Multilevel isotrigon textures

Ted Maddess

ARC Centre of Excellence in Vision Science and Centre for Visual Sciences, Research School of Biological Sciences, Australian National University, Canberra ACT 0200, Australia

Yoshinori Nagai

Center for Information Science, Kokushikan University, 4-28-1 Setagaya, Setagaya-ku, Tokyo

Jonathan D. Victor

Weill Medical College of Cornell University, New York, New York 10021

Ryan R. L. Taylor

ARC Centre of Excellence in Vision Science and Centre for Visual Sciences, Research School of Biological Sciences, Australian National University, Canberra ACT 0200, Australia

Received January 3, 2006; revised August 3, 2006; accepted August 9, 2006;
posted September 18, 2006 (Doc. ID 66921); published January 10, 2007

To date a small palette of isotrigon textures have been available to study how the brain uses higher-order spatial correlation information. We introduce several hundred new isotrigon textures. Special modulation properties are illustrated that can be used to extract neural responses to higher-order spatial correlations. We also ask how many textures make an adequate training set and how representative individual examples are of their texture class. Human discrimination of 90 of these patterns was quantified. Modeling those responses shows that humanlike performance can be obtained providing a fourth-order classifier is used, although more than one mechanism is required. © 2007 Optical Society of America

OCIS codes: 330.0330, 330.4060, 330.5000, 330.5510, 330.7310.

1. INTRODUCTION

Two-point spatial correlations inform the visual system about the (phase-independent) spatial frequency content of images; however, encoding features such as contours requires mechanisms sensitive to correlations involving three or more points.[1] Such information is characterized by third- and higher-order spatial correlations. Recent innovations in the means to quantify higher-order correlations[2] reveal that in natural images spatial correlations up to the fifth order[3] are necessary to accurately represent image regions that contain junctions or intersections. Alternative forms of analysis indicate considerable redundancies in higher-order correlations.[4,5] The best independence between visual channels that can be achieved by linear filters is accomplished by mechanisms like independent component analysis, which yield sparse Gabor-like receptive fields.[6] This optimization does not, however, capture redundancies at higher orders that are typical of natural images.[4,5] Consideration of optimal schemes to achieve better redundancy reduction by targeting these higher-order redundancies yields nonlinear transformations of images that correspond to cortical contrast gain control, complex cells, and extraclassical receptive field properties such as end stopping.[7,8] These transformations correspond to mechanisms for capturing junctions in images.[9] Good evidence has also been provided indicating that saccadic eye movements are directed at these loci.[10]

Not surprisingly, then, visual evoked potential (VEP) studies have shown that the visual cortex is sensitive to higher-order spatial correlations.[11–15] Seventy percent of single units in macaque V1 have also been shown to encode such high-order statistical information, and this ability is present in all cortical layers.[1] More recently positron emission tomography (PET)[16] and functional MRI (fMRI) studies[17,18] have revealed differential activation of brain regions during discrimination of texture patterns that differ on average only in their fourth- and higher-order spatial correlations, so-called isotrigon textures.

Until recently only 16 types of these isotrigon textures have been available for studying the properties of human texture discrimination.[13,19] Some variants of the 16 texture types have expanded our knowledge of texture discrimination.[12,20–22] We have recently shown that a larger palette of isotrigon textures might be available.[23] This paper describes several hundred new isotrigon textures and examines their discriminability. A larger palette of isotrigon textures should in principle be able to better inform models of human texture processing, since a reasonable number of discriminations is necessary to distinguish models of even modest complexity.[24] Exemplary modeling results are given. A subset of 72 of these pat-

terns is also shown to have useful modulation properties for VEP or fMRI studies. A compact modulo-linear method for describing the rules for generating many of these textures is given. The utility of the minitexture spectrum for quantifying differences between these deterministic patterns is also discussed. We also quantify how many textures make a representative training set.

2. METHODS

A. Isotrigon Textures

The second-order spatial correlation function, or autocorrelation function, of an arbitrary image sample $I(x,y)$ of $m \times n$ pixels is

$$C_{2,f}(h,\nu) = \frac{1}{N}\sum_{1}^{m}\sum_{1}^{n} I(x,y)I(x+h,y+\nu)\Delta x\Delta y. \qquad (1)$$

The factor N assumes that the area of each pixel is 1 and that $N=m\times n$. $I(x,y)$ is assumed to be zero mean, which for the textures described here is true, as they are defined to have equal numbers of pixels with values −1, 0, and 1. The Fourier transform of $C_{2,f}(h,\nu)$ is the power spectrum of $I(x,y)$. The third-order correlation function is

$$C_{3,f}(h_1,\nu_1,h_2,\nu_2) = \frac{1}{N}\sum\sum I(x,y) \times I(x+h_1,y+\nu_1)I(x+h_2,y+\nu_2)\Delta x\Delta y. \qquad (2)$$

Thus $C_{3,f}(h_1,\nu_1,h_2,\nu_2)$ is a 4D function that is an accumulation of triple products of pixels—hence the word "trigons." We will henceforth refer to the second- and third-order correlation functions as the 2CF and 3CF, respectively. The 3CF contains all the information of the 2CF. Note that for Eqs. (1) and (2) $1<x+h_n<m$ and $1<y+v_n<n$ are included in the sum, and Δx and Δy could be dropped, since they are just steps on the discrete lattice.

The utility of isotrigon textures derives from the fact that the mean 3CF of collections of these textures is everywhere 0. This is also true for patterns composed of uniformly distributed noise. Hence, if one is given a training set of two types of isotrigon textures, A and B, and then is tested on previously unseen examples drawn from types A and B, the mean 3CF of the training set can teach one nothing about how A and B differ on average up to third-order, as the average 3CF of the training patterns is not different from 0 everywhere.[25] Thus, only fourth- and higher-order correlations can be used to make these discriminations. Notice that we are defining isotrigon here in the sense of correlations as defined by Eq. (2), not in the sense of histogram statistics.[21,26,27] The meaning of higher-order statistics in these two senses diverges for images with more than two levels; so we confine ourselves to a definition in terms of Eq. (2). Patterns that have zero mean, and zero mean autocorrelations (except at the origin), but which can differ at third order and above, are said to be "isodipole." An advantage of isotrigon textures is that linearly filtered versions remain isotrigon.[15,25,28] Thus one could use isotrigon patterns to construct isotrigon patterns that have more natural $1/f$-style amplitude spectra.[29,30]

Like the original isotrigon patterns, the textures described here are formed by a cellular-automaton-based recursion process. The process proceeds much like convolution with a 2D operator except that the results at each step replace values in the initial image matrix, possibly to act as inputs later in the process. The image is initially seeded with uniformly distributed noise of N levels. In all cases presented here the operator is a 3×3 pixel domain or glider. In most cases three of the nine positions within the glider are selected as an input template [e.g., Eq. (3)]. As the glider is moved over the seed noise, the three template positions select image pixels, a, b, c to be interacted, and the result of the predefined interaction, $f(a,b,c)$, is returned to an output position that is also defined by the glider. A little-used glider is the Corners glider:

$$\text{Corners} = \begin{bmatrix} a & n & b \\ n & n & n \\ c & n & f \end{bmatrix}. \qquad (3)$$

Note that the positions labeled n take no part in the calculation; only the input points, a,b,c, and the output point f participate. The type of texture generated depends on both the glider and the rule governing the interaction of the triplets of input pixels, a,b,c. The 16 isotrigon textures used to date were binary textures $\{-1,1\}$, denoting black and white pixels, and used eight gliders and two simple rules: the Even rule, $f(a,b,c)=abc$, and the Odd rule, $f(a,b,c)=-abc$. Graphical examples of the previously used gliders are given elsewhere.[13,19,23]

B. Rules and Gliders

As shown elsewhere,[23] the Even and Odd rules are just 2 of the possible 256 rules for binary textures created from triplets of inputs. These rules can be written as the discrete equation

$$f(a,b,c) = \sum_{\alpha=0}^{n}\sum_{\beta=0}^{n}\sum_{\gamma=0}^{n} \chi_{\alpha\beta\gamma}a^{\alpha}b^{\beta}c^{\gamma}, \qquad (4)$$

where summation proceeds to $n=1$, or

$$f(a,b,c) = x_{000} + x_{100}a + x_{010}b + x_{001}c + x_{110}ab + x_{011}bc + x_{101}ac + x_{111}abc, \qquad (5)$$

where a,b,c are the input triplets. Notice that the Even and Odd rules of previous studies are the special case where all the coefficients except x_{111} are zero. Also, of these 256 rules only these two produce isotrigon patterns.[23] For the case of ternary stimuli, $\{-1,0,1\}$ or $\{0,1,2\}$ denoting black, mid-gray, and white, there are 3^{27} rules.[23] Their discrete equations contain 27 terms that can be written exactly as Eq. (4) but setting the summation limits to $n=2$, giving 27 coefficients $\{x_{000},x_{100},x_{010},\ldots,x_{122},x_{222}\}$ for 27 multinomial terms in a,b,c (Appendix A, Table 3).

Table 1. Modulo 3 Discrete Equation Form of Rules for Generating Isotrigon Textures for the Three-Input One-Output Case for Three Gray Levels[a]

	Rule Group			
Name	0	1	2	3
M_0	mod 3[2(a+b+c+1)]	mod 3[2(a−b−c+2)]	mod 3[2(−a+b−c+2)]	mod 3[2(−a−b+c+2)]
M_1	mod 3[2(a+b+c+0)]	mod 3[2(a−b−c+0)]	mod 3[2(−a+b−c+0)]	mod 3[2(−a−b+c+0)]
M_2	mod 3[2(a+b+c+2)]	mod 3[2(a−b−c+1)]	mod 3[2(−a+b−c+1)]	mod 3[2(−a−b+c+1)]
I_0	mod 3[1(a+b+c+2)]	mod 3[1(a−b−c+1)]	mod 3[1(−a+b−c+1)]	mod 3[1(−a−b+c+1)]
I_1	mod 3[1(a+b+c+0)]	mod 3[1(a−b−c+0)]	mod 3[1(−a+b−c+0)]	mod 3[1(−a−b+c+0)]
I_2	mod 3[1(a+b+c+1)]	mod 3[1(a−b−c+2)]	mod 3[1(−a+b−c+2)]	mod 3[1(−a−b+c+2)]

[a]Note that in the modulo 3 sense multiplying by 2 is the same as multiplying by −1.

