



The functional role of oriented spatial filters in the perception of mirror symmetry — psychophysics and modeling

Stéphane J.M. Rainville *, Frederick A.A. Kingdom

McGill Vision Research Unit, Room H4-14, 687 Pine Ave W., Montréal, Québec, Canada H3A 1A1

Received 19 July 1999; received in revised form 11 April 2000

Abstract

We investigated human sensitivity to vertical mirror symmetry in noise patterns filtered for narrow bands of variable orientations. Sensitivity is defined here as the amount of spatial phase randomization corresponding to 75% correct performance in a 2AFC detection task. In Experiment 1, sensitivity was found to be high for tests patterns of all orientations except those parallel to the axis of symmetry. This implies that corresponding mirror-orientations (e.g. -45° and $+45^\circ$) are combined prior to symmetry detection. In Experiment 2, observers detected symmetry in tests of variable orientation in the presence of either non-symmetric or symmetric masks filtered for orientations either parallel or perpendicular to the axis. Observers were found to be primarily affected by masks of the same orientation as the test, thus suggesting that symmetry is computed separately in distinct mirror-orientation channels. In Experiment 3, observers detected a symmetric test of variable height and width embedded in random noise. Data revealed that mirror symmetry is computed over a spatial integration region (IR) that remains approximately constant in area but whose height-to-width aspect ratio changes from 20:1 to 2:1 as orientation is varied from parallel to perpendicular to the axis. We compare human data against that of an ideal observer to identify key factors that limit visual performance and discuss the implications for the functional architecture of symmetry perception. We also propose a multi-channel model of symmetry detection that combines the output of oriented spatial filters in a simple and physiologically plausible manner. Particular emphasis is placed on the notion that changes in the shape of the IR with orientation compensate for changes in *information density* and partially equate performance across orientations. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Symmetry; Orientation; Texture; Spatial filters; Psychophysics; Ideal observer; Modeling

1. Introduction

Mirror symmetry is an image property to which the human visual system is highly sensitive. However, only a limited number of studies have taken an approach linking the perception of mirror symmetry to known low-level visual mechanisms (Dakin & Watt, 1994; Dakin & Hess, 1997; Rainville, 1999; Rainville & Kingdom, 1999b). This is somewhat surprising given that psychophysical (e.g. Campbell & Robson, 1968; Wilson, McFarlane & Phillips, 1983; Polat & Sagi, 1993), physiological (e.g. Hubel & Wiesel, 1968; DeValois, Albrecht & Thorell, 1982) and theoretical (e.g.

Marr & Hildreth, 1980; Olshausen & Field, 1996) research has converged on a model of early vision involving spatial mechanisms — or *spatial filters* — with localized, bandpass, and oriented properties. Although recent psychophysical evidence suggests that symmetry detection recruits mechanisms such as spatial filters (see below), several key issues remain. In particular, the role of orientation tuning in the computation of mirror symmetry is only partially understood primarily because past studies have relied on stimuli whose spatial structure was either orientation non-specific (isotropic) or limited to special cases of the orientation spectrum (e.g. components parallel or perpendicular to the axis of symmetry).

The purpose of this study is to provide insight into the functional role of oriented spatial filters in the perception of mirror symmetry. In the remainder of Section 1, we briefly review existing psychophysical

* Corresponding author. Present address: Center for Visual Science, 274 Meliora Hall, University of Rochester, Rochester, NY, 14627, USA. Fax: +1-716-271-3043.

E-mail address: stephane@cvs.rochester.edu (S.J.M. Rainville).

data linking the computation of symmetry to spatial filters, and in doing so, we present in more detail our motivations for this study.

1.1. Spatial filters and symmetry detection

Increasing psychophysical evidence indicates that the neural code for mirror symmetry is intimately tied to the output of mechanisms localized both in space and in spatial frequency. In an early demonstration, Julesz and Chang (1979) have shown that mirror symmetry cannot be perceived in a combined horizontally-symmetric and vertically-symmetric noise pattern, but that symmetry is perceived if the two noise patterns are filtered into spatial-frequency bands that differ by two octaves. Access to symmetry is therefore only possible if spatial-frequency decomposition precedes the computation of symmetry. Using bandlimited noise patterns of variable spatial scale, Dakin and Herbert (1998) have reported that symmetry is scale invariant because it is computed over a limited spatial integration region (IR) whose dimensions are proportional to the spatial scale of the stimulus. In addition, Dakin and Hess (1997) found that performance for detecting symmetry in bandpass noise patterns is constant for all spatial scales. Rainville and Kingdom (1999a) have shown that symmetry detection in broadband noise is optimal if contrast energy is distributed evenly across spatial scales — that is, if stimuli have power spectra that decay with the square of spatial frequency ($1/f^2$) (Brady & Field, 1995; Field & Brady, 1997). In the same study, Rainville and Kingdom also demonstrated that spatial scales are equally weighted for symmetry detection in broadband stimuli. Taken together, results from these studies suggest that coding for symmetry closely depends on spatial-filter properties and therefore takes place at fairly low-levels in the visual hierarchy.

Studies on symmetric random-dot textures also lend support to the notion that symmetry detection depends on spatial-filter properties, although the evidence is less direct. Interpreting results from experiments involving random dots in terms of spatial filtering is more difficult since random-dot stimuli afford little control over their spatial-frequency and orientation content.

Nevertheless, one study of this type reported that symmetry detection in random-dot displays is resistant to spatial jitter (Barlow & Reeves, 1979) and therefore implies filtering mechanisms operating at coarser scales than that of individual dots. In addition, several random-dot studies note that spatial disruptions near the axis have a greater impact on symmetry perception than disruptions further away from the axis (Bruce & Morgan, 1975; Barlow & Reeves, 1979; Jenkins, 1983a; Zhang, 1991; Tyler, Hardage & Miller, 1995; Wenderoth, 1996), thereby suggesting mechanisms that operate over limited spatial extents. Also, the observation

that symmetry detection is more difficult if the two symmetric halves are of opposite contrast (Zhang, 1991; Carlin, 1996; Wenderoth, 1996) implicates spatial mechanisms that, like quasi-linear spatial filters, are sensitive to phase, although see Rainville (1999), Rainville and Kingdom (submitted), and Tyler and Hardage (1996) for counter-evidence.

1.2. Oriented filters and symmetry detection

Despite evidence that symmetry perception is mediated by mechanisms such as spatial filters, comparatively few studies have established whether orientation tuning plays a role in the computation of symmetry. In a brief report dealing with this issue, Koepl and Morgan (1993) measured symmetry detection for stimuli consisting of hard-edge line segments of variable orientation and position. The authors reported that symmetry detection is independent of the orientation of the line segments. In a similar study, Wagemans, Van Gool and Van Horebeek (1991) reported a slight advantage for patterns whose hard-edge line segments have oblique orientations ($\pm 45^\circ$) with respect to the axis of symmetry. Labonte, Shapira, Cohen and Faubert (1995) showed that grouping by orientation can precede symmetry coding for stimuli composed of hard-edge line segments with variable orientation. Finally, Joung, van der Zwan and Latimer (2000) and van der Zwan, Leo, Joung, Latimer and Wenderoth (1998) reported that adapting to mirror symmetry can induce tilt after-effects, thus suggesting that oriented mechanisms are also involved in coding for symmetry. However, results obtained with textures composed of hard-edge elements such as dots or oriented segments are still difficult to relate to mechanisms such as spatial filters since, unlike bandlimited stimuli, their contrast energy is broadly distributed over scale and orientation.

Using an alternative approach to hard-edge microelements, Dakin and Hess (1997) have shown that symmetry detection in noise patterns is more difficult if stimuli are spatially filtered for orientations parallel to the axis than for orientations perpendicular to the axis. The authors have also demonstrated that symmetry detection thresholds in isotropic noise are comparable to those for noise filtered perpendicularly to the axis and therefore conclude that filters perpendicular to the axis play a predominant role in symmetry detection. The discrepancy between these results and those using oriented line segments is likely attributable to better control over stimuli's scale and orientation content. Stimuli from Dakin and Hess (1997) were also of a more textural nature since they lacked local phase alignments needed to produce features such as oriented bars and edges (Burr, Morrone & Spinelli, 1989).

In short, few studies have investigated how symmetry perception depends on the orientation *content* of the

stimulus. In addition, interpretation of their data is often complicated by the fact that hard-edge microelements distribute their contrast energy broadly over scale and orientation. When the orientation content has been narrowband, as in the study of Dakin and Hess (1997), only orientations parallel or perpendicular to the axis of symmetry have been investigated. As we illustrate in the next section, however, parallel and perpendicular orientations only constitute special cases that are not representative of the way in which the concept of orientation enters the definition of mirror symmetry.

1.3. Symmetry defined by oblique orientations

Fig. 1 depicts an image with strong components of vertical mirror symmetry (panel a) that has been passed through a bank of spatial filters tuned to different orientations such as to preserve symmetry (panels b–g). This illustrates that the spatial structure of mirror

symmetry can be construed in terms of *pairs* of mirror orientations that are at *oblique* angles to the axis. For instance, panel d illustrates how mirror symmetry is defined for oblique mirror orientations of -60 and $+60^\circ$, respectively. The only two exceptions to this general rule are orientations parallel ($\pm 0^\circ$) or perpendicular ($\pm 90^\circ$) to the axis (panels b and e), although these can be considered as special cases where geometry dictates that orientations on either side of the axis are the same.

The distinction between oblique orientations and the special cases of parallel and perpendicular orientations is important. It highlights the fact that most of the spatial structure of mirror symmetry is defined only by the *joint presence* of orientations that can differ by as much as 90° , as in the case where mirror components are at angles of -45 and $+45^\circ$ with respect to the axis. The fact that symmetry is readily perceived even in images filtered for oblique orientations suggests that coding for symmetry does not rely solely on mecha-

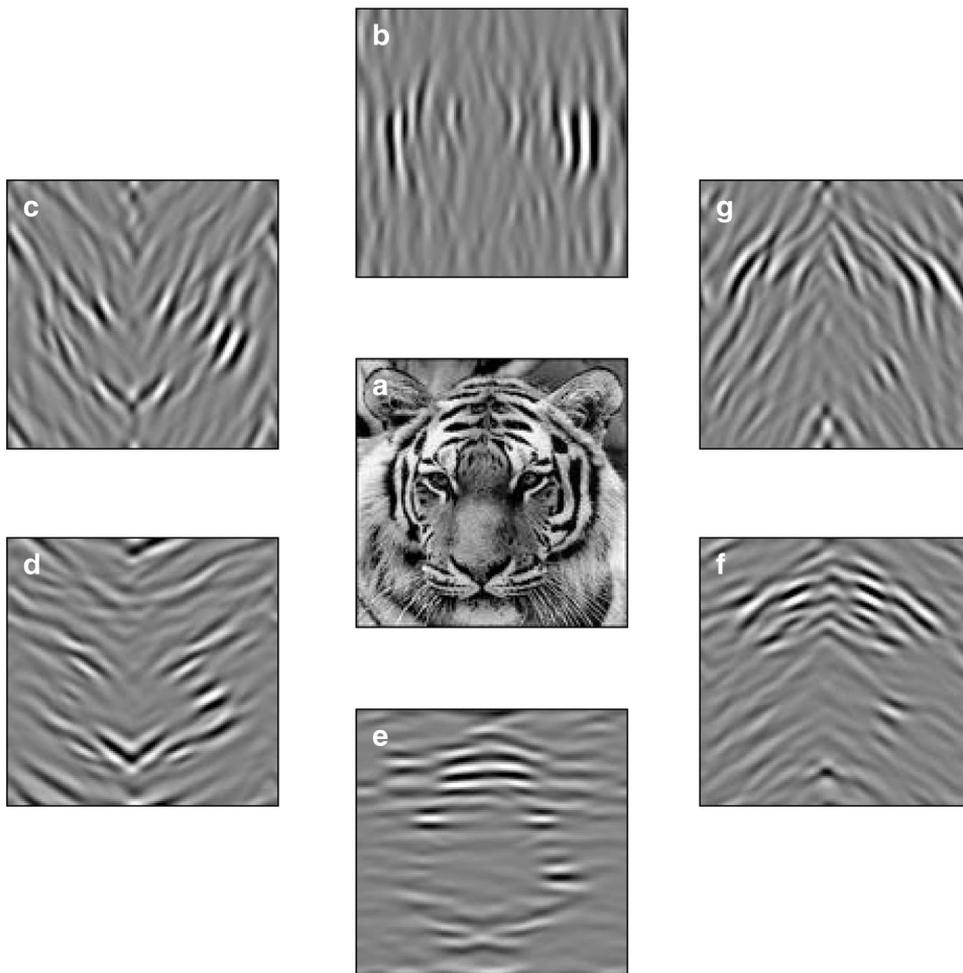


Fig. 1. Mirror symmetry and the orientation spectrum. An image with vertical mirror symmetry (panel a) is filtered with wavelet filters of various orientations such as to preserve symmetry. For orientations parallel (panel b) or perpendicular (panel e) to the axis of symmetry, a single orientation band suffices to encode symmetry. For orientations oblique to the axis (panels c, d, f, g), symmetry can only be encoded if two mirror-orientation bands are considered.

nisms tuned to orientations parallel or perpendicular to the axis. Rather, a fundamental aspect of symmetry perception involves some form of combination of orientations prior to the stage where symmetry is computed. The nature of orientation combinations in symmetry perception is precisely the issue we address in the present study. Part of this research has been presented at various meetings (Rainville & Kingdom, 1998a,b), and has also been reported in the first author's doctoral thesis (Rainville, 1999).

