
Auburn University
kcd0030@auburn.edu, tzh0059@auburn.edu

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Rihui Lan
Ocean University of China
lanrihui@ouc.edu.cn

PP1

Identifiability Analysis of Cancer Growth Models from Dose Response Curves

Often when a new cancer drug is tested, the researchers publish dose-response curves. These are measured by looking at cell viability at a specific time as a function of dose. Unfortunately, mathematical models are often parameterized using longitudinal data. In this study, we examine whether it is possible to accurately identify any cancer growth model parameters and drug efficacy parameters from dose response curves. This would allow us to use dose response curves to construct mathematical models to predict the effect of different dosing regimens. We use mathematical analysis as well as computer simulations to determine identifiability of different model parameters.

Hana Dobrovolny
Department of Physics and Astronomy
Texas Christian University
h.dobrovolny@tcu.edu

Bidhan Karki, Tiffany McHenry
Texas Christian University
bidhankarki7@gmail.com, t.y.mchenry@tcu.edu

PP1

Waves in Black Hole Geometries: An Energy-

Based Discontinuous Galerkin Method

We are interested in the behavior of waves in the vicinity of black holes, strongly influenced by the curvature of spacetime. We highlight the special case of a scalar wave equation in the vicinity of a rotating, uncharged (Kerr) black hole, on a domain that is partially contained within the event horizon. We present an energy-based discontinuous Galerkin method (EDG-GR) for the simulation of waves in black hole geometries, distinguished in that (1) it is localized, making it efficient and parallelizable; (2) accurate radiation conditions allow simulation over a smaller region while preserving accuracy, and (3) the choice of a first-order time, second-order space (FOTSOS) form greatly reduces the number of fields to be evolved.

Jaryd Domine
Southern Methodist University
jdomine@smu.edu

Thomas M. Hagstrom
Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

Stephen Lau
University of New Mexico
lau@unm.edu

PP1

An Einsum-Inspired Tensor Contraction Engine

The Einstein summation convention (einsum) allows to define complex linear algebra expressions using a concise and expressive notation. We have developed a new engine for efficient evaluation of einsum expressions. The engine achieves performance portability by relying on just-in-time generated computational primitives and optimized data layouts for intermediate tensors. This poster presents a new packing strategy for binary tensor contractions that targets high cache reuse. We also present a new memory manager for intermediate data. Through a stack-based allocation approach, our memory manager minimizes memory fragmentation and the overall memory footprint. We demonstrate the performance portability of our approach by evaluating our engine on NVIDIA Grace, AMD Ryzen, and Intel Xeon CPUs. We show that our backend achieves peak utilization of up to 70% for a set of demanding einsum expressions, while the memory manager reduces the memory footprint of intermediate data by up to 69%.

Max Engel
Friedrich-Schiller-Universität Jena
max.engel@uni-jena.de

Alexander Breuer
University of California, San Diego
alex.breuer@uni-jena.de

PP1

Improving Ocean State Estimation with Underwater Acoustics

Informed ocean state estimation is vital for understanding dynamic processes of Earths climate, yet traditional observing systems remain sparse and struggle in accessing changes of the oceans interior. Acoustic tomography offers a computationally robust approach to overcome these limitations by inferring the oceanic state on regional to basin scales through propagation of sound. This presentation describes the governing equations that define ocean dynamics and underwater acoustic propagation, emphasizing the mathematical formulation of predicted sound travel times. A 4-dimensional variational method now incorporates these travel times as a novel observational constraint. We show model-data misfit accumulation within a regional model, enhancing its ability to adjust oceanic conditions. The framework is applied to analyze the sensitivity of hydrographic states to acoustic measurements, demonstrating the potential of this approach for improving ocean state estimation.

Ivana Escobar
Oden Institute
Applied Research Laboratories
Ivana@utexas.edu

Patrick Heimbach
Oden Institute
The University of Texas at Austin Department of E
heimbach@ode.utexas.edu

Feras Habbal
Applied Research Laboratories
University of Texas at Austin
feras.habbal@arlut.utexas.edu

PP1

Event-Based Eligibility Propagation with Additional Biologically Inspired Features

Understanding how circuits of neurons learn is crucial for advancements in both neuroscience and artificial intelligence. Traditional methods like backpropagation through time (BPTT) are commonly used to study neural learning processes; however, they lack biological plausibility. Eligibility propagation (e-prop), a biologically inspired alternative proposed by Bellec et al. [Bellec et al., A solution to the learning dilemma for recurrent networks of spiking neurons, 2020], simulates learning mechanisms akin to those in animal brains. We introduce bio-inspired modifications to the e-prop algorithm that maintain learning performance while enhancing biological accuracy. We also reformulate the original model's time-driven update routine. Instead of updating synaptic states at every time step, our event-driven approach updates them only when a spike is transmitted, which improves computational efficiency. This is particularly effective in simulations of biologically realistic, sparsely connected spiking neural networks (SNNs). Our modifications, implemented within NESTa leading simulator for large-scale SNNsdemonstrate robust scalability in extensive network simulations. These advancements bridge the gap between computational efficiency and biological realism, significantly enhancing e-prop's applicability in large-scale and biologically faithful spiking neural network simulations.

Jesús Andrés Espinoza Valverde
Department of Mathematics and Science
University of Wuppertal
espinozavalverd@uni-wuppertal.de

Agnes Korcsak-Gorzo
Institute for Advanced Simulation 6, Jülich Research Centre
Dept. of Physics, RWTH Aachen University
a.korcsak-gorzo@fz-juelich.de

Jonas Stapmanns
Dept. of Physiology, University of Bern
jonas.stapmanns@unibe.ch

Hans Ekkehard Plesser
Dept. of Data Science, Norwegian University of Life Sciences
Institute for Advanced Simulation 6, Jülich Research Centre
hans.ekkehard.plesser@nmbu.no

David Dahmen
Institute for Advanced Simulation 6, Jülich Research Centre
d.dahmen@fz-juelich.de

Matthias Bolten
Dept. of Mathematics and Science, University of Wuppertal
bolten@uni-wuppertal.de

Sacha van Albada
Institute for Advanced Simulation 6, Jülich Research Centre
Institute of Zoology, University of Cologne
s.van.albada@fz-juelich.de

Markus Diesmann
Institute for Advanced Simulation 6, Jülich Research Centre
Dept. of Physics, RWTH Aachen University
diesmann@fz-juelich.de

PP1

3D Cardiac Strain Imaging of Left Ventricle

Human heart is a vital focus in medical diagnosis and scientific research, with cardiac strain imaging being a fundamental tool, particularly for the left ventricle, a critical component of the human heart. Abnormal strain patterns in specific areas of the left ventricle can indicate various health issues. While advancements in probe design and manufacturing have significantly refined 2D strain imaging, 3D cardiac strain imaging is gaining increased attention due to its enhanced research potential. Calculating 3D strain directly from 3D heart imaging data is time-consuming and challenging, particularly when converting 3D strain results into specific models. In this study, we applied a classic speckle tracking algorithm integrated with principal component analysis to calculate smoother 2D cardiac strain. Additionally, we developed a unique and intuitive method to transform 2D strain data into a 3D model, which is both highly convenient for visualization and universally applicable to strain data from any subject. To verify the accuracy of our approach, we conducted a phantom experiment that demonstrated an accuracy rate exceeding 90%. The entire process involves streaming data from a 3D probe, decomposing the data into 2D segments, and reconstructing the results into a comprehensive 3D model.

