



Is Image Steganography Natural?

Alvaro Martin¹, Guillermo Sapiro², Gadiel Seroussi
Information Theory Research Group
HP Laboratories Palo Alto
HPL-2004-39(R.1)
August 10, 2004*

image models,
steganography

Steganography is the art of secret communication. Its purpose is to hide the presence of information, using for example images as covers. We experimentally investigate if stego-images, bearing a secret message, are statistically "natural." For this purpose, we use recent results on the statistics of natural images and investigate the effect of some popular steganography techniques. We found that these fundamental statistics of natural images are, in fact, generally altered by the hidden "non-natural" information. Frequently, the change is consistently biased in a given direction. However, for the class of natural images considered, the change generally falls within the intrinsic variability of the statistics, and thus does not allow for reliable detection, unless knowledge of the data hiding process is taken into account. In the latter case, significant levels of detection are demonstrated.

* Internal Accession Date Only

¹Instituto de Computacion, Facultad de Ingenieria, Universidad de la Republica, Montevideo, Uruguay. Work done while author was with HP Laboratories Palo Alto and Electrical and Computer Engineering Dept., University of Minnesota, Minneapolis, MN 55455

²Electrical and Computer Engineering and Digital Technology Center, University of Minnesota, Minneapolis, MN
Approved for External Publication

Is Image Steganography Natural?*

*Alvaro Martín[†], Guillermo Sapiro[‡] and Gadiel Seroussi[§]

Abstract

Steganography is the art of secret communication. Its purpose is to hide the presence of information, using for example images as covers. We experimentally investigate if stego-images, bearing a secret message, are statistically “natural.” For this purpose, we use recent results on the statistics of natural images and investigate the effect of some popular steganography techniques. We found that these fundamental statistics of natural images are, in fact, generally altered by the hidden “non-natural” information. Frequently, the change is consistently biased in a given direction. However, for the class of natural images considered, the change generally falls within the intrinsic variability of the statistics, and thus does not allow for reliable detection, unless knowledge of the data hiding process is taken into account. In the latter case, significant levels of detection are demonstrated.

Index Terms— Steganography, Information Hiding, Image Models, Natural Images.

EDICS— 2-MODL Modeling, 5-AUTH Authentication and Watermarking.

1 Introduction

In steganography, we study techniques to achieve secret communication between two parties that are interested in hiding not only the content of a secret message but also the act of communicating it. To this aim, steganography algorithms (“*stego algorithms*”) embed the secret information into different types of “natural” cover data like sound, images, or video. The resulting altered data is referred to as *stego-data* and it must be perceptually indistinguishable from its natural cover. On the other hand, *stego-analysis* seeks to analyze (possibly altered) cover data to decide whether a message has been embedded in it or not. Thus, the problem can be seen as one of classification into two classes, namely, natural and stego-data.

*This work is partially supported by the Office of Naval Research grants N000140310399 and N000140310176, by the Presidential Early Career Award for Scientists and Engineers (PECASE), and a National Science Foundation CAREER Award.

[†]Instituto de Computación, Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay. Work done while the author was with Information Theory Research, Hewlett-Packard Laboratories, Palo Alto, CA 94304, and Electrical and Computer Engineering Department, University of Minnesota, Minneapolis, MN 55455.

[‡]Electrical and Computer Engineering and Digital Technology Center, University of Minnesota, Minneapolis, MN 55455.

[§]Information Theory Research, Hewlett-Packard Laboratories, Palo Alto, CA 94304.

In this paper we focus on the use of natural images, i.e., images that appear naturally in “real world” photographic scenes, as covers, and study how several recently proposed statistical models can be used for stego-analysis. These models establish particular distributions for statistics or relations among statistics which naturally induce a set of experiments. We investigate model statistics for natural and stego images and whether the models fit as accurately for stego as for natural images. The goal then is to investigate if the act of embedding (hiding) a “non-natural” message into a “natural” image, changes some of the basic statistics of the image, thereby allowing for the detection (but not necessarily interpretation) of the presence of a hidden message. For instance, we will show that a model for the distribution of the differences between adjacent pixels, which fits natural images very accurately, is not a good model for images altered by one of the stego algorithms in S-Tools [1], a popular package we included in our experiments. Other algorithms, like Jsteg [2], however, do not significantly violate this property.

We also experiment with statistics based on wavelet coefficients and block discrete cosine transform (DCT) coefficients, exploiting (partial) knowledge of the data hiding technique. Although these are basic statistics that are not specifically tailored to natural images, the results are still interesting as we found that the embedding of a hidden message bias the statistics of the image in a consistent direction.

While previous works [3, 4] had focused on rather simple image statistics, in [5], the authors proposed a stego-analysis technique based on image quality metrics while, in [6, 7], the author proposed a technique based on high order statistics of wavelet coefficients. Recently, in [8], a stego algorithm resistant to the techniques of [6, 7] was introduced. This algorithm is a modified version of the Histogram-Preserving Data Mapping (HPDM) [9], and we will refer to it as MHPDM.

One of the main conclusions of this work is that embedding a stego message generally alters the studied statistics of its cover image. Moreover, as mentioned, in some cases the hidden data biases some of the statistical parameters in a consistent direction. On the other hand, the effect is often not sufficient to “move” a significantly large set of images beyond what may be considered natural according to the studied statistical models, when the analysis is independent of the stego algorithm used. On the other hand, we demonstrate that better results, including statistically significant discrimination between natural and stego-images, can be obtained when (partial) knowledge of the

stego algorithm is used in the analysis.

The remainder of this paper is organized as follows. Section 2 briefly describes the stego algorithms that are considered in our experiments, and Section 3 introduces the models of natural images that are tested for sensitivity to steganography. Section 4 describes the general setting for the experiments, the specifics of each experiment, and the results obtained. Finally, the conclusions on the results, and directions for future research, are summarized in Section 5.

2 Steganography Algorithms

We consider three different stego algorithms in our experiments: Jsteg [2], MHPDM [8], and one of the algorithms in S-Tools [1]. Jsteg embeds a message in the least significant bit of JPEG DCT coefficients. The algorithm selected in S-Tools admits 8-bit palletized images (256 colors) as inputs, and maintains this range throughout processing. The algorithm operates in two stages. First it reduces the number of entries in the color palette of the cover image, and then it embeds a message in the least significant bits of the three RGB components, without expanding the number of colors beyond 256. Note of course that, as each RGB component of each pixel is altered independently, this technique is not directly suitable for gray-scale images since it can be detected by simply observing that some colors in the color palette are not exactly gray. We experimented with this algorithm as an example of a scheme operating in the pixel domain. To study the effects of S-Tools purely on image statistics (our main focus in the paper), the mentioned color-shift issue was bypassed by transforming RGB stego-images back to gray scale, taking the rounded luminance of each pixel.

