

CERIAS Tech Report 2001-150
Detection of Image Alterations Using Semi-Fragile Watermarks
by Eugene T. Lin and Christine I. Podilchuk and Edward J. Delp
Center for Education and Research
Information Assurance and Security
Purdue University, West Lafayette, IN 47907-2086

Detection of image alterations using semi-fragile watermarks

Eugene T. Lin^a, Christine I. Podilchuk^b, Edward J. Delp^{†a}

^a Video and Image Processing Laboratory
School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47906

^b Bell Laboratories
Lucent Technologies
Murray Hill, NJ 07974

ABSTRACT

Digital watermarking research has generally focused upon two classes of watermarks, fragile and robust. Robust watermarks are designed to be detected even after attempts are made to remove them. Fragile watermarks are used for authentication purposes and are capable of detecting even minute changes of the watermarked content. Unfortunately, neither type of watermark is ideal when considering "information preserving" transformations (such as compression) which preserve the meaning or expression of the content and "information altering" transformations (such as feature replacement) which change the expression of the content. In this paper we describe a semi-fragile watermark for still images that can detect information altering transformations even after the watermarked content is subjected to information preserving alterations.

Keywords: semi-fragile watermarks, image authentication, alteration detection

1. INTRODUCTION

The applications of digital multimedia have grown immensely in the past five years. Many inexpensive methods are readily available for synthesizing and editing multimedia information. Unlike analog information the reproduction of a digital media element is simple and robust. Furthermore a copy of a digital media element is identical to the original. Unfortunately, the very same properties that make a digital media element attractive for use and distribution also facilitate unauthorized use, misappropriation, and misrepresentation. Thus, there is great interest in developing technology that will help protect the integrity of a digital media element and the intellectual property rights of its owners.

Watermarking, which is the embedding of a signal (known as the watermark) into the original multimedia element is one method that has been proposed for the protection of digital media elements such as audio, video, and images. The embedding is performed in such a way that the watermark is imperceptible under normal observation conditions, this is possible by taking advantage of the masking and other properties of the human visual system [1]. The desired characteristics of watermarks used to convey ownership and track the distribution of a digital multimedia element differs significantly from those used to ensure the integrity of a digital work. Therefore two primary types of digital watermarks have been studied: robust watermarks and fragile watermarks. Recently, a third type of watermark (the semi-fragile watermark) that combines the characteristics of robust and fragile watermarks has also been studied.

Robust watermarks are designed to resist attempts to remove or destroy the watermark. Their primary applications are copyright protection and content tracking. Robustness of the embedded watermark to removal by signal processing operations is crucial in order to resist unintentional destruction (such as lossy compression performed by legitimate users) or willful attempts of removal (such as intentional modifications performed by a pirate.) To achieve the desired robustness, many watermarking techniques are based on spread-spectrum techniques [2, 3] where a narrowband signal (the watermark) is transmitted through a noisy wideband channel (the media element). Recent techniques—known as image-adaptive watermarks [4]—vary the embedding parameters based on local image characteristics to maximize the robustness while minimizing the distortion caused by watermark embedding.

Unlike robust watermarks, fragile watermarks are designed to be easily destroyed if the watermarked image is manipulated in the slightest manner. This property is ideal for image authentication applications, where the objective is to determine if a watermarked image has been tampered with or modified. Many fragile watermarks are also capable of localization, where the

[†]This work was partially supported by grants to EJD from Texas Instruments, Intel, and the Purdue Center for Education and Research in Information Assurance and Security. Address all correspondence to E. J. Delp, ace@ecn.purdue.edu, <http://www.ecn.purdue.edu/~ace>, or +1 765 494 1740.

areas of the watermarked image that have been tampered with can be determined and distinguished from areas where the watermarked image has not been modified. Early fragile watermarking systems embedded checksums [5] or pseudo-random sequences [6, 7] in the least-significant bit (LSB) plane of an image while more recent systems apply more sophisticated embedding mechanisms [8], including the use of cryptographic hash functions [9] to detect changes to a watermarked image.

Robust and fragile watermarks are not ideal when considering “information preserving” transformations (such as compression) which preserve the meaning or expression of the content and “information altering” transformations (such as feature replacement) which change the expression of the content. Consider a watermarked image that is tampered by the addition or removal of features and then subsequently lossy-compressed. If the embedded watermark was a robust watermark, the detector could miss the relatively small feature changes because robust watermark detection typically examines large-scale correlations and dismiss the localized feature changes as noise. If the embedded watermark was a fragile watermark, the quantization noise of the compression will cause the detector to identify the entire image as tampered, even though the modified image would appear similar to the original with the exception of the modified features.

