Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 45 (2005) 311–328
Vision Research
www.elsevier.com/locate/visres

Interaction of luminance and higher-order statistics in texture discrimination

Jonathan D. Victor [a,*], Charles Chubb [b], Mary M. Conte [a]

[a] Department of Neurology and Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021, United States
[b] Department of Cognitive Sciences, University of California at Irvine, Irvine, CA 92697, United States

Received 1 March 2004; received in revised form 11 August 2004

Abstract

Most studies of texture processing are based on textures in which individual pixel statistics are varied and spatial correlations are absent (“IID textures”), or textures in which spatial correlation structure is varied and luminance, or first-order, statistics are held constant. Here we jointly examine simple pixel statistics and fourth-order spatial correlation structure along the continuum of “even” and “odd” isodipole textures of Julesz, Gilbert and Victor [Julesz, B., Gilbert, E.N., & Victor, J.D. (1978). *Biological Cybernetics, 31*(3) 137–140], as well as their interactions. Absolute efficiency to detect either kind of statistical cue is low: approximately 0.05 for luminance statistics, and 0.004 for isodipole statistics. Above threshold, isodipole statistics must change by approximately four times the amount that pixel statistics must change to generate an equally salient texture. When pixel statistics and isodipole statistics are simultaneously varied, the two texture cues combine by probability summation and perceptual distances are approximately Euclidean. Superimposed on this picture are subtle foreground/background asymmetries that suggest properties of the visual mechanisms that are sensitive to these image statistics.

Keywords: Histogram statistics; Markov random fields; Probability summation

1. Introduction

Segregation of an image into regions that represent distinct objects or surfaces is a fundamental operation of visual processing. Typical objects and surfaces are not uniform, but are characterized by distinctive visual textures (Knill, 1998; Portilla & Simoncelli, 2000). Such texture differences, in addition to differences in luminance, are well known to provide strong cues for image segregation (Beck, Sutter, & Ivry, 1987; Bergen & Adelson, 1988; Graham, 1989; Gurnsey & Browse, 1989; Julesz, 1981).

Theories of texture segregation generally have two components: a stage in which local computations extract statistical attributes of the image, and a second stage at which these measurements are used for grouping and/or segregation (Bergen & Adelson, 1988; Bergen & Landy, 1991; Chubb & Landy, 1991; Graham, 1989; Graham, Beck, & Sutter, 1992; Grossberg & Mingolla, 1985; Malik & Perona, 1990; Wilson, 1993).

To understand the first stage of this process, many investigators have studied segregation and discrimination of textures that differ in specific kinds of statistical attributes. The first-order statistics of a texture consists of its distribution of luminance values, or, equivalently, a histogram that indicates the fraction of pixels that are assigned each gray level. For binary textures (such as the ones considered here) that have only black and white checks, specifying the first-order statistics is equivalent

* Corresponding author. Tel.: +1 212 746 2343; fax: +1 212 746 8984.
E-mail address: jdvicto@med.cornell.edu (J.D. Victor).

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.08.013

to specifying the mean luminance. The second-order statistics of a texture consist of the joint distribution of pairs of luminance values, at each vector displacement. Textures that have the same first-order statistics but not second-order statistics are typically discriminable. However, it is not straightforward to determine whether the basis for this discrimination is sensitivity to aspects of spatial structure, versus merely sensitivity to altered effective contrast. This is because second-order statistics affect the spatial frequency content of the stimulus. By virtue of the linear spatial filtering properties of the early visual system, alterations in spatial frequency content will be confounded with alterations in effective contrast.

Therefore, to constrain models for texture segregation, two kinds of stimuli have been particularly useful: "independent identically distributed" (IID) textures, and "isodipole" texture pairs. In an IID texture, the luminance in each pixel is chosen independently from a specified distribution. Textures generated in this fashion have no spatial correlations, and thus allow study of visual sensitivity to distributions of gray levels ("histogram statistics") in the absence of form cues. Discrimination between IID textures can be accounted for on the basis of sensitivity to a small number of specific histogram statistics (Chubb, Econopouly, & Landy, 1994).

Isodipole texture pairs, conversely, are typically constructed from only two luminance values, and have strong spatial correlations. Two isodipole texture pairs, by definition, have the same pairwise correlations ("dipole statistics") at all separations, but have higher-order correlations that differ. They thus cannot be distinguished by linear or quadratic mechanisms, even after linear filtering, and they also share the same effective contrast. Here we consider fourth-order correlations, and use the "even and odd" isodipole textures (Julesz, Gilbert, & Victor, 1978; Pollack, 1971a, 1971b, 1972, 1973) as exemplars (see Appendix A). Sensitivity of cortical (Purpura, Victor, & Katz, 1994) but not lateral geniculate neurons (Victor, 1986) to these fourth-order statistics has been demonstrated experimentally.

Although consideration of IID and isodipole textures is analytically convenient, few if any natural textures are defined purely on the basis of local histogram statistics, and few if any texture boundaries are examples of changes in higher-order statistics without associated changes in lower-order statistics. However, there has been little in the way of systematic study of mixtures of these statistical cues, and little is known about how they interact. This paper is intended as an initial exploration of these interactions. To examine the interaction of isodipole and luminance statistics, we construct a two-parameter space of binary textures, in which isodipole and luminance statistics can be independently specified. We consider only binary textures, and use the "even and odd" textures as a starting point. The strength of the fourth-order correlation is parameterized by α, where $\alpha = 1$ denotes the even texture, $\alpha = -1$ denotes the odd texture, and $\alpha = 0$ denotes a completely random texture. For all values of α, these textures have half of their checks black, and half white. We then modify their histogram statistics (the fraction of checks that are white) in a manner that preserves the equality of higher-order statistics. This results in a two-parameter family of textures $T(\gamma, \alpha)$ (Fig. 1A), in which a first parameter, γ, describes the extent of luminance bias, and a second parameter, α, describes the strength of the fourth-order statistics. Textures for which $\gamma = 0$ are the original isodipole textures, with correlation strength specified by α. Textures for which $\alpha = 0$ are a specific family of binary textures with luminance bias specified by γ. By examining discrimination of textures defined in the (γ, α)-plane, we can determine whether these cues interact, and if so, how.

2. Methods

2.1. Subjects and training

Studies were conducted in four normal subjects (2 male, 2 female), ages 25–51. Two subjects (AO and CC) were naïve to the purpose of the experiment. All subjects were practiced psychophysical observers in tasks involving visual textures and had visual acuities (corrected if necessary) of 20/20 or better.

2.2. Stimuli

Stimuli were presented on a PC programmed in Matlab with the PsychToolbox extensions (Brainard, 1997). Stimuli subtended an 11.6° square when viewed binocularly at 57 cm. The monitor had a luminance 57 cd/m^2 and a refresh rate of 75 Hz. Contrast was 1.0, and presentation duration was 160 ms.

Stimuli consisted of a 64×64 element array of black and white checks centered on a background of mean gray luminance. A 16×64 element rectangle within this array (the "target") was filled according to one rule for texture generation, designated the "foreground" texture. The remaining checks were colored according to a different rule, designated the "background" texture. The foreground could occupy any of four positions with equal probability: vertically oriented in the middle of the left or right half of the array, or horizontally oriented in the middle of the top or bottom half of the array (see Fig. 1C for examples).

Textures (both foreground and background) were drawn from a two-parameter family $T(\gamma, \alpha)$, illustrated in Fig. 1A and described in greater detail in Appendix B. The parameter γ denotes the extent of luminance bias. The fraction of white checks is equal to $\frac{1}{2}(1+\gamma)$ and the fraction of black checks is equal to $\frac{1}{2}(1-\gamma)$, so that $\gamma = 1$

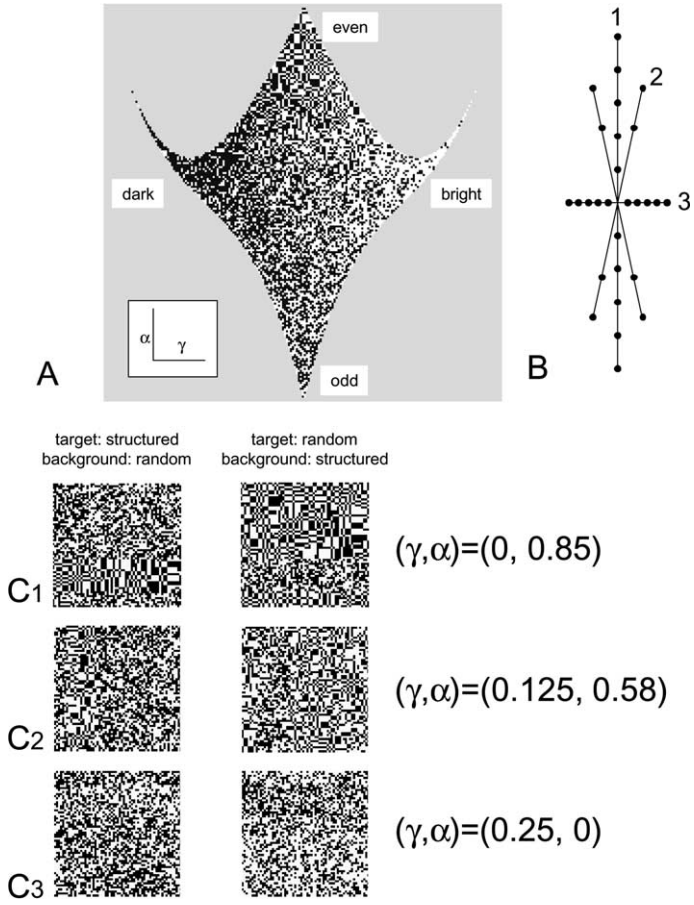

Fig. 1. (A) The two-parameter space of texture stimuli $T(\gamma,\alpha)$, in which α determines the strength of the fourth-order correlation, and γ determines the luminance bias. Although α and γ individually can range from -1 to 1, only limited values of the parameter pairs (γ,α) can be realized. This restricted two-parameter space is illustrated by the textured region of the (γ,α)-plane. At each point $T(\gamma,\alpha)$ in this region, the corresponding texture rule $T(\gamma,\alpha)$ is used. In the stimuli used in the experiment, discrete values of α and γ were used to create uniform texture arrays. The locations of the specific values used are illustrated in panel B. In panels A and B, the origin (γ,α) $(0,0)$ is at the center. (C) Stimuli corresponding to the three labeled points in panel B. As indicated, each point corresponds to a particular (γ,α)-pair, with values as indicated. For each (γ,α)-pair, two stimuli (corresponding to the two experimental conditions) are shown: one stimulus consisting of a target patch of texture $T(\gamma,\alpha)$ on a background of texture $T(0,0)$, and another stimulus consisting of a target patch of texture $T(0,0)$ on a background of texture $T(\gamma,\alpha)$. In C1((γ,α) = (0,0.85), "even"), the targets are located at the bottom of the stimuli; in C2((γ,α) = (0.125,0.58), combined "bright" and "even"), at the left of the stimuli, and in C3((γ,α) = (0.25,0), "bright"), at the top of the stimuli.

denotes a texture that is completely white, $\gamma = -1$ denotes a texture that is completely black, and $\gamma = 0$ denotes a texture in which half of the checks are white, and half are black. The second parameter, α, determines the strength of the fourth-order correlation, where $\alpha = 1$ denotes the even texture, $\alpha = -1$ denotes the odd texture, and $\alpha = 0$ denotes α texture with no fourth-order correlations. For each value of γ, variations in a leave the second- and third-order statistics unchanged. As described in Appendix B, there is a restricted two-parameter space of (γ,α)-pairs for which these properties can be realized. This two-parameter space, from which all foreground and background textures are drawn, is illustrated in Fig. 1A.