Yiyue Feng
San Diego State University
yfeng9707@sdsu.edu

PP1

Performance Evaluation of Algorithms for Tall-Skinny QR Factorization on Recent Computer Systems

QR factorization of tall and skinny matrices is a fundamental matrix computation that has several applications including the calculation of an orthogonal basis in numerical methods for solving linear systems and eigenvalue problems. Traditional algorithms like the Gram-Schmidt and Householder algorithms have been widely known and used in practical applications. On the other hand, new algorithms such as the TSQR and CholeskyQR algorithms have been developed to exploit the capabilities of modern computer systems. In this poster, we provide a brief overview of these QR factorization algorithms with highlighting their characteristics and present their performance results, measured on recent computer systems. Although our focus is on deterministic algorithms, we will also discuss potential future comparisons with randomized algorithms actively developed in the community of numerical linear algebra and high-performance computing.

Takeshi Fukaya
Hokkaido University
fukaya@iic.hokudai.ac.jp

PP1

Quantifying Spatial Heterogeneity of Syncytial Cells Using Alpha Shapes

We introduce a structural method used for quantifying the spatial heterogeneity(or clumpiness) of viral syncytial cells in a transfected bioassay. The solution lies in an interdisciplinary process based on simplicial topology being applied to a biological system. Our method revolves around using topological theories including Delaunay Tessellations and Voronoi Graphs to signify cell-cell interaction probability. The main emphasis is the subset of Delaunay Tessellation called Alpha Shapes. By applying a filtration to the overall Delaunay Tessellation, we can obtain unique Alpha Shapes that have cell-cell interactions removed. The emphasis of the filtration is to find the correct shape where there were no connection crossing syncytia, only between healthy neighborhoods of cells. The process allows for the associated alpha number to be assigned to the clumpiness. Alpha numbers can then be used to separate different bioassays, or quantify temporal changes found in a single viral transfection due to syncytia.

Anthony D. Gerg
Graduate Student, Texas Christian University
a.d.gerg@tcu.edu

Hana Dobrovolny
Department of Physics and Astronomy
Texas Christian University
h.dobrovolny@tcu.edu

PP1

Cross Interpolation for Solving High-Dimensional Tensor Differential Equations on Low-Rank Tensor Train and Tucker Manifolds

We present a novel algorithm for the time integration of nonlinear tensor differential equations (TDEs) on the tensor train and Tucker tensor low-rank manifolds. The presented methodology offers multiple advantages: (i) It offers near-optimal computational savings in terms of memory and floating-point operations by leveraging cross algorithms based on the discrete empirical interpolation method. (ii) It is robust against small singular values

and high-order explicit Runge-Kutta schemes are developed. (iii) The algorithm is easy to implement, and it does not use tangent space projections, whose implementation is intrusive. The efficiency is demonstrated in several cases, including a nonlinear 100-dimensional TDE.

Behzad Ghahremani
University of Pittsburgh
beg66@pitt.edu

PP1

A High Order Finite Difference-Discontinuous Galerkin Method for the Elastic Wave Equation

We propose a versatile and efficient hybrid framework for solving the elastic wave equation in a two-dimensional discontinuous medium. We employ the high-order discontinuous Galerkin method within a limited portion of the domain in proximity to complex geometric features. Meanwhile, we utilize the computationally efficient summation-by-parts finite difference method in a majority of the domain. To couple the two discretizations, we construct and optimize a set of accuracy and stability preserving projection operators. The overall discretization satisfies a discrete energy estimate and converges to optimal order of accuracy.

Andreas Granath
Department of Mathematics and mathematical statistics
Umeå University, Sweden
andreas.granath@umu.se

Siyang Wang
Department of Mathematics and Mathematical Statistics
Umeå University, Sweden
siyang.wang@umu.se

PP1

A New Radiosity Method for Occluded Global Illumination

The radiosity method is an approach to simulate viewpoint-independent global illumination in contrast to viewpoint-dependent methods like ray tracing. Radiosity works by expressing the pairwise light transfer between surface patches in a scene as an integral and postulating strict local radiation balance. This leads to a Fredholm integral equation of the second kind, which is numerically solved to obtain the global equilibrium. The algorithmically and mathematically most challenging aspect of the radiosity method is to handle occlusion. Occlusion calculations are algorithmically demanding since every patch in a given scene is a potential occluder for every pair of patches in the scene. Mathematically, the occlusion introduces a problematic decrease of kernel regularity in the radiosity integral equation due to kernel discontinuities depending on the whole scene. We present a reformulation of the radiosity equation as a system of coupled integral equations that do not require explicit occlusion detection. Then we use the improved regularity to solve the system of integral equations using a fast multipole type method that we tailor to the given system. We hereby aim to achieve a linear runtime and memory complexity. We show practical results of our current shared-memory parallelized implementation, which is an intermediate step towards a high performance computing implementation employing a distributed-memory parallelism.

Tim Griesbach
Universität Bonn, Germany
Institute for Numerical Simulation
tim.griesbach@uni-bonn.de

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

PP1

Optimizing Heterogeneous Porous Media Flow Simulations with Mrcm-O Preconditioning

The development of efficient and accurate solvers for large-scale porous media flow problems is of paramount importance to the industry, particularly in heterogeneous formations. Recent advances in multiscale mixed methods, specifically those based on domain decomposition techniques, have shown promise due to their natural parallelizability and effectiveness in handling complex problems. In this work, we extend these methods by utilizing the Multiscale Robin Coupled Method with Oversampling (MRCM-O) as a preconditioner for challenging industry-relevant problems. Our approach leverages the MRCM-Os iterative method with oversampling and smoothing techniques, previously demonstrated to achieve flux accuracy on the order of 1e-10 within approximately 10 iterations when tested on the SPE10 benchmark, a standard for evaluating numerical methods in porous media. This demonstrates the methods capability to handle intricate permeability variations, including high-contrast and channelized formations. We focus on the practical implementation of these advanced preconditioning techniques in real-world scenarios, emphasizing the scalability and robustness of the solver. Preliminary applications suggest that the MRCM-O-based preconditioning approach holds significant potential for improving the efficiency and accuracy of numerical solutions in industry-relevant porous media flow simulations. Ongoing work aims to further validate these findings.

Rafael Guiraldello
Piri Technologies LLC
USA
rafaeltrevisanuto@gmail.com

Dilong Zhou, Felipe Pereira
The University of Texas at Dallas
dxz200000@utdallas.edu, luisfelipe.pereira@utdallas.edu

PP1

An Interior-Point Multigrid-Based Approach for Scalable Computational Contact Mechanics

A critical aspect of modeling complex engineering systems is the interaction of physical bodies in contact. Frictionless contact problems are modeled so that the state is the minimizes an energy functional. However, generally such problems are: nonlinear, nonconvex and contain an optimization variable whose dimension is unbounded with respect to mesh refinement. We focus on the scalable solution of such large-scale contact mechanics problems on high-performance computing systems. We employ a Newton-based interior-point filter line-search method, that is one of the most robust methods for nonlinear nonconvex constrained optimization, to computationally estimate minimizers of such large-scale constrained optimization problems. The outer Newton-based interior-point loop converges rapidly; however, each step requires the solution of a large saddle-point linear system. A major challenge with the inner interior-point Newton-based linear system is that,

in addition to the general challenges of solving large-scale linear systems, it can become arbitrarily ill-conditioned as the optimizer estimate approaches the optimal point. In this talk, we detail an interior-point multigrid-based approach for solving such problems and present scaling results obtained from an implementation of said framework on a few contact mechanics example problems. The results show that the solution of various contact mechanics problems can be achieved in a manner that scales well in the large-scale regime.

