

A visual mechanism tuned to black

Charles Chubb^a, Michael S. Landy^{b,*}, John Econopoulou

^a Department of Cognitive Sciences & Institute for Mathematical Behavioral Sciences, University of California at Irvine, USA

^b Department of Psychology and Center for Neural Science, New York University, 6 Washington Place, 8th Floor, Room 961, New York, NY 10003, USA

Received 21 July 2003; received in revised form 9 July 2004

Abstract

Chubb et al. [Journal of the Optical Society of America A 11 (1994) 2350] investigated preattentive discrimination of achromatic textures comprising random mixtures of 17 Weber contrasts ranging linearly from -1 to 1 . They showed that only a single mechanism B is used to discriminate between textures whose histograms are equated in mean and in variance. Although they provided a partial characterization of B , their methods did not allow them to measure the sensitivity of B to texture mean and variance. Here, additional measurements are performed to complete the functional characterization of B . The results reveal that B (i) is strongly activated by texture elements of the lowest contrast (near -1), (ii) is slightly activated by texture elements of contrast -0.875 , and (iii) barely distinguishes the 15 contrasts ranging from -0.75 all the way up to 1 . To reflect the sharpness of its tuning to very dark, sparse elements in a predominantly bright scene, we call B the *blackshot* mechanism.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Texture perception; Texture segregation; Brightness

1. Introduction

The pioneering research of Julesz (1962, 1975, 1981), Julesz and Bergen (1983), Julesz, Gilbert, Shepp, and Frisch (1973), Julesz, Gilbert, and Victor (1978), Beck (1966, 1982) and Beck, Prazdny, and Rosenfeld (1983) revealed that human vision incorporates systems that operate preattentively to segment the visual field based not only on color and brightness differences, but also based on texture differences. These demonstrations gave rise to various closely related models of texture segregation (e.g., Bergen & Adelson, 1988; Bergen & Landy, 1991; Bovik, Clark, & Geisler, 1990; Caelli, 1985; Fogel & Sagi, 1989; Graham, 1991; Grossberg & Mingolla, 1985; Knutsson & Granlund, 1983; Landy & Bergen,

1991; Malik & Perona, 1990). All of these “back pocket” models (Chubb & Landy, 1991) proposed that preattentive texture segregation occurs in two main stages, a *measurement* stage followed by a *surveying* stage.

In the measurement stage the visual system applies to its input a battery of spatially local image transformations that we shall call *mechanisms*. In the surveying stage, boundaries are constructed between visual field regions that differ significantly in the activations they produce in one or more mechanisms.

We focus here on achromatic textures of the sort shown in Fig. 1. Each of these textures is a randomly scrambled array of squares conforming to a specified luminance histogram. In Section 3 we describe evidence suggesting that (1) the number of mechanisms used to discriminate such textures is three, (2) two of these are sensitive primarily to the first and second moments of the texture histogram, and (3) the third mechanism, B , is sensitive to many higher-order moments. Chubb,

* Corresponding author. Tel.: +1 212 998 7857; fax: +1 212 995 4349.

E-mail address: landy@nyu.edu (M.S. Landy).




Fig. 1. Examples of texture stimuli. Each texture is a scramble S_{U+m} . That is, it has a gray level histogram $U + m$, shown as an inset in each panel, where U is the uniform distribution, and m is a histogram modulator. Textures for eight different modulators are shown: (a) $m = \lambda_1$, (b) $m = -\lambda_1$, (c) $m = \lambda_2$, (d) $m = -\lambda_2$, (e) $m = \lambda_3$, (f) $m = -\lambda_3$, (g) $m = \lambda_4$, (h) $m = -\lambda_4$. Note that the textures in e–h are readily discriminable, yet they all have histograms with the same mean and variance as U .

Econopoulou, and Landy (1994) partially characterized B 's sensitivity to different texture element (texel) contrasts. Here, we finish the job.

2. Preliminaries

2.1. Texture patches

Our stimuli are arrays of small square texels painted with values from C , the set of 17 Weber contrasts $-1, -7/8, \dots, 7/8, 1$. The term “histogram” usually refers to a function that gives the number of pixels in an image that take any given value. We deviate from this usage in calling any probability distribution on C a *histogram*. U denotes the uniform histogram: i.e., $U(c) = 1/17$ for all $c \in C$.

To generate a texture patch comprising N texels, we first select a histogram p . Then we load a virtual urn with N contrasts from C in proportions conforming as closely as possible to p and draw N times from the urn *without replacement* to assign contrasts to texels. The resulting patch S_p is called the *scramble* with histogram

p . Thus, S_U is the scramble with equal proportions of all 17 contrasts.¹

2.2. Modulators and modulation spaces

A function $m : C \rightarrow \mathfrak{R}$ is called a *modulator* if each of $U + m$ and $U - m$ is a probability distribution (i.e., is non-negative for all $c \in C$ and sums to 1). In addition m is called *maximal* if $\max(|m|) = 1/17$.

In our experiments, subjects discriminate between scrambles S_{U+m} and S_{U-m} for various modulators m . Often, as a shorthand for saying that an observer correctly discriminates S_{U+m} from S_{U-m} with probability ρ we will instead say that “The observer discriminates m with success rate ρ ” or, if ρ is large, “ m is easily discriminable.”

The set of functions spanned by N linearly independent modulators is called a *modulation space* of dimension N . (Note that any modulation space contains many functions g for which $\max(g) > 1/17$ and thus are not modulators.) Ω denotes the 16-dimensional space containing all modulators.