The MHPDM algorithm [9], as well as its predecessor HPDM [8], works by altering the least significant bit of a subset of the JPEG DCT coefficients of an image. If the 64 coefficients of each DCT block are indexed from zero following the usual zig-zag order [10], only coefficients 1 through 20 are candidates for modification. The rest are left untouched, since values of coefficient with index 0 (DC) are far from being independent, and coefficients 21 through 63 are highly quantized during the JPEG process.

Both MHPDM and HPDM preserve the zero-order histograms of each DCT frequency independently. Denoting by $x_{i,j}$ the value of DCT coefficient j at block i for a given image I , and $x'_{i,j}$ the corresponding value for a stego image I' with cover I , the histograms of $\{x_{i,j}\}$ and $\{x'_{i,j}\}$ are

preserved for all fixed j in the range $0 \leq j \leq 63$. In order to do that, it is necessary that the message bit stream to be embedded in the j -coefficients have the same memoryless empirical distribution as $\{lsb(x_{i,j})\}$, where $lsb(x)$ denotes the least significant bit of x . This is done by assuming that the input message b has approximately as many zeros as ones, and processing it with an entropy decoder designed for $P(b_i = 1) = \hat{P}(lsb(x_{i,j}) = 1)$, the latter denoting the mentioned empirical distribution of the least significant bit of the j -th DCT coefficient (see, e.g., [11] for the use of arithmetic decoders for similar purposes). The value of this probability is included with the coded data, to allow for lossless decoding of the hidden data. In [8], the authors showed certain weakness of the HPDM algorithm with respect to the stego-analysis in [6, 7] (which is based on statistics of wavelets coefficients), and observed that it could be avoided by not modifying coefficients with values 0,1 and -1. This modification constitutes basically the MHPDM algorithm that we use in our experiments.

3 Models of Natural Images

Our experiments are based on statistics of wavelet coefficients and block DCT coefficients, and on three recently proposed statistical models of natural images. These models, which are briefly described below, reflect in general properties that are more global than those used in earlier stego-analysis works. For wavelet and DCT coefficient, although we do not test an explicit full-fledged model, we investigate relations between various particular statistics, and how they are altered by the introduction of stego information.

3.1 Areas of Connected Components Model

In [12, 13], it is observed that the distribution of the areas of connected components of bilevel (thresholded) images follow a power law which depends on just two parameters, an exponent α and a scaling factor C . More precisely, consider an image I whose gray levels are between 0 and N . For an integer k , define the bilevel (thresholded) images

$$I_l(i, j) = \begin{cases} 1 & \text{if } (l-1)\frac{N}{k} \leq I(i, j) \leq l\frac{N}{k}, \\ 0 & \text{otherwise.} \end{cases}$$

In [12, 13], the authors found that the total number $f(a)$ of connected components of the bilevel images I_l with area a is

$$f(a) \approx Ca^\alpha$$

Furthermore, it was experimentally found and theoretically justified [13] that the exponent α is close to -2 for natural images. We refer to this model as the *Areas Model*. We should note that this is a strongly non-local statistical model, since it looks at areas and at all bilevel images simultaneously. This is in sharp contrast with models based on individual pixels statistics, which were common in earlier works.

3.2 Adjacent Pixel Values Model

In [14, 15, 16], a statistical model for the horizontal derivative $I_x = \frac{\partial I}{\partial x}$ of an image I is introduced. Based on the *transported generator model* [17], the authors model an image as a random number of profiles of the same object, and each pixel is obtained as a linear combination of these profiles, weighted randomly. Mathematically,

$$I(z) = \sum_i a_i g(z - z_i) \tag{1}$$

where z and z_i are coordinates in \mathbb{R}^2 denoting a pixel location and an object profile location respectively, g is the profile of an object, and the coefficients a_i are random weights. Locations z_i are modeled as samples from a 2D Poisson process with uniform intensity, and weights a_i are modeled as independent and identically distributed (IID), also independent of the z_i -s.

Under this model and certain assumptions on $u(z) = \sum_i g_x^2(z - z_i)$, the authors show that the probability density function of I_x is

$$f(t) = \frac{1}{\sqrt{\pi}\Gamma(p)} \left(\frac{c}{2}\right)^{-\frac{p}{2}-\frac{1}{4}} (2)^{-p+\frac{1}{2}} t^{p-\frac{1}{2}} K_{p-\frac{1}{2}}\left(\sqrt{\frac{2}{c}}t\right), \tag{2}$$

where K is the modified Bessel function, Γ is the Gamma function, and p and c are two parameters referred to as *shape parameter* and *scale parameter* respectively. Furthermore, they show that p and c satisfy

$$p = \frac{3k_1^2}{k_2}, \quad c = \frac{k_2}{3k_1} \quad (3)$$

where

$$k_1 = E[I_x^2], \quad k_2 = E[I_x^4] \quad (4)$$

and expectation is taken according to the distribution of I_x which depends on (1).

Equation 2 applies as long as $I(z)$ is randomly generated according to Equation 1 with g, a_i , and z_i as described. Each image is modelled as a realization of a different random process $I(z)$, thus, parameters p and c may be different for each image. We will refer to this model as the *PC Model*.

3.3 Laplacian Distribution Model

In [18], the author reports on an empirically observed property of natural images referred to as *Differentially Laplacian*. It is observed that for a reasonably small constant k , and any fixed set of k^2 coefficients adding up to 0, the linear combination of k^2 pixel intensities in a $k \times k$ square, using these k^2 coefficients as weights, tends to exhibit a Laplacian-like distribution for natural images (this is related to, and generalizes, the well known Laplacian distribution of prediction errors in image coding [19]).

3.4 Wavelet Coefficients Model

We investigated relations among wavelet coefficients of natural images that may be altered by embedding a stego message. We explored estimations of mean, standard deviation, skewness, and kurtosis of several statistics calculated from Haar wavelet coefficients. We experimented with differences and sums of pairs of coefficients taken from horizontal, vertical and diagonal wavelet bands. In particular, denoting by $h_{i,j}$ a coefficient in the horizontal band of the first level decomposition of a $N \times M$ image, we found that the estimated kurtosis of $h_{i,j+1} - h_{i,j}$ with $0 \leq i < N/2, 0 < j < M/2$, is consistently altered for MHPDM stego images as described in section 4.2.4.

