

Higher-order Wavelet Statistics and their Application to Digital Forensics

Hany Farid[†] and Siwei Lyu

Abstract

We describe a statistical model for natural images that is built upon a multi-scale wavelet decomposition. The model consists of first- and higher-order statistics that capture certain statistical regularities of natural images. We show how this model can be useful in several digital forensic applications, specifically in detecting various types of digital tampering.

1. Introduction

The space of possible images is enormous - there are 256^{N^2} possible 8-bit grayscale images of size $N \times N$ pixels (with as few as $N = 10$ pixels, there are a whopping 1.3×10^{154} possible images). And yet, natural photographic images occupy only a relatively tiny portion of this space [7, 15, 4, 21]. One might expect, therefore, for natural images to exhibit statistical regularities that distinguish them from the sea of all possible images. Popular examples of statistical models include those based on power spectra [18, 7, 21], Markov random fields [11, 3, 10], or wavelets [17, 22].

The decomposition of images using basis functions that are localized in spatial position, orientation, and scale (e.g., wavelets) have proven extremely useful in image compression, image coding, noise removal, and texture synthesis. One reason is that such decompositions exhibit statistical regularities that can be exploited. In this paper we describe a statistical model for natural images that is built upon a multi-scale wavelet-like decomposition. The model consists of first- and higher-order statistics that capture the regularities that are inherent to natural images. We then show how this model differentiates between natural and un-natural images, and how it can be used in several digital forensic applications, specifically in detecting various types of digital tampering.

2. Image Statistics

The image decomposition employed here is based on separable quadrature mirror filters (QMFs) [24, 25, 23].

[†] Correspondence should be addressed to H. Farid (farid@cs.dartmouth.edu), 6211 Sudikoff Lab, Department of Computer Science, Dartmouth College, Hanover, NH 03755




Figure 1: An idealized multi-scale and orientation decomposition of frequency space. Shown, from top to bottom, are levels 0, 1, and 2, and from left to right, are the lowpass, vertical, horizontal, and diagonal subbands.

As illustrated in Figure 1, this decomposition splits the frequency space into multiple scales and orientations. This is accomplished by applying separable lowpass and highpass filters along the image axes generating a lowpass, vertical, horizontal, and diagonal subband. Subsequent scales are created by recursively filtering the lowpass subband. The vertical, horizontal, and diagonal subbands at scale $i = 1, \dots, n$ are denoted as $V_i(x, y)$, $H_i(x, y)$, and $D_i(x, y)$, respectively. Shown in Figure 2, for example, is a three-level decomposition of a “disc” image.

Given this image decomposition, the statistical model is composed of the mean, variance, skewness and kurtosis of the subband coefficients at each orientation and at scales $i = 1, \dots, n-1$. These statistics characterize the basic coefficient distributions. As shown in Figure 2, however, these statistics do not capture the strong correlations that exist across space, orientation, and scale. For example, if a large (small) coefficient is found at the finest scale this coefficient’s “parent”




Figure 2: Shown are the absolute values of the subband coefficients at three scales and three orientations for a “disc” image. The residual lowpass subband is shown in the upper-left corner.

and “grandparent” at the coarser scales tend also to be large (small). In order to capture some of these higher-order statistical correlations, we collect a second set of statistics that are based on the errors in a linear predictor of coefficient magnitude.

As described in [2], the subband coefficients are correlated to their spatial, orientation and scale neighbors. For purposes of illustration, consider first a vertical band, $V_i(x, y)$, at scale i . A linear predictor for the magnitude of these coefficients in a subset of all possible neighbors¹ is given by:

$$\begin{aligned}
 |V_i(x, y)| &= w_1|V_i(x-1, y)| + w_2|V_i(x+1, y)| \\
 &+ w_3|V_i(x, y-1)| + w_4|V_i(x, y+1)| \\
 &+ w_5|V_{i+1}(x/2, y/2)| + w_6|D_i(x, y)| \\
 &+ w_7|D_{i+1}(x/2, y/2)|,
 \end{aligned} \tag{1}$$

where $|\cdot|$ denotes absolute value and w_k are the scalar weighting values. This linear relationship is expressed more compactly in matrix form as:

$$\vec{V} = Q\vec{w}, \tag{2}$$

where the column vector $\vec{w} = (w_1 \dots w_7)^T$, the vector \vec{V} contains the coefficient magnitudes of $V_i(x, y)$ strung out into a column vector, and the