In this paper, we describe a semi-fragile watermark for still images that is capable of authenticating an image even after some degree of lossy compression has been done to the watermarked image.

1.1. Semi-fragile watermarks

A semi-fragile watermark combines the properties of fragile and robust watermarks. Like a robust watermark, a semi-fragile watermark is capable of tolerating some degree of change to the watermarked image, such as the addition of quantization noise from lossy compression. And like a fragile watermark, the semi-fragile watermark is capable of localizing regions of the image that have been tampered and distinguishing them from regions that are still authentic. Thus, a semi-fragile watermark can differentiate between localized tampering and information-preserving, lossy transformations. An example where semi-fragile techniques are useful includes a digital camera that embeds a camera-dependent signature (the watermark) into an image as the user takes a snapshot [10]. This embedded authentication watermark allows the user to show that the picture is not falsified or altered. However, it is not unreasonable to assume that the user may wish to post the digital image on the Internet, where lossy compression, such as JPEG, would be employed. Ideally, we would want the embedded watermark to also remain in the compressed image.

Because the primary applications of semi-fragile watermarks involve tamper detection and image authentication, the overall requirements of a semi-fragile watermarking system resemble those of fragile watermarking systems. The ability to identify regions of suspected alterations and distinguish those regions from other areas where there is high confidence that the watermarked image has not been damaged is crucial. Like most authentication systems, the watermark detector must not require that the original, unmarked image be available (i.e. blind detection.) And as in the general case of invisible watermarking, the protection provided by the watermark cannot degrade the value of the image.

1.2. Challenges and recent work in semi-fragile watermarks

In the development of semi-fragile watermarking systems, one should first consider the challenges that prevent naïve use of a good fragile or robust watermarking technique as a semi-fragile technique. Many fragile watermarking systems perform watermark embedding on the LSB plane and are unable to tolerate a single bit error in this bit. However, the quantization noise introduced by compression is likely to cause many least significant bits to change. Furthermore, recent fragile watermarking systems employ cryptographic hash functions which are not suitable in a semi-fragile framework. A hash function $h(x)$ will produce completely different outputs $h(x_1)$ and $h(x_2)$ if the binary inputs x_1 and x_2 are distinct but very similar. Even if some characteristic of the image that is expected to remain invariant during lossy compression was hashed, the output of the hash function would have to be embedded in a way that is resilient to errors.

The challenge of transforming a simple robust watermarking technique into a semi-fragile technique is to somehow gain the ability to detect and localize altered regions in the image. In spread-spectrum embedding, the watermark is treated as a low power signal transmitted through a very noisy channel. For reliable signal detection, long sequences are required at the receiver (watermark detector). Using sequences that are too short can result in false detections since the detector output would essentially be more dependent on the characteristics of the image and not of the watermark. However, while using longer sequences one can achieve better detection, long sequences also limit the localization sensitivity. To achieve better localization and reliable detection, a method for achieving processing gain is required.

Recent approaches for providing localized alteration detection, while allowing some robustness to image compression, include: embedding a low-information (or heavily quantized) version of the original image into the watermarked image, embedding key-dependent random patterns in blocks of the image, wavelet embedding to take advantage of the spatial and frequency information provided by the wavelet transform, and embedding multiple watermarks. These are discussed in more detail below.

The technique of embedding a low-information version of the original image is proposed in [11], this method was found to resist JPEG compression with a quality factor of 90. For some applications, this is an interesting technique to consider but the limited tolerance to lossy compression makes it impractical. Embedding information to “recover” the damage that had been inflicted upon an altered watermarked image is unnecessary for many applications that only require the damage be detected and localized.

Another technique is described in [12], where a spread-spectrum watermark is embedded in the DCT domain by a construction involving projections of image blocks to random smooth patterns. The patterns are used to construct the spread-spectrum signal that is embedded in the middle frequencies of the DCT. The technique was shown to survive moderate levels of JPEG compression. Techniques based on the wavelet transform have also been proposed to localize alterations. One wavelet-based technique based on Haar wavelets is proposed in [13], which was shown to be tolerant to high quality JPEG compression.