2. General method

2.1. Observers

The first author (SR) acted as an observer in all experiments. SR is slightly myopic but wore corrective eyewear during trials. The second author (FK) participated to Experiments 1, 2, and 4. FK is an emmetrope and needed no optical correction. A naive observer HW with corrected-to-normal vision participated in Experiment. 3.

2.2. Hardware and calibration

All experiments were carried out on a Power Macintosh 7600/120 computer driving a 17 in. Sony Multi-scan color monitor via a standard 8-bit/gun video card. We measured the monitor's luminance versus RGB function with a UDT photometer and linearized the display using the inverse of the best-fitting gamma function to the luminance data. After gamma correction, the linearized look-up table had an effective depth of 7.1 bits/gun and produced a mean luminance of 33.4 cd/m².

2.3. Stimuli

All stimuli used in the present paper shared several characteristics that are described in this section. Stimuli consisted of random-noise patterns that were mirror-symmetric about the vertical axis. Patterns were filtered for various mirror orientations with a bandpass Gaussian wedge, and orientation bandwidth was fixed to 20° (full-width at half-height) across all stimulus conditions. Stimuli were also narrowband filtered for spatial scale using an isotropic log-Gaussian filter with a 1.15 octave bandwidth (full-width at half-height) to avoid effects known to be scale specific (Dakin & Herbert, 1998; Rainville & Kingdom, 1999b).

In order to psychophysically measure observer sensitivity to symmetry, we manipulated the amount of mirror symmetry by adding variable amounts of random jitter to the phase spectrum of the stimulus. The level of possible phase randomization ranges from 0°

(perfect symmetry) to 360° (completely random). Phase randomization has been employed in other studies on symmetry perception (Dakin & Hess, 1997; Dakin & Herbert, 1998) and is increasingly common in other branches of psychophysical research (Victor & Conte, 1996; Rainville & Kingdom, 1997).

Fig. 2 shows patterns with levels of mirror symmetry ranging from perfectly symmetric (phase jitter = 0°) to an intermediate level of symmetry (phase jitter = 180°) to perfectly random (phase jitter = 360°). Notice how phase jitter affects the level of mirror symmetry without otherwise affecting the orientation or spatial-scale content of the stimuli. This is because mirror symmetry is a higher-order (i.e. phase-dependent) property of our stimulus class whereas scale and orientation are lower-order (i.e. phase-independent) properties in the sense that they are only defined by the stimulus power spectrum. Insets in the first column of Fig. 2 filter amplitude spectra for various orientation conditions.

Stimuli were scaled to 25% RMS contrast unless stated otherwise. Patterns were also windowed by a circular aperture whose diameter matched the width of the image. Further technical details on stimulus construction are given in Appendix A.

2.4. Procedure

Stimuli were computed online in the MATLAB environment and were shown on the screen using high-level interfaces from Brainard's PsychToolbox (Brainard, 1997) calling lowlevel routines from Pelli's VideoToolbox (Pelli, 1997).

We should also note that because the spatial structure of our stimuli was spatially homogenous, there were no cues other than mirror symmetry which could be used to locate the axis. This is not the case for images such as those in Fig. 1 since the axis of symmetry can be located by the change in orientation for oblique stimuli (panels c, d, g and f) but not for stimuli with structure parallel (panel b) or perpendicular to the axis (panel e). However, an unfortunate side-effect of this spatial homogeneity is that the amount of physical information present in our stimuli is not independent of orientation. We use an ideal-observer analysis to address this issue in Section 8.1 to ensure that the differential amount of information across orientations was taken into account in the analysis of the results.

In all experiments except the control for Experiment 1, stimuli consisted of a 128 × 128 pixel matrix presented at a viewing distance of 68 cm. At this distance, the spatial frequency content of the image is bracketed between 0.5 and 16.0 cpd. However, these limits were never attained since bandpass filtering ensured that most of the contrast energy was concentrated around 7.0 cpd.

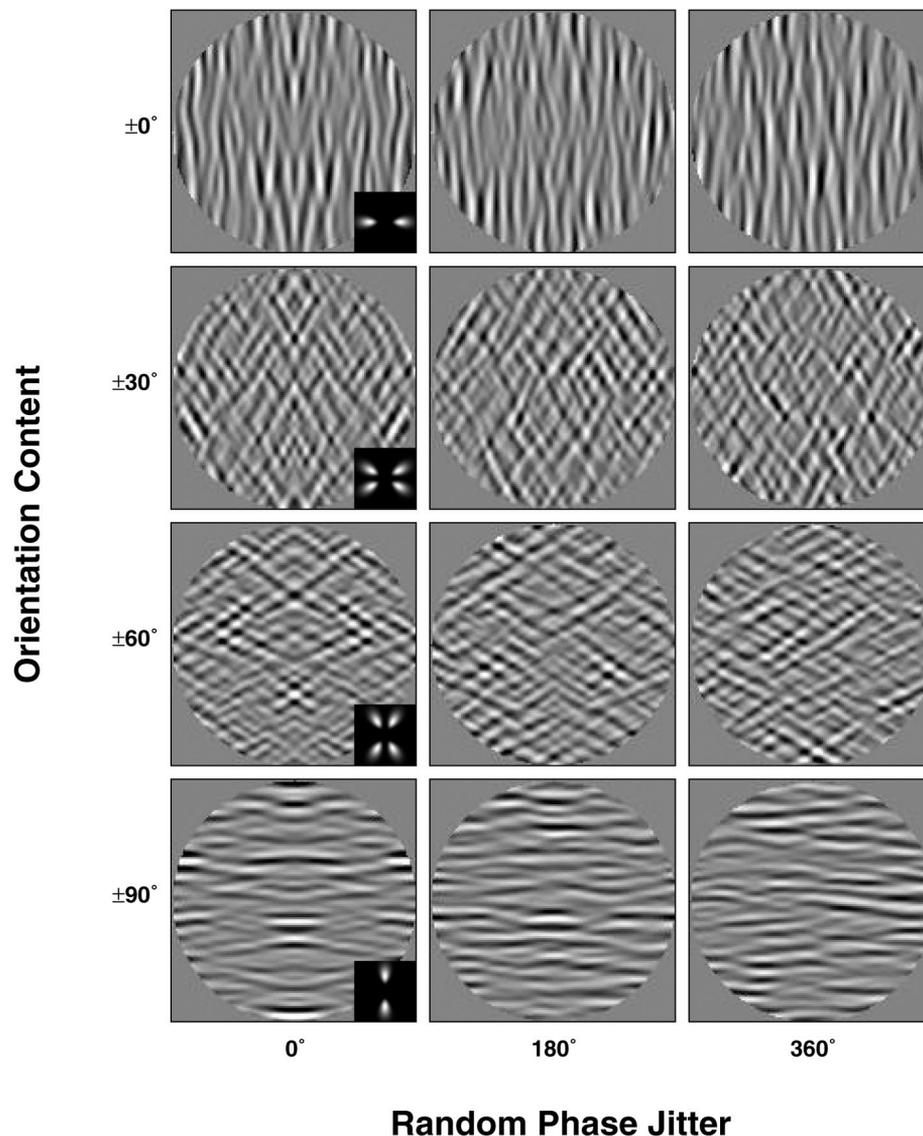


Fig. 2. Stimuli for Experiment 1. Noise patterns with varying levels of mirror symmetry are shown for various orientation bands. Columns correspond to phase randomization levels of 0° (perfect symmetry), 180° (partially symmetric), and 360° (completely random), respectively. Rows illustrate patterns filtered for mirror orientations of $\pm 0^\circ$, $\pm 30^\circ$, $\pm 60^\circ$ and $\pm 90^\circ$, respectively. The Fourier amplitude spectrum of the operators used to filter orientations are shown in the panel insets of the first column.

In all experiments, observers discriminated between symmetric and non-symmetric images in a two-alternative forced-choice (2AFC) paradigm. Observations were collected over several levels of phase jitter. We used a method of constant stimuli since adaptive routines tend to be vulnerable to the intrinsic noise of stochastic stimuli. The order of presentations was randomly interleaved from one trial to the next, and observers pressed one of two keys to report the interval that appeared more symmetric. Images were presented for 500 ms (or the equivalent of 34 screen refreshes at 68 Hz) and were separated by an ISI of 500 ms. A low-contrast fixation dot corresponding to the center of the stimulus was shown before every presentation to ensure that the axis of symmetry was foveated. Observers received auditory feedback on incorrect responses.

Forty observations were made in each run. Although the total number of runs varied across conditions, a minimum of 200 observations were collected for every data point presented in the paper. Psychometric functions were approximated by a two-parameter cumulative normal yielding the best least-squares fit to the data. The level of phase jitter corresponding to 75% correct performance was taken as our measure of the observer's sensitivity to mirror symmetry. Standard-deviation estimates for sensitivity were obtained by resampling the data 100 times using a 'bootstrap' technique (Efron & Tibshirani, 1993). Additional data were collected if the psychometric fit yielded unacceptably large error estimates.

3. Experiment 1: sensitivity versus orientation

As a first step, we measured the ability to detect symmetry as a function of the orientation content of the stimulus. The primary purpose of this manipulation was to establish empirically whether symmetry detection is possible with patterns consisting only of oblique mirror orientations. As noted in Section 1, the case of oblique orientations is particularly interesting because, unlike orientations parallel or perpendicular to the axis, it deals with orientation *combination* in the computation of mirror symmetry.

Observers detected mirror symmetry in random noise patterns filtered for narrow orientations bands. A total of seven mirror orientations were selected at equal intervals of 15° namely ± 0 , ± 15 , ± 30 , ± 45 , ± 60 , ± 75 and $\pm 90^\circ$. Examples of stimuli used in this experiment are shown in Fig. 2 for various orientations and levels of phase jitter. FK and SR acted as observers for this experiment.

We also ran a control experiment in which stimuli dimensions were doubled in size from 128×128 to 256×256 pixels. The purpose of this control was to ensure that performance was not limited by stimulus dimensions (Dakin & Herbert, 1998; Rainville & Kingdom, 1999b). To preserve the same retinal spatial frequency (i.e. 7.0 cpd), the center frequency of the narrowband filter was doubled from 14.0 to 28.0 cpi. Necessarily, the dimensions of the stimulus aperture were also doubled.

Fig. 3 plots sensitivity (filled circles) to symmetry as a function of the orientation content of the stimulus. Results from the control experiment (filled squares) are also plotted. For both observers, symmetry was easily detected over most of the orientation spectrum with the exception of orientations parallel or nearly parallel to the axis of symmetry. The difference in sensitivity between best and worst performance corresponds to roughly 100° of phase jitter, or roughly a factor of 1.5. Results from the control experiment are virtually identi-

cal for SR, but sensitivity for FK increases slightly as the number of spatial cycles is doubled.

Results show that symmetry detection is possible not only for orientations parallel or perpendicular to the axis but also for oblique mirror orientations. Since oblique mirror orientations necessarily come in pairs, results also imply that mirror orientations are somehow combined prior to the computation of symmetry. In addition, poor performance with structure parallel to the axis of symmetry reflects a genuine property of the visual system rather than an artifact of stimulus dimensions since doubling the number of spatial cycles had little effect on overall performance.

4. Experiment 2: cross-orientation masking

This experiment investigates whether coding for symmetry is confined to separate mirror-orientation channels or whether symmetry mechanisms pool across all orientations. To address the issue, we used a masking paradigm in which observers detected a test of one orientation in the presence of a mask of either the same or different orientation. Two different types of masks were used, namely random noise and perfectly symmetric noise. As we describe below, these two types of masks allow us to investigate orientation selectivity at different stages in the computation of symmetry.

Consider an image which contains both a signal (i.e. noise with variable amounts of symmetry) and a mask consisting of pure random noise. If the signal is presented in a different orientation band than the mask, anisotropic filters of the appropriate orientation can recover the signal while avoiding contamination by noise in other orientation bands. By comparison, isotropic filters would recover the signal but would also integrate noise from other orientation bands. This leads to two clear predictions. If the computation of symmetry is mediated by isotropic filters, detecting symmetry in the presence of a random mask would be expected to

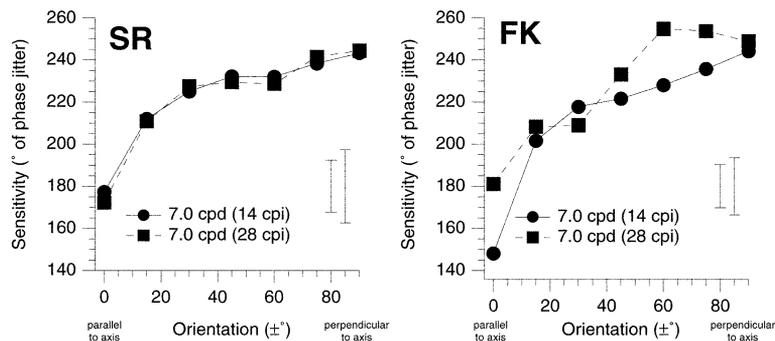


Fig. 3. Results from Experiment 1. Sensitivity to symmetry is plotted as a function of stimulus orientation for observers SR and FK. Filled circles correspond to the 14 cpi condition whereas filled squares indicate performance in the 28 cpi control condition. Maximum and mean standard deviations estimates are also shown.

