

On Advances in Statistical Modeling of Natural Images

A. Srivastava ^{*}, A. B. Lee [†], E. P. Simoncelli, [‡], S-C. Zhu [§]

Abstract. Statistical analysis of images reveals two interesting properties: (i) invariance of image statistics to scaling of images, and (ii) non-Gaussian behavior of image statistics, i.e. high kurtosis, heavy tails, and sharp central cusps. In this paper we review some recent results in statistical modeling of natural images that attempt to explain these patterns. Two categories of results are considered: (i) studies of probability models of images or image decompositions (such as Fourier or wavelet decompositions), and (ii) discoveries of underlying image manifolds while restricting to natural images. Applications of these models in areas such as texture analysis, image classification, compression, and denoising are also considered.

Keywords: natural image statistics, non-Gaussian models, scale invariance, statistical image analysis, image manifold, generalized Laplacian, Bessel K form.

1. Introduction

In recent years, applications dealing with static or dynamic images (movies) have become increasingly important and popular. Tools for image analysis, compression, denoising, transmission, and understanding have become widely involved in many scientific and commercial endeavors. For instance, compression of images is important in transmission or storage, and automatic recognition of people from their camera images has become important for security purposes. Development of tools for imaging applications starts with mathematical representations of images. Many of these applications need to “learn” some image patterns or tendencies before the specific data can be analyzed. The area of natural image statistics has resulted from efforts to observe, isolate, and explain patterns exhibited by natural images. Lately, there has been a greater emphasis on explicit probability models for images. Perhaps one reason for that is the growing appreciation for the variability exhibited by the images, and the realization that exact mathematical/physical models may not be practical, and a statistical approach needs to be adopted. Any statistical approach will need probability models that capture “essential” image variability and yet are computationally tractable. Another motivation for pursuing image statistics has been to understand sensory coding as a strategy for information storage/processing in animal vision. An understanding of image statistics may provide clues to the architecture of animal visual system (Simoncelli and Olshausen, 2001).

There are several paths to understanding image statistics. Even though images are expressed as elements of a large vector space (e.g. the space of rectangular arrays of positive numbers), referred to as the *image space* here, the subset of interesting images is rather small and restricted. So one path is to isolate this subset, called an *image manifold* here, and learn a simplistic probability model, such as a uniform or a Gaussian-like model, on it. Given the image manifolds and probability distributions on them,

^{*} Department of Statistics, Florida State University, Tallahassee, FL 32306

[†] Division of Applied Mathematics, Brown University, Providence, RI 02912

[‡] Courant institute for Mathematical Sciences, New York University, New York, NY 10003

[§] Department of Computer Science, Ohio State University, Columbus, OH 43210.



statistical tools for imaging applications follow. The other path is to derive probability models that are defined on the larger image space but put any significant mass only on the image manifold. We have divided this paper along these two categories.

In search for statistical descriptions of images, mathematical and physical ideas are not abandoned but are intimately involved. For example, harmonic analysis of image intensities is used to decompose images into individual components that better lend to the model building than the original images. Since image spaces are rather high dimensional, and common density estimation techniques apply mostly to small sizes, one usually starts by decomposing images into their components, and then learning lower order statistics of these components individually. Also, imaging is a physical process and physical considerations often contribute to the search for statistical descriptors in some form. Physical models have been the motivation of many studies that have led to interesting statistical characterizations of images.

This paper is laid out as follows: In Section 2 we start with a historical perspective on spectral analysis of images and some discoveries that followed. An important achievement was the discovery of non-Gaussianity in image statistics. In the next two sections we present some recent results categorized into two sets: Section 3 studies probability models for images decomposed into their spectral components, while Section 4 looks at methods for discovering/approximating image manifolds for natural images. A few applications of these methods are outlined in Section 5, and some open issues are discussed in Section 6.

2. Image Decompositions & Their Statistical Properties

For the purpose of statistical modeling, an image is treated as a realization of a spatial stochastic process defined on some domain in \mathbb{R}^2 . The domain is assumed to be either a (continuous) rectangular region or a finite, uniform grid. A common assumption in image modeling is that the underlying image process is stationary, i.e. image probabilities are invariant to translations in the image plane.

2.1. CLASSICAL IMAGE ANALYSIS

In case the image process is modeled as a second order spatial process, the spectral analysis becomes a natural tool. Defining the covariance function, for two arbitrary pixel values, as $C(x)$ where x is the difference between two pixel locations, one can define the power spectrum as $P(w) = \int_{\mathbb{R}^2} C(x)e^{-jwx}dx$, and where w denotes the 2D spatial frequency. Since a Fourier basis is also an eigen basis for circulant matrices, which is the case for $C(x)$ under the stationarity assumption, Fourier analysis also coincides with the popular principal component analysis (discussed later in Section 4). Also, the Fourier representation guarantees uncorrelated coefficients. Images are then represented by their Fourier coefficients, image statistics are studied via coefficient statistics, and one can inherit a stochastic model on images by imposing a random structure on the Fourier coefficients. Specification of means and covariances of the coefficients completely specifies a second order image process. Early studies of spatial power spectra indicated that the power $P(w)$ decays as $\frac{A}{|w|^{2-\eta}}$ where $|w|$ is the magnitude of the spatial frequency. This

property, called the power law spectrum for images, was first observed by television engineers in the 50's (Kretzmer, 1952; Deriugin, 1956) and discovered for natural images in late 80s by (Field, 1987) and (Burton and Moorhead, 1987). As summarized in (Mumford and Gidas, 2001), the value of η changes with the image types but is usually a small number.

Although Fourier analysis is central to classical methods, other bases have found popularity for a variety of reasons. For instance, in order to capture the locality of objects in images, decomposition of images using a wavelet basis has become an attractive tool. In particular, it is common to use Gabor wavelets (Gabor, 1946) for decomposing the observed images simultaneously in space and frequency. In addition, Marr (Marr, 1982) suggested using the Laplacian of Gaussian filter to model early vision. If one considers images as realizations on a finite, uniform grid in \mathbb{R}^2 , the image space becomes finite-dimensional, and one can linearly project images into low-dimensional subspaces that are optimal under different criteria. Some of these linear projections are covered in Section 4. Once again image statistics are studied via the statistics of the projected coefficients.

2.2. SCALE INVARIANCE OF IMAGE STATISTICS

A discovery closely related to the power law is the invariance of image statistics when the images are scaled up or down (Field, 1987; Burton et al., 1986). In other words, the marginal distributions of statistics of natural images remain unchanged if the images are scaled. The power law is a manifestation of the fractal or scale invariant nature of images. By studying the histograms of the pixel contrasts ($\log(I(x)/I_0)$) at many scales, Ruderman et al (Ruderman and Bialek, 1994) showed its invariance to scaling. Independently, Zhu et al. (Zhu and Mumford, 1997) showed a broader invariance by studying the histograms of wavelet decompositions of images. Ruderman (Ruderman, 1994; Ruderman, 1997) also provided evidence of scale invariance in natural images and proposed a physical model for explaining them. Turiel et al. (Turiel and Parga, 2000) investigated the multi-fractal structure of natural images and related it to the scale invariance. Scaling of different types of scenes was studied by Huang (Huang, 2000). In (Turiel et al., 2000), Turiel et al. showed that hyperspectral images (color images) also demonstrate multiscaling properties and the statistics are invariant to scale. Scaling of order statistics of pixel values in small windows was studied by Geman et al. (Geman and Koloydenko, 1999). In addition to pixel statistics, scaling of topological statistics obtained from morphological operations on images was demonstrated by Alvarez et al. (Alvarez et al., 1999).

It must be emphasized that only the statistics of large ensembles of images are scale invariant; statistics of individual images vary across scales. Theoretical models that seek probabilistic description of image ensembles aim for scale invariance, while application driven models that work with individual images aim to capture individual image variability.