Embedding multiple watermarks [14] has also been described as a method of authenticating an image with a degree of robustness. For example, a robust watermark can be embedded into an image to establish ownership followed by the embedding of a fragile watermark for authentication. The primary disadvantage is that if lossy compression is performed, all of the authentication information is lost whereas a semi-fragile watermark is capable of providing a degree of authenticity even after lossy compression.

2. A SEMI-FRAGILE WATERMARK

The semi-fragile watermark we propose is based on extending a simple spread-spectrum watermarking technique with a modified detector that performs correlations of pixel differences in the spatial domain. The detection process is performed on each block of the image (as opposed to using the entire image) so regions of alterations can be identified.

2.1. Watermark construction and embedding

The watermark is constructed in the DCT domain to generate a smooth watermark that will resist being damaged by JPEG compression. Pseudo-random zero-mean, unit variance Gaussian distributed numbers are used for the watermark and are located in each DCT block as shown in Figure 1. Each block has a different watermark but the watermark is distributed in the block identically. For color images only the luminance component is watermarked. Note that some coefficients of each block, including the DC coefficient, are unmarked. The DC value is unmarked because any non-zero value added is likely to be visible and its effect does not contribute significantly to the detector. The high frequency AC coefficients are unmarked because any embedding at those frequencies is likely to be destroyed by lossy JPEG compression.

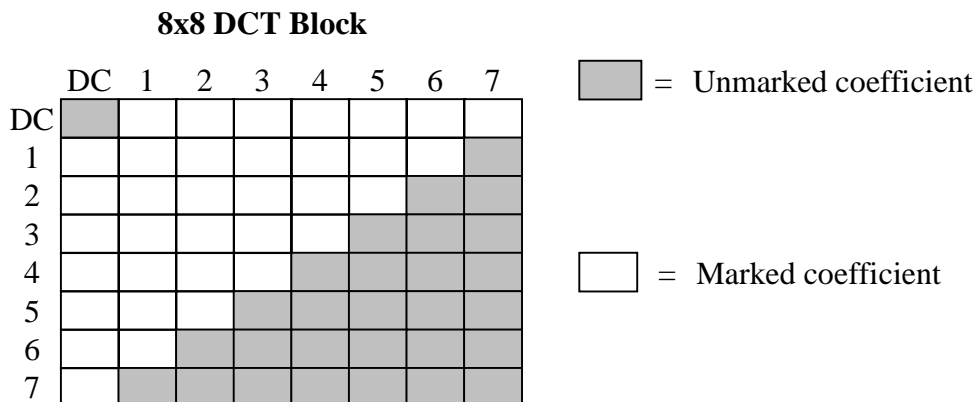


Figure 1: Watermark Generation in DCT Domain

Once the watermark has been constructed in the DCT domain, the inverse DCT is taken to produce a spatial domain watermark W . W is then embedded into the original image X to produce the watermarked image Y :

$$Y = X + \sigma W \quad (1)$$

Where σ is the watermark strength.

2.2. Watermark detection

Watermark detection is performed on a block-by-block basis for the localization of regions likely to be altered. A test statistic is generated for each block and compared with a threshold to classify the block as probably authentic or likely altered. First we describe how the detection statistic is generated for each block, and then how the classification is performed.

2.2.1. The block correlation detector

As mentioned earlier, a simple correlation detector will not be sufficient to achieve good detection because the length of the sequences to be correlated need to be much longer than the number of points available in a small block to overcome the power of the image signal.

Our detector is based on the differences of adjacent pixel values in the spatial domain. Most natural images consist of large areas of relatively smooth features with an occasional edge, so the power in a difference signal should be substantially less than the power of the signal. Therefore, unless an edge is present, the difference between adjacent pixel values will be the embedded watermark signal and a low energy random component from the image itself.