2.3. Projection

Let $\phi_1, \dots, \phi_N$ be an orthogonal basis of Φ , a subspace of Ω . For any function $f : C \rightarrow \mathfrak{R}$, the *projection* of f into Φ is

$$\text{Proj}_\Phi(f) = \sum_{i=1}^N W_i \phi_i, \quad (1)$$

where

$$W_i = \frac{\phi_i \cdot f}{\phi_i \cdot \phi_i}, \quad i = 1, \dots, N. \quad (2)$$

2.4. Some distinguished modulators

The maximal modulators $\lambda_1, \dots, \lambda_7$ (Fig. 2) play a central role here. These modulators (discrete domain analogues of the Legendre polynomials) are orthogonal. Moreover, λ_i is an i th-order polynomial, for $i = 1, \dots, 7$; thus, $\lambda_1, \dots, \lambda_i$ span the space of all i th-order polynomial modulators. Fig. 1 shows the scrambles S_{U+m} and S_{U-m} for $m = \lambda_1, \dots, \lambda_4$.

¹ Chubb et al. (1994) used independent, identically distributed (IID) textures, whose texel values are jointly independent random contrasts identically distributed as p . Unlike a scramble, an IID texture patch will have different contrasts in proportions that deviate randomly from p . We use scrambles to eliminate uncontrolled histogram variability, although with the use of textures such as ours with a reasonably large number of samples, the difference in the accuracy of the measurements is likely to be quite small.




Fig. 2. The maximal modulators $\lambda_1, \dots, \lambda_7$.

2.5. Modeling discrimination processes

We adapt a standard model (Graham, 1989) to describe scramble discrimination. Specifically, we assume human vision comprises some number, N_M , of mechanisms, $M_1, M_2, \dots, M_{N_M}$, that are differentially sensitive to different scrambles. We write $M_i(S_p)$ for the expectation of the space-average activation produced in M_i by S_p . We assume that the channel applies a static non-linearity f_i to each texel's contrast, and thus

$$M_i(S_p) = \sum_{c \in C} f_i(c) p(c) = f_i \cdot p. \quad (3)$$

We call f_i the *impact function* of M_i ; $f_i(c)$ gives the mean impact exerted by a texel of contrast c in M_i (see Chubb, 1999; Chubb et al., 1994).

For one of our stimuli, the task requires the subject to discriminate textures S_{U+m} and S_{U-m} . We assume the probability of success in such a discrimination is a function of the differential activation produced in each mechanism. For a given mechanism M_i , the differential activation is

$$\Delta_i(m) = M_i(U+m) - M_i(U-m) = 2f_i \cdot m. \quad (4)$$

We assume that the success rate is a monotonic function of the Minkowski length of the vector $(\Delta_1(m), \Delta_2(m), \dots, \Delta_{N_M}(m))$. That is, for some strictly increasing psychometric function Ψ and some exponent θ ,

$$\Pr[\text{discriminating } S_{U+m} \text{ from } S_{U-m}] = \Psi(\text{Sal}(m)), \quad (5)$$

where

$$\begin{aligned} \text{Sal}(m) &= \left(\sum_{i=1}^{N_M} |\Delta_i(m)|^\theta \right)^{1/\theta} \\ &= 2 \left(\sum_{i=1}^{N_M} |f_i \cdot m|^\theta \right)^{1/\theta} \end{aligned} \quad (6)$$

is called the *salience* of the texture difference produced by m . θ is analogous to the slope of the psychometric function and typically has values between 3 and 6 (e.g., Watson, 1979). We only assume that $\theta \geq 1$.

2.6. The indeterminacy argument

Note that if $\theta = 2$, then salience is Euclidean. In this case, for F the $17 \times N_M$ matrix whose i th column vector is f_i , Eq. (6) becomes

$$\text{Sal}(m) = 2(m^T F F^T m)^{1/2} \quad (7)$$

(treating m as a column vector). Thus, equisalience contours form parallel ellipses centered at 0. In practice it is difficult to reject the hypothesis that a given set of discrimination data conforms to such a pattern (Poirson & Wandell, 1990). Moreover, if salience is given by Eq. (7), then for any $N_M \times N_M$ isometry R (i.e., rotation and/or reflection), it is also true that

$$\text{Sal}(m) = 2(m^T F R R^T F^T m)^{1/2}. \quad (8)$$

That is, the matrix FR of candidate impact functions yields exactly the same value of salience as F for any modulator. This suggests that the impact functions $f_1, f_2, \dots, f_N$ can be determined only up to an arbitrary isometry and hence that individual mechanisms cannot be characterized. However, as we show in Sections 2.7 and 2.8, certain antecedent conditions (which hold in the current case) can be exploited to determine the impact function of an individual mechanism.

2.7. Univariate modulation spaces

A modulation space Φ is called perceptually *univariate* if there exists $f_\Phi \in \Phi$ such that

$$\text{Sal}(m) = 2 |f_\Phi \cdot m| \quad (9)$$

for any $m \in \Phi$ (note that this factor of 2 is the same one as in Eqs. (4) and (7)). In this case f_Φ is called the *discriminator* of Φ . Note that for any two modulators m_1 and m_2 in univariate modulation space Φ ,

$$\begin{aligned} \text{Sal}(m_1 + m_2) &= 2 |f_\Phi \cdot m_1 + f_\Phi \cdot m_2| \\ &= \begin{cases} \text{Sal}(m_1) + \text{Sal}(m_2) & \text{if } \text{sign}(f_\Phi \cdot m_1) = \text{sign}(f_\Phi \cdot m_2), \\ | \text{Sal}(m_1) - \text{Sal}(m_2) | & \text{if } \text{sign}(f_\Phi \cdot m_1) \neq \text{sign}(f_\Phi \cdot m_2). \end{cases} \end{aligned} \quad (10)$$

Several tests of univariance follow from Eq. (10). Suppose, for example, that modulators m_1 and m_2 support equal discrimination success rates. In this case, $\text{Sal}(m_1) = \text{Sal}(m_2)$, from which Eq. (10) implies that either

$$\text{Sal}(m_1 + m_2) = 0 \quad \text{and} \quad \text{Sal}(m_1 - m_2) = 2\text{Sal}(m_1) \quad (11)$$

or else

$$\text{Sal}(m_1 - m_2) = 0 \quad \text{and} \quad \text{Sal}(m_1 + m_2) = 2\text{Sal}(m_1). \quad (12)$$

Thus, if m_1 and m_2 yield equal discrimination performance, then one of $m_1 + m_2$ or $m_1 - m_2$ must yield chance performance. If m_1 and m_2 are orthogonal, and each alone yields easy discrimination, this *cancellation test* provides a powerful test of univariance.