3.5 DCT Coefficients Model

An additional area that we explored was the information from higher order joint statistics of DCT coefficients. We consider the collection of 64-dimensional vectors obtained by applying the DCT on 8×8 blocks of an image, and taking the absolute value of the resulting transform coefficients. We look at the absolute values of the coefficients in each 8×8 DCT block as a vector in \mathbb{R}^{64} . Each image of size $N \times M$ brings $\frac{N}{8} \frac{M}{8}$ sample vectors. We study whether the joint distribution of DCT coefficients is affected by a stego algorithm by means of statistics of the form $w \cdot v'$ where $v \in \mathbb{R}^{64}$ is a vector of absolute values of DCT coefficients, v' is its transpose, and $w \in \mathbb{R}^{64}$ is a projection vector that results from a training process. Specifically, given a JPEG image and the same image with a stego message embedded, let $J = \{j_i\}, S = \{s_i\}, 0 \leq i < \frac{N}{8} \frac{M}{8}$, be the sets of sample vectors in \mathbb{R}^{64} obtained from each image respectively. We compute a vector $w \in \mathbb{R}^{64}$ that maximizes the empirical correlation $w = \operatorname{argmax}\{\hat{\rho}(w \cdot v', I_v)\}$, where v is a sample taken from J or S , and I_v is valued 1 or -1 when v is a taken from J or S respectively. Averaging uniformly vectors w computed for several pairs of training images, we seek assigning a high weight to DCT coefficients that aid classification for many images whereas others would receive low weights. We denote W the average projection vector $W = \operatorname{mean}(w)$.

4 Experimental Results

4.1 Experimental Setting

For all experiments we used gray-scale 1536×1024 images from the popular Van Hateren's data base.¹ This is particularly appropriate for our experiments as it is comprised of numerous natural images from outdoor scenes, and the models we experiment with are specified for natural gray-scale images. This set of images is widely accepted as representative of the class of natural images (see e.g. [20]). The 12-bits pixel values of all images are proportional to the light intensities in the scenes; however, the multiplying constant need not be the same for different images. In experiments where this disparity might affect the statistics of interest, we follow [20], and use *log-contrast* images. In the log-contrast image of I , the pixel at location (i, j) is calculated as $\log^+(I(i, j)) - E(\log^+(I))$,

¹<http://hlab.phys.rug.nl/imlib/index.html>

where $\log^+(x) = \log(x + 1)$, and $E(f(I))$ denotes the arithmetic mean of $f(I(i, j))$ when (i, j) ranges over all pixel coordinates in the image.² Cases where log-contrast was used will be explicitly identified in the sequel.

We experimented with a subset, which will be denoted \mathcal{I} , of 1400 images from the Van Hateren’s data base. From this set of images we generated Jsteg and MHPDM stego images by first reducing the number of gray levels to a maximum of 256 (scaling by $255/\max(I)$ and rounding) and then compressing with JPEG and embedding a random message in JPEG DCT coefficients during the process.³ For MHPDM in particular, the message satisfied $P(b_i = 1) = \hat{P}(\text{lsb}(x_{i,j}) = 1)$ for every coefficient index j , $1 \leq j \leq 20$. The amount of information embedded was always the maximum allowed by the image, i.e., a message as long, in bits, as the number of coefficient values suitable for modification according to the stego algorithm. When we used S-Tools to generate stego data, we also started from a 256 gray level version of the original image and adjusted the length of the embedded message to avoid visually perceptible artifacts. The amount of information embedded in an image in this case was significantly smaller than for the Jsteg or MHPDM counterparts. The S-tools images were always converted back to gray level images before computing statistics, by working on the image formed by the rounded luminance values.

Since JPEG lossy compression may affect image statistics, when analyzing results for Jsteg and MHPDM we always compare stego images to clean *JPEG images*, i.e., images with no message embedded but that have been lossily compressed with JPEG (again reducing the number of gray levels to a maximum of 256 and using the same software and settings as for Jsteg and MHPDM). Similarly, we use the term *bitmap image* to refer to an image with no information embedded but whose number of gray levels has been reduced to a maximum of 256.

Some experiments rely on estimations of mean (μ), standard deviation (σ), skewness (γ_1) and kurtosis (β_2) of a random variable X based on observed samples $x_1 \dots x_n$. The skewness and kurtosis of X are defined (see e.g. [21]) as

$$\gamma_1 = \frac{E(x - \mu)^3}{\sigma^3}; \quad \beta_2 = \frac{E(x - \mu)^4}{\sigma^4}$$

² \log^+ is used to avoid problems with the logarithm of zero. The slight effect of this bias on eliminating the constant multiplier of the light intensity is secondary for the cases of interest.

³Jsteg and our implementation of MHPDM are both based on source code from the Independent JPEG Group’s JPEG software, <http://www.ijg.org/>, with the *quality setting* parameter fixed at 75%.

We use estimators respectively calculated as

$$\bar{\mu} = \frac{\sum_{i=1}^n x_i}{n}; \quad \bar{\sigma} = \left(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{\mu})^2\right)^{1/2}; \quad \bar{\gamma}_1 = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{\mu})^3}{\bar{\sigma}^3}; \quad \bar{\beta}_2 = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{\mu})^4}{\bar{\sigma}^4},$$

where x_i ranges over all sample values of interest.

4.2 Experiments

We now describe several experiments involving the different stego algorithms and the natural image statistical models described above. We also present some additional experiments targeting MHPDM stego-analysis in particular. In this case, we include also an analysis of wavelet and DCT coefficients.

4.2.1 Areas Model Parameters

We explore the effect of stego algorithms on the values (α, C) of the Areas Model parameters. We observe that the power law holds in bitmap, JPEG, and stego images and, although the parameter values are often modified for individual images, they generally remain in the (relatively large) range of values observed for natural images as we can appreciate in Figure 1. Similarly, in Figure 2 we observe that also the joint distribution of parameter values is similar for natural and stego images. Thus, the variation does not allow us to clearly distinguish between natural and stego images. Moreover, there is not a consistent bias effect, in contrast with the other models we discuss below; this characterization of natural images is mostly “randomly” modified by the stego process. Figure 3 shows the distribution of connected components areas of a particular image from \mathcal{I} as bitmap, JPEG, and covering a message embedded with Jsteg and S-Tools. We observe that the plots are very close and the values of the exponent α for the best linear fitting in each case are $-2.09, -2.06, -2.05,$ and $-2.06,$ respectively. Figure 4 shows, enclosed in a rectangular frame, parameters values for a particular image from \mathcal{I} as bitmap, JPEG, and covering a message embedded with MHPDM, Jsteg and S-Tools, together with a cloud of points obtained plotting parameters values for a subset \mathcal{I}_{1000} of 1000 JPEG images from \mathcal{I} . The variation resulting from embedding a message is rather small as compared to the universe of observed values. Figure 5 shows the movement from a point representing parameter values for a clean JPEG image to a point for the same image bearing a message embedded using MHPDM. The plot shows the movement for a subset \mathcal{I}_{200} of 200 JPEG images from \mathcal{I} and their corresponding MHPDM stego images. We

