Spatial-scale contribution to the detection of mirror symmetry in fractal noise

Stéphane J. M. Rainville and Frederick A. A. Kingdom

McGill Vision Research Unit, 687 Pine Avenue West, H4-14, Montréal, Québec H3A 1A1, Canada

Received November 23, 1998; revised manuscript received March 19, 1999; accepted May 4, 1999

We investigated how the detection of mirror symmetry depends on the distribution of contrast energy across spatial scales. Stimuli consisted of vertically symmetric noise patterns with fractal power spectra defined by $1/f^{\beta}$ slopes ($-2 \leqslant \beta \leqslant 5$). While overall rms contrast remained fixed at 25%, symmetry-detection thresholds were obtained by corrupting the signal with variable amounts of noise with identical spectral characteristics. A first experiment measured thresholds as a function of spectral slope, and performance was found to be substantially facilitated in images with power spectra that characterize natural scenes ($1.2 \leqslant \beta \leqslant 3.2$). In a second experiment, symmetry was removed from randomly chosen octave bands and replaced by noise with the same spectral profile. Results revealed that only in images with $1/f^{2}$ spectra does performance decrease by constant amounts across all frequency bands. Together, the results imply that symmetry mechanisms extract equal amounts of information from constant-octave frequency bands but lack the ability to whiten stimuli whose spectral slopes differ from those of natural scenes. Results are qualitatively well predicted by a multichannel model that (1) relies on spatial filters with equal-volume point-spread functions and constant-octave frequency bandwidths and (2) restricts the computation of symmetry to spatial regions whose dimensions are proportional to the filters' spatial scale. These findings are also consistent with the notion that mechanisms that mediate the perception of form exploit the ability of early vision to reduce second-order redundancy in natural scenes. [S0740-3232(99)00909-6]

OCIS codes: 330.1800, 330.4060, 330.5510, 330.6100, 330.6110, 330.6180.

1. INTRODUCTION

The human visual system is particularly sensitive to mirror symmetry—a property found in several objects of ecological importance, such as faces and body silhouettes. Indeed, the mere reflection of a random pattern about an axis is often sufficient to elicit compelling impressions of familiar shapes, and this suggests an important role for symmetry in the processing of visual form. However, a full explanation of the visual mechanisms that mediate the detection of mirror symmetry has yet to be cast in terms of other known properties of the visual system. In particular, the relative contribution of spatial scales to symmetry detection is a fundamental issue that remains unresolved and is addressed in the present paper.

The early stages of vision involve neurons that filter the original image at a variety of scales,[1,2] and there is increasing evidence that this spatial analysis precedes the computation of symmetry. In an early demonstration, Julesz and Chang[3] showed that the sum of vertically symmetric and horizontally symmetric white-noise patterns appears to contain no symmetry but that the two symmetries can be simultaneously perceived if the two component patterns are spatially filtered into frequency bands that differ by two octaves. In more-recent studies, Dakin and colleagues reported that performance is roughly scale invariant in bandpass noise patterns[4] and that symmetry is computed over an integration region whose dimensions are proportional to the spatial scale of the stimulus.[5] Other studies also showed that symmetry detection in random-dot displays is quite resistant to spatial jitter near the axis,[6–11] and this finding suggests filtering mechanisms operating at a coarser scale than that of the dots. The fact that detection thresholds are also dependent on the orientation content of the image[4,12–16] strengthens the evidence that symmetry detection is cortical. Moreover, symmetry detection is vulnerable to reversals in contrast polarity between mirror-image halves[17–19] and therefore relies on mechanisms that are sensitive to spatial phase (although see Tyler and Hardage[20]). Finally, there are some reaction-time data that suggest that symmetry perception relies on a coarse-to-fine analysis.[21] Taken together, these findings imply that symmetry detection is mediated by spatial filters whose properties are similar to those of simple-cell receptive fields found in primary visual cortex.[22]

In the present paper we investigate whether the mechanisms that mediate the detection of mirror symmetry rely more on some spatial scales than on others. Stated differently, we ask whether symmetry information in broadband stimuli is given equal consideration across spatial scales. There are two reasons why this issue has not yet been resolved. First, many studies of symmetry involve images composed either of randomly positioned hard-edge elements[6–9,17,23,24] (e.g., dots or lines) or white noise[10] that, owing to their broadband properties, are presumed to recruit filters from a variety of spatial tunings. In such experiments, however, the relative contributions from spatial filters of different scales are difficult to determine, primarily because these stimuli have statistically flat power spectra, which, according to current studies of suprathreshold contrast sensitivity,[25,26] tend to perceptually overemphasize high spatial frequencies. Consequently, broadband stimuli with flat power spectra give only partial insight into the weighting of spatial

scales in symmetry detection because they restrict the range of scales to which the visual system has access. In addition, sparse random-dot stimuli are not ideal for testing issues related to spatial scale since the low dot density prevents significant amounts of information from entering the small spatial integration region for symmetry reported for high frequencies.[5]

Second, studies that explicitly considered the role of spatial scales in symmetry detection have used band-limited patterns[4,5,12,13] which provide useful results on how symmetry mechanisms operate at different individual scales. However, such stimuli make the response to a broadband pattern difficult to predict because the rules by which spatial scales are combined in symmetry detection cannot be revealed by studying each scale in isolation. For instance, symmetry detection with band-limited stimuli should be more or less invariant with perceived contrast, but for broadband stimuli such as white noise, access to a desired scale may depend on the *relative* perceived contrast between scales. In addition, access to a given spatial scale in broadband stimuli may be prevented by competitive cross-scale interactions such as winner-take-all rules. Examples of this type highlight how predictions from band-limited studies may fail to generalize to broadband stimuli. In short, broadband stimuli are better suited than narrow-band patterns for determining how spatial scales interact and are combined in symmetry detection.

In this paper we emphasize the use of stimuli that give equal *perceptual* access to information contained in different frequency bands of the stimulus. Unlike white noise, the ensemble of natural scenes is characterized by a power spectrum that falls roughly as $1/f^2$ (Refs. 26–30; N.B.: $1/f^2$ in the power spectrum corresponds to $1/f$ in the amplitude spectrum), although the measured spectral slopes of individual natural scenes exhibit significant variability across images. This phenomenon was noted by Field and Brady,[31] who reported that, across studies on natural scenes, spectral slopes vary from $1/f^{1.2}$ to $1/f^{3.2}$ and that averages from different studies fall between $1/f^{1.8}$ and $1/f^{2.4}$. The spectral properties of natural images have led various researchers to suggest that the mammalian visual system operates with maximal efficiency when viewing a $1/f^2$ distribution of contrast energy across scales.[26,32,33] In particular, Brady and Field[25] and Field[26] proposed a model of contrast sensitivity in which the gain and bandwidth properties of visual cells are configured such that their expected activity profile is flat across spatial scales only when images with $1/f^2$ spectra are viewed. To a good first approximation, psychophysical and physiological evidence supports this model of suprathreshold contrast sensitivity.[1,25,31,34–36–39] Such a model, in turn, is expected to promote the maximum transfer of information to higher-level processes such as those that mediate symmetry detection.

In the following sections we report two psychophysical experiments designed to investigate the relative contributions of spatial scales to the detection of symmetry. In both experiments, stimuli consist of symmetric broadband noise spatially filtered to one of several $1/f^{\beta}$ spectral profiles, where β determines the rate of contrast-energy decay as a function of spatial frequency. In the remainder of the paper, to simplify notation we refer to β as the spectral slope of the stimulus. In the first experiment we measured symmetry-detection thresholds in noise patterns of variable spectral slope. To the extent that each spatial scale contributes to symmetry detection, performance is expected to peak in images whose spectra fall within the range that characterizes natural scenes. In the second experiment we studied the effects of randomizing the phase spectrum in one of four randomly chosen constant-octave bands of frequencies. If scales contribute equally to symmetry detection, then imposing phase randomization in $1/f^2$ images is expected to reduce performance by a constant amount irrespective of which scale is chosen. Finally, we cast our findings in terms of a model of human performance that combines across spatial scales and predicts symmetry detection in broadband noise of variable spectral slope.