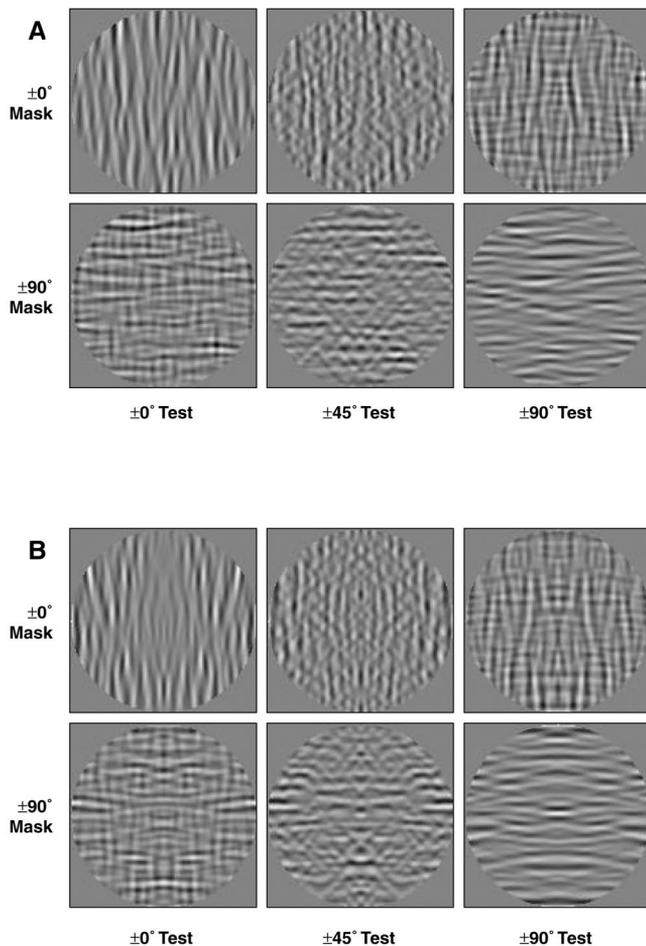


Fig. 4. Stimuli for Experiment 2. (A) Random masks: a test of variable orientation ($\pm 0^\circ$, $\pm 45^\circ$, $\pm 90^\circ$) is detected in the presence of a *random* mask filtered for orientations parallel (first row) or perpendicular (second row) to the axis of symmetry. (B) Same as A but for *symmetric* masks.

be poor irrespective of the orientation of the mask. Alternatively, if the computation of symmetry involves oriented filters, performance would be expected to be poor only in the presence of a mask of the same orientation as the signal.

While random-noise masks allow us to determine whether mirror orientations are pooled *before* symmetry is encoded, they do not test for the possibility that symmetry information is pooled across orientation *after* symmetry has been computed separately in different mirror-orientation channels. Thus, to test whether the presence of symmetry in one orientation band would affect the ability to detect symmetry in another orientation band, we measured symmetry detection in the presence of *symmetric* masks composed of perfect vertically-symmetric noise patterns filtered at orientations either parallel or perpendicular to the axis of symmetry. If symmetry information is combined across all mirror-orientation channels, the presence of the strongly symmetric mask in one orientation band may disrupt the

detection of the weaker test signal in another orientation band, much like a suprathreshold pedestal can be disruptive to the detection of a test (Legge & Foley, 1980). Conversely, if symmetry is encoded separately in different mirror-orientation channels, then the presence of pedestals in other channels could be safely ignored.

In order to accommodate the simultaneous presence of the test and the mask without exceeding the range of the lookup table, it was necessary to reduce the RMS contrast of the test from 25 to 10%. The mask was set to 20% RMS contrast. The masks were spatially filtered at mirror orientations of either $\pm 0^\circ$ (parallel to the axis) or $\pm 90^\circ$ (perpendicular to the axis). The orientation of the test was selected from one of five orientation pairs, namely $\pm 0^\circ$, $\pm 22.5^\circ$, $\pm 45^\circ$, $\pm 67.5^\circ$, or $\pm 90^\circ$. Random and symmetric masks were generated using the same procedure as the test. Examples of stimuli for tests and masks of various orientations are shown in Fig. 4. Random masks are shown in Fig. 4A and symmetric masks are depicted in Fig. 4B. FK and SR acted as observers for this experiment.

Fig. 5A plots sensitivity to symmetry as a function of test orientation for *random* (i.e. non-symmetric) masks of $\pm 0^\circ$ and $\pm 90^\circ$. Results from Experiment 1 are plotted for comparison and correspond to the 'no mask' condition. For both observers, sensitivity is nil in the presence of a random mask of the same orientation as the test. By comparison, masking has virtually no effect if the test and the mask have different orientations. Performance never quite reaches the levels of the previous experiment in which no mask was present, but this is most likely due to the fact that the contrast of the test was lowered by more than half. Also, the fact that performance is at chance when the orientation of the test and the mask coincide is proof that the mask contains enough energy to interfere with the detection of the signal.

Fig. 5B plots sensitivity to symmetry as a function of test orientation for *symmetric* masks of $\pm 0^\circ$ and $\pm 90^\circ$. Results from Experiment 1 are also replotted for comparison. For both observers, results differ substantially from those obtained with the random masks in that performance is considerably less affected by the presence of symmetric masks of any orientation. For test orientations parallel to the axis of symmetry ($\pm 0^\circ$), a symmetric mask of the same orientation appears to have a slight facilitation effect whereas a slight masking effect is noticeable for test and masks with orientations perpendicular to the axis ($\pm 90^\circ$). Overall, however, the presence of a symmetric pedestal has little influence on performance for any combination of test and mask orientations.

Results from the random-mask experiment demonstrate that symmetry detection is mediated by spatial mechanisms that are orientationally tuned since observers can separate signal and noise on the basis of orien-

tation alone. Although very small facilitation/masking effects were noticeable when symmetric masks were of the same orientation as the test, our data reveal that symmetric patterns are not very potent pedestals. Thus, data from the symmetric-mask experiment are more consistent with the notion that symmetry is computed separately in different mirror-orientation channels, although further investigation is needed to confirm this.

5. Experiment 3: the integration region for mirror symmetry

It has been reported that symmetry detection in isotropic narrowband-filtered noise is limited to an integration region (IR) covering only a few cycles of spatial scale near the axis of symmetry (Dakin & Herbert, 1998). The same authors also report that the IR is elongated by a ratio of roughly 2:1 in the direction parallel to the axis. However, since our previous exper-

iment has shown that oriented filters mediate the computation of symmetry, the dimensions of the IR measured by Dakin and Herbert (1998) for isotropic patterns may in fact reflect the composite integration region from several mirror-orientation channels. In the present experiment, we measured the dimensions of the IR as a function of the orientation content of the stimulus.

As in Experiment 1, stimuli consisted of noise patterns with variable amounts of symmetry and variable mirror orientations. However, symmetry was spatially restricted to a window that could be extended either in height (parallel to the axis) or in width (perpendicular to the axis) from the center of the image. The symmetric window was embedded in non-symmetric noise which filled the remainder of the stimulus. A smooth Gaussian spatial transition ($\sigma = 0.15^\circ$ of arc) between the symmetric and non-symmetric windows attenuated subjective contours formed at their boundary. This technique is similar to the one used by Dakin and

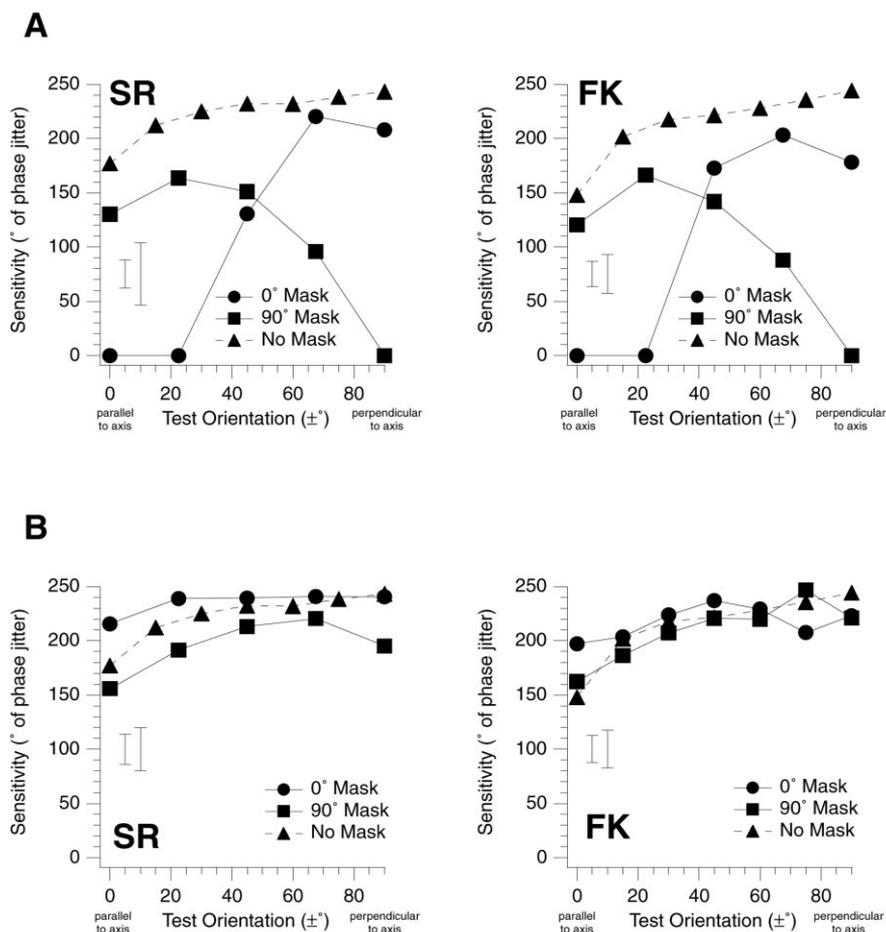


Fig. 5. Results from Experiment 2. (A) Random masks: sensitivity to symmetry in the presence of random masks of various orientations is plotted as a function of test orientation for observers SR and FK. Filled circles correspond to vertically-filtered random masks whereas filled squares indicate horizontally-filtered random masks. Results from the previous experiment are plotted as filled triangles and correspond to the 'no mask' condition. Maximum and mean standard-deviation estimates are also shown. (B) Symmetric masks: same as in A with the exception that masks are perfectly symmetric rather than random.

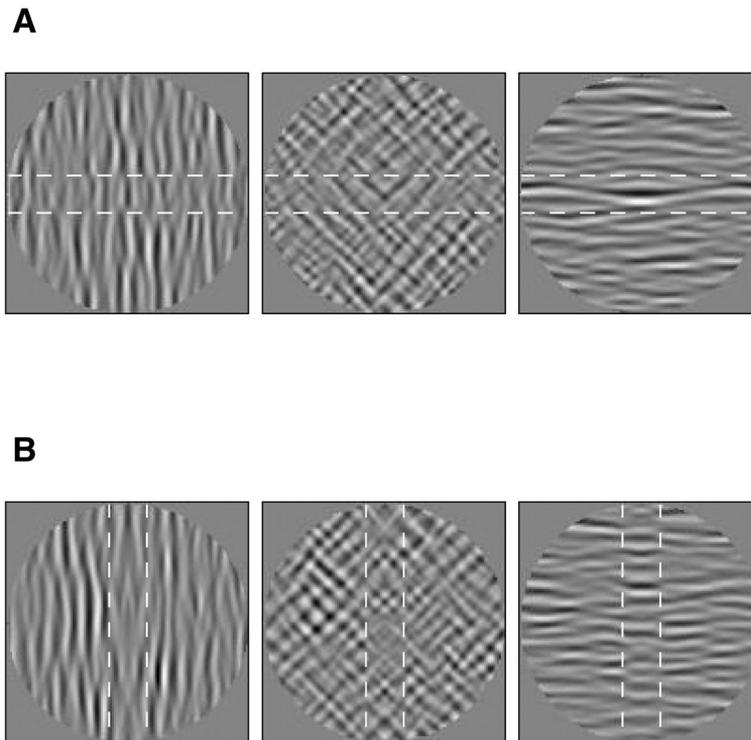


Fig. 6. Stimuli for Experiment 3. (A) Symmetric patterns filtered for orientations parallel ($\pm 0^\circ$), oblique ($\pm 45^\circ$), and perpendicular ($\pm 90^\circ$) to the axis are shown. Mirror symmetry was physically limited to a window of variable height and was embedded in random noise that filled the remainder of the display. A smooth Gaussian transition between the symmetric and non-symmetric portions preserved spatial structure and avoided the formation of subjective contours. Dotted lines are shown here for illustrative purpose but were not presented in actual trials. (B) Same as A except the symmetry window varies in width.

Herbert (1998) to measure the integration region in isotropic noise. As can be judged from Fig. 6, the Gaussian transition limits the spatial extent of mirror symmetry without noticeably affecting the local spatial structure of the texture.

The first author and a naive observer, HW, participated in the experiment. Measures of sensitivity were gathered for symmetry regions of various heights and widths ranging in log-steps from 0.02 to 2.4° of arc. This procedure was repeated for three mirror orientations, namely ± 0 , ± 45 , and $\pm 90^\circ$.