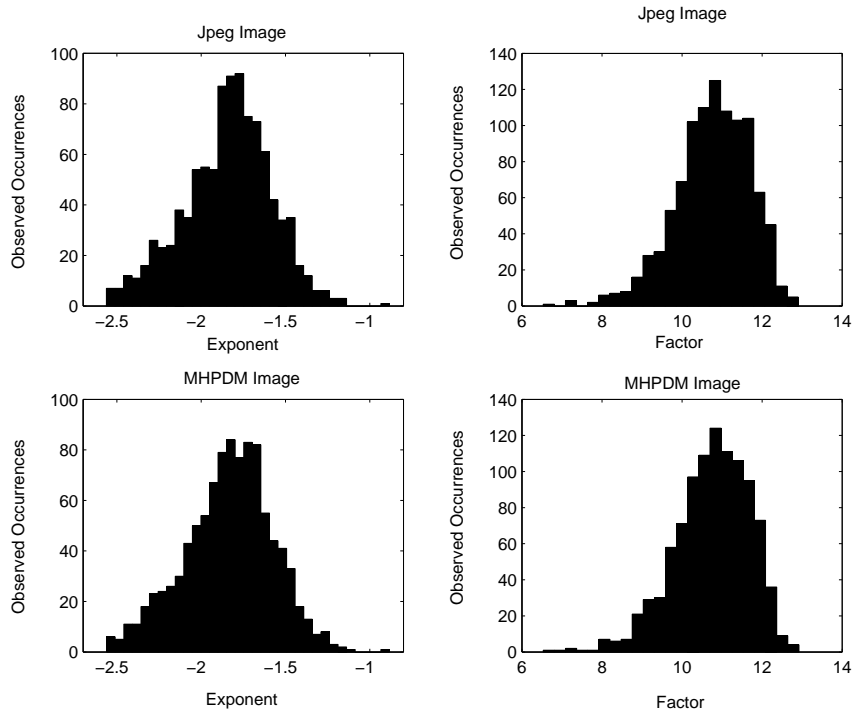


Figure 1: *Histogram of Areas Model parameters (exponent and multiplying factor in the power law) for JPEG images and MHPDM stego images. We observe that the distributions of parameter values is similar in both classes*

observe that there is no clear bias effect. Over a set of 1400 images from \mathcal{I} the percentage of images for which parameter α was incremented as a result of embedding a message was approximately 60% and similarly, approximately 64% of the images were affected by an increment of parameter C .

4.2.2 PC Model Parameters

This model has been found to fit accurately the distribution of differences between adjacent pixels. Given an image I , one can approximate I_x as a difference \hat{I}_x between adjacent pixel values.⁴ Then, empirical estimates of k_1 and k_2 , as defined in (4), are obtained by taking the arithmetic mean of \hat{I}_x over all suitable pixels x . Finally, estimates for p and c are computed according to (3), which can be substituted into (2) to obtain an estimate of $f(t)$. One can then check the fit of the distribution of I_x predicted by this estimate of (2) with the actual empirical distribution of values \hat{I}_x , and measure possible deterioration of this fit due to the introduction of stego data. Figure 6 shows the model

⁴Although there exist better approximations for I_x (see, e.g., [22]), taking pixel differences is accurate enough for this model to fit well.




Figure 2: *Joint histogram of Areas Model parameters (exponent and multiplying factor in the power law) for JPEG images and MHPDM stego images. Darker areas have a higher concentration of pairs of parameter values. We observe that stego images present a high concentration in regions where also natural images do.*




Figure 3: *Distribution of connected components areas for four versions of the same image, with and without hidden message. We observe that the exponential distribution is observed both by the original and the stego images, thereby limiting the use of this model for stego-analysis.*




Figure 4: *Cloud of Areas Model parameters (exponent and multiplying factor in the power law) values for JPEG images and the effect of hiding information on one particular image. Note that the variation due to the hiding process is rather small compared to the intrinsic variability of the parameters for this class of natural images.*




Figure 5: *Movement of Areas Model parameters (exponent and multiplying factor in the power law) values which results from embedding a stego message using MHPDM stego algorithm. Each dot in the plot represents a pair of parameter values for a MHPDM stego image and each segment connects the points given by parameter values for a JPEG image and the same image bearing a MHPDM stego message.*

fit for a given JPEG image, and independently for the same image including a message embedded with MHPDM, Jsteg and S-Tools, respectively. As observed in the figure, Jsteg and MHPDM do not produce a noticeable departure from the model. However, the algorithm from S-tools does, and an image bearing a message embedded using this algorithm can easily be detected by observing the histogram of differences between adjacent pixels and its discrepancy with the model, as estimated for the stego image.

Figure 7 shows parameters values for the same four variations of the same image as Figure 6, and also the bitmap representation, immersed in a cloud of points obtained for parameters values of the subset \mathcal{I}_{1000} of 1000 JPEG images from \mathcal{I} . Except for the values obtained for S-Tools, the rest, enclosed in a rectangular frame, show small differences as compared to the range of different values observed on JPEG images.

A closer examination of the effect on the whole data set \mathcal{I}_{1000} reveals that the parameter p is altered in a consistent direction by the MHPDM algorithm, i.e., in more than 95% cases of 1000 pairs of MHPDM / JPEG images from \mathcal{I}_{1000} , the stego algorithm causes an increase in the value of p . A histogram of relative differences of parameter p (the difference divided by the value of p for the JPEG image) is shown in Figure 8 where we observe that practically all values are positive. This consistent bias indicates a potential weakness of MHPDM with respect to stego-analysis based on this model. However, Figure 10 which plots false alarm probability vs. hit probability obtained for 1000 JPEG images and 1000 MHPDM stego images from \mathcal{I}_{1000} varying a threshold, shows a poor classification performance. The shift is not large enough to achieve significant discrimination for this class of images, as can be appreciated in Figure 8, showing relative differences smaller than 5% in most cases. This can also be observed in Figure 9 which shows the movement from a point representing parameter values for a clean JPEG image to a point for the same image bearing a message embedded using MHPDM. The plot shows the movement for a subset \mathcal{I}_{200} of 200 JPEG images from \mathcal{I} and their corresponding MHPDM stego images. It is noticeable that variation in parameter p tends to be larger for smaller values of c . This effect can be observed in Figure 9 and was corroborated for the larger set \mathcal{I}_{1000} . The indistinguishability of parameter values for both classes can also be appreciated in Figure 11 and Figure 12, showing very similar histograms of both parameters for 1000 JPEG images and 1000 MHPDM stego images from \mathcal{I}_{1000} independently and