2. GENERAL METHOD

A. Observers

The two authors participated as observers in the experiments. SR is slightly myopic and wore optical correction during trials. FK is an emmetrope and wore no optical correction.

B. Hardware and Calibration

We carried out all experiments on a Power Macintosh computer (7100/80) driving a standard 8-bit/gun video card. Stimuli were presented on a Sony Multiscan 17-in. display at a resolution of 1024×768 pixels. The linearity of the display was measured with a UDT photometer, and linear gamma correction was provided by a look-up table derived from least-squares estimates. The display had an effective look-up table depth of 7.1 bits after gamma correction and a mean luminance of 33.4 cd/m^2. To prevent pixel underflow or overflow, we thresholded intensities that fell outside the range defined by the look-up table below 0 and above 255, although these intensities were extremely rare. Routines from Pelli's VideoToolbox[40] were incorporated into the software that was used to display the stimuli.

C. Procedure

Stimuli were presented at display mean luminance and scaled to 25% rms contrast. Observers viewed stimuli from a distance of 63 cm, at which patterns subtended 3.2 deg of visual angle. Patterns had square dimensions consisting of 128 pixels on a side, and the spatial-frequency content of the images was bracketed between 0.3 and 20.0 cpd.

A two-alternative forced-choice paradigm was used in which observers were presented with a symmetry-present and a symmetry-absent stimulus on every trial. The order of presentation was randomly interleaved from one trial to the next. Observers pressed one of two keys to report the interval that appeared symmetric. Images were presented for 444 ms (or the equivalent of 30 screen refreshes at 67.6 Hz) and were separated by an interstimulus interval of 444 ms. No feedback was provided.

3. EXPERIMENT 1: SYMMETRY DETECTION AS A FUNCTION OF SPECTRAL SLOPE

Stimuli consisted of random-noise patterns that varied in spectral slope and in the amount of vertical mirror symmetry they contained. Figure 1 shows examples of stimuli for three spectral profiles (columns) and four levels of symmetry (rows).

Images were computed on-line with software developed by the first author. Vertically symmetric white-noise patterns were obtained by random sampling from a Gaussian distribution ($\mu = 0$, $\sigma = 1$) and assignment of values to pairs of pixels positioned symmetrically about the vertical axis, as given by the relationship $g(x, y) = g(-x, y)$. The origin $g(0, 0)$ was assigned to the centermost pixel in the image. The continuum of symmetry levels required for measuring detection thresholds was obtained by corrupting the symmetric signal with variable amounts of added uncorrelated Gaussian noise. We define a level of symmetry, denoted ξ here, as the proportion of total image variance taken up by the signal, in line with the notion of signal-to-noise ratio commonly found in the digital signal processing literature. ξ is given by

$$\xi = \frac{{\sigma_s}^2}{{\sigma_s}^2 + {\sigma_n}^2}, \tag{1}$$

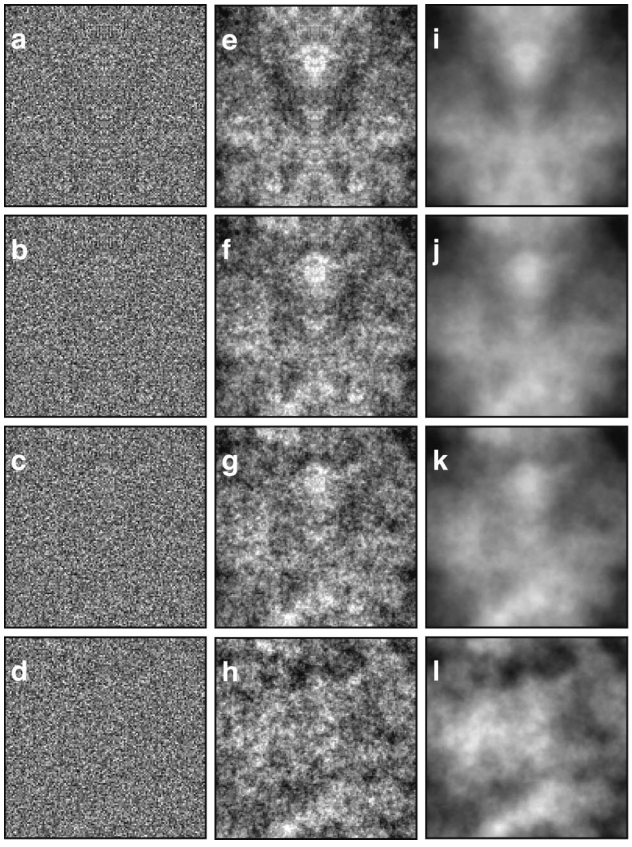

Fig. 1. Examples of symmetric broadband noise patterns used in experiment 1. Each row shows three images of equal rms contrast but that vary in the slope of their power spectra: a–d, $1/f^0$; e–h, $1/f^2$; and i–l, $1/f^4$. Each column shows four patterns with decreasing amounts of mirror symmetry: a, e, i, 1.0; b, f, j, 0.67; d, h, k, 0.33; and e, i, l, 0.0.

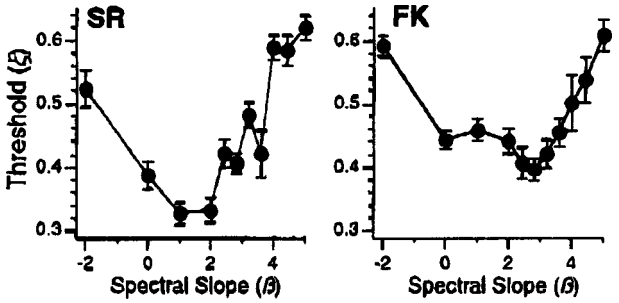

Fig. 2. Results from experiment 1. Symmetry-detection thresholds are shown as a function of spectral slope for observers SR and FK. Error bars indicate ±1 standard-deviation estimates.

where ${\sigma_s}^2$ and ${\sigma_n}^2$ correspond to the variances of the signal and the corrupting noise, respectively. Levels of symmetry necessarily fall between $\xi = 0$ (perfectly random) and $\xi = 1$ (perfectly symmetric).

Once the white-noise pattern with the desired level of symmetry was computed, the image was filtered to the desired spectral profile by computing its fast Fourier transform, multiplying that by $1/f^{\beta/2}$, and reverse transforming it to the spatial domain (N.B.: the spectral-slope constant must be divided by 2 since filtering must be carried out in the amplitude domain rather than in the power domain). Radial frequency f is defined here as $f = \sqrt{u^2 + v^2}$, where u and v represent the dimensions of a two-dimensional Cartesian frequency coordinate system. Imposing spatial filtering only after symmetry is introduced into the white-noise pattern prevents the appearance of artifacts that could have arisen if the left and right halves of the image had been filtered independently of each other.

Symmetry-detection thresholds were measured in patterns with $1/f^{\beta}$ spectral profiles selected from a range of spectral slopes $\beta(-2.0 \leqslant \beta \leqslant 5.0)$ in rough steps of 0.4. Our selection of spectral slopes concentrated on the fractal range investigated by Knill *et al.*[37] for nonsymmetric random-noise stimuli but was also extended to lower values. For each spectral slope, symmetry-detection thresholds were obtained with an adaptive Bayesian method,[41] which estimated the level of symmetry that corresponds to a performance of 81% correct. Each run consisted of 50 trials, and at least 5 runs were made for each spectral slope unless large error bars prompted additional observations. Data were pooled within each spectral slope condition, and a log-invariant Weibull function was used to fit the psychometric function.