First, we define an operator $\Delta_{COL}(\cdot)$ as the difference-of-columns, where $B(x,y)$ is an arbitrary block:

$$\Delta_{COL}(B(x,y)) = \begin{cases} B(x,y) - B(x+1,y) & \text{if } x \in \{1,2,\dots, \text{Blocksize} - 1\} \\ 0 & \text{if } x = \text{Blocksize} \end{cases} \quad (2)$$

The difference-of-rows operator $\Delta_{ROW}(\cdot)$ can be similarly defined:

$$\Delta_{ROW}(B(x,y)) = \begin{cases} B(x,y) - B(x,y+1) & \text{if } y \in \{1,2,\dots, \text{Blocksize} - 1\} \\ 0 & \text{if } y = \text{Blocksize} \end{cases} \quad (3)$$

Figure 2 below shows an example of using these operators for block size of 4:

$$B(x,y) = \begin{bmatrix} -1 & 1 & 4 & -7 \\ 3 & 3 & 5 & -1 \\ 5 & 1 & 4 & 3 \\ 1 & -4 & -5 & -3 \end{bmatrix} \Rightarrow \Delta_{COL}(B) = \begin{bmatrix} -2 & -3 & 11 & 0 \\ 0 & -2 & 6 & 0 \\ 4 & -3 & 1 & 0 \\ 5 & 1 & -2 & 0 \end{bmatrix} \quad \Delta_{ROW}(B) = \begin{bmatrix} -4 & -2 & -1 & -6 \\ -2 & 2 & 1 & -4 \\ 4 & 5 & 9 & 6 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Figure 2: Example of difference operators on block B

Let T_b denote a block of the test image and W_b denote the corresponding block of the watermark. (The detector can regenerate the watermark image given the generator key and the σ used to embed the watermark.) If either the column difference or the row difference was correlated with the corresponding watermark difference, the result would involve n^2-n points (n =block size) because the last column or row of a difference block is always zero. We can achieve a more accurate correlation, however, by considering both the row and column differences to lessen the effect of horizontal and vertical edges.

Let T_b^* be the concatenation of the column-difference and row-difference of the test image block and let W_b^* be the corresponding watermark image as shown in equation 4 below:

$$T_b^* = [\Delta_{COL}(T_b(x, y)) \quad \vdots \quad \Delta_{ROW}(T_b(x, y))] \quad (4)$$

$$W_b^* = [\Delta_{COL}(W_b(x, y)) \quad \vdots \quad \Delta_{ROW}(W_b(x, y))]$$

Correlating T_b^* and W_b^* involves $2(n^2-n)$ points, which is greater than the n^2 points for a spatial correlation of the block. The block detection statistic is then the normalized correlation ρ_b shown in equation 5. Note that the dot product is defined on vectors and not matrices, so T_b^* and W_b^* are first “reshaped” to a row or column vector. The permutation by which the reshaping is performed is not important as long as the same permutation is used for both T_b^* and W_b^* .

$$\rho_b = \frac{T_b^* \cdot W_b^*}{\sqrt{(T_b^* \cdot T_b^*)(W_b^* \cdot W_b^*)}} \quad (5)$$

In summary, by correlating the spatial difference between adjacent pixels as opposed to a direct correlation, the expected signal power arising from the original image is dramatically reduced and a small amount of accuracy is gained by virtue of using more correlation points.

2.2.2. Classifying the blocks

Once the normalized correlation statistic ρ_b has been obtained for each block (equation 5 above), it is compared to a threshold T and an outcome is chosen:

$$\begin{aligned} \rho_b \geq T: & \text{ Block is authentic} \\ \rho_b < T: & \text{ Block is altered.} \end{aligned}$$

First consider the problem of blind detection, where the object is to determine the presence of the watermark in a test image. Let X_b be a block of the original (unmarked) image and define X_b^* to be the difference block as defined by:

$$X_b^* = [\Delta_{COL}(X_b(x, y)) \quad \vdots \quad \Delta_{ROW}(X_b(x, y))] \quad (6)$$

Natural images tend to have large, relatively smooth regions with occasional edges so we expect X_b^* to be zero-mean and have very small variance unless edges or strong textures are present. Since we are attempting to detect W_b^* in a test image, X_b^* is treated as uncorrelated noise and the binary hypothesis test is H_0 : Observe X_b^* versus H_1 : Observe $W_b^* + X_b^*$. From classical detection theory [15], the optimum detectors are correlation-based. If some assumptions about W_b^* and X_b^* were made, explicit expressions for the probability of false detection and miss can be derived. However, one can immediately see that the detector will perform well in smooth regions of the image and not as well in regions with edges or textures.

If a watermarked image undergoes quantization (quantization noise Q is added) the expectation is that Q will be uncorrelated with W_b^* and the statistic will decrease, however the amplitude of Q should be sufficiently small so that ρ_b is still above the threshold T for an unaltered block. If a deliberate attempt has been made to alter the image (perhaps by adding or removing a feature), and the attacker does not know the structure of the watermark, the probability that the altered block will be correlated with W_b^* is very small and $\rho_b \approx 0$.