Tucker Hartland
Lawrence Livermore National Laboratory
hartland1@llnl.gov

Socratis Petrides
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
petrides1@llnl.gov

Tzanio Kolev, Cosmin G. Petra
Lawrence Livermore National Laboratory
kolev1@llnl.gov, petra1@llnl.gov

PP1

Data-Assimilated Time and Space Dependent Diffusion Coefficient Estimation in Marine Lakes

Marine lakes, located near the sea, present unique environments characterized by limited exchanges between the lake and the nearby sea. Understanding the dynamics of these exchanges, particularly the temporally and spatially varying diffusion coefficient, D, is crucial for modeling and predicting the lake's behavior. This poster presents an approach to estimate a time and depth dependent D. We formulated this inverse problem as an optimization problem governed by a linear reaction-diffusion equation and solved it using inexact Newton-CG. The first and second-order derivatives are obtained via adjoint calculus. Our partial differential equation (PDE) constraint is modeled using a linear reaction-diffusion equation that captures the distinct characteristics of marine lakes, including subterranean water exchanges and the absorption of solar heat by algal layers. Through empirical investigation, we demonstrate the computational trade-offs of utilizing one versus two sets of data and compare our method's performance to Physics-Informed Neural Networks (PINNs). The results reveal that the proposed method outperforms PINNs when data are scarce and the parameter-to-observation differential operator is nonlinear, making it particularly well-suited for understanding marine ecosystems but also holds promise for broader applications.

Alex Ho, François Blanchette, Noémi Petra
University of California Merced
aho38@ucmerced.edu, fblanchette@ucmerced.edu,
npetra@ucmerced.edu

Michael Dawson
University of California, Merced
mdawson@ucmerced.edu

PP1

Structure-Preserving Finite Element Schemes for the Euler-Poisson Equations

We present a structure-preserving numerical discretization for the repulsive and attractive Euler-Poisson equations that find applications in fluid-plasma and self-gravitation modeling. The scheme is fully discrete and structure-preserving in the sense that it maintains a discrete energy law as well as hyperbolic invariant domain properties such as positivity of density and a minimum principle on the specific entropy. We discuss the algorithmic details of the scheme and present numerical experiments that demonstrate the performance of the method.

Jordan N. Hoffart
Texas A&M University
jordanhoffart@tamu.edu

PP1

Developing K-Core and K-Truss Algorithms in Large Dynamic Graph

K-core and K-truss decomposition identify tightly connected components in networks, offering insights into their structure, dynamics, and resilience. These tools help in understanding network stability by focusing on the most cohesive and robust substructures. Dynamic networks are crucial for understanding real-world systems that evolve over time, such as social, biological, and communication networks. They allow for the analysis of temporal patterns, predictive modeling, and the assessment of resilience and stability. By capturing the changing nature of interactions, dynamic networks provide deeper insights into system behavior, aiding in more effective decision-making and adaptation. In this poster, I will present a novel parallel algorithm for computing K-core and K-truss decomposition in large dynamic graphs. Our algorithm leverages the distinct structural properties of the graphs, including whether they are perfect, chordal, etc. By utilizing these properties, our algorithm demonstrates significant improvements over state-of-the-art algorithms, advancing the field of dynamic graph analysis.

Farahnaz Hosseini
University of North Texas
University of North Texas
farahnazhosseini@my.unt.edu

Sanjukta Bhowmick
University of North Texas
sanjukta.bhowmick@unt.edu

PP1

Mathematical Models of the Zebrafish Segmentation Clock: Notch Signaling and Traveling Waves

The segmentation clock is a gene expression oscillator controlling rhythmic segmentation of the vertebral column during embryonic development. In zebrafish, Her1 and Her7 proteins generate oscillatory gene expression with a period of 30 minutes. These oscillations are synchronized in pre-somitic mesoderm (PSM) via Delta-Notch signaling. The period of oscillations lengthens as cells move from posterior to anterior PSM, creating traveling waves of clock gene expression. Here, we present two models to elucidate the underlying dynamics of the segmentation clock and traveling waves. The first differential equation model explores Delta proteins roles in synchronizing oscillations. Through pseudo-stochastic simulations and Genetic Algorithm optimization, this model replicates experimental observations and differentiates DeltaC and DeltaD roles. Our model reveals DeltaC protein is key for synchronizing clock oscillations, and the newly discovered positive feedback loop through the deltaC gene enhances the robustness of clock oscillations. The second model incorporates ppErk

and Dusp proteins to study traveling waves. Model simulations with 200 cells accurately capture traveling wave dynamics, including a 17-minute cycle shift for ppErk mutants. This work highlights the power of mathematical modeling, parameter estimation, and computational simulations in advancing our understanding of complex biological processes.

Xiyan (Sia) Hu
Departments of Mathematics and Computer Science
Colgate University
xhu@colgate.edu

Eslim Alpay
Department of Cell and Developmental Biology
Northwestern University Feinberg School of Medicine
eslimesra.alpay@northwestern.edu

Oriana Zinani
Division of Developmental Biology
Cincinnati Childrens Hospital Medical Center
orianazinani@gmail.com

Ahmet Ay
Departments of Biology and Mathematics
Colgate University
aay@colgate.edu

Ertugrul zbudak
Department of Cell and Developmental Biology
Northwestern University Feinberg School of Medicine
ozbudak@northwestern.edu

PP1

Acceleration of Approximate Maps for Matrices Arising in Discretized PDEs

Generally, discretization of partial differential equations (PDEs) creates a sequence of linear systems $A_k x_k = b_k, k = 1, 2, ..., N$ with well-known and structured sparsity patterns. For solving closely related systems in sequence, we can use preconditioner updates such as Sparse Approximate Map (SAM) instead of computing a preconditioner for each system from scratch. A SAM acts as a map from one matrix in the sequence to another nearby one for which we have an effective preconditioner. We seek to compute an optimal sparsity pattern to efficiently compute an effective SAM update. In this poster, we examine several sparsity patterns for computing the SAM update in an effort to characterize optimal or near-optimal sparsity patterns for linear systems arising from discretized PDEs. The allowable number of nonzeros in the sparsity pattern should strike a balance between the accuracy of the map and the cost to apply it in the iterative solver. We can show that the sparsity pattern of the exact map is a subset of the sparsity pattern of the transitive closure of a graph representation of A_k, $G(A_k)$. Additionally, we make use of the heterogeneous computing environment to accelerate the computation of the SAM. The inherently parallel nature of the SAM algorithm naturally lends itself towards efficient implementation in GPU and distributed computing systems. We present preliminary results in this area.

Rishad Islam
Lehigh University
mds222@lehigh.edu

Arielle K. Carr
Lehigh University, U.S.
arg318@lehigh.edu

Yuesheng Xu, Yaodong Sheng
Lehigh University
yux220@lehigh.edu, yas616@lehigh.edu

PP1

Managing and Analyzing Large-Scale Seismic Data Using Array Coherence

As sensor technologies advance, the need to efficiently manage and interactively query vast datasets for further analysis becomes increasingly critical. Emerging seismic sensing technologies, such as distributed acoustic sensing (DAS) and low-cost accelerometer nodes, have made it easier than ever to continuously collect high-resolution, large-scale seismic data. However, public seismology data archives are struggling to accommodate these massive datasets, and research teams face significant challenges in exploring, transferring, visualizing, and rapidly interpreting such large volumes of data. In this study, we demonstrate how array coherence can be leveraged to effectively manage the surplus of data, serving as a robust method for signal detection and identifying regions of interest for further analysis. Array coherence, which measures the similarity between sensor pairs within an array, has proven to be particularly effective in enhancing the detection of weak or emergent signals in noisy environments. We explore the parallels between array coherence analysis and principal component analysis, identifying computational efficiencies that can be transferred between these methods. Additionally, we investigate the use of QR decomposition to improve the localization of signals detected through array coherence, offering a more precise approach to signal detection and analysis in large-scale seismic datasets.