Figure 2 shows symmetry-detection thresholds as a function of spectral slope for observers SR and FK. Despite some variability across observers, detection thresholds clearly follow a distinct U-shaped pattern as a function of spectral slope. For SR, best performance was achieved with patterns whose spectral slopes fell between 1 and 2. By comparison, FK's performance was optimal for spectral slopes near 2.8. FK's performance was generally worse than SR's by a constant of ~0.05. The magnitude of the effect between best and worse performance for each observer is significant since correct performance was maintained at 81% correct despite an approximate 20% decrease in signal energy and a simultaneous 20% increase in the energy of the corrupting noise. From

these data, we note that performance is substantially facilitated for spectral slopes that fall within the range that characterizes natural scenes but that facilitation is quite broadly tuned.

4. EXPERIMENT 2: NARROW-BAND PHASE RANDOMIZATION

Stimuli for experiment 2 were identical to the ones used in experiment 1, with the exception that the symmetric signal in one of four randomly selected frequency bands was replaced by random noise with the same spectral profile. This technique is equivalent to randomizing the phase of all frequency components within the chosen band. Phase randomization is a technique that has been used in other studies on symmetry perception[4,5,13] and that is increasingly common in other branches of psychophysical research.[42,43] Our band-selective phase-randomization paradigm has a considerable advantage over conventional bandpass stimuli in that it allows us to study how a single frequency band contributes to symmetry detection in the presence of other spatial scales. This is important because effects that are due to scale combination or cross-scale contrast differences could not be studied in the context of band-limited stimuli.

Gaussian white-noise patterns with specified levels of symmetry were obtained exactly as outlined in experiment 1, and the stimuli were notch filtered with one of four idealized isotropic filters that tiled frequency space in equal nonoverlapping strips of 1.2 octaves. These notch filters, denoted here by $H(f)$, are given by

$$H(f) = \begin{cases} 0 & \text{if } f_{\text{low}} \leqslant f < f_{\text{high}} \\ 1 & \text{otherwise} \end{cases}, \qquad (2)$$

where f_{low} and f_{high} define the lower and upper cutoff frequencies, respectively, of the notch. Table 1 lists, in retinal spatial frequencies, the center frequencies as well as the f_{low} and f_{high} cutoffs of the four notch filters used in this experiment.

The vacant frequency band left by notch filtering was replaced by random (i.e., nonsymmetric) white Gaussian noise filtered with the complement of the corresponding notch filter. The complement of the notch filter in Eq. (2) is simply given by $H'(f) = 1 - H(f)$.

In the first column of Fig. 3 we show four $1/f^2$ noise patterns, each of which lacks mirror symmetry in one of four octave bands of frequencies. The second and third columns show the symmetric and phase-randomized frequency bands, respectively. The symmetric component consists of a notch-filtered $1/f^2$ noise pattern with variable amounts of symmetry, just as in experiment 1. The phase-randomized component consists of perfectly random bandpass noise. Visual inspection of patterns in the first column of Fig. 3 reveals that their statistics are indistinguishable from those in the middle column of Fig. 1, except for the absence of symmetry at one of four spatial scales. This confirms that phase randomization removes symmetry without introducing artifacts into the power spectrum.

Throughout the experiment, symmetry levels were fixed to the detection thresholds (81% correct) measured in experiment 1 under the corresponding spectral-slope conditions. Patterns were presented by a method of constant stimuli in which the band lacking symmetry was randomly chosen from one trial to the next. To detect potential ceiling effects and to verify that the adaptive procedure in experiment 1 had approximated the point 81% correct performance, we designed one fifth of the trials to contain patterns in which symmetry was not replaced. The number of spectral slopes was limited to four key conditions ($\beta = -2, 0, 2, 4$) because the space of all possible conditions is too large to explore should all spectral slopes of experiment 1 be used.

Table 1. Spatial-Frequency Parameters of Notch Filters Used in Experiment 2[a]

Band	f_c (cpd)	f_{low} (cpd)	f_{high} (cpd)
1	1.7	1.25	2.2
2	3.4	2.5	4.7
3	7.0	5.0	9.7
4	14.0	10.0	19.7

[a] Bands are numbered from 1 to 4 in increasing order of center frequency f_c. Lower- and upper-frequency cutoffs are given by f_{low} and f_{high}, respectively. All spatial frequencies are expressed in cycles per degree (cpd).

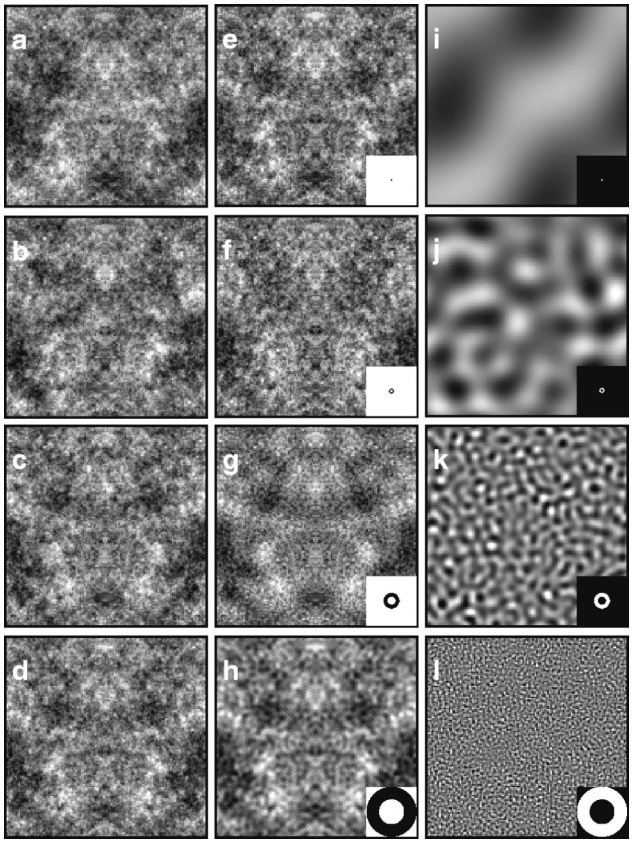

Fig. 3. Examples of stimuli used in experiment 2. Left column: Patterns are identical to those of experiment 1, except that information in one of four possible octave frequency bands (rows) is phase randomized. Middle column, symmetric (i.e., nonrandomized) component. Right column, phase-randomized components. Patterns from the second and third columns are added together to produce patterns in the first columns. Insets, complementary band–reject and band–pass filter pairs in the Fourier domain.

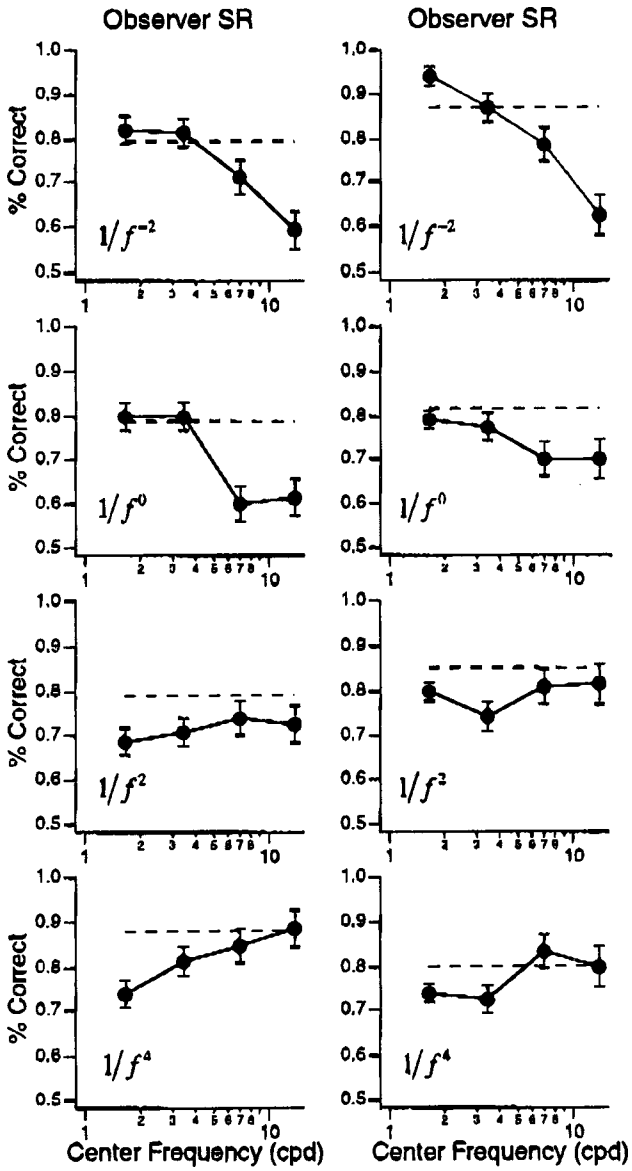

Fig. 4. Results from Experiment 2. Performance is shown for two observers (columns) and four spectral slopes (rows). Graphs plot percent-correct performance as a function of the center frequency of the phase-randomized band. Spectral profiles are indicated along the rows. Dashed lines, performance with symmetry present in all four bands. Error bars, ±1 binomial standard-deviation estimates.

