Steganalysis of Watermarking Techniques using Image Quality Metrics

İsmail Avcıbaş[a], Nasir Memon[a], Bülent Sankur[b]

[a]Department of Computer and Information Science, Polytechnic University, Brooklyn, NY, USA
[b]Department of Electrical and Electronics Engineering, Boğaziçi University, Bebek, İstanbul, Turkey

ABSTRACT

In this paper, we present techniques for steganalysis of images that have been potentially subjected to a watermarking algorithm. Our hypothesis is that a particular watermarking scheme leaves statistical evidence or structure that can be exploited for detection with the aid of proper selection of image features and multivariate regression analysis. We use some sophisticated image quality metrics as the feature set to distinguish between watermarked and unwatermarked images. To identify specific quality measures, which provide the best discriminative power, we use analysis of variance (ANOVA) techniques. The multivariate regression analysis is used on the selected quality metrics to build the optimal classifier using images and their blurred versions. The idea behind blurring is that the distance between an unwatermarked image and its blurred version is less than the distance between a watermarked image and its blurred version. Simulation results with a specific feature set and a well-known and commercially available watermarking technique indicates that our approach is able to accurately distinguish between watermarked and unwatermarked images.

Keywords: Steganalysis, watermarking, image quality measures, multivariate regression analysis.

1. INTRODUCTION

Steganography refers to the science of "invisible" communication. Unlike cryptography, where the goal is to secure communications from an eavesdropper, steganographic techniques strive to hide the very presence of the message itself from an observer. Although steganography is an ancient subject, the modern formulation of it is often given in terms of the *prisoner's problem* [1] where Alice and Bob are two inmates who wish to communicate in order to hatch an escape plan. However, all communication between them is examined by the warden, Wendy, who will put them in solitary confinement at the slightest suspicion of trouble. Specifically, in the general model for steganography, we have Alice wishing to send a *secret message m* to Bob. In order to do so she "embeds" *m* into a *cover-object c,* to obtain the *stego-object s.* The stego-object *s* is then sent through the public channel. The warden Wendy who is free to examine all messages exchanged between Alice and Bob can be *passive* or *active.* A passive warden simply examines the message and tries to determine if it potentially contains a hidden message. If it appears that it does, then she takes appropriate action else she lets the message through without alteration. An active warden on the other hand can alter messages deliberately, even though she does not see any trace of a hidden message, in order to foil any secret communication that can nevertheless be occurring between Alice and Bob.

It should be noted that the main goal of steganography is to communicate securely in a completely undetectable manner. That is, Wendy should not be able to distinguish in any sense between cover-objects (objects not containing any secret message) and stego-objects (objects containing a secret message). In this context, "*steganalysis*" refers to the body of techniques that are designed to distinguish between cover-objects and stego-objects. In this paper, we present a steganalysis technique for detecting the presence of *digital watermarks* in still images, using image quality metrics. Although we focus on images, the techniques we discuss would also be applicable to audio and video data.

A digital watermark is an imperceptible signal added to digital content that can be later detected or extracted in order to make some assertion about the content. Given the proliferation of content in digital form, recent years have seen an increasing interest in digital watermarking. In the past few years, many different watermarking algorithms have been proposed for different applications. Although the main applications for digital watermarking appear to be copyright protection and digital rights management, watermarks have also been proposed for secret communication, that is, steganography. However, there has been very little effort aimed at analyzing or evaluating the effectiveness of watermarking techniques for steganographic applications. Instead, most work had focused on analyzing or evaluating the watermarking algorithms for their robustness against various kinds of attacks that try to remove or destroy them. However, if digital watermarks are to be used in steganography applications, detection of their presence by an unauthorized agent defeats their very purpose. Even in applications that do not require hidden communication, but only robustness, we note that it would be desirable to first detect the possible presence of a watermark before trying to remove or manipulate it. This means that a given signal would have to be first analyzed for the presence of a watermark. Based on this analysis there could then be attempts made to remove the watermark.