The specific values of (γ,α)-pairs used in these experiments are shown in Fig. 1B. Along each of the cardinal axes, five values were chosen, equally spaced on a linear scale up to a maximum value of $|\gamma| = 0.25$ or $|\alpha| = 0.85$. Along the diagonals, one set of values was chosen at $|\gamma| = 0.125$ and $|\alpha| = 0.58$, and a second set was chosen at 70% of the distance from the origin to these points ($|\gamma| = 0.09$ and $|\alpha| = 0.41$). This yielded 20 (γ,α)-pairs along the axes, and eight along the diagonals (see Fig. 1B). For each (γ,α)-pair, stimuli were constructed with that texture as foreground and the random ((γ,α) = (0,0)) texture as background, and with the assignments of the textures reversed (40 assignments along the axes, 16 along the diagonals)—see examples in Fig. 1C. In each experimental block, these assignments were presented with the target in each of the four possible positions, yielding 160 stimuli along the cardinal axes, and 64 along the diagonals. Two realizations of the stimuli along the diagonals were included in each block, thus yielding a total of $288 = 160 + 2 \times 64 =$ trials per block.

The scan of a CRT necessarily induces small correlations in luminance between adjacent pixels along the raster lines, which, to a first approximation, is akin to an anisotropic blurring filter. However, we doubt that these artifacts played a significant role in these experiments, since (a) the texture checks were approximately 8 display pixels wide, (b) there was no difference in performance for targets that were horizontal or vertical, and (c) all of the stimuli had the same (flat) spectrum.

2.3. Procedure

The subject's task was to identify the position of the target (a four-alternative forced choice texture segregation task). Subjects were told that it was equally likely to be in any of four locations (top, right, bottom, left), and were instructed to maintain central fixation on a one-pixel dot, rather than to attempt to scan the stimulus. This task is essentially identical to the one used by Chubb and coworkers in the study of IID textures (Chubb, Landy, & Econopouly, 2004). Subjects practiced to stability (2–3 h) prior to data collection. Feedback on error was provided with a tone in practice and data collection runs. After performance stabilized, blocks of the 288 trials described above (with trials presented in randomized order) were presented. We collected data from 15 such blocks (4320 trials per subject), grouped into three or four experimental sessions.

3. Results

3.1. Performance along cardinal and diagonal axes

We first describe performance in the texture segregation task along the cardinal axes of the stimulus space of Fig. 1. Along the horizontal axis γ, target and background differed only in first-order (luminance) histogram statistics but not in fourth-order statistics. Along the vertical axis α, target and background differed in fourth-order statistics but not in luminance.

Performance for the four observers is shown in the insets of Fig. 2. The sigmoidal dependence of fraction correct on the strength of the statistical cue (α or γ) suggested fitting these data with Weibull functions (Weibull, 1951). Since chance performance was 0.25, the model along the isodipole axis was

$$f_{\mathrm{model}}(0,\alpha) = \begin{cases} 1 - \frac{3}{4}\exp\left(-\left|\frac{\alpha}{\alpha_{\mathrm{iso}^+}}\right|^b\right), \\ 1 - \frac{3}{4}\exp\left(-\left|\frac{\alpha}{\alpha_{\mathrm{iso}^-}}\right|^b\right) \end{cases} \tag{1}$$

with the first alternative (scale parameter a_{iso^+}) chosen for $\alpha > 0$, and the second alternative (scale parameter a_{iso^-}) chosen for $\alpha < 0$. Similarly, the model along the luminance axis was

$$f_{\mathrm{model}}(\gamma,0) = \begin{cases} 1 - \frac{3}{4}\exp\left(-\left|\frac{\gamma}{a_{\mathrm{lum}^+}}\right|^b\right), \\ 1 - \frac{3}{4}\exp\left(-\left|\frac{\gamma}{a_{\mathrm{lum}^-}}\right|^b\right) \end{cases} \tag{2}$$

with the first alternative (scale parameter a_{lum^+}) chosen for $\gamma > 0$, and the second alternative (scale parameter a_{lum^-}) chosen for $\gamma < 0$. Thus, we constructed fits separately along the positive. and negative axes, for each of the four cardinal directions. The model was determined by the Weibull exponent b and the four scale parameters a_{iso^+}, a_{iso^-}, a_{lum^+}, and a_{lum^-}. We used a separate set of scale parameters for the trials in which (i) the target was structured and the background was random (black in Fig. 2), or (ii) the target was random and the background was structured (gray in Fig. 2). However, a single Weibull exponent b was used for all eight curves within each subject. Thus, for each subject, nine parameters were simultaneously determined via maximum likelihood: two sets of values of the four scale parameters, and a single value of the Weibull exponent b. This fit was constrained by 40 measurements: 4 directions $\times$ 5 distances $\times$ 2 target conditions. Allowing the Weibull exponent to vary according to direction and/or target condition produced only minimal improvement in the fits.

The fitted parameters are shown in Table 1, along with confidence limits as determined by a bootstrap procedure. To calculate bootstrap confidence limits, we proceeded as follows: in a given single simulation (of the 1000 performed for each subject), for each condition (five conditions along each main half-axis and two conditions along each half-diagonal), we took the simulated data in that condition from a binomial distribution based on the observed proportion correct, and number of trials in that condition. The surrogate data were then fitted by maximum likelihood, and confidence limits were then calculated from these surrogate fits via the standard bootstrap procedure (Efron & Tibshirani, 1998).

The scale parameters in the positive and negative isodipole (α)-direction, a_{iso^+} and a_{iso^-} were three to fourfold higher than the scale parameters in the positive and negative luminance (γ)-direction, a_{lum^+} and a_{lum^-}, indicating a corresponding ratio of sensitivities. There was no consistent difference in the scale parameter along the positive versus the negative γ-direction. However, for three of the four subjects, (AO, MC, and CFC), the scale parameter was significantly smaller in the positive isodipole (even, $\alpha > 0$) direction than in the negative isodipole (odd, $\alpha < 0$) direction for the random-target conditions, indicating a lower threshold for detecting a random target on an even background, than on an

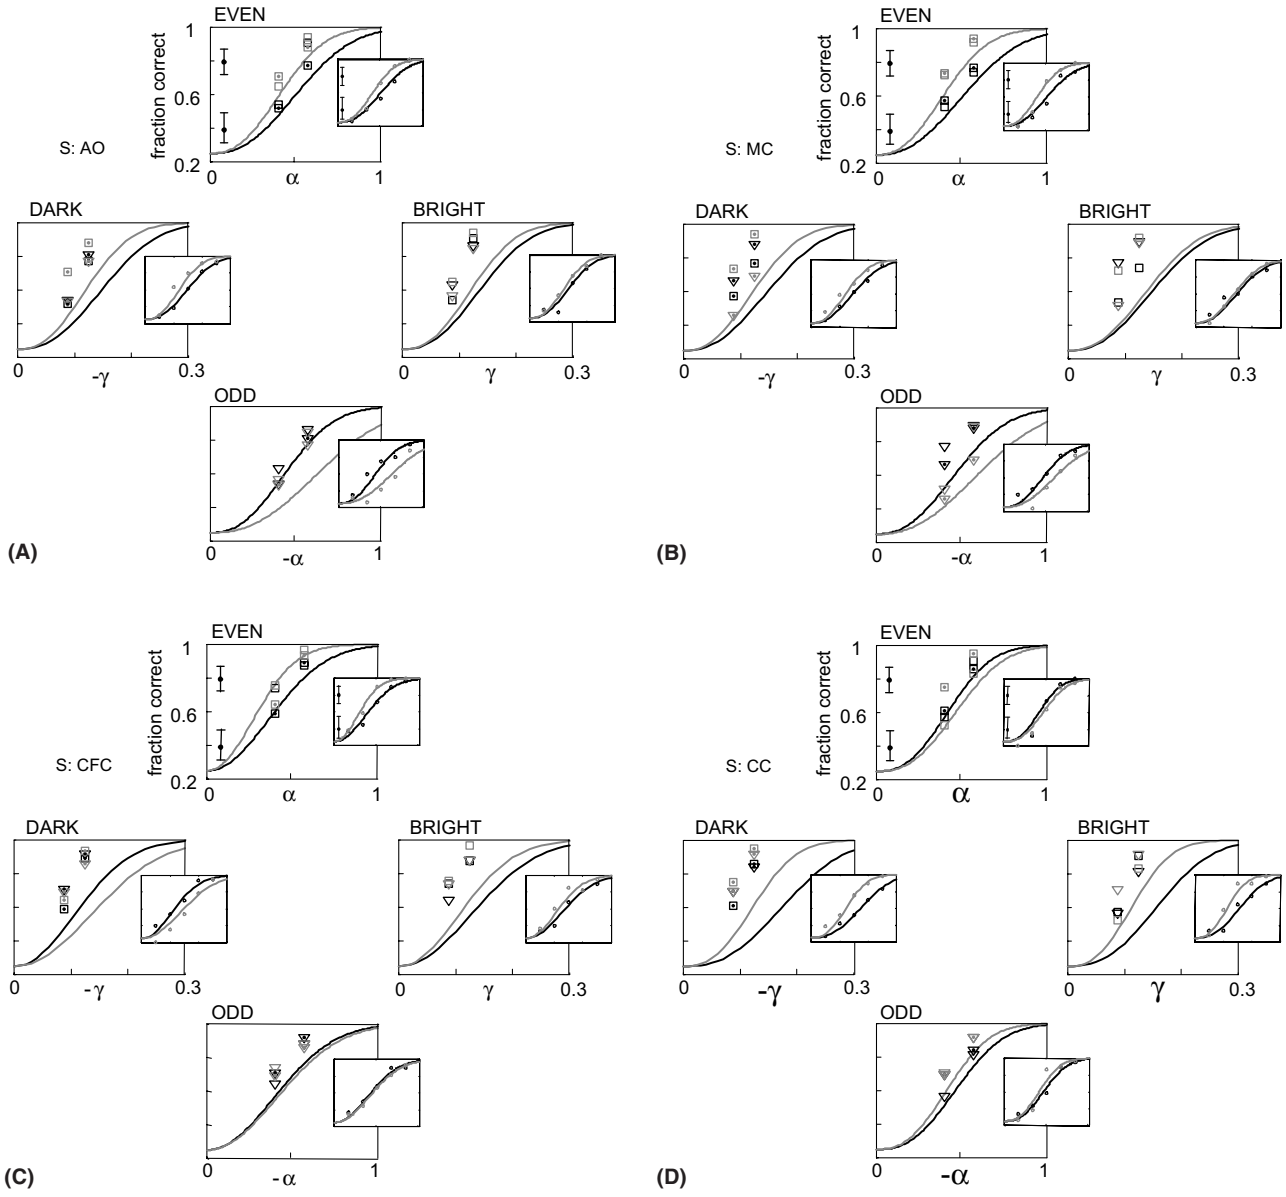

Fig. 2. Psychophysical performance on the texture segregation task for the four subjects. *Insets* show observed fraction correct for stimuli along each of the cardinal axes (data points), along with the Weibull functions (curves) fit via maximum likelihood. Black curves correspond to conditions in which the target was structured and the background was random; gray curves correspond to conditions in which the target was random and the background was structured. *Main graphs* compare these Weibull functions (smooth curves) with the observed fraction correct for stimuli along adjacent diagonal directions. The directions corresponding to the isolated data points are designated as follows: squares, $\alpha > 0$; triangles, $\alpha < 0$; symbols without a central dot, $\gamma > 0$; symbols with a central dot, $\gamma < 0$. Error bars (shown only in the upper main graph and upper inset) indicate 95% confidence limits determined by binomial statistics.

odd background. This difference reversed for two subjects (AO and MC) when the target was structured and the background was even.