Figure 4 shows results for the two observers (columns) and the four spectral slopes (rows).

Results are encouragingly consistent between the two observers. For a spectral slope of 2, performance drops by approximately 5–10%, regardless of the band in which symmetry was replaced, with no significant rise or decrease in performance across bands. However, for spectral slopes of −2 and 0, performance is virtually unchanged by the absence of symmetry in low-frequency bands but is considerably affected when symmetry is lacking at high frequencies. For a spectral slope of 4, both observers show marked reduction in performance at low frequencies, whereas performance at high frequencies is comparable with that of the all-band condition shown by the dashed lines. Moreover, the dashed lines confirm that the adaptive routine in experiment 1 has, for the most part, correctly homed in on the 81% mark, although this is not a critical issue since the thresholds of experiment 1 were used here only to avoid floor or ceiling effects, which would have masked the trend in the data.

5. DISCUSSION OF RESULTS

Data from the first experiment are compatible with the notion that symmetry detection in broadband images relies on a wide range of spatial scales, provided that these scales are roughly equated for visibility. The fact that performance with patterns whose spectra fall within the spectral range of natural scenes is better than when either high spatial frequencies (e.g., $1/f^0$) or low spatial frequencies (e.g., $1/f^4$) are perceptually dominant also opens up the possibility that symmetry mechanisms combine across scales to achieve better performance. However, the first experiment by itself does not prove that access to multiple spatial scales is responsible for better detection under $1/f^2$ conditions. Since each spectral profile used in this experiment corresponds to a unique level of contrast energy at each spatial scale, it is conceivable (although perhaps improbable) that our U-shaped data depend on the activity of a single channel rather than on a pooling of channels. In particular, the amount of contrast energy in the middle-frequency range of the stimulus peaks near $1/f^2$ conditions because middle frequencies have little or no energy in very shallow (e.g., $1/f^{-2}$) or very steep (e.g., $1/f^4$) spectral profiles. Unfortunately, changing viewing distance would only partially solve the problem because the middle-frequency peak in energy for $1/f^2$ profiles is a property of the stimulus rather than of the visual system. Should we obtain similar results at a different viewing distance, another narrow-band channel could be invoked, and the single-band argument could never be refuted.

Second, even if better performance with $1/f^2$ images were mediated by a pooling of information across scales, the first experiment offers little indication as to how much each spatial scale contributes to performance since results can reflect only the envelope of combined channel responses. In addition, provided that all spatial scales contribute positively to the detection of symmetry, performance is expected to peak at $1/f^2$ irrespective of how spatial scales are weighted in symmetry detection. This point is admittedly counterintuitive, but it necessarily follows from the observation that spatial channels such as those presumed to mediate early vision have thresholds below which stimulus energy is ineffective.[2] Thresholding nonlinearities impose a fundamental limit to the amount of information that can be transferred to higher stages of processing, and no postthreshold whitening scheme could recover the lost information.

For this reason, the value of experiment 1 lies primarily in the demonstration that symmetry detection in noise is optimal for patterns with natural spectral slopes and that higher-level processes such as symmetry detection are not immune to changes in the relative amounts of contrast energy across scales.

Results from the second experiment offer direct evidence that constant-octave frequency bands contribute equally to symmetry detection in broadband images, pro-

vided that they are equated in contrast energy. There are many implications to these findings. First, we can refute the possibility raised above that a single narrowband channel tuned to middle frequencies could account for the data of our first experiment since our results reveal that symmetry detection relies equally on a wide range of octave frequency bands. Our results also suggest that symmetry mechanisms are capable of considering information from different scales in parallel since stimulus presentations were too brief to use attention to scan across scales for bands containing symmetry. This argument is also strengthened by the fact that the frequency band that did not contain symmetry was randomly selected on every trial and that observers could not have performed the task reliably by focusing attention exclusively at a single scale.

The combined findings of experiments 1 and 2 can be interpreted as evidence that spatial scales in symmetry detection are not combined according to some form of winner-take-all strategy. If winner-take-all rules governed scale combination, performance in experiment 1 would have remained approximately constant as a function of spectral slope since the benefit of pooling information across scales would have been lost. Although our methodology does not allow us to determine the precise rules by which scales are combined in the computation of symmetry, the gain in performance obtained when several scales are accessible is compatible with a cooperative scheme such as linear or probability summation.[44,45]

6. MODELING

In this section we first present a brief discussion of what distinguishes human observers from model observers with various capabilities and, in doing so, identify key factors that limit human performance for detecting symmetry. Following this analysis, we propose a model of human performance that combines across spatial channels to predict symmetry detection in broadband noise of variable spectral slope. Finally, we implement our model and compare its performance with human data from experiments 1 and 2.

A. Model Observers versus Human Observers

Model observers are useful tools for understanding the physical content of a stimulus and, as such, they highlight factors that limit human performance. Two separate and crucial differences between model observers and human observers are considered here, namely, the extent of the spatial region over which symmetry is computed and the ability to whiten stimuli. We discuss each of these in turn.

1. Spatial Region of Integration for Mirror Symmetry

Model observers rely on unlimited resources and can therefore consider the entire spatial extent of the stimulus in their computations. By comparison, human observers detecting mirror symmetry in bandpass noise are reported to use information only over a few cycles of spatial scale.[5] To illustrate the effects of the spatial region of integration on performance, consider how four model observers, each equipped with a single bandpass channel, would perform in a symmetry-detection task such as those employed in this study. This situation is illustrated schematically in Fig. 5A, in which a stimulus of finite spatial dimensions is processed by four arrays of spatial filters—that is, four channels—each tuned to one of four different spatial scales (panels a through d) and where each circle corresponds to the spatial extent of an individual filter. To maximize access to spatial information, each model observer would recruit all available spatial filters. Consequently, the model observer using the high-frequency channel (e.g., panel a) would invariably perform better than model observers using lower-frequency channels (e.g., panels b, c, and d) since the high-frequency channel has a comparatively large number of nonoverlapping filters per fixed unit area and can thus encode a large number of independent bits of information. In other words, model observers reveal that the spatial density of information is proportional to filter scale and that high-frequency channels are expected to be more informative than low-frequency channels if the entire spatial extent of the stimulus is taken into account.

Unlike the model observers described above, the human visual system may elect not to integrate across the entire spatial extent of the image. The fact that the dimensions of the integration region for symmetry in human vision are proportional to the scale of the stimulus is a case in point[5] and imposes a fundamental limit on how much symmetry information is integrated by various visual channels. This is illustrated in Fig. 5A, where spatial filters shaded in gray are those recruited by an integration region whose dimensions are proportional to filter scale. Because of the scaling properties of the integration region in human vision, the number of recruited filters is constant across channels and, by the same virtue, the amount of information is also kept constant across scales. This follows from the theoretical notion that self-similar filters integrate a constant number of spatial cycles, irrespective of their scale, but we have also verified this claim by using numerical simulations.