A general underlying idea behind watermarking is to create a watermarked signal that is *perceptually identical but statistically different* from the host signal. A decoder uses this statistical difference in order to detect the watermark. However, the very same statistical difference that is created could be potentially exploited to determine if a given image contains a watermark. In this paper, we show that addition of a watermark leaves unique artifacts, which can be detected using sophisticated image quality measures. The rest if this paper is organized as follows. In Section 2, we discuss the selection of the image quality measures to be used in the steganalysis and the rationale of utilizing concurrently more than one quality measure. Section 3 describes the regression analysis to build a composite measure of quality to indicate the presence or absence of a watermark. Statistical tests and experiments are given in Section 4 and, finally, conclusions are drawn in Section 5.

2. CHOICE OF IMAGE QUALITY MEASURES

As stated in the introduction, the main goal of this paper is to develop a discriminator for watermark presence in still images, using an appropriate set of image quality measures. Image quality measurement continues to be the subject of intensive research and experimentation [5, 10, 11]. Objective image quality measures are based on image features, a functional of which, correlates well with subjective judgment, that is, the degree of (dis)satisfaction of an observer [7]. The interest in developing objective measures for assessing multimedia data lies in the fact that subjective measurements are costly, time-consuming and not easily reproducible. Objective measures are also utilized in performance prediction of vision algorithms against quality loss due to sensor inadequacy or compression artifacts [8]. In this paper, however, we want to exploit image quality measures, not as predictors of subjective image quality or algorithmic performance, but as steganalysis tools, that is, as detectors of watermark presence.

A good image quality measure should be accurate, consistent and monotonic in predicting quality. In the context of steganalysis, *prediction accuracy* can be interpreted as the ability of the measure to detect the presence of watermark with minimum error on average. Similarly, *prediction monotonicity* signifies that image quality measure scores should ideally be monotonic in their relationship to the strength of the watermark signal. Finally, *prediction consistency* relates to the quality measure's ability to provide consistently accurate predictions for a large set of watermarking techniques and image types. This implies that the spread of quality scores due to image variety or watermarking method should not eclipse the score differences arising from watermarking artifacts.

The steganalysis technique we develop is based on regression analysis of a number of "prominent" image quality measures. Hence, we seek quality measures that are sensitive specifically to watermarking and blurring effects. In other words, those measures for which the variability in score data can be explained better because of treatment rather then as random variations due to the image set. The reason why we consider blur in addition to watermarking will be more evident in Section 3. It suffices to say for the moment that the idea behind detection of watermark presence is to obtain consistent distances for images containing a watermark and those without, *with respect to a common reference processing*. The reference processing is blurring.

In order to identify specific quality measures that are useful in steganalysis, we use ANOVA (Analysis of Variance) [3] test. ANOVA helps us distinguish measures that are most consistent and accurate vis-à-vis the effects of watermarking and the effects of blurring. More specifically, we consider a set of training images, and the resulting quality degradation measures from watermarking and from blurring. These quality scores are then subjected to a statistical test to determine if the fluctuations of the measures result from image variety or whether they arise due to treatment effects, that is, blurring and watermarking.

As for the choice of quality measures we refer to [12] where some 26 measures had been investigated to predict compression, blur and noise artifacts. From these set of 26 measures, we aim to select the ones that best serve our purpose. The rationale of using several quality measures is that different measures respond with differing sensitivities to artifacts and distortions. For example, some measures like mean square error respond more to additive noise, others like spectral phase or mean square HVS-weighted (Human Visual System) error are more sensitive to pure blur, while the gradient measure reacts to distortions concentrated around edges and textures. However, we want our steganalyzer to be able to work with a variety of watermarking algorithms. Recall that some watermarking algorithms inject 'noise' in block DCT coefficients, another in a narrow-band of global DCT or Fourier coefficients, still others operate in selected localities in the spatial domain. Thus, a multitude of quality features are needed so that the steganalyzer has the chance to probe all features in an image that are significantly impacted by the watermarking process.

Table I: One-Way ANOVA results (F scores) for the image database used in training.