C. Basic Stimulus Parameters

Visual stimuli were presented at a frame rate of 101.5 Hz on a 19 in. Barco monitor (CCID 7551, Kortrjk, Belgium) controlled by an AT Vista graphics board (Truevision, Shadeland Station, Indiana). Data analysis and display routines were written in Matlab. The viewing distance was 60 cm. The achromatic (color temperature 6500 K) texture patterns were displayed at mean luminance of 42 cd m^{-2}. The CRT supplied the primary illumination in the room.

Textures were presented at the CRT center at three sizes: 8, 16, and 32 pixels square. The textures were presented embedded in a surround of uniformly distributed ternary noise that was 64×52 pixels. Part of the reason for doing this was that some of the textures look distinctly 3D when they are presented against a mean luminance grey. At the CRT center each square pixel subtended 0.56° on a side. The stimulus and the textured surround had a smooth temporal onset and offset with contrast rising and falling sigmoidally from 0 to 1 and back to 0. This temporal envelope is described in more detail elsewhere.[19] The total duration of the window was 0.3 s, and of the contrast 1 portion 0.2 s.

3. RESULTS

We have reported that a ternary variant of the method of Gilbert[31] yielded a group of six rules that generate isotrigon textures for ternary textures.[23] We refer to this as group 0. Appendix A, Table 3 gives the original discrete function form, and the leftmost column of Table 1 shows a shorter form that will be introduced below. Subsequent research has found three further groups of six isotrigon rules. Thus, while there are 2 rules for binary isotrigon textures [the Even rule, which stipulates that $d = \text{mod } 2(a+b+c)$, and the Odd rule, which stipulates that $d = \text{mod } 2(a+b+c+1)$], there are 24 rules in the ternary case. Other, often larger, sets of isotrigon rules are also discussed below. Figure 1 shows examples of the 24 texture types generated by the Box glider and the 24 isotrigon rules. For a given glider the six rules of each group can be divided into two sets of three rules, each set of three generating superficially similar textures, the M and the I subgroups (Fig. 1). The reasoning for these names will be given below once a few concepts have been developed.

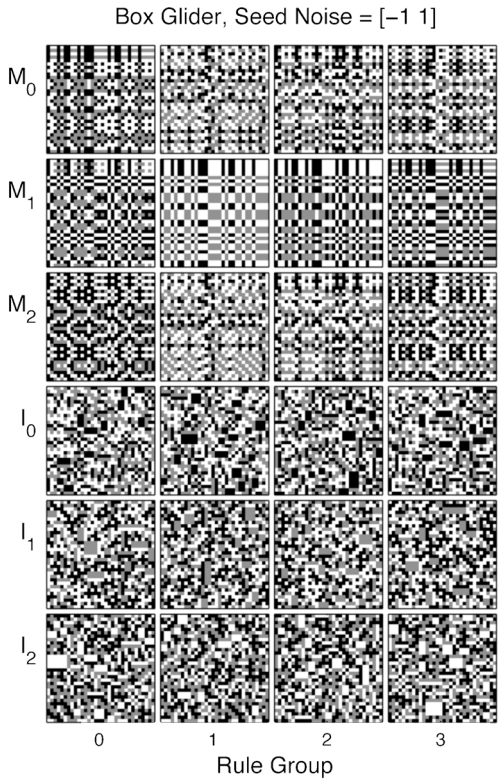

Fig. 1. Examples of textures generated with the Box glider and the four groups of six rules for combining the inputs during the recursion process. The 24 rules are given in a succinct modulo 3 form in Table 1. In each case the patterns were generated from the same matrix of evenly distributed binary noise. Although some of the patterns look similar, their minitexture content is quite disparate (Table 4, Appendix A).

We found that for each of these four groups of rules the sets of 27 discrete equation coefficients were simpler for two of the six rules in having many $x_{\alpha\beta\gamma}$ coefficients set to 0 (Appendix A, Table 3). We refer to these as the M_1 and I_1 rules (Figs. 1–3). In the case of the M rules, the simpler rule generates the most structured textures, particularly when the seed noise for the recursion process is set to −1 or 1 with equal probability instead of the usual ternary noise {−1, 0, 1} (cf. Figs. 1 and 2). As we will see, a minority of these binary seeded textures, but not the ternary seed textures, can be detected by special second-order

rules but are isotrigon by the methods described in relation to Eqs. (1) and (2), which have been described previously[23] and which were used to generate Fig. 6 below.

The 27-term discrete equation form of the texture-generating rules is cumbersome. One can also specify the rules in terms of 27 substitutions. The starting point for this method is that for a ternary texture composed of $\{-1,0,1\}$ pixels there are exactly 27 possible input triplets $\{[-1,-1,-1],[-1\ -1\ 0],\ldots[0,0,0],\ldots[1,1,1]\}$, and so one can define a 27-element substitution rule that dictates that the output, $f(a,b,c)$, for a given input triplet is one of -1, 0, or 1. Examples of substitution rules for the textures presented in Fig. 7 below are given in Table 6 of Appendix A. Those rules are more compact, as they result from gliders with only two inputs. Examples of the longer tables of substitution rules for textures created from the input gliders are given elsewhere.[23] The mathematical relationship between these substitution rules and the multinomial rule and its derivation has also been given elsewhere.[23] Nevertheless, like the multinomial discrete equation version, 27 terms are required. The six substitution rules for group 0 have been given before.[23]

We have found that a simpler form of the isotrigon rules is possible if modulo 3 arithmetic is used, yielding compact linear functions of a,b,c. The resulting 24 mod linear isotrigon discrete functions or rules are given in Table 1. The rules can be defined for the textures being composed of $\{0,1,2\}$ or $\{-1,0,1\}$ values; we have used the former here, as it is more in keeping with the modulo 3 idea. Which representation is better is moot, given that achromatic textures taking the values $0,1,2$ will present contrasts of $-1,0,1$ to the eye.

Notice that all rules are of the form $pa+qb+rc+sd=u$, where the addition is modulo 3, and all letters are integers modulo 3, with a, b, c, and d representing pixel values and p, q, r, and u defining the glider rule. Each of the rules in Table 1 has all of p, q, r, $s\in\{1,\ 2=-1\}$. Note that

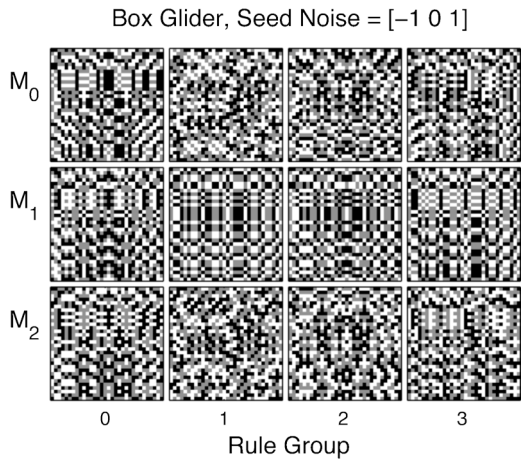

Fig. 2. M rule Box textures generated with ternary seed noise. Table 4, Appendix A, shows that only the M_0, group 1 patterns are not changed from the case of using a binary seed. As shown in Table 4, the I textures are the same for both types of initial noise, each having the correct number and characteristic set of 243 minitextures.

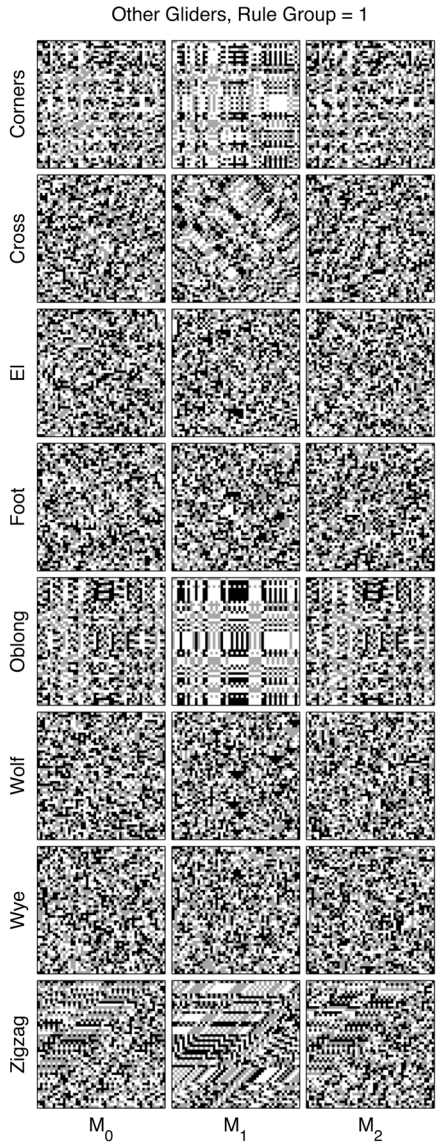

Fig. 3. Textures generated from the group 1 M rules by using eight other gilders. The Corners glider is given in Eq. (3), and the others can be found in Ref. 19.

for $d\in\{0,1,2\}$ and $s\in\{1,2\}$, $\bmod 3(sd)=k$ iff $d=\bmod 3(sk)$. Thus, any of these rules can be written as $sd=\bmod 3[u-(pa+qb+rc)]$ or, equivalently, $d=\bmod 3\{s[u-(pa+qb+rc)]\}$. The expressions appearing in Table 1 are generators of d of this sort. There are six permutations of the three-element set $\{0,1,2\}$ (modulo 3), and each of these can in turn be written as linear substitutions, e.g., $x\rightarrow x+1$, $x\rightarrow x+2$, $x\rightarrow 2x$, $x\rightarrow 2x+1$, $x\rightarrow 2x+2$, $x\rightarrow x$. Each of these substitutions acts of the above equation, mapping onto an equation of the same form.

The next few sections are dedicated to issues relating to how the textures can be identified and how many textures make a representative training set. As will be shown, textures of about 33 pixels square or greater are each good representatives of their texture group, and a training set of a dozen or so are sufficiently large to ensure that the mean 3CF (and so also the 2CF) are not significantly different from zero at all points.