2. Whitening the Power Spectrum

To maximize the performance of a model observer one can endow it with the ability to whiten the power spectrum of the stimulus, a process that consists in removing second-order correlations such as to weigh all bits of information equally. To illustrate this we consider a model observer that has simultaneous access to four constant-octave channels, each tuned to a different spatial scale. To achieve maximal performance this model observer could adjust the gain of each channel based on the contrast structure of the stimulus such as to equate response across channels. This adaptive whitening strategy is highly efficient since it provides equal access to all information while rendering the model observer insensitive to spectral slope (i.e., the relative amount of contrast energy in each frequency band). However, instead of adaptively adjusting the gains of channels to match the power spectrum of any stimulus, a model observer could tune its channels to the statistics of a particular ensemble of images with similar power spectra. Although such an observer would forgo the ability to whiten images whose power spectra deviate too much from what it expects, this

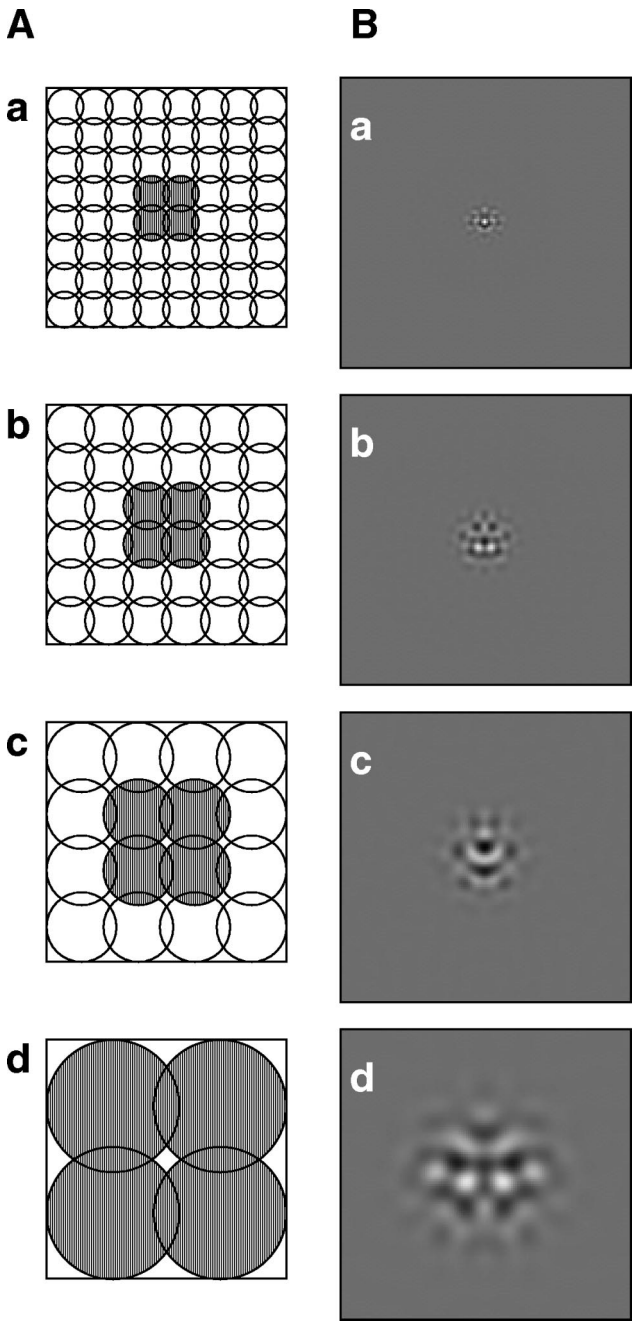

Fig. 5. A, Schematic representation of four bandpass channels (a–d) viewing an image of finite dimensions. Individual circles represent the spatial extent of filters that constitute each channel. Circles shaded in gray represent filters recruited by an integration whose dimensions are proportional to the filter scale. B, Output from four channels (a–d) viewing a perfectly symmetric $1/f^2$ noise pattern. Channels are spatially limited by a Gaussian region of integration, and underlying filters have constant-octave frequency bandwidths and constant-volume point-spread functions.

may not matter much if deviations from the expected slope are small. In addition, the implementation of such an observer may prove to be simpler because the relative gain of channels can be fixed rather than adaptive.

It appears that the human visual system has adopted the fixed whitening strategy mentioned above, because others and ourselves have shown that human performance is not immune to variations in spectral slope and that access to information is maximized primarily for stimuli whose spectral slopes are similar to the ones that characterize natural scenes. This nonadaptive whitening strategy is illustrated in Fig. 5B, where we show the raw response of four bandpass channels to perfectly symmetric $1/f^2$ noise. Channels are spatially restricted by an integration region whose dimensions are proportional to the scale of the analysis, and their underlying spatial filters have constant-volume point-spread functions and constant-octave frequency bandwidths that collectively act as a fixed whitening filter tuned to images with $1/f^2$ power spectra. Note that all four channels respond with roughly equal amplitude and that this response is independent of the area covered by the integration region. If spectral slopes deviate too much from $1/f^2$, these channels do not respond equally.

In this subsection we have compared the performance of model and human observers to isolate two key components that determine human performance, namely, the integration region for symmetry and the inability to whiten images whose power spectra deviate substantially from those of natural scenes. In the following subsection, we incorporate these and other known components of human vision into a model of symmetry detection in broadband noise patterns of variable spectral slopes.

B. Model of Human Performance

Our model observer consists of four linear channels composed of spatial filters that tile the frequency domain into isotropic log-Gaussian bands of 1.1 octaves. Channels have peak frequencies that are logarithmically spaced from 5.7 to 45.3 cpi, and their underlying filters have equal peak amplitudes in the Fourier domain. The Fourier kernel of filters from the nth channel is given by

$$H_n(u, v) = \exp\left\{ -\frac{1}{2}\left[\frac{\ln(f/f_n)}{\ln(f_n\sigma)}\right]^2\right\}, \tag{3}$$

where f_n is the channel's peak frequency and σ determines the channel's frequency bandwidth in octaves. The spatial profile of filters from the nth channel is given by the reverse Fourier transform of H, namely, $h_n(x, y) = F^{-1}[H_n(u, v)]$. Since all frequency components of H are set to cosine phase, filters obtained by the reverse Fourier transform are spatially localized.

In the spatial domain, each channel is restricted by a Gaussian window that limits the region over which symmetry is integrated. We define the window w of the nth channel as

$$w_n(x, y) = \exp\left[-\frac{(x - x_0)^2}{2\sigma_n{}^2}\right]\exp\left[-\frac{(y - y_0)^2}{2\sigma_n{}^2}\right], \tag{4}$$

where x_0 and y_0 determine the window's spatial location and σ_n determines the window's width and height dimensions. To maintain a constant amount of information across scales, we keep the dimensions of the integration region proportional to the filter scale; that is, σ_n is inversely proportional to the channel peak frequency as given by $\sigma_n = k/f_n$, where k is a scaling constant that for the purpose of our model is set to 1. Although Dakin and Herbert[5] report that the integration region is elongated

by approximately a factor of 2:1 along the axis of symmetry, our model assumes a circular window since the precise shape of the integration region is not crucial to the model's behavior.

Having specified the nature of the channels, their underlying filters, and the spatial region that they cover, we can obtain the response of the nth channel C for stimulus I by

$$C_n(x, y) = w_n(x, y)[h_n(x, y) \otimes I(x, y)], \tag{5}$$

where $\otimes$ denotes the convolution operator. Figure 5B (panels a–d) illustrates the raw output of each of the four channels when presented with a perfectly symmetric $1/f^2$ noise pattern. Note how the response amplitude is similar across channels and also how the integration region decreases as a function of spatial frequency to let through only ~4 spatial cycles of information before tapering off.