Quality Measure	*Watermark F score*	*Blurring F score*
Mean Square Error	**513.7**	57.81
MSE in Lab Space	35.63	49.03
Minkowsky Metric p=1 (Mean Abs. Distance)	309.4	122.9
Maximum Difference	4.931	40.28
Sorted Maximum Difference	10.18	39.36
Czenakowski.	45.24	24.71
Neighborhood Distance	22.06	0.34
Multiresolution Distance Measure	46.69	**134.3**
Structural Content	**370.8**	157.9
Cross Correlation	**573.6**	71.57
Image Fidelity	62.3	47.48
Angle Mean	0.72	4.41
Angle Standard Deviation	24.12	81.14
Pratt Measure	23.96	72.18
Gradient MSE	**35.52**	**368.1**
Spectral Magnitude	718.3	81.14
Spectral Phase	239.4	384.5
Weighted Spectral Distance	164.2	**809.3**
Median Block Spectral Magnitude	1069	22.05
Median Block Spectral Phase	1287	16.85
Median Block Weighted Spectral Distance	**1443**	28.59
Context Probabilistic Measure	0.075	2.29
f-divergence Variational Distance	0.0614	2.81
f-divergence Hellinger Distance	0.2962	3.18
f-divergence Matusita Distance	0.3667	3.62
Normalized Absolute Error (HVS)	155.9	**589.9**
Normalized Mean Square Error(HVS)	109.4	69.66
HVS Based L2	**450.9**	19.01
Multi Channel HVS based Measure	66.75	76.8

For ANOVA, a set of test images were watermarked at strength level 4 from the choice of 1 to 4 for Digimarc [14], denoted as μ_1, μ_2, μ_3, μ_4. The data given to the ANOVA algorithm consisted of four vectors, each of dimension M, where M is the number of images. More specifically, consider a typical quality measure, say $D(\mu)$, where the parametric dependence upon the watermark strength is shown with μ.. The M-dimensional vector **D** reads as: $\mathbf{D}(\mu) = [D(1|\mu) \;....\; D(M|\mu]^T$. Similarly, the images blurred with a Gaussian filter were analyzed. In this case the quality measure vector **D** can be denoted as $D(\sigma) = [D(1|\sigma) \ldots D(M|\sigma)]^T$, where the σ parameter indicates the spread factor of the blurring filter. Three blur filters were used.

In Table 1, we show how the quality measures scored in the ANOVA test when responding to watermarking and blurring effects. The selected subset of image quality measures with respect to their discriminative power, as evident by their high F scores in Table 1, were the following: 1) Mean Square Error, 2) Multiresolution Distance Measure, 3) Structural Content, 4) Cross Correlation, 5) Weighted Spectral Distance, 6) Median Block Weighted Spectral Distance, 7) Normalized Absolute Error (HVS), 8) HVS Based L2, and 9) Gradient Measure. These 'winning' measures are shown in bold in Table 1.

To further illustrate the qualities desired by a measure in its ability to discriminate between watermarked and non-watermarked images, we use box plots. Box plots are good graphical aids in giving an idea on the effectiveness of different measures. For example if the boxes are thinner, it implies that the measure has greater prediction accuracy and it suffers less outliers. In Figure 1, we exemplify box plots of a good and a poor measure for watermark detection. For box plot visualization, the data has been appropriately scaled without any loss of information.