Fig. 7A plots sensitivity as a function of the height of the symmetric window for both observers. Data for windows of variable widths are plotted in Fig. 7B. Solid lines are the best-fitting log-cumulative-Gaussians to the data. Results are consistent between the two observers, although HW was overall less resistant to phase jitter than SR and integrated over a wider spatial region. In line with results reported by Dakin and Herbert (1998), sensitivity improved as the spatial extent of the symmetry window increased, but performance eventually asymptoted as the window exceeded the spatial limits of the IR. In agreement with results from Experiment 1, absolute sensitivity was high for patterns filtered to orientations perpendicular to the axis of symmetry ($\pm 90^\circ$) somewhat lower for oblique

orientations ($\pm 45^\circ$) and poor for orientations parallel to the axis ($\pm 0^\circ$).

The key result of the experiment lies in how the limits of the IR changed as a function of orientation. For convenience, we define the limits of the IR as the spatial extent corresponding to the 95% knee-point in the best-fitting log-cumulative-normal. For both observers, the width of the integration region *increased* by a factor of about three as orientation was varied from parallel to perpendicular. By comparison, the height of the IR *decreased* by about a factor of three as orientation was varied from parallel to perpendicular. This is summarized in Fig. 8A which plots the logarithm of height and width of the IR as a function of orientation. Note also from Fig. 8A that the logarithm of the IRs dimensions vary approximately linearly with orientation, thereby indicating the dimensions of the IR are an exponential function of orientation. Fig. 8B shows example stimuli viewed through a Gaussian window whose dimensions are the logarithmic mean of the IR dimensions measured for SR and HW under the corresponding orientation condition. Although Fig. 8B shows how the dimensions of the IR change rather dramatically with orientation, it is for illustrative purposes only; our measurements do not allow us to determine whether a Gaussian window is an appropriate model for the IR.

Table 1 shows the dimensions of the integration region across mirror orientations for both observers as well as their logarithmic means. The product of height and width as a function of orientation for both observers reveals that area is approximately constant across orientations, although the area is slightly larger for orientations perpendicular to the axis of symmetry. It is also worth noting that the IR is always elongated in the direction parallel to the axis, but that its aspect ratio is of about 20:1 for parallel orientations, roughly 3:1 for oblique orientations ($\pm 45^\circ$), and approximately 2:1 for perpendicular orientations.

6. Summary of results

The main empirical findings of this study are the following:

- Sensitivity to symmetry is constant over nearly the entire spectrum of mirror orientations except for

mirror orientations parallel to the axis of symmetry where sensitivity drops.

- Symmetry detection in the presence of a random mask is only possible if the mask differs in orientation from the test.
- Symmetry detection is largely unaffected by the presence of a symmetric mask for any combination of test and mask orientation. However, slight facilitation/masking effects are noticeable for tests and masks of the same orientation.
- The dimensions of the IR vary considerably as a function of mirror orientation. The width of the IR increases by a factor of about three and the height decreases by a factor of about three as mirror orientations change from parallel to perpendicular to the axis of symmetry.
- The height and width of the IR vary as exponential functions of orientation.
- The area covered by the IR remains approximately constant with orientation, although the area tends to increase slightly as orientation changes from parallel to perpendicular to the axis.

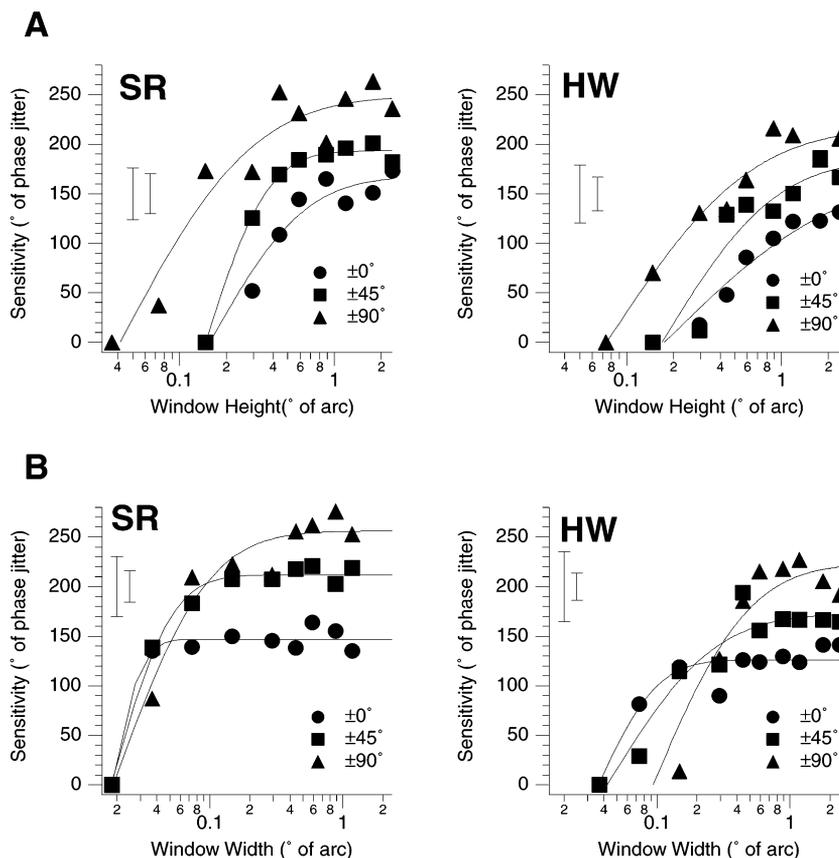


Fig. 7. Results of Experiment 3. (A) Sensitivity to symmetry is plotted as a function of the height (in $^\circ$ of arc) of the symmetry window for observers SR and HW. Filled circles correspond to patterns with orientations parallel to the axis ($\pm 0^\circ$), filled squares denote oblique orientations ($\pm 45^\circ$) and filled triangles indicate orientations perpendicular to the axis ($\pm 90^\circ$). Solid curves are the best-fitting log-cumulative-Gaussian to the data. Maximum and mean standard-deviation estimates are shown in each panel. (B) Same as A but for symmetry windows of variable widths.

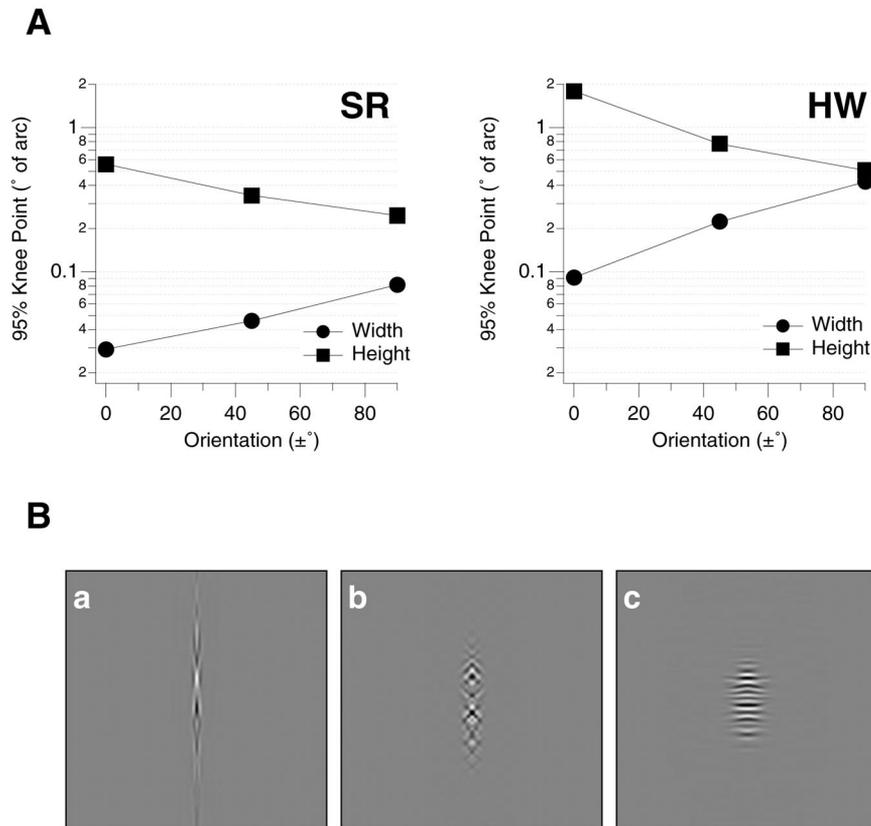


Fig. 8. Summary of results from Experiment 3. (A) The 95% knee-points in the height and width of the integration region are plotted as a function of texture orientation for observers SR and HW. Filled circles correspond to critical width whereas filled squares indicate critical height. (B) Integration regions computed from logarithmic means between observers SR and HW are drawn for the three orientations tested.

Table 1
Dimensions and area of the integration region^a

Orientation (°)	Height (° of arc)			Width (° of arc)			Area (° of arc ²)		
	SR	HW	Mean	SR	HW	Mean	SR	HW	Mean
±0	0.56	1.79	1.00	0.03	0.09	0.05	0.02	0.16	0.05
±45	0.34	0.77	0.51	0.05	0.22	0.10	0.02	0.17	0.05
±90	0.25	0.51	0.35	0.08	0.42	0.19	0.02	0.21	0.07

^a Estimates of the visual angle corresponding to the 95% height and width knee-points in the integration region for observers SR and HW are shown for three mirror orientations. The logarithmic mean between SR and HW is also shown. Note that the dimensions of the integration region change considerably as a function of orientation but that total area (defined by the product of height and width) remains approximately constant across orientations.

7. Discussion of results

In this section, we focus on the implications of our empirical findings on the architecture of symmetry perception. These implications are derived quite independently from Section 8 in which we present an ideal-observer analysis and a multi-channel model of symmetry detection.

7.1. Combining oblique orientations

Results from our first experiment are in line with the findings of Dakin and Hess (1997), namely that symmetry detection in noise patterns is more resistant to phase jitter if noise patterns are filtered for perpendicular orientations rather than for parallel orientations. However, our experiment produced the novel finding that

performance with oblique orientations remains high and does not collapse to chance as models which rely only on orientations parallel and perpendicular to the axis would predict — e.g. the quasi-linear model in Dakin and Hess (1997). Our results therefore imply that a fundamental property of symmetry perception in human vision involves mechanisms that integrate across mirror orientations that can differ by as much as 90°, as with patterns whose mirror orientations are at angles of -45 and 45° with respect to the axis of symmetry.

Although absolute sensitivity varies with orientation, we have shown that observers can compute mirror symmetry from structure that can take on any orientation with respect to the axis. From an ecological standpoint, this is a sound strategy since natural scenes, unlike our artificial stimuli, offer no guarantee that symmetry is defined well by all orientations. Including oblique orientations in the computation of symmetry can not only increase signal-to-noise ratio but can also reduce the risk of missing critical symmetric signals defined predominantly by orientations oblique to the axis.

7.2. Oriented versus non-oriented filters in symmetry perception

Although our first experiment showed that oblique orientations are combined in the computation of symmetry, it did not reveal whether mechanisms mediating symmetry perception are orientationally tuned. This is because the combination of orientations can be achieved either through non-oriented filters or through pooling across oriented mechanisms. However, results from our second experiment clearly favor the interpretation that oriented mechanisms mediate symmetry perception since observers were capable of separating the signal from the noise on the basis of orientation alone. The finding that a random mask perpendicular to the axis had little effect on symmetry detection for orientation parallel to the axis also rules out the possibility that symmetry detection is mediated solely by perpendicular filters with broad orientation bandwidths. The demonstration that oriented spatial filters are involved also argues against an early sub-cortical locus for symmetry mechanisms since pathways such as LGN afferents consist mainly of concentric non-oriented receptive fields (Maffei & Fiorentini, 1972). Rather, results reveal that oriented mechanisms mediate symmetry perception, and that oriented mechanisms covering the entire orientation spectrum are involved. This point towards a cortical locus such as primary visual cortex or beyond where an abundance of neurons with oriented receptive fields are reported (Hubel & Wiesel, 1968). Evidence that symmetry detection is cortical is also reported by Joung et al. (2000) and van der Zwan et al. (1998), although their evidence concerns the neural representation of the axis of mirror symmetry rather than the filtering mechanisms

that precede the computation of symmetry.

A possible objection to our conclusions could be that masks whose mirror orientations differed from those of the test had no effect because observers used attention to select only those orientations containing the signal. While observers may have used this strategy, it does not detract from the fact that the stimulus must necessarily be decomposed into different orientations bands in order for attention to successfully isolate key orientations from others. The mere fact that we can detect symmetry while ignoring the potentially deleterious effects of other orientation bands is good evidence that oriented filters are involved prior to the computation of mirror symmetry.

7.3. Combining symmetry channels

From our results obtained with symmetric masks, we noted that slight facilitation effects are discernable for tests and masks that are both parallel to the axis, but are absent for tests and masks that are orthogonal. Although the effect is admittedly small, it is consistent with the notion that symmetry is computed independently in separate mirror-orientation channels. Little is known about symmetry channels, and even less is known about how they are ultimately combined, but it appears likely that channel combination depends on stimulus properties. For instance, Rainville and Kingdom (1999b) showed that symmetry detection in broadband noise patterns is consistent with probability summation (Graham & Robson, 1987; Graham, 1989) across symmetry channels tuned to various spatial scales. In the experiments of Rainville and Kingdom (1999b), pooling across scales led to better performance because the stimuli contained useful information at several spatial scales. In the symmetric-mask experiment of the present paper, however, the signal was contained only in one or two orientation bands while other orientations (i.e. masks/pedestals) were strictly uninformative. In this particular case, blindly pooling across symmetry channels could in fact decrease performance since symmetric masks could serve as pedestals that would reduce the visual system's ability to discriminate small changes in symmetry. A more optimal solution could be to pool across channels deemed to be informative and simply ignore channels contributing only noise. As evidenced by results of Experiment 2, the fact that observers are impervious to masks/pedestals of orientations that differ from the signal may be indicative of a more adaptive pooling strategy that, contrary to strict probability summation, allows one to reject uninformative channels in order maximize signal-to-noise ratio. In short, our conclusion that symmetry is computed in separate mirror-orientation channels is not incompatible with the requirement of a unified sensory experience, but it highlights that rules governing channel combination are likely to be flexible and task dependent.