A. Minitextures

We previously introduced concept of minitextures and their spectra to compare binary textures.[19] The basic idea is to partition the texture into a tessellation of smaller domains, typically 3×3 pixels each. For three-level noise, with an equal probability of each level, there will be $3^9=19{,}683$ unique minitextures. If at least one pixel is determined, this falls to $3^8=6561$ possible minitextures. Since all the textures described here are deterministic, they therefore all have 3^8 or fewer minitextures no matter how large the texture example. For different gliders different patterns of the initial noise contribute deterministically,[19] and so particular numbers of minitextures result for particular gliders. The Box glider produces particularly constrained textures with a maximum of $3^{(M+N-1)}=3^5=243$ minitextures, where M and N are the dimensions of the minitextures (for other gliders, consult the Table A1 in Ref. 19 but replace the 2 in the base of each power function with a 3). Note that Box textures produced by different rules each have different sets of 243 minitextures. One can thus form a histogram of the actual minitextures found in a texture. The histogram is formed with 19,683 bins, one for every possible minitexture, providing a minitexture spectrum that identifies the texture type[19] (Figs. 4D and 4E). Table 4, Appendix A, shows the number of minitextures that are shared between the texture types shown in Fig. 1. Although some of the textures are superficially similar, it is worth noting that none of the M textures share any minitextures with any other M texture within a rule group, and similarly for the I rules. Also, only limited overlap occurs between minitextures present in textures from different groups. So, for example, the M_0 and M_2 textures from groups 0 and 1 share no minitextures; groups 2 and 3 share only six (nine if the seed noise is $\{-1,0,1\}$).

The textures of Fig. 1 were generated by using binary seed noise because we had noticed that when we did this some of the patterns became simpler in appearance. We checked the average 2CF and 3CF functions of all these patterns (and all the texture types described in this paper) by our previous method[23] and used Subsection 3.C on convergence functions (Fig. 6 below) to confirm that the simpler versions are also isotrigon (but see Subsection 3.F on modeling). Only some of the M_0, M_1, and M_2 rules become simpler (in having fewer minitextures) when they are initialized with binary noise as quantified in Table 4, Appendix A. Notice that the number of unique minitextures falls as low as 111, again no matter how large a texture example is generated. Figure 2 shows examples of the M textures generated by using three-level seed noise. Figure 3 shows textures generated for the other gliders for rules M_0 to M_2.

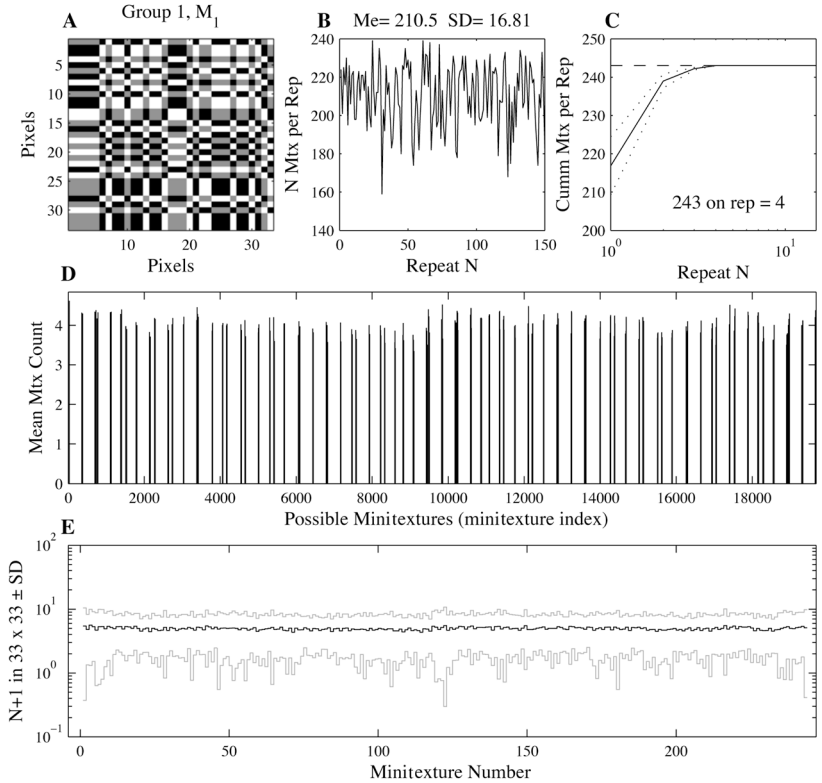

Fig. 4. Minitexture spectra example. All the figures are for the Box M_1 textures of the type shown in A, which is 33×33 pixels in size. All texture examples were generated from ternary noise. B, Number of minitextures (mtx) found in each of 150 examples, each 33×33 pixels; the mean±SD is 210.5±16.81. Note that the maximum number is 243 minitextures. C, Mean±SD cumulative number of minitextures found; all possible minitextures are observed by an average of four examples (mean, solid curve; ±1 SD, dotted curves). To compute the SD the 150 examples were divided into 10 sets of 15, and the cumulative function was computed 10 times. The horizontal dashed line indicates 243, the maximum number of observable minitextures in any Box texture. D, Raw minitexture spectrum, showing all $3^9=19{,}683$ possible bins, 243 of which are filled. The ordinate shows the mean minitexture count. E, 243 nonzero bins of D concatenated (black curve) displayed with ±1 SD. Overall each example is highly representative of the parent texture type from which it is drawn. Minitexture spectra for all glider types were similarly flat. Nonisotrigon patterns can have very uneven spectra.

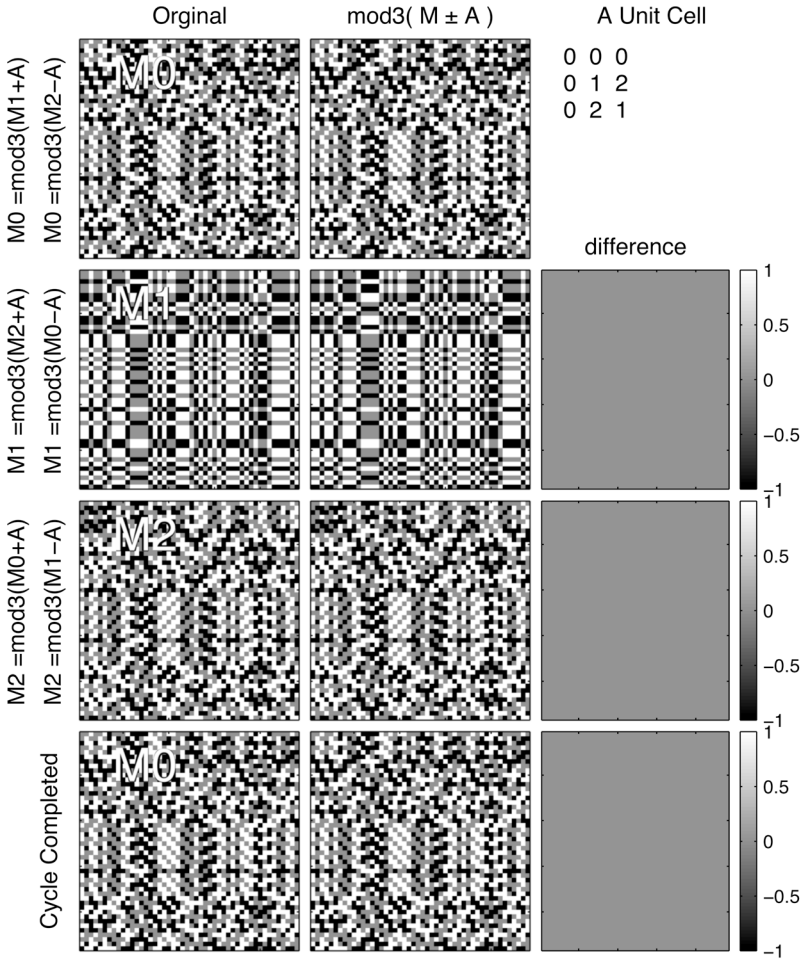

Fig. 5. Illustration of the special matrix modulation properties of textures created from M rules. For all four rule groups there is an increment matrix, A, which is a regular tessellation of a unit cell (Table 2), the unit cell for group 0 rules being shown at top right. Left column, titled Original, three textures created from one (ternary) seed noise matrix by applying rules M_0, M_1, M_2, and the Box glider. Middle column, titled mod 3(M±A), textures obtained by incrementing (or decrementing) by A as shown on the ordinate labels. Right column, (with the unit cell at the top) shows that there is no difference between the ordinals and the increment derived textures. It is worth noting that, in modulating, the patterns commonly undergo local orientation changes.

B. Modulation Properties

An interesting feature of these textures was that a simple operation can convert some texture classes to others. The simplest case is for the textures generated by the I rules. The operation is simply to add 1 to all the pixels and then take modulo 3. Then I_0 is converted to I_1, then to I_2, and then back to I_0. Thus, the name I implies "increment" textures, and this property also explains our presentation order, where I_1 has the simple discrete function (Table 3, Appendix A) and the other two are obtained by increment. The most obvious feature of the I Box textures is their rectangular domains, which are all of one intensity level. If black, gray, and white correspond to values of 0, 1, 2, then it is easy to see that black rectangles transform to gray and so on. While the transformation of the rectangles is simple, recall that the various I textures share few minitextures, so the remaining parts of the texture are quite dissimilar.

It is important to note that if the same seed noise is used to generate, say, an I_0 Box texture and an I_1 Box texture and then I_0 is transformed to $I_1' = \mathrm{mod}\,3(I_0+1)$, then $I_1 \neq I_1'$; however, the minitexture spectra of I_1 and I_1' are identical, indicating that they are of the same texture type. The fact that the same seed noise does not generate increment pairs is evident from Fig. 1, where all the examples were generated from the same matrix of seed noise. The M textures can also be switched between groups by incrementing and taking modulo 3, but they have another increment property that guarantees that these special increment pairs are equal to what would be obtained by generating each from the same seed noise matrix. It is worth noting, given the deterministic nature of these textures, that if a given seed noise is used to generate I_0, and then that seed noise is incremented mod 3 and used to generate an II texture, that this texture will be identical to the texture I_1 derived by incrementing I_0.

In this, possibly more useful, special form of modulation a regular increment matrix, A, consisting of $\{0,1,2\}$ elements is added, and modulo 3 taken, to transform M_0 to M_1, M_1 to M_2, and M_2 to M_0. This process is illustrated in Fig. 5. The name of these rules therefore derives from this matrix incrementing. The increment matrix differs for the four groups of rules. They are all regular matrices, however, and so each can be described by its unit cell

(Table 2). In fact, the original binary Box Even and Box Odd patterns were related by an increment matrix like that for group 0 here, but where the unit cell contains a 1 instead of a 2 and the operation is modulo 2. This property was exploited usefully for generating stimuli for evoked potential studies of texture discrimination[11,13,14] in that a component derived from modulation of the higher-order statistics of the textures can be measured. We expect that the modulation properties described here will prove to be similarly useful, possibly more so given the diversity of modulation possibilities.