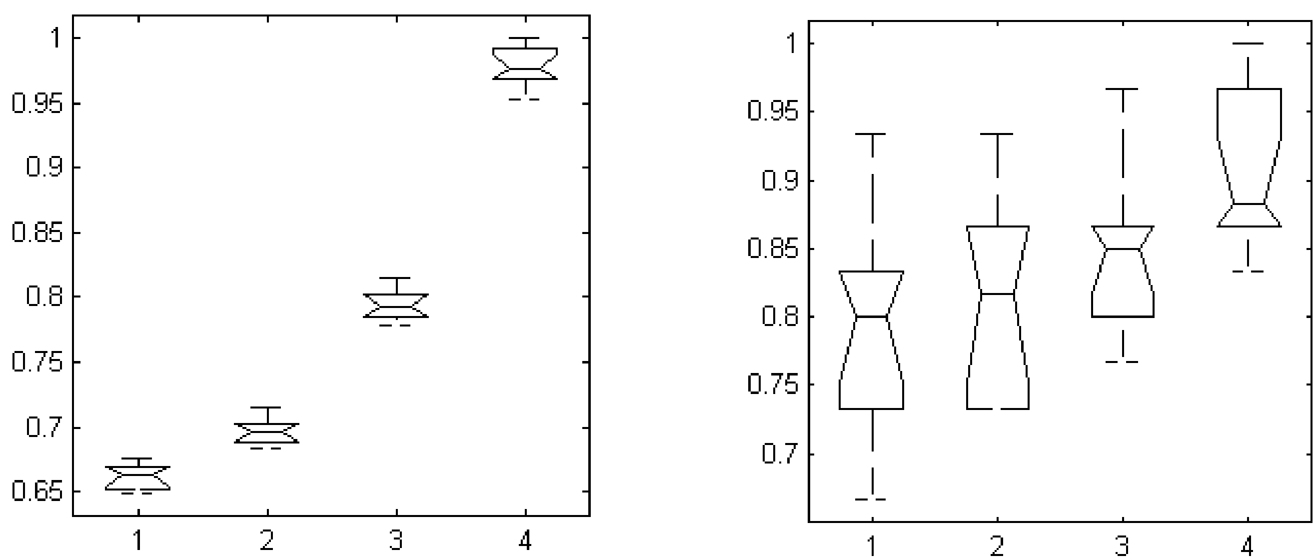

(a) (b)

Figure 1: Box plot examples of two image quality measures a) Median Block Weighted Spectral Distance measure with high discriminative power, F=1443, p=0; b) Maximum Difference measure with low discriminative power, F=4.93, p=0.006. Horizontal axis represents embedded watermark strength levels. Vertical axis is the image quality score normalized to the [0,1] interval.

3. REGRESSION ANALYSIS OF THE QUALITY MEASURES

The steganalysis we propose is based on the stratagem that a watermarked and blurred image would differ significantly from a non-watermarked but blurred image. In other words, both the watermarked and non-watermarked images are compared against the common reference of their blurred images. Watermark embedding can broadly be thought of as structural noise injection. It has been observed that blurring an image with no watermark causes changes in the quality measures differently than the changes brought about on non-watermarked images. This differential behavior is in part due to the fact that watermarking is not in general a global operation, but is local in nature. The watermark is either injected locally, e.g., on a block basis, or the watermark signal is subjected to a perceptual mask. In any case, we consistently obtained statistically different quality scores from blurred-and-watermarked images and from blurred-but-non-watermarked sources. In conclusion for the hypothesis test we use the various measured quality scores, which are either due to the difference between an originally non-watermarked image and its blurred version, or the difference between watermarked image and its blurred version.

In the design phase of the steganalyzer, we regressed the normalized quality measure scores to, respectively, -1 and 1, depending upon whether an image did not or did contain a watermark. Similarly, image quality scores were calculated between the original images and their blurred version. In the regression model [3], we expressed each decision label y in a sample of n observations as a linear function of the image quality measure scores x's plus a random error, ε :

$$\begin{aligned} y_1 &= \beta_1 x_{11} + \beta_2 x_{12} + \ldots + \beta_q x_{1q} + \varepsilon_1 \\ y_2 &= \beta_1 x_{21} + \beta_2 x_{22} + \ldots + \beta_q x_{2q} + \varepsilon_2 \\ &\vdots \\ y_n &= \beta_1 x_{n1} + \beta_2 x_{n2} + \ldots + \beta_q x_{nq} + \varepsilon_n \end{aligned}$$

In this expression, x_{ij} denotes the quality measure score, where the first index indicates the i'th image and the second one the quality measure. The total number of quality measures considered is denoted by q. The β's denote the regression coefficients. The complete statement of the standard linear model is

$$y = \mathbf{X}_{nxq}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \quad \textit{such that} \quad \begin{cases} rank(\mathbf{X}) = q \\ E[\boldsymbol{\varepsilon}] = \mathbf{0} \\ Cov[\boldsymbol{\varepsilon}] = \sigma^2 \mathbf{I} \end{cases}$$

The corresponding optimal MMSE linear predictor $\boldsymbol{\beta}$ can be obtained by

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{X}^T \mathbf{X}\right)^{-1}\left(\mathbf{X}^T \mathbf{y}\right).$$

Once the prediction coefficients are obtained in the training phase, these coefficients can be used in the testing phase. Given an image in the test phase, first it is blurred and the q image quality measure scores are obtained using the image and its blurred version. Then using the prediction coefficients, these scores are regressed to the output value. If the output exceeds the threshold 0 then the decision is that the image contains watermark, otherwise the decision is that the image does not contain watermark. That is

$$\hat{y} = \hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 + \ldots + \hat{\beta}_q x_q$$

for $\hat{y} \geq 0$ the image contains watermark, and for $\hat{y} < 0$ it does not.