¹The particular choice of spatial, orientation and scale neighbors was motivated by the observations of [2] and modified to include non-casual neighbors.

columns of the matrix Q contain the neighboring coefficient magnitudes as specified in Equation (1) also strung out into column vectors. The coefficients are determined by minimizing the quadratic error function:

$$E(\vec{w}) = \|\vec{V} - Q\vec{w}\|^2, \tag{3}$$

where $\|\cdot\|$ is the vector two-norm. This error function is minimized by differentiating with respect to \vec{w} :

$$\frac{dE(\vec{w})}{d\vec{w}} = 2Q^T(\vec{V} - Q\vec{w}), \tag{4}$$

setting the result equal to zero, and solving for \vec{w} to yield:

$$\vec{w} = (Q^T Q)^{-1} Q^T \vec{V}. \tag{5}$$

Given the large number of constraints (one per image pixel) in only seven unknowns, it is generally safe to assume that the 7×7 matrix $Q^T Q$ will be invertible.

Once the coefficients of the linear predictor are estimated, the log error between the actual coefficients and the predicted coefficients is computed as:

$$\vec{E} = \log_2(\vec{V}) - \log_2(|Q\vec{w}|), \tag{6}$$

where the log is computed point-wise on each vector component. It is from this error that additional statistics are collected, namely the mean, variance, skewness, and kurtosis. This process is repeated for each vertical subband at scales $i = 1, \dots, n-1$, where at each scale a new linear predictor is estimated. A similar process is repeated for the horizontal and diagonal subbands. The linear predictor for the horizontal subbands is of the form:

$$\begin{aligned}
 |H_i(x, y)| &= w_1|H_i(x-1, y)| + w_2|H_i(x+1, y)| \\
 &+ w_3|H_i(x, y-1)| + w_4|H_i(x, y+1)| \\
 &+ w_5|H_{i+1}(x/2, y/2)| + w_6|D_i(x, y)| \\
 &+ w_7|D_{i+1}(x/2, y/2)|,
 \end{aligned} \tag{7}$$

and for the diagonal subbands:

$$\begin{aligned}
 |D_i(x, y)| &= w_1|D_i(x-1, y)| + w_2|D_i(x+1, y)| \\
 &+ w_3|D_i(x, y-1)| + w_4|D_i(x, y+1)| \\
 &+ w_5|D_{i+1}(x/2, y/2)| + w_6|H_i(x, y)| \\
 &+ w_7|V_i(x, y)|.
 \end{aligned} \tag{8}$$

The same error metric, Equation (6), and error statistics computed for the vertical subbands, are computed for the horizontal and diagonal bands, for a total of $12(n-1)$ error statistics. Combining these statistics with the $12(n-1)$ coefficient statistics yields a total of $24(n-1)$ statistics that form a feature vector of “natural images”.





Figure 3: Natural and un-natural images. From top to bottom are: natural, noise, fractal, and disc images. Statistics from 1000 natural images and 100 of each type of un-natural image are collected, see Figure 4




Figure 4: Statistics for natural (●), and unnatural (gray) images projected onto a 3-D linear subspace. The gray symbols correspond to the noise (×), fractal (□), and disc (○) images, see Figure 3.

2.1. Natural vs. Un-Natural

Shown in the top portion of Figure 3 are four images taken from a database of natural images². These images span decades of digital and traditional photography and consist of a broad range of indoor and outdoor scenes. Statistics from 1000 images are collected as follows. Each 8-bit per channel RGB image is cropped to a central 640×480 pixel area. Each image is then converted from RGB to gray-scale³ ($\text{gray} = 0.299R + 0.587G + 0.114B$). A four-level, three-orientation QMF pyramid is constructed for each image, from which a 72-length feature vector of coefficient and error statistics is collected, Section 2. To reduce sensitivity to noise in the linear predictor, only coefficient magnitudes greater than 1.0 are considered.