The main graphs of Fig. 2 compare performance along the diagonal directions of the stimulus space (isolated symbols) with the performance along the adjacent cardinal directions (smooth curves taken from the corresponding insets). In the conditions represented by the isolated symbols, target and background differed both in local histogram statistics and fourth-order statistics. In all cases, performance was better than predicted from the individual cues in isolation. This relationship held in all four diagonal directions and in all four subjects. These data thus indicate that isodipole (α) and luminance (γ) cues could be combined.

3.2. Error patterns

For three of the four subjects, we collected a detailed log of the error patterns in some or all of the blocks.

Table 1
Maximum-likelihood estimates and 95% confidence limits on model parameters of Eq. (4) for the four subjects

Parameter	Subject AO	Subject MC	Subject CFC	Subject CC
Random background				
Structured target				
$a_{\text{lum+}}$	0.140	0.152	0.143	0.164
	(0.128, 0.152)	(0.132, 0.169)	(0.129, 0.159)	(0.149, 0.183)
$a_{\text{lum−}}$	0.161	0.154	0.125	0.177
	(0.146, 0.175)	(0.139, 0.174)	(0.112, 0.137)	(0.159, 0.201)
$a_{\text{iso+}}$	0.533	0.570	0.438	0.447
	(0.497, 0.573)	(0.526, 0.623)	(0.401, 0.477)	(0.419, 0.488)
$a_{\text{iso−}}$	0.490	0.447	0.444	0.477
	(0.452, 0.529)	(0.404, 0.495)	(0.406, 0.485)	(0.447, 0.525)
Random target				
Structured background				
$a_{\text{lum+}}$	0.124	0.137	0.121	0.124
	(0.113, 0.136)	(0.125, 0.155)	(0.108, 0.134)	(0.112, 0.138)
$a_{\text{lum−}}$	0.129	0.136	0.156	0.120
	(0.117, 0.141)	(0.119, 0.150)	(0.142, 0.172)	(0.094, 0.132)
$a_{\text{iso+}}$	0.444	0.409	0.358	0.493
	(0.411, 0.478)	(0.379, 0.493)	(0.328, 0.387)	(0.458, 0.563)
$a_{\text{iso−}}$	0.649	0.646	0.454	0.438
	(0.597, 0.709)	(0.590, 0.712)	(0.413, 0.492)	(0.406, 0.504)
Weibull exponent b	2.570	2.368	2.131	2.564
	(2.328, 2.860)	(2.005, 2.694)	(1.937, 2.357)	(2.233, 2.839)
Minkowski exponent m	2.350	2.114	2.143	2.386
	(2.047, 2.857)	(1.838, 2.721)	(1.849, 2.539)	(2.029, 3.068)

Confidence limits (95%) were determined by a bootstrap.

When wrong, subjects chose the position opposite to that of the target disproportionately often. If errors were made at random, each false position would account for 1/3 of the errors. However, we found that the opposite position was selected in 42% of the error trials (AO), 43% of the error trials (MC), and 45% of the error trials (CC). This pattern of performance, which was stable over time, suggests that subjects were able to detect the orientation but not position of the target boundary. No subject had a clear bias for left versus right-sided targets, but two subjects (AO and CC) had a response bias directed towards the bottom target position and one subject (MC) had a smaller response bias directed towards the upper target position. These biases were independent of the parameters (γ, α) of the target.

3.3. Analysis of cue combination

To analyze how the isodipole and luminance cues combine, we considered a range of combination rules, based on Minkowski geometries. In a combination rule corresponding to a Minkowski geometry with exponent m (Poirson & Wandell, 1990), the distance from the origin to a general point (γ, α), here denoted $d(\gamma, \alpha)$, is postulated to be related to the distances along the axes, $d(0, \alpha)$ and $d(\gamma, 0)$ by the following rule:

$$d(\gamma, \alpha) = \left(d(\gamma, 0)^m + d(0, \alpha)^m\right)^{1/m}. \quad (3)$$

For $m = 2$, Eq. (3) is the Pythagorean rule, and the distances correspond to a Euclidean distance. In the limit $m \to \infty$, Eq. (3) states that $d(\gamma, \alpha)$ is the maximum of $d(0, \alpha)$ and $d(\gamma, 0)$, i.e., there is no subthreshold summation. At the other extreme ($m = 1$), the distances combine according to "city block" geometry.

To carry out this analysis, we fitted the psychophysical performance data to Weibull functions along the examined eight directions. The scale parameter of the Weibull functions along the four diagonal directions was determined by a Minkowski combination of the scale parameters a_{lum} and a_{iso} along the adjacent coordinate axes. Thus, the observed fraction correct data $f(\gamma, \alpha)$ were fit to the following model:

$$f_{\text{model}}(\gamma, \alpha) = 1 - \frac{3}{4}\exp\left(-\left(\left|\frac{\gamma}{a_{\text{lum}}}\right|^m + \left|\frac{\alpha}{a_{\text{iso}}}\right|^m\right)^{b/m}\right). \quad (4)$$

The scale parameter a_{lum} is chosen to be one of the two values, a_{lum^+} or a_{lum^-}, depending on the sign of γ, and similarly, the scale parameter a_{iso} is chosen to be a_{iso^+}, and a_{iso^-}, depending on the sign of α. The parameters a_{lum^+}, a_{lum^-}, a_{iso^+}, and a_{iso^-} are also allowed to depend on whether the target was structured or random. As in the cardinal-axis fits, only a single value of the Weibull exponent b and the Minkowski exponent m were used. Thus, a tenth parameter (the Minkowski exponent) was added to the nine parameters determined by the cardinal-axis fits. This additional parameter was constrained

by 16 data points (performance in the four diagonal directions, at two levels, and for two target conditions). Values of m obtained by a maximum-likelihood fitting procedure, and confidence limits (determined as above via a bootstrap procedure) are given in the final row of Table 1.

Across subjects, the results are remarkably consistent. In all the four subjects, the Minkowski exponent m clusters around values just slightly higher than 2. Thus, the two cues appear to combine in a Euclidean fashion.

Moreover, the Weibull exponent b and the Minkowski exponent m are statistically indistinguishable in all four subjects. Equality of these exponents (for any geometry specified by m) corresponds to probability summation. This can be seen as follows. If $m = b$, Eq. (4) can be rewritten

$$f_{\mathrm{model}}(\gamma,\alpha) = 1 - \frac{3}{4}\exp\left(-\left|\frac{\gamma}{a_{\mathrm{lum}}}\right|^b\right)\exp\left(-\left|\frac{\alpha}{a_{\mathrm{iso}}}\right|^b\right). \tag{5}$$

Since this is a 4-alternative forced choice paradigm, the model fraction correct f_{model} is related to the model hit rate h_{model} by $f_{\mathrm{model}} = h_{\mathrm{model}} + \frac{1}{4}(1 - h_{\mathrm{model}}) = \frac{1}{4} + \frac{3}{4}h_{\mathrm{model}}$. Substituting this into Eq. (5) leads to

$$1 - h_{\mathrm{model}}(\gamma,\alpha) = \exp\left(-\left|\frac{\gamma}{a_{\mathrm{lum}}}\right|^b\right)\exp\left(-\left|\frac{\alpha}{a_{\mathrm{iso}}}\right|^b\right). \tag{6}$$

It follows (if $m = b$) that

$$1 - h_{\mathrm{model}}(\gamma,\alpha) = (1 - h_{\mathrm{model}}(\gamma,0))(1 - h_{\mathrm{model}}(0,\alpha)), \tag{7}$$

or

$$h_{\mathrm{model}}(\gamma,\alpha) = h_{\mathrm{model}}(\gamma,0) + h_{\mathrm{model}}(0,\alpha) - h_{\mathrm{model}}(\gamma,0)h_{\mathrm{model}}(0,\alpha) \tag{8}$$

as required for probability summation of the contributions of the two components.

3.4. Isodiscrimination contours

The Minkowski exponent determines the shape of the iso-discrimination contours. In the Euclidean case ($m = 2$), each quadrant of the iso-discrimination contour is a sector of an ellipse. In the limit $m \to \infty$, the iso-discrimination contours form rectangles aligned with the cardinal axes. For $m = 1$, the iso-discrimination contours are quadrilaterals, with their vertices on the cardinal axes.

In general, the relative extent of the iso-discrimination contour along each axis also depends on the criterion fraction correct. However, when $m = b$, the shape of the iso-discrimination contour (but not its size) is independent of the choice of criterion. When $m = b$, the fraction correct along any 0 direction $\gamma = z\gamma_0$, $\alpha = z\alpha_0$ given by

$$f_{\mathrm{model}}(z\gamma_0, z\alpha_0) = 1 - \frac{3}{4}\exp\left(\left|\frac{z}{a(\gamma_0,\alpha_0)}\right|^b\right) \tag{9}$$

with

$$\left|\frac{1}{a(\gamma_0,\alpha_0)}\right|^b = \left|\frac{\gamma_0}{a_{\mathrm{lum}}}\right|^b + \left|\frac{\alpha_0}{a_{\mathrm{iso}}}\right|^b. \tag{10}$$

Eq. (10) states that the shape of the psychophysical function is characterized by the same Weibull exponent b, for all the directions in the (γ,α) plane. Consequently, the distance from the origin to an iso-discrimination contour in the direction $\gamma = z\gamma_0$, $\alpha = z\alpha_0$ is proportional to $a(\gamma_0,\alpha_0)$. The proportionality constant depends on the criterion fraction correct and b, but not on the direction.

