

Visually Distinct Patterns with Matching Subband Statistics

Joshua Gluckman, *Member, IEEE*

Abstract—A commonly used representation of a visual pattern is a statistical distribution measured from the output of a bank of filters (Gaussian, Laplacian, Gabor, etc.). Both marginal and joint distributions of filter responses have been advocated and effectively used for a variety of vision tasks, including texture classification, texture synthesis, object detection, and image retrieval. This paper examines the ability of these representations to discriminate between an arbitrary pair of visual stimuli. Examples of patterns are derived that provably possess the same marginal and joint statistical properties, yet are “visually distinct.” This is accomplished by showing sufficient conditions for matching the first k moments of the marginal distributions of a pair of images. Then, given a set of filters, we show how to match the marginal statistics of the subband images formed through convolution with the filter set. Next, joint statistics are examined and images with similar joint distributions of subband responses are shown. Finally, distinct periodic patterns are derived that possess approximately the same subband statistics for any arbitrary filter set.

Index Terms—Texture, statistical models, feature representation, moments.

1 INTRODUCTION

THE ability of a visual system to discriminate among a multitude of stimuli ultimately depends on the underlying representation of a visual pattern. Computing a large number of statistical measures from a set of filtered images is one commonly used representation. Many successful methods for object recognition, object detection, image retrieval, texture synthesis, and texture recognition have been developed based on such a representation [1], [2], [3], [4], [5], [6], [7], [8] (to name a few). In this class of methods, a set of subband images is created through convolution with a bank of filters, Gaussian, Laplacian, Gabor, etc., then statistical measures are computed from the subband images. A variety of statistical measures have been proposed, including parametric models, moments, entropies, histograms, and joint histograms. Detection and recognition are performed by classifying the statistical representation of a novel image while texture synthesis is performed by randomly sampling from the ensemble of images that match the statistics of a particular pattern. In order to understand the ability of these methods to scale to a large number of visual patterns, it is important to examine the ability of the underlying representation to discriminate between any pair of visual stimuli.

In this paper, we address the use of marginal statistics (histograms) and joint statistics (multidimensional histograms) of images filtered with Gaussian, Laplacian, derivative, and Gabor functions. In particular, we derive examples of visual patterns that are “distinct,” yet *provably* map to the same marginal probability density function (up to the first k moments) of a large filter bank. Second, we show the existence of “distinct” images with similar joint statistical

distributions. And third, we show the existence of periodic textures whose marginal and joint histograms have the same moments up to order 5 for any arbitrary set of filters.

The fact that these patterns exist does not imply that marginal and joint distributions are not important or useful. To the contrary, they have been used to successfully synthesize and classify a variety of textures and patterns. However, the results in this paper do suggest the need for certain classes of filters and the need for additional statistics in order to represent all “distinct” patterns.

2 STATISTICAL REPRESENTATIONS

There is a long history of using statistics to represent textures. One of the first models for representing a texture was proposed by Julesz, who suggested that cooccurrence statistics of N -tuples of pixels might explain texture perception [9]. The model was later disproved by Julesz and others by constructing stochastic processes that shared the same *ensemble* cooccurrence statistics for $N = 2$ and $N = 3$, yet produced patterns that were distinct to a human observer [10], [11]. These results have also been generalized by Gilbert for $N > 3$ [12]. Yellott has also constructed examples of distinct deterministic patterns that share the same *sample* cooccurrence statistics for $N = 2$ [13].

Since the existence of these counter-examples, much of the research on texture modeling has focused on filtering theory. Faugeras and Pratt first suggested representing textures by the marginal statistics of subband images after applying a bank of filters [14]. This line of research has been motivated by both neurophysiologic evidence that the human visual system decomposes the retinal image into subband channels and psychophysical evidence that textures sharing similar marginal distributions are difficult to discriminate [15], [16], [17].

Much of the early filter-based work concentrated on designing an effective set of filters for texture classification

• The author is with the Department of Computer Science, Polytechnic University, Brooklyn, NY 11201. E-mail: jgluckma@poly.edu.

Manuscript received 26 Sept. 2003; revised 9 July 2004; accepted 12 July 2004; published online 13 Dec. 2004.

Recommended for acceptance by J. Goutsias.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0293-0903.

and segmentation (see [18] for a review). Most of these methods first filter an input image to produce a set of subband images (or filter responses) and then compute a measure of local energy at each spatial location and for each subband image. However, the estimates of local energy typically only account for the low order moments, such as mean and variance, of the filter outputs [19], [20].

More recently, the existence of higher order moments (non-Gaussian statistics) in the filter response histograms of natural imagery has been observed [6], [21], [22], [23], [24], [25]. Partly motivated by this discovery, texture modeling and synthesis returned to histograms. Heeger and Bergen developed an algorithm for synthesizing textures that match the histograms of a target texture across the levels of a Steerable pyramid (oriented filters) [26]. Zhu et al. advocated using marginal distributions (histograms) and demonstrated the synthesis of a variety of visual patterns and textures by adaptively choosing a small set of filters so as to best represent the pattern [22]. An important theoretical result linking these filtering methods to Markov random field models was shown by Zhu et al. [27]. Parametric models, such as generalized Laplacian distributions and Bessel K forms, provide a good fit to the histograms of subband images and have also been successfully used for texture modeling [6], [21], [28].

Textures have also been represented using joint distributions (multidimensional histograms). Work by De Bonet and Viola has shown texture synthesis using the joint statistics of filter outputs [7]. In addition, Portilla and Simoncelli have also expressed the need for joint statistics when synthesizing textures [5]. Recently, good results on texture classification have been obtained using “textons,” which are clusters in the joint statistical space of filter outputs [29], [30], [31]. Joint distributions have also been used for object recognition, image retrieval, and segmentation [2], [3], [32].

While it is clear that statistical measures of subband images are a useful representation of visual patterns, questions regarding the expressive power of any particular statistic and any particular set of filters need to be answered in order to design efficient and robust visual systems.

Previously, Zhu et al. have shown that the histograms of filtered images are sufficient to uniquely characterize any pattern [22]. However, the proof requires an uncountably infinite number of filters and an exact match between the distributions. Hadjidemetriou et al. derived the class of continuous transformations that preserve the histogram of an unfiltered image [33]. They also studied the histograms of Gaussian filtered images and were able to demonstrate sensitivity to many texture properties [34]. Srivastava et al. derived a link between subband histograms and the “distinctness” and “sparseness” of image features [28].

The main contribution of this paper is to show the existence of “visually distinct” patterns that map to “similar” marginal and joint subband distributions for several classes of commonly used filters. We use the term “similar” because we only require the first k moments of the distributions to be the same and assume that the moments beyond k cannot be reliably estimated in the presence of noise.

The outline of this paper is as follows: In Section 3, marginal and joint statistics are defined and a specific class of

functions is defined for which we derive, in subsequent sections, conditions for matching the marginal and joint statistics. In Section 4, sufficient conditions are derived for matching the marginal statistics of these functions prior to filtering. In Section 5, this result is extended to ensure that the marginal statistics match after convolution with a bank of filters. Then, we discuss how to construct visually distinct patterns that have matching marginal statistics for some specific filter sets. In Section 6, sufficient conditions are shown for matching the joint subband statistics of the class of functions under consideration. Finally, in Section 7, an example is shown of a periodic texture with both matching marginal and joint statistics for any arbitrary filter set.

3 MARGINAL AND JOINT DISTRIBUTIONS

Given an image $f(x, y)$ defined on the continuous image domain $\Omega \subset \mathbb{R}^2$ and a set of filters $g = \{g_1, \dots, g_k\}$, the subband images $\{f_1, \dots, f_k\}$ are formed by the convolution $f_\alpha = g_\alpha * f$. The marginal probability distributions of the filter outputs consider each subband image in isolation and are defined by

$$p(i; f_\alpha) = \frac{1}{|\Omega|} \int_{\Omega} \delta(f_\alpha - i) d\Omega, \quad (1)$$

where δ is the Dirac delta function. For a fixed image size, the marginal distribution is equivalent to the histogram. Joint distributions are functions of the output of several or all filters and given by

$$p(i_1, i_2, \dots, i_k; f_1, \dots, f_k) = \frac{1}{|\Omega|} \int_{\Omega} \prod_{\alpha=1}^k \delta(f_\alpha - i_\alpha) d\Omega. \quad (2)$$

For a fixed set of filters, the joint distribution contains more information than the marginal distributions. However, it is difficult to reliably estimate joint distributions of high dimension from a single image due to insufficient data.

The first question addressed in this paper is whether distinct patterns f and f' exist that have similar marginal distributions for a particular set of filters g , meaning $p(f_\alpha) \approx p(f'_\alpha), \forall \alpha : g_\alpha \in g$. The second question addressed is whether distinct patterns f and f' exist that have similar joint distributions for a particular set of filters g , meaning $p(f_1, \dots, f_k) \approx p(f'_1, \dots, f'_k)$. The third question addressed is whether distinct periodic patterns exist that have similar marginal and joint distributions for an arbitrary filter set.

To answer the questions posed in this paper, we consider a particular class of functions. These functions are the sum of an arbitrary but fixed number of cosine functions occurring at integer frequencies with an arbitrary but fixed set of weights.

Definition 1. Given a set of n weights $w = (w_1, \dots, w_n) : w_i \in \mathbb{R}$, define $F(w)$ to be the set of functions such that for all $u = (u_1, \dots, u_n) \in \mathbb{Z}^n$ and $v = (v_1, \dots, v_n) \in \mathbb{Z}^n$,

$$\sum_{i=1}^n w_i \cos(u_i x + v_i y) \in F(w).$$

The reason for choosing these functions is two-fold. First, they are simple enough that we can analyze the relationship between the functions and the subband statistics. Second,

they are flexible enough that we can create visually distinct patterns. Although these functions may not be able to represent natural looking textures, they are sufficient for deriving complex artificial patterns.