Abdul Hafiz S. Issah
Colorado School of Mines
aissah@mines.edu

PP1

Masking Satellite Trails in Astronomical Images

The Condor Array Telescope is an astronomical telescope made up of an array of six refracting telescopes, optimized for the detection of low-surface-brightness features and point sources. The telescope is capable of taking astronomical images at a very rapid cadence, such as 60 seconds or even less with a limited amount of sky-noise. Sky-noise can be described by any type of contaminants in astrophotography, such as hot/bad pixels, tracking errors, satellite trials and any other interferences. Though Condor has implemented software to minimize the amount of errors in the images it captures, there remains some contaminants in some of its images. One of the most frequent ones that remain are satellite trails, described by a thin line of pixels that stretch across the image. In order to maximize the amount of data gained from the images, our goal is to identify and track these satellites, calculating their trajectories and matching up the trail to the trajectory. Using a list of TLEs, two-line element sets that describe Earths orbital elements at a given time, or epoch, we can gather data about each satellite, getting its right ascension and declination and matching it up with the trail in the image. We wish to work further on this to mask over the satellite trail pixels in the image using the calculated trajectory and making this entire process autonomous.

Mandy D. Joseph, Mandy D. Joseph

Stony Brook University
mandy.joseph@stonybrook.edu,
mandy.joseph@stonybrook.edu

PP1

Using Low-Rank Methods and the Operator Fourier Transform to Simulate Waves

This poster considers low-rank techniques for evaluation of the operator Fourier transform (OFT). The OFT was recently introduced by Cubillos and Jimenez (arXiv:2407.09436) as a building block for fast, direct solution algorithms for the Helmholtz equation in one, two and three-dimensional inhomogeneous unbounded media. A crucial aspect of our methodology is low-rank step truncation, which optimizes the computational process by selectively truncating less significant solution components. This ensures that, for certain problems, the computational load is small without sacrificing accuracy. We also explore the use of different contours for inversion and how they affect the overall accuracy and cost. Throughout we use Summation-By-Parts (SBP) discretization, a technique that ensures numerical stability and accuracy in the discretization of differential operators. We explore the advantages and disadvantages, when using low-rank techniques, of imposing boundary conditions through the so-called simultaneous approximation technique and the projection method.

Daniel Appelo, Jack Kelley
Virginia Tech
appelo@vt.edu, jackkelley@vt.edu

PP1

Stabilizing Physics-Informed Neural Networks for Stiff Differential Equations: Re-Spacing Layer

Approximating solutions for stiff differential equations, which exhibit abrupt changes in specific regions, using physics-informed neural networks (PINNs) is a challenging task. Typically, PINN training involves using a higher density of samples concentrated around areas of rapid change to capture the sharp gradients. However, this approach often leads to data imbalance, which in turn causes slower convergence and diminishes the prediction performance. To address these issues, we introduce Re-spacing layer (RS-layer), a pre-trained encoding layer designed to map the skewed distribution of sampling points onto a uniform distribution while preserving desirable statistical characteristics of the input data for more effective PINN training. We demonstrate that RS-layer enhances PINN training by regularizing the solution gradient in the transformed space. Our numerical experiments validate the effectiveness of our approach, showing that RS-layer not only speeds up convergence but also improves accuracy.

Eunsuh Kim
POSTECH
Department of Mathematics
s2077@postech.ac.kr

Minseok Choi
POSTECH Mathematics Institute, Korea
mchoi@postech.ac.kr

PP1

Layer Potential Methods for Harmonic Functions on Finitely-Connected Tori

We develop layer potential methods to represent harmonic functions on finitely-connected tori. Extending the results for Euclidean domains, we represent the doubly-periodic Green's function using the Jacobi theta function and analyze the resulting single- and double-layer potentials. The cases where the boundary has one connected component and multiple connected components are handled separately. We apply our developed theory to approximate solutions of the Dirichlet and Neumann boundary value problems, as well as the Steklov eigenvalue problem. We implement the developed methods and demonstrate the theory with several numerical examples.

Bohyun Kim
University of Utah
bohyun.kim@utah.edu

Braxton Osting
University of Utah
Mathematics
braxton.osting@utah.edu

PP1

Evaluating and Mapping Demographic Disparities with Traffic Stop Data

According to the Stanford Open Policing Project (SOPP), there are over 50,000 traffic stops conducted in the United States on an average day, yet there is no prescribed means to record and report those stops. This prevents the contextualization of news-worthy events pursuant to traffic stops, such as police violence, within the broader scope of regional and national averages and trends. It also prevents analysis that could either improve trust in public institutions or highlight areas for improvement of policing activities. To help bridge this gap, we use traffic stop records collected by the SOPP along with census records, geospatial data, and traffic measurements to model the number and demographic composition of annual traffic stops at the county level across the U.S. We then use this model to conduct a rolling geographic analysis of expected traffic stops by region in order to identify outliers with respect to both the regional population and incidents of police violence pursuant to traffic stops within the region.

Chad Kite
San Francisco State University
ckite@mail.sfsu.edu

PP1

Compiling Einsum Trees for GPUs

The Einstein summation convention (einsum) encodes operations on one or more input tensors. In practice, given a suitable contraction path, einsums are evaluated in a binary fashion, computing the contraction of two tensors at a time. In contrast to matrix multiplications, general binary tensor contractions that appear in the evaluation of einsum must address high-dimensional data layouts. This work presents a compilation approach for the fast evaluation of einsum trees on GPUs. We show that it is crucial to optimize the data layout of the binary contractions' output tensors, which in turn form the input for subsequent contractions. Furthermore, we use the Triton language and compiler to define and generate tensor contraction kernels. The poster concludes with a presentation of the throughput of our GPU-accelerated einsum tree approach and a

comparison of the obtained runtimes with state-of-the-art tensor contraction routines.

Max Koch
Friedrich Schiller University Jena
m.koch@uni-jena.de

Alexander Breuer
University of California, San Diego
alex.breuer@uni-jena.de

PP1

Fast and Accurate Continuous Fmm Stray Field Evaluation for Fem Micromagnetics

Micromagnetic simulation is an indispensable tool for developing novel magnetic materials and devices. As applications increase in complexity, simulation meshes increase in size leading to extremely long runtimes. The main bottleneck is the stray field interaction, whose direct evaluation cost scales $O(N^2)$ with the mesh size N, although specialized algorithms generally achieve $O(N \log N)$ scaling [Abert, Micromagnetics and spintronics: models and numerical methods, 2019]. For very large meshes, the Fast Multipole Method (FMM) [Greengard, A fast algorithm for particle simulations, 1987] with $O(N)$ scaling promises unparalleled performance, also due to the fact that it has been implemented efficiently on GPUs [Gumerov, Fast multipole methods on graphics processors, 2008]. In this work, we aim to implement a high performance stray field evaluation routine on GPUs for use in our finite-element micromagnetic simulation package magnum.pi [Abert, magnum.fe: A micromagnetic finite-element simulation code based on FEniCS, 2013]. As opposed to other works [Kritsikis, Beyond first-order finite element schemes in micromagnetics, 2014], we compute multipole coefficients directly from the mesh geometry, without introducing point charges. Recently, we completed a CPU implementation which already speeds up simulations in magnum.pi significantly. With GPU parallelism and further optimizations in the future, this will enable simulations of much larger systems than previously possible.

Robert Kraft
University of Vienna
Research Platform MMM
Mathematics-Magnetism-Materials
robert.kraft@univie.ac.at

Claas Abert
University of Vienna
Physics of Functional Materials
claas.abert@univie.ac.at

Dieter Suess
Vienna University of Technology
dieter.suess@univie.ac.at

PP1

A Study on the Effect of Graph Structure Versus Feature Quality for GCN

Graph Convolution Networks (GCN) have become increasingly popular too for node classification. In a GCN the input includes both a graph, which is on a non-Euclidean space and a list of features associated with each node. We investigate which of these two inputs has the most effect on the accuracy of the results. Our experiments on synthetic and real-world graphs demonstrate that the alignment of the graph structure with the distribution of classes is the strongest factor in determining the accuracy of the classification. Based on these observations we provide two complementary algorithms to improve the accuracy (i) by rewiring the graph to match the structure with the distribution of classes, and (ii) by creating a similarity graph using the feature vectors that can be combined with the original graph to increase the effect of the features. Our results show that by combining these two corrections the classification results can significantly improve over state-of-the-art GCN methods. This research will not only help to improve accuracy of any GNN model but also help to alleviate the problem of training of small, noisy datasets by building an effective graph structure for the input data.