The *titration method* gives a second test and also enables one to determine the discriminator f_Φ if Φ is univariate. Let $\phi_1, \phi_2, \dots, \phi_n$ be an orthogonal basis of Φ , and suppose ϕ_1 has been scaled to yield threshold discrimination performance. Without loss of generality, we set $f_\Phi \cdot \phi_1 = 1$. For any other ϕ_k ($k = 2, 3, \dots, n$), the locus of points (x, y) for which $x\phi_k + y\phi_1$ supports threshold performance is called the *titration line* of ϕ_k with ϕ_1 . Eq. (10) implies that this line must be straight with x -intercept x_0 satisfying $f_\Phi \cdot \phi_k = x_0^{-1}$. Chubb et al. (1994) describe methods for efficiently estimating the titration line. If all titration lines are convincingly straight, univariance is supported, in which case Eqs. (1) and (2) can be used to estimate f_Φ (note that $\text{Proj}_\Phi(f_\Phi) = f_\Phi$).

2.8. Isolating a single mechanism

As proven in Appendix A, if Φ is perceptually univariate with discriminator f_Φ , then for each of $i = 1, 2, \dots, N_M$, $\text{Proj}_\Phi(f_i)$ must be equal to $k_i f_\Phi$ for some scalar k_i . Suppose, then, that a modulation space Φ of moderately high dimension N (e.g., $N \geq 3$) is found to be perceptually univariate. In this case, we naturally infer that only a single mechanism senses any differences between the scrambles generated by $m \in \Phi$. Although this conclusion is not logically necessary, the only other possibility is that multiple mechanisms each project to exactly the same line through the origin in Φ , which seems implausible.

3. Empirical background

λ_1 is easily discriminable (compare Fig. 1a and b), as is λ_2 (Fig. 1c and d). For λ_1 , discrimination is driven by a difference in texture brightness, which seems to be controlled by histogram mean, whereas for λ_2 , discrimination seems driven not by brightness but by the difference in some statistic like histogram variance. Intuition thus suggests that the mechanism used to discriminate λ_1 is different from that used to discriminate λ_2 . This intuition is supported by the cancellation test which strongly rejects the hypothesis that the space spanned by λ_1 and λ_2 is perceptually univariate. It follows that at least two mechanisms are used to discriminate modulators varying in λ_1 and λ_2 . Call these mechanisms M_1 and M_2 and their impact functions f_1 and f_2 .

Chubb et al. (1994) hypothesized that only one mechanism is differentially sensitive across the 14-dimensional space $A_{3,\dots,16}$ orthogonal to both λ_1 and λ_2 .

(For any $m \in A_{3,\dots,16}$, $U + m$ has the same mean and variance as U .) They used the titration method to test the perceptual univariance of the five-dimensional space $A_{3,\dots,7}$ spanned by λ_k , $k = 3, 4, \dots, 7$ (as a surrogate for $A_{3,\dots,16}$). With very small measurement error, the perceptual univariance of $A_{3,\dots,7}$ was confirmed for one observer and barely rejected for another. Fig. 3 plots the (nearly identical) estimates of the discriminator $f_{3,\dots,7}$ for both observers, scaled to be maximal modulators. Using the logic of Section 2.8, they concluded that discrimination in $A_{3,\dots,7}$ (and thus almost certainly in $A_{3,\dots,16}$) is accomplished by a single mechanism B with impact function f_B whose projection into $A_{3,\dots,7}$ is a scaled version of $f_{3,\dots,7}$.

3.1. Other mechanisms are slightly sensitive to differences in $A_{3,\dots,7}$

Contrary to this conclusion, however, mechanisms other than B seem to have slight differential sensitivity to some $m \in A_{3,\dots,7}$. In particular, modest failures of the cancellation test are obtained for λ_3 and λ_4 . Why did the results of Chubb et al. (1994) fail to implicate more than one mechanism? First (this may be the whole story), the cancellation test provides a more stringent test of perceptual univariance than the titration method. A second possibility is the following. It seems likely that texture discrimination processes are more plastic than Eq. (6) suggests. Suppose observers can attentionally weight the differential activations $\Delta_k(m)$, $k = 1, 2, \dots, N_M$, to optimize performance in discriminating m . If so, then because B is the single mech-




Fig. 3. The discriminator $f_{3,\dots,7}$. This function was found by Chubb et al. (1994) to govern discrimination within the modulator space $A_{3,\dots,7}$. Specifically, for all modulators $m \in A_{3,\dots,7}$, the probability of discriminating m was found to be a function of $|f_{3,\dots,7} \cdot m|$.

anism whose response varies most dramatically across $A_{3,\dots,7}$, participants in the experiments of Chubb et al. (1994) may have chosen to “tune in” to $\Delta_B(m)$ and ignore the typically much smaller differential activations of other mechanisms. It should be noted, however, that the logic underlying the current experiments holds regardless of whether or not observers used such a strategy.

3.2. Implications

These findings imply that there are at most three scramble-sensitive mechanisms. If B is distinct from M_1 and M_2 , then there are exactly three. However, the results of Chubb et al. (1994) do not rule out the possibility that B is one and the same as one of M_1 or M_2 . This issue was resolved by Chubb, Landy, Nam, Bindman, and Sperling (2004) who showed that for any three distinct $i, j, k \in \{1, 2, 3, 4\}$, discrimination is above chance for any maximal modulator m in the space $A_{i,j,k}$ spanned by λ_i , λ_j , and λ_k . If M_1 and M_2 were the only mechanisms, then there would have to exist an m in the three-dimensional space $A_{i,j,k}$ orthogonal to both f_1 and f_2 , implying that $\text{Sal}(m) = 0$ and hence that m should yield chance performance. The failure to find such an m in any of the four subspaces $A_{i,j,k}$ thus argues that M_1 , M_2 and B are distinct. Chubb et al. (2004) also show that the space $A_{1,2,3,4}$ spanned by λ_1 , λ_2 , λ_3 and λ_4 does indeed contain a maximal modulator m yielding chance performance, bolstering the conclusion that human vision has only the three scramble-discriminating mechanisms, M_1 , M_2 and B . Note finally that each of f_1 and f_2 can be well-approximated by a quadratic function. Otherwise $A_{3,\dots,7}$ should deviate more strongly from perceptual univariance.