Simple increment matrices also exist for the Oblong and Corners textures. This occurs because these textures are just tessellations (or multiplexes) of Box textures. The Oblong glider is obtained from the Corners [Eq. (3)] by swapping the second and third rows. For Oblong textures every alternate column is drawn from one of two independently generated Box patterns; the independent Box patterns are therefore interdigitated columnwise, ABAB, with A indicating a column from Box texture A, and B indicating a column from Box texture B. Corners textures are composed of four independent Box textures

$$\begin{matrix} A & B \\ C & D \end{matrix}$$

as if two Oblong textures were interdigitated rowwise. This can be understood by inspection of the gliders themselves and consideration of the recursion process that generates these textures.

Textures formed by the M rules using gliders other than the Box, Oblong, and Corners gliders (Fig. 3) do not have increment matrices that can be described by a tessellated unit cell. One can, however, generate an increment matrix for these textures by computing two textures, M_i and M_{i+1}, from the same seed noise and then subtracting them. This matrix will serve to transform M_0 to M_1, M_1 to M_2, and M_2 back to M_0. The existence of increment matrices to cycle among the rules can also be understood from considering the modulo 3 linear form of the rules, $pa+qb+rc+sd=u$. This follows from observing that the above rules are linear (even if they were not related to one another by a linear substitution; see above). Fix p, q, r, and s. A solution for $u=0$ can be added to any solution for $u=1$ to get another solution for $u=1$. The increment matrices are just solutions for $u=1$ or $u=2$ obtained with rows and columns of the seed matrix initialized to 0.

Table 2. Increment Matrix Unit Cells

$$\text{Group}_0=\begin{bmatrix}0 & 0\\ 0 & 2\end{bmatrix}$$

$$\text{Group}_1=\begin{bmatrix}0 & 0 & 0\\ 0 & 1 & 2\\ 0 & 2 & 1\end{bmatrix}$$

$$\text{Group}_2=\begin{bmatrix}0 & 0\\ 0 & 1\\ 0 & 2\end{bmatrix}$$

$$\text{Group}_3=\begin{bmatrix}0 & 0 & 0\\ 0 & 1 & 2\end{bmatrix}$$

C. Convergence Properties

The utility of isotrigon textures derives from the fact that an ensemble average has a 3CF that everywhere approaches 0 (and so too for the 2CF, except at the origin by definition). Questions therefore arise as to how many examples are needed to reach an average 3CF that does not differ from zero and how representative of a given texture class a particular example is. We have already had a taste of this in Fig. 4. Figure 4B shows the number of minitextures seen in 150 examples, each 33×33 pixels. The mean$\pm$SD is 210 ± 16.8 minitextures per example. Figure 4C shows the mean$\pm$SD (dotted curves) cumulative number of each minitexture. On average the total possible number of minitextures that can be seen in these textures is reached by four examples. This arises because the chance of encountering a given minitexture is quite uniform, as shown by the very flat minitexture frequency spectrum (Figs. 4D and 4E). These results show that not only is every possible minitexture seen in almost every 33 square pixel example, but that an average of just over 4 of every possible minitextures is seen in every single texture example (Fig. 4D). This may seem at odds with the mean occurrence rate's being about 210 per example, but this means that in some cases particular minitextures tend not to occur, leading to high variance for some minitextures (Fig. 4E). The spectrum is calculated only for contiguous minitextures; if the overlapping case is considered, about nine times more minitextures are seen in every example. Contiguous minitextures are more obviously instantaneously observable, so we consider only those minitextures here. Clearly, even for modest-sized textures, every example is a good representative of its set.

We examined empirically how many textures are needed to obtain a mean 2CF or 3CF that is not different from 0. To do this we created 200 examples of 33×33 pixel textures. This was done for 14 texture types, including some not discussed so far. We then used a method given in more detail elsewhere to examine the convergence.[23] The 2CF and 3CF of each texture were computed as previously described. The sets of 200 CFs were divided into 8 subsets of 25. For each set of 25 CFs we computed the pointwise mean and standard deviation for each of the first 3, 4, 5, 6…25 CFs. The resulting mean and SD arrays were then divided pointwise to yield functions with the same forms as the CFs but where the coefficients were t statistics (z scores if N were infinity). The maximum t statistic in each CF was then found. This was repeated for the eight blocks of examples so that standard deviations in the maximum t values could be estimated. Figure 6 shows how the maximum t values declined with the number of textures considered. We refer to these plots as "convergence" functions. When every t value was less than 2, the mean CF was taken to be not significantly different from 0. For all texture types examined, this level was reached when at least 12 CFs were av-

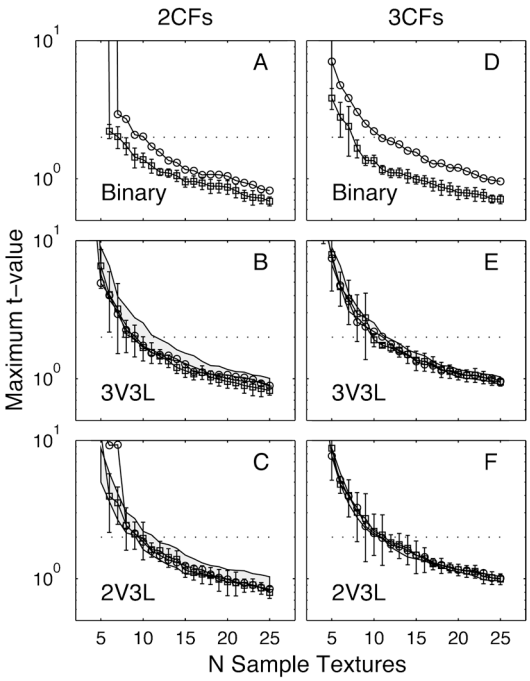

Fig. 6. Convergence functions exploring how many examples are necessary to obtain mean 2CFs and 3CFs that do not differ from zero for textures and noise. For each texture type 200 examples were created, and these were blocked into 8 sets of 25 examples. The 2CF and 3CF of each were computed, and then the mean and SD for each of the 8 sets of 25. The ratio of the resulting functions was computed pointwise, yielding a t value for each coefficient in the original CF. This was repeated for the first $3,4,5,6\ldots25$ textures in each set of 25. The maximum t statistic was then found in each case. The process was repeated for the ternary seed noise matrices used to generate each texture. A, D, Plots of the maximum t value for binary (black and white) Box textures (squares) for the 2CFs (A) and 3CFs (D), and for the random noise (circles). The textures actually proceeded more rapidly to 0 than the random textures. B, E, Similar plots for Group 0, M_1 Box textures, perhaps the most similar to the binary Box textures. These (squares) converged at a rate that was not appreciably different than for the random ternary noise (circles). The gray regions are the envelopes enclosing the convergence functions for M_1 textures for the all gliders shown in Fig. 3. C, F, Convergence function for 2V3L textures shown in Fig. 7. Error bars are SD.

eraged. Thus 12 would seem to be a reasonable minimum number of training textures.

We also examined the convergence functions for the 2CFs and 3CFs of the seed noise used to generate the textures (circles, Fig. 6). Thus we could compare how the textures converged compared with the random textures. The first texture examined was the original binary Box texture. Interestingly the textures converged quicker than the binary (black and white) noise (Figs. 6A and 6D). Aside from those binary textures, the new textures described so far in this paper can be described as three-input-variable, three-level, textures (3V3L). The square symbols in Figs. 6B and 6E correspond to group 0, M_1 Box textures, the most similar to the binary Box textures. These converged at a rate that was not appreciably different than for the random ternary noise. The gray regions in Figs. 6B and 6D are the envelope enclosing the convergence functions for all M_1 textures for the gliders shown in Fig. 3. Figures 6C and 6F are for another type of isotrigon texture, the two-input-variable, three-level textures (2V3L).

D. Other Textures

We used the previously described method for generating interesting rules to attempt to develop some rules for the 2V3L case. This generated a set of 12 rules. Empirical testing, using the t-function method described above, indicated that only four of these rules did actually produce interesting textures. In the two-input case there are nine possible input pairs from $\{[-1,-1],[-1\,0],\ldots[0,0],\ldots[1,1]\}$. The discrete equations have the form

$$f(a,b)=\sum_{\alpha=0}^{2}\sum_{\beta=0}^{2}\chi_{\alpha\beta\gamma}a^{\alpha}b^{\beta}. \tag{6}$$

The discrete equations, and their simpler modulo 3 variants, are given in Table 5, Appendix A, and the equivalent substitution rules in Table 6. These rules correspond to the modulo 3 linear rules described earlier of the form $pa+qb+rc=u$. Notionally this should provide 24 rules, where $p,q,r=\pm1$ and the parameter u is -1, 0, or 1. Half of these rules are not unique on multiplication with -1, and of those only four have $u=0$, and they are the isotrigon rules in Tables 5 and 6. Overall there are only $3^9=19{,}683$ possible rules in the 2VL3 case. We checked the t functions for all of them (2 months computer time), and only those four rules were isotrigon. There are many possible two-input gliders defined on a 3×3 domain. We examined 21 gliders. Four that provide more structured looking textures for each rule are shown in Fig. 7.

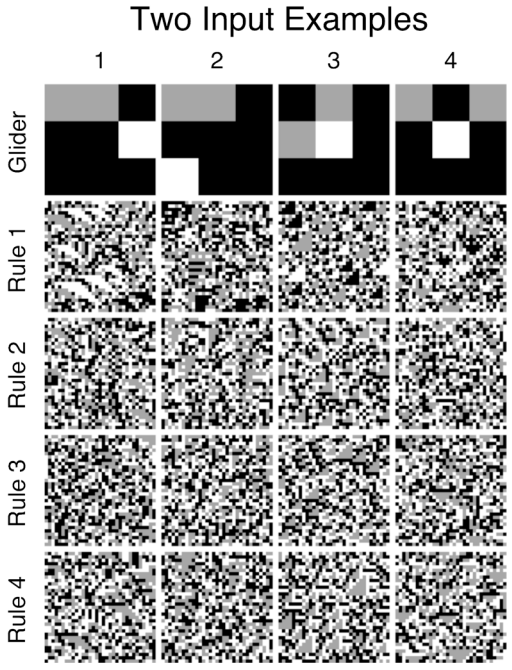

Fig. 7. Examples of 2V3L isotrigon textures. We examined textures created from the 4 rules (Table 5) and 21 gliders defined on a 3×3 pixel domain. The four gliders shown in the top row gave reasonably structured looking patterns. Note that these representations of the gliders are greatly magnified, as their block pixels are in fact the size of the pixels in the texture examples; gray pixels denote inputs (a,b), and the white pixels denote the location where the output, f, should be inserted during recursive generation of the textures. Interestingly, many 3VL3 rules generate a texture very similar to that of the three-level triangles of glider 3, rule 1, here but none of those 3V3L textures is isotrigon,[23] (Fig. 9).

