MEDIA PROCESSING AND A MODIFIED WATERMARKING SCHEME BASED ON THE SINGULAR VALUE DECOMPOSITION *

KATHERINE KEEGAN†, DAVID MELENDEZ‡, AND JENNIFER ZHENG§

Abstract. We introduce the singular value decomposition (SVD) along with some of its key properties, illustrate its utility in image processing and audio processing, and show how this relates to watermarking and digital ownership protection. After establishing experimental results regarding image processing, we propose a modified version of a watermarking scheme introduced in *Jain, Arora, and Panigrahi, A reliable SVD based watermarking scheme, CoRR abs/0808.0309 (2008)* which offers improved robustness and imperceptibility properties.

Key words. singular value decomposition, linear algebra, media processing, digital ownership protection, watermarking

1. Introduction.

In the 21st century, most media consumption and transmission occurs digitally. This introduces a major challenge: how can one prove their ownership of a piece of digital media, such as a song, photograph, or book? Luckily, the nature of digital information requires that such media is usually stored as a data matrix of numerical values, which allows for the employment of mathematical methods in digital ownership protection. To this end, we introduce the concept of watermarking schemes, which allows one to embed identifying ownership information into a piece of media with little to no perceptible change in the media.

To evaluate these schemes, we have some criteria which must be ensured. First, we want our scheme to be imperceptible: the embedding of a watermark should not leave any significant observable impact on the media. We also want our scheme to be secure: a third party should not be able to manipulate our scheme and try to establish themselves as the false owner of our media. Finally, we want to maintain some robustness against distortions: ideally, our watermarked media could undergo some kind of compression, distortion, or even an intentional attack, and we could still extract our watermark and reasonably prove our ownership. We evaluate the performance of three schemes according to these criteria, including two existing watermarking schemes based upon the SVD and our proposed modification of one of these schemes.

The code used to produce the results in this paper can be found in the GitHub repository at https://github.com/S-I-SVD/Randomized-SVD.

1.1. SVD background.

The singular value decomposition, or the SVD, was first discovered independently by Eugenio Beltrami and Camille Jordan in the 1870s as they were tackling problems related to bilinear forms in linear algebra. Since then, it has become one of the most useful tools in linear algebra, seeing applications in widely disparate fields.

The function of the SVD is to factorize a matrix into a product of three matrices, which can then be decomposed further into a sum of rank-1 matrices. The SVD is closely related to the eigenvalue decomposition (EVD): given a symmetric real $n \times n$

*Submitted to the editors 04/13/2021

Funding: This work was done during the summer@ICERM2020 program at the Institute for Computational and Experimental Research in Mathematics at Brown University. All authors contributed equally to this work.

†Mary Baldwin University, Staunton, VA (katie@katiekeegan.org).

‡University of Central Florida, Orlando, FL (davidmelendez@Knights.ucf.edu).

§Emory University, Atlanta, GA (jenniferzyy1012@gmail.com).

matrix X (i.e. for which $XX^\top = X^\top X$), there exists an orthogonal matrix P and a diagonal matrix D such that $X = PDP^{-1}$. In the EVD, the columns of P are eigenvectors of X that form an orthonormal basis for $\mathbb{R}^n$, and the diagonal entries of D contain the corresponding eigenvalues [9]. However, the EVD can only be performed on diagonalizable symmetric square matrices. To generalize the EVD, we have the SVD. Note that although the SVD exists for complex matrices, we will be focusing only on real matrices in our paper.

Let A be an $m \times n$ matrix with rank r. The definition of the SVD tells us that A can be decomposed into three matrices: U, an $m \times m$ orthogonal matrix; V, an $n \times n$ orthogonal matrix; and Σ, a diagonal matrix; such that $A = U\Sigma V^\top$. This Σ matrix consists of ordered positive diagonal entries $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_r$, with all of the remaining diagonal entries being zero. These σ entries are known as the singular values, and the columns of U and V are called the left and right singular vectors, respectively.

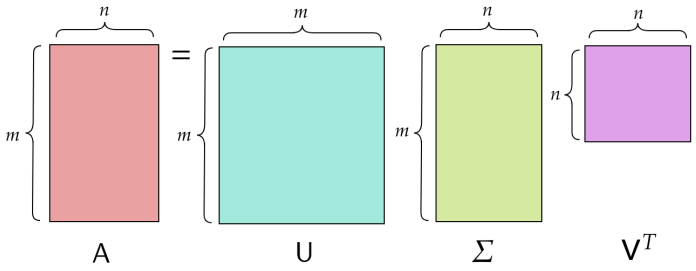

Figure 1.1: SVD Matrices

The singular values of A are the square roots of each of the eigenvalues of $AA^\top$ or $A^\top A$. The left singular vectors are the eigenvectors of $AA^\top$, while the right singular vectors are the eigenvectors of $A^\top A$.

As briefly mentioned earlier, one of the most valuable properties of the SVD is that it is applicable to all matrices. This means that our matrix A can be any size or represent any kind of information, and we would still be able to obtain its fundamental SVD matrices. Based on this fact, we can apply SVD techniques to any data that can be represented as a matrix, which we will show in Section 2 is of great utility for digital ownership protection.

An important application of the SVD is in approximating matrices. For an $m \times n$ matrix A with rank r, we take each left and right singular vector and the corresponding singular value to factor A as

$$A = U\Sigma V^\top = \sigma_1 u_1 v_1^\top + \sigma_2 u_2 v_2^\top + \sigma_3 u_3 v_3^\top + \cdots + \sigma_r u_r v_r^\top \tag{1.1}$$

using the SVD, where $u_1, u_2, \ldots, u_r$ are the columns of U and $v_1, v_2, \ldots, v_r$ are the columns of V. Let A_k be the sum of the first k terms in (1.1) for $k \leq r$.

We define $\|A\|_2$ to be the ℓ^2 norm, also called the spectral norm of the matrix A, where

$$\|A\|_2 = \max_{x \neq 0} \frac{\|Ax\|_2}{\|x\|_2} = \sigma_1.$$

We define $\|A\|_F$ to be the Frobenius norm of the matrix A [13], where

$$\|A\|_F = \sqrt{\sigma_1^2 + \cdots + \sigma_r^2}.$$

The Eckart-Young Theorem states that for any B with rank k, we have

$$\|A - B\|_F \geq \|A - A_k\|_F .$$

This also holds true for

$$\|A - B\|_2 \geq \|A - A_k\|_2 .$$

In other words, A_k is the closest rank k approximation to A. This means that any low-rank approximation for a matrix found using the SVD would be the best approximation possible at that particular rank [12].

1.2. Image Processing. Before discussing how we can utilize the SVD to protect ownership of media, we must first introduce how the SVD is applied to media processing.

A grayscale image of size $m \times n$ is represented by an $m \times n$ array of luminosity values and can be operated on directly. A color image is represented by an $m \times n \times c$ array, where c is 3 or 4. This can be thought of as an $m \times n$ array of c-tuples representing the red, green, and blue, and potentially alpha components of each pixel. Such a color image is represented by a $3m \times n$ matrix, produced by "stacking" the color channels.

A video can be thought of as an array of k images, each of which is represented by an $m \times n \times c$ array, where c is the number of color channels. Such a video can be represented as a $cmn \times k$ matrix, where each of the k columns is a "flattened" image.

Both of these procedures are visualized in Figure 1.2. For more information about the matrix representation of images and videos as well as the usage of the SVD in media processing, please see [6] or [4].