The same set of statistics are collected from “unnatural” images, samples of which are also shown in Figure 3 - from top to bottom are: (1) noise images which are scrambled versions of the “natural” images shown in the same figure; (2) fractal images with a $1/\omega^p$ power spectrum and random phase ($p \in [1, 2]$); and (3) disc images consisting of overlapping anti-aliased discs of variable size radii. These unnatural images were chosen to be consistent with various sta-

²Images were downloaded from: philip.greenspun.com and reproduced here with permission from Philip Greenspun.

³Images are converted from RGB to grayscale to simplify the analysis. There is no doubt that there are strong statistical correlations between the color channels and we plan to explore these in future work.

tistical properties of natural images. The noise images retain the intensity distribution of the natural images; the fractal images have roughly the same Fourier energy; and the disc images have similar phase statistics.

Shown in Figure 4 is the projection of 1000 natural and 300 un-natural 72-D feature vectors onto a three-dimensional linear subspace. This lower-dimensional space is the result of projecting the original 72-D vectors onto the top three principle components, as computed from a principle component analysis (Appendix A) - these top three components capture 75% of the total variance. Even in this reduced space, there is a clear separation between the natural and un-natural images, suggesting that this statistical model captures certain aspects of natural images.

3. Digital Forensics

In our current digital age, it is startling to see the ease with which digital media can and is being manipulated to alter our sense of reality. Whether it be a Hollywood studio, a national news organization, or an average computer user, the images and sounds that are being created can no longer be unquestionably believed. The courts, in particular, are wholly unprepared to contend with the sophisticated digital technology that allows even the most novice of users to alter our sense of reality.

We have employed the statistical model outlined above for use in digital forensics. Our overall goal is the development of computational techniques for determining whether a digital audio signal, image, or video has been tampered with from the time of its recording. In contrast to digital watermarking or signatures, the basic approach taken here is to characterize the statistics of natural images, build models for various types of tampering, and then construct classification schemes for detecting deviations from the expected statistics. Below we describe three such applications: (1) detecting hidden messages (steganography); (2) differentiating between computer graphic and natural images; and (3) detecting re-broadcast attacks against biometric identification systems.

3.1. Steganography

Information hiding techniques (steganography) have recently received quite a bit of attention (see [14, 12, 19] for general reviews). Applications of this technology include unobtrusive military and intelligence communication and covert criminal communication.

With digital images as carriers, detecting the presence of hidden messages poses significant challenges to the legal community. Although the presence of




Figure 5: Shown is a cover image and a steg image containing an embedded message. Also shown is the the 64×64 message (at scale), and the absolute value of the difference between the cover and steg image (renormalized into the range $[0,255]$ for display purposes).

embedded messages is, by design, imperceptible to the human eye, it may nevertheless disturb the statistics of an image. Previous approaches to detecting such deviations [13, 26, 20] typically examine first-order statistical distributions of intensity or transform coefficients (e.g., discrete cosine transform, DCT). In contrast, we employ the higher-order statistical models described above for the purposes of detecting the presence of a hidden message within a digital image (see also [9] for related work on higher-order DCT image statistics).

Shown in Figure 5 is a 640×480 cover and 64×64 message image, and the result of embedding the message into the cover image. In this example, the mean of the absolute value of the difference between the cover and steg image is 6.2 intensity values with a standard deviation of 6.3. For display purposes the difference image is renormalized into the same intensity range as the cover image, $[0, 255]$.