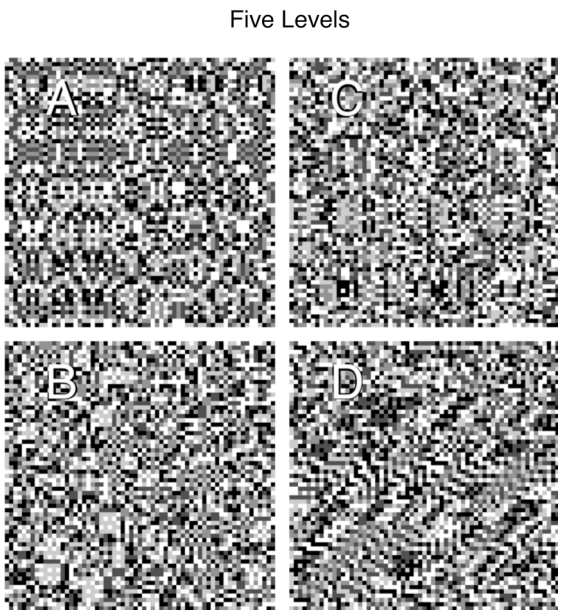

Fig. 8. Examples of three-input, five-level isotrigon textures (3V5L) showing a range of behaviors. B, D were made with the Box glider, C the Oblong glider, and D with Zigzag.

We also examined the three input cases for four and five gray levels. We have so far identified 2 groups of 24 isotrigon rules for the 3V4L case, and 24 rules for the 3V5L case. The discrete equations have the same form as Eq. (4), but the summations run to $n=3$ or $n=4$. Many of the textures look familiar (e.g., Fig. 8A), while others are less so (e.g., Fig. 8B). Interesting modulation relations exist between many of the textures, like those of the 3V3L M subgroup of textures. It is worth noting that the linear substitutions that determine the rules of Table 1 do not apply to the four- and five-level cases.

E. Texture Discrimination

A previous study of discrimination of the limited set of binary isotrigon textures showed that relatively simple models could mimic human performance.[19] That study compared the performance of formally second- and fourth-order discriminant models relative to that of human subjects. Predictably, only fourth-order models could successfully discriminate the isotrigon textures, and some models emulated human performance well. We therefore decided to examine the performance of a larger set of related models and compare that with our larger set of human texture discrimination data.

Consistent with that study, discrimination experiments were done here by using textures presented at three sizes, 8, 16, and 32 pixels square (Section 2). The subjects' task was to discriminate the textures from random ternary noise that was presented in half the trials. Note that noise patterns with uniformly distributed numbers of gray levels are isotrigon with the present textures. If, as in the present case, the contrasts levels are −1, 0, and 1, then like the textures the mean 2CFs and 3CFs of the noise patterns are 0 everywhere (except at the origin of the 2CF). Subjects fixated a small spot at the CRT center. Blocks of trials were conducted for each texture, the order of the blocks being random. Before each trial subjects were given a paper example to study, which showed 12 example textures of the texture to be tested. Each example was 33 pixels square and was presented against a medium grey background. Subjects were at liberty to study the examples for as long as they wished, which was typically less than a minute. Sample and test stimuli were seeded with binary noise, and so their minitexture composition matches that of Table 4 (Appendix A). The texture comparison experiments were done in blocks of 16 trials for each size, beginning with 32 pixels square and finishing with 8 pixels square. Subjects had a rest break of many seconds between these blocks of 16 presentations. A tone indicated an incorrect choice; so there was thus scope for learning within the 48 presentations per texture type. Two subjects (TM, YN) completed 16 trials at each scale for rules groups 0, 2, and 3. Two other subjects (IH, KO) completed the group 2 comparisons.

The resulting data are shown in Fig. 9. There seemed to be an obvious practice effect in that the results for 16 pixel square textures were sometimes better than for the 32 pixel square textures. Overall, M textures gave better performance than I textures, and the Box textures were the most discriminable.

F. Some Models

Our objective in these models is to examine the simplest possible models that can mimic human performance. A previous modeling study based on human discrimination of 16 binary isotrigon texture classes indicated that models that contained a fourth-order comparison of about four receptive field (RF) types that spanned a small range of sizes (~2) and orientations (~2) could mimic human texture discrimination of those isotrigon textures.[19] Those models all had three stages: (1) convolution with the RF operators on a small training set of textures, (2) computing variance measures in the RF outputs, and (3) computing the best linear or quadratic discriminant classifier model[32] based on the RF variance measures. Models using five types of RF variance measure, two global and three local measures, were compared. Since the inputs to the discriminant process were second order (i.e., variance), the application of linear or quadratic discriminant models yielded classification rules based on second- or fourth-order correlation information. Predictably, no second-order model performed well, since isotrigon textures by definition cannot be discriminated on average by second- or third-order statistics. Overall the basic objective was not to find physiologically realistic models *per se*, but rather to use nonparametric statistical models to place quantitative limits on the types of required computation. We did, however, demonstrate that a cascade of two second-order processes, as occurs in all these models, is physiologically plausible by using only rectification, addition, and subtraction.[19]

The power to discriminate between models that yield similar results was limited in our former study by the small number of texture comparisons. Armed with a larger data set, 30 per rule group, we decided to examine similar models again. In the first set of models the discriminant functions were optimized for each texture. Thus, although the RF inputs and variance calculations were the same, the weights for the classifier models were

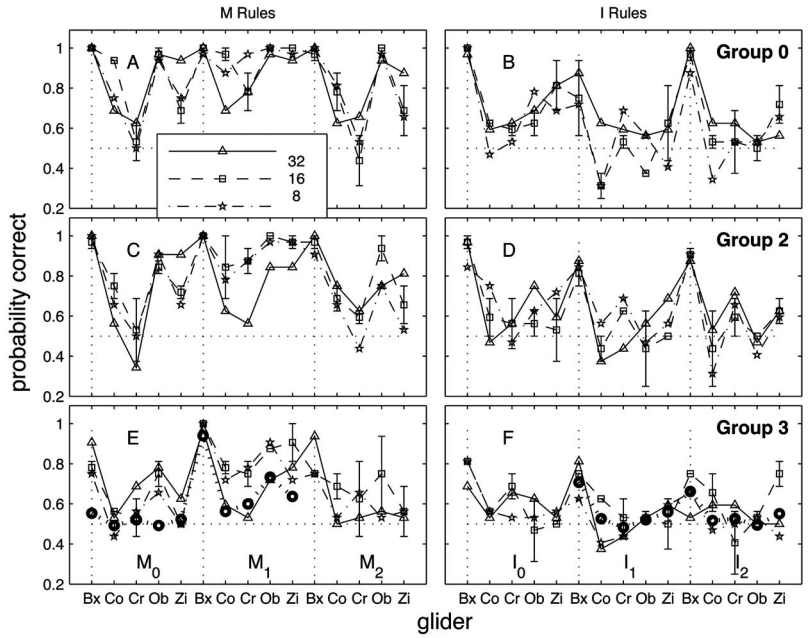

Fig. 9. Human discrimination of 90 isotrigon patterns from random ternary noise patterns. The test contained 16 trials for each of 3 texture sizes: 8×8, 16×16, and 32×32 pixels. All patterns were presented at the center of a display surrounded by random ternary noise. Test patterns were made by using three of the four possible rule groups (Table 1). The abscissa marks the gliders used: Bx, Box; Co, Corners; Cr, Cross; Ob, Oblong; Zi, Zigzag. The gliders repeat for each of the six rules per group, from left to right: M_0, M_1, M_2, I_0, I_1, I_2; vertical dotted lines demark the Box textures. The dotted horizontal line in each panel marks chance performance ($p=0.5$). Representative error bars (SE) for one each of the stimulus sizes are presented in each. Note that performance for 8×8 textures is similar to 16×16 and often better than for 32×32, apparently due to a practice effect.

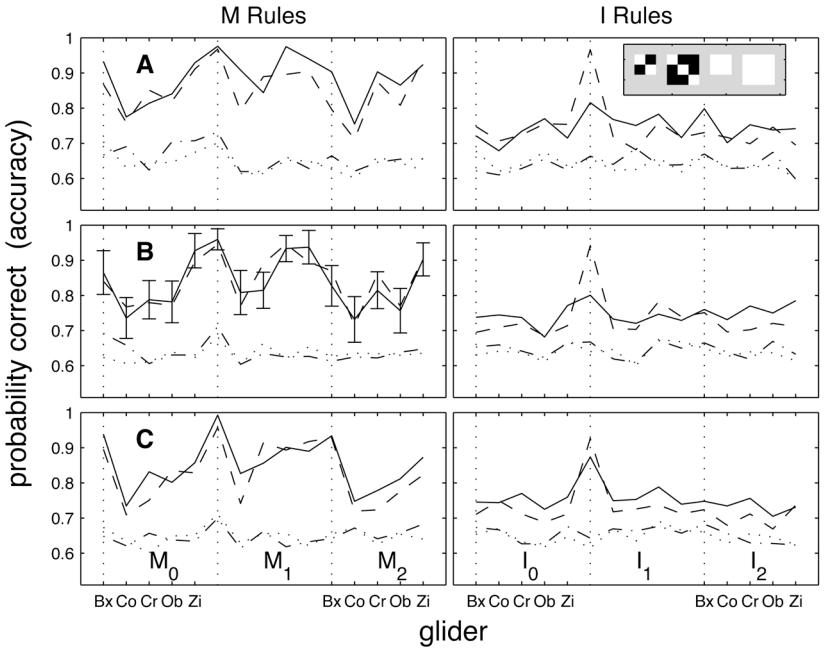

Fig. 10. Accuracy functions from discriminant models based upon the variance in the responses of four receptive field types (see text) in response to textures versus ternary noise, i.e., the same comparison as in Fig. 9. Row A, Group 0 textures, B, Group 1; C, for Group 3; otherwise the same model form was used for all textures. The four receptive field types used in the model are shown in the inset at top right. The receptive field sizes can be gaged from the fact that the pixels in the two oriented versions are the size of the pixels in the textures. Two types of variance, the standard variance (SV) and the Allan variance (AV) were used. The dotted (AV) and dashed-dot (SV) lines hovering below 70% correct are the outputs of models where the best linear discriminant classifier operated on the variance measures. Note that none of these second-order models is able to discriminate the isotrigon textures efficiently. Thus, none of these second-order models performed like humans. In models where the best quadratic classifier (based on weighted products of the variances, hence fourth order) was used, performance for M textures was reminiscent of that of the human subjects, being high for Box and Oblong, but low for Cross. For the I textures, accuracy less clearly reflected that of humans.