Thus, when the Minkowski and Weibull exponents are identical (or nearly so), the shape of the iso-discrimination contour is a particularly useful way to summarize the perceptual data. The iso-discrimination contours are plotted in Fig. 3 for the four subjects, for a fraction correct of 62.5% (halfway between chance and perfect performance). As is expected from the finding that $b \approx m \approx 2$, these contours are very nearly elliptical.

The main deviation from ellipticity in Fig. 3 was that the extent of the iso-discrimination contours along the positive and negative axes were not the same, and also depended on whether the target was structured and the background was random (black curves), or whether the target was random and the background was structured (gray curves). In subjects AO and MC, an even texture ($\alpha > 0$) was easier to discriminate as background than as foreground, while an odd texture ($\alpha < 0$) was easier to discriminate as foreground than as background. Subject CFC showed the foreground versus background asymmetry only for even textures. Subject CC did not show these asymmetries. Asymmetries were also seen along the γ-axis, but these were less consistent: bright texture ($\gamma > 0$) easier to discriminate as background than foreground for subject CC; dark texture ($\gamma < 0$) easier to discriminate as background than foreground for subjects AO and CC; but easier to discriminate as foreground than background for subject CFC.

3.5. Intrinsic discriminability of the textures

The psychophysical findings have several features. (i) Sensitivity to first-order histogram statistics γ is three to four times greater than sensitivity to local fourth-order statistics α (compare values of a_{lum^+} and a_{lum^-} to a_{iso^+} and a_{iso^-} in Table 1). (ii) Three of the four subjects show an asymmetry between the even ($+\alpha$) and odd ($-\alpha$) directions (subjects AO, MC, and CFC). (iii) In the subjects that show this asymmetry, the asymmetry is in opposite directions for a structured target on a random

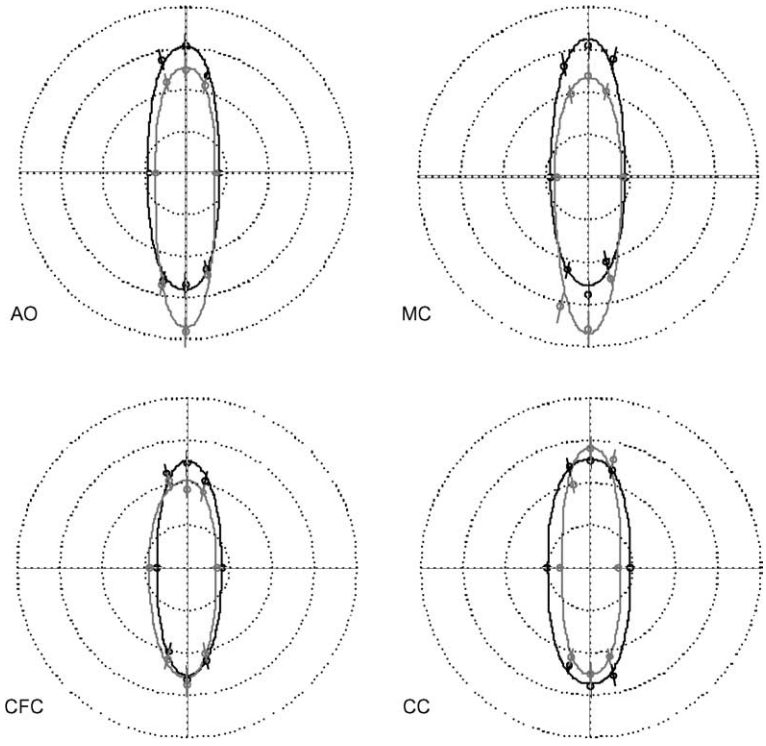

Fig. 3. Isodiscrimination contours for the four subjects. Data points indicate the coordinates (γ, α) for 62.5% correct for discriminating a structured target from a random background (black), or a random target from a structured background (gray), as interpolated from the psychometric functions of Fig. 2, fitted separately along each direction. The error bars correspond to 95% confidence limits. The curves correspond to the model of Eq. (4).

background, versus a random target on a structured background. (iv) In all subjects, the cues combine in an approximately Euclidean fashion. We next determine to what extent these findings might be due to the statistical properties of the stimuli themselves, rather than revealing the limitations of visual processing.

Standard information-theoretic analysis (Cover & Thomas, 1991; Latham, 2002) provides a way to quantify the extent to which two probability distributions are discriminable. Suppose P and Q are two probability distributions on a set of elements x_j, with p_j representing the frequency of x_j in P, and q_j representing the frequency of x_j in Q. Consider an observer who expects that a stream of observations is drawn from Q (''the background''), and whose task is to detect when the source of the observations is switched to P (''the foreground''). Under these circumstances, the minimum number of observations required to reach a criterion level of performance is proportional to the Kullback–Leibler divergence of P and Q, $K(P\|Q)$. This quantity is defined by

$$K(P\|Q) = \sum_j p_j \log_2\left(\frac{p_j}{q_j}\right). \tag{11}$$

The above definition can be extended to the situation in which P and Q represent texture ensembles, and the ''elements'' x_j represent indefinitely large samples of the textures. As we show in Appendix B, when the textures are generated from Markov random fields (as they are here), the Kullback–Leibler divergence (per unit area) can be evaluated in closed form from the 2×2 block probabilities. Thus, lines of constant intrinsic discriminability correspond to lines of constant Kullback–Leibler divergence.

Fig. 4 shows contour lines for the Kullback–Leibler divergences corresponding to discriminating a structured texture from a random background (black), or a random texture from a structured background (gray), plotted in the (γ, α)-space of the stimuli. Note that the central contour lines are nearly circular (see also Eq. (38) of Appendix B), and there is little deviation between the black and gray contours. Thus, in contrast to our experimental results, lines of constant intrinsic discriminability contours do not show (i) elongation along the α-axis, (ii) asymmetry between the positive and negative directions along either the α-axis or the γ-axis, or (iii) a foreground-background asymmetry. However, the analysis of intrinsic discriminability does predict the observed Euclidean combination of cues within the range examined.

4. Discussion

In this report, we systematically investigate discrimination of textures that differ both in histogram statistics

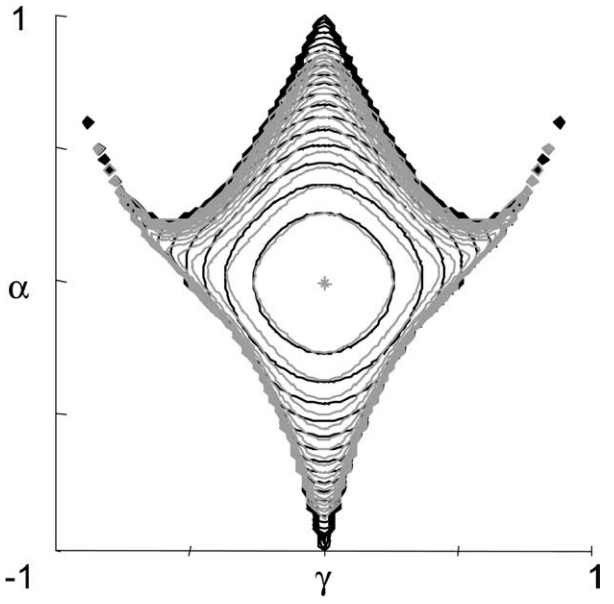

Fig. 4. Intrinsic discriminability of textures. Contour lines are lines of equal Kullback–Leibler divergences corresponding to discriminating a structured target from a random background (black), or a random target from a structured background (gray). Each contour line represents 0.05 bits.

and correlation structure. We begin by summarizing our findings, and then consider the possible mechanisms that may underlie our findings.

We found that the absolute sensitivity to differences in luminance statistics (as parameterized by γ) was approximately four times greater than sensitivity to differences in isodipole statistics (as parameterized by α). Luminance and isodipole cues combined according to probability summation in all subjects (concordance of the Minkowski exponent m and the Weibull exponent b). We note, however, that consistency with probability summation does not imply that probability summation is actually performed; other mechanisms might lead to the same behavioral outcome.

Perceptual distances combined in a manner rigorously consistent with a Euclidean distance in two subjects; in the other two subjects, the confidence limits for the Minkowski exponent extended to 2.05, but not 2. Experimental identification of deviations from Euclidean distances can be difficult (Poirson & Wandell, 1990), even when such deviations are present. However, no subject's data were consistent with (95% confidence limits) a Minkowski exponent less than 1.8, or greater than 3.1.

We also found modest asymmetries in the salience of a texture patch, contingent on the sign of the difference in isodipole statistics: A random ($\alpha = \gamma = 0$) patch on an even ($\alpha > 0$) background was more readily detected than an even patch on a random background. The opposite was true for the odd ($\alpha < 0$) textures. Foreground-background asymmetries for first-order histogram statistics were present but less prominent or consistent.

4.1. Efficiency of the human observer

The above analysis of the intrinsic discriminability of the textures shows that these features of the data reflect aspects of human performance, not of the textures themselves. However, this analysis, strictly speaking, is not an "ideal observer." analysis. This is because we examined the intrinsic discriminability of the textures, not the ideal strategy for performing this task. A true ideal observer knows the 56 possible pairings of foreground and background statistics, and simply inspects the stimulus to determine (in a Bayesian fashion) which is most likely. Such a calculation is laborious and gives little insight into human performance, since the strategy (though ideal) is highly implausible.

Nevertheless, an estimate of the absolute efficiency of the human observer can be obtained in two ways. In the first approach, we choose texture parameters that led to close to perfect performance among the subjects (Fig. 2) for the entire $16 \times 64 = 1024$-check target: $\alpha \approx 0.8$ for even versus random, and $\gamma \approx 0.25$ for luminance-biased versus random. Such near-perfect performance requires approximately 2 bits of information (distinguishing between one of four equally-likely alternatives). Based on the calculation of Eq. (34), we can determine the size ($n_1 \times n_2$) of a texture patch required to provide 2 bits of information in the Kullback–Leibler sense, based on the mismatch between two Markov random fields. The ratio of this critical value of $n_1 \times n_2$ (the size of the target required by information-theoretic criteria) to $16 \times 64 = 1024$ (the size of the target required by the human observer), is thus an estimate of the human observer's absolute efficiency. These calculations are summarized in Table 2. Note that for luminance statistics, the absolute efficiency is about 4%. This is about an order of magnitude lower than the efficiency of a human observer for detecting differences in dot density (Barlow, 1978). For isodipole statistics, the absolute efficiency is even lower, about 0.3%.