Our model assumes that mirror symmetry is computed separately for each channel and that measures of symmetry are subsequently combined by use of a probability summation that weighs each channel by its contrast response. Our measure of symmetry for the nth channel is denoted by S_n and is given by

$$S_n = \frac{\sum_{y} \sum_{x=1}^{X/2-1} |C_n(x, y) C_n(-x, y)|}{\left[\sum_{y} \sum_{x=1}^{X/2-1} C_n(x, y)^2 \sum_{y} \sum_{x=1}^{X/2-1} C_n(-x, y)^2 \right]^{1/2}}. \tag{6}$$

The measure S_n consists essentially of a cross correlation between the spatial structure on one side of the axis and the mirror reflection of the other side. We are aware that measuring symmetry in this way is mathematically convenient[6] and that more physiologically plausible implementations of symmetry-detection mechanisms (including our own) have been proposed.[4,8,12] However, the focus of our model is on scale combination rather than on the mechanisms that actually carry out the computation of symmetry. Using cross correlation takes away little from the relevance of our model to scale combination in symmetry perception, but the replacement of S_n by other measures of symmetry can be accommodated easily if strict compliance with physiological principles is desired.

We have shown in both experiments that human observers lack the ability to whiten stimuli; this follows directly from the fact that human performance varies as a function of spectral slope. Our model, however, would be insensitive to contrast differences in various frequency bands if it simply measured symmetry separately for each channel and then combined these measures across channels without first weighting them. For instance, if one channel received less contrast energy than a second channel, the two channels would each return a measure of symmetry bounded between 0 and 1, but these measures would not reflect whether contrast activates one channel more strongly than the other. Thus, for model performance to depend on spectral slope, S_n must necessarily be weighted by some function on the response magnitude of channel C_n before the stage where S_n is pooled across channels. To do this weighting we need to define a metric for the magnitude of channel response.

Ideally, a metric of channel response would reflect local rather than global amplitude. To illustrate this distinction we consider again Fig. 5B. Although the four channels respond with the same amplitude to $1/f^2$ noise, their global contrast energy is not the same because low-frequency channels have larger integration regions and consequently integrate energy over larger spatial extents than high-frequency channels. For this reason, a metric based on mean energy (e.g., rms) would be inappropriate since it would report that filters bounded by the integration region do not respond equally to $1/f^2$ noise. By comparison, a metric such as Michelson contrast would correctly show that filters in the integration region respond equally to $1/f^2$ noise, but because this metric uses minimum and maximum values in its computations it is inherently unreliable with stochastic stimuli such as random noise. A better metric of channel response consists in computing what we call local energy density (LED), a measure that essentially integrates contrast energy over a region of specified dimensions and divides by the area of this region. We define LED for the nth channel as

$$\mathrm{LED}_n = \frac{\sum\sum [C_n(x, y) - \bar{C}_n]^2}{{\sigma_n}^2}, \tag{7}$$

where contrast energy is integrated over the whole spatial extent of channel C_n and is divided by a factor that is proportional to the area of the integration region, namely, ${\sigma_n}^2$. Using numerical simulations, we verified that the LED metric is constant across channels for images with $1/f^2$ power spectra. Our LED metric is similar to rms contrast with the exception that it takes into consideration the area over which contrast energy is distributed, much like the "rectified contrast spectrum" metric proposed by Field and Brady.[31]

The purpose of using LED_n to weigh S_n is to tune our model specifically to images with $1/f^2$ power spectra. However, it also makes our model heavily dependent on the overall contrast of the stimulus. For instance, the model would erroneously report that the amount of symmetry in a stimulus increased by a factor of 2 if its overall contrast were doubled. To prevent this, our model weighs S_n by relative LED across channels rather than by the absolute LED. We denote the relative LED for the nth channel by α_n, a measure obtained by normalizing the LED of each channel by the sum of LED's from all channels, as given by

$$\alpha_n = \frac{\mathrm{LED}_n}{\sum_{n} \mathrm{LED}_n}. \tag{8}$$

This procedure effectively acts as a broadly tuned gain control mechanism that scales the outputs of all channels by a constant such that their cumulative output sums to 1.

To combine measures of symmetry S_n across channels and form a composite measure S, we first weigh S_n by relative LED_n, namely, α_n. We then use the Quick pooling rule,[46] which assumes an exponent p of 2.0 that corresponds to probability summation given by

$$S = \sqrt[p]{\frac{\sum_n (\alpha_n S_n)^p}{N}}. \quad (9)$$

We have no *a priori* reasons for selecting probability summation over, say, linear summation ($p = 1.0$) apart from the fact that probability summation across channels is frequently encountered in studies of vision.[44,45] However, as we mentioned above, competitive interactions between channels such as winner-take-all rules ($p = \infty$) are unlikely since the win of a single band is not expected to lead to better performance even when all spatial scales are equally visible. The scalar value S constitutes the model's judgment of mirror symmetry and falls on a continuum bounded by 0 (perfectly random) and 1 (perfectly symmetric).

C. Simulations

We implemented our model in the MATLAB environment and presented our model observer with stimuli generated in the same manner as for experiments 1 and 2. For each stimulus the model returned a single value S [see Eq. (9)] that corresponded to its judgment of symmetry. We describe the simulations and their results for each of the experiments.

For the simulation of experiment 1 the model measured symmetry in 250 symmetric–nonsymmetric stimuli pairs for each cell of a 15×7 matrix containing 15 spectral slopes (β) and 7 levels of symmetry (ξ). The 14 spectral slopes covered the range $-2 \leqslant \beta \leqslant 5$ in equal steps of 0.5, and the 7 levels of symmetry covered the range $0 \leqslant \xi \leqslant 1$ in equal steps of 0.5. Overall, the model judged symmetry in 58,702 images. Once the simulation was completed, we computed d' for each cell of the 15×7 matrix by comparing judgments for symmetric and nonsymmetric stimuli. d' is a dimensionless measure of the ability of the model to discriminate between symmetric and nonsymmetric images that bears a monotonic relationship to percent correct.[47,48] For each spectral slope we fitted a separate exponential function to d' versus ξ and estimated the level of symmetry that corresponded to 81% correct performance in a two-alternative forced-choice task, namely, $d' = 1.24$. Model performance for this experiment is plotted in Fig. 6A and should be compared with human performance as shown in Fig. 2.

Results from this simulation reveal that the model shows significant facilitation for images with $1/f^2$ power spectra and performs more poorly for spectral slopes that are either steeper or shallower than $1/f^2$. Although the peak in model performance is considerably sharper than in human performance, the model's symmetry-detection thresholds fall roughly within the same range as human data ($0.35 \leqslant \xi \leqslant 0.65$) and is optimally sensitive to images with power spectra similar to those of natural scenes.

For experiment 2 the model measured symmetry in 250 symmetric–nonsymmetric stimuli pairs for each cell of a 4×4 matrix that correspond, respectively, to 4 spectral slopes and 4 phase-randomized frequency bands. In all, the model evaluated symmetry for 8,000 images. The four spectral slopes are those of experiment 2, namely, $\beta = -2, 0, 2, 4$. The four frequency bands are exactly 1 octave wide, and their center frequencies are logarithmically spaced from 5.7 to 45.3 cpi. These center frequencies coincide with those of the four channels used in our model to speed up computations, but the qualitative behavior of the model would not be affected if more channels were involved, provided that their center frequencies were logarithmically spaced. As with experiment 2, only one of the four octave bands was phase randomized for each stimulus presentation. For the purpose of comparison, we also included a condition in which none of the bands was phase randomized. To maintain model performance in roughly the same range as that of human observers, we set the level of symmetry for the entire simulation to $\xi = 0.5$. Once the simulation was complete, values of d' between judgments made on symmetric and nonsymmetric images were computed separately for each cell of the 4×4 matrix. Values of d' were also mea-