3. EVALUATION OF THE SEMI-FRAGILE WATERMARK

3.1. Evaluation Procedure

To evaluate the proposed technique, a synthetic image (“*gradient*”) and real images (“*girls*”, “*sign*”, and “*money*”) were altered (see Figure 3 and Figure 4 below) prior to watermarking. A difference image between the (unmarked) altered and original versions of each image was constructed to establish the “truth” since knowledge of whether a particular block

contains alterations is essential in evaluating the detector. Obviously, the watermark detector itself will not have access to these difference images.

Once a difference image has been constructed, the original image was watermarked ($\sigma=5.0$, which was nearly invisible.) Then, for every pixel location where the unmarked original differs from the (unmarked) altered image, the corresponding pixel was copied from the altered image to the watermarked image, overwriting the watermarked image contents at these locations. Finally, this watermarked image and several compressed versions of it were created. This procedure simulates an attacker acquiring a watermarked image, altering it, and then lossy-compressing the altered image.

The detector generates a separate statistic ρ_b for each block. The values of ρ_b for altered and unaltered blocks determines the detectability. Let ρ_U be the set of detection statistics (ρ_b) for all unaltered blocks, and ρ_A be the set of ρ_b for altered blocks. If the means $E[\rho_U] = E[\rho_b | \text{Block } b \text{ is unaltered}]$ and $E[\rho_A] = E[\rho_b | \text{Block } b \text{ is altered}]$ are distant, then good detectability is achieved if the corresponding variances $\text{Var}[\rho_U]$ and $\text{Var}[\rho_A]$ are sufficiently small. However, if the means are similar, then ρ_b cannot be used to reliably distinguish the unaltered and altered blocks (regardless of any chosen threshold T .)

Table 1 below shows some selected values for $E[\rho_U]$, $E[\rho_A]$, $\text{Var}[\rho_U]$, and $\text{Var}[\rho_A]$. Figure 5 shows the same results in graphical form. The detection block size is 16×16 pixels. The detector should have little difficulty discerning the unaltered and altered blocks for lightly compressed “gradient” and “sign” images, as the difference in the means are large. However, effects of the edges in “sign” and the textures in “money” can be seen by the low difference of means for even high-quality compression. Out of the three real images examined, we expect “girls” to exhibit the best performance and “money” the worst.

The detector performance, measured as the percentage of correctly classified blocks, varies as the detection threshold is changed. Figure 6 shows the percentage of correct detections for the four images at various compression levels (see Table 1 for the bit rates at each JPEG compression level). For moderate compression, at least 75% correct block detection was achieved for all images using a threshold of $T=0.1$. It is interesting to note that the performance does not decrease uniformly as the amount of compression is increased; some edges present in the image that cause detector errors may be “softened” by lossy compression, yielding better detection.

Figure 7 shows an example illustrating the performance of the detector (see the caption for the embedding and detection parameters). Most of the detection errors occur near the edges in the image or in textured areas, as expected. Some of the detection misses near the edge of an altered region may be unreasonable for any semi-fragile watermark detector to find—a block is considered “altered” even if a single pixel value within that block was changed from the original. A semi-fragile detector with single-pixel sensitivity and resilience to lossy compression may not be feasible.

4. CONCLUSIONS AND FURTHER WORK

A semi-fragile watermark was described which could identify altered regions within a watermarked image with 75% accuracy under moderate compression and with near 90% accuracy under light compression. The detector was based on correlation of spatial-domain differences, which takes advantage of the fact that most natural images consist of large regions that are relatively smooth. Edges and textures increase the probability of erroneous detection.


While the presence of edges and texture can be a problem for the detector, regions of the image that are highly textured can also be more strongly watermarked due to the masking effects of the human visual system. Therefore, using the proposed watermarking technique with some kind of visual model should increase performance.

It should be noted that the size of the block chosen to determine whether portions of the image has been tampered also greatly affects the performance results. The block size studied here (16×16) is fairly small and for many practical applications, it may be sufficient to detect tampering on a much coarser scale (larger block size) where the accuracy results can be greatly improved.