Given an image we would like to determine if it contains a hidden message. To do so, we first collect statistics from images with and without hidden messages. For purposes of illustration, messages are embedded using Jsteg⁴ into the same 1000 natural images used in the previous section. A message consists of a 64×64 pixel region of a random image chosen from the same image database. The same transformation, decomposition, and collection of statistics as described in Section 2.1 is performed on the 1000 natural images and 1000 images with embedded messages. A Fisher linear discriminant (FLD) [8, 5] is then trained on 750 of the natural and 750 of the steg images, and then tested on the remaining 250 natural and 250 steg images (Appendix B). In the training stage, the FLD correctly classifies 99.5% of the natural and 98.3% of the steg images - a threshold was selected so as to afford a less than 0.5% false positive rate. In the testing stage, 98.9% of the natural and 97.6% of the steg images are correctly classified. These values are averaged over 100 random training/testing splits.

The statistical model captures certain aspects of natural images that are sufficiently specific so as to easily detect images that contain hidden messages that are themselves imperceptible to the human eye. In addition to these results, we have shown in previous work that both a Fisher linear discriminant and a support vector machine are effective in discriminating between natural and steg images across a range of different image formats and embedding programs (see [6, 16] for more details).

3.2. Computer Graphics or Photograph?

In 1996 the Child Pornography Prevention Act was passed which, in part, prohibited any image that *appears to be or conveys the impression* of someone under 18 engaged in sexually explicit conduct. This law made illegal computer generated pictures that only appear to show minors involved in sexual activity. In 2002, the United States Supreme Court struck down this law in their 6-3 ruling in *Ashcroft v. Free Speech Coalition* - the court said language in the 1996 child pornography law was unconstitutionally vague and far-reaching.

This ruling makes it considerably more difficult for law enforcement agencies to prosecute child pornography crimes, since it is always possible to claim that any image is computer generated. To this end, it would be helpful to be able to distinguish between natural and computer generated images. Given that




Figure 6: Computer graphic images⁵.

current computer graphics software can often create images that look to the human eye like a photograph, this classification poses considerable challenges. Shown in Figure 6 are four images⁵ from a database of 500 computer graphic images. These images were generated from a number of different software packages: 3D Studio Max, Maya, SoftImage 3D, Lightwave 3D, Imagine, and Alias PowerAnimation.

Given an image we would like to determine if it is computer graphic (CG) or natural. To do so, we first collect statistics from natural and computer graphic images. The same transformation, decomposition, and collection of statistics as described in Section 2.1 is performed on the same 1000 natural images from previous sections and 500 CG images. A Fisher linear discriminant (FLD) is trained on 750 of the natural and 350 of the CG images, and then tested on the remaining 250 natural and 150 CG images (Appendix B). In the training stage, the FLD correctly classifies 99.5% of the natural and 36.9% of the CG images - a threshold was selected so as to afford a less than 0.5% false positive rate. In the testing stage, 98.7% of the natural and 35.4% of the CG images are correctly classified. These values are averaged over 100 random training/testing splits.

The statistical model of natural images is sufficiently specific so that slightly more than one-third of computer graphic images can be reliably distinguished from their natural counterparts.

3.3. Live or Re-Broadcast?

With biometric identification systems (e.g., face, iris, or voice) becoming more popular and effective, these

⁴Jsteg V4, by Derek Upham, is available at www.nic.funet.fi/pub/crypt/steganography

⁵Graphic images are credited to Asier H. Laviña (top) and Frank Vitale (bottom) and reproduced here with their permission.



Figure 7: Shown is the original image (left) and the image after being printed and re-scanned (right). The mean of the absolute value of the difference between these images is 19 with a standard deviation of 17.5 (on a scale of [0, 255]).

systems have started to become vulnerable to attack. A simple and effective attack on, for example, a face recognition system would be to place in front of the camera a photograph of an individual. As in the previous two sections we detect such an attack by quantifying the perturbations in the image statistics caused by, in this case, the re-broadcasting of an image.