Another way to compare human performance to that of the ideal observer is to determine the level of statistical structure that would be required for the ideal observer to perform at near perfect levels, given access to a target of size $16 \times 64 = 1024$. Since the above analysis indicates that the values of α and γ will be small, we can use Eq. (38) of Appendix B to determine the values of α and γ at which $1024K_0$ provides the requisite 2 bits of information. This analysis indicates that near-perfect performance would be expected when (γ, α) lies on a circle of radius 0.052. Such textures are virtually indistinguishable from random for a human observer, even with scrutiny. Human performance becomes close to perfect only with a fivefold higher level of structure along the luminance (γ) axis, or a 15-fold higher level of structure along the isodipole (α) axis. Thus, looked at either way, human absolute efficiency is very low

Table 2
Efficiency analysis

Target (γ_P, α_P)	Background (γ_Q, α_Q)	K_0	Critical $n_1 n_2$	Estimated absolute efficiency
(0, ±0.8)	(0,0)	0.531	3.8	0.0037
(0,0)	(0, ±0.8)	0.737	2.7	0.0027
(±0.25, 0)	(0,0)	0.046	43.8	0.0429
(0,0)	(±0.25, 0)	0.047	42.9	0.0419

The first two columns specify texture pairs for which human observers performed at close to perfect levels. The column labeled K_0 represents the discriminability of the texture pairs, quantified as the Kullback–Leibler divergence per unit check (Eq. (35)). The fourth column is the number of checks required for the Kullback–Leibler divergence to reach 2 bits. The fifth column is an estimated absolute efficiency, equal to the ratio of the number of critical checks to the number of checks required by the human observer, 1024.

along the luminance axis, and several times lower along the isodipole axis.

4.2. What image statistics are used to perform these discriminations?

By definition, an IID texture is determined by the probabilities assigned to each luminance value i.e., a histogram of unit area. Studies of IID textures have shown that a rather complete account of segregation can be framed in terms of three specific statistics, i.e., linear functionals of this histogram (Chubb et al., 2004). The first two statistics may correspond approximately to mean and variance. Although the sensitivity functions characterizing these two statistics have not been measured, it has been established that (1) neither of these two statistics is significantly sensitive to third- or higher-order moments of the texture histogram, and (2) these two statistics jointly provide sensitivity to independent variations of both texture mean and texture variance. For our purposes, it suffices to observe that at least one of these first two statistics is sensitive to overall texture brightness as reflected by the histogram mean. The third statistic, "blackshot", is sensitive to contrasts near the black end of the range, and insensitive to the lighter contrasts. In this account, these three statistics are independently measured by the visual system. Two textures are discriminable when they are sufficiently different along at least one or more of these measures. However, the possibility that subthreshold differences along two or more measures can combine to result in discrimination has not been explicitly ruled out.

Spatial correlations are intrinsically much more difficult to parameterize than univariate luminance distributions. Consequently, our understanding of the role of spatial correlations underlying texture discrimination is far less comprehensive. Since spatial filtering confounds histogram statistics and second-order correlations, it is convenient to focus on isodipole texture pairs, namely, texture pairs that share the same luminance distribution and power spectrum. Investigations of discrimination of isodipole texture pairs have indicated that spatial arrangement of higher-order correlations, not order of correlation per se, determines whether textures can be discriminated (Victor & Conte, 1991). An extensive set of discriminability data can be accounted for on the basis of a local analysis that detects aligned contours in a nonlinear fashion (Victor & Conte, 1991), in accord with other observations of local grouping operations (Field, Hayes, & Hess, 1993).

The iso-discrimination contours observed in Fig. 3 are well fit by ellipses whose major and minor axes (corresponding to the principal components of salience) are aligned with the α and γ axes. Although it is parsimonious to conclude that the underlying statistics being extracted by human vision are α and γ, the data allow other possibilities. As observed by Poirson and Wandell (1990), if the cue-combination rule is Euclidean (as is approximately true in the current study), then the axes of the underlying mechanisms are only defined up to an arbitrary rotation of the principal components. Moreover, the statistics used by the visual system might be distinct from those used to define α and γ, but highly correlated with these parameters for the specific textures used here.

4.2.1. Inferences from studies with IID textures

Experiments with IID textures (Chubb et al., 2004) suggest that human vision has at least two systems, at least one brightness-sensitive system as well as the blackshot system, that are relevant to the judgments in the current experiment. Suppose that when applied to an IID texture T, the brightness-sensitive system extracts a statistic $M(T)$ approximately equal to the mean Weber contrast of the texture. By contrast, the blackshot system extracts a statistic $B(T)$ approximately equal to the proportion of checks in the texture that are black (i.e., that have Weber contrast close to -1.0) [1]. For the binary IID textures $T(\gamma, 0)$ the statistics $B(T)$ and $M(T)$ are highly (negatively) correlated, since all checks are either black or white. Thus the current study does

[1] It should be noted, however, that previous analyses of the "brightness" and "blackshot" systems has used only IID textures in which black checks occurred very sparsely. It is not clear that the documented behavior of the blackshot system will generalize to binary IID textures in which the proportion of black checks is much higher.

not permit us to dissociate the influences on performance of the brightness and blackshot systems. Rather, we are restricted to making inferences about the combined influence on performance exerted by the amalgam of blackshot and brightness systems (the "B&B amalgam").

Because the brightness and blackshot systems have only been investigated using IID textures, nothing is known about their possible sensitivity to spatial structure. Our results bear on this point. Suppose that the components of the B&B amalgam were sensitive to the statistical features controlled by α. For example, suppose that the occurrence in a texture T of an all-black, two-by-two cluster of checks increases the level of blackshot by more than the increase produced by four isolated black checks. Let performance be near threshold for discriminating a texture $T(\gamma_0, 0)$ from $T(0,0)$. Since $T(0,0)$ is IID and $T(\gamma_0, 0)$ is nearly so, we assume that the B&B amalgam alone drives discrimination between them. If γ is positive (i.e., if there are fewer black checks in $T(\gamma_0, 0)$ than in $T(0,0)$), then $B(T(\gamma_0, 0)) < B(T(0,0))$. Under these assumptions, $B(T(\gamma_0, \alpha))$ will be an increasing function of α (since raising α increases the frequency of 2×2 rectangles of black checks). That is, for $\alpha > 0$, $B(T(\gamma_0, \alpha))$ moves towards $B(T(0,0))$, leading to poorer performance. Similarly, for $\alpha < 0$, $B(T(\gamma_0, \alpha))$ moves away from $B(T(0,0))$, leading to improved discrimination of $T(\gamma_0, \alpha)$ from $T(0,0)$.

Thus, under the hypothetical scenario in which the blackshot mechanism is selectively activated by 2×2 rectangles of black checks, an iso-discrimination contour running through $T(\gamma_0, 0)$ would be a line of positive slope (rather than being vertical as observed in the current study). A similar prediction would follow from the hypothesis that the brightness mechanism was influenced by α, or, indeed, for any influence of α on the components of the B&B amalgam that does not cancel. Since we observe nearly perfect alignment of the major and minor axes of the threshold contours with the coordinate axes in Fig. 3, the most parsimonious conclusion is that neither the brightness nor the blackshot mechanism is differentially sensitive to the aspects of spatial structure controlled by α.

4.2.2. No detectable role for histogram statistics on a larger scale

The stimuli are designed so that variations in the isodipole-strength parameter α do not change the fraction of checks that are black or white. Thus, mechanisms sensitive to the histogram statistics of the luminances of individual checks cannot be affected by changes in α. However, it is possible that mechanisms sensitive to the histogram statistics of the checks also analyze the image in larger local chunks, e.g., via the analysis of the distribution of the total luminance in patches of various sizes. Since alterations in α affect local correlations, these alterations necessarily affect the distribution of the summed luminances in larger patches of the texture, e.g., those containing $m \times n$ checks. Mean and variance of these distributions are unaffected by α, because of the isodipole property. In principle, the blackshot mechanism of Chubb and coworkers (Chubb et al., 2004), which is sensitive to the fraction of the histogram distribution that is very close to black, might provide a discrimination cue were it able to operate on the statistics of these larger patches. However, as outlined below, the basic idea of the preceding section shows that this is not a likely contributor to our results.

Regions containing only black are encountered with increasing frequency as the value of γ becomes increasingly negative, but only if α is positive (i.e., even). This is because when $\alpha < 0$, the presence of three black checks within a 2×2 patch biases the fourth check to be white—thus reducing the likelihood of an all-black 2×2 patch. The effect of nonzero values of α on the probability of larger all-black patches is in the same direction, but larger—because the above phenomenon operates within each 2×2 sub-patch. Thus, the signature of a blackshot mechanism selective for larger patches is that discrimination would be enhanced specifically in the quadrant of the stimulus space in which $\alpha > 0$ and $\gamma < 0$. Since no such asymmetry is observed, either in the psychometric functions (Fig. 2) or in the manner in which the data points deviate from the iso-discrimination ellipses (Fig. 3), we conclude that there is no evidence for this kind of mechanism.

4.3. Foreground/background asymmetry

A salient feature of the data for three of our observers (AO, CFC, and MC) is that for textures $T(0, \alpha)$ tending toward even from IID ($\alpha > 0$), a structured target $T(0, \alpha)$ on a random background $T(0,0)$ is more difficult to localize than a random target $T(0,0)$ on a structured background $T(0, \alpha)$. For MC and AO, this pattern is reversed for textures tending toward odd from IID: that is, for $\alpha < 0$, a target $T(0, \alpha)$ on a background $T(0,0)$ is easier to localize than a target $T(0,0)$ on a background of $(0, \alpha)$. Thus, there appears to be a foreground-background asymmetry, whose signature depends on the signature of the fourth-order correlation α.

4.3.1. Role of figure/ground

To discuss the possible sources of this asymmetry, let us assume that discrimination of the textures $T(0, \alpha)$ from $T(0,0)$ is based on some statistic Q, for which $Q(T(0, \alpha))$ is a monotonic function of α. Since the target is much smaller than the rest of the array, we associate it with "figure", and we associate the remainder of the stimulus with "ground". It is of course possible that the visual system simply has a bias to construe regions higher in Q as ground. However, there are at least two

other possibilities that are less ad hoc. Previous authors (Gurnsey & Browse, 1989; Treisman & Gormican, 1988) have observed that if salience of some feature can be rated on a scale, then the component of an image in which this rating is further from zero has a bias to be figure, and the component closer to zero has a bias to be ground. Now consider the possibility that the observer's subjective notion of "random" is not located at $T(0,0)$, but rather, at $T(0,\alpha_{\mathrm{random}})$ for some $\alpha_{\mathrm{random}} > 0$. (That is, a texture that is objectively somewhat "even" appears subjectively random.) In this case, $T(0,\alpha)$ a would have a bias to appear as background if

$$|\, Q(T(0,\alpha)) - Q(T(0,\alpha_{\mathrm{random}}))\,| < |\, Q(T(0,0)) - Q(T(0,\alpha_{\mathrm{random}}))\,| \quad (12)$$

and as foreground under other circumstances. The effect would be maximal when $\alpha \approx \alpha_{\mathrm{random}}$, since at this point, $T(0,\alpha)$ is subjectively random, while $T(0,0)$ is subjectively "odd". Moreover, consistent with our observations, the effect would invert when $\alpha < 0$.