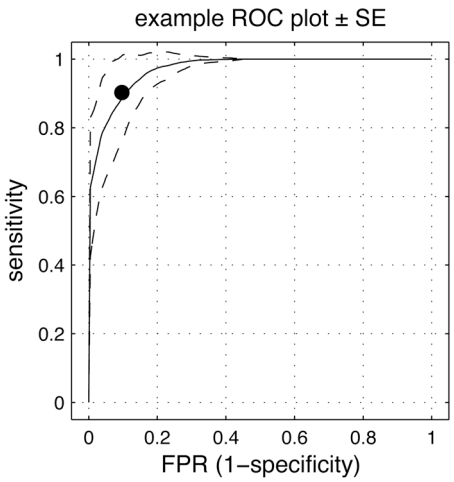

Fig. 11. Example receiver operator characteristic (ROC) plot. The plot was produced with the Box M_2 point in Fig. 10 for SV. The plot shows the mean ROC plot ± SE from bootstrap estimates ($N=50$). The solid dot is the mean accuracy derived from 50 bootstrap estimates of the accuracy. This process was repeated to derive each point in Figs. 10 and 12.

optimized for each texture type. The same RFs and variance measures were computed for ternary noise input, and these were contrasted with the similar measures from the textures in forming the discriminant classifiers.

Figure 10 shows outputs for a single model applied to group 0, 1, and 2 textures. The quantity plotted on each ordinate is derived from receiver operator characteristic (ROC) plots (Fig. 11) of classifier performance. The ordinate on these ROC plots is the sensitivity, in this case the probability of correctly identifying a texture when a texture is presented. The abscissa is the false positive rate (FPR), the rate at which a texture is reported when the stimulus is a noise pattern. The specificity is 1 − FPR, or the probability of correctly identifying noise when noise is presented. Since there is no greater penalty for either type of error, we can take the closest approach to the top left corner of the plot by the ROC curve, i.e., the point of highest simultaneous sensitivity and specificity, as a measure of classifier performance. This is sometimes referred to as the accuracy, the quantity plotted in Figs. 10 and 12. In fact we made bootstrap estimates ($N=50$) of the standard errors in the ROC plots and the accuracies.

The best performing models here were similar to the best performing model for binary textures,[19] having two unoriented and two oriented RFs, one each at sizes 2×2 pixels and 3×3 pixels (inset, top right of Fig. 10). The dashed and dotted curves hovering around probability $p=0.65$ are the output of the best performing second-order models. Overall, these second-order models performed very poorly, and nothing like humans. This agrees with our previous findings on binary isotrigon textures.[19] Two curves are presented for two linear and two quadratic discriminant models. The two quadratic curves (solid, dashed) are for global variance measures, the standard variance (SV) and the Allan variance[19] (AV); the three local variance models used previously gave only poor discrimination, <0.7, and so are not shown. For the M rules the model accuracy was qualitatively similar to that of humans (cf. Fig. 9), with good performance for Box (Bx) textures and Oblong (Ob), and poor performance for Corners (Co). We examined 25 models ($\times5$ variance measures) that encompassed receptive field side lengths of 2, 3, 4, 5, 6, 10, and 12 and that were oriented (±45°) or unoriented. Overall we could get good discrimination of any of the I textures only when the 2×2 RFs were included. In general we kept to four or fewer RF types because the number of coefficients in the quadratic discriminant case grows as the square of the number of input RFs, and we did not want to overmodel the data. For the present models discrimination was good for only the I_1 Box pattern and not I_0 or I_2.

The results from the 25 sets of models suggested that perhaps more than one mechanism had to operate to yield humanlike performance for both M and I textures. One simple way to address this idea is to examine how well a single model would perform on all textures. By this we mean not just a model of similar structure (i.e., the same RF inputs and variance model) but the same set of weights in the classifier model. To do this we selected a particular texture type to train the classifiers (linear and quadratic), then took the mean coefficients from 50 bootstrap repeats, and applied that classifier to all textures. Figure 12 shows selected results from 47 model sets ($\times5$ variance measures), which were informed by the earlier 25 model sets. For Figs. 12A–12D the training texture was the M_0 Box pattern. For Figs. 12A and 12B the four model RFs had sizes 2×2 and 6×6 and orientations ±45° (insets top right), Fig. 12A being for group 1 rules and 12(*B*) for group 2 rules. Several such models gave similarly good performance, providing the training texture was of an M type. None of these models worked well for I textures (Figs. 12A–12D). Figure 12C shows the performance for the model of Fig. 10 for group 0 textures. The interesting feature here is the results that were significantly less than 0.5 for M Cross patterns. In fact, this result was occasionally observed for some subjects, leading to a few very low mean values for these textures in Fig. 9. Adding a third scale (4×4) of oriented RFs moderated this effect (Fig. 12D). To obtain high accuracies for any I texture, we had to train on that I texture. Occasionally this gave reasonable performance when the (quadratic) classifier trained on I textures was applied to M textures (Fig. 12E, RFs sized 3×3 and 6×6 oriented at ±45°). In one case, training on the I_1 Box patterns actually gave performance ≤0.5 for all M textures (Fig. 12F); the model was similar to those of Figs. 12B and 12C in terms of the RFs applied. Overall the results suggest that three to four different models might need to operate simultaneously to capture human performance.

A surprising result from the modeling study was that some textures could be resolved with a second-order mechanism. These were some Box and Cross M textures generated from binary rather than ternary seed noise. For particular models only some of these textures produced high accuracies ($p=1.0$), while others produced very low accuracies ($p<0.2$), and Box and Cross patterns could be among either the very low or very high performers for particular models. This odd variability in the mean accuracies across textures was reproducible (SD <0.07). So, while these textures had passed the test of rapidly converging correlation functions, some RF and variance

measure combinations were able to discriminate these textures by using a linear classifier. For that reason all the modeling results reported in Figs. 10–12 were produced with textures seeded with ternary noise. None of the 25 models considered here (×5 variance measures) gave good discrimination (>0.7) of any of these textures. Since the psychophysical data of Fig. 9 had been gathered with patterns created with binary seeds, we examined discrimination of 20 group 1 textures seeded with ternary noise. The mean results from 21 subjects are shown as the thicker dotted curves and circles of Figs. 9(E) and 9(F). The results were very similar to those for the binary seeded textures.

4. DISCUSSION

We present simple modulo 3 linear versions of 24 rules (Table 1) that generate ternary isotrigon textures for three-variable inputs (3V3L). In combination with 9 glider patterns these generate 216 texture classes (examples Figs. 1–3). One can think of rules defined here that are of the form $pa+qb+rc+sd=u$ as specifying a 3D subspace of a 4D space, whose direction cosines are p, q, r, and s and whose offset from the origin is u. A further four rules for generating isotrigon textures from two inputs are also presented (2V3L). At least 4 two-input gliders give quite different and well-structured patterns yielding at least a further 16 isotrigon textures for study (Fig. 7). The possibility of four- and five-level (Fig. 8) isotrigon textures was also demonstrated. One variant of the four-level case is the complex case using the basis set 1, −1, i, $-i$. This permits textures to be made that could mimic interactions across the dimensions of a 2D perceptual space such as the isoluminant color plane.

Of the 3V3L textures, 90 were examined here for their

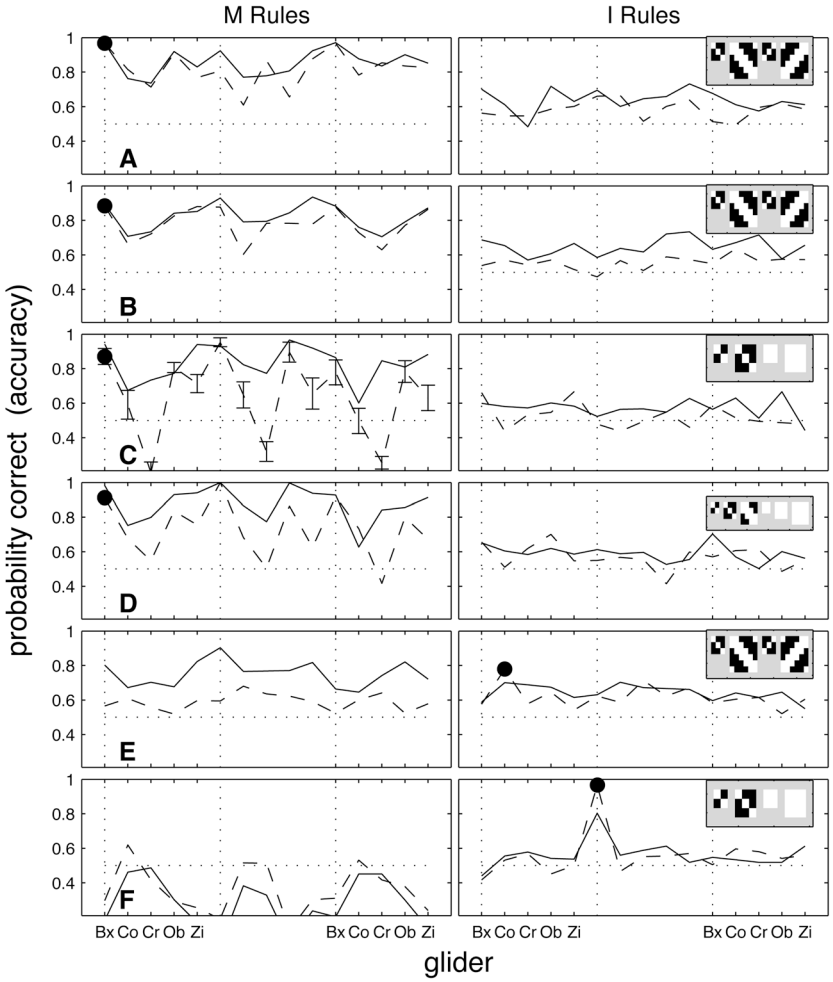

Fig. 12. Model outputs where the best single quadratic classifier model was obtained for one texture versus random comparison and was then applied to all other texture versus random comparisons. Thus the models assume one fixed mechanism for all textures rather than one mechanism with adjustable decision weights as in Fig. 10. In all cases the training texture is marked with a dot. The SV (solid lines) and AV (dashed) were applied to the RF outputs as the inputs to the quadratic classifier. Linear classifiers were also computed in all cases, but all those accuracies were below 0.65 and so are not shown. The receptive fields for each model are illustrated in the small inset panel in the top right of each I rule accuracy plot. Notice that scale is slightly smaller for the receptive fields of D, the largest of which are 4×4 pixels. A and B compare accuracies for the same four RFs (see text) applied to group 1 and 2 textures, respectively. C, results for the models of Fig. 10; the low accuracies for Cross textures were actually observed for some subjects, which contributed to the low p values for those textures in Fig. 9. D, same as C, but data from two larger receptive fields were added, smoothing out the large variations. For E and F the training was done on I textures (dot), which improved accuracy only for that I texture and occasionally (F) gave poor performance for the M textures. Overall a single, simple, texture discrimination mechanism looks improbable. Figure 10 indicates that even a single mechanism that is reprogrammable with new weights in the classifier stage is unlikely to be sufficiently complex to emulate human performance on all textures.