We collect statistics from natural images and the same images after having been printed on a laser printer and re-scanned with a flat-bed scanner (printing and scanning are done at 72 dpi), Figure 7. The same transformation, decomposition, and collection of statistics as described in Section 2.1 is performed on the 1000 natural images and 200 re-broadcast images. A Fisher linear discriminant (FLD) is then trained on 750 of the natural and 150 of the re-broadcast images, and then tested on the remaining 250 natural and 50 re-broadcast images (Appendix B). In the training stage, the FLD correctly classifies 99.5% of the natural and 100% of the re-broadcast images - a threshold was selected so as to afford a less than 0.5% false positive rate. In the testing stage, 99.5% of the natural and 99.8% of the re-broadcast images are correctly classified. These values are averaged over 100 random training/testing splits.

It is not surprising that the dithering that results from printing significantly disturbs the image statistics. There is little doubt, however, that detecting a re-broadcast image will become increasingly more difficult with higher-quality printers.

3.4. Summary

We have shown in the previous three sections that the proposed statistical model of Section 2 is effective

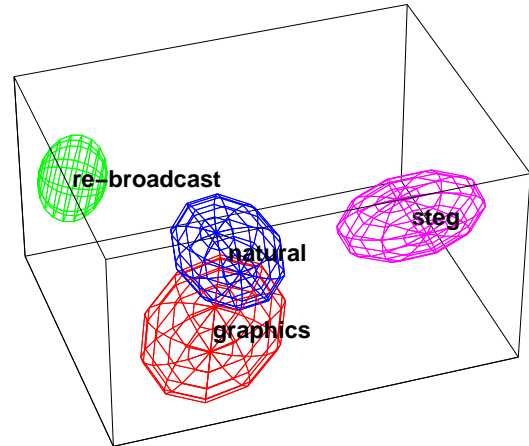


Figure 8: Results from a four-way classifier of 1000 natural, 1000 steg, 500 graphic, and 200 re-broadcast images. The bounding ellipsoids are rendered so as to capture 1.5 times the standard deviation of the data projected onto each of the principle axis (individual data points are not shown).

in distinguishing natural images from images with hidden messages, computer graphic images, and re-broadcast images. This classification was performed using a linear classifier (FLD) where each class of images was treated separately. Shown in Figure 8 are results from a four-class FLD, where all four classes of images were simultaneously classified. The bounding ellipsoids are rendered so as to capture 1.5 times the standard deviation of the data projected onto each of the principle axis (for clarity, the individual data points are not shown). Note how the four clusters are reasonably well separated, suggesting that the statistical disturbance between these classes of images are distinct. This type of separation should prove to be particularly valuable in a digital forensic setting.

We next wondered what aspects of the statistical model are essential for distinguishing natural from un-natural images. For example, would classification accuracy be as good with only the coefficient statistics, or are the higher-order error statistics the critical component? To this end, we re-trained each FLD of Sections 3.1-3.3 using either only the coefficient statistics or only the error statistics. With only the coefficient statistics: (1) 99.5% of the natural and 97.0% of the steg images are correctly classified; (2) 99.5% of the natural and 6.0% of the computer graphic images are correctly classified; and (3) 99.5% of the natural and 99.8% of the re-broadcast images are correctly classified. On the other hand, with only the higher-order

error statistics: (1) 99.5% of the natural and only 7.5% of the steg images are correctly classified; (2) 99.5% of the natural and 16.2% of the computer graphic images are correctly classified; and (3) 99.5% of the natural and 38.8% of the re-broadcast images are correctly classified. It seems then that in some cases the coefficient statistics are sufficient for classification (steg, re-broadcast), and in others (computer graphics), both coefficient and error statistics are needed. Further study is needed to determine if certain spatial scales or orientations are more influential than others.

4. Discussion

We have presented a statistical model for natural images consisting of first- and higher-order wavelet statistics. This model seems to capture certain statistical regularities of natural images. We have shown how this model can be useful in several digital forensic applications. From the point of view of digital forensics, the classification techniques outlined here will almost certainly benefit from a more flexible classifier, for example, a non-linear support vector machine. The statistical model is, of course, vulnerable to counter-attack. It may be possible, for example, to alter an image so that its first- and higher-order statistics are consistent with those of natural images. The development of such counter-measures will in turn lead to better models, and so on.