2.3. NON-GAUSSIANITY OF MARGINAL STATISTICS

Classical methods assume that images are second order processes but the observations do not support this assumption. Higher order statistics of natural images were found to exhibit interesting patterns and the researchers focused next on these higher order

moments. One implication is that image statistics do not follow Gaussian distribution and require higher order statistics. For example, a popular mechanism for decomposing images locally - in space and frequency - using wavelet transforms leads to coefficients that are quite non-Gaussian, i.e. the histograms of wavelet coefficients display heavy tails, sharp cusps at the median and large correlations across different scales. To our knowledge, Field (Field, 1987) was the earliest to highlight the highly kurtotic shapes of wavelet filter responses. Mallat (Mallat, 1989) pointed out that coefficients of multi-scale, orthonormal wavelet decompositions of images could be described by generalized Laplacian density (given later in Section 3.2). This non-Gaussian behavior of images has also been studied and modeled by Ruderman (Ruderman, 1994), Simoncelli & Adelson (Simoncelli and Adelson, 1996), Moulin & Liu (Moulin and Liu, 1999), and Wainwright (Wainwright and Simoncelli, 2000). Recent work of Thomson (Thomson, 2001) studies the statistics of natural images using phase-only second spectrum, a fourth order statistic, and demonstrates both the power law and the scale invariance for this statistic. In fact, projection onto any localized zero-mean linear kernel seems to produce kurtotic responses (Zetzsche, 1997). Huang (Huang, 2000) showed that images when filtered by 8×8 random mean-0 filters have distributions with high kurtosis, a sharp cusp at zero and long exponential tails. This suggests a role for linear decompositions that maximize the kurtosis or some another measure of non-Gaussianity. Such efforts (Bell and Sejnowski, 1997; Olshausen and Field, 1996; van Hateren, 1998) have resulted in bases that are spatially oriented with (spatial) frequency bandwidths being roughly one octave, similar to many multiscale decompositions. Similar results were obtained by minimizing the independence of coefficients, under linear decompositions, leading to independent component analysis (Cardoso, 1989; Comon, 1994; Hyvarinen et al., 2001). These observations justify widespread use of orthonormal wavelets in general image analysis applications. Use of Gabor wavelets is also motivated by the fact that the receptive fields of simple cells in the visual cortex of animals have been found to resemble Gabor functions (Miller, 1994). Note, however, that most wavelet decompositions of images are based on separable application of one-dimension filters, which leads to non-oriented (mixed diagonal) subbands. Alternate representations that provide better orientation decomposition, and thus higher kurtosis response, include (Daugman, 1985; Watson, 1987a; Simoncelli et al., 1992; Donoho and Flesia, 2001).

2.4. NON-GAUSSIANITY OF JOINT STATISTICS

In addition to the non-Gaussian behavior of marginal statistics, a number of authors have studied joint statistics of filter responses. In particular, the local structures present in most images lead to dependencies in the responses of local linear operators. A study of joint histograms of wavelet coefficients shows dependency across scales, orientations and positions. Zetzsche and colleagues noted that the joint density of projections onto even- and odd-symmetric Gabor filters have circular symmetry. If the density were Gaussian, this would correspond to independent (correlated) marginals, but the highly kurtotic marginals imply that the responses are strongly dependent (Wegmann and Zetzsche, 1990). Shapiro developed a heuristic method for taking advantage of joint dependencies between wavelet coefficients that revolutionized the field of image compression (Shapiro, 1993). Simoncelli and colleagues studied and modeled the dependency between responses

to pairs of bandpass filters and found that the amplitudes are strongly correlated, even when the signed responses are not (Simoncelli, 1997; Buccigrossi and Simoncelli, 1999). This is illustrated in Figure 1, which shows conditional histograms for several pairs of coefficients. Note that unlike second-order correlations, these dependencies can not be eliminated with a linear transform.




Figure 1. Bivariate histograms of wavelet coefficients associated with different basis functions. Top row shows contour plots, with lines at equal intervals of log probability. Left two cases are for different spatial offsets (same scale and orientation), third is for different orientation (same scale and nearly same position), and the rightmost corresponds to a pair at adjacent scales (same orientation and nearly same position). Bottom row shows some conditional distributions: brightness corresponds to larger frequency.

Observed statistics of an image patch can be described by joint histograms of coefficients under several filters. For jointly Gaussian coefficients equiprobable contours in joint histograms would all be ellipsoidal, but in the case of natural images the 2D and 3D contour surfaces display striking polyhedra-like shapes. Huang (Huang, 2000) showed that the peaks and cusps in the contours correspond to occurrences of simple geometries in images with partially constant intensity regions and sharp discontinuities. These results points to the ubiquity of “object-like” structures in natural images and underline the importance of object-based models. Grenander (Grenander and Srivastava, 2001) also has attributed non-Gaussianity to the presence of objects in images and has used that idea for model building. Lee et al. have shown that images synthesized from occlusion models (Section 3.1) show irregular polyhedra-like shapes in contour surfaces of histograms.

3. Emerging Statistical Models on Image Space

Earliest, and still widely used, probability models for images were based on Markov random field models (MRFs) (Winkler, 1995). An image field is as a collection of random variables, each denoting a pixel value on a uniformly spaced grid in the image plane. In MRFs, the conditional probability of a pixel value given the remaining image is reduced to a conditional probability given a neighborhood of that pixel. Efficiency results if

the neighborhood of a pixel is small, and furthermore, stationarity holds if the same conditional density is used at all pixels. Ising and Potts model are the simplest examples of this family. Besag (Besag, 1974; Besag, 1986) expressed the joint density of image pixels as a product of conditional densities, and ignored the normalizer to obtain a pseudo-likelihood formulation. Clifford-Hammersely theorem, see for example (Winkler, 1995), states that full conditionals completely specify the joint density function (under a positivity assumption) and enabled the analysis of images using a Gibbs sampler. Geman and Geman (Geman and Geman, 1984) utilized the equivalence of MRFs and Gibbs distributions to sample from these distributions. Kersten (Kersten, 1987) worked on computing the conditional entropies of the pixel values, given the neighboring pixels in an MRF framework. Zhu et al. (Zhu and Mumford, 1997) used a Gibbs model on images and estimated model parameters using a minimax entropy criterion. Let $H(I)$ be a concatenation of the histograms of coefficients, under several wavelet bases, then the maximum entropy probability takes the form: $P(I|\lambda) \propto e^{-\langle \lambda, H(I) \rangle}$. The vector λ is estimated by setting the mean histograms to equal the observed histograms, i.e. $\int H(I)P(I|\lambda)dI = H_{obs}$, and reduces to the maximum likelihood estimation of λ under $P(I_{obs}|\lambda)$.

3.1. MODELS MOTIVATED BY PHYSICS

A number of researchers have studied image statistics from a physical viewpoint, trying to capture the phenomena that generates images in the first place. A common theme to all these models is random placements of planar shapes (lines, templates, objects, etc) in images according to a Poisson process. Different models differ in their choice of shapes (e.g. primitive versus advanced), and their interaction (e.g. some favor occlusion while others favor superposition). Here we summarize a few of these models:

1. Superposition Models: In order to capture scale invariance and non-Gaussianity of images, Mumford et al. (Mumford and Gidas, 2001) utilized a family of infinitely divisible distributions. They showed that such distributions arise when images are modeled as superpositions of random placements of objects. To achieve self-similarity, it was proposed that the sizes of objects present in images be distributed according to density function Zr^{-3} over a subset of \mathbb{R}_+ (Z is the normalizing constant). Using this model, Chi (Chi, 1998) described a Poisson placement of objects with sizes sampled according to the $1/r^3$ -law. Additionally, he assumed a surface process that models the texture variation inside the 2D views of the objects. Grenander et al. (Bitouk et al., 2001; Grenander and Srivastava, 2001) have also assumed that images are made up of 2D appearances of the objects g_i s placed at homogeneous Poisson points z_i s. The objects are chosen randomly from an object space (arbitrary sizes, shapes or texture). A weighted superposition, with weights (or contrasts) given by independent standard normals a_i s, is used to form an image; it allows for expressing the linearly filtered images as a superposition of filtered objects. This formulation leads to an analytical (parametric) probability *for individual images* with the final form given in Section 3.2.