7.4. The spatial region of integration for mirror symmetry

There are several implications to the findings from the experiment in which we measured the spatial IR for mirror symmetry. First, results indicate that orientation channels remain separate at least up to the stage where the dimensions of the IR are determined. This necessarily follows from the fact that the shape of the IR would be nearly invariant across orientations if orientations were combined before the IR was defined. These results are also consistent with the notion that orientation channels remain separate beyond the stage where symmetry is computed since recombining across orientations would be impractical at best given that the IRs for different orientations are quite mismatched in shape.

Second, the fact that the dimensions of the IR critically depend on stimulus orientation suggests that the computation of symmetry is a fairly low-level process rather than a higher-level representation operating on complex objects instead of local cues such as orientation. This argument is similar to the one put forward by Dakin and Herbert (1998) and Rainville and Kingdom (1999b) with respect to change in IR dimensions as a function of spatial scale. As we note in Section 9, however, this does not rule out that high-level symmetry-detection mechanisms may be activated when presented with more complex stimuli.

Third, the observation that the IR is always elongated in the direction parallel to the axis is important because it provides further evidence that the dimensions of the IR are not simply a by-product of the aspect ratio of underlying spatial filters. If filter aspect ratio solely determined the dimensions of the IR, one would expect that width should exceed height in the $\pm 90^\circ$ condition and that width and height should be nearly equal in the $\pm 45^\circ$ condition. Data clearly show that this is not the case. While the change in width and height measured as a function of orientation could be attributable in part to the spatial aspect-ratio of underlying filters, the elongated shape of the IR argues in favor of *preferential connectivity between filters positioned along the axis of symmetry*. This last point is critical and is addressed further in Section 8.2.

Fourth, it is somewhat surprising that the IR extends only a short distance away from the axis of symmetry. Although this strategy is not optimal for patterns such as stimuli used in this study, it may be more optimal in the context of natural scenes where symmetry tends to be spatially localized within objects and likely degrades as a function of the distance between corresponding symmetric points. Thus, integrating mirror symmetry over a large portion of a natural scene would likely reduce signal-to-noise ratio.

8. Modeling

In the subsections that follow, we go one step further in trying to better understand the factors that limit human performance in the experiments reported in the first part of this paper. As a first step, we compare human data in Experiment 1 against that of the ideal observer in order to determine whether performance should be attributed to stimulus physics or limitations of human vision. This exercise will prove instructive, especially in cases where human performance deviates from ideal-observer predictions. As a second step, we explore how information density changes with orientation and how the IR effectively places a limit on human performance in Experiment 1. As a third and final step, we present a multi-channel model of symmetry detection that accounts for our results and is well suited to physiological implementation. Our model is then compared against existing models of symmetry detection.

8.1. Ideal observer considerations

Stimuli in this study were constructed by filtering symmetric noise patterns for two narrow orientation bands that could overlap to different extents depending on the chosen orientation. To quantify the amount of information for stimuli of different orientations, we rely on a cross-correlation operator that compares the spatial contents from one side of the axis with the mirror reflection of the other side. In noise patterns such as the ones used in the present study, cross-correlation is a reasonable approximation to the ideal observer (Tapiovaara, 1990; Rainville & Kingdom, 1999b). We define the cross-correlator c_{ij} for vertical mirror symmetry as

$$c_{ij} = \frac{\sum_y \sum_{x=1}^{X/2-1} g_{ij}(x, y) \cdot g_{ij}(x, y)}{\sqrt{\sum_y \sum_{x=1}^{X/2-1} g_{ij}(x, y)^2 \cdot \sum_y \sum_{x=1}^{X/2-1} g_{ij}(x, y)^2}} \quad (1)$$

where g is an instance of a symmetric stimulus for the i th orientation level and of the j th phase jitter level. In a MATLAB numerical simulation, we computed cross-correlation on 250 symmetric/non-symmetric stimuli pairs for each cell of a 9×9 matrix containing nine orientations equally spaced between ± 0 and $\pm 90^\circ$ and nine phase-jitter levels equally spaced between 0 and 360° . In total, cross-correlation values were obtained for 40 500 images. d' (d-prime) values were computed independently for each cell of the 9×9 matrix by comparing cross-correlation means and variances for symmetric and non-symmetric stimuli (Green & Swets, 1988). For each orientation, the decrease in d' as a function of phase jitter was independently fitted with a cumulative error function, and the phase-jitter value corresponding to a d' value of 1.76 was taken as the model's tolerance to phase jitter, as opposed to a d'

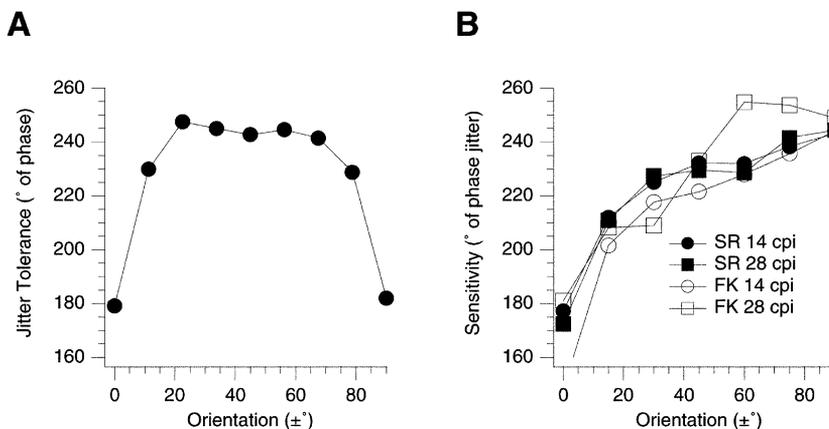


Fig. 9. Cross-correlator and human performance versus orientation. (A) Maximum tolerable phase-jitter for a cross-correlator performing at 95% correct. (B) Human performance for observer SR and FK performing at 75% correct in Experiment 1.

value of 0.95 for human observers. There is nothing special about a d' value of 1.76 apart from the fact that it brought cross-correlator and human data within approximately the same performance range and prevented ceiling effects that are inevitable at extremely high sensitivities. Therefore, it is only meaningful to compare *relative* sensitivities between cross-correlator results and human data, not absolute sensitivities.

Fig. 9A plots the cross-correlator's tolerable phase jitter as a function of stimulus orientation. For comparison, we have also plotted human data from the two observers in Experiment 1. This is shown in Fig. 9B.

The ideal observer fails to predict the pattern of human performance in two essential respects. First, ideal-observer performance is identical for the ± 0 and $\pm 90^\circ$ conditions whereas human sensitivity for $\pm 90^\circ$ is considerably higher than at $\pm 0^\circ$. This shows that the difference in human sensitivity for ± 0 and $\pm 90^\circ$ is a result of neural factors rather than a consequence of stimulus properties. Second, the ideal observer performs better for oblique orientations (e.g. $\pm 45^\circ$) than for perpendicular orientations ($\pm 90^\circ$). This indicates that the stimulus carries more information in the oblique condition than in the perpendicular one. However, the fact that human performance with oblique orientations is either the same or lower than for perpendicular orientations implies that human vision extracts symmetry more efficiently in the perpendicular condition than in the oblique condition.

8.2. Orientation, axis proximity and information density

Fig. 10 illustrates how an image of finite size can be tiled spatially by oriented spatial filters (ellipses). The axis for vertical mirror symmetry is shown by the dashed line. The figure's purpose is to provide the reader with a sense of the spatial distribution of infor-

mation at different orientations when represented by wavelet-like filters of a given scale. Each filter conveys information about a different part of visual space, and as long as filters do not overlap significantly, the amount of information contributed by one filter is independent of its neighbors. Therefore the amount of information per unit area — or *information density* — can be understood in terms of the number of non-over-

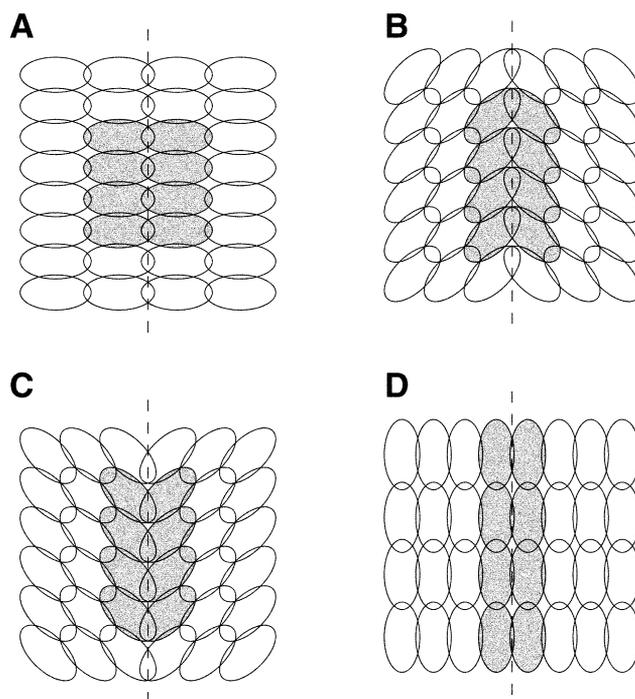


Fig. 10. Spatial distribution of information across orientations. Panels A–D show the spatial distribution of information across an image of finite dimensions for orientation channels tuned to ± 0 , -45 , $+45$, and $\pm 90^\circ$, respectively. Channels are assumed to be composed of collections of self-similar filters shown as ellipses with 2:1 aspect ratio. Shaded filters are equal in number and show how spatial integration must change with orientation to include a fixed amount of information.

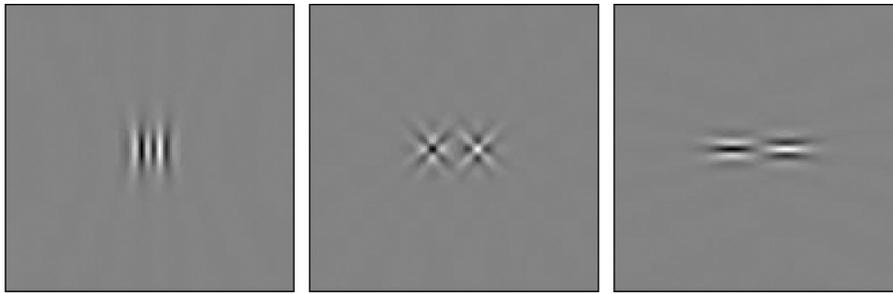


Fig. 11. Basic symmetry-detection units for three different mirror orientations. Units consist of pairs of adjacent spatial filters that have oriented mirror-symmetric profiles linked by mutually inhibitory connections. When units are positioned over symmetric portions of an image, filter outputs cancel each other perfectly whereas some residual activity remains if units are positioned over non-symmetric portions of an image. Note that units detecting oblique mirror symmetry actually consist of the sum of four filters (panels b and c) in order to match the spatial structure of our stimuli.

lapping filters that can be packed within a given region of space.

Consider panels a and d of Fig. 10. The fact that the total number of non-overlapping filters (i.e. overall information density) is the same for both orientations confirms our ideal observer's first finding, namely that the amount of information in the parallel and perpendicular conditions is the same. What does change with orientation, however, is information density along the x and y axes of the image. Information density along the y axis is higher for horizontal filters than for vertical filters since, in this example, there are eight horizontal filters per column as opposed to four for vertical filters. The opposite is true along the x axis where for any given row, vertical filters outnumber horizontal ones.

In Experiment 3, we have shown that symmetry detection is restricted to an area that extend only a short distance away from the axis. If we impose the same constraint of axis proximity onto filters of Fig. 10, filters perpendicular to the axis, by virtue of their high density along the y dimension, have an intrinsic advantage over filters parallel to the axis. In order to integrate the same amount of information (filters shaded in gray), the height of the IR must change with orientation. Otherwise, should the height of the IR remain fixed, filters parallel to the axis would be at a severe disadvantage with respect to filters of other orientations. From this perspective, it is remarkable that the height of the IR measured for human observers *does* in fact stretch significantly with orientation, as though human vision attempts to compensate for the lack of near-axis information in the parallel condition. *However, as we show in the next section, stretching the height of the IR by the same amount measured for human observers is insufficient to fully equate performance between the parallel and perpendicular conditions.*

Inspection of Fig. 10 also confirms that if one takes the entire stimulus into account, oblique conditions are physically more informative than either the perpendicular or parallel conditions. The reason is simply that our

oblique stimuli have two orientation components on each side of the axis (i.e. panels B and C combined) whereas the parallel (panel A) and perpendicular conditions (panel D) have only one (recall from Section 1 that parallel and perpendicular conditions are special cases). However, if only filters adjacent to the axis of symmetry are considered, then oblique orientations have a lower information density along the axis than orientations perpendicular to the axis. Thus, the numerosity advantage that oblique filters have over perpendicular ones is offset by the fact that the latter have a higher information density along the axis. This is admittedly a subtle but nonetheless important point since, as we show in the next section, it accounts for our findings that human performance in the oblique condition is in fact no better than in the perpendicular condition.