4.3.2. Role of induced long-range or higher-order correlations

Another possible contribution to the foreground-background asymmetry is that detection or measurement of positive values of Q is an accelerating function of the size of the region in which it is displayed. This size dependence might arise because, as area increases, there is more opportunity for long edge-like structures (Victor & Conte, 1991), or large blob-like structures, to be present. These induced features, rather than the four-check local statistics per se, appear to be important for the isodipole discrimination task. When an even texture serves as the target, it occupies only 25% of the display area, and there is less of a chance for such structures to be present, than when it serves as the background.

5. Summary

Despite the impressive ability of the visual system to make use of image statistics, much previous work (Caelli & Julesz, 1978; Caelli, Julesz, & Gilbert, 1978; Chubb & Landy, 1991; Julesz, Gilbert, Shepp, & Frisch, 1973; Julesz et al., 1978; Victor & Conte, 1991, 1996) shows that only specific kinds of image statistics are in fact used to discriminate textures. Laboratory analysis of these statistical cues is facilitated by considering histogram statistics and spatial correlation structure separately, as isolated in IID and isodipole stimuli. In naturally occurring textures, these analytically-convenient kinds of cues are combined, rather than present in isolation. Thus, one might suspect that the IID and isodipole paradigms have only limited relevance to natural textures. The present study encourages optimism that this is not the case. When IID and isodipole statistics are combined, they exert independent influences on texture discriminability: discrimination deviates at most slightly from a Euclidean combination rule, and probability summation of the two cues predicts psychophysical discriminability.

Acknowledgements

Portions of this work were presented at the 2003 meeting of the Vision Sciences Society, Sarasota, FL. This work was supported by NIH NEI EY7977. We thank Rick Gurnsey for several insights and helpful discussion.

Appendix A. Historical note: the even and odd isodipole textures

We recently learned from R. Gurnsey that the even and odd isodipole textures were introduced by Irwin Pollack in 1971 (Pollack, 1971a, 1971b, 1972, 1973), several years before the paper (Julesz et al., 1978) often cited in this regard. Pollack (Pollack, 1972, 1973) parametrically varied the textures (in a manner corresponding precisely to "propagated decorrelation", our parameter α) to obtain thresholds as a function of the strength of higher-order correlations. Pollack cites Julesz (1962) for the basic scheme for construction of two-dimensional Markov textures. However, although the scheme that Julesz (1962) introduced for construction of two-dimensional Markov textures *could* have produced the even and odd isodipole textures, Julesz did not pursue this route. One possible reason was the influence of a theorem of Rosenblatt and Slepian (1962), which indicated that *one-dimensional* textures produced in this fashion would necessarily be trivial.

Appendix B. Detailed description of textures and their construction

Here we describe how we create textures $T(\gamma,\alpha)$ specified by a luminance bias, parameterized by γ, and an even–odd (fourth-order isodipole) bias, parameterized by α. We then discuss some properties of, and calculate information-theoretic quantities from, these textures.

γ is the difference between the fraction of checks that are white and the fraction of checks that are black. $\gamma = 0$ specifies a texture that has an equal number of white and black checks; $\gamma = 1$ specifies a texture that has only white checks, and $\gamma = -1$ specifies a texture that has only black checks. α specifies the difference between the fraction 2×2 of blocks that contain an *even* number of white checks, and the fraction of 2×2 blocks that contain an *odd* number of white checks. Thus, $\alpha = 1$ specifies a texture in which all 2×2 blocks contain an *even*

number of white checks, and $\alpha = -1$ specifies a texture in which all 2×2 blocks contain an *odd* number of white checks.

To provide a rationale for the strategy used, we first point out that this construction is not a trivial matter. Naïvely, one might hope that textures corresponding to $\alpha = 0$ (no fourth-order correlations) but $\gamma \neq 0$ (luminance bias) would be some kind of IID texture. But this is not the case, for the following reason. In an IID texture that conforms to a particular value of γ, $\frac{1}{2}(1+\gamma)$ of the checks are white, and $\frac{1}{2}(1-\gamma)$ of the checks are black. Since (by the definition of IID) all of the checks of a 2×2 block are determined independently, the fraction of 2×2 blocks with n white checks (denoted w_n, for $n = 0, 1, 2, 3, 4$) is readily determined

$$w_n = \binom{4}{n} \left(\frac{1-\gamma}{2}\right)^{4-n} \left(\frac{1+\gamma}{2}\right)^{n} \tag{13}$$

Here, $\binom{m}{n} \equiv \frac{m!}{n!(m-n)!}$ is the binomial coefficient, namely, the number of ways of choosing n distinguishable objects (here, the checks that are colored white) out of m (here, 4, the number of checks in a 2×2 block). The value of α is determined by the imbalance between the fraction of 2×2 blocks with an even number of white checks ($w_0 + w_2 + w_4$) and the fraction of 2×2 blocks with an odd number of white checks ($w_1 + w_3$). That is

$$\alpha = \sum_{n=0}^{4} (-1)^n w_n = \gamma^4,$$

where the final equality follows from Eq. (13) via straightforward algebra. The above analysis shows that the introduction of low-order statistical structure ($\gamma \neq 0$) in the naïve fashion results in alteration of high-order statistics ($\alpha \neq 0$). Thus, it is not clear how to achieve a texture with arbitrarily specified values of α and γ. Moreover, if it is possible to achieve a particular (γ, α)-pair, then it is likely that there is more than one texture that does so. These textures will differ in long-range correlations and/or correlations of order <4, that are not explicitly specified by α and γ. It is not immediately clear whether, among these possibilities, any particular texture is the most "natural."

For these reasons, we appeal to the maximum-entropy formalism (Zhu, Wu, & Mumford, 1998) for textures. Informally, a "maximum-entropy" texture is a texture for which certain statistics are specified explicitly (the "constraints"), and higher-order statistics are determined by maximizing the entropy of the process that generates the texture, subject to these constraints. The rationale for maximum entropy (see (Zhu et al., 1998) for further details) is that this procedure does not introduce correlations of higher order or longer range than the constraints, except as is necessary to conform to the constraints. The maximum-entropy formalism also has advantages for modeling, in that the only information available to the observer is contained in the constraints. However (Zhu et al., 1998), maximum-entropy constructions involving statistics of order greater than 2 (as in the present case) are not automatic or straightforward.

The use of maximum-entropy distributions as a way to study interactions of statistics of various orders is also motivated by the theory of information geometry (Amari, 2001). A pivotal result of this theory is that maximum entropy distributions combine constraints of different orders without introducing new dependencies. Thus, the approach presented here is a natural one for studying interaction of texture statistics in general. However, there is no guarantee that maximum entropy distributions can always be constructed as explicitly as in the present case.

For our purposes, the constraints are statements about the probabilities of the various 2×2 blocks and subsets thereof. Most of the constraints involve γ, the parameter that determines the luminance bias. We require that the probability of a white check (in any position) is

$$p_\gamma = \tfrac{1}{2}(1+\gamma) \tag{14}$$

We also require second-order constraints, so that the resulting textures will be isodipole. In particular, we require that the probability that any designated pair of checks is white is p_γ^2. This, along with the fact that every check pair either contains two white checks, one white and one black check, or two black checks, determines the rest of the probabilities for two-check blocks: $p_\gamma(1-p_\gamma)$ that a designated check is white and the other check is black, and $(1-p_\gamma)^2$ that both checks are black. Together, these second-order constraints guarantee that the second-order statistics of the texture to be generated match an IID texture based on p_γ and $1-p_\gamma$, i.e., the texture is locally isodipole. (The maximum entropy formalism will then guarantee that the texture is globally isodipole). We also require that the local third-order statistics match those of an IID texture. The rationale for this is that we do not want our construction to create third-order structure, since no third-order structure is present in the IID textures or the isodipole textures $\gamma = 0, \alpha \neq 0$. Thus, within a 2×2 block, we require that the probability that any designated triplet of checks is white is p_γ^3, that the probability that any designated two checks are white and one is black is $p_\gamma^2(1-p_\gamma)$, that the probability that any designated two checks are black and one is white is $p_\gamma(1-p_\gamma)^2$, and that the probability that any designated triplet of checks is black is $(1-p_\gamma)^3$. The third-order constraints imply the second-order constraints, and the second-order constraints imply the first-order constraints. This is because higher-order probabilities determine lower-order probabilities. For example, every pair of checks can be extended to

a triple by either adding a white check or a black check. Thus, the probability of a white pair must be equal to the probability of a white triplet, plus the probability that two designated checks are white, and the third is black. To these constraints involving γ, we now add a single constraint that specifies the fourth-order correlation structure α: the difference between the probability that the number of white checks in a 2×2 block is even, and the probability that the number of white checks in a 2×2 block is odd, must be α.

We next observe that all of these constraints are satisfied by assigning the 2×2 block probabilities as follows:

$$g_h(\gamma,\alpha)=\frac{1}{16}((1-\gamma)^{4-h}(1+\gamma)^h+(-1)^h(\alpha-\gamma^4)), \quad (15)$$

where h is the number of checks within the 2×2 block that are white. To see that the constraints related to γ are satisfied, we observe from Eq. (15) that

$$\begin{aligned} &g_{h+1}(\gamma,\alpha)+g_h(\gamma,\alpha)\\ &\quad=\frac{(1-\gamma)^{3-h}(1+\gamma)^{h+1}+(1-\gamma)^{4-h}(1+\gamma)^h}{16}\\ &\quad=\frac{(1-\gamma)^{3-h}(1+\gamma)^h}{8}=(1-p_\gamma)^{3-h}p_\gamma^h, \end{aligned}$$

where the final equality follows from Eq. (14). The left hand side indicates the decomposition of a third-order statistic into a pair of fourth-order statistics. For example, with $h=3$, it corresponds to the probability of a white triple at a designated location within the 2×2 block, since this is equal to the probability that all four checks are white, $g_4(\gamma,\alpha)$, plus the probability that three designated checks are white and one is black, $g_3(\gamma,\alpha)$. The right hand side indicates that the assigned values for g (Eq. (15)) conform to the desired constraints.

To see that the constraint related to α is satisfied by Eq. (15), we note that the difference between the probability that a 2×2 block contains an even number of white checks, and the probability that it contains an odd number of white checks, is given by

$$\sum_{h=0}^{4}(-1)^h\binom{4}{h}g_h(\gamma,\alpha), \quad (16)$$

since there are $\binom{4}{h}$ configurations of 2×2 blocks that have exactly h white checks. Substitution of the value of Eq. (15) for $g_h(\gamma,\alpha)$ into Eq. (16) yields identically α. Moreover, as a general consequence of the properties of the maximum-entropy approach, the solution represented by Eq. (15) must be unique.

Provided that

$$\gamma^4-(1-|\gamma|)^4\leqslant\alpha\leqslant\gamma^4+(1-|\gamma|)^4 \quad (17)$$

all of the expressions on the right hand side of Eq. (15) are in the range $[0,1]$. (This follows from the observation that the extrema of the first term of Eq. (15) must occur at $h=0$ or $h=4$. When equality holds, one or more of the $g_h(\gamma,\alpha)$ equal zero, but no two consecutive values are zero.)