4 MOMENTS OF IMAGES

In this section, we derive sufficient conditions to ensure that a pair of functions $f, f' \in F(w)$ have the same or similar histograms, $p(f) = p(f')$. The conditions are independent of the weights w and only depend on the choice of integer frequency pairs $\{(u_1, v_1), \dots, (u_n, v_n)\}$ for f and $\{(u'_1, v'_1), \dots, (u'_n, v'_n)\}$ for f' .

The density (1) and (2) are not defined directly on the image f , therefore, we work with the moments of the distributions. Let M^k be the k th moment of the distribution (1),

$$M_k(f) = \int i^k p(i) di = \frac{1}{|\Omega|} \int_{\Omega} f(x, y)^k d\Omega.$$

Provided the existence of the moment generating function, the moments uniquely represent the probability distribution [35]. For images, only a finite number of moments are needed to represent the distribution due to the presence of limited dynamic range, quantization, and noise. Thus, sufficient conditions are derived for ensuring the first k moments are identical for a pair of functions f and f' in $F(w)$, implying that $p(f) \approx p(f')$. Note that we are not advocating the use of moments as a representation, but as a tool for examining the expressive power of marginal and joint statistics.

Dropping the normalization constant, the moments for a function in $F(w)$ are

$$M_k = \int_{\Omega} \left(\sum_{i=1}^n w_i \cos(z_i) \right)^k d\Omega, \quad (3)$$

where $z_i = u_i x + v_i y$. Applying the multinomial theorem,¹

$$M_k = \int_{\Omega} \sum_{k_1, k_2, \dots, k_n} \frac{k!}{k_1! k_2! \dots k_n!} \prod_{i=1}^n w_i^{k_i} \cos^{k_i}(z_i) d\Omega,$$

where the sum is taken over all of the nonnegative integers such that $\sum_i^n k_i = k$. To simplify the notation, this is written as

$$M_k = \int_{\Omega} \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \prod_{i=1}^n \cos^{k_i}(z_i) d\Omega,$$

with $\kappa = (k_1, \dots, k_n)$ and $K = \{\kappa : \forall i k_i \geq 0, \sum_i^n k_i = k\}$. In order to remove the exponentiation, the following substitution is applied:

$$\cos^a(z) = \sum_{i=0}^a C(i, a) \cos(iz), \quad a \in \mathbb{Z}.$$

We are not concerned with the actual values of $C(i, a)$ except to note that they are independent of z and have the following property:

$$C(i, a) \neq 0 \Leftrightarrow \text{parity}(i) = \text{parity}(a). \quad (4)$$

Applying the substitution,

$$M_k = \int_{\Omega} \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \prod_{i=1}^n \left(\sum_{r_i=0}^{k_i} C(r_i, k_i) \cos(r_i z_i) \right) d\Omega.$$

Interchanging the product and the summation results in

$$M_k = \int_{\Omega} \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \sum_{r \in R_{\kappa}} \prod_{i=1}^n C(r_i, k_i) \cos(r_i z_i) d\Omega,$$

where $r = \{r_1, \dots, r_n\}$ and the summation is taken over the set $R_{\kappa} = \{r : 0 \leq r_i \leq k_i, \forall i\}$. Again simplifying the notation and moving the integral inside the summation,

$$M_k = \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \sum_{r \in R_{\kappa}} \mathcal{C}(r, \kappa) \int_{\Omega} \prod_{i=1}^n \cos(r_i z_i) d\Omega, \quad (5)$$

where $\mathcal{C}(r, \kappa) = \prod_{i=1}^n C(r_i, k_i)$. From (4), $\mathcal{C}(r, \kappa)$ has the following property:

$$\mathcal{C}(r, \kappa) \neq 0 \Leftrightarrow \forall i \text{ parity}(r_i) = \text{parity}(k_i). \quad (6)$$

The product is removed by repeatedly applying by the following trigonometric substitution:

$$\cos(a) \cos(b) = \frac{1}{2} (\cos(a+b) + \cos(a-b)),$$

resulting in

$$M_k = \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \sum_{r \in R_{\kappa}} \mathcal{C}(r, \kappa) \int_{\Omega} \frac{1}{2^{n-1}} \sum_{d_2, \dots, d_n} \cos\left(\sum_{i=1}^n d_i r_i z_i\right) d\Omega, \quad (7)$$

where $d_1 = 1$ and $d_i = \pm 1$ for $i \geq 2$. The left summation inside the integral is taken over the 2^{n-1} possible assignments of d_i . Note that each integral is of the form

$$\int_{\Omega} \cos(r_1 z_1 \pm r_2 z_2 \pm \dots \pm r_n z_n) d\Omega.$$

Recall that $z_i = u_i x + v_i y$ and $r_i, u_i, v_i \in \mathbb{Z}$; therefore, if the integration is performed over $(-\pi, \pi)$ in both the x and y direction, each integral evaluates to 0, unless all of the terms sum to 0, in which case the integral evaluates to $|\Omega| = 4\pi^2$. Thus,

$$\int_{\Omega} \cos\left(\sum_{i=1}^n d_i r_i z_i\right) d\Omega = \begin{cases} 4\pi^2 & \text{if } \sum_i^n d_i r_i z_i = 0 \\ 0 & \text{otherwise.} \end{cases}$$

From this result, the integral is removed and

$$M_k = \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \sum_{r \in R_{\kappa}} \mathcal{C}(r, \kappa) \frac{1}{2^{n-1}} t(u, v, r). \quad (8)$$

Note that the $4\pi^2$ is removed because it cancels with the normalization constant that was dropped in (3). The function $t(u, v, r)$ is the number of solutions to the 2^{n-1} equations of the form

$$r_1 z_1 \pm r_2 z_2 \pm \dots \pm r_n z_n = 0,$$

which is an integer in the range $(0, 2^{n-1})$. Note that t is a function of u, v, r because a particular equation,

$$d_1 r_1 z_1 + \dots + d_n r_n z_n = 0,$$

1. A generalization of the binomial theorem.

is satisfied if and only if

$$\sum_{i=1}^n d_i r_i u_i = 0 \wedge \sum_{i=1}^n d_i r_i v_i = 0. \quad (9)$$

Equation (8) separates out the dependence of the moment function on the integer frequencies u and v and will be used to prove several results linking the marginal distributions to the image frequencies. From (9), the function $t(u, v, r)$ is determined by the solutions to a pair of Diophantine equations.²

Definition 2. Given a function f with the integer frequencies (u_1, \dots, u_n) and (v_1, \dots, v_n) , define $A(f) = \{(m_1, \dots, m_n)\} \subseteq \mathbb{Z}^n$ to be the set of solutions to the Diophantine equations

$$\begin{aligned} u_1 m_1 + u_2 m_2 + \dots + u_n m_n &= 0 \\ v_1 m_1 + v_2 m_2 + \dots + v_n m_n &= 0. \end{aligned} \quad (10)$$

Now, we are ready to show sufficient conditions under which a pair of functions f and f' have the same moments.

Theorem 1. Let f and f' be functions in $F(w)$. If $\forall c \in A(f) \ominus A(f')$, there exists an integer \hat{k} such that $\sum_i |c_i| > \hat{k}$, then $\forall k \leq \hat{k}$, $M_k(f) = M_k(f')$.

Proof. We will prove the contrapositive statement. Thus, we assume there exists a $k \leq \hat{k}$ such that $M_k(f) \neq M_k(f')$. From (8), there must exist a $\kappa = (k_1, \dots, k_n)$ and an $r = (r_1, \dots, r_n)$ such that $t(u, v, r) \neq t(u', v', r)$, where $\sum_{i=1}^n k_i = k$ and $r_i \leq k_i$. Without loss of generality, assume that $t(u, v, r) > t(u', v', r)$, which implies there must exist a (d_1, \dots, d_n) where $d_i \in \{-1, 1\}$ such that

$$\begin{aligned} d_1 r_1 (u_1 x + v_1 y) + \dots + d_n r_n (u_n x + v_n y) &= 0 \\ \wedge \\ d_1 r_1 (u'_1 x + v'_1 y) + \dots + d_n r_n (u'_n x + v'_n y) &\neq 0. \end{aligned}$$

Define $c = (c_1, \dots, c_n)$ such that $c_i = d_i r_i$. Clearly, $c \in A(f)$ and $c \notin A(f')$, thus $\exists c \in A(f) \ominus A(f')$.³ Recall that $r_i \leq k_i$ and $\sum_i k_i = k$ and $k \leq \hat{k}$, hence $\sum_i |c_i| \leq \hat{k}$ and we have proven the contrapositive statement. \square

Corollary 1. If $A(f) \ominus A(f') = \emptyset$, then for all k the moments $M_k(f) = M_k(f')$.

Proof. This must follow directly from the theorem to avoid a contradiction. \square

Theorem 1 relates the histogram to the integer frequencies through the solution set of a pair of Diophantine equations. This theorem states that a pair of functions with the same solution set up to a magnitude of k have identical moments up to k . Given any pair of functions in $F(w)$, the number of identical moments are found by generating the solutions to the Diophantine equations until a member of $A(f) \ominus A(f')$ is found. Note that this process will always terminate, provided $A(f) \neq A(f')$, which is the case if (u, v) and (u', v') are linearly independent. From the corollary, if the frequencies are related by a linear transform, then all of

the moments of f and f' will be the same and $p(f) = p(f')$. By relaxing the constraint that all of the moments must match, functions can be generated that are not related by an affine transformation yet have $p(f) \approx p(f')$, which, in the presence of quantization and noise, will be statistically indistinguishable.

5 MOMENTS OF FILTERED IMAGES

Next, we consider the marginal distributions of the subband images of f and f' . Given an arbitrary filter g , the subband images are formed by the convolution $g * f$ and $g * f'$. A sufficient condition for matching the moments of the subband images for a particular filter is given by the following theorem:

Theorem 2. Let f and f' be functions in $F(w)$. Define \hat{k} to be $< \sum_i^n |c_i|$ for all $c \in A(f) \ominus A(f')$. Let g be any symmetric or antisymmetric function with the Fourier transform G . If $\forall i, G(u_i, v_i) = G(u'_i, v'_i)$, then $\forall k : k \leq \hat{k}$, $M_k(g * f) = M_k(g * f')$.