probability of discrimination from random ternary noise (Fig. 9). An interesting feature of those data was that accuracy improved as the subject went from the larger to the smaller textures. This undoubtedly indicated that some learning was occurring but nevertheless suggests that interactions much beyond 8 pixels are not that useful in classifying these textures, in agreement with previous studies of briefly presented (preattentive) isotrigon textures.[12,19] Certainly longer-range interactions are important for longer inspection times.[22] The rates of learning isotrigon textures vary, and it would be interesting to attempt perceptual learning experiments with those that are more difficult.

Our modeling studies confirm that fourth-order interactions between a small number of receptive fields that collectively span orientation space can mimic human performance, at least for the M textures (Fig. 10). Further modeling results suggested that a single classification scheme cannot explain all human performance and so either a more-complex model (Fig. 12) or a few simple models acting in parallel are required. One way to pursue the issue of one more-complex model or multiple independent models is to examine the covariance within the texture discrimination functions of a relatively large number of subjects and textures. This approach follows that of Peterzell and Teller.[33] We are currently investigating this approach.

We examined some blob-detector models as suggested by Tyler.[22] Here we created models based on as many as seven scales of unoriented RFs. The larger models could have been overmodeling the data. The outcome was that performance increased with increasing number of RF scales but the accuracy functions were uniformly high, i.e., not like human performance. This confirmed the similar result for binary isotrigon stimuli,[19] Fig. 11.

Another issue addressed here was the convergence of the average second- and third-order correlation functions toward being 0 everywhere (Fig. 6). This was examined in two ways. First, examining the 3×3 minitextures present in the simpler Box textures showed that 33×33 pixel examples will contain, on average, about four examples for every minitexture possible (Fig. 4(E)), and that within four such sample textures every possible minitexture will certainly be seen. It is worth noting that there was very little overlap in minitexture content, even for superficially similar textures (Table 4, Appendix A). In the second method we computed the mean and standard deviation in collections of second- and third-order correlation functions to compute t statistics for each coefficient of these functions. The maxim t value was always less than 2 for an average of 10–15 example textures. Overall this means that even fairly small ensembles of texture examples are highly representative of their class, both in terms of their minitexture content and the fact that their average correlation functions converge to 0 at least as rapidly as random ternary noise. Tyler[21] has suggested that textures should be evaluated, and even defined, based upon the ensemble statics of smaller exemplars drawn from a parent texture example. In some ways this is related to the idea of characterizing deterministic textures in terms of their minitextures, but Tyler's method can be extended to natural textures; see also Ref. 34. lt would be interesting to evaluate the Tyler model to see if it can predict human-like performance for the present textures.

Table 3. Group 0 Discrete Equation [Eq. (5)] Coefficients

$x_{\alpha\beta\gamma}$	$a^\alpha b^\beta c^\gamma$	M_0	M_1	M_2	I_0	I_1	I_2
x_{000}	1	−8	**0**	8	−8	**0**	8
x_{100}	a	4	−8	4	−4	8	−4
x_{010}	b	4	−8	4	−4	8	−4
x_{001}	c	4	−8	4	−4	8	−4
x_{110}	ab	6	**0**	−6	6	**0**	−6
x_{101}	ac	6	**0**	−6	6	**0**	−6
x_{011}	bc	6	**0**	−6	6	**0**	−6
x_{200}	a^2	12	**0**	−12	12	**0**	−12
x_{020}	b^2	12	**0**	−12	12	**0**	−12
x_{002}	c^2	12	**0**	−12	12	**0**	−12
x_{210}	a^2b	−6	12	−6	6	−12	6
x_{201}	a^2c	−6	12	−6	6	−12	6
x_{120}	ab^2	−6	12	−6	6	−12	6
x_{021}	b^2c	−6	12	−6	6	−12	6
x_{102}	ac^2	−6	12	−6	6	−12	6
x_{012}	bc^2	−6	12	−6	6	−12	6
$\boldsymbol{x_{111}}$	$\boldsymbol{abc}$	**−3**	**6**	**−3**	**3**	**−6**	**3**
x_{220}	a^2b^2	−18	**0**	18	−18	**0**	18
x_{202}	a^2c^2	−18	**0**	18	−18	**0**	18
x_{022}	b^2c^2	−18	**0**	18	−18	**0**	18
x_{211}	a^2bc	−9	**0**	9	−9	**0**	9
x_{121}	ab^2c	−9	**0**	9	−9	**0**	9
x_{112}	abc^2	−9	**0**	9	−9	**0**	9
x_{221}	a^2b^2c	9	−18	9	−9	18	−9
x_{212}	a^2bc^2	9	−18	9	−9	18	−9
x_{122}	ab^2c^2	9	−18	9	−9	18	−9
x_{222}	$a^2b^2c^2$	27	**0**	−27	27	**0**	−27

The small number of minitextures in some textures highlights a powerful way of discriminating random from deterministic textures, which is the ratio of the number of minitextures found compared with the number expected for a random texture. For the three-level case of Box textures versus random textures the maximum ratio for $M\times N$ minitextures is $3^{N+M-1}/3^{NM}=1/3^{(N-1)(M-1)}$, which becomes vanishingly small as N and M grow larger.[19] For even the most random looking textures shown here (e.g., Foot or Wye in Fig. 3) the ratio for 6×6 minitextures is $3^{2N+2M-1}/3^{NM}=6.27\times10^{-7}$. In some ways it seems extraordinary that we cannot discriminate from noise (see Fig. 9) textures that would occur naturally only at less than 1 chance in a million, even when 25 contiguous 6×6 pixel examples are found in a 33×33 pixel example.

A potentially useful property of some of the newly presented textures was their capacity to morph from one class to another by simply incrementing the value of certain pixels (text associated with Fig. 5). Victor and collaborators have exploited a similar property of the binary Box textures to evoke neuronal[1] and VEP response components that can be identified with calculation of higher-order correlation information.[14,35] Hopefully the similar properties of a much larger collection of isotrigon texture will permit a more detailed investigation of neural properties.

The fact that some of the textures created from binary seed noise could be discriminated from noise on the basis

Table 4. Number of Minitextures Shared between the Texture Types Shown in Fig. 1

		Group											
Group	Rule	0						1					
		Rule											
		M_0	M_1	M_2	I_0	I_1	I_2	M_0	M_1	M_2	I_0	I_1	I_2
0	M_0	**219**	0	0	9	0	0	0	9	0	9	0	0
	M_1	0	**111**	0	0	7	0	0	5	0	0	7	0
	M_2	0	0	**219**	0	0	9	0	9	0	0	0	9
	I_0	9	0	0	**243**	0	0	0	8	0	9	0	0
	I_1	0	7	0	0	**243**	0	0	7	0	0	9	0
	I_2	0	0	9	0	0	**243**	0	8	0	0	0	9
1	M_0	0	0	0	0	0	0	**243**	0	0	0	0	0
	M_1	9	5	9	8	7	8	0	**111**	0	8	7	8
	M_2	0	0	0	0	0	0	0	0	**243**	0	0	0
	I_0	9	0	0	9	0	0	0	8	0	**243**	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	**243**	0
	I_2	0	0	9	0	0	9	0	8	0	0	0	**243**
2	M_0	0	0	0	0	0	0	0	**0**	0	0	0	0
	M_1	15	17	15	6	7	6	0	17	0	6	7	6
	M_2	0	0	0	0	0	0	0	**0**	0	0	0	0
	I_0	9	0	0	9	0	0	0	8	0	9	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	9	0
	I_2	0	0	9	0	0	9	0	8	0	0	0	9
3	M_0	0	0	0	0	0	0	0	**0**	0	0	0	0
	M_1	15	17	15	6	7	6	0	17	0	6	7	6
	M_2	0	0	0	0	0	0	0	**0**	0	0	0	0
	I_0	9	0	0	9	0	0	0	8	0	9	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	9	0
	I_2	0	0	9	0	0	9	0	8	0	0	0	9
		Group											
		2						3					
		Rule											
		M_0	M_1	M_2	I_0	I_1	I_2	M_0	M_1	M_2	I_0	I_1	I_2
0	M_0	0	15	0	9	0	0	0	15	0	9	0	0
	M_1	0	17	0	0	7	0	0	17	0	0	7	0
	M_2	0	15	0	0	0	9	0	15	0	0	0	9
	I_0	0	6	0	9	0	0	0	6	0	9	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	9	0
	I_2	0	6	0	0	0	0	0	6	0	0	0	9
1	M_0	0	0	0	0	0	0	0	0	0	0	0	0
	M_1	**0**	17	**0**	8	7	8	**0**	17	**0**	8	7	8
	M_2	0	0	0	0	0	0	0	0	0	0	0	0
	I_0	0	6	0	9	0	0	0	6	0	9	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	9	0
	I_2	0	6	0	0	0	9	0	6	0	0	0	9
2	M_0	**216**	0	0	0	0	0	6	4	6	0	0	0
	M_1	0	**111**	0	6	7	6	4	5	4	6	7	6
	M_2	0	0	**216**	0	0	0	6	4	6	0	0	0
	I_0	0	6	0	**243**	0	0	0	6	0	9	0	0
	I_1	0	7	0	0	**243**	0	0	7	0	0	9	0
	I_2	0	6	0	0	0	**243**	0	6	0	0	0	9
3	M_0	6	4	6	0	0	0	**216**	0	0	0	0	0
	M_1	4	5	4	6	7	6	0	**111**	0	6	7	6
	M_2	6	4	6	0	0	0	0	0	**216**	0	0	0
	I_0	0	6	0	9	0	0	0	6	0	**243**	0	0
	I_1	0	7	0	0	9	0	0	7	0	0	**243**	0
	I_2	0	6	0	0	0	9	0	6	0	0	0	**243**

of second-order statistics for some RF types means that claims about isotrigon properties need to be checked rigorously. We are currently comparing discrimination of textures generated from binary and ternary seeds. In principle these may provide a sensitive measure of what neurons respond to, given that one can now produce superficially similar textures that are, or are not, isotrigon. The fact that for some rules how you initialize them matters, while for others it does not, is a statement that some rules are invertible, and others are not. For deterministic rules, one can see that "invertible" and "ergodic" are equivalent. An even stronger statement is true: all configurations of a given size are either allowed or excluded, and the allowed ones have equal probability.[31]