Turja Kundu, Sanjukta Bhowmick
University of North Texas
turjakundu@my.unt.edu, sanjukta.bhowmick@unt.edu

PP1

Unifying Nonlinearly Constrained Optimization

Nonlinearly constrained optimization problems arise in a broad rage of applications, including optimal experimental design, the control and operation of the power-grid, and the analysis of experimental campaigns. We introduce the basic building block of iterative solvers for nonlinearly constrained optimization problems. We show that these building blocks can be presented as a double loop framework that allows us to express a broad range of state-of-the-art nonlinear optimization solvers within a common framework. We have implemented this framework in Uno, a modern, lightweight and extensible C++ solver that unifies the workflow of most derivative-based iterative nonlinear optimization solvers. Uno is meant to enable researchers to experiment with novel optimization strategies while leveraging established subproblem solvers and interfaces to modeling languages.

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

PP1

Tail Behavior-Aware Reduced Order Models for Reservoir Simulation

Over-pressurization in subsurface reservoirs increases the risk of leakage and induced seismicity. Standard reduced order models (ROMs) aim to accurately resolve the mean model behavior, but may be inaccurate for distribution tails corresponding to extreme (but important) model behavior. We develop ROM methods capable of capturing such tail behavior. To achieve satisfactory accuracy for distribution tails using the standard proper orthogonal decomposition (POD) requires that tail samples are contained in the snapshots, and that a large number of POD modes is used. We propose two changes to the standard POD by i) using a weighted inner product that emphasizes relevant tail samples; (ii) generating extreme samples to be added as snapshots.

Mingxin Li
Courant Institute
New York University
ml6223@nyu.edu

Georg Stadler

Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

PP1

Data-Driven Solver Selection for Sparse Linear Matrices

Sparse linear systems sit at the core of many computational problems, and their solution strongly correlates to overall execution time and solution quality. However, with the constant increase of the number of linear solver and preconditioner implementations available across a plethora of numerical libraries, choosing the most efficient combination for a given problem (in terms of time-to-solution or another problem-dependent metric) is a challenging task. Indeed, even selecting a numerically stable combination may seem to be an unsurmountable endeavor, especially for a novice user. In this work, we compare classic machine learning approaches to the solver-preconditioner selection problem and develop a novel pipeline through a combination of embedding and linear modelling techniques. We then apply the developed model to choose optimal solvers for a given input matrix across a selection of Krylov solver implementations and preconditioners from the PETSc framework over different datasets. We then use different metrics to analyze the results, since the raw accuracy value can be misleading, given the imbalance in the data. Current results show similar or up to ca. 15% better performance in the NDCG score, as well as reduced variance across the different metrics vs. top black-box approaches on a broad range of systems.

Hayden Liu Weng
Technical University of Munich
h.liu@tum.de

Hans Joachim Bungartz
Technical University Munich
bungartz@tum.de

Felix Dietrich
Technical University of Munich
felix.dietrich@tum.de

PP1

Optimal Experiment Design and Image Reconstruction Using Generative Methods

Reconstructing images from noisy and indirect measurements is an ill-posed inverse problem critical for many medical and geophysical imaging applications. Optimizing the experiment design through slight improvements in the information content or small reductions in the measurement costs is beneficial in many applications. As such, optimizing the measurement process by finding the best measurements is often desirable, though this presents a difficult and intractable problem. As a possible approach, this project investigates the potential of a conditional continuous normalizing flow (CCNF) and learned binary mask model to reconstruct images from noisy, indirect, and optimized measurements. We simulated the process of taking noisy and indirect measurements using the MNIST dataset as a proxy for MRI brain scans due to the comparatively small image sizes of the MNIST dataset. The measurement optimization was performed by applying a learned binary mask to the images. Our model approach used CCNF, in conjunction with an encoder/decoder framework, to reconstruct the original images in the data space and by processing the noisy measurements and conditionals (original quality images) from the latent space. The model produced semi-legible reconstructed images, indicating that while the approach holds potential, the results also highlighted the need for further refinement of the model to improve the quality of image reconstructions.

Spiros K. Manolas
Stony Brook University
spirosmanolas99@gmail.com

Anish Mitagar
University of Massachusetts Amherst
amitagar@umass.edu

Nela Riddle
Indiana University
nmriddle@iu.edu

Nicole Yang
Emory University
tianjiao.yang@emory.edu

PP1

Jexpresso: a General Purpose, Julia Package to Solve Arbitrary Systems of PDEs with Minimal User Intervention

Jexpresso is a new multi-physics, general-purpose solver designed to numerically tackle arbitrary systems of PDEs while streamlining the setup process for users specific physical problems. Although the beta-version presented in the poster primarily employs spectral elements and finite differences, its architecture is designed so that users can incorporate additional grid-based numerical methods without modifying the existing code framework. Jexpresso runs on CPUs and GPUs. It relies on MPI.jl for paralellization and P4est.jl for grid partitioning, load balancing, and adaptive mesh refinement. The user can define a system of PDEs by specifying the vectors of the unknowns, fluxes, and sources; when necessary, incorporating a constitutive law (e.g., an ideal gas equation of state). If the equations are formulated in flux form, then nothing else than the above vectors are necessary from the users perspective. Otherwise, if the equations are expressed in advective form, the user is also required to provide the non-constant matrix of coefficients. The user provides a boundary conditions file specifying the conditions exactly as they might be written on paper; however, the underlying code that processes them remains hidden from view. Standard benchmarks and performance results are shown in the poster. The source code is available on Github at https://github.com/smarras79/Jexpresso.

Simone Marras
New Jersey Institute of Technology
smarras@njit.edu

PP1

Dimensionality Reduction of Large Network Datasets Using Parallel Tempering

Networks play a crucial role in agent-based modeling and in the study of epidemics, critical infrastructure, and various other fields. Recent improvements in data collection have led to an exponential increase in the size and complexity of network datasets. This trend, combined with growing interest in high-fidelity data, poses significant computational challenges for analytics and simulation at large-scales. To overcome these challenges, researchers often use smaller

synthetic datasets. The networks in these datasets are typically generated using random graph models like ErdosRnyi, BarabsiAlbert and Watts-Strogatz, or by applying various dimensionality reduction techniques. However, these methods often fail to replicate the intricate topologies of real-world networks, leading to inaccuracies in simulation results. We propose a graph generating procedure that can produce graphs with specified metrics. Starting with a smaller, randomly generated seed graph, our method applies an algorithm based on parallel tempering to iteratively make modifications to match the desired properties of a real-world target graph. We validate our approach on large-scale network datasets, comparing the effectiveness and computational efficiency of our procedure with other dimensionality reduction methods. Additionally, we examine and assess the dynamics of models simulated on both original networks and their miniatures.