For (γ,α)-pairs within the limits specified by Eq. (17), we next create a Markov random field in which the 2×2 blocks have the probabilities specified by Eq. (15). In the Markov random field formalism of Zhu et al. (1998), the "neighborhood" associated with a given check will be the eight checks that are adjacent to it, or make contact with a corner, and the "cliques" consist of (a) all the isolated checks, and (b) all 2×2 blocks. This Markov random field will necessarily be a maximum-entropy texture on the entire plane, constrained by the above 2×2 block probabilities.

The Markov random field generation process is recursive. The recursion is initialized by assigning the colors of the checks in the first row and first column according to p_γ (Eq. (14)). To generate the interior of the texture, proceed as follows. Suppose that the check at location (i,j) is to be assigned, and that checks at locations $(i-1,j)$,$(i,j-1)$, and $(i-1,j-1)$ have already been assigned. If h of these checks are white, then completion of the 2×2 block with an additional check will lead to a block with either h white checks (if the check at (i,j) is colored black), or to a block with $h+1$ white checks (if the check at (i,j) is colored white). Since the unconditional probabilities of these blocks must have the ratio $\frac{g_h(\gamma,\alpha)}{g_{h+1}(\gamma,\alpha)}$ and must sum to $g_h(\gamma,\alpha)+g_{h+1}(\gamma,\alpha)$, it follows that the conditional probability that a check is colored white, given that h of the remaining three checks in its block are also white, is given by

$$\begin{aligned} &P\{(i,j)\text{is white} \mid h \text{ of } (i,j-1),(i-1,j),\\ &\quad(i-1,j-1) \text{ are white}\}=\frac{g_{h+1}(\gamma,\alpha)}{g_h(\gamma,\alpha)+g_{h+1}(\gamma,\alpha)}. \end{aligned} \quad (18)$$

According to the remarks following Eq. (17), the denominator is always positive, so the propagation rule is always well-defined.

Textures generated in this fashion have isotropic statistics, even though the above method for generating them proceeds in a directional fashion (e.g., from left to right, and top to bottom). The reason for this is as follows. Since the block probabilities are symmetric, the same generation process could have been used, but proceeding in a different direction. This alternative generation procedure would necessarily have the same statistics, since it too would be a Markov random field, and hence, maximum-entropy with the same constraints—because of the uniqueness of maximum-entropy processes with identical constraints (Zhu et al., 1998). A similar argument demonstrates that the checks along any row, or any column, are all independently colored. This is because any such strip could be used to initialize

the texture, with the rule (18) used to propagate the texture in both directions from the initializing strip. Were it the case that such textures had statistics that are distinct from the one described above, the uniqueness of maximum-entropy processes would again be violated.

Finally, we note that for $\gamma = 0$, the full range α can be realized (Eq. (17)). Textures generated in this fashion are generalizations of the "propagated" decorrelated isodipole textures that have been previously studied (Pollack, 1971a, 1971b, 1972, 1973; Victor, 1985; Victor & Conte, 1989).

B.1. Kullback–Leibler divergences

The Kullback–Leibler divergence $K(P\|Q)$ is a natural measure of the extent to which a probability distribution P can be distinguished from a probability distribution Q. More precisely, (Cover & Thomas, 1991; Latham, 2002), assume that P and Q are two probability distributions on a set of elements $X = \{x_j\}$, with p_j representing the frequency of x_j in P, and q_j representing the frequency of x_j in Q. An observer, who expects that a stream of observations is drawn from Q, is given the task of detecting when the source of the observations is switched to P. Under these circumstances, the number of bits per observation is given by the Kullback–Leibler divergence of P and Q, $K(P\|Q)$, namely,

$$K(P\|Q) = \sum_j p_j \log_2\left(\frac{p_j}{q_j}\right). \tag{19}$$

Strictly speaking, this definition applies only to a finite set of elements X. For probability distributions on infinite sets (such as the set of textures on an infinite lattice), the Kullback–Leibler divergence is expected to grow without bound. However, as we show here, this definition has a natural extension to textures based on Markov random fields based on block neighborhoods (Zhu et al., 1998). More precisely, for such textures, the Kullback–Leibler divergence between two textures is asymptotically proportional to their area. This proportionality constant, which depends on the textures, thus indicates the extent to which an ideal observer can discriminate these textures, per unit area. Moreover, for Markov random fields, the Kullback–Leibler divergence (per unit area) can be readily calculated from the 2×2 block probabilities.

The key observation concerns a one-dimensional analogue of Markov random fields, namely, Markov chains. A Markov chain on r symbols $S = \{s_1, s_2, \ldots, s_r\}$ is defined by a transition matrix M. M is a $r \times r$ matrix whose typical element m_{jk} is the conditional probability that the symbol s_j is followed by the symbol s_k. Entries in M are necessarily within the interval $[0,1]$, and each row of M sums to 1 (since each symbol must have a successor). Moreover, M has a unique left (row) eigenvector $\vec{m} = (m_1^*, m_2^*, \ldots, m_r^*)$ whose entries sum to 1, and for which $\vec{m}M = \vec{m}$. This eigenvector is the probability distribution that is the stable result of repeated applications of M.

We now consider the quantity

$$U(P_n\|Q_n) = \sum_{x \in W_n} P_n(x) \log_2 Q_n(x), \tag{20}$$

where W_n is the set of words composed of n symbols drawn from S, and the probabilities $P_n(x)$ and $Q_n(x)$ indicate the probabilities assigned by Markov chains P and Q to the word x. For example, if x is the word $s_{j_1} s_{j_2} \ldots s_{j_n}$, then

$$P_n(x) = p_{j_1} P_{j_1 j_2} P_{j_2 j_3} \ldots P_{j_{n-1} j_n} \tag{21}$$

and similarly for $Q_n(x)$.

To establish the behavior of the sum (20), we note that for a Markov chain, the probability of a word x can be determined recursively from three sub-words: the word Lx, obtained by retaining the $n-1$ leftmost (first) elements of x, the word Rx, obtained by retaining the $n-1$ rightmost (last) element of x, and the word LRx (or RLx), obtained by retaining the $n-2$ middle elements of x. The recursive relationship is

$$P_n(x) = \frac{P_{n-1}(Lx) P_{n-1}(Rx)}{P_{n-2}(LRx)}. \tag{22}$$

This follows from Eq. (21), since $P_{n-1}(Lx) = p_{j_1} P_{j_1 j_2} \ldots P_{j_{n-2} j_{n-1}}$, $P_{n-1}(Rx) = p_{j_2} P_{j_2 j_3} \ldots P_{j_{n-1} j_n}$, and $P_{n-2}(LRx) = p_{j_2} P_{j_2 j_3} \ldots P_{j_{n-2} j_{n-1}}$. We now use Eq. (22) and the parallel expression for $Q_n(x)$ to calculate (for $n \geqslant 2$):

$$\begin{aligned} U(P_n\|Q_n) &= \sum_{x \in W_n} \frac{P_{n-1}(Lx) P_{n-1}(Rx)}{P_{n-2}(LRx)} \\ &\quad \times \log_2 \frac{Q_{n-1}(Lx) Q_{n-1}(Rx)}{Q_{n-2}(LRx)} \\ &= \sum_{x \in W_n} \frac{P_{n-1}(Lx) P_{n-1}(Rx)}{P_{n-2}(LRx)} \log_2 Q_{n-1}(Lx) \\ &\quad + \sum_{x \in W_n} \frac{P_{n-1}(Lx) P_{n-1}(Rx)}{P_{n-2}(LRx)} \log_2 Q_{n-1}(Rx) \\ &\quad - \sum_{x \in W_n} \frac{P_{n-1}(Lx) P_{n-1}(Rx)}{P_{n-2}(LRx)} \log_2 Q_{n-2}(LRx). \end{aligned} \tag{23}$$

Each of these terms simplifies in a similar manner. We consider the first term. A summation over all words x of n symbols can be considered as an outer sum over all words Lx of length $n-1$, and an inner sum over all possible final symbols x_{j_n}. The product $P_{n-1}(Lx) \log_2 Q_{n-1}(Lx)$ is independent of the choice of the final symbol x_{j_n}, since it is deleted by L. Because of the construction of the Markov chain, the quotient $\frac{P_{n-1}(Rx)}{P_{n-2}(LRx)} = P_{j_{n-1} j_n}$. Since each row of the transition matrix $P_{j_{n-1} j_n}$ must sum to 1, we find

$$\sum_{x\in W_n}\frac{P_{n-1}(Lx)P_{n-1}(Rx)}{P_{n-2}(LRx)}\log_2 Q_{n-1}(Lx)$$
$$=\sum_{x\in W_{n-1}}P_{n-1}(Lx)\log_2 Q_{n-1}(Lx)$$
$$=U(P_{n-1}\|Q_{n-1}). \quad (24)$$

Similar reasoning leads to

$$\sum_{x\in W_n}\frac{P_{n-1}(Lx)P_{n-1}(Rx)}{P_{n-2}(LRx)}\log_2 Q_{n-1}(Rx)$$
$$=\sum_{Rx\in W_{n-1}}P_{n-1}(Rx)\log_2 Q_{n-1}(Rx)=U(P_{n-1}\|Q_{n-1}) \quad (25)$$

and, by writing

$$\frac{P_{n-1}(Lx)P_{n-1}(Rx)}{P_{n-2}(LRx)}\log_2 Q_{n-2}(LRx)$$
$$=\left(\frac{P_{n-1}(Lx)}{P_{n-2}(LRx)}\right)\left(\frac{P_{n-1}(Rx)}{P_{n-2}(LRx)}\right)P_{n-2}(LRx)\log_2 Q_{n-2}(LRx), \quad (26)$$

to

$$\sum_{x\in W_n}\frac{P_{n-1}(Lx)P_{n-1}(Rx)}{P_{n-2}(LRx)}\log_2 Q_{n-2}(LRx)$$
$$=\sum_{LRx\in W_{n-2}}P_{n-2}(LRx)\log_2 Q_{n-1}(LRx)=U(P_{n-2}\|Q_{n-2}). \quad (27)$$

Substitution of the above three results (Eqs. (24), (25), (27)) into Eq. (23) yields

$$U(P_n\|Q_n)=2U(P_{n-1}\|Q_{n-1})-U(P_{n-2}\|Q_{n-2}). \quad (28)$$

Eq. (28) states that for successive values of n, the sequence $U(P_n\|Q_n)$ forms an arithmetic progression with constant difference $U(P_2\|Q_2)-U(P_1\|Q_1)$. In particular, it follows that

$$\lim_{n\to\infty}\frac{U(P_n\|Q_n)}{n}=U(P_2\|Q_2)-U(P_1\|Q_1). \quad (29)$$

The right hand side can be calculated explicitly from the probabilities of words of length 1 and 2, respectively.