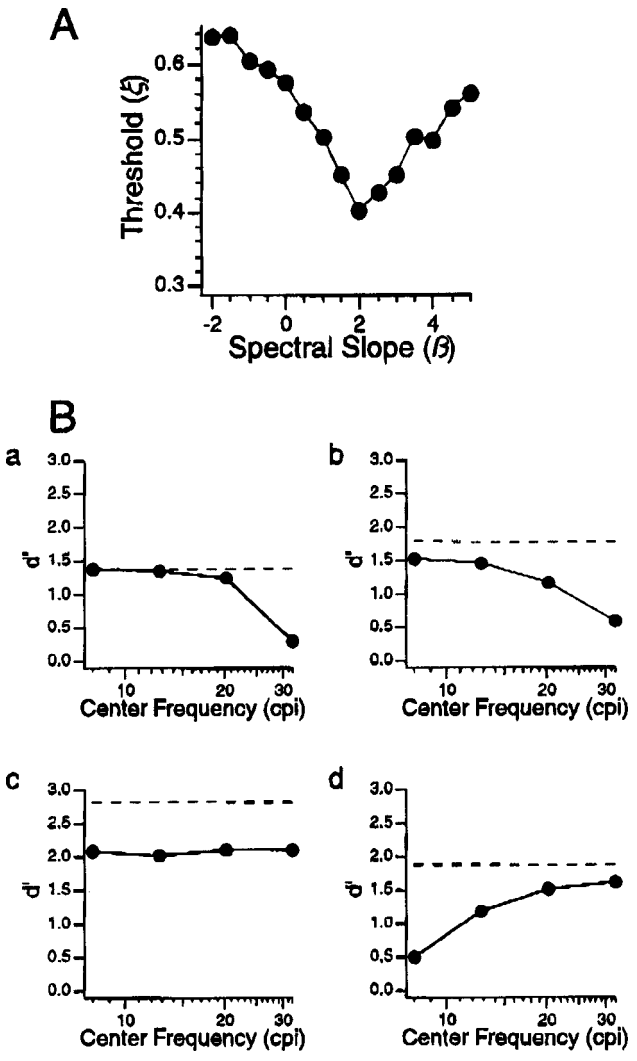

Fig. 6. Modeling results. A, Simulation for experiment 1. Symmetry-detection thresholds are plotted as a function of spectral slope. Results should be compared with human performance (Fig. 2). B, Simulation results for experiment 2. Model performance (d') plotted as a function of the center frequency of the randomized octave band. Dashed line indicate reference performance when no bands are phase randomized. Performance is shown for power spectra of $1/f^{-2}$, $1/f^0$, $1/f^2$, and $1/f^4$.

sured for the comparison condition in which none of the bands was phase randomized. Model performance for this simulation is shown in Fig. 6B and should be compared with human performance in Fig. 4.

From this simulation we note that the model performs in a similar way to human subjects. For spectral slopes that are shallower than 2, the model overemphasizes high spatial frequencies and does not have the capability to whiten the image to equate performance across channels. For spectral slopes greater than 2, the model relies too heavily on low frequencies and again lacks the ability to whiten the stimulus. However, for a spectral slope of 2, phase randomizing of constant-octave frequency bands produces equal deficits in performance across all channels. Note that, unlike for human observers, performance is overall better for $1/f^2$ condition. This, however, is simply a consequence of the fact that the level of symmetry for the simulation was fixed to $\xi = 0.5$ for all spectral slopes, whereas the levels of symmetry for human observers were set separately for each spectral slope to avoid floor and ceiling effects to which d' is immune.

In this section we have shown that the main results from both experiments are qualitatively well predicted by a model that combines across spatial channels with constant-octave bandwidths and equal peak sensitivities but that integrates symmetry over a region whose dimensions are proportional to spatial scale. Our model also correctly predicts that symmetry detection in broadband images relies equally on all spatial scales, provided that the contrast structure of the stimulus matches the spectral profile expected from natural scenes. Although our model makes several assumptions about early spatial vision, we believe it nonetheless captures some of the key aspects of the spatial mechanisms that mediate the detection of symmetry in broadband noise patterns.

7. GENERAL SUMMARY

The present study has produced the following findings:

- Symmetry detection is optimal in noise patterns whose spectral slopes fall within the range that characterizes natural scenes ($1.2 \leqslant \beta \leqslant 3.2$).
- Symmetry detection in noise patterns is affected equally by the phase randomization of constant-octave frequency bands only for $1/f^2$ power spectra.
- Together, our results imply that symmetry mechanisms extract equal amounts of information across scales but that the visual system lacks the ability to whiten images whose spectral slopes deviate substantially from those of natural scenes.
- Symmetry detection in broadband noise is qualitatively well predicted by a multichannel model in which contrast-weighted symmetry judgments from all channels are combined through probability summation into a single metric. The key components of the model are (a) an integration region whose dimensions are proportional to filter scale and that keeps a constant amount of information across scales and (b) linear spatial filters with constant-volume point-spread functions and constant-octave bandwidths that determine access to spatial scales and effectively tune the model to images with $1/f^2$ power spectra.

8. GENERAL DISCUSSION

Results from this study are consistent with the notion that the detection of symmetry is efficient and is intimately tied to the spatial characteristics of filtering mechanisms that mediate early vision. As mentioned in Section 1, several authors have suggested that early visual mechanisms are optimally tuned to second-order statistics (i.e., the power spectrum) of natural scenes. In particular, since spatial filters have roughly equal Fourier amplitudes and constant-octave bandwidths, filter outputs are approximately constant across spatial scales in images with $1/f^2$ spectral profiles. Although these constant outputs ensure access to a vast selection of spatial scales, little would be accomplished if subsequent stages of processing did not use spatial scales equally in their analysis. In line with this idea, we have shown that symmetry detection is optimal with patterns whose power spectra approximately match the second-order statistics of the ensemble of natural scenes and that symmetry mechanisms weigh constant-octave frequency bands equally. Consequently, symmetry detection takes advantage of the efficiency of lower-level mechanisms tuned to the statistics of natural scenes.

At first glance, our proposal that symmetry detection is efficient is not well supported by the fact that the integration region has been shown to vary with the spatial scale of the stimulus.[5] In fact, a constraint of maximal efficiency such as an ideal observer (see Section 6) would demand that symmetry be computed over the entire visual field rather than over roughly 4 cycles of spatial scale. However, for a biological visual system such as ours, efficiency is perhaps better expressed as the optimal compromise between selecting mechanisms that encode key information for survival and limiting the neural costs of implementation. Restricting the region of integration for symmetry to a few cycles at each spatial scale may reduce neural costs and may also be optimal in the sense that it distributes the information load evenly across spatial channels. In addition, encoding symmetry information over large regions is perhaps not the optimal strategy because mirror symmetry in natural scenes is imperfect and is likely limited to spatial regions near the axis.

It is perhaps worth restating that channels with constant-octave bands of frequencies are not expected to contribute equally to the detection of features that, unlike symmetry, are computed over the entire extent of the stimulus. As we mentioned in Section 6, high spatial frequencies are intrinsically more informative than low frequencies simply because the spatial density of information is proportional to the scale of the filter used. In the case of symmetry, the integration region depends on the spatial frequency of the stimulus and serves to keep the amount of information constant across scales. However, should a similar experiment be carried out with a feature other than mirror symmetry, performance might still be optimal with $1/f^2$ profiles, but high frequencies might be overrepresented if neural computations were not re-

stricted to a fixed number of cycles at each spatial scale. Perhaps constructing a stimulus that physically limits the information content to a few cycles at each spatial scale would partially solve this problem.

9. CONCLUSIONS

The present study has shown that symmetry detection in broadband noise is well predicted by a model that relies on a bank of linear spatial filters of various scales. Although such a linear filtering stream may be an important component of symmetry detection, it is difficult to generalize our findings to other classes of symmetric stimuli since random noise lacks the higher-order statistics needed to reveal mechanisms that may detect symmetry through nonlinear filtering or feature matching. Challenges for the future are to characterize the properties of the various mechanisms that mediate the computation of mirror symmetry and, ultimately, to integrate them into a comprehensive account of the perception of symmetry and visual form in general.

ACKNOWLEDGMENTS

The authors are grateful to David Field and Steve Dakin for their useful suggestions and comments on this study. This research was funded by a Canadian Natural Science and Engineering Research Council postgraduate scholarship to S. J. M. Rainville and grants from the Canadian Natural Science and Engineering Research Council and the Canadian Medical Research Council to F. A. A. Kingdom.