Acknowledgments

This work was supported by an Alfred P. Sloan Fellowship, a National Science Foundation CAREER Award (IIS-99-83806), a Department of Justice Grant (2000-DT-CS-K001), and a departmental National Science Foundation Infrastructure Grant (EIA-98-02068).

Appendix A - PCA

Denote column vectors $\vec{x}_i \in \mathcal{R}^n$, $i = 1, \dots, N$ as the original feature vectors. The overall mean is:

$$\vec{\mu} = \frac{1}{N} \sum_{i=1}^N \vec{x}_i \quad (9)$$

The zero-meaned data is packed into a $n \times N$ matrix:

$$M = (\vec{x}_1 - \vec{\mu} \quad \vec{x}_2 - \vec{\mu} \quad \dots \quad \vec{x}_N - \vec{\mu}) \quad (10)$$

If the dimensionality n of \vec{x}_i is smaller than the number of data points N , as in our case, then the $n \times n$ (scaled) covariance matrix is computed as:

$$C = MM^t \quad (11)$$

The principle components are the eigenvectors \vec{e}_j of the covariance matrix (i.e., $C\vec{e}_j = \lambda_j\vec{e}_j$), where the eigenvalue, λ_j is proportional to the variance of the original data along the j^{th} eigenvector. The dimensionality of each \vec{x}_i is reduced from n to p by projecting (via an inner product) each \vec{x}_i onto the top p eigenvalue-eigenvectors. The resulting p -dimensional vector is the reduced-dimension representation.

Appendix B - FLD

For simplicity a two-class FLD is described - the extension to multiple classes is straight-forward. Denote column vectors \vec{x}_i , $i = 1, \dots, N_x$ and \vec{y}_j , $j = 1, \dots, N_y$ as training exemplars from each of two classes. The within-class means are defined as:

$$\vec{\mu}_x = \frac{1}{N_x} \sum_{i=1}^{N_x} \vec{x}_i, \quad \text{and} \quad \vec{\mu}_y = \frac{1}{N_y} \sum_{j=1}^{N_y} \vec{y}_j. \quad (12)$$

The between-class mean is defined as:

$$\vec{\mu} = \frac{1}{N_x + N_y} \left(\sum_{i=1}^{N_x} \vec{x}_i + \sum_{j=1}^{N_y} \vec{y}_j \right). \quad (13)$$

The within-class scatter matrix is defined as:

$$S_w = M_x M_x^T + M_y M_y^T, \quad (14)$$

where, the i^{th} column of matrix M_x contains the zero-meaned i^{th} exemplar given by $\vec{x}_i - \vec{\mu}_x$. Similarly, the j^{th} column of matrix M_y contains $\vec{y}_j - \vec{\mu}_y$. The between-class scatter matrix is defined as:

$$S_b = N_x(\vec{\mu}_x - \vec{\mu})(\vec{\mu}_x - \vec{\mu})^T + N_y(\vec{\mu}_y - \vec{\mu})(\vec{\mu}_y - \vec{\mu})^T. \quad (15)$$

Let \vec{e} be the maximal generalized eigenvalue-eigenvector of S_b and S_w (i.e., $S_b\vec{e} = \lambda S_w\vec{e}$). The training exemplars \vec{x}_i and \vec{y}_j are projected onto the one-dimensional linear subspace defined by \vec{e} (i.e., $\vec{x}_i^T\vec{e}$ and $\vec{y}_j^T\vec{e}$). This projection simultaneously reduces the within-class scatter while increasing the between-class scatter. Once the FLD projection axis is determined from the training set, a novel exemplar, \vec{z} , from the testing set is classified by first projecting onto the same subspace, $\vec{z}^T\vec{e}$. In the simplest case, the class to which this exemplar belongs is determined via a simple threshold. In the case of a two-class FLD, we are guaranteed to be able to project onto a one-dimensional subspace (i.e., there will be at most one non-zero eigenvalue). In the case of a N -class FLD, the projection may be onto as high as a $N - 1$ -dimensional subspace.