3.3. Toward a more complete characterization of B

Assuming $\lambda_8, \lambda_9, \dots, \lambda_{16}$ contribute little to f_B , then for unknown weights W_1 , W_2 , and $W_{3,\dots,7}$

$$f_B = W_1\lambda_1 + W_2\lambda_2 + W_{3,\dots,7}f_{3,\dots,7} + \left[\begin{array}{c} \text{minor contributions} \\ \text{from higher order terms} \end{array} \right]. \quad (13)$$

Our goal is to crystallize our picture of f_B by measuring W_1 and W_2 . Note again that we can only determine f_B up to an unknown multiplicative constant.

4. Methods

4.1. Logic of the methods

Suppose we are assured that a particular modulator m , yielding threshold success rate ρ , “isolates B ” in the sense that $\Delta_1(m) = \Delta_2(m) = 0$, implying that m -discrimi-

nation is accomplished by B alone. In this case, for low amplitude modulators δ , even though $f_1 \cdot \delta$ and $f_2 \cdot \delta$ may be non-zero, we expect that B will continue to predominate in discriminating $m + \delta$. Thus, we might hope to infer B 's sensitivity to λ_k (for either $k = 1$ or $k = 2$) by adding a small amount of λ_k to m and observing its effect on discrimination.

Toward this end, we must first choose a modulator m likely to do a good job of isolating B . Certainly m should be drawn from $A_{3,\dots,7}$. Indeed if $A_{3,\dots,7}$ were perfectly univariate, we might choose m indiscriminately from $A_{3,\dots,7}$. However, this is not the case (see Section 3.1). We opt for $m = f_{3,\dots,7}$ in hopes of minimizing the ratios $|f_k \cdot m|/|f_{3,\dots,7} \cdot m|$ ($k = 1, 2$). Across all m of some fixed norm this choice maximizes the denominator; without knowing more about f_1 and f_2 this is the best we can do. We assume that $f_{3,\dots,7} \cdot f_k \approx 0$ for $k = 1, 2$.

The first step in measuring $f_B \cdot \lambda_1$ and $f_B \cdot \lambda_2$ is to observe the success rate at discriminating $Af_{3,\dots,7}$ for various histogram amplitudes A . We use a 4-AFC task (chance = 0.25). These data are used to estimate the psychometric function Ψ of Eq. (5) which we model as a Weibull function. It is convenient to choose the arbitrary scale factor for salience by setting $|f_B \cdot f_{3,\dots,7}| = 1$. Under this convention

$$\text{Sal}(Af_{3,\dots,7}) = 2 |Af_B \cdot f_{3,\dots,7}| = 2A. \quad (14)$$

For any $p \in [0.25, 1]$, we write A_p for the amplitude such that $A_p f_{3,\dots,7}$ yields discrimination success rate p , implying that $2A_p = \text{Sal}(A_p f_{3,\dots,7}) = \Psi^{-1}(p)$.

Next, for threshold probability ρ (we use $\rho = 0.6$) we fix a base modulator $b = A_\rho f_{3,\dots,7}$ and proceed to measure the effect on performance of perturbing b in four ways. For small modulation amplitudes ε_1 and ε_2 , we measure performance at discriminating

$$\begin{aligned} b_1^- &= b - \varepsilon_1 \lambda_1, \\ b_1^+ &= b + \varepsilon_1 \lambda_1, \\ b_2^- &= b - \varepsilon_2 \lambda_2, \quad \text{and} \\ b_2^+ &= b + \varepsilon_2 \lambda_2. \end{aligned} \quad (15)$$

Let ρ_1^- , ρ_1^+ , ρ_2^- and ρ_2^+ be the obtained success rates.

The differential activations produced in B by b_k^- , b and b_k^+ are

$$\begin{aligned} \Delta_B(b_k^-) &= 2f_B \cdot b_k^- = 2(A_\rho - \varepsilon_k f_B \cdot \lambda_k), \\ \Delta_B(b) &= 2f_B \cdot b = 2A_\rho, \quad \text{and} \\ \Delta_B(b_k^+) &= 2f_B \cdot b_k^+ = 2(A_\rho + \varepsilon_k f_B \cdot \lambda_k). \end{aligned} \quad (16)$$

We choose ε_k small enough to be sure that $|\varepsilon_k f_B \cdot \lambda_k| < A_\rho$, implying that all three of $\Delta_B(b_k^-)$, $\Delta_B(b)$, and $\Delta_B(b_k^+)$ in Eq. (16) are positive.