8.3. A multi-orientation model for symmetry coding

The purpose of the present section is to demonstrate how the pattern of results obtained in Experiments 1–3 are consistent with a model that combines the output of oriented spatial filters to encode mirror symmetry in a simple and physiologically plausible way. To our knowledge, this is the first model of symmetry detection that explicitly includes oriented filters that together cover the entire orientation spectrum.

Our model is built on symmetry-detection units (SDUs) that, via oriented spatial filters, are tuned to one of several possible mirror orientations. Examples of SDUs are shown in Fig. 11 for three different mirror orientations (± 0 , ± 45 and 90° , respectively). Each SDU has a minimum of two oriented filters that are spatially adjacent (like the shaded filters shown in Fig. 10) and of opposite polarities. Spatial filters are lined-up along the dimension perpendicular to the axis of symmetry. For simplicity, we have also assumed that filters are matched to the corresponding orientation of our stimuli; thus filters in the $\pm 45^\circ$ condition are tuned

to *two* orientations bands whereas filters in the ± 0 and $\pm 90^\circ$ conditions are tuned only to one orientation band.

The principle by which SDUs encode mirror symmetry is quite simple. In cases where the axis of symmetry falls halfway between filters of opposite polarity, filter responses are of equal magnitudes but of opposite signs. Thus, in this situation, filter outputs cancel each other out. On the other hand, if the same SDU is centered on non-symmetric portions of an image, filters responses are different and their sum leads to a non-zero SDU output that signals the absence of mirror symmetry. In short, filtering an image with a SDU produces near-zero responses where local mirror-symmetry exists, and non-zero responses everywhere else.

Applying a pointwise rectifying nonlinearity to SDU outputs discards the sign. Gurnsey, Herbert and Kenemy (1998) have recently proposed a similar symmetry-detection scheme, although their model filters neither for spatial scale nor orientation.

Fig. 12 shows schematically how an arbitrary stimulus with vertical mirror symmetry (A) is filtered with SDUs of three different mirror-orientations (B) whose outputs are full-wave rectified to produce two-dimensional symmetry maps (C). These symmetry maps can be collapsed onto a one-dimensional profile (D) by summing along the dimension parallel to the axis over any desired height. Note how, for all mirror orientations, the presence and the location of the axis of symmetry in the original image (A) are respectively

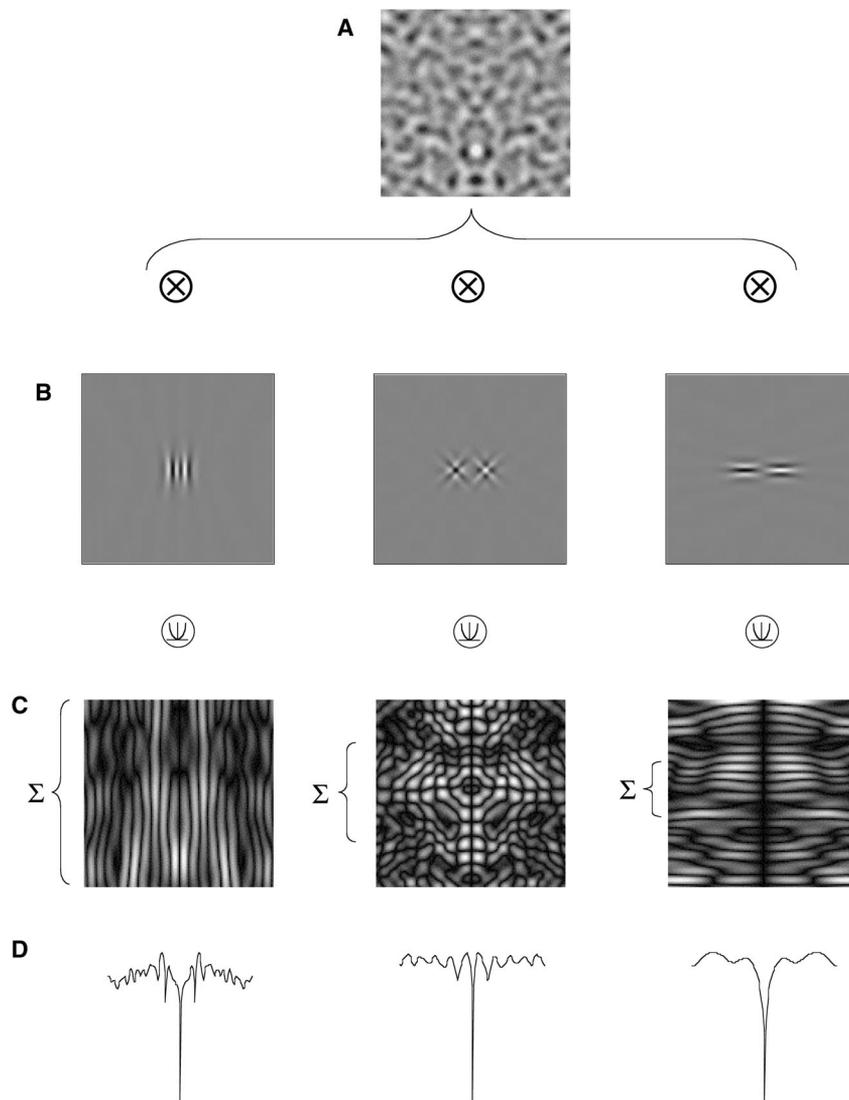


Fig. 12. Model of symmetry-detection with multiple orientation channels. A given stimulus (A) is filtered by symmetry-detection units (B) whose combined output is passed through a squaring nonlinearity. Two-dimensional maps of potential regions of symmetry (dark) and non-symmetry (light) are produced (C). Summing in the direction parallel to the axis of symmetry over variable heights results in one-dimensional profiles whose dip encode both the magnitude and location of the axis of symmetry.

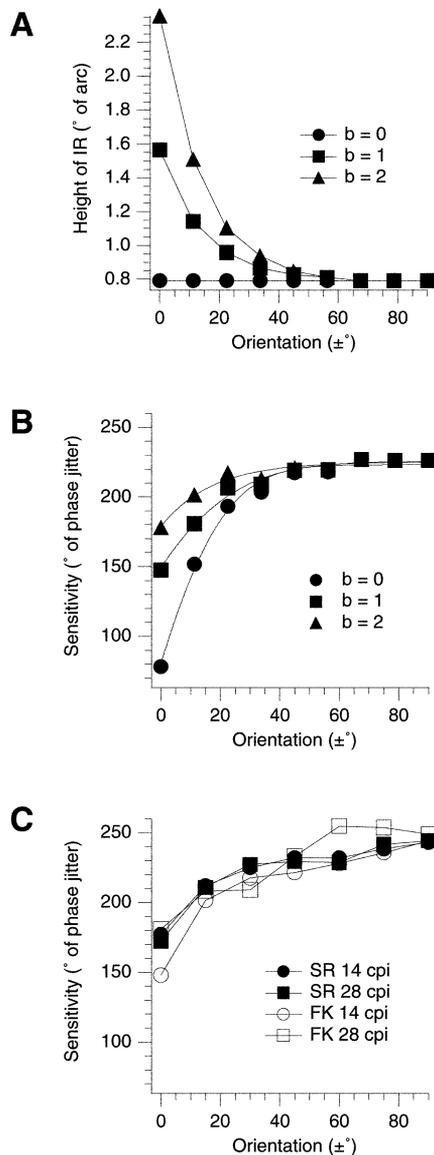


Fig. 13. Model simulations. (A) Plots of the height of the integration region for three ‘stretching’ constants. (Filled circles) Height remains constant with orientation ($b = 0$). (Filled squares) Height decreases exponentially by a factor of two ($b = 1$). (Filled triangles) Height decreases exponentially by a factor of three ($b = 2$). (B) Model performance for three ‘stretching’ constants. Symbols correspond to the ‘stretching’ condition in panel A. (C) All data from all human observers in Experiment 1 are replotted in order to compare model performance.

encoded by the magnitude and location of through activity in (D). Further technical details on the specifics of the model are given in Appendix B.

Data from Experiment 3 revealed that the height of the IR varies *exponentially* with stimulus orientation and that height changes by roughly a factor of three. With the help of our model, we can explore how changing the height of the IR affects performance. In three separate MATLAB simulations, we subjected our model to the same stimulus set used in Experiment 1

and in our ideal observer analysis. In all simulations, the height of the model’s IR decayed exponentially with stimulus orientation. Across simulations, however, we varied the spatial extent over which height was allowed to stretch. This is illustrated in Fig. 13A where the height of the model’s IR is plotted as a function of orientation for three different ‘stretching’ constants. In one simulation, no stretching was allowed, and so the height of the models’ IR remained constant with orientation. In the other two simulations, height in the parallel condition was allowed to stretch by factors of two and three, respectively. The height of the IR, denoted by Y_i , is given by

$$Y_i = a\lambda(1 + b \exp[-c\theta_i]) \quad (2)$$

where λ is the spatial scale and b is a ‘stretching’ constant that determines the amplitude by which the height of the IR changes with orientation. In our three simulation, b took values of 0, 1 and 2, respectively. The multiplicative constant a determines the number of spatial cycles included in the ‘non-stretched’ IR and was empirically set to 5 in line with results of Experiment 3. The constant c determines the rate at which the IR’s height changes with orientation and was empirically set to 4.

For each simulation, our model made symmetry judgements on 250 symmetric/nonsymmetric stimuli pairs for each cell of a 9×9 matrix containing nine orientations equally spaced between 0 and 90° and nine phase-jitter levels equally spaced between 0 and 360° . In total, symmetry judgements were obtained for 121 500 images. d' values were computed independently for each cell of the 9×9 matrix by comparing means and variances of the model’s response to symmetric and non-symmetric stimuli (Green & Swets, 1988). For each orientation, the decrease in d' as a function of phase jitter was independently fitted with a cumulative error function. The phase-jitter value corresponding to the same discriminability achieved by human observers (75% correct, or $d' = 0.95$) was taken as our model’s tolerance to phase jitter.

Fig. 13B plots the model’s tolerable phase jitter as a function of stimulus orientation for each of the three allowed stretching extents. For comparison, Fig. 13C plots all available data from human observers in Experiment 1. Note that for all stretching constants, the model shows the deficit for parallel orientations that is characteristic of human performance. However, if the IR is allowed to increase its height as stimulus orientations is made more parallel to the axis, the model’s performance deficit is considerably reduced. In the case where stretching is the same as that measured psychophysically in Experiment 3 (approximately 3:1), the model provides good qualitative agreement with human data. It should also be noted that, unlike the ideal observer (see Section 8.1), our model’s performance is

not shaped as an inverse-U but rather decays monotonically as the orientation content is varied from perpendicular to parallel to the axis.

We should note immediately that agreement between our model and our data does not constitute proof that symmetry detection in the human visual system is implemented in the manner we have just described. In fact, there is an infinity of possible implementations which would yield similar results. However, the key components of our model, namely the change in the shape of the IR with orientation, is inspired by psychophysical data and it is therefore likely that other model implementations which take this into consideration would behave similarly. In fact, a cross-correlation operator would perform in much the same way as our model if it viewed the stimulus through an aperture that mimicked the stretching of the IR, but its physiological plausibility would be questionable.

In Section 8.2 and in the model simulations reported here, we have shown that our results with human observers are predicted well by a model that recruits oriented spatial filters near the axis of symmetry. In addition, our simulations reveal that the performance deficit for orientations parallel to the axis is considerably larger if the height of the IR is not allowed to stretch to compensate for lower information density. In the next section, we briefly compare our model to existing computational models of mirror symmetry and focus particularly on the issue of orientation tuning.

8.4. Other models of mirror-symmetry detection

While several computational models are successful in detecting and localizing mirror symmetry in an image, most do so without relying on mechanisms with spatial properties similar to those thought to mediate early stages of human vision. As Dakin and Hess (1997) succinctly put it, 'In terms of Marr's (1982) levels of representation, there has been an emphasis on algorithms for symmetry detection in the absence of implementational considerations'. One such class of models detects symmetry by computing a point-by-point cross-correlation between the two symmetric halves of the image (e.g. Barlow & Reeves, 1979; Pintsov, 1989; Gurnsey et al., 1998). Although this approach has the virtue of simplicity, computations operate at the level of individual pixels (or averaged regions) and therefore do not reflect known neural transformations such as bandpass filtering for spatial scale or orientation. Because of this, such models make few predictions with respect to orientation except perhaps from their implicit assumption of isotropy. Consequently, these models fail to predict our results.

A second class of models lies at the opposite end of the spectrum from the first and requires complex representations and/or grouping rules from which mirror symmetry can then be extracted (e.g. Pashler, 1990; Wagemans, 1993; Pani, 1994; Zabrodsky & Algom, 1994; Labonte et al., 1995). In particular, the model of Labonte et al. (1995) implements a form of orientation filtering which segregates clusters of compatible line segments into different orientation groups using a relaxation-labeling procedure. However, this algorithm assumes orientation labeling is a computational non-problem and only succeeds if the orientation of line segments is known a priori. Also, because our stimuli consist only of filtered noise, they lack most of the features that high-level spatial grouping or segregation strategies require to succeed in detecting mirror symmetry. Although models of this sort may capture some higher-level aspects of human vision, it remains to be shown how they could be implemented using simple visual mechanisms such as oriented spatial filters.