4. SIMULATION RESULTS

The watermarking techniques we used were the following: 1) Photoshop plug-in Digimarc [14], Cox's technique [16], and the technique from Swiss Federal Institute of Technology, PGS [13]. One reason for the selection of these techniques was their free availability on the Internet and they were all popularly known algorithms. The other reason was that with these techniques it was possible to embed watermarks at different strengths.

We used an image database from [15] for the simulations. The database contained a variety of images including computer generated images, images with bright colors, images with reduced and dark colors, images with textures and fine details like lines and edges, and well-known images like Lena, peppers. Twelve images were used in the training and ten images (different from training images) were used in testing.

Table II. Detection performance for individual and pooled watermarking algorithms

Technique	*False Alarm*	*Miss*	*Correct Detection*	*% Performance*
Digimarc	3	0	17	85
Cox	4	1	15	75
PGS	0	0	20	100
All 3 watermark algorithms	4	2	34	85

We performed two types of training and testing for steganalyzer, one for the individual watermarking algorithms and the other for the ensemble of algorithms. In the individual case, the 12 training and 10 test images were watermarked with separate watermarking algorithms (Digimarc, Cox and PGS). They were compared against their non-watermarked versions in the test and training phases. Thus, three sets of regression coefficients were obtained, $\beta_{digimarc}$, β_{cox}, and β_{PGS}, one for each of the watermarking methods. In the mixed case, all the images marked with the three watermarking algorithms were pooled in one set (a set of 36 for training, another set of 30 for testing). The corresponding regression vector is referred as simply β. The results of individual and mixed tests are given in Tables II.

5. CONCLUSIONS

In this paper, we have addressed the problem of steganalysis of watermarked images. That is, we develop techniques for discriminating between watermarked and non-watermarked images. Our approach is based on the hypothesis that a particular watermarking scheme leaves statistical evidence or structure that can be exploited for detection with the aid of proper selection of image features and multivariate regression analysis. We used image quality metrics as the feature set to distinguish between watermarked and non-watermarked images. To identify specific quality measures, which provide the best discriminative power, we used analysis of variance (ANOVA) techniques. After selecting an appropriate feature set, we used multivariate regression techniques to get an optimal classifier using an image and its blurred version. Simulation results with a specific feature set and a well-known and commercially available watermarking technique indicates that our approach is able to accurately distinguish between watermarked and non-watermarked images.

6. APPENDIX

We give brief descriptions of the selected measures in this Appendix. We denote multispectral components of an image at the pixel position *i*, *j*, and in band *k* as $C_k(i,j)$, where $k = 1,...,K$. The boldface symbols $\mathbf{C}(i,j)$ and $\hat{\mathbf{C}}(i,j)$ indicate the multispectral pixel vectors at position (i,j). The following measures are functions of $\mathbf{C}$ and $\hat{\mathbf{C}}$, D($\mathbf{C}$, $\hat{\mathbf{C}}$), and their type is differentiated by a subscript.

Mean Square And Multiresolution Error Measures

These measures calculate the distortion between two images on the basis of their pixelwise differences. The L_γ norm of the dissimilarity of two images can be calculated by taking the Minkowsky average of the pixel differences spatially and then chromatically (that is over the bands):

$$D_{mean-square} = \frac{1}{K}\sum_{k=1}^{K}\left\{\frac{1}{N^2}\sum_{i,j=1}^{N}\left|C_k(i,j) - \hat{C}_k(i,j)\right|^{\gamma}\right\}^{1/\gamma}$$