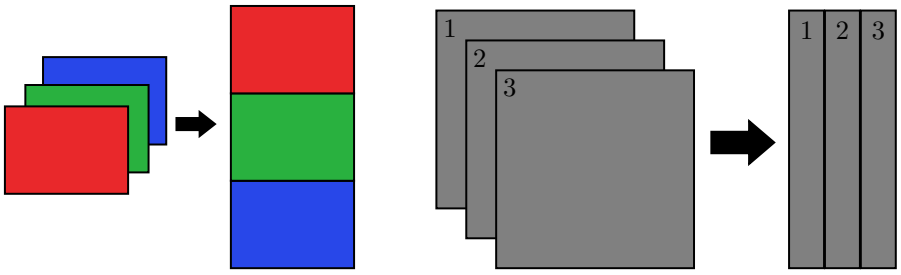

(a) Color image (b) Video

Figure 1.2: Matrix representation of images and videos

The Eckart-Young Theorem provides us with a remarkably simple image compression algorithm, in which one only needs to compute a low-rank approximation of the image. The only complication here is that one must choose the rank of the approximation, and the desirable choice varies with the complexity of the image. Since the error in a rank k approximation A_k of a matrix A depends on the singular value distribution of $A - A_k$, or equivalently the singular values $\sigma_{k+1}, \sigma_{k+2}, \ldots, \sigma_r$ of A, one may use the singular value distribution of A to choose an appropriate approximation rank.

The images that we are interested in compressing are rarely that simple. However, many images can be approximated surprisingly well by a low-rank compression using the SVD. In the language of linear algebra, we say that many images have very low intrinsic rank.

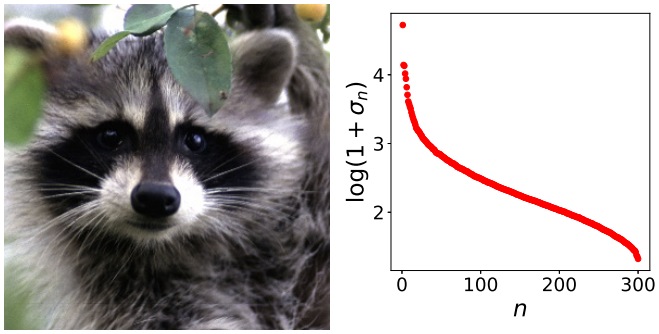

(a) Original Image [5] (b) Singular Value Log Plot

(c) Rank 5 (d) Rank 10 (e) Rank 15 (f) Rank 50

Figure 1.3: Examples of Image Compression

Figure 1.3 provides a clearer illustration of this. The original image is of rank 300. However, we see that a rank 50 approximation computed using the SVD is almost indistinguishable from the original. Most importantly, each reconstruction is simply the sum of a certain number of rank-1 matrices: for example, Figure 1.3c is the sum of 5 rank-1 matrices, and so on for Figures 1.3d-1.3f. The relative weight that each rank-1 matrix contributes to the reconstruction of the original image is plotted using the log plot in Figure 1.3b, where n refers to the nth singular value. This plot shows that the first several rank-1 matrices will contain the most dominant aspects of the original image, which we can visually confirm with the examples.

1.3. Singular Value Modification. We have seen how image compression using the SVD works. Knowing this, what would happen if we changed the SVD matrices somehow? Specifically, how would uniformly applying some mapping to the singular values, such as multiplying or adding by a scalar, impact a reconstructed image? Multiplication of the singular values by some scalar would just be equivalent to multiplying the entire matrix by that scalar. However, adding a scalar to each singular value will influence the resulting image while still preserving some characteristics, as seen in Figure 1.4 [1].

[1] Landscape image from personal collection

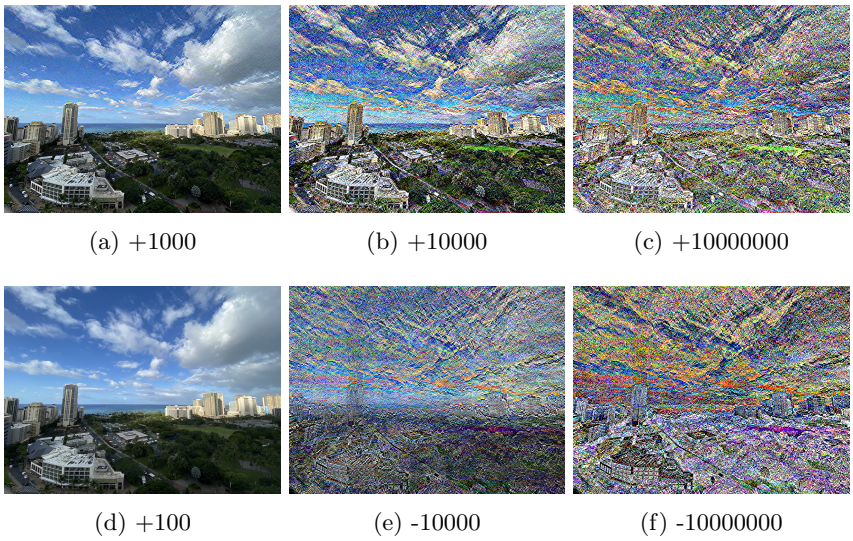

(a) +1000 (b) +10000 (c) +10000000

(d) +100 (e) -10000 (f) -10000000

Figure 1.4: Modifying Singular Values: Addition of a Scalar to each singular value

In this example, we see that adding relatively large scalars, such as 100, to each singular value does not leave a noticeable impact on the image, while adding an even larger scalar, such as 100000, can essentially result in the entries going out of range, with few intelligible trace of the original image in the reconstruction. The sign of the scalar does not seem to have a significant effect on the nature of the reconstruction. The reconstruction of singular values by adding in values leaves an irregular layer of vibrantly colored pixels as artifacts.

Note that the effect of adding a scalar to the singular values is dependent on the magnitude of the singular values themselves, and that the qualitative results will vary highly depending on the distribution of the singular values specific to a given matrix.

The most important result from studying this mapping is that smaller scalars have a minimal impact on the reconstruction in comparison to larger scalars, which leave noticeable artifacts. This will explain the choice of scaling value for watermark embedding in Section 2.

In all of these cases, we notice that when the singular values are modified, many of the dominant characteristics of the original image are preserved, while some are altered depending on the respective mapping. This property of singular value modification will be crucial in the successful embedding of a watermark, as we will see later on.

2. Watermarking Background. With an understanding of the SVD and its use in computing low-rank approximations, we can now explore SVD-based digital ownership protection methods. In particular, various watermarking schemes have been developed that utilize the properties of the SVD matrices to "hide" some watermark within the original data. As mentioned in 1, some primary considerations in evaluating such schemes include perceptibility, security, and robustness against distortions.

A watermarking scheme is given by an embedding function $\mathcal{E}$ and an extraction function $\mathcal{X}$. The function $\mathcal{E}$ takes a data matrix A, a watermark matrix W, and a scaling factor α, and returns the watermarked matrix A_W along with a key K. The

scaling factor α determines the "intensity" with which to embed W into A. As we will see in section 4, increasing α makes watermarks more visible but also more robust against distortions.

2.1. Watermarking Security Evaluation. Before proceeding further into an analysis of watermarking schemes, we first elaborate upon this security criterion. In order for a watermarking scheme to be employable in proof of ownership, it is vital that given a watermarked image A_W, an adversary cannot produce a phony original image A_p and watermark W_p such that the watermark W_p can be extracted from A. This process was utilized in [8], and we use it to evaluate the security for each of the watermarking schemes we analyze here. A basic test of security, then, is given by the procedure described in Algorithm 2.1, using an image A, two watermarks W_1, W_2, a user chosen scalar α, a watermark embedding function $\mathcal{E}$, and an extraction function $\mathcal{X}$.