To extend this result to the two-dimensional context of Markov random fields, we consider the quantity

$$U(P_{n_1,n_2}\|Q_{n_1,n_2})=\sum_{x\in B_{n1,n2}}P_{n_1,n_2}(x)\log_2 Q_{n_1,n_2}(x). \quad (30)$$

This is analogous to Eq. (20), but here, B_{n_1,n_2} represents the set of $n_1\times n_2$ blocks, and $P_{n_1,n_2}(x)$ and $Q_{n_1,n_2}(x)$ represent the probabilities assigned to the block x by the Markov random fields P and Q. We now allow one of the block dimensions (say, n_2) to approach infinity, while the other block dimension remains finite. A block of size $n_1\times n_2$ can be construed as a sequence of n_2 blocks, each of size $n_1\times 1$. That is, with n_1 held fixed, blocks of size $n_1\times n_2$ can be considered to constitute a *one*-dimensional Markov chain, whose *symbols* are the blocks of size $n_1\times 1$. The one-dimensional result (Eq. (29)) therefore applies, and we find

$$\lim_{n_2\to\infty}\frac{U(P_{n_1,n_2}\|Q_{n_1,n_2})}{n_2}=U(P_{n_1,2}\|Q_{n_1,2})-U(P_{n_1,1}\|Q_{n_1,1}). \quad (31)$$

We next consider the limiting behavior of this expression as n_1 approaches infinity. The one-dimensional result (Eq. (29)) can now be applied to the right-hand side of Eq. (31) along the first dimension, yielding

$$\lim_{n_1\to\infty}\lim_{n_2\to\infty}\frac{U(P_{n_1,n_2}\|Q_{n_1,n_2})}{n_1n_2}$$
$$=U(P_{2,2}\|Q_{2,2})-U(P_{1,2}\|Q_{1,2})-U(P_{2,1}\|Q_{2,1})$$
$$+U(P_{1,1}\|Q_{1,1}). \quad (32)$$

Note that since the right-hand side of this result is symmetric in the dimensions, the order of limits on the left does not matter. For $P=Q$, the above expression is the entropy per pixel of the texture P.

Finally, since the Kullback–Leibler divergence of P and Q, $K(P\|Q)$ (Eq. (19)) is related to the quantity $U(P\|Q)$ (Eq. (20)) by

$$K(P_{n_1,n_2}\|Q_{n_1,n_2})=U(P_{n_1,n_2}\|P_{n_1,n_2})-U(P_{n_1n_2}\|Q_{n_1,n_2}), \quad (33)$$

we find

$$\lim_{n_1\to\infty}\lim_{n_2\to\infty}\frac{K(P_{n_1,n_2}\|Q_{n_1,n_2})}{n_1n_2}$$
$$=K(P_{2,2}\|P_{2,2})-K(P_{1,2}\|P_{1,2})-K(P_{2,1}\|P_{2,1})$$
$$+K(P_{1,1}\|P_{1,1}). \quad (34)$$

This equation, along the probabilities of Eq. (15) that characterized our stimuli (i.e., using $g_h(\gamma_p,\alpha_p)$ for P and $g_h(\gamma_Q,\alpha_Q)$ for Q), generated the graph of Fig. 4. For convenience, we denote this quantity, which represents the discriminability per unit pixel, by K_0:

$$K_0(P,Q)=\lim_{n_1\to\infty}\lim_{n_2\to\infty}\frac{K(P_{n_1,n_2}\|Q_{n_1,n_2})}{n_1n_2}. \quad (35)$$

For probability distributions that are nearly unbiased, the expressions (20) and (30) for $U(P\|Q)$ have a simple approximate form, which in turn provides for a simple approximation for Eq. (34). Let $P=\frac{1}{N}+\delta P$ and $Q=\frac{1}{N}+\delta Q$, where N is the number words or blocks. The approximation $\log_2(1+u)=\frac{1}{\ln 2}(u-\frac{u^2}{2}+\cdots)$ leads to

$$U(P\|Q)\approx\frac{1}{2N\ln 2}\left(-\sum(\delta Q)^2+2\sum(\delta P)(\delta Q)\right) \quad (36)$$

and, via Eq. (33), to

$$K(P\|Q)\approx\frac{1}{2N\ln 2}\sum(\delta P-\delta Q)^2, \quad (37)$$

where the errors in the approximations are at most third-order in δP and δQ.

The analysis up to this point has exploited the assumption that P and Q are Markov random fields, but has not made use of the particular assignments for block probabilities. We now consider the particular case that the block probabilities for P and Q are given by the rule (15), (i.e., $g_h(\gamma_P, \alpha_P)$ determines P and $g_h(\gamma_Q, \alpha_Q)$ determines Q). In this case, combining Eqs. (15), (34) and (37) leads to

$$K_0((\gamma_P, \alpha_P), (\gamma_Q, \alpha_Q)) = \lim_{n_1 \to \infty} \lim_{n_2 \to \infty} \frac{K(P_{n_1,n_2} \| Q_{n_1,n_2})}{n_1 n_2} \approx \frac{1}{2 \ln 2} ((\gamma_P - \gamma_Q)^2 + (\alpha_P - \alpha_Q)^2), \quad (38)$$

where, as in Eqs. (36) and (37), the error in the approximation is third-order in the deviations of α and γ from zero. Thus, as suggested by Fig. 4, contour lines of the Kullback–Leibler divergence are expected to be circular when α and γ are small.

References

Amari, S.-I. (2001). Information geometry on hierarchy of probability distributions. *IEEE Transactions of Information Theory, 47*, 1701–1711.

Barlow, H. B. (1978). The efficiency of detecting changes of density in random dot patterns. *Vision Research, 18*, 637–650.

Beck, J., Sutter, A., & Ivry, R. (1987). Spatial frequency channels and perceptual grouping in texture segregation. *Computer Vision, Graphics, and Image Processing, 37*, 299–325.

Bergen, J. R., & Adelson, E. H. (1988). Early vision and texture perception. *Nature, 333*, 363–364.

Bergen, J. R., & Landy, M. (1991). The computational modeling of visual texture segregation. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 253–271). Cambridge, MA: MIT Press.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.

Caelli, T., & Julesz, B. (1978). On perceptual analyzers underlying visual texture discrimination: Part I. *Biological Cybernetics, 28*(3), 167–175.

Caelli, T., Julesz, B., & Gilbert, E. (1978). On perceptual analyzers underlying visual texture discrimination: Part II. *Biological Cybernetics, 29*(4), 201–214.

Chubb, C., Econopouly, J., & Landy, M. S. (1994). Histogram contrast analysis and the visual segregation of IID textures. *Journal of Optical Society America A, 11*(9), 2350–2374.

Chubb, C., & Landy, M. (1991). Orthogonal distribution analysis: A new approach to the study of texture perception. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 291–301). Cambridge, MA: MIT Press.

Chubb, C., Landy, M. S., & Econopouly, J. (2004). A visual mechanism tuned to black. *Vision Research, 44*(27), 3223–3232.

Cover, T. M., & Thomas, J. A. (1991). Elements of information theory. In *Wiley series in telecomunications* (pp. 52). New York: Wiley.

Efron, B., & Tibshirani, R. J. (1998). In *An Introduction to the Bootstrap. Monographs on statistics and applied probability* (57, pp. 436). Boca Raton, FL: Chapman & Hall/CRC Press.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local association field. *Vision Research, 33*(2), 173–193.

Graham, N. (1989). *Visual pattern analyzers*. Oxford: Clarendon Press.

Graham, N., Beck, J., & Sutter, A. (1992). Nonlinear processes in spatial-frequency channel models of perceived texture segregation: Effects of sign and amount of contrast. *Vision Research, 32*(4), 719–743.

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics, 38*(2), 141–171.

Gurnsey, R., & Browse, R. A. (1989). Asymmetries in visual texture discrimination. *Spatial Vision, 4*(1), 31–44.

Julesz, B. (1962). Visual pattern discrimination. *IRE Transactions of Information Theory, IT-8*, 84–92.

Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature, 290*(5802), 91–97.

Julesz, B., Gilbert, E. N., Shepp, L. A., & Frisch, H. L. (1973). Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited. *Perception, 2*(4), 391–405.

Julesz, B., Gilbert, E. N., & Victor, J. D. (1978). Visual discrimination of textures with identical third-order statistics. *Biological Cybernetics, 31*(3), 137–140.

Knill, D. C. (1998). Surface orientation from texture: Ideal observers, generic observers and the information content of texture cues. *Vision Research, 38*(11), 1655–1682.

Latham, P. E. (2002). Available: <http://culture.neurobio.ucla.edu/pel/critics/bound.ps>. Department of Neurobiology, UCLA, Los Angeles, CA.

Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of Optical Society America A, 7*(5), 923–932.

Poirson, A. B., & Wandell, B. A. (1990). The ellipsoidal representation of spectral sensitivity. *Vision Research, 30*(4), 647–652.

Pollack, I. (1971a). Detection of one-, two-, and three-dimensional Markov constraints in visual displays. *Acta Psychologica, 35*, 219–232.

Pollack, I. (1971b). Perception of two-dimensional Markov constraints within visual displays. *Perception and Psychophysics, 9*, 461–464.

Pollack, I. (1972). Visual discrimination thresholds for one- and two-dimensional Markov spatial constraints. *Perception and Psychophysics, 12*(2A), 161–167.

Pollack, I. (1973). Discrimination of third-order Markov constraints within visual displays. *Perception and Psychophysics, 13*(2), 276–280.

Portilla, J., & Simoncelli, E. P. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision, 40*(1), 49–71.

Purpura, K. P., Victor, J. D., & Katz, E. (1994). Striate cortex extracts higher-order spatial correlations from visual textures. *Proceedings of the National Academy of Sciences, USA, 91*(18), 8482–8486.

Rosenblatt, M., & Slepian, D. (1962). *N*th order Markov chains with any set of *N* variables independent. *Journal of Society of Industry of Applied Mathematics, 10*, 537–549.

Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95*(1), 15–48.

Victor, J. D. (1985). Complex visual textures as a tool for studying the VEP. *Vision Research, 25*(12), 1811–1827.

Victor, J. D. (1986). Isolation of components due to intracortical processing in the visual evoked potential. *Proceedings of the National Academy of Sciences, USA, 83*(20), 7984–7988.

Victor, J. D., & Conte, M. M. (1989). Cortical interactions in texture processing: Scale and dynamics. *Visual Neuroscience, 2*(3), 297–313.

Victor, J. D., & Conte, M. M. (1991). Spatial organization of nonlinear interactions in form perception. *Vision Research, 31*(9), 1457–1488.

Victor, J. D., & Conte, M. M. (1996). The role of high-order phase correlations in texture processing. *Vision Research, 36*(11), 1615–1631.

Weibull, W. (1951). A statistical distribution function of wide applicability. *Journal of Applied Mechanics, 18*, 292–297.

Wilson, H. R. (1993). Nonlinear processes in visual pattern discrimination. *Proceedings of the National Academy of Sciences, USA, 90*(21), 9785–9790.

Zhu, S. C., Wu, Y., & Mumford, D. (1998). Filters, random fields and maximum entropy (FRAME): Towards a unified theory for texture modeling. *International Journal of Computer Vision, 27*(2), 107–126.