Moreover, for $j, k = 1, 2$, the differential activations produced in M_j by b_k^- , b and b_k^+ are

$$\begin{aligned}
\Delta_j(b_k^-) &= 2f_j \cdot b_k^- = -2\varepsilon_k f_j \cdot \lambda_k, \\
\Delta_j(b) &= 2f_j \cdot b = 0, \quad \text{and} \\
\Delta_j(b_k^+) &= 2f_j \cdot b_k^+ = 2\varepsilon_k f_j \cdot \lambda_k.
\end{aligned} \tag{17}$$

Note (with reference to Eqs. (5) and (6)) that $2A_{\rho_k^-} = \Psi^{-1}(\rho_k^-) = \text{Sal}(b_k^-)$; similarly, $2A_{\rho_k^+} = \text{Sal}(b_k^+)$. Thus substituting the results of Eqs. (16) and (17) into Eq. (6) we find

$$A_{\rho_k^-} = (|A_{\rho_k^-} - \varepsilon_k f_B \cdot \lambda_k|^{\theta} + |\varepsilon_k f_1 \cdot \lambda_k|^{\theta} + |\varepsilon_k f_2 \cdot \lambda_k|^{\theta})^{1/\theta}, \tag{18}$$

and

$$A_{\rho_k^+} = (|A_{\rho_k^+} + \varepsilon_k f_B \cdot \lambda_k|^{\theta} + |\varepsilon_k f_1 \cdot \lambda_k|^{\theta} + |\varepsilon_k f_2 \cdot \lambda_k|^{\theta})^{1/\theta}. \tag{19}$$

Observe that if θ were equal to 1, then $A_{\rho_k^+} - A_{\rho_k^-}$ would be exactly equal to $2\varepsilon_k f_B \cdot \lambda_k$. This would also be true if θ were equal to ∞ , assuming we have chosen ε_1 and ε_2 small enough that $|\Delta_1(m)| > |\Delta_2(m)|$ and $|\Delta_2(m)| > |\Delta_1(m)|$ for all of $m = b_1^-, b_1^+, b_2^-$ and b_2^+ . Thus, if the true value of θ were either 1 or ∞ , then

$$f_B \cdot \lambda_k = \frac{A_{\rho_k^+} - A_{\rho_k^-}}{2\varepsilon_k}. \tag{20}$$

This equality breaks down, however, for values of θ between these extremes. Simulations reveal that the estimate of $f_B \cdot \lambda_k$ provided by Eq. (20) is poorest if the true value of θ is 2, in which case the true value of $f_B \cdot \lambda_k$ is given by

$$f_B \cdot \lambda_k = \frac{A_{\rho_k^+}^2 - A_{\rho_k^-}^2}{4\varepsilon_k A_{\rho_k}}. \tag{21}$$

Thus, Eqs. (20) and (21) provide bounds on $f_B \cdot \lambda_k$ whose difference reflects possible error due to uncertainty about θ . As we shall see, Eqs. (20) and (21) yield very similar estimates of f_B .

4.2. Selection of ε_1 and ε_2

Pilot experiments were used to select appropriate values for ε_1 and ε_2 . Observers are quite sensitive to pure λ_1 modulations, and it is evident that most of this sensitivity is due to mechanisms other than B . To prevent b_1^- and b_1^+ from producing too much differential activation in these other mechanisms, we must choose ε_1 to be small. On the other hand, ε_1 needs to be large enough to enable us to detect a B -driven performance difference for b_1^- versus b_1^+ if such a difference exists. In a pilot experiment, we varied ε_1 to find a value that was large enough so that performance of b_1^- and b_1^+ differed, but the mean performance was close to that of b , and $\varepsilon_1 = 0.06$ achieved the desired compromise. Observers are less sensitive to λ_2 than to λ_1 , which enabled us to use a larger value of $\varepsilon_2 = 0.15$.

4.3. Observers

The observers were the three authors. All have normal or corrected-to-normal vision.

4.4. Linearization

The 17 luminances were roughly $n \times 6 \text{ cd/m}^2$, for $n = 0, \dots, 16$. More precisely, the lowest luminance was 0.5 cd/m^2 ; the highest was 95.5 cd/m^2 . Linearization was achieved by eye as follows. The screen displayed a square wave with bars alternating between two fine-grained patterns. One pattern was a checkerboard of lum_{low} and lum_{high} . The other pattern contained three intensities, lum_{low} , lum_{high} and lum_{mid} (half the area had luminance lum_{mid} , 1/4 had lum_{low} and 1/4 had lum_{high}). (We used a three-luminance pattern rather than a uniform field of luminance lum_{mid} to control for possible spatial non-linearities in the display (Klein, Hu, & Carney, 1996; Mulligan & Stone, 1989).) The screen was viewed from a far enough distance that the fine pattern granularity was invisible. At this distance, the square wave's spatial frequency was roughly 6 c/deg . Since the pattern grain could not be resolved, the square wave was visible only if alternating bars differed in mean luminance. An observer adjusted lum_{mid} to make the square wave vanish, thus setting it to the average of lum_{low} and lum_{high} . We generated a lookup table by repeating this procedure with different values of lum_{low} and lum_{high} to determine the value of lum_{mid} lying midway between lum_{low} and lum_{high} . A smooth function was fit to the resulting data, and 17 equally spaced values were used.

4.5. Stimuli and task

A stimulus comprised a square of 68×68 texels. At the viewing distance of 69 cm , the display subtended 12.7° . A bar of one texture S_{U+m} was placed on a background of S_{U-m} , or vice versa, in one of four possible positions (Fig. 4). The stimulus was displayed for 200 ms following a fixation display. Stimuli were viewed binocularly. The observer used the arrow keys to indicate target bar location and then received audible feedback.

4.6. Conditions

The experiment had two phases. In phase 1, performance at discriminating $Af_{3,\dots,7}$ was measured for $A = 0.125, 0.25, \dots, 1$. Data were collected in five blocks, each comprising 160 randomly sequenced trials, 20 for each value of A . In 10 trials the target bar was filled with $S_{U+Af_{3,\dots,7}}$ and the background with $S_{U-Af_{3,\dots,7}}$; in the other 10 it was the other way around. The data from the two trial types were pooled, yielding 100 trials per condition. A Weibull function, scaled for chance performance of 0.25 was fit (using a maximum likelihood




Fig. 4. The four-alternative forced-choice task. On a given trial, for some modulator m , one of these four stimulus configurations was presented with a bar of texture S_{U+m} on a background of S_{U-m} or vice versa.

procedure) to the resulting data as follows. Following Eq. (5), and continuing to express salience in multiples of $|f_B \cdot f_{3,\dots,7}|$, we derive

$$\begin{aligned} \Psi(\text{Sal}(Af_{3,\dots,7})) &= \Psi(2A | f_B \cdot f_{3,\dots,7} |) = \Psi(2A) \\ &= (1/4) + (3/4)(1 - 2^{-(2A/\alpha)^\beta}). \end{aligned} \quad (22)$$

Letting α absorb the factor of 2 on the right side of Eq. (22) yields

$$\Psi(2A) = (1/4) + (3/4)(1 - 2^{-(A/\alpha)^\beta}). \quad (23)$$

This was used to estimate $A_{0.6} = \Psi^{-1}(0.6)/2$.