Algorithm 2.1 Watermarking Security Test

1: $A_{W_1}, K_1 \leftarrow \mathcal{E}(A, W_1, \alpha)$
2: $A_{W_2}, K_2 \leftarrow \mathcal{E}(A, W_2, \alpha)$
3: $\widetilde{W_2} \leftarrow \mathcal{X}(A_{W_1}, K_2)$
4: Compare $\widetilde{W_2}$ to W_2

Examples of $\mathcal{E}$ and $\mathcal{X}$ are given in section 2.2 and section 2.3.

2.2. The Liu & Tan Watermarking Scheme. The Liu & Tan watermarking scheme, described in great detail in [11], is a widely-cited and foundational scheme in SVD-based watermarking techniques. This method involves embedding a watermark into a matrix's singular values and then computing the SVD of this new matrix as a means of embedding information, as is outlined in the following algorithm:

Algorithm 2.2 Liu & Tan: Embedding $(A_W, K) = \mathcal{E}(A, W, \alpha)$

1: Take the SVD of data matrix: $A \rightarrow USV^\top$
2: Add αW to S and compute the SVD of their sum: $S + \alpha W \rightarrow U_W S_W V_W^\top$
3: Replace S with S_W to reconstruct watermarked matrix: $A_W \leftarrow US_W V^\top$
4: Save the keys: $K \leftarrow (S, U_W, V_W, \alpha)$
5: **return** (A_W, K)

This scheme relies on replacing the singular values of A with the singular values of $S + \alpha W$ in order to reconstruct a watermarked piece of media.

In Section 3, we visually see how when the singular values are slightly modified in some way, the major characteristics of an image are still preserved. The final step of the Liu & Tan embedding process is essentially taking advantage of this useful aspect of the SVD.

The Liu & Tan scheme requires knowledge of U_W, S, V_W, and α, which can be thought of as our keys to prove ownership.

To attempt to extract our watermark from a (possibly distorted) watermarked image $\widetilde{A}_W$, we perform the following steps: The Liu & Tan watermarking scheme changes the singular value spectrum of the watermarked image while leaving the principal components unaffected. Because of this, embedding a watermark in an image only enhances some of the already-present features of the image, making the

Algorithm 2.3 Liu & Tan: Extraction $\widetilde{W} = \mathcal{X}(\widetilde{A}_W, K)$

1: $(S, U_W, V_W, \alpha) \leftarrow K$
2: Take the SVD of $\widetilde{A}_W$: $\widetilde{A}_W \rightarrow \widetilde{U}\widetilde{S}_W\widetilde{V}^\top$
3: Construct $\widetilde{D}$: $\widetilde{D} \leftarrow U_W\widetilde{S}_W V_W^\top$
4: Reconstruct watermarked matrix: $\widetilde{W} \leftarrow \frac{1}{\alpha}(\widetilde{D} - S)$
5: **return** $\widetilde{W}$

watermark remarkably imperceptible for reasonably high values of α. Using the image and watermark shown in Figure 2.1, this is demonstrated in Figure 2.2.

(a) Original image (b) Watermark [2]

Figure 2.1: Raccoon and fox images used in experiments

(a) $\alpha = 1$ (b) $\alpha = 0.5$ (c) $\alpha = 0.25$ (d) $\alpha = 0.1$

Figure 2.2: Images watermarked using Liu & Tan scheme [11]

In Figure 2.3, we test the Liu & Tan watermarking scheme's security using the basic security test outlined in Section 2.1 (Algorithm 2.1). As noted in [14], we observe a major risk.

(a) Phony watermark (W_2) [15] (b) Watermarked image (A_{W_1}) (c) Phony extracted ($\widetilde{W}_2$)

Figure 2.3: Basic security test for Liu & Tan watermarking scheme [11], using Algorithm 2.1

As we can see in this example, a phony watermark is able to be extracted from the watermarked image with near perfect accuracy. This essentially means that in this scheme, anyone could claim ownership by extracting their own watermark from the media. For example, someone could claim the rights to an image by claiming that their phony watermark was the correct one, and since any watermark can be extracted almost perfectly from the watermarked image, there is no way for the true owner to indisputably prove their ownership. Naturally, this renders the Liu & Tan watermarking scheme unusable as a reliable method of digital ownership protection.

2.3. The Jain et al. Watermarking Scheme. In [8], Jain et al. propose a modification to the Liu & Tan watermarking scheme, where instead of embedding the entire watermark in the singular value matrix of an image, we only embed the principal components of the watermark.

Algorithm 2.4 Jain et al.: Embedding $(A_W, K) = \mathcal{E}(A, W, \alpha)$

1: Take the SVD of data matrix: $A \to USV^\top$
2: Take the SVD of watermark matrix: $W \to U_W S_W V_W^\top$
3: Add principal components to singular values of data matrix: $S_1 \leftarrow S + \alpha U_W S_W$
4: Replace S with S_1 to reconstruct watermarked matrix:
4: $A_W \leftarrow US_1V^\top$
5: Save the keys: $K \leftarrow (A, V_W, \alpha)$
6: **return** (A_W, K)

The corresponding extraction algorithm proposed by Jain et al. is listed as below.

Algorithm 2.5 Jain et al.: Extraction $\widetilde{W} = \mathcal{X}(\widetilde{A}_W, K)$

1: $(A, V_W, \alpha) \leftarrow K$
2: Subtract A from $\widetilde{A}_W$: $D = \widetilde{A}_W - A = \alpha U U_W S_W V^\top$
3: Reconstruct watermarked matrix: $\widetilde{W} \leftarrow \alpha^{-1} U^\top D V V_W^\top$
4: **return** $\widetilde{W}$

As in Figure 2.2, we test various values of α for embedding the fox image into the raccoon image using the Jain et al. watermarking scheme and display the results in Figure 2.4.

(a) $\alpha = 1$ (b) $\alpha = 0.5$ (c) $\alpha = 0.25$ (d) $\alpha = 0.1$ (e) $\alpha = 0.01$

Figure 2.4: Images watermarked using Jain et al. scheme [8]

Figure 2.4 shows that the Jain et al. scheme requires very low values of α in order to retain the qualities of the original image, and we will see later that it is also less robust to distortions. However, it does have a major advantage over the Liu & Tan watermarking scheme: a much-improved level of security. Here, as in Figure 2.3, we perform the basic security test for the Jain et al. watermarking scheme. As we can see in Figure 2.5, extracting the incorrect watermark from the image (employing the security test described in Algorithm 2.1 with the Jain embedding and extraction functions) produces an image that barely resembles the watermark.

(a) Watermarked image (b) Watermark extracted (c) Phony watermark extracted

Figure 2.5: Basic security test for Jain et al. watermarking scheme [8] as in Figure 2.3, using Algorithm 2.1

The idea of the Jain et al. scheme is that in contrast with the Liu & Tan scheme, in which a scaled version of the entire W is added into S, we are instead embedding a scaled version of only partial information about W, i.e. its principal components $U_W S_W$. Thus, someone trying to prove ownership would need the complementary information V_W in order to correctly extract the watermark.