Proof. Case 1: The function g is symmetric. The symmetry implies that G is real valued, hence from the properties of the Fourier transform, it can be shown that

$$\begin{aligned} g * f &= \sum_{i=1}^n G(u_i, v_i) w_i \cos(ux + vy) \\ g * f' &= \sum_{i=1}^n G(u'_i, v'_i) w_i \cos(u'x + v'y). \end{aligned}$$

If $G(u_i, v_i) = G(u'_i, v'_i)$ defines $m = (m_1, \dots, m_n) : m_i = G(u_i, v_i) w_i$, thus $g * f \in F(m)$ and $g * f' \in F(m)$. Note that the integer frequencies have remained the same, therefore, $A(g * f) = A(f)$ and $A(g * f') = A(f')$. This implies that $\hat{k} < \sum_i^n |c_i|$ for all $c \in A(g * f) \ominus A(g * f')$, thus $\forall k : k \leq \hat{k}$, $M_k(g * f) = M_k(g * f')$ follows from Theorem 1.

Case 2: The function g is antisymmetric. The anti-symmetry implies that G is imaginary, hence, from the properties of the Fourier transform, it can be shown that

$$\begin{aligned} g * f &= \sum_{i=1}^n G(u_i, v_i) w_i \sin(ux + vy) \\ g * f' &= \sum_{i=1}^n G(u'_i, v'_i) w_i \sin(u'x + v'y). \end{aligned}$$

The sine functions are cosine functions with a $\pi/2$ shift. Moments are invariant to constant shifts, hence, case 2 follows from case 1. \square

Theorem 2 extends Theorem 1 to include the subband images, provided the frequencies of f and f' are chosen such that each pair are on the iso-contours of the Fourier transform of the filter used to create the subbands. Together, Theorems 1 and 2 provide sufficient conditions for ensuring the first k moments of the histograms of a set of subband images match for a pair of functions f and f' and a set of filters $g = \{g_1, \dots\}$. We are not concerned with the actual value of k , except to note that it can always be defined large enough that the moments above k are lost in quantization and noise, thus we claim $p_\alpha(f) = p_\alpha(f'), \forall \alpha : g_\alpha \in g$. To construct examples of f and

2. An equation in which only integer solutions are allowed.

3. \ominus is the symmetric set difference operator, which is the element belonging to one but not both sets.

f' with matching moments, the frequencies (u, v) and (u', v') must be chosen carefully such that:

1. The solutions to the Diophantine equations (10) of magnitude $\leq k$ are identical.
2. For each filter $g_\alpha \in g$ with Fourier transform G_α , we have $\forall i, G_\alpha(u_i, v_i) = G_\alpha(u'_i, v'_i)$.

The general approach taken to construct examples of f and f' is to define two sets of frequencies, $u_a = \{(u_1, v_1), \dots, (u_n, v_n)\}$ and $u_b = \{(u_{n+1}, v_{n+1}), \dots, (u_m, v_m)\}$. Then, define $f = f_a + f_b$, where f_a is constructed from the frequency u_a and f_b from u_b . The weights $(w_1, \dots, w_n, w_{n+1}, \dots, w_m)$ are set to 1. Similarly, $f' = f_a + R(f_b)$ is defined where R is a rotation of the frequencies u_b and the weights are also set to 1. The exact choice of R will depend on the filters, which is addressed later. The key is to choose the frequencies such that u_a, u_b do not "interact" and $u_a, R(u_b)$ do not "interact" to contribute to any of the moments $\leq k$. Formally, if $c = (c_1, \dots, c_{n+m}) \in A(f) : \sum |c_i| \leq k$, then we require $(c_1, \dots, c_n) \in A(f_a) \vee (c_{n+1}, \dots, c_m) \in A(f_b)$. Likewise, if $c \in A(f') : \sum |c_i| \leq k$, then we require $(c_1, \dots, c_n) \in A(f_a) \vee (c_{n+1}, \dots, c_m) \in A(R(f_b))$. These properties imply that $c \notin A(f) \ominus A(f')$. Rather than iterating through possible choices of u_a and u_b , simple geometric heuristics are applied to avoid interacting with any small moments. For example, ensuring that no two members of u_b and $R(u_b)$ are parallel to any two members of u_a . The frequencies are also chosen such that "visually distinct" patterns emerge. This is accomplished by placing members of u_b close to members of u_a , thus creating interference patterns that will not exist between $R(u_b)$ and u_a . Note that although R is an affine transform, the functions f and f' will, in general, not be related by an affine transformation because they share a set of fixed frequencies f_a . To handle certain classes of filters we need one more theorem.

Theorem 3. For a function $f \in F(w)$, if $\forall i, u_i$ are odd or $\forall i, v_i$ are odd, then $\forall k$ such that k is odd, $M_k(f) = 0$.

Proof. Consider an arbitrary odd moment k and an arbitrary κ and r in the summation of (8). To prove that $M_k = 0$, we only need to show that $t(u, v, r)$ is zero for the arbitrary term.

Define the set O_κ to contain the odd parity terms of κ . Define the set O_r to contain the odd parity terms of r . Define the set O_u to contain the odd parity terms of $d_i r_i u_i$. The fact that k is odd and $\sum_i^n k_i = k$ implies that an odd number of odd parity k_i terms must exist, hence $|O_\kappa|$ is odd. We need only consider the terms in (8) where $\mathcal{C}(r, \kappa) \neq 0$, hence, from (6), we also know that $|O_r|$ is odd. Now, if all of the u_i are odd, the fact that $|O_r|$ is odd implies that $|O_u|$ is odd. If $|O_u|$ is odd, then $\sum_i^n d_i r_i u_i \neq 0$, hence $t(u, v, r) = 0$ for the arbitrary term and, therefore, $M_k = 0$. The same holds true if the v_i are odd. \square

A filtering operation does not alter the frequencies of a function, it only effects the weights. Thus, by using odd parity frequencies, we have a way of ensuring all of the odd moments are zero for all of the subband images of f and f' .

Now, we are ready to consider examples of f and f' for several classes of commonly used filters. In order to create example patterns, discrete images must be formed. However, the theorems have assumed a continuous image model and will not necessarily be true for sampled functions. We work around this problem by randomly sampling the

functions inside each pixel in which case the expected value of the first k moments match.

5.1 Rotationally Symmetric Filters

Rotationally symmetric filters include Gaussian, Laplacian of Gaussian, and other center surround kernels. These filters are defined with a scale parameter σ , which determines the size of the kernel. For example, the Gaussian function is

$$g(x, y; \sigma) = ce^{-(x+y)^2/2\sigma^2},$$

where c is a normalization constant. They are commonly used for many classification, detection, and recognition tasks in computer vision. These filters offer the advantage of being invariant to rotations of a visual pattern and have been used for texture classification and image retrieval [3] [34].

The common feature of this class of filters is a rotationally symmetric Fourier transform. Any rotation R that maps the integer frequencies to another set of integer frequencies results in exactly the same Fourier coefficients. Integer rotations are easily achieved by interchanging the u and v frequencies. Noninteger rotations can also be applied. Although the coefficients of the Fourier transform will not be identical, the difference is very small for all but the lowest frequencies and the effect on the first k moments is negligible. Thus, $f = f_a + f_b$ and $f' = f_a + R(f_b)$ have the same weights after convolution with any rotationally symmetric filter applied at any scale, even if an infinite number of scales is used.

Fig. 1a and Fig. 1b show examples of such patterns sampled with 256×256 pixels. Next to each pattern is a figure showing the distribution of the frequencies used to construct the pattern. In Fig. 1a, the fixed frequencies u_a are denoted by circles, the rotated frequencies u_b are denoted by plus signs, and the first 35 moments are identical. In Fig. 1b, the frequencies are divided into several groups and different integer rotations are applied to each. Only the first four moments are identical, however, the moments beyond four are highly correlated and nearly identical for the two patterns. This is evident by the similarity of the histograms, which are dependent on all of the moments. It is not the exact value of k that is important, but the similarity of the histograms that makes it difficult for marginal statistics to distinguish these patterns. In fact, this example is used to demonstrate that only the first few moments, usually four, need to be equated in order to produce similar marginal distributions.