A version of this paper with color images is available at <http://www.ece.purdue.edu/~ace>

JPEG Compress Q (%)	Image Bit Rate (bits/pixel)	Unaltered Blocks			Altered Blocks		
		Actual Number of Unaltered Blocks Present	Sample Mean ρ_b of Unaltered Blocks	Sample Variance of ρ_b for Unaltered Blocks	Actual Number of Altered Blocks Present	Sample Mean ρ_b of Altered Blocks	Sample Variance of ρ_b for Altered Blocks
GRADIENT							
Uncompress	24	382	0.9779	0.00257	300	0.1106	0.00037
90	1.1173	382	0.8634	0.00227	300	0.0962	0.00032
70	0.5504	382	0.4983	0.00131	300	0.0477	0.00016
50	0.3617	382	0.3100	0.00081	300	0.0290	0.00010
30	0.2425	382	0.1199	0.00031	300	0.0131	0.00004
GIRLS							
Uncompress	24	4753	0.5608	0.00012	951	0.0670	0.00007
90	2.1843	4753	0.5452	0.00011	951	0.0646	0.00007
70	1.1072	4753	0.3337	0.00007	951	0.0385	0.00004
50	0.7627	4753	0.2313	0.00005	951	0.0239	0.00003
30	0.5072	4753	0.1322	0.00003	951	0.0114	0.00001
SIGN							
Uncompress	24	1459	0.2377	0.00016	77	0.1598	0.00210
90	2.9761	1459	0.2210	0.00015	77	0.0692	0.00091
70	1.3825	1459	0.1990	0.00014	77	0.0561	0.00074
50	0.8617	1459	0.1857	0.00013	77	0.0399	0.00053
30	0.5726	1459	0.0954	0.00007	77	0.0124	0.00016
MONEY							
Uncompress	24	427	0.2407	0.00057	143	0.0189	0.00013
90	3.7061	427	0.2330	0.00055	143	0.0174	0.00012
70	2.0879	427	0.1699	0.00040	143	0.0119	0.00008
50	1.5433	427	0.1338	0.00031	143	0.0139	0.00010
30	1.1203	427	0.1002	0.00024	143	0.0104	0.00007

Table 1: Block Statistics for detector (embedding $\sigma=5.0$, detection blocksize=16x16)



Original Image "gradient"



Original Image "girls"




Original Image "sign"




Original Image "money"

Figure 3: Original Images



Altered "gradient"



Altered "girls"



Altered "sign"



Altered "money"