References

- [1] R.J. Anderson and F.A.P. Petitcolas. On the limits of steganography. *IEEE Journal on Selected Areas in Communications*, 16(4):474–481, 1998.
- [2] R.W. Buccigrossi and E.P. Simoncelli. Image compression via joint statistical characterization in the wavelet domain. *IEEE Transactions on Image Processing*, 8(12):1688–1701, 1999.
- [3] G. Cross and A. Jain. Markov random field texture models. *IEEE Transactions on Pattern and Machine Intelligence*, 5:25–39, 1983.
- [4] J.G. Daugman. Entropy reduction and decorrelation in visual coding by oriented neural receptive fields. *IEEE Transaction on Biomedical Engineering*, 36(1):107–114, 1989.
- [5] R. Duda and P. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.
- [6] H. Farid. Detecting hidden messages using higher-order statistical models. In *International Conference on Image Processing*, Rochester, New York, 2002.
- [7] D.J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379–2394, 1987.
- [8] R. Fisher. The use of multiple measures in taxonomic problems. *Annals of Eugenics*, 7:179–188, 1936.
- [9] J. Fridrich and M. Goljan. Practical steganalysis: State of the art. In *SPIE Photonics West, Electronic Imaging*, San Jose, CA, 2002.
- [10] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern and Machine Intelligence*, 6:721–741, 1984.
- [11] M. Hassner and J. Sklansky. The use of Markov random fields as models of texture. *Computer Graphics and Image Processing*, 12:357–370, 1980.
- [12] N. Johnson and S. Jajodia. Exploring steganography: seeing the unseen. *IEEE Computer*, pages 26–34, 1998.
- [13] N. Johnson and S. Jajodia. Steganalysis of images created using current steganography software. *Lecture notes in Computer Science*, pages 273–289, 1998.
- [14] D. Kahn. The history of steganography. In *Proceedings of Information Hiding, First International Workshop*, Cambridge, UK, 1996.
- [15] D. Kersten. Predictability and redundancy of natural images. *Journal of the Optical Society of America A*, 4(12):2395–2400, 1987.
- [16] S. Lyu and H. Farid. Detecting hidden messages using higher-order statistics and support vector machines. In *5th International Workshop on Information Hiding*, Noordwijkerhout, The Netherlands, 2002.
- [17] S. G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern and Machine Intelligence*, 11:674–693, 1989.
- [18] A. P. Pentland. Fractal based description of natural scenes. *IEEE Transactions on Pattern and Machine Intelligence*, 6(6):661–674, 1984.
- [19] E.A.P. Petitcolas, R.J. Anderson, and M.G. Kuhn. Information hiding - a survey. *Proceedings of the IEEE*, 87(7):1062–1078, 1999.
- [20] N. Provos and P. Honeyman. Detecting steganographic content on the internet. Technical Report CITI 01-1a, University of Michigan, 2001.
- [21] D.L. Ruderman and W. Bialek. Statistics of natural image: Scaling in the woods. *Phys. Rev. Letters*, 73(6):814–817, 1994.
- [22] E.P. Simoncelli. Modeling the joint statistics of images in the wavelet domain. In *Proceedings SPIE, 44th Annual Meeting*, volume 3813, Denver, CO, USA, 1999.
- [23] E.P. Simoncelli and E.H. Adelson. *Subband image coding*, chapter Subband transforms, pages 143–192. Kluwer Academic Publishers, 1990.
- [24] P.P. Vaidyanathan. Quadrature mirror filter banks, M-band extensions and perfect reconstruction techniques. *IEEE ASSP Magazine*, pages 4–20, 1987.
- [25] M. Vetterli. A theory of multirate filter banks. *IEEE Transactions on ASSP*, 35(3):356–372, 1987.
- [26] A. Westfeld and A. Pfitzmann. Attacks on steganographic systems. In *Proceedings of Information Hiding, Third International Workshop*, Dresden, Germany, 1999.