2. Occlusion Models: In more sophisticated models, objects are placed randomly as earlier but now the objects in the front occlude those in the back. An example is the dead leaves (also called a random collage) model which assumes that images are collages of approximately independent objects that occlude each other. One gen-

erates images from the model by placing an ordered set of elementary 2D shapes g_i s (in layers) at locations z_i s with sizes r_i s that are sampled from a probability density. For gray level images, the random sets or “leaves” $T_i = g_i(z - z_i)$ are furthermore associated with intensities a_i . The leaves are typically placed front-to-back¹ according to $I^{(i)}(z) = \begin{cases} a_i & \text{if } z \notin T_j, \forall j < i \\ I^{(i-1)}(z) & \text{otherwise} \end{cases}$, where $i = 1, 2, \dots$. Left two panels in Figure 2 show two samples of a dead leaves model with elliptical and square random shapes, respectively. The final model is defined in the limit $\lim_{i \rightarrow \infty} I^{(i)}(z)$, or equivalently, when




Figure 2. Left two pictures: Two different versions of the dead leaves model with elliptical (*left*) and square (*middle*) random shapes, respectively. Right panel: Derivative statistic and scaling of a dead leaves model with disks distributed according to a cubic law of sizes. The curves correspond to the marginal distribution of the differences between adjacent pixels after contrast normalization at four different scales ($N = 1, 2, 4, 8$).

the finite image domain $\Omega \subset \mathbb{R}^2$ is completely covered. The dead leaves model dates back to Matheron (Matheron, 1975) and Serra (Serra, 1982) in mathematical morphology. They showed that the probability of any compact set $K \subset \Omega$ belonging to the same leaf and not being occluded by other leaves equals the ratio $\frac{E[\nu(X_0 \ominus \check{K})]}{E[\nu(X_0 \oplus \check{K})]}$. Here, X_0 is the random shape used in the model, $E[\nu(\cdot)]$ is the expected Lebesgue measure, and $X_0 \ominus \check{K} = \{x : x + K \in X_0\}$ and $X_0 \oplus \check{K} = \{x : X_0 \cap (x + K) \neq \emptyset\}$ represent the *erosion* and the *dilation*, respectively, of X_0 by K . Early applications of this model include studies of inhomogeneous materials and the size distribution of grains in powder (Jeulin et al., 1995). More recently, researchers have applied different versions of the dead leaves model to natural scenes. In (Ruderman, 1997), Ruderman proposed that a random collage description with statistically independent ‘objects’ can explain scaling of correlations in images. Using an occluding object model and the difference function of hand-segmented images, he infers a power law for the probability of two points belonging to the same object. This result is consistent with (Alvarez et al., 1999) that uses level sets and the original dead leaves formalism of Matheron and Serra. Here, Alvarez et al analyze the morphological structure of natural images — in particular, the effect of occlusions on the area, perimeter and intercept lengths of homogeneous, connected image regions.

Lee et al. (Lee and Mumford, 2001) have used a dead leaves model with statistically independent objects and the measure of a marked Poisson process to show that $1/r^3$

¹ see (Kendall and Thonnes, 1998) for a discussion on “forward” and “backward” dead leaves algorithms

power law in object sizes also leads to approximate full scale invariance when occlusions are taken into account. Provided that synthesized images are block averaged, their statistics match those observed in large data sets of calibrated natural images. Scale invariance of histograms of filter responses, full co-occurrence statistics of two pixels, and joint statistics of Haar wavelet coefficients are studied in this paper. Figure 2 (right panel) shows a histogram of the derivative statistic, under scaling, in a dead leaves model with smoothing.

In addition, there have been some analogies drawn from other sciences to explain self-similarity of natural images. For example, Ruderman et al. (Ruderman, 1994; Ruderman and Bialek, 1994) pointed to the statistics of turbulent flows. In (Turiel et al., 1998), Turiel et al. formally connected the statistical modeling of turbulent flows to that of images. In turbulent flows, if the second moment of the energy dissipation is proportional to r^τ for some τ , where r is the box size, then self-similarity holds. Using a measure based on edges to compute energy dissipation in images, the authors showed self-similarity of edge statistics.

3.2. ANALYTICAL DENSITIES: UNIVARIATE CASE

An important requisite of statistical models is that they be computationally efficient for real-time, or near real-time, applications, and one obvious way to accomplish that is via parametric densities. We want to represent statistical nature of images by means of parametric densities, using only limited parameters. In this section, we discuss a general family of parametric densities that seems to capture the variability in low-dimensional representations of images. Rather than explaining scale-invariance of a large collection of images, the goal here is to *capture variability of individual images*. For an image I and a linear filter F , this section deals with characterizing the marginal density of the pixel values in $I * F$.

1. **Generalized Laplacian model:** Marginal densities of image coefficients are well modeled by a generalized Laplacian density (also called generalized Gaussian or stretched exponential) $f_1(x; c, p) = \frac{e^{-|x/c|^p}}{Z_1(p, c)}$, (Mallat, 1989; Simoncelli and Adelson, 1996; Moulin and Liu, 1999), where the normalizer is $Z_1(p, c) = 2\frac{c}{p}\Gamma(\frac{1}{p})$. The parameters, $\{p, c\}$, may be fit to the subbands of specific images using maximum likelihood or the method of moments. Another way of estimating them is via the (linear) regression of $\log(\log(h(x) + h(-x)) - 2\log(h(0)))$ versus $\log(|x|)$, where $h(x)$ is the histogram value at x and x is the variable for bin centers. Values for the exponent p are typically within the range $[0.5, 0.8]$, and the width parameter c varies monotonically with the size of the basis functions, producing higher variance for coarser-scale components (Simoncelli and Adelson, 1996).

2. **Gaussian Scale Mixtures:** Observed statistics of filtered images have pointed towards a rich family of univariate densities that are highly kurtotic and heavy tailed. In statistics literature, there is a wide usage of the variables defined as *normal variance-mean mixture*; X is called a normal variance-mean mixture if the conditional density function of X given u is normal with mean $\mu + u\beta$ and variance $u\Delta$, and u is called the mixing variable. Generalized hyperbolic distributions were introduced by Barndorff-Nielsen (Barndorff-Nielsen, 1977) as specific normal variance-mean mixtures that result when the mixing variable u is of certain class. For $\mu = \beta = 0$ and $\Delta = 1$, the resulting family is also called Gaussian scale mixture (Andrews and Mallows, 1974) and has seen

applications in financial mathematics (Bollerslev et al., 1994) and speech processing (Brehm and Stammer, 1987). Furthermore, if u is a scaled Gamma density then Bessel K density results.

Under Grenander’s formulation of a superposition model (item 1, Section 3.1), if the random variable $u(z) \equiv \sum_{i=1}^n ((g_i * F)(z - z_i))^2$ is assumed to have a scaled-Gamma density, then the univariate density of the filtered pixel has been shown in (Grenander and Srivastava, 2001) to be: for $p > 0$, $c > 0$, $f_2(x; p, c) = \frac{1}{Z_2(p, c)} |x|^{p-0.5} K_{(p-0.5)}(\sqrt{\frac{2}{c}} |x|)$, where K is the modified Bessel function of third kind, and Z_2 is the normalizing constant. (Its characteristic function is of the form $1/(1 + 0.5c\omega^2)^p$ which for $p = 1$ becomes the characteristic function of a Laplacian density.) This parametric family of densities has been called *Bessel K forms* with (p, c) referred to as *Bessel parameters*. Earlier, Wainwright et al. (Wainwright et al., 2001) also investigated the application of the Gaussian scale mixture family to image modeling and referred to f_2 as the *symmetrized Gamma* density (without reaching its analytic, parametric form). As described in (Grenander and Srivastava, 2001), p and c are easily estimated from the observed data, with $\hat{p} = \frac{3}{SK(I^{(j)}) - 3}$ and $\hat{c} = \frac{SV(I^{(j)})}{\hat{p}}$, where SK is the sample kurtosis and SV is the sample variance of the filtered image. A distinct advantage of this model is that parameters can be related mathematically to the physical characteristics of the objects present in the image. If a filter F is applied to an image I to extract some specific feature—vertical edges, say—then, the resulting p has been shown in (Srivastava et al., 2002) to depend on two factors: (i) distinctness and (ii) frequency of occurrence of that feature in I . Objects with sharper, distinct edges have low p values, while scenes with many objects have large p values.

Shown in Figure 3 are some examples of estimating this density function: natural images from van Hateren database are filtered using arbitrary Gabor filters (not shown), and the resulting pixel values are used to form the histogram $h(x)$. Bottom panels of this figure show the plots of $\log(h(x))$, $\log(f_1(x))$, and $\log(f_2(x))$ with parameters estimated from the corresponding images.