Figure 4: Altered Images

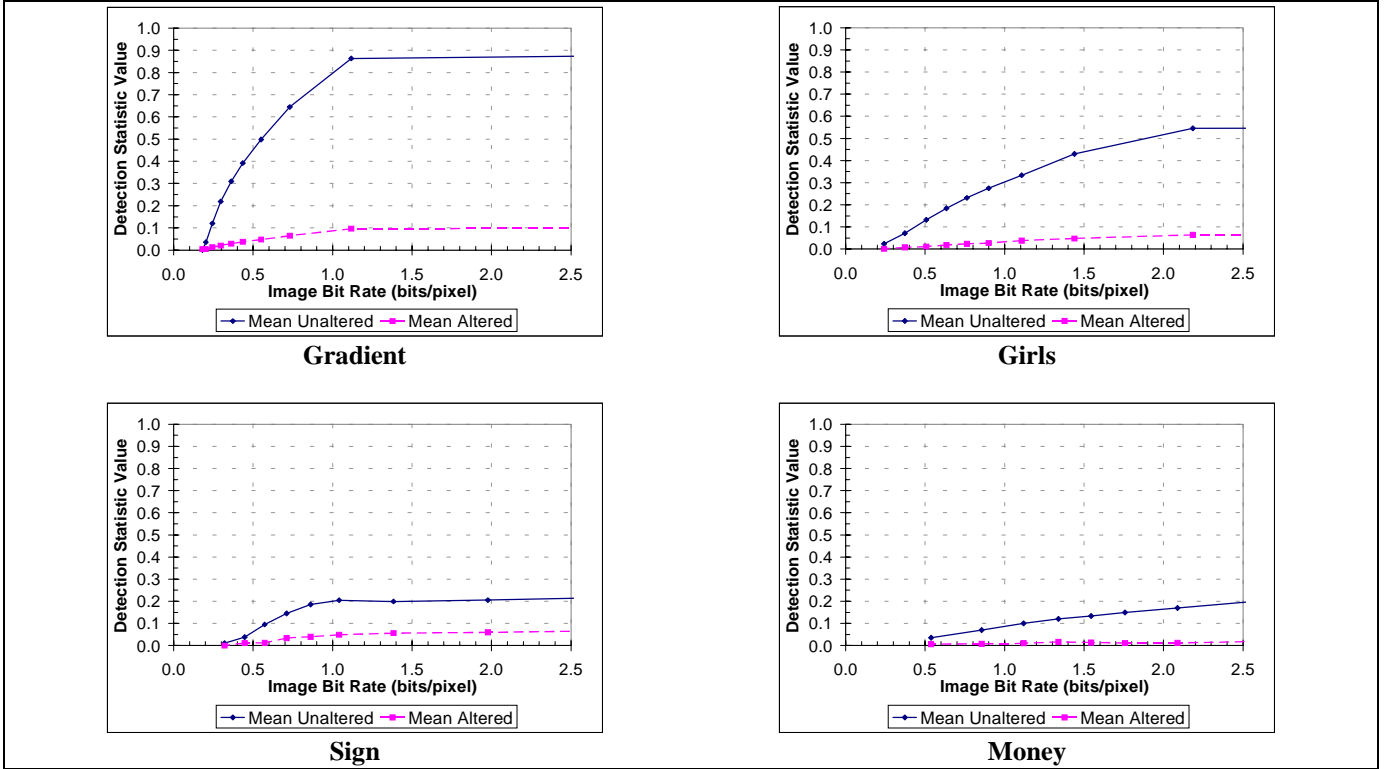


Figure 5: Detectability of altered blocks (embedding $\sigma=5.0$, detection blocksize=16x16)