Jorge A. Martinez Ortiz
Michigan State University
martjor@msu.edu

David Butts
Los Alamos National Laboratory
buttsdav@lanl.gov

Nathan Brugnone
Two Six Technologies
nathan.brugnone@twosixtech.com

Samuel Bollman, Michael Murillo
Michigan State University
bollma13@msu.edu, murillom@msu.edu

PP1

Structural and Practical Identifiability Analysis of Models for Syncytia Growth

Syncytia are multinucleated cells that can occur due to virus infection of cells. Mathematical models in the form of ordinary differential equations can be used to simulate the growth syncytia. Several novel ODE models can explain syncytia growth. Before employing these models on actual data, it is essential to analyze their structural(theoretical) and practical identifiability. Structural identifiability is an inherent property of each model and its parameters, referring to our ability to determine parameter values for the model. Practical Identifiability analysis of a model is concerned with accurately determining parameter values given experimental error. Obtaining accurate parameter values allows us to make conclusions about our data within the context of our model that can provide insight into the nature of the spread of syncytia. These two techniques allow us to determine whether or not the parameters of a model are identifiable with the data we plan to collect. Consequentially, we can plan experiments adequately to truly parameterize the data in the contexts of our models to help make better conclusions.

Gabriel S. Mccarthy
Texas Christian University
Department of Physics and Astronomy
G.S.MCCARTHY@tcu.edu

Hana Dobrovolny
Department of Physics and Astronomy
Texas Christian University
h.dobrovolny@tcu.edu

PP1

Applications of the Local Macroscopic Conservative (LoMaC) Low Rank Tensor Method to Collisional Vlasov Dynamics

Numerical simulations of the Vlasov-Poisson system plays a fundamental role in understanding the dynamics of plasma in many applications of science. However, many numerical simulations of the system are often infeasible due to high-resolution requirements and the exponential scaling of computational cost with respect to dimension, while preserving physical invariants of the equation. Our project aims to simulate the Vlasov-Poisson system with collision using the Local Macroscopic Conservative (LoMaC) low rank tensor method which alleviates these challenges of the system. Recently, the LoMaC method [Wei Guo, A Local Macroscopic Conservative (LoMaC) low rank tensor method with the discontinuous Galerkin method for the Vlasov dynamics] was introduced to simulate the Vlasov-Poisson system. In the method, the macroscopic system of the conservation laws are simultaneously evolved using the kinetic flux vector splitting. Then, the low rank solution is projected onto a subspace shared with the macroscopic observables. This method was used to simulate collisionless plasma under the effects of Landau damping. In our project, we will investigate the capabilities of the LoMaC method in simulating the VP system with collision. In this project, we will use the method to observe the VP system (with Landau damping) with three different collision operators: the Bhatnagar-Gross-Krook (BGK) operator, the linear Fokker-Planck operator, and the Landau collision operator.

Austin M. Nelson
Texas Tech Univeristy
nelsonaustin905@gmail.com

Wei Guo
Texas Tech University
weimath.guo@ttu.edu

PP1

Exploring Sensitivity to Uncertainty in Nonlinear Debye Media

We present a study on the time-domain propagation of electromagnetic waves in dielectric materials modeled by a nonlinear Debye medium with random perturbations. Polynomial Chaos Expansions are employed to transform the random nonlinear Debye polarization model into a deterministic framework. We investigate the sensitivity of nonlinear properties to uncertainty, particularly when the amplitude of the input signal is large. Given the challenges in manufacturing where uncertainties can cause optimal parameters to vary and potentially disrupt nonlinear effects, our approach incorporates these uncertainties within the simulation. This can enable the model-based design identification of realizable materials that maintain their desired effects despite variations. The findings from this study contribute to a deeper understanding of wave propagation in complex media, with potential implications for applications in optical communications, material science, and electromagnetic wave control.

Emmanuel E. Oguadimma, Nathan L. Gibson
Oregon State University

oguadime@oregonstate.edu, gibsonn@oregonstate.edu

PP1

Separable Physics-Informed Neural Networks for Solving the BGK Model of the Boltzmann Equation

In this study, we introduce Separable Physics-informed Neural Networks (SPINNs) as a new method for effectively solving the BGK model of the Boltzmann equation. While the mesh-free nature of Physics-informed Neural Networks (PINNs) offers significant advantages for handling high-dimensional partial differential equations (PDEs), applying quadrature rules for accurate integral evaluation in the BGK operator can compromise these benefits and increase computational costs. To address this issue, we leverage the canonical polyadic decomposition structure of SPINNs and the linear nature of moment calculation to significantly reduce the computational expense associated with quadrature rules. However, the multi-scale nature of the particle density function poses challenges for precisely approximating macroscopic moments using neural networks. To overcome this, we fuses SPINNs with Gaussian functions and utilizes a relative loss approach. This modification enables SPINNs to decay as rapidly as Maxwellian distributions, enhancing the accuracy of macroscopic moment approximations. The relative loss design ensures that both large and small-scale features are effectively captured by the SPINNs. The effectiveness of our approach is validated through six numerical experiments, including a complex 3D Riemann problem. These experiments demonstrate the potential of our method to efficiently and accurately tackle challenges in computational physics.

Jaemin Oh
Korea Advanced Institute of Science and Technology
Department of Mathematical Sciences
jaeminoh@kaist.ac.kr

Seung Yeon Cho
Gyeongsang National University
Department of Mathematics
chosy89@gnu.ac.kr

Seok-Bae Yun
Sungkyunkwan University
South Korea
sbyun01@skku.edu

Eunbyung Park
Sungkyunkwan University
Department of Electrical and Computer Engineering
epark@g.skku.edu

YoungJoon Hong
Seoul National University
hongyj@snu.ac.kr

PP1

Mixed Precision Iterative Refinement for Linear Inverse Problems

Many problems in science and engineering give rise to linear systems of equations that are commonly referred to as large-scale linear discrete ill-posed problems. The matrices that define these problems are typically severely ill-conditioned and may be rank deficient. Because of this, the solution of linear discrete ill-posed problems may not exist or be extremely sensitive to perturbations caused by error in the available data. These difficulties can be reduced by applying regularization. We explore the connections between iterated Tikhonov regularization and iterative refinement on the Tikhonov problem in mixed precision using a filter factor analysis.

Lucas Onisk
Emory University
lonisk@emory.edu

James G. Nagy
Emory University
Department of Math and Computer Science
jnagy@emory.edu

PP1

The Better Scientific Software Fellowship Program

As high-performance computers advance, scientists and engineers face increasing complexity in computational models, computer architectures and emerging workflows. In order to succeed now and into the future, software developers need a community where information on the best practices for all aspects of software development flow freely. With this vision, the Better Scientific Software (BSSw) Fellowship Program (https://bssw.io/fellowship) launched in 2018. The BSSw Fellowship Program fosters and promotes practices, processes and tools to improve productivity and software sustainability of scientific codes. Each year, a new class of fellows and honorable mentions are chosen to seed this community of like-minded individuals and highlight their contributions. Our Fellowship alumni serve as leaders, mentors and consultants who share best practices for computational software development and sustainability. This poster highlights the programs current and past Fellows to encourage viewers to benefit from our resources and join in our efforts.

Erik Palmer
Lawrence Berkeley National Laboratory
epalmer@lbl.gov

Elsa Gonsiorowski
Lawrence Livermore National Laboratory
gonsiorowski1@llnl.gov

Mary Ann E. Leung
Sustainable Horizons Institute
mleung@shinstitute.org

PP1

Comparative Analysis of Prominent Ai Models for Enhanced Diabetes Diagnosis Prediction

With a 700% rise in diabetes among youth in the next 35 years predicted, the escalating prevalence of diabetes worldwide calls for modern research methodologies and techniques within the field - specifically, implementing Artificial Intelligence and Machine Learning for diabetes prediction. This research project explores multiple advanced algorithms to predict diabetes, leveraging data analysis using Python libraries to conduct a comparative analysis of seven different algorithmic models: Artificial Neural Network (ANN), Support Vector Machine (SVM), K-Nearest Neighbors (KNN), Decision Tree, Random Forest, XGBoost, and Gaussian Classifier. The objective of this project is to find the most effective AI model for predicting diabetes and to identify the specific health features

that play a pivotal role in its diagnosis. Data was drawn from the Center for Disease Control (CDC), utilizing a pre-balanced dataset composed of nutritional and demographic variables. The analysis focused on metrics such as precision, recall, and accuracy to determine each model's performance - ultimately revealing that the XGBoost and Random Forest models displayed superior performance in terms of accuracy with 4% higher accuracy than other models. With AI being the future of medical care, this research offers insights into the potential for AI-driven tools to enhance predictive diagnostics for patient care.