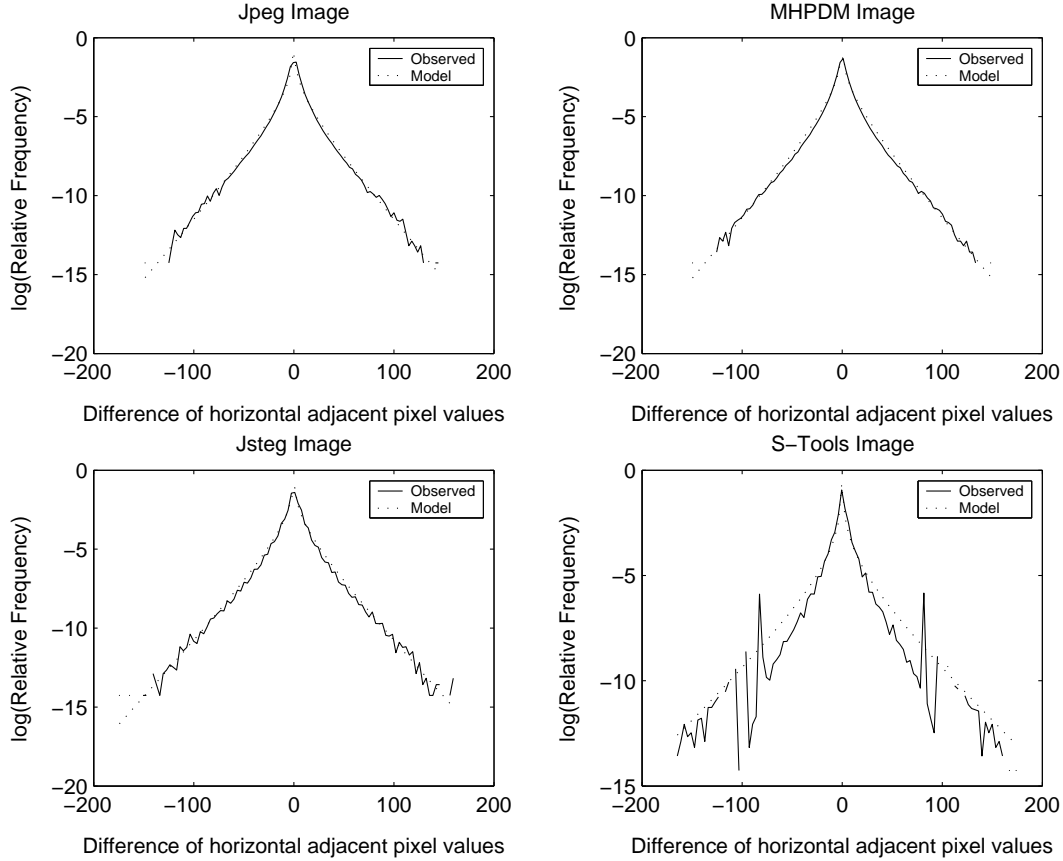


Figure 6: *PC Model fit independently to four versions of the same image. When the message is embedded using the S-Tools algorithm, this can be easily detected due to its discrepancy with the model. For MHPDM and Jsteg stego algorithms, there is a good fit to the natural images model.*

jointly respectively.

4.2.3 Differentially Laplacian Model

For the Differentially Laplacian Model experiments we select k^2-1 coefficients pseudo-randomly with a uniform distribution in the interval $(-1,1)$ and choose one more coefficient so that the overall coefficient sum is zero. As previously observed for the PC Model, the fit of the Differentially Laplacian Model does not deteriorate significantly when hidden data is embedded. This was the case observed for several values of parameter k and different images.

Also, for a fixed linear combination T , if we denote by $T(I) = \{T(block_{i,j}(I))\}$ where $block_{i,j}(I)$ varies along all $k \times k$ blocks of a partition of I , estimations of mean, standard deviation and kurtosis of $T(I)$ do not aid in the classification process.




Figure 7: *Cloud of PC Model parameters values for JPEG images and the effect of hiding information on one particular image. Enclosed in a frame are values for bitmap, JPEG, Jsteg and MHPDM versions of the same image. The value for the same image processed with S-Tools, outside the frame, clearly shows that S-Tools produces a large deviation on parameter values, while MHPDM and Jsteg maintain relatively small distances.*




Figure 8: *Histogram of relative differences between parameter p of PC Model for a MHPDM image and its corresponding JPEG image. We observe that values of relative differences (difference divided by the value of p for JPEG image) are mostly smaller than 5% and practically all values are positive. This consistent bias of parameter p , although relatively small, indicates a potential weakness of MHPDM with respect to stego-analysis based on this model*

Figure 13 shows a log normalized histogram of the values calculated for a fixed 5×5 linear combination of pixels values, on a JPEG image and the same image including a message embedded with MHPDM, Jsteg, and S-Tools. The four plots are very similar.

4.2.4 Statistics of Wavelet Coefficients

In this subsection, we analyze statistics from Wavelet Coefficients Model. The investigation focused on the MHPDM algorithm and used log-contrast images. We found that the estimated kurtosis of $h_{i,j+1} - h_{i,j}$, the difference between adjacent coefficient in the horizontal band of the first level decomposition, is consistently altered for stego images. Stego images showed a higher kurtosis than their corresponding JPEG images in more than 95% cases of the set \mathcal{I}_{1000} of 1000 pairs of images from \mathcal{I} . However, the kurtosis variability in this class of natural images is once again quite large, and it seems difficult to fix a threshold that could reliably discriminate between the two groups.

Figure 14 shows the estimated kurtosis of $h_{i,j+1} - h_{i,j}$ for 20 JPEG and MHPDM stego images from \mathcal{I} . Crosses representing kurtosis of stego images appear always above dots corresponding to JPEG images, but it seems difficult to choose a threshold that would separate the two series of




Figure 9: *Movement of PC Model parameters values which results from embedding a stego message using MHPDM stego algorithm. Each dot in the plot represents a pair of parameter values for a MHPDM stego image and each segment connects the points given by parameter values for a JPEG image and the same image bearing a MHPDM stego message.*

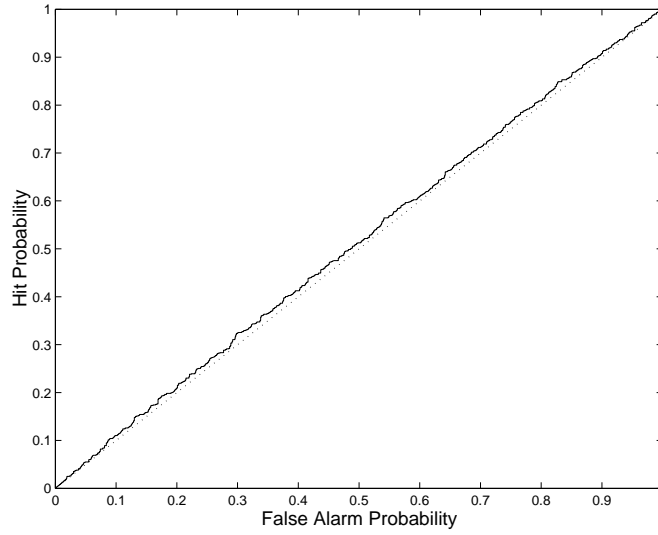


Figure 10: *Classification performance using a threshold for parameter p of PC Model. The plot is obtained varying a threshold that is compared to parameter p for 1000 JPEG images and 1000 MHPDM stego images. The reason for the poor classification performance is that, although consistently positive, the shift produced on parameter p by the stego algorithm is relatively small.*