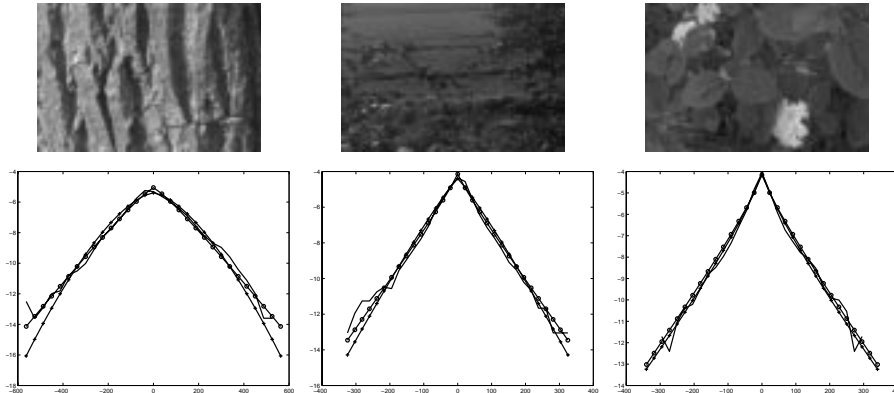


Figure 3. Estimated Bessel K and generalized Laplacian densities compared to the observed histograms for the images in top. Plain lines are histograms, lines with large beads are Bessel K, and lines with small beads are generalized Laplacian. These densities are plotted on a log scale.

There are several ways to judge the performance of any proposed probability model. The simplest idea is to compare the observed frequencies to the frequencies predicted by the model, using any metric on the space of probability distributions. Two commonly used metrics are the Kolmogorov-Smirnov distance: $\max_{x \in \mathbb{R}} |\int_{-\infty}^x (f_1(y) - f_2(y)) dy|$ and the Kullback-Leibler divergence: $\int_{\mathbb{R}} f_1(x) \log(\frac{f_1(x)}{f_2(x)}) dx$. Both the generalized Laplacian and the Bessel K form perform well in capturing univariate image statistics under these metrics. In general, for small values of p (sharper features and less number of objects in image) the generalized Laplacian performs better while for larger p the Bessel K density better matches the observed densities.

3.3. BIVARIATE PROBABILITY MODELS

So far we have discussed only univariate models but the complexity of observed images points to statistical interactions of high orders. As a first extension, we look at the models for capturing pairwise interactions between image representations, i.e. the bivariate probability densities. For instance, these models may involve a pair of wavelet coefficients at different scales, positions, or orientations.

1. Variance Dependency and Gaussian Scale Mixtures: As described in Section 2.4, pairs of wavelet coefficients corresponding to basis functions at nearby positions, orientations, or scales are not independent, even when the responses are (second-order) decorrelated. Specifically, the conditional variance of any given coefficient depends quite strongly on the values of surrounding coefficients (Simoncelli, 1997). As one considers coefficients that are more distant (either in spatial position, orientation, or scale), the dependency becomes weaker. This dependency appears to be ubiquitous, and may be observed across a wide variety of image types.

What sort of probability model can serve to explain the observations of variance dependency? One candidate is the Gaussian scale mixture model described earlier for the univariate case. In this scheme, wavelet coefficients are modeled as the product of a Gaussian random variable and a hidden “multiplier” random variable (same as mixing variable u earlier) that controls the variance. To explain pairwise statistics shown in Figure 1, one can prescribe a relationship between the hidden multiplier variables of neighboring coefficients. That is, the hidden variables now depend on one another, generating the variance scaling seen in the data. Modeling of this dependency remains an ongoing topic of research. One possibility is to link these variables together in a Markov tree, in which each multiplier is independent of all others when conditioned on its parent and children in the tree (Wainwright et al., 2001).

2. Bivariate Extension of Bessel K Forms: The univariate Bessel K form has been extended by Grenander to specify bivariate densities of the filtered images. The basic idea is to model all linear combinations of the filtered pixels by Bessel K forms and then invoke the Cramer-Wold device, see (Billingsley, 1995) for a definition, which states that specification of densities of all linear combinations of a number of random variables specifies uniquely their joint density. The use of the Cramer-Wold device here assumes that there exists a 2D distribution whose Radon-like marginals (for half spaces, not on lines) behave according to Bessel K densities. An imposition of Bessel K forms on marginals seems to agree well with the data qualitatively but its performance remains to be quantified over large datasets.

Let $I_1 = I * F^{(1)}$ and $I_2 = I * F^{(2)}$ for two wavelet filters $F^{(1)}$ and $F^{(2)}$, and for $a_1, a_2 \in \mathbb{R}$, let $J(a_1, a_2, z) \equiv a_1 I_1(z) + a_2 I_2(z)$. The Cramer-Wold idea is to compute the characteristic functions of $J(a_1, a_2, z)$, for all pairs (a_1, a_2) , using Bessel parameter estimation, and then take an inverse Fourier transform to obtain an estimate of the joint density $f(I_1(z), I_2(z))$. This estimate is parametrized by eight joint moments: $\mu_{0,2}, \mu_{1,1}, \mu_{2,0}, \mu_{4,0}, \mu_{3,1}, \mu_{2,2}, \mu_{1,3}, \mu_{0,4}$, where $\mu_{i,j} = \int \int I_1^i I_2^j f(I_1, I_2) dI_1 dI_2$. Shown in Figure 4 is an example of this bivariate density estimation. For the image shown in left, consider the bivariate density of its two filtered versions under two Gabor filters at the same scale but orientations 20 degrees apart. Top two panels show the mesh and the contour plots of the estimated density while the bottom panels show the observed bivariate histogram. The densities are all plotted on a log scale.




Figure 4. Bivariate extension of Bessel K form: For the image shown in left panel, the estimated (top) and observed (bottom) bivariate densities plotted as meshes and contours, on a log scale.

4. Discovering Image Manifolds

It has been well highlighted that in the space of rectangular arrays of positive numbers, only a small subset has images of natural scenes. One seeks to isolate and characterize this subset for use in image analysis applications. The main idea is to identify this set as a low-dimensional, differentiable manifold and use its geometry to characterize images. Having defined this manifold, a simplistic probability model can help capture the image variability. We now present a summary of some commonly used methods for estimating image manifolds:

1. **Approximating Image Manifolds by Linear/Locally-Linear Subspaces:** Perhaps the easiest technique to approximate the image manifold is to fit a linear subspace through the observations. The fitting criterion will specify the precise subspace that is selected. For example, if one intends to minimize the cumulative residual error (Euclidean distance between the observation points and the subspace), then the optimal subspace is the dominant (or principal) subspace of the data matrix, and is easily computed using eigen decomposition or singular value decomposition. A common application of PCA is in recognition of people from their facial images (Kirby and Sirovich, 1990) or study of natural images (Hancock et al., 1992). Instead, if the goal is to minimize the statistical correlation between the projected components, or to make them as independent as possible, then the independent component basis results (Comon, 1994; Bell and Sejnowski,

1997; van Hateren, 1998). Other criteria lead to similar formulations of the subspace basis such as sparsity (Olshausen and Field, 1997), Fisher discrimination (Belhumeur et al., 1997), and non-negative factorization (Lee and Seung, 1999). The use of sparseness is often motivated by the scale invariance of natural images.

Approximating the image manifold by a flat subspace is clearly limiting in general situations. In (Zetzsche and Rohrbein, 2001), the authors argue that linear processing of images leaves substantial dependencies between the components, and a nonlinear technique is required. One extension is to seek a “local linear embedding” approximation of the image manifold by fitting neighboring images by low dimensional subspaces (Roweis and Saul, 2000; Tenenbaum et al., 2000). Definition of a neighborhood is through Euclidean metric but that deserves further study. Another idea is to combine local basis elements, such as wavelet bases, into higher level structures that provide a better representation of the image manifold. For instance, Zhu et al (Zhu et al., 2002) have combined placements of transformed basis elements to form structures called *textons* in order to better characterize images and their manifolds. Shown in the left panel of Figure 5 is an example of building a “star” texton to match the given image.