Rudy Pathak
Clements High School
rudyrip101@gmail.com

PP1

Uncertainty Quantification in Data Fusion of Radiograph Diagnostics

The reconstruction of density from radiographic diagnostics is an ill-posed inverse problem, characterized by instability in the transformation processes. This presentation introduces a Bayesian approach that provides a framework for both accurate mean reconstruction and uncertainty quantification. Additionally, we will explore how incorporating high-accuracy or known data points can further improve the numerical results, ultimately enhancing both the accuracy and credibility of radiographic diagnostics.

Robert M. Pawloski
Arizona State University
rpawlosk@asu.edu

PP1

Approximation of the Compressible Euler Equations in the Low Mach Limit

We study a second order, invariant domain preserving approximation of the compressible Euler equations in the low Mach Regime. This approximation gives rise to an explicit, consistent, and conservative scheme that satisfies the necessary entropy inequalities to have physically relevant solutions. The scheme works well in the supersonic regime but suffers in the low Mach regime where stiffness in the pressure term results in a very restrictive cfl condition. We see that as the Mach number decreases, the number of time steps required to reach a target time drastically increases.

David Pecoraro, Jean-Luc Guermond, Matthias Maier
Texas A&M University
pecoraro27@gmail.com, guermond@tamu.edu, maier@tamu.edu

PP1

Real-Time Aerodynamic Sensing for Hypersonics Via Strain-Based Inverse Maps with Optimal Sensor Placement

We develop a real-time sensing strategy to infer aerodynamic surface loads from sparse measurements of the structural strain. This approach targets hypersonic flight conditions, where the harsh aerothermal environment inhibits direct measurement of the aerodynamic quantities necessary for guidance, navigation, and control. The inference problem is formulated as a least-squares problem with a linear constraint arising from the discretization of the governing elasticity partial differential equation. The resulting inverse map, or estimator, can be pre-computed to enable real-time estimates of the surface pressure and corresponding integrated quantities. This inverse approach enables explicit uncertainty quantification, which provides insights for optimal sensor placement. For a limited number of spatially distributed strain sensors, we seek to minimize the uncertainty in the estimated quantities of interest in the presence of sensor noise. Specifically, we formulate an optimization problem to identify the sensor configuration that minimizes the desired information criteria of the estimator covariance matrix. We additionally propose sensor selections considering robustness to sensor failures. Numerical studies demonstrate the estimator performance and sensor selection for the Initial Concept 3.X (IC3X) conceptual hypersonic vehicle.

Julie Pham
University of Texas, Austin
jvpham@utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@oden.utexas.edu

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

PP1

Simulating Subsurface Flow with Differentiable Programming

Differentiable programming is a computational approach that enables automatic differentiation, a crucial piece of machine learning on more general computer codes, including physics simulators. Unfortunately, current differentiable programming codes are limited to single-machine programming and do not scale to parallel computing on high-performance architectures. We have developed a differentiable programming framework for extreme-scale scientific computing in the Julia ecosystem with a focus on geophysics applications such as subsurface flow and transport, reservoir pressure management, and carbon sequestration. This Julia framework combined with MPI and PETSc enables a seamless integration of huge, complex scientific simulations with machine learning tools. Beyond our machine learning goals, this project will dramatically lower the barrier for entry to use advanced high-performance computing methods for, e.g., inverse analysis and uncertainly quantification that rely on adjoint methods, opening new doors for carbon sequestration decision support. Doing so will open new doors at the intersection of differentiable programming and massively parallel partial differential equation solvers.

Daniel O'Malley, Harun Rashid, Aleksandra Pachalieva
Los Alamos National Laboratory
omalled@gmail.com, hrashid@lanl.gov, apachalieva@lanl.gov

PP1

Sensitivity Analysis for Storm Surge Modeling and Prediction

Predicting storm surge elevation and velocity with certainty is crucial for assessing coastal hazards and understanding potential impacts of climate change. The accuracy of said predictions is influenced by a number of inputs that can vary spatially and temporally, such as sea surface

height, bottom surface roughness, and wind surface stress. Understanding uncertainties associated with these inputs is critical for making accurate forecasts and can help models better estimate storm surge. Sensitivity analysis is an important first step toward elucidating how predictions are impacted by input uncertainties and reducing uncertainty in surge models overall. Here, we explore sensitivity analysis methods and how they can serve to improve computational modeling of storm surges.

Haley Rosso, Talea Mayo
Emory University
haley.rosso@emory.edu, talea.mayo@emory.edu

Rebekah White
Sandia National Laboratories
rebwhit@sandia.gov

PP1

An Optimal O(N) Helmholtz Solver for Complex Geometry Using WaveHoltz and Overset Grids

An efficient and high-order accurate solver for Helmholtz problems on complex geometry is described. The scheme is based on the WaveHoltz algorithm which computes time-harmonic solutions of the Helmholtz equation by time-filtering solutions of an associated initial-boundary-value problem for wave equation. Problem domains with complex geometric configurations are treated with overset grids. The wave equation is solved efficiently with implicit time-stepping using as few as five time-steps per period, independent of the mesh size. When multigrid is used to solve the implicit time-stepping equations, the cost of the resulting WaveHoltz scheme for overset grids scales linearly with the number of grid points N (at fixed frequency) and is thus optimal in CPU-time and memory usage as the mesh is refined. Numerical results are given for problems in two and three space dimensions, and using second and fourth-order accuracy in space, and they show the potential of the approach to solve a wide range of large-scale problems.

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

PP1

A Kernel Based SIR Model for Infectious Spread in Space and Time

In this study, we consider an integro-differential model designed to describe the spatial spread of an epidemic. Building on the classical SIR model, our approach leverages similar probabilistic arguments to define infection rates while still producing a diffusive-like spread. A key feature of the model is the flexibility it offers through the selection of different kernels, enabling precise control over the dynamics of disease propagation. Like a generalized elliptic operator, the kernel can generate anisotropic spread and is allowed to vary through space and time. Through simulations, we illustrate various infection spread patterns and explore the qualitative behavior of the model under different scenarios.

Alex Sheranko, Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
asheran1@umbc.edu, draga@umbc.edu

PP1

Analyzing Dynamical Differences Between Vaccinated and Unvaccinated Rsv Patients

Abstract: Respiratory Syncytial Virus (RSV) is a common respiratory virus that can cause serious illness in infants and the elderly. Vaccines for RSV have recently been introduced and have been shown to reduce the severity of the disease, although there has been limited examination of how viral dynamics differ between vaccinated and unvaccinated individuals. Here, we use data from the MVA-BN-RSV Phase II vaccine study to quantify the differences in dynamics between those who have been vaccinated versus those who were unvaccinated. We use an ordinary differential equation model to fit the data, finding that vaccinated patients experience a higher variance in response as compared to the placebo group. This result could have implications for the use of MVA-BN-RSV in vaccination.