3. Our Modified Watermarking Scheme. When considering the matrix operations involved in Jain et al. embedding process, we can see that

$$\begin{aligned} A_W &= US_1V^\top \\ &= U(S + \alpha U_W S_W)V^\top \\ &= USV^\top + \alpha U U_W S_W V^\top \\ &= A + \alpha U U_W S_W V^\top. \end{aligned} \tag{3.1}$$

Despite the Jain et al. scheme's improved security, we see that this scheme is equivalent to just adding a scalar multiple of the matrix $UU_WS_WV^\top$ to the original matrix A. This tells us that the perceptibility of the watermark in an image will increase linearly with α.

This additive term, $\alpha UU_WS_WV^\top$, presents us with something to modify and hopefully improve. It turns out that by replacing this term with $\alpha U_WS_WV^\top$, we can construct a new watermarking scheme that preserves some desirable properties (security and robustness to distortions) of the Jain watermarking scheme while also increasing imperceptibility of the embedded watermark. Again, the formulation of this modified scheme as being a matrix added into an image means that the perceptibility of the watermark in the image will be directly proportional to the α we choose. From now on, we will refer to this new process as our Modified Jain watermarking scheme. The embedding and extraction processes are outlined below.

Algorithm 3.1 Modified Jain: Embedding $(A_W, K) = \mathcal{E}(A, W, \alpha)$

1: Take the SVD of cover matrix: $A \to USV^\top$
2: Take the SVD of watermark matrix: $W \to U_WS_WV_W^\top$
3: Add additive term to cover matrix: $A_W = A + \alpha U_WS_WV^\top$
4: Save the keys: $K \leftarrow (A, V, \alpha)$
5: **return** (A_W, K)

Algorithm 3.2 Modified Jain: Extraction $\widetilde{W} = \mathcal{X}(\widetilde{A}_W, K)$

1: $(A, V_W, \alpha) \leftarrow K$
2: Use keys to reconstruct watermark from potentially distorted $\widetilde{A}_W$:
$\widetilde{W} \leftarrow \frac{(\widetilde{A}_W - A)VV_W^\top}{\alpha}$
3: **return** $\widetilde{W}$

Figure 3.1 shows the results of embedding the fox watermark in the raccoon image using various scaling factors α. We notice that for this modified scheme, $\alpha = 0.1$ or $\alpha = 0.25$ will suffice in maintaining a reasonable degree of imperceptibility, and even slightly larger values are still not too noticeable.

(a) $\alpha = 1$ (b) $\alpha = 0.5$ (c) $\alpha = 0.25$ (d) $\alpha = 0.1$

Figure 3.1: Images watermarked using the modified Jain watermarking scheme

When we apply the basic security test described in Algorithm 2.1 with the modified Jain embedding and extraction functions, we obtain the following results:

(a) Watermarked image (b) Watermark (c) Phony watermark extracted

Figure 3.2: Basic security test for Modified Jain watermarking scheme as in Figure 2.3, using Algorithm 2.1

As in the Jain et al. scheme, the extracted phony watermark is significantly distorted, and our modified scheme passes the basic security test.

3.1. Comparing Jain et al. and Modified Jain Watermarking Schemes. The above experimental results seem to indicate that the modified Jain watermarking scheme preserves the desirable robustness properties of the Jain et al. watermarking scheme while improving on the imperceptibility of the embedded watermark. In this section, we support this observation with some theoretical results and some additional experimental results.

Recall that the Jain et al. watermarking scheme has an embedding function

$$\mathcal{E}_J : (A, W, \alpha) \;\mapsto\; \begin{array}{l} A_J = A + \alpha U U_W S_W V^\top \\ K = (A, V_W, \alpha), \end{array}$$

and extraction function

$$\mathcal{X}_J : (A_J^*, K) \;\mapsto\; W_J^* = \alpha^{-1} U^\top (A_J^* - A) V V_W^\top.$$

Also recall that the modified Jain et al. watermarking scheme has an embedding

function

$$\mathcal{E}_M : (A, W, \alpha) \;\mapsto\; \begin{array}{c} A_M = A + \alpha U_W S_W V^\top \\ K = (A, V_W, \alpha), \end{array}$$

and extraction function

$$\mathcal{X}_M : (A_M^*, K) \;\mapsto\; W_M^* = \alpha^{-1}(A_M^* - A)VV_W^\top.$$

Using the above, we now show that the Frobenius norm error incurred in the extracted watermark in response to an additive perturbation P on the watermarked matrix is the same for the modified and unmodified Jain et al. scheme, and this error can be expressed in terms of $\|P\|$ and α, where $\|\cdot\|$ is any unitarily invariant norm.

PROPOSITION 3.1. *Let A, W, and P be matrices, let $\alpha \in \mathbb{R}$, and let $\|\cdot\|$ be a unitarily invariant matrix norm. With notation as above, Let*

$$\begin{aligned} (A_J, K_J) &= \mathcal{E}_J(A, W, \alpha), \\ (A_M, K_M) &= \mathcal{E}_M(A, W, \alpha), \\ \widetilde{W_J} &= \mathcal{X}_J(A_J + P, K_J), \text{ and} \\ \widetilde{W_M} &= \mathcal{X}_M(A_M + P, K_M). \end{aligned}$$

Then we have

$$\left\|W - \widetilde{W_J}\right\| = \left\|W - \widetilde{W_M}\right\| = \frac{\|P\|}{|\alpha|}.$$

Proof. Apply the SVD to obtain $A = USV^\top$ and $W = U_W S_W V_W^\top$. Under the Jain et al. watermarking scheme, if we attempt to extract the distorted watermark $\widetilde{W_J}$ from the perturbed matrix $\widetilde{A_J} = A_J + P$, we find that

$$\begin{aligned} W_J^* &= \alpha^{-1}(U^\top(\widetilde{A_J} - A)VV_W^\top \\ &= \alpha^{-1}U^\top(A_J + P - A)VV_W^\top \\ &= \alpha^{-1}U^\top(A + \alpha U U_W S_W V^\top + P - A)VV_W^\top \\ &= \alpha^{-1}U^\top(\alpha U U_W S_W V^\top + P)VV_W^\top \\ &= U_W S_W V_W^\top + \alpha^{-1}U^\top PVV_W^\top \\ &= W + \alpha^{-1}U^\top PVV_W^\top \end{aligned}$$

Computing the error with respect to the norm $\|\cdot\|$, we find that

$$\left\|\widetilde{W_J} - W\right\| = \left\|\alpha^{-1}U^\top PVV_W^\top\right\| = \frac{\|P\|}{|\alpha|},$$

since the matrices U, V, and V_W are orthogonal.

Similarly, for the modified Jain watermarking scheme, we find that

$$\begin{aligned}\widetilde{W_J} &= \alpha^{-1}(\widetilde{A_M} - A)VV_W^\top \\ &= \alpha^{-1}(A_M + P - A)VV_W^\top \\ &= \alpha^{-1}(A + \alpha U_W S_W V^\top + P - A)VV_W^\top \\ &= \alpha^{-1}(\alpha U_W S_W V^\top + P)VV_W^\top \\ &= U_W S_W V_W^\top + \alpha^{-1}PVV_W^\top \\ &= W + \alpha^{-1}PVV_W^\top,\end{aligned}$$

and

$$\left\|\widetilde{W_M} - W\right\| = \left\|\alpha^{-1}PVV_W^\top\right\| = \frac{\|P\|}{|\alpha|},$$

since V, V_W are orthogonal matrices. ∎

In particular, this result holds when $\|\cdot\|$ is the Frobenius or ℓ^2 norm.

Note that an additive perturbation to the watermarked matrix results in an additive perturbation to the extracted watermark. The proof of Proposition 3.1 also shows that the error in the extracted watermark increases with $\|P\|_F$ and decreases as α increases, as expected.