In phase 2, for $b = A_{0.6}f_{3,\dots,7}$, $\varepsilon_1 = 0.06$, $\varepsilon_2 = 0.15$, performance was assessed at discriminating each of b_1^- , b_1^+ , b_2^- , b_2^+ (Eq. (15)), and also b (to refine the estimate of $\Psi(b) = 0.6$ derived in phase 1). Each observer ran 16 blocks. Each block had 20 trials in each of these five conditions (with 10 trials of each of the two background/target-bar configurations in each condition), yielding a total of 320 trials per condition.

5. Results

The estimates, derived from phase 1, of psychometric function parameters α (mean) and β (slope) (Eq. (23)) as

well as of $A_{0.6} = \Psi^{-1}(0.6)$ are given in Table 1 for each observer. The results of phase 2 are shown in Fig. 5. Results for λ_1 are shown in the left column. $A_{\rho_1^-}$, A_ρ and $A_{\rho_1^+}$ are plotted against the values $-\varepsilon_1$, 0, and ε_1 (for $\varepsilon_1 = 0.06$). The corresponding results for λ_2 (for $\varepsilon_2 = 0.15$) are shown in the right column.

If performance in these conditions were determined by B alone, then the three points in each plot would lie on a straight line. The tendency of these curves toward upward concavity reflects the influence of mechanisms M_1 and M_2 , which increase the discriminability of b_1^- , b_1^+ , b_2^- and b_2^+ in comparison to b , to which M_1 and M_2 are blind.

Table 1
Parameters of the fit psychometric function

	α	β	$A_{0.6}$
JE	0.8222	2.7609	0.7936
CC	0.5343	2.1593	0.5107
MSL	0.7697	2.2194	0.7365

α is the threshold amplitude. β is the slope. $A_{0.6} = \Psi^{-1}(0.6)$ is the estimate of the amplitude resulting in a proportion correct of 0.6.




Fig. 5. Experimental results. Left-hand column: $A_{\rho_1^-}$, A_ρ and $A_{\rho_1^+}$ are plotted against the values $-\varepsilon_1$, 0, and ε_1 . That is, for each value of λ_1 amplitude, we plot the amount of $f_{3,\dots,7}$ alone that would have resulted in equivalent performance. If the Minkowski exponent is 1 or ∞ , the slope of the best-fitting straight line through these three points provides an estimate of $2f_B \cdot \lambda_1$. Right-hand column: analogous results for λ_2 .

Table 2

Estimates of the first- and second-order components of the impact function f_B along with 95% confidence intervals for each

	JE	CC	MSL
$f_B \cdot \lambda_1$ (Eq. (20))	1.1192	1.0628	0.5041
95% c.i.	± 0.7260	± 0.6078	± 0.8601
$f_B \cdot \lambda_1$ (Eq. (21))	1.2897	1.1392	0.5365
95% c.i.	± 0.8859	± 0.6664	± 0.8714
$f_B \cdot \lambda_2$ (Eq. (20))	-1.3133	-1.1060	-1.2374
95% c.i.	± 0.3015	± 0.2560	± 0.3974
$f_B \cdot \lambda_2$ (Eq. (21))	-1.5233	-1.2126	-1.4089
95% c.i.	± 0.4068	± 0.3278	± 0.4200

Estimates derived from Eq. (20) are appropriate if Minkowski exponent θ is either 1 or ∞ . Deviations from this estimate are maximal if $\theta = 2$; in this case, the correct estimator of $f_B \cdot \lambda_k$ ($k = 1, 2$) is given by Eq. (21). The confidence intervals were computed using a bootstrap (Efron and Tibshirani, 1993) calculated by resampling the raw data.

Eqs. (20) and (21) were used to obtain bounds on $f_B \cdot \lambda_k$, $k = 1, 2$. The resulting estimates are presented in Table 2. For JE and CC, (nearly identical) estimates of $f_{3,\dots,7}$ were derived by Chubb et al. (1994) (compare the o's and +s in Fig. 3). Thus, our new estimates of

$\lambda_1 \cdot f_B$ and $\lambda_2 \cdot f_B$ enable a full seventh-order polynomial approximation of f_B for each of these two observers. For observer MSL, we estimate f_B using the current results plus the solid curve in Fig. 3 (the average of the $f_{3,\dots,7}$ estimates for JE and CC).

Our estimates of $f_B \cdot \lambda_1$ and $f_B \cdot \lambda_2$ were obtained under the arbitrary but convenient assumption that $f_B \cdot f_{3,\dots,7} = 1$. Thus, Eqs. (1) and (2) yield

$$f_B \approx \frac{f_{3,\dots,7}}{f_{3,\dots,7} \cdot f_{3,\dots,7}} + \frac{f_B \cdot \lambda_1}{\lambda_1 \cdot \lambda_1} \lambda_1 + \frac{f_B \cdot \lambda_2}{\lambda_2 \cdot \lambda_2} \lambda_2. \quad (24)$$

Fig. 6 shows the resulting estimates of f_B . The solid lines use the values of $f_B \cdot \lambda_k$ from Eq. (20), and the dashed lines use the values from Eq. (21). The similarity of these two curves shows that the possible variation in our estimate of f_B due to uncertainty about the Minkowski pooling exponent θ (Eq. (6)) is negligible.