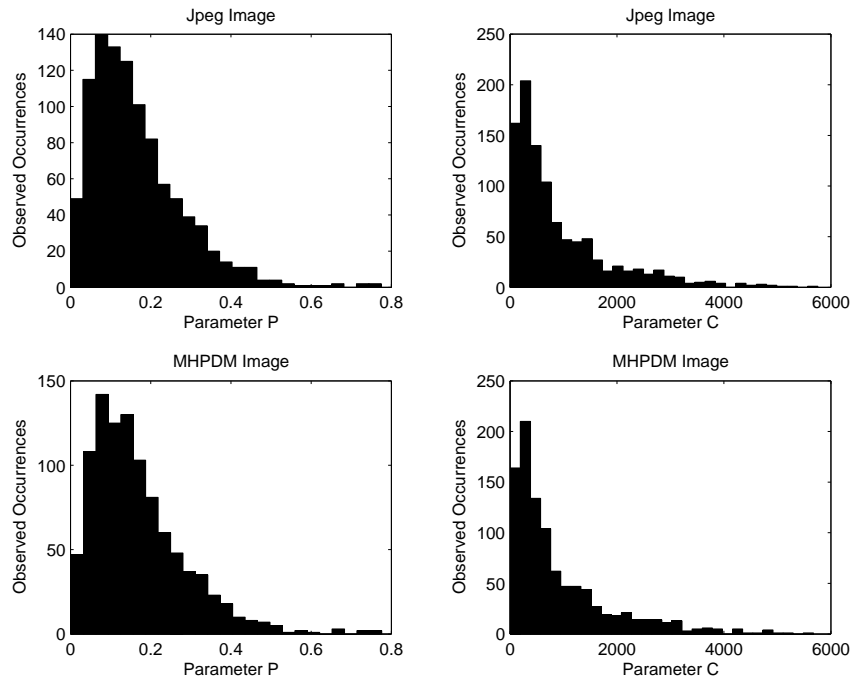


Figure 11: *Histograms of PC Model parameters for JPEG and MHPDM images. We observe that the parameters distributions are similar for natural and stego images.*

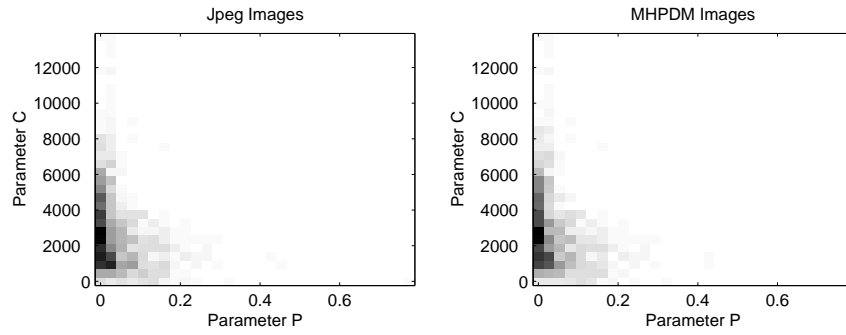


Figure 12: *Joint histogram of PC Model parameters for JPEG images and MHPDM images. Darker areas have a higher concentration of pairs of parameter values. We observe that stego images present a high concentration in regions where also natural images do.*

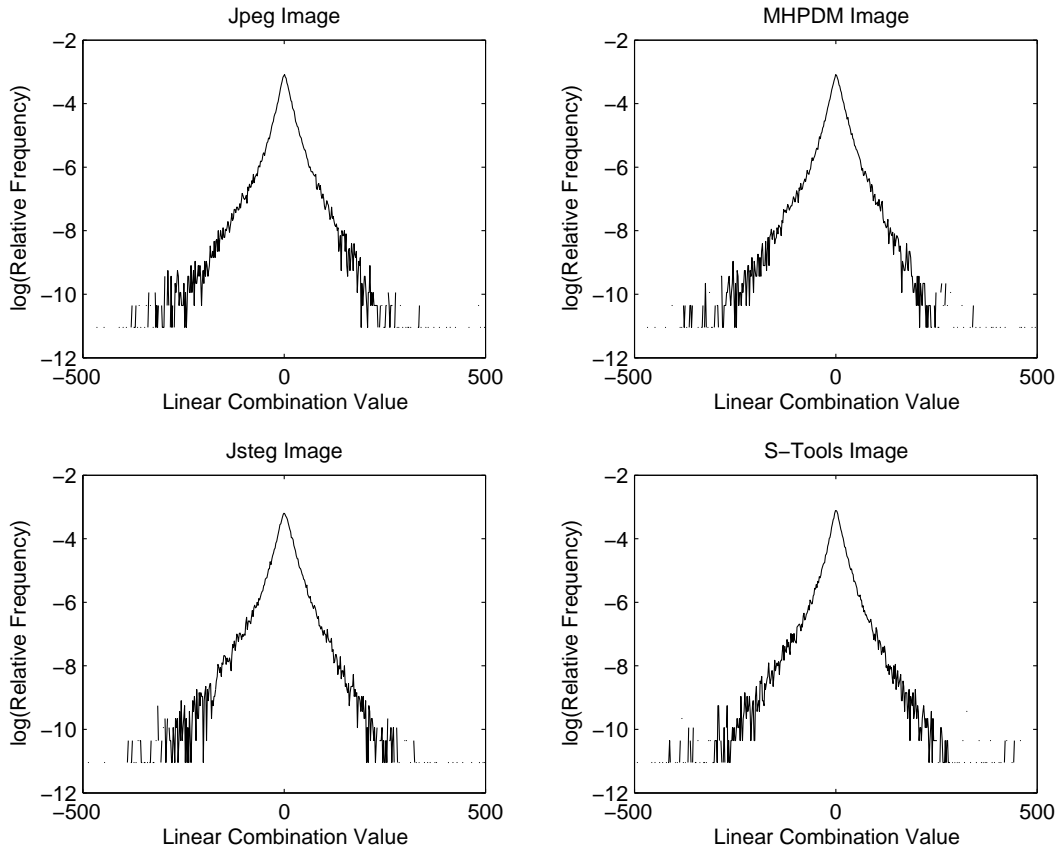


Figure 13: *Distribution of linear combination values for the Differentially Laplacian Model for the four classes of images. All log normalized histograms were calculated using a fixed linear combination of pixels values in a 5×5 neighborhood. We observe very similar distributions, indicating that the model is not powerful enough to detect stego images.*




Figure 14: *Estimated kurtosis of differences of adjacent horizontal Haar wavelet coefficients shown for 20 JPEG images and the same images with a message embedded with MHPDM stego algorithm. We observe a clear bias in the stego image.*

values precisely.