Each pattern is convolved with both Gaussian and Laplacian of Gaussian kernels at scales $\sigma = \{1, 2, \dots, 8\}$, for a total of 16 subband images. For each pair of subband images, histograms are formed with 100 bins and the L_1 norm is computed. For all histograms, the difference is between 1 and 3 percent of the number of pixels—a negligible amount. Next to the patterns in Fig. 1a and Fig. 1b are example histograms for the Gaussian and Laplacian filters that resulted in the largest L_1 norm. Note that the choice of norm is not important. All of the histograms are nearly identical. We only use the L_1 norm to determine which ones to display. From the similarity of the histograms, it is clear that no statistical test, such as L_1 , L_2 , χ^2 , or K-L divergence, will be able to reliably

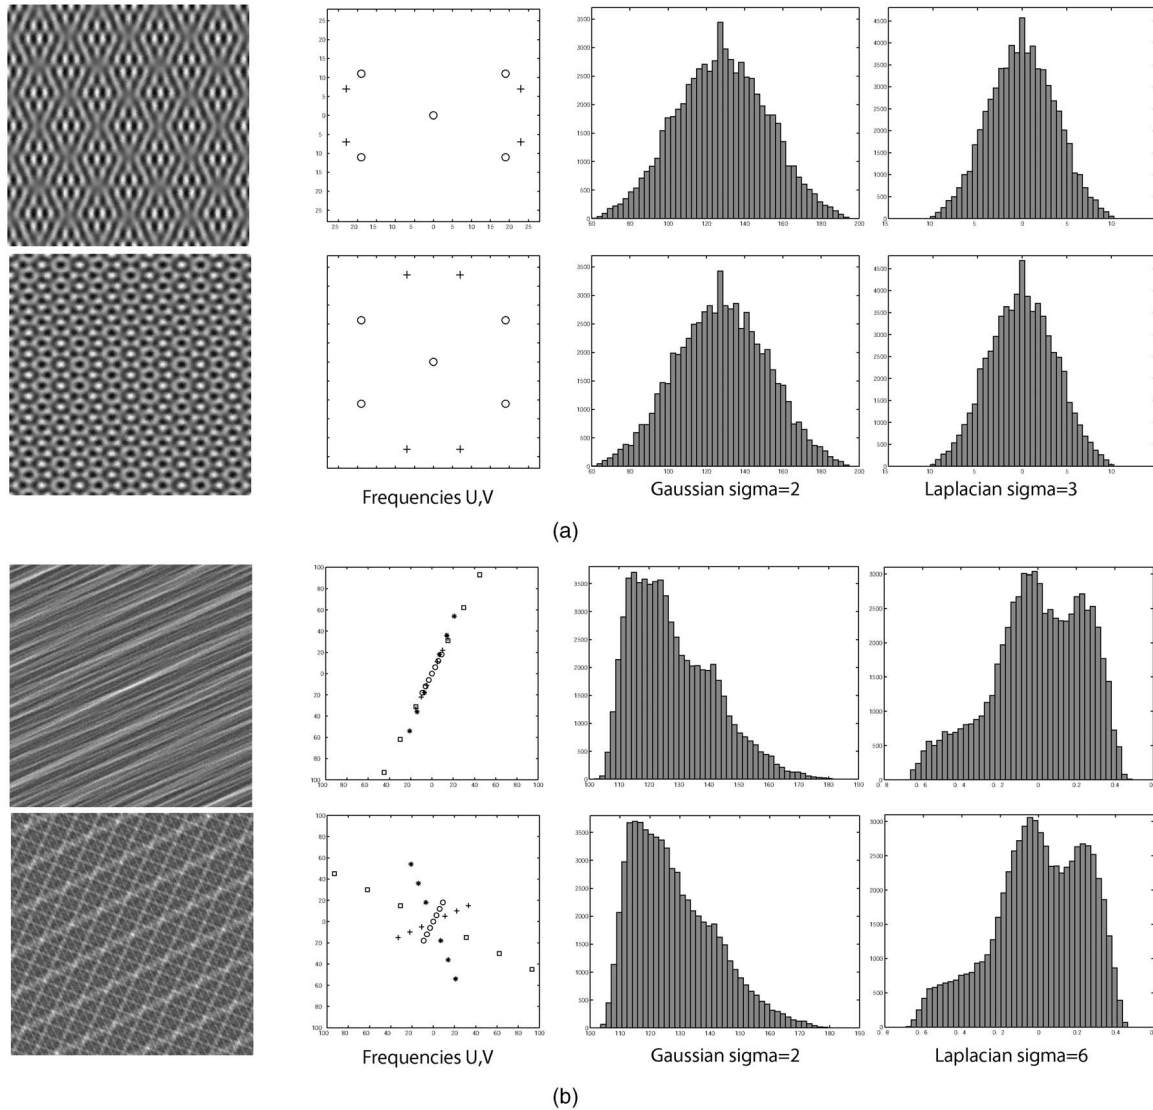


Fig. 1. Visually distinct patterns with matching subband statistics for rotationally symmetric filters. Each pair of images in the left column is constructed with identical moments up to $k = 35$ for (a) and $k = 4$ for (b). Next to each image is a graph of the frequencies (Fourier transform) that make up the pattern. Visually distinct differences are created by grouping the frequencies into sets and applying rotations independently to each set. Only two sets are used in (a), while several are used in (b). Both Gaussian and Laplacian of Gaussian filters were applied at eight different scales for a total of 16 subband images. All of the histograms of the subbands are nearly identical. Examples of the subband histograms are shown on the left. Given the similarity, it is unlikely that any metric would be able to reliably differentiate these patterns from their marginal statistics.

differentiate these patterns from their histograms using rotationally symmetric filters.

5.2 Derivative Filters

Another commonly used class of filters is the derivatives of the Gaussian function. They have three parameters: the order of the spatial derivative in the x and y directions and the scale of the Gaussian,

$$g(x, y; m, n, \sigma) = \frac{\partial^{m+n}}{\partial x^m \partial y^n} e^{-(x+y)^2/2\sigma^2}.$$

Even order derivative filters are symmetric, while odd order derivative filters are antisymmetric. The Fourier transform of a derivative filter is a polynomial in u, v modulated by a Gaussian and given by

$$G(u, v) = (2\pi i u)^m (2\pi i v)^n e^{-(u+v)^2/2\sigma^2}.$$

It is straightforward to show that the coefficients of a derivative of Gaussian are such that $G(u, v) = G(-u, v)$ for symmetric filters and $G(u, v) = -G(-u, v)$ for antisymmetric filters. A reflection transformation preserves the coefficients of a symmetric filter, thus $f(u, v) = f_a(u, v) + f_b(u, v)$ and $f'(u, v) = f_a(u, v) + f_b(-u, v)$. To handle antisymmetric filters, an additional constraint is needed. For an antisymmetric derivative filter g , we have $g * f = g * f_a(u, v) + g * f_b(u, v)$ and $g * f' = g * f_a(u, v) - g * f_b(-u, v)$. This causes a change of sign in the odd moments, resulting in a reflected histogram. However, if all of the odd moments are zero, then the histogram is symmetric and invariant to a reflection. By Theorem 3, imposing the constraint that either the u or v frequencies of f_b are odd forces the odd moments to zero. Example patterns are shown in Fig. 2a and Fig. 2b. For Fig. 2a, the first nine moments match and, for Fig. 2b, the first 39 moments match. A total of 80 subband images are formed using

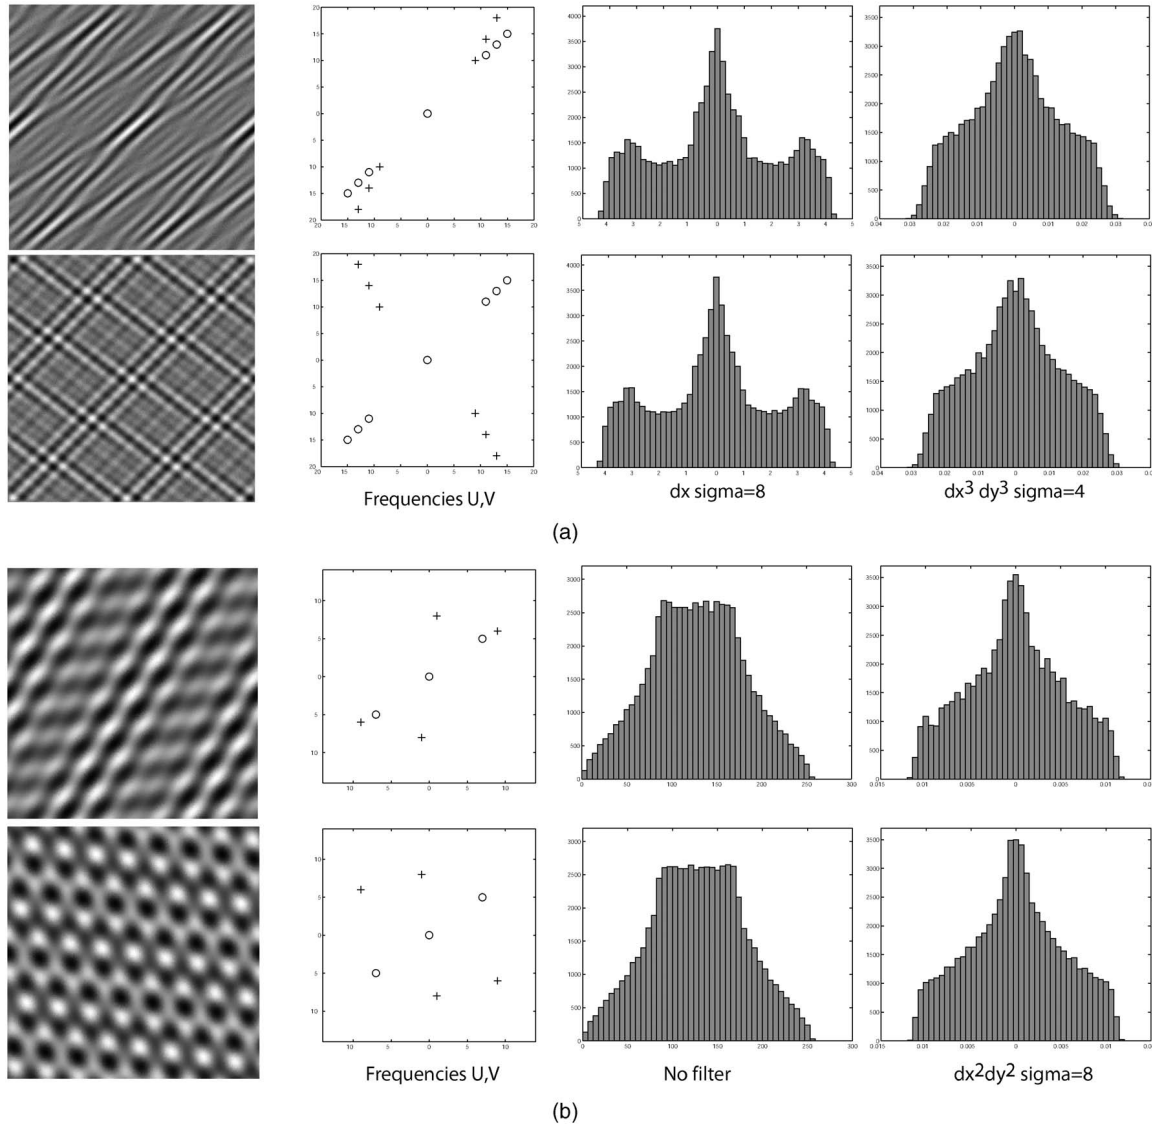


Fig. 2. Visually distinct patterns with matching subband statistics for rotationally symmetric and Gaussian derivative filters of all orders. Each pair of images in the left column is constructed with identical moments up to $k = 9$ for (a) and $k = 39$ for (b). Next to each image is a graph of the frequencies (Fourier transform) that make up the pattern. Visually distinct differences are created by grouping the frequencies into sets and applying reflections independently to each set. Derivative filters at eight different scales and up to order three were applied, creating a total of 80 subband images for each pattern. All of the histograms of the subbands are nearly identical. Because a reflection transformation of the frequencies is also an integer rotation, the histograms of rotationally symmetric filters are also nearly identical. Examples of the subband histograms with the largest L_1 difference are shown on the left. Again, given the similarity, it is unlikely that any metric would be able to reliably differentiate these patterns from their marginal statistics.