6. Discussion

As Fig. 6 makes clear, f_B is essentially flat across the entire range of texel contrasts from -0.75 to 1.0. This means that for any scramble S_p , $B(S_p)$ should be approximately proportional to $p(-1)$, the fraction of S_p 's texels whose contrast is -1 (or very near it). We call the visual statistic sensed by B *blackshot* to emphasize B 's exquisite sensitivity to the relatively sparse spattering of very dark elements (like buckshot) across a predominantly bright and variegated field.

We have suggested that blackshot constitutes a brightness-coding mechanism distinct from those used to discriminate overall texture brightness and contrast. Observers are nearly veridical in judging mean texture brightness (Nam & Chubb, 2000), suggesting the existence of a nearly linear impact function f_1 . On the other hand, observer judgments of texture contrast are primarily influenced by the distribution of texel contrasts below the mean (Chubb & Nam, 2000), consistent with the existence of a second impact function, f_2 , that is similar to half-wave rectification, but does not emphasize the darkest texels nearly so much as f_B .

Whittle (1986) provides a precedent for the current result (see also Kingdom & Moulden, 1991). In that study, observers attempted to judge which of two small, uniform target squares was higher in luminance. The targets could be either brighter or darker than the background. For targets near the background in luminance, discrimination sensitivity followed a "dipper function": the threshold difference in luminance between the two squares at first decreased with deviation of target luminances from background, then began steadily to increase, following Weber's law. For targets greater than the background in luminance, this pattern persisted over the full range of target luminances tested. However, the results were different for targets




Fig. 6. The estimated blackshot impact function f_B for three subjects. For each subject, two functions are shown corresponding to θ values of 1 or ∞ (solid line) or 2 (dashed line).

of luminance lower than the background. Regardless of background luminance, as spot luminances were decreased to values near 0, observers became very sensitive to small differences in luminance between the two spots.

Thus, like the current study, Whittle's (1986) results implicate a system that enables fine discrimination between targets with contrasts near -1 , even though the targets are set against a much brighter background. The current study makes it clear that this system operates preattentively, in spatially parallel fashion to sense blackshot: the varying concentrations of the darkest patches throughout the visual field, regardless of the overall luminance of the background.

Our result may also relate to the way the darkest pixels are grouped in random textures. Frisch and Julesz (1966) found that random black-and-white textures underwent the most frequent figure-ground reversals when 40% of the texels were black. This might suggest that black texels are more capable of grouping, and hence of being treated as foreground, as compared to white texels.

It was mentioned earlier that Chubb et al. (2004) demonstrated the existence of texture scramble metamers, consistent with the notion that the perceptual space for these scramble textures is three-dimensional. In a convincing demonstration of this phenomenon, they were able to construct a texture consisting of only three discrete gray levels that was metameric to S_U (the texture where all 17 gray levels occur in equal numbers). Indeed, observers who are not informed of the difference between the two textures in this display fail to notice any inhomogeneity even after prolonged, free viewing. However, there are limits to this behavior. For example, for histograms p and q defined by setting $p(-3/8) = p(3/8) = 0.5$, and $q(0) = 0.75$ and $q(-3/4) = q(3/4) = 0.125$, scrambles S_p and S_q are easily discriminable despite the fact that they have equal mean (0) and variance (9/64), and equal blackshot (since neither scramble contains any contrasts close enough to -1 to excite the blackshot mechanism). Moreover, it proves impossible to adjust the probabilities of contrasts $-3/4$, 0 and $3/4$ to achieve a scramble metameric to S_p . This argues that if histogram entropy is made very low, then mechanisms other than M_1 , M_2 and B come into play enabling scramble discrimination, even in a brief flash with a post-stimulus masker.

One might well wonder about the adaptive function of the blackshot mechanism. We speculate that the primary function of this system has nothing to do with texture discrimination per se. Rather we suggest that the purpose of the blackshot system is to enable vision to be useful in shaded areas in an otherwise brightly illuminated field of view. If one had to rely on a system whose activation were linear over the range of luminances in

the entire scene, then most shaded areas would appear uniformly black. To draw any visual distinctions within a darkly shaded region of an otherwise bright scene requires a special-purpose system like the blackshot system, sharply tuned to fine differences between intensities that are very low in comparison to the ambient light level.

Acknowledgments

MSL was supported in part by NIH grant EY08266 and Human Frontier Science Program grant RG0109/1999-B.

Appendix A. Our claim in Section 2.8 to have isolated and functionally characterized a single visual mechanism B is based on the following theorem.

Theorem. Let Ω be a Minkowskian modulation space of perceptual dimension N with impact functions f_i , $i = 1, \dots, N$ and exponent θ . Then for any subspace $\Phi \subset \Omega$ the following conditions are equivalent:

- (i) Φ is perceptually univariate with discriminator f .
- (ii) There exist weights k_j , $j = 1, \dots, N$ such that:

$$\begin{aligned} \text{(a)} \quad & \sum_{j=1}^N |k_j|^\theta = 1; \\ \text{(b)} \quad & \text{for all } \phi \in \Phi, |f_j \cdot \phi| = |k_j f \cdot \phi| \quad (j = 1, \dots, N). \end{aligned} \quad (\text{A.1})$$

Proof that (ii) implies (i). If (ii) holds, then immediately, for any $\phi \in \Phi$,

$$\begin{aligned} \text{Sal}(\phi) &= 2 \left(\sum_{j=1}^N |f_j \cdot \phi|^\theta \right)^{1/\theta} = 2 \left(\sum_{j=1}^N |k_j f \cdot \phi|^\theta \right)^{1/\theta} \\ &= 2 |f \cdot \phi| \left(\sum_{j=1}^N |k_j|^\theta \right)^{1/\theta} = 2 |f \cdot \phi|, \end{aligned} \quad (\text{A.2})$$

implying (i).