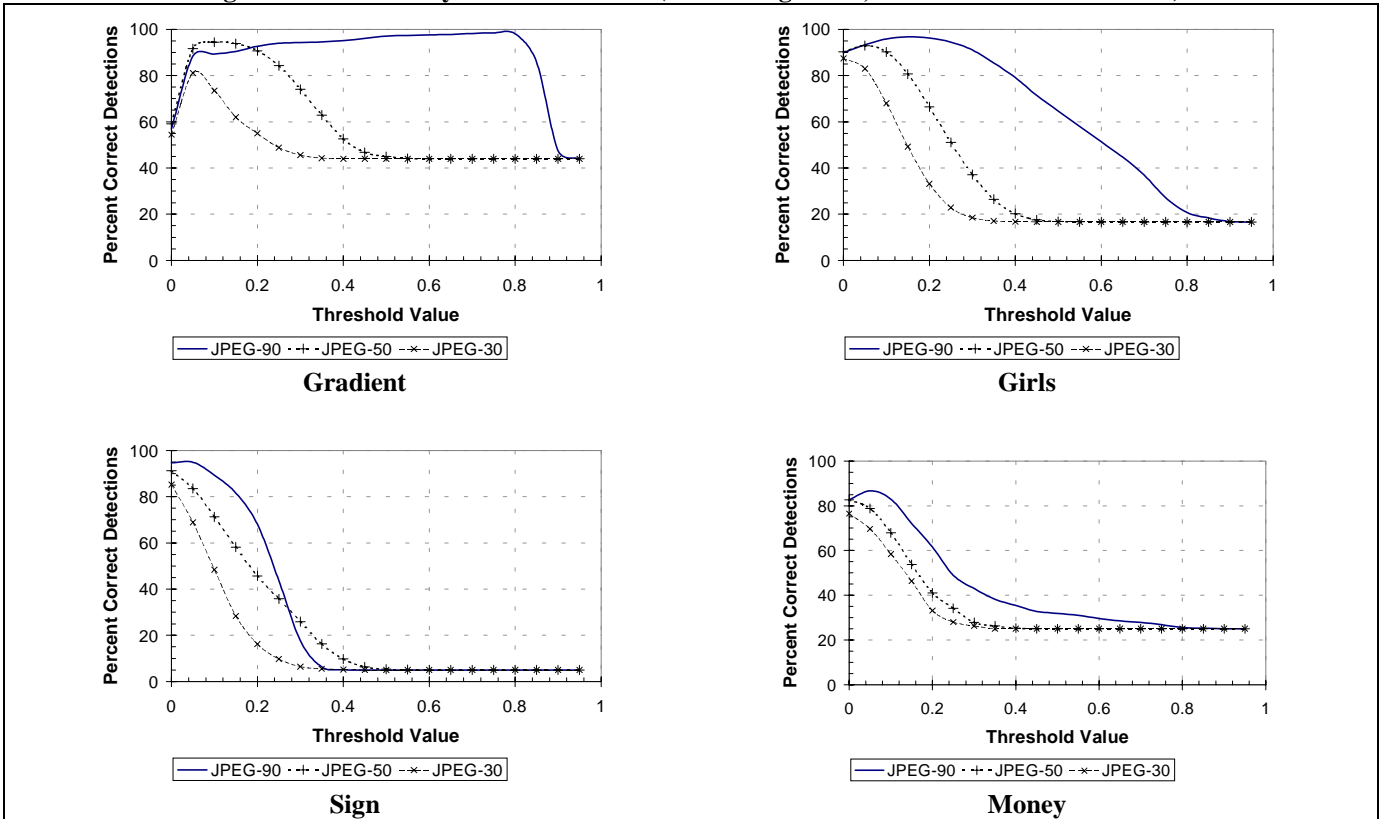


Figure 6: Detector Performance



Figure 7: Illustration of Detection Performance ($\sigma=5.0$, $T=0.1$, blocksize=16x16, JPEG-60 bitrate=0.90 bpp, 93% correct detection, 4% false positive, 17% misses.) A box indicates an altered block correctly identified, X indicates false positive, and X within a box indicates a miss.

REFERENCES

- [1] R. B. Wolfgang, C. I. Podilchuk, and E. J. Delp, "Perceptual Watermarks for Digital Images and Video," *Proceedings of the IEEE*, Vol. 87, No. 7, July 1999, pp. 1108-1126.
- [2] I. Cox, J. Kilian, T. Leighton, and T. Shamoan, "Secure spread spectrum watermarking for multimedia," *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1673-1687, December 1997.
- [3] A. Tirkel and T. Hall, "Advanced spread spectrum watermarking," *Proceedings of the ACM Multimedia and Security Workshop*, pp. 37-42, Orlando, Florida, October 1999.
- [4] C. I. Podilchuk and W. Zeng, "Image adaptive watermarking using visual models", *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 16, no. 4, pp. 525-539, May 1998.
- [5] S. Walton, "Information authentication for a slippery new age," *Dr. Dobbs Journal*, vol. 20, no. 4, pp. 18-26, April 1995.
- [6] R. Wolfgang and E. J. Delp, "A watermark for digital images," *Proceedings of the IEEE International Conference on Image Processing*, vol. 3, pp.219-222, 1996.
- [7] R. B. Wolfgang and E. J. Delp, "Fragile Watermarking Using the VW2D Watermark," *Proceedings of the SPIE/IS&T International Conference on Security and Watermarking of Multimedia Contents*, vol. 3657, January 25 - 27, 1999, San Jose, CA, pp. 204-213.
- [8] M. Yeung and F. Mintzer, "Invisible watermarking for image verification," *Journal of Electronic Imaging*, vol. 7, no. 3, pp. 578-591, July 1998.
- [9] P. Wong, "A watermark for image integrity and ownership verification," *Final Program and Proceedings of the IS&T PICS 99*, pp. 374-379, Savannah, Georgia, April 1999.
- [10] G. Friedman, "The trustworthy digital camera: Restoring credibility to the photographic image," *IEEE Transactions on Consumer Electronics*, vol. 39, pp. 905-910, November 1993.
- [11] J. Fridrich, and M. Goljan, "Images with self-correcting capabilities," *Proceedings of the IEEE International Conference on Image Processing*, Kobe, Japan, October 1999.
- [12] J. Fridrich, "Image watermarking for tamper detection," *Proceedings of the IEEE International Conference on Image Processing*, vol. 2, pp. 404-408, Chicago, Illinois, October 1998.
- [13] D. Kundur and D. Hatzinakos, "Towards a telltale watermarking technique for tamper-proofing," *Proceedings of the IEEE International Conference on Image Processing*, vol. 2, pp. 409-413, Chicago, Illinois, October 1998.
- [14] F. Mintzer and G. Braudaway, "If one watermark is good, are more better?," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, vol. 4, Phoenix, Arizona, May 1999.
- [15] H. Van Trees, *Detection, Estimation, and Modulation Theory*, John Wiley & Sons, 1968.