Address correspondence to S. Rainville; e-mail, stephane@jiffy.vision.mcgill.ca.

REFERENCES

1. R. L. De Valois, D. G. Albrecht, and L. G. Thorell, "Spatial frequency selectivity of cells in macaque visual cortex," Vision Res. **22**, 545–559 (1982).
2. F. W. Campbell and J. G. Robson, "Application of Fourier analysis to the visibility of gratings," J. Physiol. (London) **197**, 551–566 (1968).
3. B. Julesz and J. Chang, "Symmetry perception and spatial-frequency channels," Perception **8**, 711–718 (1979).
4. S. C. Dakin and R. F. Hess, "The spatial mechanisms mediating symmetry perception," Vision Res. **37**, 2915–2930 (1997).
5. S. C. Dakin and A. M. Herbert, "The spatial region of integration for visual symmetry detection," Proc. R. Soc. London, Ser. B **265**, 659–664 (1998).
6. B. Jenkins, "Component processes in the perception of bilaterally symmetric dot textures," Percept. Psychophys. **34**, 433–440 (1983).
7. H. B. Barlow and B. C. Reeves, "The versatility and absolute efficiency of detecting mirror symmetry in random dot displays," Vision Res. **19**, 783–793 (1979).
8. S. C. Dakin and R. J. Watt, "Detection of bilateral symmetry using spatial filters," Spatial Vision **8**, 393–413 (1994).
9. J. Wagemans, L. Van Gool, V. Swinnen, and J. Van Horebeek, "Higher-order structure in regularity detection," Vision Res. **33**, 1067–1088 (1993).
10. C. W. Tyler, L. Hardage, and R. T. Miller, "Multiple mechanisms for the detection of mirror symmetry," Spatial Vision **9**, 79–100 (1995).
11. V. G. Bruce and M. J. Morgan, "Violations of symmetry and repetition in visual patterns," Perception **4**, 239–249 (1975).
12. S. J. M. Rainville and F. A. A. Kingdom, "Does oblique structure support the detection of mirror symmetry?" Invest. Ophthalmol. **39**, S170 (1998).
13. S. J. M. Rainville and F. A. A. Kingdom, "From spatial filters to mirror symmetry: new findings and new model," Perception **27**, 58a (1998).
14. U. Koeppl, "Local orientation versus local position as determinants of perceived symmetry," Perception **22**, 111 (1993).
15. P. Wenderoth, "The salience of vertical symmetry," Perception **23**, 221–236 (1994).
16. J. Wagemans, L. Van Gool, and J. Van Horebeek, "Orientation selective channels in symmetry detection: effects of cooperation and attention," in *Channels in the Visual Nervous System: Neurophysiology, Psychophysics and Models*, B. E. Blum, ed. (Freund, London, 1991), pp. 425–445.
17. P. Wenderoth, "The effects of the contrast polarity of dot-pair partners on the detection of bilateral symmetry," Perception **25**, 757–771 (1996).
18. L. Zhang, "Symmetry perception in human vision," Ph.D dissertation (University of Trieste, Trieste, Italy, 1991).
19. P. Carlin, "On symmetry in visual perception," Ph.D dissertation (University of Stirling, Stirling, Scotland, 1996).
20. C. W. Tyler and L. Hardage, "Mirror symmetry detection: predominance of second-order pattern processing throughout the visual field," in *Human Symmetry Perception and Its Computational Analysis*, C. W. Tyler, ed. (VSP, Utrecht, The Netherlands, 1996).
21. S. Oomes, "Human visual perception of spatial structure: symmetry, orientation, and attitude," Ph.D dissertation (Catholic University of Nijmegen, Nijmegen, The Netherlands, 1998).
22. D. H. Hubel and T. N. Wiesel, "Receptive fields and functional architecture of monkey striate cortex," J. Physiol. (London) **195**, 215–243 (1968).
23. F. Labonte, Y. Shapira, P. Cohen, and J. Faubert, "A model for global symmetry detection in dense images," Spatial Vision **9**, 33–55 (1995).
24. B. Jenkins, "Redundancy in the perception of bilateral symmetry in dot textures," Percept. Psychophys. **32**, 171–177 (1982).
25. N. Brady and D. J. Field, "What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns," Vision Res. **35**, 739–756 (1995).
26. D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," J. Opt. Soc. Am. A **4**, 2379–9234 (1987).
27. G. J. Burton and I. R. Moorhead, "Color and spatial structure natural scenes," Appl. Opt. **26**, 157–170 (1987).
28. D. J. Tolhurst, Y. Tadmor, and T. Chao, "Amplitude spectra of natural images," Ophthal. Physiol. Opt. **12**, 229–232 (1992).
29. D. L. Ruderman and W. Bialeck, "Statistics of natural images: scaling in the woods," Phys. Rev. Lett. **73**, 814–817 (1994).
30. D. J. Field, "Scale-invariance and self-similar 'wavelet' transforms: an analysis of natural scenes and mammalian visual systems," in *Wavelets, Fractals, and Fourier Transforms*, M. Farge, J. C. R. Hunt, and J. C. Vassilicos, eds. (Clarendon, Oxford, 1993), pp. 151–193.
31. D. J. Field and N. Brady, "Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes," Vision Res. **37**, 3367–3383 (1997).
32. J. J. Atick and A. N. Redlich, "Towards a theory of early visual processing," Neural Comput. **2**, 308–320 (1990).
33. M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: a fresh view of inhibition in the retina," Proc. R. Soc. London, Ser. B **216**, 427–59 (1982).
34. C. Blakemore and F. W. Campbell, "On the existence of neurones in the human vision system selectively sensitive to the orientation and size of retinal images," J. Physiol. (London) **203**, 237–260 (1969).
35. H. R. Wilson, D. K. McFarlane, and G. C. Phillips, "Spatial frequency tuning of orientation selective units estimated by oblique masking," Vision Res. **23**, 873–882 (1983).

36. L. Croner and E. Kaplan, "Receptive fields of P and M ganglion cells across the primate retina," Vision Res. **35**, 7–24 (1995).
37. D. C. Knill, D. Field, and D. Kersten, "Human discrimination of fractal images," J. Opt. Soc. Am. A **7**, 1113–1123 (1990).
38. Y. Tadmor and D. J. Tolhurst, "Discrimination of changes in the second-order statistics of natural and synthetic images," Vision Res. **34**, 541–554 (1994).
39. D. J. Tolhurst and Y. Tadmor, "Band-limited contrast in natural images explains the detectability of changes in the amplitude spectra," Vision Res. **37**, 3203–3215 (1997).
40. D. G. Pelli, "The VideoToolbox software for visual psychophysics: transforming numbers into movies," Spatial Vision **10**, 437–442 (1997).
41. A. B. Watson and D. G. Pelli, "QUEST: A Bayesian adaptive psychometric method," Percept. Psychophys. **33**, 113–120 (1983).
42. J. D. Victor and M. M. Conte, "The role of high-order phase correlations in texture processing," Vision Res. **36**, 1615–1631 (1996).
43. S. J. M. Rainville and F. A. A. Kingdom, "Is motion perception sensitive to local phase structures?" Invest. Ophthalmol. **38**, S215 (1997).
44. N. V. S. Graham, *Visual Pattern Analyzers* (Oxford U. Press, New York, 1989).
45. N. Graham and J. G. Robson, "Summation of very close spatial frequencies: the importance of spatial probability summation," Vision Res. **27**, 815–826 (1987).
46. R. F. J. Quick, "A vector-magniture model of contrast detection," Kybernetik **16**, 65–67 (1974).
47. N. A. Macmillan and C. D. Creelman, *Detection Theory: A User's Guide* (Cambridge U. Press, Cambridge, 1991).
48. D. M. Green and J. A. Swets, *Signal Detection Theory and Psychophysics* (Peninsula, Los Altos, Calif., 1988).