Next, informed by the experimental examples in Sections 3 and 2.3, we examine how the perceptibly of the watermark differs for the modified and unmodified Jain et al. watermarking schemes, similar to the error estimation for the Liu & Tan watermarking scheme given in [11].

Examining errors in terms of a unitarily invariant norm $\|\cdot\|$, we find that

$$\|A_J - A\| = \left\|\alpha U U_W S_W V^\top\right\| = \alpha \|W\|,$$

and

$$\|A_M - A\| = \left\|\alpha U_W S_W V^\top\right\| = \alpha \|W\|,$$

since the matrices U, U_W, and V are orthogonal. Thus, the error between the watermarked and original matrix for the modified and unmodified Jain et al. watermarking schemes are the same. However, the removal of the multiplier U in the additive term $\alpha U U_W S_W V^\top$ seems to have desirable effects on the imperceptibility of the watermark in the image. Thus, we use other methods to examine the relationships between A_J and A, and A_M and A.

As in [3], we define the correlation between two matrices X and Y to be

$$\operatorname{corr}(X, Y) = \frac{\langle X, Y\rangle_F}{\|X\|_F \|Y\|_F},$$

where $\langle X, Y\rangle_F$ is the Frobenius inner product $\langle X, Y\rangle_F = \operatorname{Tr}(X^\top Y)$. Note that $-1 \leq \operatorname{corr}(X, Y) \leq 1$, and if X is a scalar multiple of Y, then $\operatorname{corr}(X, Y) = 1$. This is the cosine of the angle between X and Y. In Figure 3.3, we display the results of some experiments in which we embedded some watermarks into images with various scaling factors α and examined the correlation between the resulting watermarked image and the original image.

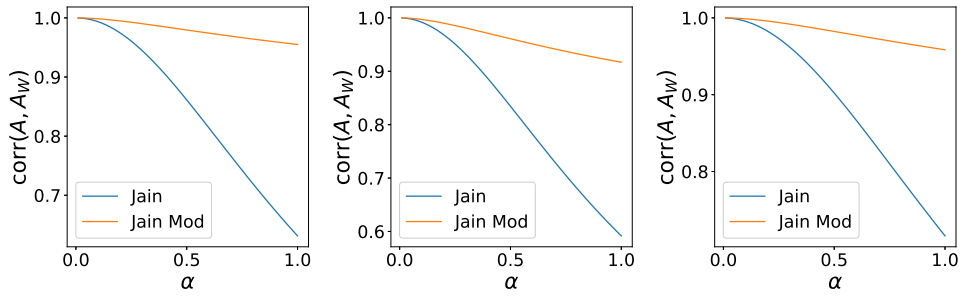

(a) Raccoon watermarked by fox (b) Raccoon watermarked by noise (c) Fox watermarked by husky

Figure 3.3: Correlations between image (A) and watermarked image (A_W) after watermark embedding with various scaling factors α

The correlation between the watermarked and original images remains above 0.95 for the modified Jain scheme, while it drops quickly for the unmodified version. This means that with the modified Jain scheme, the watermarked image is approximately a scalar multiple of the original image. This, combined with the results on the Frobenius norm error between the watermarked image and the original image, indicate that the modified Jain et al. watermarking scheme has better imperceptibility than the unmodified for any given scaling factor α, which is consistent with observation.

Although the error in the extracted watermark after additive perturbation is not different between the modified and unmodified Jain et al. schemes, the property of the modified Jain scheme described in the previous paragraph allows us to use larger scaling factors α in practice without a negative impact on perceptibility, resulting in less error in the extracted watermark.

4. Robustness Against Distortions. There are many of ways that a watermarked piece of media could become somehow distorted in its lifetime. It could undergo a deliberate cyber attack by some malicious actor. It might be shortened somehow, whether through the clipping of a video clip or the cropping of an image. It may even undergo SVD-based low rank compression. Thus, another major task in designing an watermarking scheme is in ensuring that a watermarked media can undergo some distortion and still contain proof of ownership to a reasonable degree.

As images are digitally transmitted, it is common that they might undergo low rank compression. While we already know that many natural images are often of intrinsically low-rank structure and can thus be compressed to a much lower rank, we are curious about how compression will affect extraction accuracy in watermarking schemes. We provide some examples in Figure 4.1, where the tree image is the watermark being embedded into the landscape image [2].

[1] Tree image from personal collection

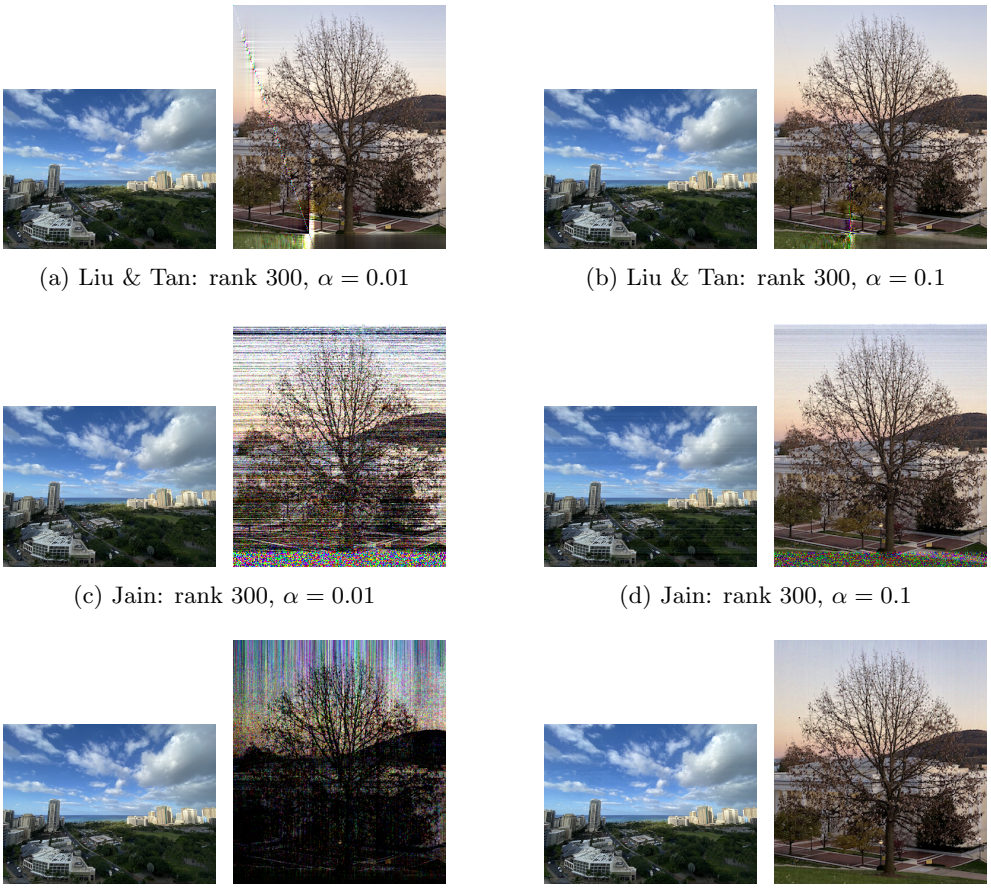

(a) Liu & Tan: rank 300, $\alpha = 0.01$

(b) Liu & Tan: rank 300, $\alpha = 0.1$

(c) Jain: rank 300, $\alpha = 0.01$

(d) Jain: rank 300, $\alpha = 0.1$

(e) Modified Jain: rank 300, $\alpha = 0.01$

(f) Modified Jain: rank 300, $\alpha = 0.1$

Figure 4.1: Robustness Against Low-Rank Compression

Low-rank compression is a one example of a common distortion. The plots in Figure 4.2 illustrate how as the rank is compressed using the SVD along the x-axis, the error according to the Frobenius norm of the extracted watermark compared with the original watermark is increased.