2. Deformable Template Characterization of Image Manifold: Grenander’s pattern theoretic framework for image understanding is motivated by physical models and leads to probabilistic, algebraic representations of scenes. Objects appearing in the scenes are represented by 3D models of shape, texture, and reflectivity, and their occurrences in scenes are captured by group transformations on the typical occurrences or templates. A strong feature of this approach is the logical separation between the actual, physical scenes and the resulting images. Variability of scenes is better modeled in a 3D coordinate system as it is guided by the physical principles, rather than the image space where the Euclidean representations do not work well. Since the two, 3D scenes and 2D images, are related by a projection (orthographic or perspective) map, manifolds formed by images can be generated by projecting the manifolds formed using 3D representations. As described in (Grenander, 1993; Miller and Younes, 2002), occurrences of physical objects in 3D system are modeled as the orbits of groups acting on the configurations of objects in the scenes. These orbits are projected into the image space to form manifolds on which the images lie. Let C^α be a CAD model of a 3D object labeled by α and let S be the group of transformations (3D rotation, translation, deformation, etc.) that change the object occurrences in a scene. Let, for $s \in S$, sC^α denote the action of s on the template. Then, $\{sC^\alpha : s \in S\}$ is a 3D orbit associated with object α . Further, if T is the imaging projection from a 3D scene to the image plane, then $\{T(sC^\alpha) : s \in S\}$ is the image manifold generated by this object, and dimension of this image manifold equals $\dim(S)$. Geometry of this manifold, such as the tangent spaces or the exponential maps, can also be obtained by projecting their counterparts in the bigger space.

3. Empirical Evidence of Manifolds for Image Primitives: In contrast to a physical specification, as advocated by the deformable template theory, one can also study image manifolds directly using the observed images. This image based search is also inspired by Marr’s idea ((Marr, 1982)) of representing early vision by converting an array of raw intensities into a symbolic representation — a so called “primal sketch” — with primitives such as edges, bars, blobs and terminations as basic elements. An important issue in computer and human is: How are Marr’s primitives represented geometrically and statistically in the state space of image data? As stated in (Pedersen and Lee, 2002; Lee et al., 2002),

the geometrical model that arises from Marr’s hypothesis is a set of continuous manifolds of the general form $M(s) = [F^{(1)}(\cdot) * sC^\alpha, F^{(2)}(\cdot) * sC^\alpha, \dots]$, $F^{(j)}(j = 1, \dots, n)$ form a bank of filters, and sC^α is an image of a primitive α parameterized by $s \in S$. Dimension of resulting image manifold $M(s)$ is determined by an intrinsic dimensionality of the primitives; it is usually a small number (2 for contrast-normalized edges, 3 for bars and circular curves etc.). Furthermore, the manifolds generated by planar primitives form an hierarchical structure: The 2D manifold of straight edges, for example, is a subset of both the 3D manifold of bars and the 3D manifold of circular edges.

In (Lee et al., 2002), Lee et al. found that the state space of 3×3 natural image patches is extremely sparse with the patches densely concentrated around a non-linear manifold of blurred edges. For the top 20% highest contrast patches in natural images, half of the data lie within a neighborhood of the edge manifold that occupies only 9% of the total volume. Estimated probability density, as a function of the distance to the manifold, takes the form $f(\text{dist}) \sim 1/\text{dist}^\beta$ with $\beta = 2.5$; it has an *infinite density* at the image manifold where $\text{dist} = 0$. The paper (Pedersen and Lee, 2002) extends these results by considering filtered patches with the filters being (up to third order) derivatives of Gaussian kernels. Figure 5 right panel shows the estimated probability density as a function of the distance to the edge manifold for multiscale representations of natural images. The density is approximately scale invariant and seems to converge toward the functional form $p(\text{dist}) \sim 1/\text{dist}^{0.7}$.

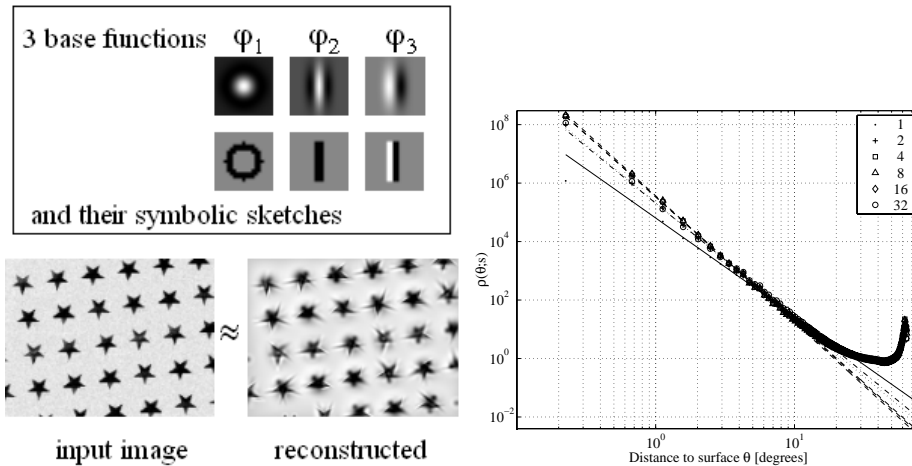


Figure 5. Left panel: To model an image of stars, some wavelet bases are combined to form a star texton. Right Panel: Natural images have an infinite probability density at a manifold in state space that corresponds to blurred edges. The figure shows the empirical probability density as a function of the distance to the manifold of Gaussian scale-space edges in jet space, with curves corresponding to jet representations of natural images at scales $s = 1, 2, 4, 8, 16, 32$.

5. Applications of Statistical Models

Main reason for developing formal statistical models is to apply them in image processing/analysis applications. There is a large array of applications that continually benefit from advances in modeling of image statistics. Here we have selected an important subset: 1. **Texture Synthesis:** Use of newer statistical models has revolutionized the area of texture analysis. In 1980 Faugeras et al. (Faugeras and Pratt, 1980) suggested using the marginals of filtered images for texture representations. Bergen et al (Bergen and Landy, 1991), Chubb et al. (Chubb et al., 1994), and Heeger et al. (Heeger and Bergen, 1995) also advocated the use of histograms. Zhu et al. (Zhu et al., 1997) showed that marginal distributions of filtered images, obtained using a collection of filters, sufficiently characterize homogeneous textures. Choice of histograms implies that only the frequencies of occurrences are retained and the locations are discarded. Exploiting periodicity of these textures, one extracts features using wavelet decompositions at several scales and orientations, and uses them to represent images. Many schemes have been proposed, two of them are: (i) Julesz ensemble that impose equivalence on all images that lead to the same histograms, and (ii) Gibbs model stated earlier in Section 3. Shown in Figure 6 are some examples of texture synthesized under these models: the top row shows the real images and the bottom row shows corresponding synthesized textures. It must be noted that this synthesis framework holds well even when the raw histograms are replaced by parametric analytical forms, as shown in (Srivastava et al., 2002).




Figure 6. Texture synthesis using Julesz ensemble model (first two columns) and the Gibbs model (last two columns). Top panels: observed real images and bottom panels: synthesized images.

Beyond the use of marginals, Portilla et al. (Portilla and Simoncelli, 2000) developed a model based on a characterization of joint statistics of filter responses. Specifically, they measured the correlations of raw coefficients as well as the correlations of their magnitudes, and developed an efficient algorithm for synthesizing random images, subject to these constraints, by iteratively updating the image and forcing the constraints.

2. **Image Compression:** Compression seems to be a natural application for emerging statistical image models. In a typical implementation, the image is decomposed using a linear basis, the coefficients of this representation are quantized to discrete values, and these values are then efficiently encoded by taking advantage of their relative probability

of occurrence. In this context, the statistical models considered here provide estimates of these probabilities. Most widely used compression scheme for images is the JPEG standard, which is based on a block-by-block frequency decomposition. In early 80s, it was recognized that multi-scale wavelet-style representations offer more flexibility and better compression performance (Burt and Adelson, 1983; Vetterli, 1984). Early coders were based on marginal models of coefficients, but this changed abruptly with the development of the first contextual coder by Shapiro (Shapiro, 1993). This coder, and many that followed (Rinaldo and Calvagno, 1995; Chrysafis and Ortega, 1997), took heuristic advantage of the joint statistical properties of wavelets. Some subsequent coders have been based more explicitly on such models (LoPresto et al., 1997; Buccigrossi and Simoncelli, 1999).