Proof that (i) implies (ii). Suppose (i) holds. We first show that in this case, for $j = 1, \dots, N$, $\text{Proj}_\Phi(f_j) = k_j f$ for some $k_j \in \mathbb{R}$. Suppose the contrary. That is, suppose that for some j , $\text{Proj}_\Phi(f_j)$ is not a scalar multiple of f . Then for some non-zero function g_j orthogonal to f , and some $k_j \in \mathbb{R}$,

$$\text{Proj}_\Phi(f_j) = k_j f + g_j. \quad (\text{A.3})$$

Because $f \in \Phi$ and $\text{Proj}_\Phi(f_j) \in \Phi$, we note that $g_j = \text{Proj}_\Phi(f_j) - k_j f$ is also an element of Φ , implying that

$$\text{Sal}_\Omega(g_j) = 2 |f \cdot g_j| = 0. \quad (\text{A.4})$$

However, it is also true that $g_j \in \Omega$, implying that

$$\text{Sal}_\Omega(g_j) = 2 \left(\sum_{j=1}^N |f_j \cdot g_j|^\theta \right)^{1/\theta} \geq 2 |f_j \cdot g_j| > 0, \quad (\text{A.5})$$

contradicting Eq. (A.4). This proves that if (i) holds, then $\text{Proj}_\Phi(f_j) = k_j f$ for $j = 1, \dots, N$. Note that furthermore in this case, for any $\phi \in \Phi$,

$$\begin{aligned} |f \cdot \phi| &= \frac{\text{Sal}_\Omega(\phi)}{2} = \left(\sum_{j=1}^N |f_j \cdot \phi|^\theta \right)^{1/\theta} \\ &= \left(\sum_{j=1}^N |k_j f \cdot \phi|^\theta \right)^{1/\theta} = |f \cdot \phi| \left(\sum_{j=1}^N |k_j|^\theta \right)^{1/\theta}, \end{aligned} \quad (\text{A.6})$$

implying that $\sum_{j=1}^N |k_j|^\theta = 1$. This completes the proof that (i) implies (ii).

References

- Beck, J. (1966). Perceptual grouping produced by changes in orientation and shape. *Science*, 154, 538–540.
- Beck, J. (1982). Textural segmentation. In J. Beck (Ed.), *Organization and representation in perception* (pp. 285–317). Hillsdale, NJ: Erlbaum.
- Beck, J., Prazdny, K., & Rosenfeld, A. (1983). A theory of textural segmentation. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 1–38). New York: Academic Press.
- Bergen, J. R., & Adelson, E. H. (1988). Visual texture segmentation based on energy measures. *Journal of the Optical Society of America A*, 3, 98–101.
- Bergen, J. R., & Landy, M. S. (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.
- Bovik, A. C., Clark, M., & Geisler, W. S. (1990). Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 55–73.
- Caelli, T. (1985). Three processing characteristics of visual texture segmentation. *Spatial Vision*, 1, 19–30.
- Chubb, C. (1999). Texture-based methods for analyzing elementary visual substances. *Journal of Mathematical Psychology*, 43, 539–567.
- Chubb, C., Econopoulou, J., & Landy, M. S. (1994). Histogram contrast analysis and the visual segregation of IID textures. *Journal of the Optical Society of America A*, 11, 2350–2374.
- Chubb, C., & Landy, M. S. (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., Nam, J.-H., Bindman, D. R., & Sperling, G. (2004). The three dimensions for encoding contrast in simple textures [abstract]. *Journal of Vision*, 4(8), 713.
- Chubb, C., & Nam, J.-H. (2000). The variance of high contrast texture is sensed using negative half-wave rectification. *Vision Research*, 40, 1677–1694.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman Hall.
- Fogel, I., & Sagi, D. (1989). Gabor filters as texture discriminators. *Biological Cybernetics*, 61, 103–113.
- Frisch, H. L., & Julesz, B. (1966). Figure-ground perception and random geometry. *Perception and Psychophysics*, 1, 389–398.
- Graham, N. (1989). *Visual pattern analyzers*. New York: Oxford.
- Graham, N. (1991). Complex channels, early local nonlinearities, and normalization in texture segregation. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 273–290). Cambridge, MA: MIT Press.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics*, 38, 141–171.
- Julesz, B. (1962). Visual pattern discrimination. *IRE Transactions on Information Theory*, IT-8, 84–92.
- Julesz, B. (1975). Experiments in the visual perception of texture. *Scientific American* (April), 34–43.
- Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature*, 290, 91–97.
- Julesz, B., & Bergen, J. R. (1983). Textons, the fundamental elements in preattentive vision and perception of textures. *Bell System Technical Journal*, 62, 1619–1645.
- 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. *Perception*, 2, 391–405.
- Julesz, B., Gilbert, E. N., & Victor, J. D. (1978). Visual discrimination of textures with identical third order statistics. *Biological Cybernetics*, 31, 137–149.
- Kingdom, F., & Moulden, B. (1991). A model for contrast discrimination with incremental and decremental test patches. *Vision Research*, 31, 851–858.
- Klein, S. A., Hu, Q. J., & Carney, T. (1996). The adjacent pixel nonlinearity: Problems and solutions. *Vision Research*, 36, 3167–3181.
- Knutsson, H., & Granlund, G. H. (1983). Texture analysis using two-dimensional quadrature filters. In *Proceedings of the IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management* (pp. 206–213). Silver Spring, MD: IEEE Computer Society.
- Landy, M. S., & Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research*, 31, 679–691.
- Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, 7, 923–932.
- Mulligan, J. B., & Stone, L. S. (1989). Halftoning method for the generation of motion stimuli. *Journal of the Optical Society of America A*, 6, 1217–1227.
- Nam, J.-H., & Chubb, C. (2000). Texture luminance judgments are approximately veridical. *Vision Research*, 40, 1695–1709.
- Poirson, A. B., & Wandell, B. A. (1990). The ellipsoidal representation of spectral sensitivity. *Vision Research*, 30, 647–652.
- Watson, A. B. (1979). Probability summation over time. *Vision Research*, 19, 515–522.
- Whittle, P. (1986). Increments and decrements: Luminance discrimination. *Vision Research*, 26, 1493–1507.