derivative filters of all orders between zero and three and each filter is applied at scales $\sigma = \{1, 2, \dots, 8\}$. Again, all histograms matched within an L_1 difference of 1 to 3 percent. The histogram with the largest L_1 norm and the histogram of the unfiltered image are shown. Although derivative filters have convenient mathematical properties, such being orthogonal and separable, they are clearly not able to distinguish these patterns. Note that higher order derivatives or additional scales will not help.

5.3 Gabor Filters

Gabor functions are a class of oriented filters and are extensively used in texture modeling. These functions come in symmetric/antisymmetric pairs where each pair is a

cosine/sine wave modulated by a Gaussian function. For example, a symmetric or real Gabor function takes the form

$$g(x, y; \sigma_a, \sigma_b, \theta, u) = ce^{-(x^2/\sigma_a^2 + y^2/\sigma_b^2)} \cos(2\pi u).$$

The parameters are the scale and aspect ratio of the Gaussian, which are controlled by σ_a and σ_b . The orientation parameter θ rotates the xy coordinate system and u is the frequency parameter.

The Fourier transform of a Gabor function is a Gaussian function centered at the frequency of the cosine/sine wave. A bank of Gabor filters is usually designed by choosing parameters that tile the frequency plane, with Gaussian functions occurring at regular intervals in orientation and scale. Given a bank of Gabor filters with orientations

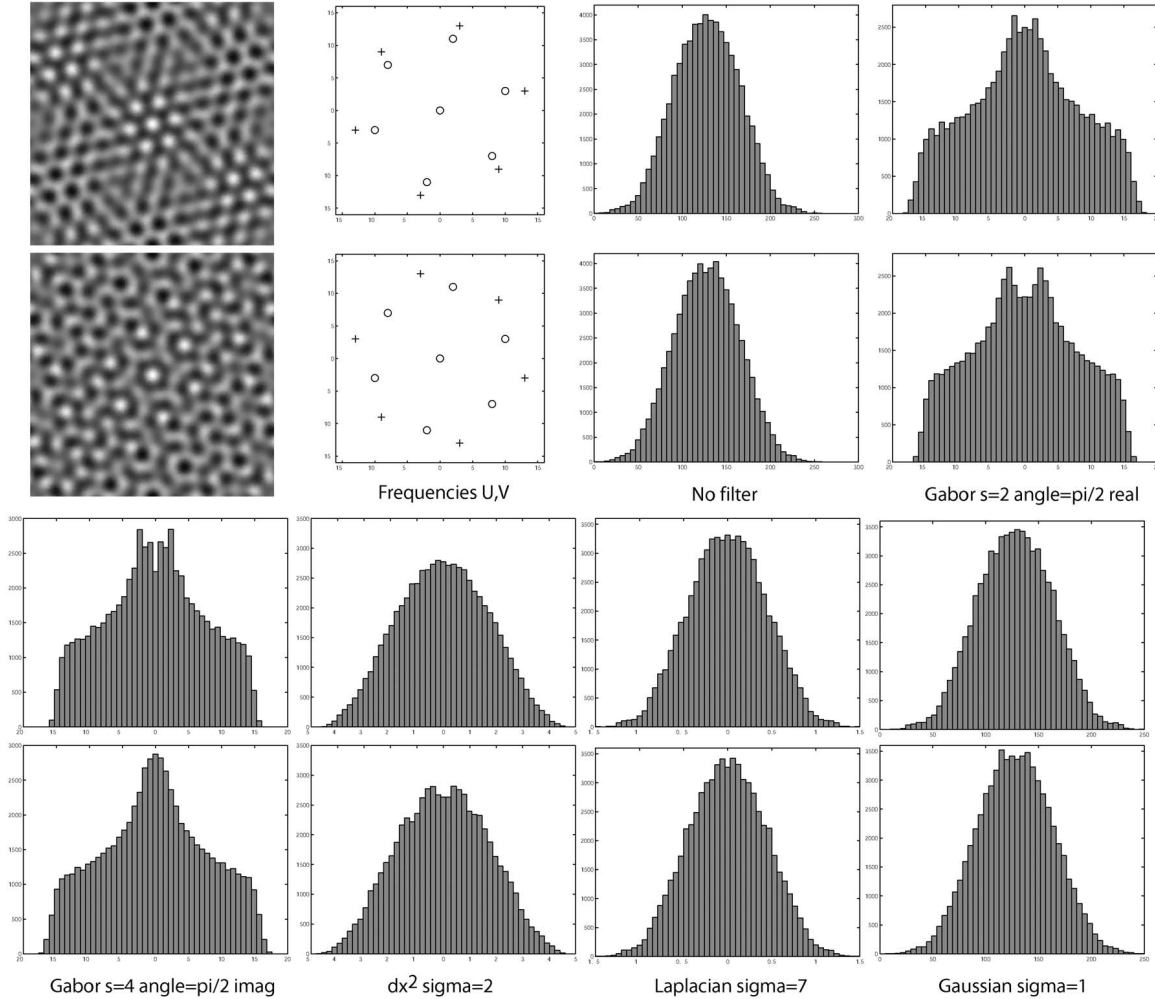


Fig. 3. Visually distinct patterns with matching subband statistics for rotationally symmetric, Gaussian derivative filters of all orders and Gabor filters at all scales with an angular separation of $\psi = 30$ degrees. The images are constructed with identical moments up to $k = 5$. Next to each image is a graph of the frequencies (Fourier transform) that make up the pattern. For each image, subbands were formed using 48 Gabor filters (six orientations and four scales), 80 derivative filters (eight scales and up to order three), and 16 Gaussian and Laplacian filters. All 144 histograms are nearly identical. A sampling of the histograms for a variety of filters is shown. Again, given the similarity, it is unlikely that any metric would be able to reliably differentiate these patterns from their marginal statistics. The reason subband statistics have a difficult time is that filters essentially perform local feature analysis. A close inspection of these images shows nearly identical distribution of local features. They are visually distinct due to a higher level grouping of these local features. Yet, multiresolution analysis will not be of help because the frequencies of these images occur at exactly the same scale.

$\{\theta_1, \theta_2, \dots\}$ and a spacing of ψ between consecutive orientations, the Fourier coefficients of all of the filters at orientation θ_i have the following property for any angle ϕ :

$$G(r \cos \theta_i + \phi, r \sin \theta_i + \phi) = G(r \cos \theta_i - \phi, r \sin \theta_i - \phi). \tag{11}$$

Therefore, if we define the frequencies of f_b such that they have the orientations $\{\psi/2, 5\psi/2, 9\psi/2, \dots\}$ and R is a rotation of angle ψ , the coefficients of f_b for an arbitrary filter will be identical to $R(f_b)$. For example, consider a bank of filters with orientations $\pi i/6$ for $i = 0 \dots 5$, where $\psi = \pi/6$. The orientations of the frequencies of f_b are defined by $(1 + 4j)\pi/12$ for $j = 0 \dots 5$ and the orientations of $R(f_b)$ are defined by $(3 + 4k)\pi/12$ for $k = 0 \dots 5$. For an arbitrary filter at orientation $\pi i/6$, the Fourier coefficient of the frequency location with orientation $(1 + 4j)\pi/12$ matches that of $(3 + 4k)\pi/12$, where $k = (5 - j + i) \bmod 6$.

This follows from the property (11) of the Gabor functions. Such a pattern is discussed in the following subsection.

5.4 Combining Filters

Function pairs that share the same marginal statistics for combinations of these filters also exist. The patterns derived for derivative and Gabor filters also have the same distributions for rotationally symmetric filters. However, the transformation R for derivative filters is a reflection which will not satisfy the constraint on R for Gabor filters. But, the rotation for Gabor filters is also a reflection, thus the constraints for Gabor filters imply the constraints for the symmetric derivative filters. To satisfy the constraints for the antisymmetric derivative filters, the odd frequency constraint must be added to the Gabor filters. Fig. 3 shows an example of a pair of patterns that have been defined this way for a bank of Gabor filters at orientations $\{0, \pi/6, \pi/3, \dots\}$. The patterns only share the first five moments, however,

they still have nearly identical histograms. Each pattern has been convolved with a total of 48 Gabor filters (six orientations and four scales), 80 derivative filters, and 16 rotationally symmetric filters. The marginal densities of all 144 subband images matched to within 6 percent. A sampling of the histograms is shown in Fig. 3.

It is important to note that these patterns could be differentiated by increasing the number of orientations, but not the number of scales. Although a new pair of patterns could be defined given the new set of orientations, as the number of orientations increases, eventually the "visual distinctness" will vanish. This is a marked difference from the rotationally symmetric and the derivative filters in that increasing the number of these filters does not help. The key difference is the orientation selectivity of the Gabor functions.