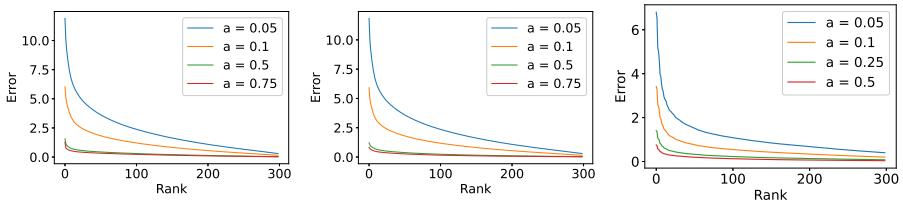

(a) Liu & Tan

(b) Jain

(c) Modified Jain

Figure 4.2: Robustness Against Low-Rank Compression Plots

Images can also be cropped, which eliminates entries from its matrix and complicate watermark extraction. We again provide examples of how watermarking schemes might withstand this in Figure 4.3 by changing 20 rows or columns on the right or bottom of the watermarked images to zeros as a means of simulating cropping.

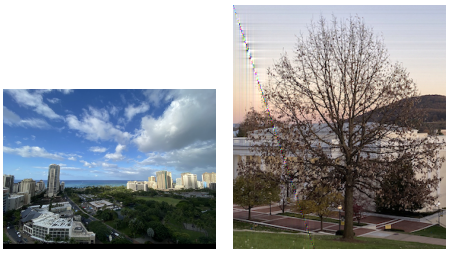

(a) Liu & Tan: 20 rows from bottom, $\alpha = 0.01$

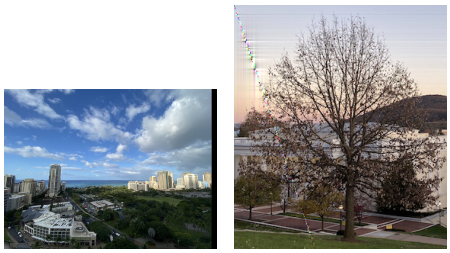

(b) Liu & Tan: 20 columns from right, $\alpha = 0.01$

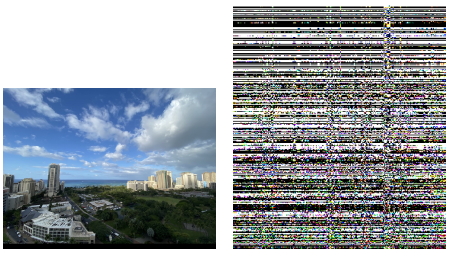

(c) Jain: 20 rows from bottom, $\alpha = 0.01$

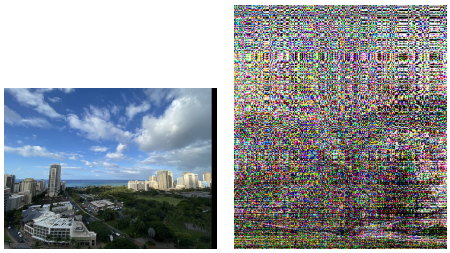

(d) Jain: 20 columns from right, $\alpha = 0.01$

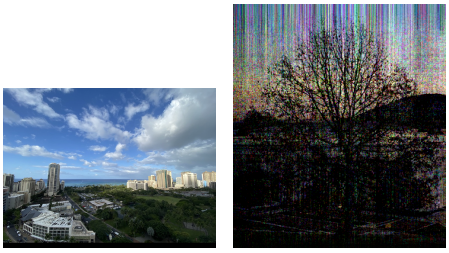

(e) Modified Jain: 20 rows from bottom, $\alpha = 0.01$

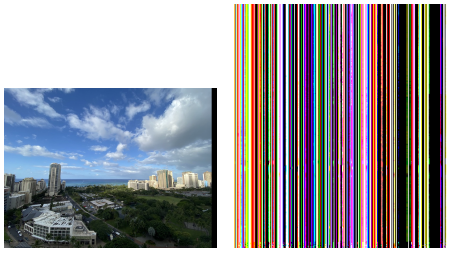

(f) Modified Jain: 20 columns from right, $\alpha = 0.01$

Figure 4.3: Robustness Against Cropping

We see from Figure 4.3 that the Liu & Tan scheme does exhibit robustness against cropping, although this is somewhat misleading, as its failure of the basic security test would suggest. For the Jain and Modified Jain schemes, even just 20 rows removed renders the extracted watermarks almost unrecognizable. Note that the side from which rows or columns were removed appears to matter.

However, note that the experiments presented in Figure 4.3 represent a somewhat idealized situation for cropping, in which we have knowledge of the original dimensions of the watermarked image and can fill in the remaining entries of the cropped image with zeros to match the dimension. What if a user were to receive a cropped version

of a watermarked image not knowing the original dimensions of the image, or knowing the original dimensions but not the side from which the image was cropped? This indicates a limitation to SVD-based watermarking scheme robustness and presents a potential challenge in practical implementation.

We also test how each scheme responds to rotations in Figure 4.4.