We point out that, although similar wavelet statistics did not produce better results, this consistent bias can be regarded as a weakness of MHPDM and further research may exploit more sophisticated wavelet statistics seeking better results.

4.2.5 Comparing DCT Coefficients

We applied the DCT Coefficients Model to MHPDM stego algorithm using log-contrast images. The fact that the histogram of each coefficient is preserved separately by MHPDM opens the possibility that some joint distribution might be altered, thus aiding in stego-analysis.

We recall that under this model, we compute a projection vector $W \in \mathbb{R}^{64}$ from the space of 64 absolute values of DCT coefficients to a one dimensional space, seeking projected values highly correlated with a variable indicating whether a vector of coefficients belongs to a natural or a stego image. For this experiment, we classify an image I by calculating the arithmetic average $mean\{W \cdot v_i'\}$, where $v_i \in \mathbb{R}^{64}$ ranges over vectors of absolute values of DCT coefficients of I , and finally using a threshold for the decision that is fixed according to a trade off between false alarms and hit probabilities (i.e. respectively the probability of incorrectly classifying a natural image as stego and the probability of correctly classifying a stego image as such). The averaged projections were consistently higher for stego images than their corresponding JPEG images in




Figure 15: *Projection of DCT coefficients shown for 60 JPEG images and the same images with a message embedded with MHPDM stego algorithm. Although the values are similar for the natural and stego images, there is a consistent bias, the value for the stego image being always larger than the one for the corresponding natural one.*

more than 99% cases of a subset \mathcal{I}_{test} of 1000 pairs of test images from \mathcal{I} with a training subset \mathcal{I}_{train} of 400 pairs of images from \mathcal{I} . Results for a subset of the testing set are shown in Figure 15 where crosses representing values for stego images appear always above circles representing JPEG images. This bias once again indicates a clear modification by MHPDM of the statistics of natural images. However, once again the variability for this large class of natural images is significant and it seems difficult to fix an absolute threshold that would work well for most pairs at the same time. However, as shown in Section 4.2.6, when partial knowledge of the stego scheme is incorporated into the analysis, the statistical changes produced by the hidden message allow for significant discrimination between natural and stego images.

4.2.6 Coefficients Correlations Estimation: Exploiting algorithm knowledge in stego-analysis

Empirical correlations of DCT coefficients vary considerably among natural images. However, it is also possible to look at empirical correlations between empirical correlations for different pairs of coefficients. That is, images that have high correlation between coefficients, say a and b , might also have high correlation between a different pair of carefully chosen coefficients a' and b' , with high probability. This fact can be exploited particularly for the MHPDM algorithm if we consider that only coefficients with indices 1 through 20 are modified, i.e. we strongly use our knowledge about MHPDM to stego-analyze it. Based on a set of log-contrast training images, for each pair of absolute values of DCT coefficients a and b in $A = \{1..20\}$, we get an estimation $\hat{\rho}(|a|, |b|)$ of $\rho(|a|, |b|)$ (the empirical correlation between $|a|$ and $|b|$) based on the empirical correlations between pairs of absolute values of DCT coefficients taken from the set $B = \{0, 21..63\}$ and use $\hat{\rho}(|a|, |b|) - \rho(|a|, |b|)$ as a feature for classification. To calculate $\hat{\rho}(|a|, |b|)$, we determine the projection from the vector v of values $v_i = \rho(|a'|, |b'|)$ to a one-dimensional space that maximizes the empirical correlation with $\hat{\rho}(|a|, |b|)$. The set of pairs of coefficients (a', b') is the set of all possible pairs of coefficients from a subset $B' \subset B$, where highly quantized coefficients are discarded. Once this projection w is determined, we use a linear fitting from $w \cdot v'$ to $\hat{\rho}(|a|, |b|)$ over the set of training images and use this polynomial to calculate $\hat{\rho}(|a|, |b|)$. Having determined the estimator $\hat{\rho}$ of ρ for all pairs $(|a|, |b|)$ we can calculate features $\hat{\rho}(|a|, |b|) - \rho(|a|, |b|)$ for the set of training images and determine a projection from the space of features to a one dimensional space that maximizes the empirical correlation with a variable valued 1 for natural images and -1 for stego images. Classifying an image consists of comparing the projection of its features with a given threshold, which is chosen to determine an operating point in the “hit/false alarm” plane, as described below. This technique achieved the best classifications results. Figure 16 shows false alarm probability vs. hit probability for an experiment on a subset $\mathcal{I}_{train'}$ of 800 training pairs of JPEG/stego images from \mathcal{I} and a subset $\mathcal{I}_{test'}$ of 600 test pairs from \mathcal{I} . We observe that the test described achieves a significant classification performance.

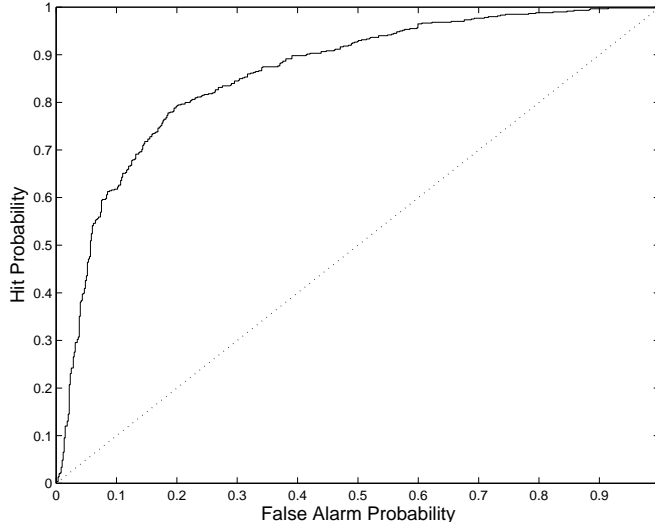


Figure 16: *Classification performance using Coefficients Correlations Estimation. The graph shows that this technique, that uses information about the stego algorithm, can detect MHPDM stego images with significant accuracy.*

5 Conclusions

We have studied the effect of applying popular steganography algorithms on different statistical models of natural images. On one hand, we observed that some popular stego algorithms consistently bias these statistics for some of the most fundamental models. On the other hand, the intrinsic variability of these statistics is so high, for the class of images studied, that this bias induced by hiding “unnatural” information is not sufficient in general to move the results outside of the “natural” range, unless knowledge of the embedding algorithm is available and exploited. The best classification results were obtained in the latter case.