3. Image Classification: An interesting application is to classify images into some pre-defined categories. If certain lower order statistics of images are found sufficient for this purpose, some level of efficiency can be achieved. We start the discussion with the classical methods for classification.

- **Multiple Hypothesis Testing:** A classical technique for classification, given the probability models associated with the classes, is hypothesis selection (Grenander et al., 2000). Given an observation, the goal is to select the hypothesis that is most likely to have generated that observation. Let H_1, H_2, \dots, H_n correspond to the n image classes, and let $P(I|H_i)$ be the likelihood of image I belonging to class H_i , then hypothesis selection can be accomplished as a sequence of binary

$$\text{hypothesis tests: } \frac{P(I|H_i)}{P(I|H_j)} \underset{H_j}{\underset{H_i}{\geq}} \nu_{ij}, \text{ where } \nu_{ij} \text{ is a threshold value generally taken to be}$$

one. Neyman-Pearson lemma provides optimality to hypothesis testing in a certain sense. In case of a Bayesian selection, the threshold ν_{ij} is given by the ratio of priors $\nu_{i,j} = P(H_j)/P(H_i)$. For all these tests, one needs the likelihood function $P(I|H_i)$ which will depend upon the choice of statistical models. For example, if a parametric form is chosen to model the image probability, then the classes can be directly related to the parameter values. Typical values (“average values”) of the parameter represent a class and the likelihood $P(I|H_i)$ is written in the parametric form with the typical parameter values for each class. Deformable templates (Section 4), has been successfully applied to object recognition. The inference is obtained by hypothesis testing for α in presence of nuisance variables $s \in S$. Here, the likelihood is computed via the nuisance integral $P(I|H_i) = \int_S P(I, s|H_i)\gamma(ds)$. Nuisance variable estimation on a group S and estimation error bounds are derived in (Grenander et al., 1998), while the hypothesis testing for α and recognition error bounds are derived in (Grenander et al., 2000).

- **Metrics for Image Comparison:** A broader goal in image analysis is to quantify differences between two given images. Given such a metric one can perform image clustering, image retrieval, image classification, and even recognition of objects in given images. If the probability models are parametric, one can derive a distance measure on the image space that takes a parametric form. If $f(x|p_1, c_1)$ and $f(x|p_2, c_2)$ are two univariate density functions parameterized by (p_1, c_1) and

(p_2, c_2) respectively, then the distance measure takes the form: $d(p_1, c_2, p_2, c_2) = \tilde{d}(f(x|p_1, c_1), f(x|p_2, c_2))$, where \tilde{d} is a metric on the space of univariate densities. Several forms have been proposed for \tilde{d} including geodesic length (Riemannian metric), Earth Mover's distance (Cohen and Guibas, 1999), Kullback-Leibler divergence, Renyi's α -divergence (Hero et al., 2001), Jensen-Renyi divergence (He et al., 2002; Hamza et al., 2001), χ^2 -distance, and the L^p norm for $p = 1, 2, \dots$. The choice of metric depends upon the application and the desired computational simplicity. Srivastava et al. (Srivastava et al., 2002) have derived a parametric form of L^2 metric between two Bessel K forms. Two densities can be compared directly using their estimated parameters without requiring to compute full histograms.

4. Image Denoising: A common approach here is to decompose image into bands of spatial frequency and to threshold the coefficients after some nonlinear transformation, proposed first by Bayer et al. (Bayer and Powell, 1986) and later by Donoho (Donoho, 1995). The nonlinear transformation is used essentially to shrink all wavelet coefficients towards zero. This shrinking is based on thresholding which can be implemented as a hard threshold or a soft threshold. For applying statistical approaches, such as Bayesian or MAP techniques, an explicit prior model on the image pixels is required. Simoncelli et al. (Simoncelli and Adelson, 1996) studied image denoising in a Bayesian framework while Moulin et al. (Moulin and Liu, 1999) reported the use of generalized Laplacian models in statistical denoising approaches. The fact that statistics of decomposed images are easier to characterize, than that of the original images, has led to many pyramid-based approaches to image denoising. Images are decomposed into multi-scale representations and statistics of coefficients are used for denoising (Leporini et al., 1999; Portilla et al., 2001; Simoncelli, 1999) in individual frequency bands. Also, the Gaussian scale mixture model for joint statistics can be used in a Bayesian framework (Portilla et al., 2001), producing results that are significantly better than those achieved with a marginal model.

5. Other Applications: There are a number of other image analysis/synthesis applications that have benefited from statistical ideas. (Yendrikhovskij, 2001) uses a clustering of color statistics to model the perceived color environment and to compute categories of coloring. A bidirectional radiosity distribution function (BRDF) of an object completely specifies its appearance under arbitrary illumination conditions. Dror et al. (Dror et al., 2001) have shown that statistics of these illumination maps are similar to the statistics of natural images, and hence the proposed image models can apply there as well (Weiss, 2001).

6. Discussion

In this paper we have discussed some recent advances in statistical modeling of natural images. Not only these models provide a better match than the traditional models but also have led to significant improvements in a number of imaging applications.

Although substantial progress has been made over the past twenty years in understanding complex statistical properties of natural images, we are still quite far from a full probability model. For example, samples drawn from existing models are unable to capture the variety and complexity of real world images, except in the restricted case of

homogeneous textures. Beyond univariate and bivariate densities of the image statistics, the computational complexity increases exponentially. An important question is: In the context of a specific imaging application, say face recognition from video images, what order densities are required to ensure a reasonable success? For homogeneous textures, univariate models have been successful but how much more is needed for more general applications. The notion of sufficient statistics needs to be made precise for different application contexts. Even among the proposed models, several issues remain open. For instance, one issue in dead leaves model is how to incorporate texture and dependencies of objects. At this point, there are few analytical results for realistic dead leaves models.

Many of the models described in this paper model statistics of an ensemble of images; their applications for analysis of individual images needs to be clarified.

Beyond applications, such as synthesis and compression of images, an important reason for developing statistical models is image understanding, an area with many outstanding problems. One such problem is: Given an image outside a rectangle, find the most probable interpolation of it inside of the rectangle. Lack of Markovity disallows the classical harmonic analysis, and points to more powerful pattern-theoretic structures underlying the image ensembles.

Acknowledgments

We thank the creators of van Hateren database for making their database available to public. We express gratitude to our collaborators Ulf Grenander, Kim S. Pedersen and David Mumford for their help in this research. Many thanks to Hamid Krim for his invitation and support in writing this paper. AS has been supported in part by the grants NSF DMS-0101429, ARO DAAD19-99-1-0267, and NMA 201-01-2010. EPS was supported in part by NSF CAREER grant MIP-9796040, the Alfred P. Sloan Foundation, and the Howard Hughes Medical Institute.

References

- Alvarez, L., Y. Gousseau, and J.-M. Morel: 1999, ‘The size of objects in natural and artificial images’. In: P. W. H. et al. (ed.): *Advances in Imaging and Electron Physics*. Academic Press.
- Andrews, D. and C. Mallows: 1974, ‘Scale mixtures of normal distributions’. *Journal of the Royal Statistical Society* **36**, 99–102.
- Barndorff-Nielsen, O.: 1977, ‘Exponentially decreasing distribution for the logarithm of a particle size’. *Journal of Royal Statistical Society A* **353**, 401–419.
- Bayer, B. E. and P. G. Powell: 1986, ‘A method for the digital enhancement of unsharp, grainy photographic images’. *Advances in Computer Vision and Image Processing* **2**, 31–88.
- Belhumeur, P. N., J. P. Hefanaha, and D. J. Kriegman: 1997, ‘Eigenfaces vs. fisherfaces: Recognition using class specific linear projection’. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(7), 711–720.
- Bell, A. J. and T. J. Sejnowski: 1997, “‘The “independent components” of natural scenes are edge filters’”. *Vision Research* **37**(23), 3327–3338.
- Bergen, J. R. and M. S. Landy: 1991, *Computational modeling of visual texture segregation*, pp. 253–271. MIT Press, Cambridge, MA.
- Besag, J.: 1974, ‘Spatial interaction and statistical analysis of lattice systems (with discussion)’. *J. Royal Statist. Soc. series B*, vol. **36**, 192–326.