A third class of symmetry-detection models includes those which relate the computation of mirror symmetry to early spatial mechanisms (such as filters) that are psychophysically, physiologically, and theoretically motivated. One model of symmetry detection proposed by Osorio (1996) capitalizes on the fact that convolving an image with a pair of quadrature Gabor filters oriented parallel to the axis of symmetry generates energy peaks that correspond to the presence of symmetry. Such peaks arise because spatial harmonics *parallel to the axis* fall into local cosine phase alignment ($90/270^\circ$) at the axis. According to the model, the position and the magnitude of energy peaks can be used to determine the location of the axis as well as the amount of local symmetry. However, because such a scheme is only defined for orientations parallel to the axis, the model does not generalize to stimuli consisting only of oblique or perpendicular orientations. Given that human performance is poorer with structure parallel to the axis than with structure of any other orientation — see Experiment 1 and Dakin and Hess (1997) — it is unlikely that symmetry is mediated by mechanisms concerned with peaks in quadrature energy.

Dakin and colleagues propose several versions of a model which computes symmetry using the co-alignment of blob clusters obtained through half-wave rectification and thresholding of the output of spatial filters. In its simplest form, the model operates on filter outputs that are either isotropic or oriented perpendicular to the axis of symmetry (Dakin & Watt, 1994). However, the joint findings of Experiment 1 and 2 of this study are incompatible with both the isotropic and perpendicular versions of the model since our data demonstrate that symmetry mechanisms combine across oblique mirror orientations but do so while

relying on *oriented* spatial filters. In a more elaborate version (Dakin & Hess, 1997), the model filters the image *only* at orientations parallel and perpendicular to the axis of symmetry. For orientations parallel to the axis, the model limits the region of integration by a Gaussian window of variable width. According to the authors' own account, this model does not provide an explanation for the observed deficit in performance with orientations parallel to the axis. In addition, the model's lack of sensitivity to oblique orientations is not supported by our data from the present paper. Finally, the authors propose a nonlinear model in which the image passes through an initial isotropic filtering stage, undergoes half-wave rectification, and is refiltered for orientations perpendicular to the axis of symmetry. While the initial isotropic stage renders the model sensitive to oblique mirror orientations, the half-wave rectification stage makes the model susceptible to masking from oblique orientations since, in those conditions, non-Fourier energy is introduced at parallel and perpendicular orientations. In addition, because this nonlinear model consists of a single filter-rectify-filter channel, it does not allow the IR to change its shape as a function of texture orientation. Finally, despite a sandwiched nonlinearity, an isotropic filter followed by a single oriented filter does not allow symmetry to be encoded separately in different orientation channels. These predictions are also not supported by our data reported in this paper.

We should also note that in a previous study (Rainville & Kingdom, 1999b), we have proposed a multi-scale model of symmetry detection inspired from early spatial vision. While our previous model provides a good account of symmetry detection in broadband noise for variable distributions of contrast energy across scales (i.e. power spectra), it relies exclusively on isotropic filters and therefore cannot account for the computation of symmetry in separate orientation channels. It is straightforward, however, to include the multiple-orientation channels to the multi-scale model we have proposed in Rainville and Kingdom (1999b), and we anticipate to do so in a future study.

In summary, models of symmetry detection vary in the degree to which they incorporate properties of early spatial vision, but none provides a full account of the psychophysical results presented in this paper. In particular, most models of symmetry detection fail to predict our data since they rely predominantly on a single channel (either isotropic or oriented) rather than on multiple orientation channels. In this modeling section, we have proposed a new class of symmetry-detection models that include multiple orientation channels. Our model predicts many empirical findings on symmetry perception (including ours) and also retains compatibility with many of the main principles governing early spatial vision.

9. General discussion

9.1. Attention and uncertainty with respect to axis location

In their study on symmetry detection in orientation-filtered noise, Dakin and Hess (1997) reported that none of their models account for the comparatively poor performance of human observers under conditions where stimuli are filtered for structure parallel to the axis. To explain these data, the authors propose that the detection of mirror symmetry in such patterns involves a strong attentional component and tested their hypothesis by randomizing the position of the axis within a region of finite width. When observers did not know the position of the axis, performance was greatly reduced for structure parallel to the axis but were not affected as much for patterns with structure perpendicular to the axis. The authors interpreted this as evidence for a differential role of attention across stimulus orientation. Our model on the other hand not only successfully accounts for the deficit in performance for orientation parallel to the axis but also accounts for human performance for stimuli with structure oblique to the axis. Although our model does not include any attentional component, it correctly predicts that detecting mirror symmetry in patterns with parallel structure to the axis is more difficult than with perpendicular structure when the position of the axis is not known. The key in explaining performance deficits for orientations parallel to the axis, we argue, lies in how the visual system handles the lower information density that characterizes that particular stimulus. Specifically, we have suggested that the visual system fails to sufficiently stretch the IR vertically to compensate for the intrinsically lower information density for vertically oriented symmetry information lying close to the axis. In short, attentional factors are not needed to account for the performance deficit measured for orientations perpendicular to the axis.

9.2. Axis orientation

Given that stimuli used in the present study were always vertically symmetric, it is reasonable to ask whether our results are expected to generalize to stimuli in which the axis of symmetry had some other orientation. This concern is particularly important since several studies report a marked advantage for vertical symmetry over horizontal or oblique mirror symmetry (Palmer & Hemenway, 1978; Barlow & Reeves, 1979; Wenderoth, 1994), although other studies present evidence to the contrary (Corballis & Roldan, 1975; Fisher & Bornstein, 1982; Jenkins, 1983b; Locher & Wagemans, 1993). However, in the study most closely related to ours, Dakin and Hess (1997) measured symmetry

detection in noise patterns filtered for structure either parallel or perpendicular to the axis. The authors also reported that the deficit for structure parallel to the axis is present irrespective of whether the axis of symmetry is vertical or horizontal with respect to gravity. Although this provides some indication that our results would generalize to other axis orientations, the issue remains open for further empirical investigation.

9.3. One or many integration regions for mirror symmetry?

Our results and those of other studies (see Section 1) reveal that the dimensions of the IR are highly dependent on the spatial structure of the stimulus. In the present paper, we have shown that the IR is also highly dependent on the stimulus' orientation content. But how is this change in the shape of the IR implemented in human vision? One possibility is that the dimensions of the IR are hard-wired for channels of given orientation and spatial scale. However, a more intriguing possibility is that the IR is adaptive and acts to regulate the flow of spatial information sent to subsequent stages involved in symmetry detection. The notion of an IR that adapts its shape and size to match stimulus properties remains speculative at this point, but if the changes in the shape of the IR is the result of an adaptive dynamic process, one should be able to measure its time course.

9.4. Specialized mechanisms for symmetry detection?

One potential criticism of our model reported in Section 8 is that it implements neural connections that are specific only to the detection of mirror symmetry. In fact, such an argument could be made for most computational models designed for specific purposes such as, for example, discriminating between spatial frequencies (Wilson & Gelb, 1983). In the case of mirror symmetry, however, some argument can be made for specialized mechanisms dedicated to its detection. Symmetry has been shown to be of considerable importance in the ecology of several species (Swaddle & Cuthill, 1994; Møller, 1995; Horridge, 1996). The fact that several of those species have considerably simpler visual systems than ours suggests that neural connections which implement the detection of symmetry are a fairly low-level and hard-wired process rather than one concerned with the explicit representation of primitives, surfaces and objects (Marr, 1982).

Another argument for specialization is that mirror symmetry cannot be detected unless the axis is foveated (Gurnsey et al., 1998) — although see Gurnsey and Sally (submitted). If symmetry were detected by general-purpose mechanisms, why are human observers incapable of detecting symmetry when the axis falls a

few arc minutes away from fixation? An explanation for the detection of vertical mirror symmetry should involve specific neural connections implemented near fixation. Such specific neural connections between the left and right visual hemifields are found in the brain's corpus callosum, and empirical evidence for the 'callosal' hypothesis for symmetry perception has been provided by Herbert and Humphrey (1996) among others. However, the callosal hypothesis cannot account for the fact that *horizontal* symmetry is easily detected at fixation (Dakin & Hess, 1997). Another possibility is that symmetry detection depends on the foveal specialization of linking mechanisms such as those that mediate the detection of co-aligned features (Hess & Dakin, 1997). The computation of feature co-alignment is an explicit property of symmetry models by Dakin and colleagues (Dakin & Watt, 1994; Dakin & Hess, 1997) and is also implicit in our model presented in Section 8.3. Evidence that contour detection weakly interacts with symmetry detection is provided by van der Zwan et al. (1998), but more research is needed to determine how much the neural representation of mirror symmetry has in common with that of other visual features.

9.5. Information density and the IR

In the present study, we have argued that changes in the shape of the IR at least partially compensates for changes in *information density* as stimuli varied in their orientation content (see Section 8.2). In a previous study (Rainville & Kingdom, 1999b), we have applied a similar argument to the findings of Dakin and Herbert (1998) concerning spatial scale, namely that the size of the IR compensates for lower information density intrinsic to low-frequency stimuli. From these studies, however, it is difficult to determine directly whether information density is the primary factor that determines the size and shape of the IR since information density covaried either with orientation or spatial scale. Fig. 14 shows how one may test for the effects of density without altering either the spatial-frequency or orientation content of the stimulus.

The three textures of Fig. 14 consist of a variable number of vertical Gabor microelements that were linearly superimposed. For all three textures, Gabor microelements are positioned symmetrically about the vertical axis *except* those whose center falls within the central 1/8th of image width. Because of the bandpass nature of the microelements, the three textures are of the same spatial scale and orientation content. Consequently, any change in the IR must be attributed to some stimulus property other than spatial scale such as the spatial density of information. In the limit, if the texture contains many elements (panel a), its information density approaches that of filtered noise. Under

such conditions, as we have verified empirically in the present study, the perception of mirror symmetry is impossible without scrutiny because the IR does not extend from the axis beyond a few cycles of spatial scale. However, if information density is reduced (panel b and c), the perception of mirror symmetry becomes possible, in agreement with our prediction that the IR ‘stretches’ to compensate for the lower information density.

Fig. 14 also highlights that simple filtered noise patterns may not be sufficient to fully reveal the spatial properties of mechanisms underlying the detection of mirror symmetry. It is possible that under appropriate (possibly more naturalistic) stimulus conditions, symmetry detection is not limited to the narrow IR we have measured in Experiment 3. If such is the case, then we must ask what are the spatial properties of the IR in the low-density condition. Tyler and Hardage (1996) report that symmetry perception in low-density symmetric patterns composed of Gaussian blobs can take place over large regions of space and that envelope information is a key component which mediates performance. However, it is still not empirically known how the IR changes shape for sparse textures such as the ones portrayed in Fig. 14C, nor is it known whether results obtained with dense random noise generalize to IRs that are larger than the ones predicted by the nominal spatial scale of the stimulus. In particular, is symmetry still computed independently in separate orientation/spatial-frequency channels or does some form of pooling take place prior to symmetry detection? We are currently in the process of reporting the results from our investigations on these and other issues (Rainville & Kingdom, 1999a, 2000; Rainville & Kingdom, in preparation; Rainville & Kingdom, submitted).

10. Summary and conclusions

In the present study, we have reported psychophysical evidence that links known visual mechanisms —

namely oriented spatial filters — to the computation of mirror symmetry in the human visual system. Our first experiment revealed that symmetry detection is possible irrespective of the orientation content of our stimuli. In particular, the finding that symmetry is well detected even for stimuli with structure oblique to the axis of symmetry implies that corresponding mirror-orientations are combined prior to the stage where symmetry is computed. Our second experiment revealed that masking interferes with the detection of a symmetric test only if the orientation of the mask and the test coincide, thereby suggesting that symmetry is computed largely independently in separate mirror-orientation channels. Our third finding is that the dimensions of the IR for mirror symmetry vary substantially with the orientation content of the stimulus. In subsequent sections, we developed an ideal observer analysis in which we identified neural factors (such as the dimensions of the IR) that limit human performance. We incorporated these factors into a successful multi-channel model of symmetry detection that combines the output from oriented filters in a simple and physiologically plausible manner. With the help of this model, we were able to verify our hypothesis that changes in the dimensions of the IR measured in our third experiment compensate for changes in information density and thus partially equate performance across stimulus orientations.

Acknowledgements

Support for this research has been provided by a PGS B Fellowship from the Natural Science and Engineering Research Council of Canada to SR and by grants from the Medical Research Council of Canada and the Natural Science and Engineering Research Council of Canada to FK. This research was also supported by US Public Health Service grants EY-4885 and EY-1319.

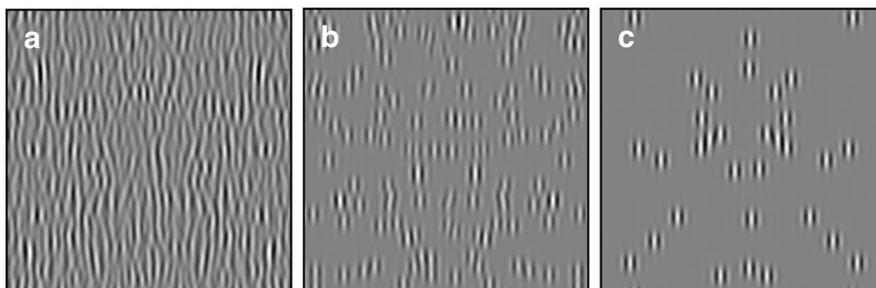


Fig. 14. Symmetric textures of identical scale and orientation but varying density. Textures are composed of the linear superposition of a variable number of vertical Gabor microelements. All microelements are positioned symmetrically except those falling within the central 1/8th of the stimulus' width.