A possible criticism of the textures presented here is that they are not natural. It is true that they are not natural in the sense of not having the characteristic $1/f$ amplitude spectrum.[29,30] They are natural in the sense of having corners and intersections. Recent studies show that to accurately reconstruct corners and intersections fourth- to sixth-order correlation information is required.[3] This just confirms the common sense statement of Purpura *et al.*[1] that information about three or four relationships is necessary to encode information such as contours. Another way to make the patterns more natural is to linearly filter them in order to produce a $1/f$ spectrum. In that case the patterns remain isotrigon with similarly filtered noise.[15,25,28]

APPENDIX A

In Table 3, note that, although the discrete equation has the same form whether the seed data are $\{-1,0,1\}$ or $\{0,1,2\}$, the actual coefficients will differ. The presented $x_{\alpha\beta\gamma}$ are those for the seed=$\{-1,0,1\}$ case. They typically have fractional values but are here multiplied by 8 to make them all integers for ease of presentation. Note that the full discrete equation is the sum of all the 27 multinomial terms: $f(a,b,c)=x_{000}+x_{100}a+\cdots+x_{111}abc+\cdots+x_{122}ab^2c^2+x_{222}a^2b^2c^2$. The equations for M_1 and I_1 are the simplest, given their large number of 0-valued coefficients.

In Table 4, textures were all generated with binary $\{-1,1\}$ seed noise. When the textures are seeded with ternary noise $\{-1,0,1\}$ the diagonals are all 243, and off-diagonal entries corresponding to values larger than 9 below are all 27. For those with 9 or fewer here, the number is 9. The 8 bolded zeros in the table have 27 intersections in the ternary case. Tables 5 and 6 show substitution rules related to Fig. 7.

Table 5. Examples of Substitution Rules (Fig. 7)

Name	Discrete Equation	Modulo 3 Version
$Rule_1$	$-a-b+1.5a^2b+1.5ab^2$	mod 3(2(a+b))
$Rule_2$	$a+b-1.5a^2b-1.5ab^2$	mod 3(1(a+b))
$Rule_3$	$a-b+1.5a^2b-1.5ab^2$	mod 3(2(−a+b))
$Rule_4$	$-a+b-1.5a^2b+1.5ab^2$	mod 3(1(−a+b))

Table 6. Equivalent Substitution Rules (Fig. 7)

Inputs	Outputs (f)			
a,b	$Rule_1$	$Rule_2$	$Rule_3$	$Rule_4$
−1,−1	−1	1	0	0
−1,0	1	−1	−1	1
−1,1	0	0	1	−1
0,−1	1	−1	1	−1
0,0	0	0	0	0
0,1	−1	1	−1	1
1,−1	0	0	−1	1
1,0	−1	1	1	−1
1,1	1	−1	0	0

ACKNOWLEDGMENTS

This research was supported by the Centre for Visual Sciences and the Australian Research Council through the ARC Centre of Excellence in Vision Science (CE0561903).

Corresponding author T. Maddess's e-mail address is ted.maddess@anu.edu.au; phone, 612-6249-4099; fax, 612-6249-3808.

REFERENCES

1. K. P. Purpura, J. D. Victor, and E. Katz, "Striate cortex extracts higher-order spatial correlations from visual textures," Phys. Rev. A **91**, 8482–8486 (1994).
2. M. O. Franz and B. Schölkopf, "Implicit Weiner series Part 1: cross-correlation vs. regression in reproducing kernel Hilbert spaces," Technical Report TR-114 (Max-Planck-Institut für biologische Kybernetik, 2003).
3. M. Franz and B. Schölkopf, "Implicit Weiner series for higher-order image analysis," in *Advances in Neural Information Processing Systems, 17: Proceedings of the 2004 Conference*, L. K. Saul, Y. Weiss, and L. Bottou, eds., (MIT Press, 2005), pp. 1–8.
4. C. Zetzsche, E. Barth, and B. Wegmann, "The importance of intrinsically two-dimensional image features in biological vision and picture coding," in *Digital Images and Human Vision*, A. Watson, ed. (MIT Press, 1993), pp. 109–138.
5. C. Zetzsche, G. Krieger, and B. Wegmann, "The atoms of vision: Cartesian or polar?" J. Opt. Soc. Am. A **16**, 1554–1565 (1999).
6. B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," Nature **381**, 607–609 (1996).
7. C. Zetzsche and U. Nuding, "Nonlinear and higher-order approaches to the encoding of natural scenes," Network **16**, 191–221 (2005).
8. C. Zetzsche and F. Rohrbein, "Nonlinear and extra-classical receptive field properties and the statistics of natural scenes," Network **12**, 331–350 (2001).
9. G. Krieger and C. Zetzsche, "Nonlinear image operators for the evaluation of local intrinsic dimensionality," IEEE Trans. Image Process. **5**, 1026–1042 (1996).
10. G. Krieger, I. Rentschler, G. Hauske, K. Schill, and C. Zetzsche, "Object and scene analysis by saccadic eye-

movements: an investigation with higher-order statistics," Spatial Vis. **13**, 201–214 (2000).
11. J. D. Victor and V. Zemon, "The human visual evoked potential: analysis of components due to elementary and complex aspects of form," Vision Res. **25**, 1829–1842 (1985).
12. J. D. Victor and M. M. Conte, "Cortical interactions in texture processing: scale and dynamics," Visual Neurosci. **2**, 297–313 (1989).
13. J. D. Victor and M. M. Conte, "Spatial organization of nonlinear interactions in form perception," Vision Res. **31**, 1457–1488 (1991).
14. J. D. Victor, "Complex visual textures as a tool for studying the VEP," Vision Res. **25**, 1811–1827 (1985).
15. J. D. Victor and M. M. Conte, "The role of high-order phase correlations in texture processing," Vision Res. **36**, 1615–1631 (1996).
16. L. L. Beason-Held, K. P. Purpura, J. W. Van Meter, N. P. Azari, D. J. Mangot, L. M. Optican, M. J. Mentis, G. E. Alexander, C. L. Grady, B. Horwitz, S. I. Rapoport, and M. B. Schapiro, "PET reveals occipitotemporal pathway activation during elementary form perception in humans," Visual Neurosci. **15**, 503–510 (1998).
17. L. L. Beason-Held, K. P. Purpura, J. S. Krasuski, J. M. Maisog, E. M. Daly, D. J. Mangot, R. E. Desmond, L. M. Optican, M. B. Schapiro, and J. W. VanMeter, "Cortical regions involved in visual texture perception: a fMRI study," Brain Res. Cognit. Brain Res. **7**, 111–118 (1998).
18. L. L. Beason-Held, K. P. Purpura, J. S. Krasuski, R. E. Desmond, D. J. Mangot, E. M. Daly, L. M. Optican, S. I. Rapoport, and J. W. VanMeter, "Striate cortex in humans demonstrates the relationship between activation and variations in visual form," Exp. Brain Res. **130**, 221–226 (2000).
19. T. Maddess and Y. Nagai, "Discriminating isotrigon textures," Vision Res. **41**, 3837–3860 (2001).
20. J. D. Victor and M. M. Conte, "Visual working memory for image statistics," Vision Res. **44**, 541–556 (2004).
21. C. W. Tyler, "Theory of texture discrimination based on higher-order perturbations in individual texture samples," Vision Res. **44**, 2179–2186 (2004).
22. C. W. Tyler, "Beyond fourth-order texture discrimination: generation of extreme-order and statistically-balanced textures," Vision Res. **44**, 2187–2199 (2004).
23. T. Maddess, Y. Nagai, A. C. James, and A. Ankiewcz, "Binary and ternary textures containing higher-order spatial correlations," Vision Res. **44**, 1093–1113 (2004).
24. A. Watson and A. Ahumada, "A standard model for foveal detection of spatial contrast," J. Vision **5**, 717–740 (2005).
25. J. D. Victor, "Images, statistics and textures: implications of triple correlation uniqueness for texture statistics and the Julesz conjecture: comment," J. Opt. Soc. Am. A **11**, 1680–1684 (1994).
26. C. Chubb and J. I. Yellott, "Every discrete, finite image is uniquely determined by its dipole histogram," Vision Res. **40**, 485–492 (2000).
27. C. Chubb and J. I. Yellott, "Dipole statistics of discrete finite images: two visually motivated representation theorems," J. Opt. Soc. Am. A **19**, 825–832 (2002).
28. B. Julesz, E. N. Gilbert, and J. D. Victor, "Visual discrimination of textures with identical third-order statistics," Biol. Cybern. **31**, 137–140 (1978).
29. D. L. Ruderman, "Origins of scaling in natural images," Vision Res. **37**, 3385–3398 (1997).
30. D. Ruderman and W. Bialek, "Statistics of natural images: scaling in the woods," Phys. Rev. Lett. **73**, 814–817 (1994).
31. E. N. Gilbert, "Random colorings of a lattice on squares in the plane," SIAM J. Algebraic Discrete Methods **1**, 152–159 (1980).
32. R. A. Johnson and D. W. Wichern, *Applied Multivariate Statistical Analysis*, 3rd ed. (Prentice Hall, 1992).
33. D. Peterzell and D. Teller, "Individual differences in contrast sensitivity functions: the lowest spatial frequency channels," Vision Res. **36**, 3077–3085 (1995).
34. J. D. Victor, C. Chubb, and M. M. Conte, "Interaction of luminance and higher-order statistics in texture discrimination," Vision Res. **45**, 311–328 (2005).
35. J. D. Victor and M. M. Conte, "Motion mechanisms have only limited access to form information," Vision Res. **30**, 289–301 (1990).