Arjan Suri
Texas Christian University
arjan.suri17@gmail.com

Hana Dobrovolny
Department of Physics and Astronomy
Texas Christian University
h.dobrovolny@tcu.edu

PP1

Examining Complex Disease Spread Dynamics in Schools and Workplaces Using An Exascale Agent-Based Framework

Schools often serve as critical hubs for the spread of infectious diseases, with children acting as transmission vectors. Epidemiological models are indispensable for understanding disease dynamics within these environments and informing effective public health decisions. To examine critical factors affecting disease propagation in schools, we use ExaEpi, a high-performance agent-based simulation framework for exascale computing. Built on AMReX, a scientific computing library tailored for solving Eulerian and particle-based models on heterogeneous architectures, ExaEpi enables the incorporation of intricate interactions within schools (e.g., teacher-student) or workplaces (e.g., sub-department) across extensive populations and regions, such as California. It also features a range of public health interventions, including quarantines, remote working, and school dismissals, allowing for a detailed examination of how closures can significantly alter transmission patterns. Additionally, ExaEpi models multiple diseases concurrently, offering insights into how agent behavior adapts to threats and impacts transmission patterns. By simulating concurrent outbreaks and evaluating interventions, ExaEpi helps officials develop adaptive strategies, refine measures, and optimize resources. Its a powerful tool for public health planning, especially in managing complex disease scenarios. Prepared by LLNL under contract DE-AC52-07NA27344. LLNL-ABS-868593

Gbocho Masato Terasaki
UC Merced
gterasaki@ucmerced.edu

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Livermore, CA

ghosh5@llnl.gov

Ali Navid
LLNL
navid1@llnl.gov

PP1

Daggermpi: A Straight-Forward Approach to the Mpi Standard with Data Dependencies in Julia

The poster will present an implementation of the MPI extension for the Julia package Dagger, focused on optimizing dense linear algebra computations through a tile-based approach. Dagger is a task-based native runtime system that transparently manages task execution and communication, offering an unified Application Programming Interface (API) across distributed environments and architectures without requiring modifications to user code. By leveraging Dagger's capabilities, our implementation efficiently handles the inherent complexities of task scheduling and data movement. The poster will highlight the advantages of using Dagger for tile-based algorithms, emphasizing its ease of use and flexibility in diverse computing environments. Through benchmark comparisons and case studies, we demonstrate how this approach achieves high performance and scalability, making it an ideal solution for modern computational challenges in scientific and engineering applications. The main goal of the poster is to show how the abstraction layer, called datadeps, enables the user to leverage the MPI standard seamlessly by just annotating the data dependencies of function arguments, never having to call MPI routines directly.

Felipe A. Tome, Rabab Alomairy
Massachusetts Institute of Technology (MIT)
felipe0@mit.edu, rabab.alomairy@mit.edu

Julian Samaroo
Massachusetts Institute of Technology
jpsamaroo@gmail.com

Alan Edelman
Department of Mathematics
Massachusetts Institute of Technology and Julia Computing
edelman@mit.edu

PP1

A Time-Frequency Method for the Acoustic Wave Equation

Time-frequency methods for solving the acoustic wave equation on an exterior domain have many advantages over alternative direct-in-time approaches. However, these methods can be expensive and challenging to use in domains with trapping regions. The slow decay of the trapped waves severely reduces the region of analyticity of the solution in frequency space while simultaneously making the inverse Fourier transform integral highly oscillatory. Our approach improves analyticity properties through a complex deformation. Afterwards, the resulting oscillatory integral is handled through a fast sinc transform method. This combination results in a highly accurate and computationally efficient method for solving the wave equation, even when decay over time is slow.

Wietse Vaes
University of Washington
wietsev@uw.edu

Heather D. Wilber
University of Washington, U.S.
hdw27@uw.edu

Gunnar Martinnson
University of Texas
pgm@oden.utexas.edu

Abinand Gopal
UC Davis
gopal@ucdavis.edu

PP1

Multiscale Neural Networks for Approximating Green's Functions

Solving Partial Differential Equations (PDEs) using neural networks has been widely applied in the fields of physics, biology, and engineering. Learning Green's functions is an effective method for solving PDEs with the same differential operator. However, Green's functions are challenging to learn due to their poor regularity, which necessitates larger neural network sizes and longer training times. In this poster, we focus on learning Green's functions using multiscale neural networks. Based on theoretical analysis and experiments, we find that the multiscale approach can reduce the required neural network size and increase the training speed.

Tianshi Xu
Emory University
xuxx1180@gmail.com

Yahong Yang
Penn State University
yxy5498@psu.edu

Wenrui Hao
Pennsylvania State University
wxh64@psu.edu

Rui Peng Li
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

Yuanzhe Xi
Department of Mathematics
Emory University
yuanzhe.xi@emory.edu

PP1

Block-Structured Operator Inference for Coupled Multiphysics Model Reduction

We present a block-structured extension of the operator inference model reduction method. The governing equations for coupled dynamical systems often exhibit a block structure that can be exploited by separating operator inferences least squares inference step into subproblems, enabling the imposition of distinct model forms for each physics regime. This separation also permits us to tailor the regularization of each operator block independently where necessary. Block-structured reduced-order models thus provide the potential for computational savings in both the offline learning and online prediction phases, along

with improved robustness of the least squares calculations. Numerical experiments for high-fidelity coupled aerostructural dynamics demonstrate that block-structured operator inference can improve online prediction time while preserving comparable accuracy to monolithic operator inference.

Benjamin Zastrow
University of Texas at Austin
bzastrow@utexas.edu

Anirban Chaudhuri
The University of Texas at Austin
achaudhuri@utexas.edu

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

PP1

Boosting the Efficiency of the Differential Algebra-Based Fast Multipole Method Operators Using Symbolic Calculation

The Fast Multipole Method (FMM) calculates the pairwise interaction between particles with an efficiency that scales linearly with the number of particles. The strategy involves grouping particles based on their location and density and approximating interactions using expansions. Differential Algebra (DA) calculates the high-order derivatives of a function at a given point, providing a convenient way to construct expansions in the FMM. However, DA-based FMM operators have relatively low efficiency compared to their peers using other mathematical techniques. In this report, we demonstrate that the efficiency of DA-based FMM operators can be significantly improved using symbolic calculation. To construct an FMM operator, the same process that consists of a series of DA calculations needs to be repeated numerous times for different initial conditions. In numerical calculations, the initial conditions are often lost in the obtained operator. Using symbolic calculation, we derive explicit expressions that show how the operator depends on the initial conditions by performing the process once. These expressions can then be used to construct the operator for different initial conditions, avoiding time-consuming DA calculations. We have developed a C++ library for symbolic DA calculation. We will present how we obtain the symbolic DA-based FMM operators and the extent to which efficiency is boosted with numerical examples.

He Zhang
Jefferson Lab
hezhang@jlab.org

PP1

Persistent Directed Flag Laplacian (pdfl)-Based Machine Learning for ProteinLigand Binding Affinity Prediction

Directionality in molecular and biomolecular networks plays a significant role in the accurate representation of the complex, dynamic, and asymmetrical nature of interactions present in protein-ligand binding, signal transduction, and biological pathways. Most traditional techniques of topological data analysis (TDA), such as persistent homology (PH) and persistent Laplacian (PL), overlook this aspect in their standard form. To address this, we present the persistent directed flag Laplacian (PDFL), which incorporates directed flag complexes to account for edges with directionality originated from polarization, gene regulation, heterogeneous interactions, etc. This study marks the first application of the PDFL, providing an in-depth analysis of spectral graph theory combined with machine learning. Besides its superior accuracy and reliability, the PDFL model offers simplicity by requiring only raw inputs without complex data processing. The multi-kernel PDFL model was validated for its scoring power against other state-of-art methods on three popular benchmarks, namely PDBBind v2007, v2013, and v2016. Computational results indicate that the proposed PDFL model outperforms competitors in protein-ligand binding affinity predictions, indicating that PDFL is a promising tool for protein engineering, drug discovery, and general applications in science and engineering.

Mushal Zia
Michigan State University
ziamusha@msu.edu