It is also important to note that the subband distributions are defined over the extent of the pattern and, thus, are global statistics. Therefore, it may also be possible to differentiate these patterns from local statistics.

6 JOINT STATISTICS

While marginal statistics consider each subband image in isolation, joint statistics are multidimensional histograms constructed from a set of subband images. Several researchers have developed techniques for texture synthesis based on the joint statistics of filter outputs [5], [7]. Recently, good results on texture classification have been obtained using "textons," which are clusters in the joint statistical space of filter outputs [29], [30], [31]. Joint distributions have also been used for object recognition, image retrieval, and segmentation [2], [3], [32].

Just as marginal statistics can be studied using moments, joint statistics can be studied using multivariate moments. Given m subband images formed from the filter set $g = \{g_1, \dots, g_m\}$, an m -dimensional probability distribution is defined by (2). The order $k = \sum_{j=1}^m k_j$ moments of the m -dimensional distribution are defined directly on the images by

$$M_{k_1 \dots k_m} = \frac{1}{|\Omega|} \int_{\Omega} \prod_{j=1}^m f_j(x, y)^{k_j} d\Omega, \quad (12)$$

where $f_j(x, y)$ is the j th subband image formed by convolution with the filter g_j .

Next, we show that image functions f and f' sharing identical moments up to \hat{k} for a set of filters $g = \{g_1, \dots, g_m\}$ also have identical multivariate moments up to order \hat{k} . We assume that the frequencies of f and f' are chosen such that $G_{\alpha}(u_i, v_i) = G_{\alpha}(u'_i, v'_i)$, where G_{α} is the Fourier transform of filter $g_{\alpha} \in g$. Therefore, the subband images

$$f_j(x, y) = \sum_{i=1}^n w_i \cos(u_i x + v_i y)$$

and

$$f'_j(x, y) = \sum_{i=1}^n w_i \cos(u'_i x + v'_i y)$$

have the same weights w . Although the frequencies differ in f and f' , they have been defined such that for all $c \in A(f) \ominus A(f')$, $\hat{k} < \sum_i^n |c_i|$. This is the constraint that ensures the first k moments of the marginals are identical.

From (5), each term in the product of (12) can be expressed as

$$f(x, y)^a = \sum_{\kappa \in K} \mathcal{F}(\kappa, w) \sum_{r \in R_{\kappa}} \mathcal{C}(r, \kappa) \prod_{i=1}^n \cos(r_i z_i), \quad (13)$$

where $0 \leq r_i \leq a$. The coefficient functions \mathcal{F} and \mathcal{C} in the summation are only dependent on w and a . Because the subband images f_j and f'_j share the same weights, the coefficient functions \mathcal{F} and \mathcal{C} are identical for $f_j(x, y)^a$ and $f'_j(x, y)^a$. Using this property, the coefficients are ignored and (13) is redefined as

$$f(x, y)^a = \sum_{K_a} \left(\prod_{i=1}^n \cos(r_i z_i) \right).$$

Actually, the r_i vary in each term of the summation, however, we are only concerned with the property $0 \leq r_i \leq a$, therefore, we leave this slight error in the notation. Substituting this expression into (12), we have

$$M_{k_1 \dots k_m} = \frac{1}{|\Omega|} \int_{\Omega} \prod_{j=1}^m \left(\sum_{K_{k_j}} \left(\prod_{i=1}^n \cos(r_i z_i) \right) \right) d\Omega,$$

with $0 \leq r_i \leq k_j$. By summing over yet a different set, this expression can be written as

$$M_{k_1 \dots k_m} = \frac{1}{|\Omega|} \int_{\Omega} \sum_K \left(\prod_{j=1}^m \prod_{i=1}^n \cos(r_{j,i} z_i) \right) d\Omega,$$

with $\forall i, 0 \leq r_{j,i} \leq k_j$. The double subscript is used to denote the dependence of r on k_j . Using the same trigonometric substitution from Section 3, each term in the summation is again an integral of the form

$$\int_{\Omega} \cos(r_1 z_1 \pm r_2 z_2 \pm \dots \pm r_n z_n) d\Omega, \quad (14)$$

except that $r_i = \sum_{j=1}^m r_{j,i}$. Just as in the case of the single variate moments, the multivariate moments of f and f' differ only according to the solution set of the Diophantine equations defined by $A(f)$ and $A(f')$.

Because $\forall i, 0 \leq r_{j,i} \leq k_j$, the r_i in (14) are constrained by $r_i \leq \sum_{j=1}^m k_j = k$. The implication is that if f and f' share the same first \hat{k} single variate (marginal statistics) moments, then they will also share the same multivariate moments such that $\sum_{j=1}^m k_j \leq \hat{k}$. Thus, Theorem 3 can be broadened to include multivariate moments.

Theorem 4. *Let f and f' be functions in $F(w)$. Define \hat{k} to be $< \sum_i^n |c_i|$ for all $c \in A(f) \ominus A(f')$. Let $\{g_1, \dots, g_m\}$ be a set of symmetric or antisymmetric functions with the Fourier transforms $\{G_1, \dots, G_m\}$. If $\forall i \forall j, G_j(u_i, v_i) = G_j(u'_i, v'_i)$, then $\forall k: k \leq \hat{k}, M_{k_1 \dots k_m}(f) = M_{k_1 \dots k_m}(f')$, where $k \geq \sum_{j=1}^m k_j$.*

Proof. The proof is omitted because it is nearly identical to that of Theorem 2. \square

Like the single variate moments, matching the multivariate moments requires choosing the frequencies of f and f' such that:

1. The solutions to the Diophantine equations (10) of magnitude $\leq k$ are identical.
2. For each frequency pair (u_i, v_i) and (u'_i, v'_i) , we have $\forall j, G_j(u_i, v_i) = G_j(u'_i, v'_i)$, where G_j is the Fourier transform of filter g_j .

There is one subtle difference between these requirements and those stated for the marginal statistics. These require $G_j(u_i, v_i) = G_j(u'_i, v'_i)$ for a fixed ordering of the frequencies over all of the filters, whereas, for the marginal statistics, the ordering could be changed for each filter.

For a set of rotationally symmetric filters or derivative filters, the second condition is satisfied in the same manner as described in Sections 5.1 and 5.2. Therefore, the example patterns shown in Fig. 1 and Fig. 2 have matching multivariate moments up to order k , where k is the number of matching single variate moments. Note that the number of multivariate moments grows combinatorial. For m filters, the number of multivariate moments up to order k is $\sum_{i=0}^k \binom{m+i-1}{i}$. Thus, even though the joint distributions reside in an m dimensional space, the number of constraints imposed by the matching multivariate moments is large.

Finding image patterns with matching joint statistics for oriented filters, such as Gabor functions, is more difficult. The patterns derived for matching the marginal statistics of Gabor filters (Fig. 3) do not satisfy $G_j(u_i, v_i) = G_j(u'_i, v'_i)$ for a fixed ordering of the frequencies. In fact, such an ordering does not exist. One need only consider three Gabor functions to show that $G_1(u, v) = G_1(u', v')$, $G_2(u, v) = G_2(u', v')$, and $G_3(u, v) = G_3(u', v')$ implies $(u, v) = (u', v')$. In order to find patterns with matching joint statistics, we take a slightly different route.

Again, we define two sets of frequencies where we hold one set constant and move the second set. However, this time we do not rotate the frequencies, but slightly perturb them, which ensures $G_j(u_i, v_i) \approx G_j(u'_i, v'_i)$. The frequencies are also chosen such that they have the same solution set, $A(f)$ and $A(f')$, for some k . An example is shown in Fig. 4. By using just four frequencies, we can ensure that the solutions of $A(f)$ and $A(f')$ are identical up to a magnitude of $k = 43$.

Because joint distributions are of high dimension, techniques for reducing the complexity must be employed in order to compare the two patterns in Fig. 4. For example, very coarse histograms could be constructed. Other techniques which have been used for texture classification include linear subspace fitting [30], adaptive binning [32], and vector quantization [29], [31]. Because of the excellent classification results reported by Varma and Zisserman, we use their vector quantization method [31]. Each image is filtered with a bank of even/odd Gabor filters at four scales and six orientations for a total of 48 filters. Due to the sparseness of the frequencies, many of the filters have a zero response. Therefore, we only use the 12 filters with the largest average response. Next, k-means clustering is performed on the 12-dimensional space of the joint filter responses of each pattern. The number of cluster centers is 10 for each pattern, resulting in a total of 20 clusters. After




Fig. 4. Patterns with matching joint statistics. Although similar in appearance, these images are visually distinct, yet they have nearly the same joint statistics. Below each image is a plot of the frequencies used to construct the patterns. The histograms show a nearly identical distribution of cluster centers computed from the joint distribution of a set of Gabor filter responses.

clustering, each pixel in each image is mapped to the closest cluster center by Euclidean distance. Finally, a histogram of cluster center responses is constructed for each image. These histograms are representations of the joint distributions. Fig. 4 shows the histograms for each image and they are very similar. The L_1 difference between the two histograms is approximately 4 percent.

7 PERIODIC PATTERNS

Most of the patterns presented in the previous sections vary over the extent of the image and are not periodic. Because these patterns are deterministic and lack periodicity, they are not necessarily spatially homogeneous. As a result, global statistics must be used in order to match the subband histograms. However, periodic images can be created by tiling patterns with matching statistics. In this section, we show how to create such images.

Given a pair of patterns, f and f' , with matching statistics, a pair of periodic images with matching statistics can be constructed by tiling f and f' each $s \times s$ times. This is equivalent to scaling the integer frequencies of f and f' by the integer s . Because the Diophantine solution set is preserved when the frequencies are scaled by an integer, the statistics of the tiled versions of f and f' also match. The problem is that the patterns described in the previous

sections have a resolution of 256×256 pixels, thus tiling will create large images of size $(256s) \times (256s)$.