- Besag, J.: 1986, 'On the statistical analysis of dirty pictures'. *J. Royal Statistical Society B* **48**, No.3(48, No.3), 259–302.
- Billingsley, P.: 1995, *Probability and Measure*. Wiley series in probability and mathematical statistics.
- Bitouk, D., U. Grenander, M. I. Miller, and P. Tyagi: 2001, 'Fisher Information Measures for ATR in Clutter'. In: *Proceedings of SPIE, Automatic Target Recognition XI*, Vol. 4379. pp. 560–572.
- Bollerslev, T., K. Engle, and D. Nelson: 1994, *ARCH models*, Vol. IV, pp. 2959–3038. North-Holland, Amsterdam.
- Brehm, H. and W. Stammers: 1987, 'Description and generation of spherically invariant speech-model signals'. *Signal Processing* **12**, 119–141.
- Buccigrossi, R. W. and E. P. Simoncelli: 1999, 'Image Compression via Joint Statistical Characterization in the Wavelet Domain'. *IEEE Transactions on Image Processing* **8**(12), 1688–1701.
- Burt, P. J. and E. H. Adelson: 1983, 'The Laplacian pyramid as a compact image code'. *IEEE Transactions on Communications* **31**(4), 532–540.
- Burton, G. J., N. D. Haig, and I. R. Moorhead: 1986, 'A self-similar stack model for human and machine vision'. *Biological Cybernetics* **53**(6), 397–403.
- Burton, G. J. and I. R. Moorhead: 1987, 'Color and Spatial structures in natural scenes'. *Applied Optics* **26**(1), 157–170.
- Cardoso, J.-F.: 1989, 'Source separation using higher order moments'. In: *Proceedings of ICASSP*. pp. 2109–2112.
- Chi, Z.: 1998, 'Probability models for complex systems'. *Ph.D. thesis, Division of Applied Mathematics, Brown University*.
- Chrysafis, C. and A. Ortega: 1997, 'Efficient Context-Based Entropy Coding for Lossy Wavelet Image Coding'. In: *Data Compression Conf. Snowbird, Utah*.
- Chubb, C., J. Econopouly, and M. S. Landy: 1994, 'Histogram contrast analysis and the visual segregation of IID textures'. *J. Opt. Soc. Am. A* **11**, 2350–2374.
- Cohen, S. and L. Guibas: 1999, 'The Earth Mover's Distance under transformation sets'. In: *Proceedings of Seventh IEEE International conference on computer vision*, Vol. 2. pp. 1076–1083.
- Comon, P.: 1994, 'Independent component analysis, a new concept?'. *Signal Processing, Special issue on higher-order statistics* **36**(3).
- Daugman, J.: 1985, 'Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters'. *Journal of the Optical Society of America A* **2**(7), 23–26.
- Deriugin, N. G.: 1956, 'The power spectrum and the correlation function of the television signal'. *Telecommunications* **1**, 1–12.
- Donoho, D.: 1995, 'Denoising by soft-thresholding'. *IEEE Trans. Info. Theory* **43**, 613–627.
- Donoho, D. L. and A. G. Flesia: 2001, 'Can Recent Innovations in Harmonic Analysis 'Explain' Key Findings in Natural Image Statistics'. *Network: Computation in Neural Systems* **12**(3), 371–393.
- Dror, R. O., T. K. Leung, E. H. Adelson, and A. S. Willsky: 2001, 'Statistics of real-world illumination'. In: *Proceedings of 2001 IEEE Conference on CVPR*, Vol. 2. pp. 164–171.
- Faugeras, O. D. and W. K. Pratt: 1980, 'Decorrelation methods of texture feature extraction'. *IEEE Pat. Anal. Mach. Intell.* **2**(4), 323–332.
- Field, D. J.: 1987, 'Relations between the statistics of natural images and the response properties of cortical cells'. *J. of Optical Society of America* **4**(12), 2379–2394.
- Gabor, D.: 1946, 'Theory of Communications'. *Journal of IEE (London)* **93**, 429–457.
- Geman, D. and A. Koloydenko: 1999, 'Invariant statistics and coding of natural microimages'. In: *Proc. of the IEEE Workshop on Statistical and Computational Theories of Vision*.
- Geman, S. and D. Geman: 1984, 'Stochastic Relaxation, Gibbs Distributions, and the Bayesian restoration of Images'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6**(6), 721–741.
- Grenander, U.: 1993, *General Pattern Theory*. Oxford University Press.
- Grenander, U., M. I. Miller, and A. Srivastava: 1998, 'Hilbert-Schmidt Lower Bounds for Estimators on Matrix Lie Groups for ATR'. *IEEE Transactions on PAMI* **20**(8), 790–802.
- Grenander, U. and A. Srivastava: 2001, 'Probability models for clutter in natural images'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(4), 424–429.

- Grenander, U., A. Srivastava, and M. I. Miller: 2000, 'Asymptotic Performance Analysis of Bayesian Object Recognition'. *IEEE Transactions of Information Theory* **46**(4), 1658–1666.
- Hamza, A. B., Y. He, and H. Krim: 2001, 'An information divergence measure for ISAR image registration'. In: *Proc. of IEEE workshop on statistical signal processing*.
- Hancock, P. J. B., R. J. Baddeley, and L. S. Smith: 1992, 'The principal components of natural images'. *Network* **3**, 61–70.
- He, Y., A. B. Hamza, and H. Krim: 2002, 'A Generalized Divergence Measure for Robust Image Registration'. *IEEE Transactions on Signal Processing to appear*.
- Heeger, D. J. and J. R. Bergen: 1995, 'Pyramid-based texture analysis/synthesis'. In: *Proceedings of SIGGRAPH*. pp. 229–238.
- Hero, A. O., B. Ma, O. Michel, and J. Gorman: 2001, 'Alpha-divergence for Classification, Indexing and Retrieval'. *Communication and Signal Processing Laboratory, Technical Report CSPL-328, U. of Mich.*
- Huang, J.: 2000, *Statistics of Natural Images and Models*. PhD thesis, Division of Applied Mathematics, Brown University, RI.
- Hyvarinen, A., J. Karhunen, and E. Oja: 2001, *Independent Component Analysis*. John Wiley and Sons.
- Jeulin, D., I. Terol-Villalobos, and A. Dubus: 1995, 'Morphological analysis of UO₂ powder using a dead leaves model'. *Microscopy, microanalysis, microstructure* **6**, 371–384.
- Kendall, W. S. and E. Thonnes: 1998, 'Perfect simulation in stochastic geometry'. *Preprint 323, Department of Statistics, University of Warwick, UK*.
- Kersten, D.: 1987, 'Predictability and Redundancy of Natural Images'. *Journal of Optical Society of America* **4**(12), 2395–2400.
- Kirby, M. and L. Sirovich: 1990, 'Application of the Karhunen-Loeve Procedure for the characterization of Human faces'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(1), 103–108.
- Kretzmer, E. R.: 1952, 'Statistics of television signals'. *Bell Syst. Tech. Journal* **31**, 751–763.
- Lee, A. B. and D. Mumford: 2001, 'Occlusion Models for Natural Images: A Statistical Study of Scale-Invariant Dead Leaves Model'. *International Journal of Computer Vision* **41**(1,2).
- Lee, A. B., K. S. Pedersen, and D. Mumford: 2002, 'The nonlinear statistics of high-contrast patches in natural images'. *Int'l J. of Computer Vision* **In press**.
- Lee, D. D. and H. S. Seung: 1999, 'Learning the parts of objects by non-negative matrix factorization'. *Nature* **401**, 788–791.
- Leporini, D., J.-C. Pesquet, and H. Krim: 1999, 'Best basis representations with prior statistical models'. In: P. Muller and B. Vidakovic (eds.): *Lecture notes in Statistics: Bayesian inferences in wavelet based models*. Springer Verlag.
- LoPresto, S. M., K. Ramchandran, and M. T. Orchard: 1997, 'Wavelet image coding based on a new generalized gaussian mixture model'. In: *Data Compression Conf. Snowbird, Utah*.
- Mallat, S. G.: 1989, 'A Theory for multiresolution signal decomposition: The wavelet representation'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, 674–693.
- Marr, D.: 1982, *VISION: A computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman and Company.
- Matheron, G.: 1975, *Random Sets and Integral Geometry*. John Wiley and Sons.
- Miller, K. D.: 1994, 'A Model for the development of simple-cell receptive fields and the arrangement of orientation columns through activity dependent competition between on- and off-center inputs'. *Journal of Neuroscience* **14**, 409–441.
- Miller, M. I. and L. Younes: 2002, 'Group Actions, Homeomorphisms, and Matching: A General Framework'. *International Journal of Computer Vision* **41**(1/2), 61–84.
- Moulin, P. and J. Liu: 1999, 'Analysis of multiresolution image denoising schemes using a generalized Gaussian and Complexity Priors'. *IEEE Trans. Info. Theory* **45**, 909–919.
- Mumford, D. and B. Gidas: 2001, 'Stochastic Models for Generic Images'. *Quarterly of Applied Mathematics* **59**(1), 85–111.
- Olshausen, B. A. and D. J. Field: 1996, 'Emergence of simple-cell receptive field properties by learning a sparse code for natural images'. *Nature* **381**, 607–609.
- Olshausen, B. A. and D. J. Field: 1997, 'Sparse coding with an overcomplete basis set: A strategy employed by V1?'. *Vision Research* **37**(23), 3311–3325.