For $\gamma = 2$, one obtains the well-known Mean Square Error (MSE) expression. An alternative measure that resembles image perception in the human visual system more closely can be obtained by assigning larger weights to low resolutions and smaller weights to the detail image [6]. For the r^{th} resolution level one would have than 2^{2r-2} blocks of size ($\frac{N}{2^{r-1}}$ x $\frac{N}{2^{r-1}}$), characterized by the block average gray levels g_{ij}, $i, j = 1,...,2^{2r-2}$. The distortion at this level is $d_r = \frac{1}{2^r}\frac{1}{2^{2r-2}}\sum_{i,j=1}^{2^{r-1}}\left|g_{ij} - \hat{g}_{ij}\right|$ where 2^{r-1} is the number of blocks along either the i and j indices and where the average difference in gray level at the resolution r has weight $\frac{1}{2^r}$. For multispectral images one can extend this definition by summing the multiresolution distances d_r^k over the bands:

$$D_{multiresolution} = \frac{1}{K}\sum_{k=1}^{K}\sum_{r=1}^{R} d_r^k$$

where d_r^k is the multiresolution distance in the k^{th} band.

Structural Content and Correlation Measures

The closeness between two digital images can also be quantified in terms of correlation function [5]. These measures measure the similarity between two images; hence in this sense they are complementary to the difference-based measures. The Structural Content and Normalized Cross-Correlation measures are defined as follows:

Structural Content:

$$D_{structural\_content} = \frac{1}{K}\sum_{k=1}^{K}\frac{\sum_{i,j=0}^{N-1} C_k(i,j)^2}{\sum_{i,j=0}^{N-1} \hat{C}_k(i,j)^2},$$

Normalized Cross-Correlation measure:

$$D_{cross\_correlation} = \frac{1}{K}\sum_{k=1}^{K}\frac{\sum_{i,j=0}^{N-1} C_k(i,j)\hat{C}_k(i,j)}{\sum_{i,j=0}^{N-1} C_k(i,j)^2},$$

Weighted Spectral Distance and Median Block Weighted Spectral Distance:

In this category we consider the distortion penalty functions obtained from the complex Fourier spectrum of images. Consider the phase and magnitude spectra, that is $\varphi(u,v) = \arctan(\Gamma(u,v))$ and $M(u,v) = |\Gamma(u,v)|$, where $\Gamma_k(u,v)$ denotes the Discrete Fourier Transforms (DFT) of the kth band of the image. The weighted spectral distortion

$$D_{spectral\_dis\tan ce} = \frac{1}{N^2}\left(\lambda\sum_{u,v=1}^{N}\left|\varphi(u,v) - \hat{\varphi}(u,v)\right|^2 + (1-\lambda)\sum_{u,v=1}^{N}\left|M(u,v) - \hat{M}(u,v)\right|^2\right)$$

where λ is to be chosen to balance the make commensurate the contributions of the phase and magnitude errors. Due to the localized nature of distortion and/or the non-stationary image field, Minkowsky averaging of block spectral distortions can be advantageous. Thus the image is divided into L blocks of size b x b, and blockwise spectral distortions can be computed. If the DFT of the lth block of the kth band image $C_k^l(m,n)$ is denoted as $\Gamma_k^l(u,v)$, where $u,v = 1...b$ and $l = 1,...,L$, then for each such block, one can use:

$$J^l = \lambda \frac{1}{K}\sum_{k=1}^{K}\left(\sum_{u,v=1}^{b}\left(\left|\Gamma_k^l(u,v)\right| - \left|\hat{\Gamma}_k^l(u,v)\right|\right)^{\gamma}\right)^{1/\gamma} + (1-\lambda)\frac{1}{K}\sum_{k=1}^{K}\left(\sum_{u,v=1}^{b}\left(\left|\phi_k^l(u,v)\right| - \left|\hat{\phi}_k^l(u,v)\right|\right)^{\gamma}\right)^{1/\gamma}$$

Among possible rank order operations on the block spectral differences the median has proven useful, which corresponds to the $D_{block\_spectral\_distance} = \underset{l=1\ldots L}{Median}\, J^l$ measure in Table 1. The norm parameter set at γ=2 and block sizes of 32x32 or 64x64 yielded higher F scores.