In order to construct periodic images of a reasonable size, say 256×256 pixels, we must find examples of patterns with matching subband statistics that are considerably smaller. This is not a simple task because the size of the tile limits the range of frequencies that can be used. As the size of the tile shrinks, there are fewer frequencies to choose among and less room to move them. In fact, as the size of the tile gets very small, just the first few moments of a small filter set is enough to uniquely determine the functions in $F(w)$. However, for reasonably sized tiles, such as 32×32 pixels, we can construct repeated patterns with approximately matching subband statistics, both marginal and joint.

One possibility for constructing a repeated pattern is to use the same method as used to construct the patterns in Fig. 4, but with much smaller frequencies. The problem is that these patterns are created by slightly perturbing some frequencies. When the frequencies are scaled by the tiling operation, these perturbations correspond to large changes. Therefore, rather than moving the frequencies, another method for constructing visually distinct periodic patterns with matching subband statistics is needed.

Again, we use two sets of frequencies, $u_a = \{(u_1, v_1), \dots, (u_n, v_n)\}$ and $u_b = \{(u_{n+1}, v_{n+1}), \dots, (u_m, v_m)\}$, and define f_a and f_b in terms of u_a and u_b , respectively. Except, rather than applying a rotation, a phase shift corresponding to a spatial translation (δ_x, δ_y) is applied to one set of frequencies. Thus, the two patterns are defined as $f(x, y) = f_a(x, y) + f_b(x, y)$ and $f'(x, y) = f_a(x, y) + f_b(x + \delta_x, y + \delta_y)$.

Because the exact same frequencies are used for both f and f' , any arbitrary filter set can be applied and the first k moments (marginal and joint) will still be the same. We only show this for the marginal moments, but the derivation for the joint moments is analogous. Let g be any arbitrary filter with magnitude w_i and phase ϕ_i at frequency (u_i, v_i) . The convolution of f and f' with g is

$$g * f = \sum_{i=1}^n w_i \cos(u_i x + v_i y + \phi_i) + \sum_{i=n+1}^m w_i \cos(u_i x + v_i y + \phi_i)$$

and

$$g * f' = \sum_{i=1}^n w_i \cos(u_i x + v_i y + \phi_i) + \sum_{i=n+1}^m w_i \cos(u_i(x + \delta_x) + v_i(y + \delta_y) + \phi_i).$$

Because $g * f$ and $g * f'$ have the same weights w_i , to show that they have the same moments, we only need to consider arbitrary terms t and t' in (7). To simplify the notation, define $r = [r_1, \dots, r_m]^T$, $u = [u_1, \dots, u_m]^T$ and $v = [v_1, \dots, v_m]^T$. Also define $r' = [r_{n+1}, \dots, r_m]^T$, $u' = [u_{n+1}, \dots, u_m]^T$, and $v' = [v_{n+1}, \dots, v_m]^T$. Now,

$$t = \int_{\Omega} \cos(r \cdot (ux + vy + \phi)) d\Omega \quad (15)$$

and

$$t' = \int_{\Omega} \cos(r \cdot (ux + vy + \phi) + r' \cdot (u'\delta_x + v'\delta_y)) d\Omega. \quad (16)$$

Provided the frequencies are chosen such that the small magnitude Diophantine solutions cannot be formed from frequencies contained in both u_a and u_b , the small moments will match. This is stated in the following theorem:

Theorem 5. *Let $A(f)$ be the Diophantine solution set. If*

$$\forall c = (c_1, \dots, c_{n+m}) \in A(f) : \sum |c_i| \leq k \Rightarrow (c_1, \dots, c_n) = (0, \dots, 0) \vee (c_{n+1}, \dots, c_m) = (0, \dots, 0),$$

then $M_k(f) = M_k(f')$.

Proof. It is sufficient to show that the arbitrary terms t and t' in the k th moment are equal. For the k th moment, we have the constraint that $\sum |r_i| \leq k$ and, therefore, we only need to consider the following:

Case 1: $(r_1, \dots, r_n) \notin A(f)$. Both t and t' can be rewritten as $t = \int_{\Omega} \cos(ax + by + c) d\Omega$ and $t' = \int_{\Omega} \cos(ax + by + c') d\Omega$, where a and b are integers with at least one nonzero and the integrals evaluate to $t = t' = 0$.

Case 2: $(r_1, \dots, r_n) \in A(f)$ and $(r_{n+1}, \dots, r_m) = (0, \dots, 0)$. $(r_{n+1}, \dots, r_m) = (0, \dots, 0)$ implies $r' = 0$ in (16), which implies $t = t'$.

Case 3: $(r_1, \dots, r_n) \in A(f)$ and $(r_1, \dots, r_n) = (0, \dots, 0)$. This implies that $(0, \dots, 0, r_{n+1}, \dots, r_m) \in A(f)$, which implies that $r' \cdot u = 0$ and $r' \cdot v = 0$, which implies that $t = t'$. \square

Theorem 5 states sufficient conditions for ensuring that f and f' have the same marginal and joint moments up to order k for any arbitrary filter set. Thus, it is a stronger result than the previous theorems. However, because the frequencies are not moved, the visual "distinctness" of f and f' is considerably less. On the other hand, not moving the frequencies permits the use of smaller frequencies, thus periodic patterns can be formed.

Fig. 5 shows an example of a pair of periodic images whose tiles f and f' are chosen in this manner for $k = 5$. Each image is constructed from 8×8 tiles of size 32×32 pixels. The locations of the two sets of frequencies that are used to construct the tiles is shown in Fig. 5a. The two sets of frequencies form squares with one square twice the size of the other. Increasing the difference in scale between the squares increases k , but the two patterns quickly lose their distinctness.

Because the patterns in Fig. 5b have approximately the same subband statistics (up to $k = 5$), we should expect texture synthesis algorithms based on marginal and joint subband statistics to produce similar results for these two patterns. Fig. 5c shows the result of applying the texture synthesis method of Portilla and Simoncelli [5]. This method measures a set of marginal and joint subband statistics from a Steerable pyramid and then generates a random image with approximately matching statistics. As expected, the randomly generated images have a similar appearance despite the fact that the input images are visually distinct. Because the frequencies are confined to two regions of the frequency plane, increasing the number of orientations and scales does not help.

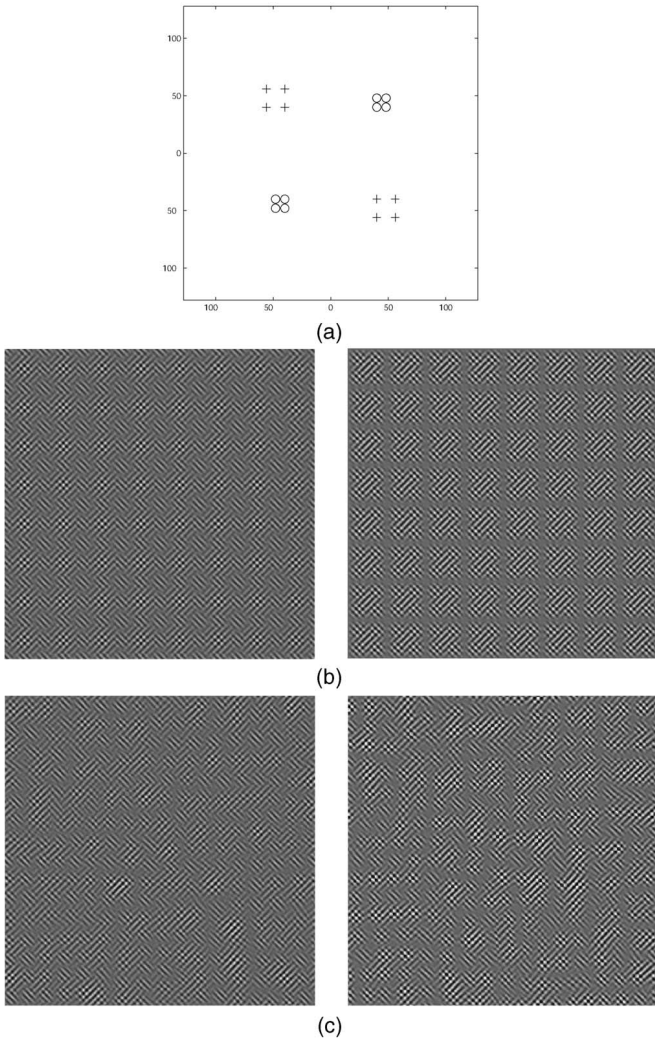


Fig. 5. Periodic patterns with matching subband statistics. (a) The patterns are constructed from the same two sets of frequencies denoted by +s and os. For one of the patterns, a phase shift corresponding to a spatial translation is applied to one set of frequencies. (b) Periodic images are formed by tiling the patterns. These images have matching marginal and joint moments up to order 5 for any arbitrary filter set. (c) The result from applying a texture synthesis algorithm based on marginal and joint subband statistics. While the patterns in (b) are visually distinct, the images generated by the synthesis algorithm are not distinct.

8 DISCUSSION

To summarize, we have shown the existence of distinct visual patterns that possess the same subband statistics. In particular, we have shown patterns that can not be distinguished from the marginal statistics of Gaussian, Laplacian, and derivative filters, even if these filters are applied at an infinite number of scales. We have also shown how to construct patterns that can not be distinguished from the marginal statistics of Gabor filters at a fixed set of orientations. Furthermore, distinct patterns have been derived that can not easily be differentiated from the joint statistics of these filters. And, finally, we showed the existence of periodic patterns that map to the same moments up to order 5, both marginal and joint, for any arbitrary filter set.

As previously mentioned, Zhu et al. have pointed out that with a large enough set of filters, the marginal

distributions alone can represent any pattern [22]. However, this relies upon exact knowledge of the distribution and, thus, all of the moments. It seems that if we allow for some degree of error in the moments, perhaps due to quantization or noise, then it becomes difficult to differentiate all distinct patterns from marginal distributions. In work on texture synthesis, Heeger and Bergen as well as Portilla and Simoncelli have noted that many complex textures are not easily modeled by marginal distributions, at least using a particular filter set [5], [26]. The results in this paper seem to confirm this and, in the case of the patterns in Fig. 5, extend this to the case of any filter set.