- Pedersen, K. S. and A. B. Lee: 2002, 'Toward a full probability model of edges in natural images'. In: *Proc. of ECCV'02*.
- Portilla, J. and E. P. Simoncelli: 2000, 'A Parametric Texture Model based on Joint Statistics of Complex Wavelet Coefficients'. *International Journal of Computer Vision* **40**(1), 49–70.
- Portilla, J., V. Strela, M. Wainwright, and E. Simoncelli: 2001, 'Adaptive Wiener Denoising Using a Gaussian Scale Mixture Model in the Wavelet Domain'. In: *Proc 8th IEEE Int'l Conf on Image Proc.* Thessaloniki, Greece, pp. 37–40, IEEE Computer Society.
- Rinaldo, R. and G. Calvagno: 1995, 'Image Coding by Block Prediction of Multiresolution Subimages'. *IEEE Trans Im Proc.*
- Roweis, S. T. and L. K. Saul: 2000, 'Nonlinear Dimensionality reduction by locally linear embedding'. *Science* **290**, 2323–2326.
- Ruderman, D. L.: 1994, 'The Statistics of Natural Images'. *Network* **5**, 517–548.
- Ruderman, D. L.: 1997, 'Origins of Scaling in Natural Images'. *Vision Research* **37**(23), 3385–3398.
- Ruderman, D. L. and W. Bialek: 1994, 'Scaling of Natural Images: Scaling in the Woods'. *Physical Review Letters* **73**(6), 814–817.
- Serra, J. P.: 1982, *Image Analysis and Mathematical Morphology*. Academic Press, London.
- Shapiro, J.: 1993, 'Embedded Image Coding Using Zerotrees of Wavelet Coefficients'. *IEEE Trans Sig Proc* **41**(12), 3445–3462.
- Simoncelli, E. and B. Olshausen: 2001, 'Natural Image Statistics and Neural Representation'. *Annual Review of Neuroscience* **24**, 1193–1216.
- Simoncelli, E. P.: 1997, 'Statistical Models for Images: Compression, Restoration and Synthesis'. In: *Proc 31st Asilomar Conf on Signals, Systems and Computers*. Pacific Grove, CA, pp. 673–678, IEEE Computer Society.
- Simoncelli, E. P.: 1999, 'Bayesian Denoising of Visual Images in the Wavelet Domain'. In: *Bayesian Inference in Wavelet Based Models: Lecture Notes in Statistics, Volume 41*. Springer Verlag.
- Simoncelli, E. P. and E. H. Adelson: 1996, 'Noise Removal via Bayesian Wavelet Coring'. In: *Third Int'l Conf on Image Proc*, Vol. I. Lausanne, pp. 379–382, IEEE Sig Proc Society.
- Simoncelli, E. P. and R. T. Buccigrossi: 1997, 'Embedded Wavelet Image Compression Based On A Joint Probability Model'. In: *Proceedings of ICIP (1)*. pp. 640–643.
- Simoncelli, E. P., W. T. Freeman, E. H. Adelson, and D. J. Heeger: 1992, 'Shiftable Multi-scale Transforms'. *IEEE Trans Information Theory* **38**(2), 587–607.
- Srivastava, A., X. Liu, and U. Grenander: September, 2002, 'Universal Analytical Forms for Modeling Image Probability'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(9), to appear.
- Tenenbaum, J. B., V. Silva, , and J. C. Langford: 2000, 'A global geometric framework for nonlinear dimensionality reduction'. *Science* **290**, 2319–2323.
- Thomson, M. G. A.: 2001, 'Beats, Kurtosis, and Visual Coding'. *Network: Computation in Neural Systems* **12**(3), 271–287.
- Turiel, A., G. Mato, and N. Parga: 1998, 'The self-similarity properties of natural images resemble those of turbulent flows'. *Physical Review Letters* **80**(5), 1098–1101.
- Turiel, A. and N. Parga: 2000, 'The multi-fractal structure of contrast changes in natural images: from sharp edges to textures'. *Neural Computation* **12**, 763–793.
- Turiel, A., N. Parga, D. L. Ruderman, and T. W. Cronin: 2000, 'Multiscaling and information content of natural color images'. *Physical Review E* **62**(1), 1138–1148.
- van Hateren, J. A.: 1998, 'Independent component filters of natural images compared with simple cells in primary visual cortex'. *Proc. Royal Statistical Society of London (B)* **265**, 359:366.
- Vetterli, M.: 1984, 'Multidimensional subband coding: some theory and algorithms'. *Signal Processing* **6**(2), 97–112.
- Wainwright, M. J. and E. P. Simoncelli: 2000, 'Scale Mixtures of Gaussians and the Statistics of Natural Images'. *Advances in Neural Information Processing Systems, Editors: S. A. Solla, T. K. Leen, and K.-R. Muller* pp. 855–861.
- Wainwright, M. J., E. P. Simoncelli, and A. S. Willsky: 2001, 'Random Cascades on Wavelet Trees and Their Use in Modeling and Analyzing Natural Imagery'. *Applied and Computational Harmonic Analysis* **11**(1), 89–123.

- Watson, A. B.: 1987a, 'The Cortex transform: rapid computation of simulated neural images'. *Comp. Vis. Graphics Image Proc.* **39**, 311–327.
- Wegmann, B. and C. Zetsche: 1990, 'Statistical Dependence between Orientation Filter Outputs used in a Human Vision based image Code'. In: *Proceedings of Visual Communication and Image Processing*, Vol. 1360. pp. 909–922, Society of Photo-Optical Instrumentation Engineers.
- Weiss, Y.: 2001, 'Deriving intrinsic images from image sequences'. In: *Proceedings Eight IEEE International Conference on Computer Vision*, Vol. 2. pp. 68–75.
- Winkler, G.: 1995, *Image Analysis, Random Fields and Dynamic Monte Carlo Methods*. Springer.
- Yendrikhovskij, S. N.: 2001, 'Computing color categories from statistics of natural images'. *Journal of Imaging Science and Technology* **45**(5), 409–417.
- Zetsche, C.: 1997, 'Polyspectra of Natural Images'. In: *Presented at Natural Scene Statistics meeting*.
- Zetsche, C. and F. Rohrbein: 2001, 'Nonlinear and extra-classical receptive field properties and the statistics of natural images'. *Network: Computation in Neural Systems* **12**(3), 331:350.
- Zhu, S. C., C. E. Guo, Y. N. Wu, and Y. Z. Wang: 2002, 'What are Textons?'. In: *Proc. of European Conf on Computer Vision*.
- Zhu, S. C. and D. Mumford: 1997, 'Prior Learning and Gibbs Reaction-Diffusion'. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **19**(11), 1236–1250.
- Zhu, S. C., Y. N. Wu, and D. Mumford: 1997, "Minimax entropy principles and its application to texture modeling". *Neural Computation* **9**(8), 1627–1660.