Normalized Absolute HVS Error and Mean Square HVS Errors:

The incorporation of HVS (Human Visual System) model into objective measures [2], [FRES] has led to a better correlation with the subjective ratings in multimedia. It is conjectured therefore that in steganalysis tasks they may have as well some relevance. We assume that the human visual system can be modeled as a band-pass filter with a transfer function in polar coordinates,

$$H(\rho) = \begin{cases} 0.05e^{\rho^{0.554}} & \rho < 7 \\ e^{-9\left[\left|\log_{10}\rho - \log_{10} 9\right|\right]^{2.3}} & \rho \geq 7 \end{cases}$$

where $\rho = \left(u^2 + v^2\right)^{1/2}$. Once images are processed with such a spectral mask and inverse DCT transformed, the various measures in the pixel-difference group can be adapted. Thus one has:

Normalized Absolute HVS Error:

$$D_{absolute\_HVS} = \frac{1}{K}\sum_{k=1}^{K}\frac{\sum_{i,j=0}^{N-1}\left|U\{C_k(i,j)\} - U\{\hat{C}_k(i,j)\}\right|}{\sum_{i,j=0}^{N-1}\left|U\{C_k(i,j)\}\right|}$$

Normalized Mean Square HVS Error:

$$D_{MSE\_HVS} = \frac{1}{K}\sum_{k=1}^{K}\frac{\sum_{i,j=0}^{N-1}\left[U\{C_k(i,j)\} - U\{\hat{C}_k(i,j)\}\right]^2}{\sum_{i,j=0}^{N-1}\left[U\{C_k(i,j)\}\right]^2}$$

7. REFERENCES

1. Gustavus J. Simmons, Prisoners' Problem and the Subliminal Channel (The), CRYPTO83 - Advances in Cryptology, August 22-24. 1984. pp. 51-67.
2. A. B. Watson, Ed., *Digital Images and Human Vision*, Cambridge, MA, MIT Press (1993).
3. A. C. Rencher, *Methods of Multivariate Analysis*, New York, John Wiley (1995).
4. A. M. Eskicioğlu, "Application of Multidimensional Quality Measures to Reconstructed Medical Images", *Opt. Eng*. 35(3) 778-785 (1996).
5. A. M. Eskicioğlu, P. S. Fisher, "Image Quality Measures and Their Performance", *IEEE Trans. Commun*., 43(12), 2959-2965 (1995).
6. P. Juffs, E. Beggs, F. Deravi, "A Multiresolution Distance Measure for Images", *IEEE Signal Processing Letters*, 5(6), 138-140 (1998).

7. S. Daly, "The visible differences predictor: An algorithm for the assessment of image fidelity", in *Digital Images and Human Vision,* A. B. Watson, ed., Cambridge, MA, MIT Press, 179-205 (1993).
8. C.E. Halford, K.A. Krapels, R.G. Driggers, E.E. Burroughs, Developing Operational Performance Metrics Using Image Comparison Metrics and the Concept of Degradation Space, Optical Engineering, 38 (5), 836-844, 1999.
9. W. K. Pratt, *Digital Image Processing*, New York, Wiley (1978).
10. C. B. Lambrecht, Ed., "Special Issue on Image and Video Quality Metrics", *Signal Processing*, vol. 70, (1998).
11. İ. Avcibas, B. Sankur, "Statistical Analysis of Image Quality Measures", *EUSIPCO'2000*: 10th European Conf. Signal Processing, 2181-2184, Tampere, Finland, 5-8 September 2000.
12. İ. Avcibas, B. Sankur, K. Sayood, "Statistical Evaluation of Image Quality Measures", Journal of Electronic Imaging, under review.
13. Copyright © 1998, Swiss Federal Institute of Technology, Lausanne, Switzerland, (EPFL). http://www\WM\AMT\Watermarking SoftwareJK_PGS.htm
14. PictureMarc Embed Watermark, v 1.00.45, Copyright 1996, Digimarc Corporation.
15. http://www.cl.cam.ac.uk/~fapp2/watermarking/benchmark/image_database.html.
16. Ingemar J. Cox, Joe Kilian, Tom Leighton, and Talal Shamoon. A secure, robust watermark for multimedia. In Anderson [1], pages 183-206.