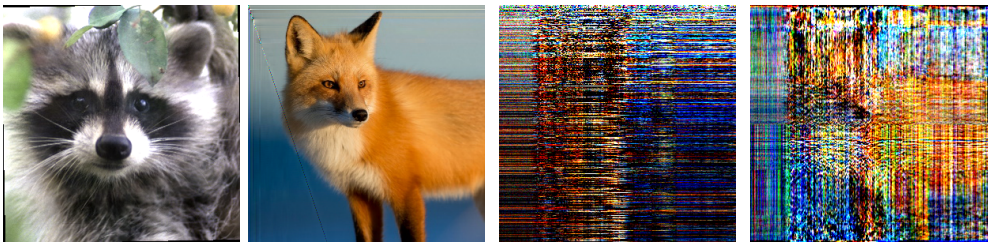

(a) 1° Rotation (b) 1° Liu & Tan (c) 1° Jain (d) 1° Modified Jain

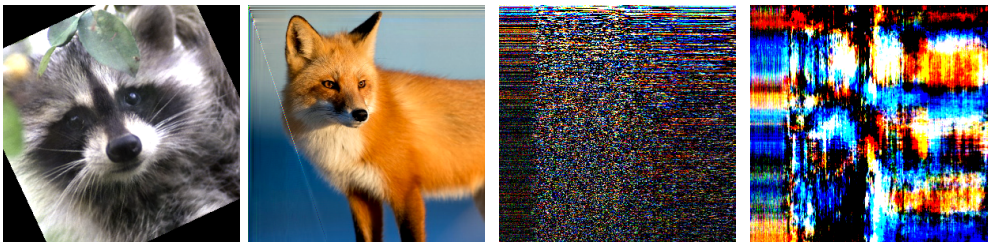

(e) 25° Rotation (f) 25° Liu & Tan (g) 25° Jain (h) 25° Modified Jain

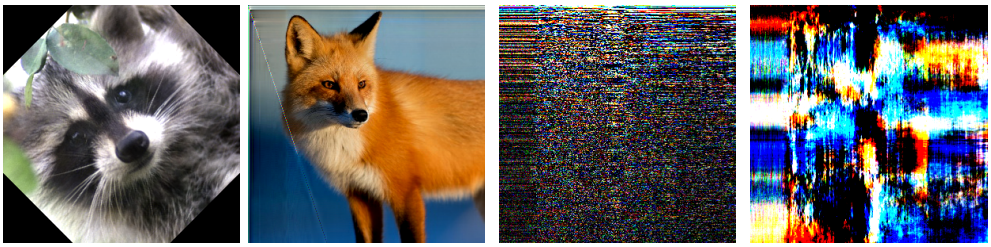

(i) 45° Rotation (j) 45° Liu & Tan (k) 45° Jain (l) 45° Modified Jain

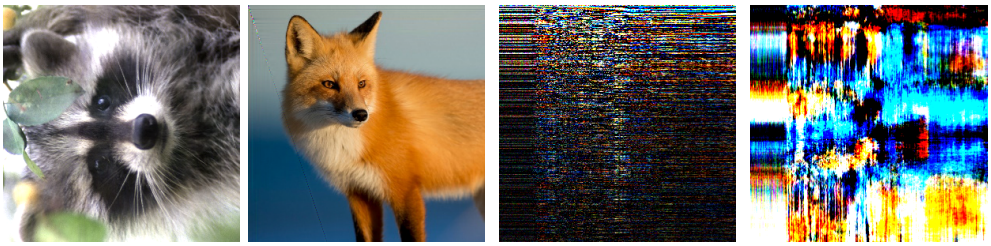

(m) 90° Rotation (n) 90° Liu & Tan (o) 90° Jain (p) 90° Modified Jain

Figure 4.4: Robustness Against Rotations

Note that the Liu & Tan extracted watermarks are extremely accurate. Both the

Jain and Modified Jain extracted watermarks are barely identifiable, even for just a one degree rotation. This represents a limitation of our modified scheme for this particular attack.

5. Audio Watermarking. A cornerstone property of the SVD is that it can be applied to any matrix. Similarly, these watermarking schemes are not limited to images and can also used for audio media such as songs or recordings. This is useful in trademarking music and finding proof of plagiarism in the entertainment industry.

Mono audio signals, in their raw form, come in as an array of N points sampled with frequency f_s (typically 44 100 Hz). To convert audio signals into matrices, we use the Short Time Fourier Transform (STFT), a time-frequency analysis technique. The output is a rectangular matrix where the entry in coordinate (f, t) is a complex number which describes the amplitude and phase of frequency f at time t. For visualization, a spectrogram can be obtained from the STFT of a signal by plotting the magnitude squared of the entries of the STFT.

Figure 5.1 displays an example audio signal and its associated spectrogram, with time on the horizontal axis and frequency on the vertical axis.

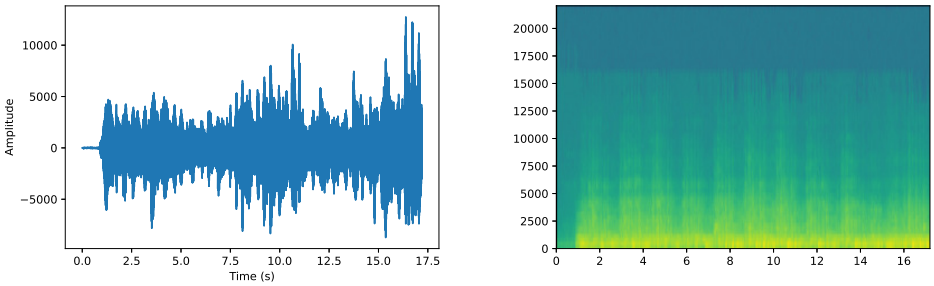

(a) Audio signal

(b) Spectrogram

Figure 5.1: Visualization of Frequency vs. Time

We use a Short Time Fourier Transform (STFT) with window size 256 and shift 128 to convert audio signals into matrices. [10, 7] further describe the STFT.

As a brief example of the application of our modified scheme in audio ownership protection, we can experiment with a 17-second clip from Bach Cello Suite No. 1 [1]. Just like how the modification of singular values were explored in Section 3 for images, scaling the singular values of an audio file results in the altered spectrograms seen in Figure 5.2. We multiply each singular value by a scalar p and observe that the volume of the audio is louder as p is adjusted to be greater.

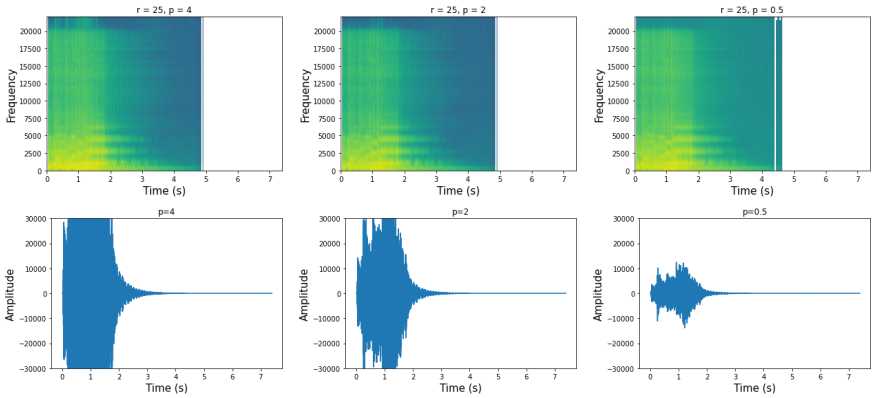

(a) p = 4 (b) p = 2 (c) p = 0.5

Figure 5.2: Modification of singular values by a linear scalar p for audio represented by their spectrograms and audio signals

While it is difficult to tell from visual inspection, it turns out that scaling the audio by a factor less than one results in a lowering of volume.

Now, when we embed a 7-second clip of news introduction music into the 17-second cello music clip according to the Modified Jain scheme with various values of α, we obtain the results shown in Figure 5.3:

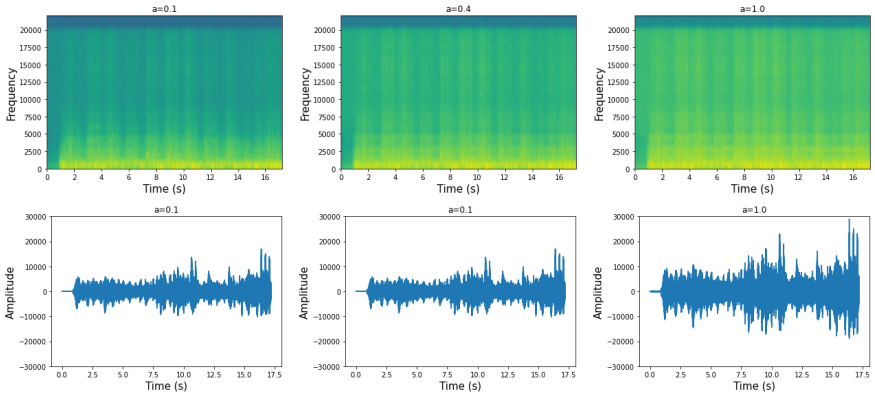

Figure 5.3: Audio watermark using the Modified Jain scheme with difference values of α represented by their spectrograms and audio signals

Finally, to determine the robustness of the Modified Jain watermarking algorithm, we apply four distortions to the watermarked music clip using external audio software: to add reverb, to lower all pitches, and to remove the noise. These results are shown in Figure 5.4 where the top two rows are the signals and spectrograms of the distorted watermarked audios, and the bottom two rows are the signals and spectrograms of the extracted watermark audios.