These experimental results lead us to conclusions in two directions. First, regarding faithful models of natural images, it seems that the reported efforts so far are not sufficient to clearly exclude some “non-natural” images, for example those obtained by artificially embedding hidden messages. Thus, there seems to be a need for further refinement of these models. Second, in the stego arena, it is obvious that stego-analysis is a “cat and mouse” game: Knowing the stego algorithm, a technique can be devised to attack it; and knowing the attack, the stego algorithm can be further modified to mislead the detection procedure. An example is given by Farid’s stego-analysis approach [6, 7], which was overcome by MHPDM, which in turn, seems to be broken by the

results in Section 4.2.6. It would therefore be desirable to have a more fundamental approach to the stego capacity in natural images, preferably based on universal properties and independent of the particular algorithm of choice. Some analysis has been done in this direction in [23, 24, 25, 26, 27]. An approach based on universal modeling and simulation [28, 11, 29, 30, 31] is currently being pursued. Results on this approach are reported elsewhere [32].

References

- [1] A. Brown, “S-Tools for Windows.” Shareware, <ftp://ftp.ntua.gr/pub/encrypt/mirrors/idea.sec.dsi.unimi.it/code/s-tools4.zip>, 1994.
- [2] D. Upham, “JPEG-JSTEG - Modifications of the independent JPEG groups JPEG software for 1-bit steganography in JFIF output files.” <ftp://ftp.funet.fi/pub/encrypt/steganography/>.
- [3] N. F. Johnson, Z. Duric, and S. G. Jajodia, *Information Hiding : Steganography and Watermarking - Attacks and Countermeasures*. Kluwer Academic, Feb. 2001.
- [4] P. Wayner, *Disappearing Cryptography, Second Edition - Information Hiding: Steganography and Watermarking*. Morgan Kaufmann, Apr. 2002.
- [5] I. Avcibas, N. Memon, and B. Sankur, “Steganalysis using image quality metrics,” *IEEE Transactions on Image Processing*, vol. 12, pp. 221–229, Feb. 2003.
- [6] H. Farid, “Detecting hidden messages using higher-order statistical models,” in *Proc. ICIP2002*, vol. 2, (Rochester, NY), pp. II–905–II–908, 2002.
- [7] H. Farid, “Detecting hidden messages using higher-order statistics and support vector machines,” in *Proc. of the 5th International Workshop of Information Hiding*, (Noordwijkerhout, The Netherlands), Oct. 2002.
- [8] R. Tzschoppe, R. Bäuml, J. B. Huber, and A. Kaup, “Steganographic system based on higher-order statistics,” in *Proc. SPIE Vol. 5020, Security and Watermarking of Multimedia Contents V*, (Santa Clara, CA), 2003.

- [9] J. J. Eggers, R. Bäuml, and B. Girod, “A communications approach to image steganography,” in *Proc. SPIE Vol. 4675, Security and Watermarking of Multimedia Contents IV*, (San Jose, CA), Jan. 2002.
- [10] J. L. Mitchell and W. B. Pennebaker, *JPEG Still Image Data Compression Standard*. Van Nostrand Reinhold, 1993.
- [11] T. S. Han and M. Hoshi, “Interval algorithm for random number generation,” *IEEE Transactions on Information Theory*, vol. 43, pp. 599–611, Mar. 1997.
- [12] Y. Gousseau and J. M. Morel, “Are natural images of bounded variation,” *SIAM Journal on Mathematical Analysis*, vol. 33, no. 3, pp. 634–648, 2001.
- [13] L. Alvarez, Y. Gousseau, and J. M. Morel, “The size of objects in natural images.” CMLA preprint, Ecole Normale Sup. -Cachan, 1999.
- [14] U. Grenander and A. Srivastava, “Probability models for clutter in natural images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, pp. 424–429, Apr. 2001.
- [15] A. Srivastava, X. Liu, and U. Grenander, “Universal analytical forms for modeling image probabilities,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 1200–1214, Sept. 2002.
- [16] U. Grenander, “Towards a theory of natural scenes,” tech. rep., Brown University, Providence, RI, 2003. <http://www.dam.brown.edu/ptg/publications.shtml>.
- [17] U. Grenander, M. I. Miller, and P. Tyagi, “Transported generator clutter models,” monograph of center for imaging sciences, Johns Hopkins University, 1999.
- [18] M. L. Green, “Statistics of images, the TV algorithm of Rudin-Osher-Fatemi for image denoising and an improved denoising algorithm,” tech. rep., University of California, Los Angeles, CA, Oct. 2002. <http://www.math.ucla.edu/imagers/htmls/reports.html>.
- [19] A. Netravali and J. O. Limb, “Picture coding: A review,” *Proc. IEEE*, vol. 68, pp. 366–406, 1980.

- [20] J. Huang and D. Mumford, “Statistics of natural images and models,” in *Proc. Computer Vision and Pattern Recognition-Volume 1*, (Fort Collins, Colorado), June 1999.
- [21] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 10th printing*, vol. 55 of *National Bureau of Standards Applied Mathematics*. Washington D.C.: U.S. Government Printing Office, Dec. 1972.
- [22] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D*, vol. 60, pp. 259–268, 1992.
- [23] C. Cachin, “An information-theoretic model for steganography,” *Lecture Notes in Computer Science*, vol. 1525, pp. 306–318, 1998.
- [24] P. Moulin and J. A. O’Sullivan, “Information-theoretic analysis of information hiding,” *IEEE Transactions on Information Theory*, vol. 49, pp. 563–593, Mar. 2003.
- [25] P. Moulin and M. K. Mihcak, “The parallel-gaussian watermarking game.” To appear in *IEEE Transactions on Information Theory*, Feb. 2004, June 2001.
- [26] P. Moulin, M. K. Mihcak, and G.-I. Lin, “An information–theoretic model for image watermarking and data hiding,” in *Proc. ICIP2000*, (Vancouver, B.C.), Sept. 2000.
- [27] P. Moulin and M. K. Mihcak, “A framework for evaluating the data-hiding capacity of image sources,” *IEEE Transactions on Image Processing*, vol. 11, pp. 1029–1042, Sept. 2002.
- [28] T. S. Han and S. Verdú, “Approximation theory of output statistics,” *IEEE Transactions on Information Theory*, vol. 39, pp. 752–772, May 1993.
- [29] Y. Steinberg and S. Verdú, “Simulation of random processes and rate-distortion theory,” *IEEE Transactions on Information Theory*, vol. 42, pp. 63–86, Jan. 1996.
- [30] N. Merhav and M. J. Weinberger, “On universal simulation of information sources using training data,” *IEEE Transactions on Information Theory*, vol. 50, Jan. 2004.
- [31] G. Seroussi, “Universal types and simulation of individual sequences.” *Proc. LATIN’2004*, Buenos Aires, Argentina, April 2004.

- [32] A. Martin, G. Seroussi, and G. Sapiro, “A universal approach to steganography.” in preparation.