However, it should be noted that the examples in Fig. 1, Fig. 2, Fig. 3, and Fig. 4 are not homogenous textures, but, rather, visual patterns. Thus, they are examples of the failure of these statistics to capture visual appearance, rather than model stochastic textures. The periodic texture in Fig. 5 demonstrates the failure of marginal and joint statistics to model all deterministic periodic patterns. Of course, there is the possibility that these patterns could be differentiated if the sixth and higher order moments could be reliably estimated. In addition, note that the histograms are global statistics and, thus, leave the possibility that these patterns could be classified from local statistics.

At this point, it is worthwhile to speculate as to how these patterns might be represented statistically. Clearly, the patterns in Fig. 1 and Fig. 2 could easily be differentiated using marginal or joint statistics of oriented filters. And, as previously mentioned, if orientations with a finer tuning are used, then the patterns in Fig. 3 could be distinguished. In order to represent patterns such as those in Fig. 4, a set of highly oriented filters would be needed. For example, Gabor filters with highly anisotropic elements. Other transforms that are constructed from such elements include the Ridglet and Curvlet transforms [36], [37]. However, arbitrarily adding more orientations could result in over fitting of other patterns.

The pairs of patterns shown in Fig. 1, Fig. 2, Fig. 3, and Fig. 4 share a common characteristic. Locally, each pair is made up of very similar features. However, globally, they differ in the larger structures formed by interference patterns. It is these interference patterns that are difficult to detect because they only effect the very high moments, which typically cannot be reliably estimated in the presence of noise. An alternative to using highly anisotropic filters to detect these interference patterns is to use a cascaded approach, where a second stage of filtering is applied after a nonlinear operation on the first stage. Such "filter-rectify-filter" approaches have been used to model human vision and are sometimes called second order texture analyzers [38] [39] [40]. A nonlinear operation, such as squaring in the spatial domain, is equivalent to convolution in the frequency domain. Thus, a second stage of oriented filters responds to the orientation of frequency offsets, which cause the interference patterns. Such statistics have been used for representing textures and are incorporated into the texture synthesis method of Portilla and Simoncelli [5]. The periodic patterns in Fig. 5 are an interesting example because even highly anisotropic filters, and some second order texture analyzers, fail to differentiate these. However,

a nonlinear operation such as contrast normalization may be able to separate the patterns.

Finally, we would like to mention that these patterns are deterministic and designed to have matching sample statistics. They are not examples drawn from a distribution with matching ensemble statistics, such as the isodipole textures that disproved the Julesz conjecture [11]. Therefore, the examples in this paper are more similar to the counterexamples of Yellott which are distinct patterns with matching sample second order statistics [13].

ACKNOWLEDGMENTS

Parts of this paper first appeared in the IEEE International Conference on Computer Vision in 2003. This work was supported by US National Science Foundation ITR award number IIS-0219078. A number of helpful suggestions were given by the anonymous reviewers.

REFERENCES

- [1] M. Mandal and T. Aboulnasr, "Image Indexing Using Moments and Wavelets," *IEEE Trans. Consumer Electronics*, vol. 42, no. 3, pp. 557-565, 1996.
- [2] B. Schiele and J. Crowley, "Recognition without Correspondence Using Multidimensional Receptive Field Histograms," *Int'l J. Computer Vision*, vol. 36, no. 1, 2000.
- [3] C. Schmid, "Constructing Models for Content-Based Image Retrieval," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '01)*, 2001.
- [4] H. Schneiderman and T. Kanade, "A Statistical Method for 3D Object Detection Applied to Faces and Cars," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '00)*, 2000.
- [5] J. Portilla and E. Simoncelli, "A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients," *Int'l J. Computer Vision*, vol. 40, no. 1, pp. 49-71, 2000.
- [6] E. Simoncelli and E. Adelson, "Noise Removal via Bayesian Wavelet Coding," *Int'l Conf. Image Processing*, 1996.
- [7] J. De Bonet and P. Viola, "A Non-Parametric Multi-Scale Statistical Model for Natural Images," *Advances in Neural Information Processing*, vol. 10, 1997.
- [8] X. Liu and D. Wang, "Texture Classification Using Spectral Histograms," *IEEE Trans. Image Processing*, vol. 12, no. 6, 2003.
- [9] B. Julesz, "Visual Pattern Discrimination," *Trans. Information Theory*, no. 8, pp. 84-92, 1962.
- [10] T. Caelli and B. Julesz, "On Perceptual Analyzers Underlying Visual Texture Discrimination," *Part I: Biological Cybernetics*, vol. 28, pp. 167-175, 1978.
- [11] B. Julesz, E. Gilbert, and J. Victor, "Visual Discrimination of Textures with Identical Third-Order Statistics," *Biological Cybernetics*, vol. 31, 1978.
- [12] E. Gilbert, "Random Colorings of a Lattice on Squares in the Plane," *SIAM J. Algorithms of Discrete Methods*, vol. 1, pp. 152-159, 1980.
- [13] J. Yellott, "Implications of Triple Correlation Uniqueness for Texture Statistics and the Julesz Conjecture," *J. Optical Soc. Am. A*, vol. 10, no. 5, pp. 777-793, 1993.
- [14] O. Faugeras and W. Pratt, "Decorrelation Methods of Texture Feature Extraction," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 2, no. 4, pp. 323-332, July 1980.
- [15] M. Silverman, D. Grosz, R. DeValois, and S. Elfar, "Spatial-Frequency Organization in Primate Striate Cortex," *Proc. Nat'l Academy of Sciences of the USA*, 1989.
- [16] C. Chubb and M.S. Landy, "Orthogonal Distribution Analysis: A New Approach to the Study of Texture Perception," *Computational Models of Visual Processing*, M.S. Landy, ed., 1991.
- [17] J.R. Bergen and E.H. Adelson, "Theories of Visual Texture," *Spatial Vision*, D. Regan, ed., 1991.
- [18] T. Randen and J.H. Husoy, "Filtering for Texture Classification: A Comparative Study," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 4, Apr. 1999.
- [19] W. Ma and B. Manjunath, "Texture Features and Learning Similarity," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '96)*, 1996.
- [20] A. Bovik, M. Clark, and W. Geisler, "Multichannel Texture Analysis Using Localized Spatial Filters," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 1, Jan. 1990.
- [21] S.G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, July 1989.
- [22] S.C. Zhu, Y. Wu, and D. Mumford, "Filters, Random-Fields and Maximum-Entropy Frame: Towards a Unified Theory for Texture Modeling," *Int'l J. Computer Vision*, vol. 27, no. 2, pp. 107-126, 1998.
- [23] D.L. Ruderman, "The Statistics of Natural Images," *Network*, 1994.
- [24] C. Zetsche, "Polyspectra of Natural Images," *Natural Scene Statistics Meeting*, 1997.
- [25] J. Huang and D. Mumford, "Statistics of Natural Images and Models," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '99)*, 1999.
- [26] D. Heeger and J. Bergen, "Pyramid-Based Texture Analysis/Synthesis," *Proc. SIGGRAPH Conf. '95*, 1995.
- [27] S.C. Zhu, X. Liu, and Y. Wu, "Exploring Texture Ensembles by Efficient Markov Chain Monte Carlo," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 6, June 2000.
- [28] A. Srivastava, X. Liu, and U. Grenander, "Universal Analytical Forms for Modeling Image Probabilities," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 9, Sept. 2002.
- [29] T. Leung and J. Malik, "Representing and Recognizing the Visual Appearance of Materials Using Three-Dimensional Textons," *Int'l J. Computer Vision*, Dec. 1999.
- [30] K. Dana and S. Nayar, "Histogram Model for 3D Textures," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '98)*, 1998.
- [31] M. Varma and A. Zisserman, "Classifying Images of Materials: Achieving Viewpoint and Illumination Independence," *Proc. European Conf. Computer Visions*, 2002.
- [32] S. Konishi and A. Yuille, "Statistical Cues for Domain Specific Image Segmentation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '00)*, 2000.
- [33] E. Hadjidemetriou, M. Grossberg, and S. Nayar, "Histogram Preserving Image Transformations," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '00)*, 2000.
- [34] E. Hadjidemetriou, M. Grossberg, and S. Nayar, "Spatial Information in Multiresolution Histograms," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '01)*, 2001.
- [35] P. Meyer, *Introductory Probability and Statistical Applications*, second ed. Addison Wesley, 1970.
- [36] E. Candes and D. Donoho, "Ridgelets: A Key to Higher-Dimensional Intermittency?" *Philosophical Trans. Royal Soc. London A*, vol. 357, pp. 2495-2509, 1999.
- [37] J. Starck, E. Candes, and D. Donoho, "The Curvelet Transform for Image Denoising," *IEEE Trans. Image Processing*, vol. 11, pp. 670-684, 2001.
- [38] N. Graham, J. Beck, and A. Sutter, "Nonlinear Processes in Spatial-Frequency Channel Models of Perceived Texture Segregation," *Vision Research*, vol. 32, pp. 719-743, 1991.
- [39] A. Sutter, G. Sperling, and C. Chubb, "Measuring the Spatial Frequency Selectivity of Second-Order Texture Mechanisms," *Vision Research*, vol. 35, no. 7, 1995.
- [40] A. Schofield and M. Georgeson, "Sensitivity to Contrast Modulation: The Spatial Frequency Dependence of Second-Order Vision," *Vision Research*, vol. 43, pp. 243-259, 2003.



Joshua Gluckman received the BS degree in economics from the University of Virginia (1992), the MS degree in computer science from the College of William and Mary (1995), and the PhD degree in computer science from Columbia University (2000). Currently, he is an assistant professor of computer science at the Polytechnic University in Brooklyn, New York. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.