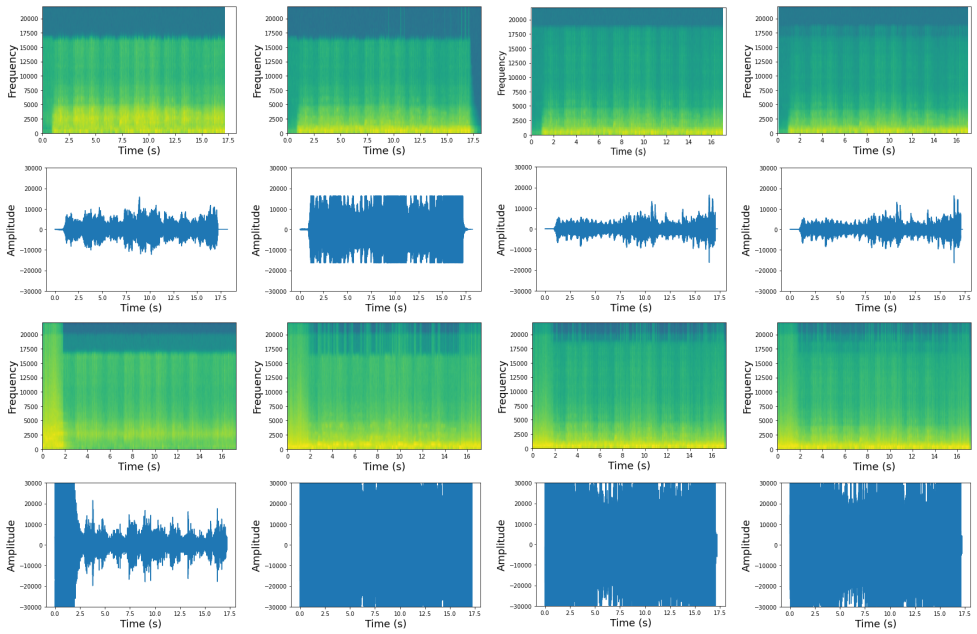

(a) Removing Low Frequencies (b) Adding Reverb Effects (c) Lowering by 1 semitone (d) Lowering by 3 semitones

Figure 5.4: Distortions on Watermarked Audio using the Modified Jain Watermarking Scheme and the Extracted Watermarks

6. Further Research Areas. We plan to explore what constitutes an ideal watermark: is there a certain type of image or media that we can choose to ensure the best robustness and security? In [11], Liu & Tan suggest that watermarks with random entries are particularly desirable for their watermarking scheme. We have done some experiments with this that suggest that there are certain benefits in security, but at the cost of extraction accuracy. We hope to further refine these results in future work.

We are also interested in using alternative SVD algorithms to develop an optimal watermarking scheme. It has been both mathematically and experimentally shown that the randomized and compressed SVD algorithms can dramatically decrease computing time while maintaining high performance and accuracy, especially for large matrices [6], [4]. We hope to run experiments in order to quantify this speed-up in the embedding and extraction process, as well as determine whether using such algorithms has any significant impact on properties like extraction accuracy or security.

One key limitation to the approach taken in this paper is that the basic security test outlined gives us a necessary but not necessarily sufficient condition for a watermark scheme to be secure. It will be important to explore more sophisticated notions of watermark security in future work, such as issues related to embedding multiple watermarks in the same image.

Finally, a key limitation to our proposed watermarking scheme is robustness against distortions such as clipping and rotation. Addressing these limitations is a natural next step in this work.

7. Conclusions. The unique properties of the SVD have been utilized in many fields, and are especially useful in protecting digital media. In this paper, we analyzed existing SVD-based watermarking schemes and propose our own modified version of the scheme put forth by Jain et al. in [8]. While the Liu-Tan watermarking scheme [11] is extremely robust to distortions but lacking in terms of ownership protection, the Jain watermarking scheme [8] gives up robustness to distortions and imperceptibility for stronger security. Our proposed modification to the Jain watermarking scheme preserves the security properties of the Jain watermarking scheme while improving imperceptibility.

We also introduce various means of quantifying robustness against distortion and attacks. Our goal for our modified scheme is to optimize the convenience and security of SVD-based watermarking schemes to meet the rapidly-expanding need for digital ownership protection.

REFERENCES

[1] J. S. Bach, *BACH, J.S.: Cello Suite No. 1, BWV 1007*, Naxos Digital Services US Inc.

[2] J. Beaufort, *Red Fox*. Licensed under CC0 1.0.

[3] M. W. Berry, Z. Drmac, and E. R. Jessup, *Matrices, vector spaces, and information retrieval*, SIAM Review, 41 (1999), pp. 335–362.

[4] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Cambridge University Press, 2019, https://doi.org/10.1017/9781108380690.

[5] B. Buchanan, *Young Raccoon in Crab Apple Tree*. licensed under Public Domain Mark 1.0.

[6] N. Erichson, S. L. Brunton, and J. Kutz, *Compressed singular value decomposition for image and video processing*, in 2017 IEEE International Conference on Computer Vision Workshop (ICCVW), Los Alamitos, CA, USA, oct 2017, IEEE Computer Society, pp. 1880–1888, https://doi.org/10.1109/ICCVW.2017.222, https://doi.ieeecomputersociety.org/10.1109/ICCVW.2017.222.

[7] E. Jacobsen and R. Lyons, *The sliding DFT*, IEEE Signal Processing Magazine, 20 (2003), pp. 74–80, https://doi.org/10.1109/MSP.2003.1184347.

[8] C. Jain, S. Arora, and P. K. Panigrahi, *A reliable SVD based watermarking schem [sic]*, CoRR, abs/0808.0309 (2008), http://arxiv.org/abs/0808.0309, https://arxiv.org/abs/0808.0309.

[9] D. Kalman, *A singularly valuable decomposition: The SVD of a matrix*, The College Mathematics Journal, 27 (1996), pp. 2–23, https://doi.org/10.1080/07468342.1996.11973744.

[10] N. Kehtarnavaz, *Chapter 7 - frequency domain processing*, in Digital Signal Processing System Design (Second Edition), N. Kehtarnavaz, ed., Academic Press, Burlington, second edition ed., 2008, pp. 175 – 196, https://doi.org/https://doi.org/10.1016/B978-0-12-374490-6.00007-6, http://www.sciencedirect.com/science/article/pii/B9780123744906000076.

[11] R. Liu and T. Tan, *An SVD-based watermarking scheme for protecting rightful ownership*, IEEE Transactions on Multimedia, 4 (2002), pp. 121–128, https://doi.org/10.1109/6046.985560.

[12] C. D. Martin and M. A. Porter, *The extraordinary SVD*, The American Mathematical Monthly, 119 (2012), pp. 838–851, https://doi.org/10.4169/amer.math.monthly.119.10.838, https://www.tandfonline.com/doi/abs/10.4169/amer.math.monthly.119.10.838, https://arxiv.org/abs/https://www.tandfonline.com/doi/pdf/10.4169/amer.math.monthly.119.10.838.

[13] G. Strang, *Linear Algebra and Learning from Data*, Wellesley-Cambridge Press, 2019, https://books.google.com/books?id=L0Y_wQEACAAJ.

[14] L. Zhang and A. Li, *Robust watermarking scheme based on singular value of decomposition in DWT domain*, in 2009 Asia-Pacific Conference on Information Processing, vol. 2, 2009, pp. 19–22.

[15] A. S. Öckel, *Siberian Husky Dog Pet*. Licensed under CC0 1.0.