Appendix A. Details of stimulus construction

Each stimulus was constructed in Fourier space which we denote here by G . For every frequency component $G(u, v)$, real and imaginary coefficients were randomly sampled from a Gaussian distribution ($\mu = 0$, $\sigma = 1$), where u and v represent the two dimensions of a Cartesian spatial-frequency coordinate system. Mirror symmetry was introduced in the stimulus by imposing the relationship

$$G(u, v) = G^*(-u, v) \quad (\text{A1})$$

where G^* denotes the complex conjugate of G . Eq. (A1) is strictly equivalent to the relationship $g(x, y) = g(-x, y)$ which defines perfect vertical symmetry in the spatial domain. Building stimuli in the Fourier domain using Eq. (A1) also provides a significant computational advantage in that spatial filtering operations are reduced to a single inverse Fourier transformation. However, Eq. (A1) is also of some theoretical significance for vision (see Section 1) since it illustrates how mirror symmetry can be construed in terms of pairs of oblique sinusoids with mirror orientations, equal amplitudes, and identical phases. From Eq. (A1), it is also possible to see that frequencies with orientations strictly orthogonal to the axis (i.e. $G(0, v)$ for all v) are not relevant since they are, in a sense, always symmetric about the vertical axis. By comparison, frequencies that are strictly parallel to the axis (i.e. $G(u, 0)$ for all u) are important in determining the symmetry of the pattern. Unlike all other frequencies however, they are required to be either in cosine (0°) or anti-cosine (180°) phase. For symmetric random-noise patterns such as the ones in this study, orientations parallel to the axis must be randomly and independently assigned to cosine or anti-cosine phases.

To ensure our stimuli possessed no imaginary components in the spatial domain (like all images that are physically realizable), we imposed the additional constraint of conjugate symmetry. Conjugate symmetry is defined as $G(u, v) = G^*(-u, -v)$ and automatically appears in the Fourier transform of any real signal (Bracewell, 1986). However, when constructing stimuli directly in Fourier space, conjugate symmetry must be specified explicitly.

The orientation content of our stimuli was determined by a filter T of the form

$$T(u, v) = \exp\left[-\frac{1}{2}\left(\frac{\theta - \theta_0}{\sigma_0}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{\theta + \theta_0}{\sigma_0}\right)^2\right] \quad (\text{A2})$$

where θ_0 defines peak orientation and σ_0 defines orientation bandwidth. Transforming Cartesian frequency coordinates into orientation coordinates is easily accomplished by $\theta(u, v) = \tan^{-1}(u/v)$ although care must be taken to avoid division by zero and orientation wrap-around artifacts. The structure of this filter follows

naturally from the definition of mirror symmetry in the Fourier domain. Filter T essentially lets through two narrow bands of spatial frequencies that consist of mirror orientations (e.g. $+30^\circ$ and -30°). If the two orientation bands overlap, their contents are simply added together. In the special cases where mirror orientations are either parallel (0°) or perpendicular (90°) to the axis of symmetry, the filter lets through only a single orientation band as can be seen from panels a and d of Fig. 2.

In the spatial-frequency domain, stimuli were band-pass-filtered with a log-Gaussian filter S given by

$$S(u, v) = \exp\left[-\frac{1}{2}\left(\frac{\ln(f/f_0)}{\ln(f_0\sigma_s)}\right)^2\right] \quad (\text{A3})$$

where f_0 determines the band's center frequency and where σ_s governs its bandwidth. Radial spatial frequency is denoted here by f and is given by $f = \sqrt{u^2 + v^2}$. For all stimuli, σ_s was set to 1.4, which corresponds to a full-width-at-half-height of 1.15 octaves. Restricting spatial scale in this way allows us to recruit only the subset of visual mechanisms tuned to a particular scale and also minimizes the influence of factors that are known to be scale dependent in symmetry perception (Dakin & Herbert, 1998; Rainville & Kingdom, 1999b).

Since the operations of filters T and S are linear they can be combined into a single filtering operator H given by $H(u, v) = T(u, v) \cdot S(u, v)$. Note that H can be characterized as a zero phase-shift filter since it operates on the power spectrum without affecting the phase structure of the stimulus. To vary the degree of symmetry, we randomized the phase structure of our stimuli by adding a phase-angle offset randomly sampled from a uniform distribution of variable width. The uniform distribution was centered on 0° of phase, and the amount of jitter was determined by the distribution's width which varied between 0° of phase and 360° of phase. Once all operations were completed in Fourier space, the stimulus G was reverse-Fourier transformed to the spatial domain g .

Appendix B. Model details

Our model assumes that its filters are matched to the stimulus. From Eqs. (A2) and (A3), we can recover the Fourier-domain expression of filter H for the i th stimulus orientation, as given by

$$H_i(u, v) = \exp\left[-\frac{1}{2}\left(\frac{\ln(f/f_s)}{\ln(f_s\sigma_s)}\right)^2\right] \cdot \left\{ \exp\left[-\frac{1}{2}\left(\frac{\theta - \theta_i}{\sigma_i}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{\theta + \theta_i}{\sigma_i}\right)^2\right] \right\} \quad (\text{B1})$$

The corresponding spatial impulse response h_i can therefore be computed using the reverse Fourier trans-

form $h_i(x, y) = F^{-1}\{H_i(u, v)\}$. The impulse response h_i is spatially localized because H_i is real and all Fourier components are therefore in cosine phase. Symmetry detection units (SDUs) m for the i th orientation are obtained by

$$m_i(x, y) = h_i(x, y) \otimes [\delta(x - \Delta x_i, 0) - \delta(x + \Delta x_i, 0)] \quad (\text{B2})$$

where \otimes denotes the convolution operator and δ represents the Kronecker delta which effectively determines filter positions. Mutual inhibition can be conveniently expressed into a single operator whose filter components have opposite polarities. Also, note that the lateral separation Δx_i between the two filters should increase as stimulus orientation becomes increasingly perpendicular to the axis of symmetry in order to prevent a loss of information due to filter overlap. In our model, filter separation Δx_i is given by

$$\Delta x_i = \lambda/2 + \lambda \sin \theta_i \quad (\text{B3})$$

where λ is the spatial scale of the filters. For the purpose of our model, the scale of the symmetry detection units matched that of the stimuli. Note that the increase of lateral filter separation with a change in orientation from parallel to perpendicular is not only theoretically motivated but also reflects the approximate threefold increase in the width of the IR we have measured psychophysically.

Once the symmetry-detection units have been defined and the height of the IR has been determined, we can compute the response r of symmetry-detection units m of the i th orientation for stimulus g at any spatial location using

$$r_i(x, y) = \sum_{y-Y_i/2}^{y+Y_i/2} [g(x, y) \otimes m_i(x, y)]^2 \quad (\text{B4})$$

where the output of the convolution between the stimulus and the symmetry units is squared and summed along the height of the IR obtained by Eq. (2) in Section 8.3.

References

- Barlow, H. B., & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research*, 19(7), 783–793.
- Bracewell, R. N. (1986). *The fourier transform and its applications* (2nd ed.). Montreal: McGraw-Hill International Editions.
- Brady, N., & Field, D. J. (1995). What's constant in contrast constancy? The effects of scaling on the perceived contrast of band-pass patterns. *Vision Research*, 35(6), 739–756.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 443–446.
- Bruce, V. G., & Morgan, M. J. (1975). Violations of symmetry and repetition in visual patterns. *Perception*, 4(3), 239–249.
- Burr, D. C., Morrone, M. C., & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research*, 29(4), 419–431.

- Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, 197, 551–566.
- Carlin, P. (1996). *On symmetry in visual perception*. Unpublished Ph.D dissertation, University of Stirling, Stirling.
- Corballis, M. C., & Roldan, C. E. (1975). Detection of symmetry as a function of angular orientation. *Journal of Experimental Psychology: Human Perception & Performance*, 104(1), 221–230.
- Dakin, S. C., & Herbert, A. M. (1998). The spatial region of integration for visual symmetry detection. *Proceedings of the Royal Society of London — Series B: Biological Sciences*, 265(1397), 659–664.
- Dakin, S. C., & Hess, R. F. (1997). The spatial mechanisms mediating symmetry perception. *Vision Research*, 37(20), 2915–2930.
- Dakin, S. C., & Watt, R. J. (1994). Detection of bilateral symmetry using spatial filters. Special issue: the perception of symmetry: part 1. Theoretical aspects. *Spatial Vision*, 8(4), 393–413.
- DeValois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22, 545–559.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.
- Field, D. J., & Brady, N. (1997). Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. *Vision Research*, 37(23), 3367–3383.
- Fisher, C. B., & Bornstein, M. H. (1982). Identification of symmetry: effects of stimulus orientation and head position. *Perception & Psychophysics*, 32(5), 443–448.
- Graham, N., & Robson, J. G. (1987). Summation of very close spatial frequencies: The importance of spatial probability summation. *Vision Research*, 27, 815–826.
- Graham, N. V. S. (1989). *Visual pattern analyzers*. New York: Oxford University Press.
- Green, D. M., & Swets, J. A. (1988). *Signal detection theory and psychophysics*. Los Altos, CA: Peninsula.
- Gurnsey, R., Herbert, A. M., & Kenemy, J. (1998). Bilateral symmetry embedded in noise is detected accurately only at fixation. *Vision Research*, 38, 3795–3803.
- Gurnsey, R., & Sally, S. Symmetry detection across the visual field. *Spatial Vision* (submitted).
- Herbert, A. M., & Humphrey, G. K. (1996). Bilateral symmetry detection: testing a 'callosal' hypothesis. *Perception*, 25(4), 463–480.
- Hess, R. F., & Dakin, S. C. (1997). Absence of contour linking in peripheral vision. *Nature*, 390(6660), 602–604.
- Horridge, G. A. (1996). The honeybee (*apis mellifera*) detects bilateral symmetry and discriminates its axis. *Journal of Insect Physiology*, 42, 755–764.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–243.
- Jenkins, B. (1983a). Component processes in the perception of bilaterally symmetric dot textures. *Perception and Psychophysics*, 34(5), 433–440.
- Jenkins, B. (1983b). Spatial limits to the detection of transpositional symmetry in dynamic dot textures. *Journal of Experimental Psychology: Human Perception & Performance*, 9(2), 258–269.
- Joung, W., van der Zwan, R., & Latimer, C. R. (2000). Tilt aftereffects generated by bilaterally symmetrical patterns. *Spatial Vision*, 13(1), 107–128.
- Julesz, B., & Chang, J. (1979). Symmetry perception and spatial-frequency channels. *Perception*, 8(6), 711–718.
- Koepl, U., & Morgan, M. (1993). Local orientation versus local position as determinants of perceived symmetry. *Perception*, 22, 111.
- Labonte, F., Shapira, Y., Cohen, P., & Faubert, J. (1995). A model for global symmetry detection in dense images. Special issue: the

- perception of symmetry: II. Empirical aspects. *Spatial Vision*, 9(1), 33–55.
- Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, 70, 1458–1470.
- Locher, P. J., & Wagemans, J. (1993). Effects of element type and spatial grouping on symmetry detection. *Perception*, 22(5), 565–587.
- Maffei, L., & Fiorentini, A. (1972). Retinogeniculate convergence and analysis of contrast. *Journal of Neurophysiology*, 35, 65–72.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London Series B*, 207, 187–217.
- Møller, A. P. (1995). Bumblebee preference for symmetrical flowers. *Proceedings of the National Academy of Science USA*, 92, 371–376.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(13), 607–609.
- Osorio, D. (1996). Symmetry detection by categorization of spatial phase, a model. *Proceedings of the Royal Society of London — Series B*, 263(1366), 105–110.
- Palmer, S. E., & Hemenway, K. (1978). Orientation and symmetry: effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception & Performance*, 4(4), 691–702.
- Pani, J. R. (1994). The generalized cone in human spatial organization. Special issue: the perception of symmetry: part I. Theoretical aspects. *Spatial Vision*, 8(4), 491–501.
- Pashler, H. (1990). Coordinate frame for symmetry detection and object recognition. *Journal of Experimental Psychology: Human Perception & Performance*, 16(1), 150–163.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Pintsov, D. A. (1989). Invariant pattern recognition, symmetry, and radon transforms. *Journal of the Optical Society of America*, 6(10), 1544–1554.
- Polat, U., & Sagi, D. (1993). Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Research*, 33(7), 993–999.
- Rainville, S. J. M. (1999). *The spatial mechanisms mediating the perception of mirror symmetry in human vision*. Unpublished Ph.D. Thesis, McGill University, Montréal, Québec, Canada.
- Rainville, S. J. M., & Kingdom, F. A. A. (1997). Is motion perception sensitive to local phase structures? *Investigative Ophthalmology*, 38(4), S215.
- Rainville, S. J. M., & Kingdom, F. A. A. (1998a). Does oblique structure support the detection of mirror symmetry? *Investigative Ophthalmology*, 39(4), 170.
- Rainville, S. J. M., & Kingdom, F. A. A. (1998b). From spatial filters to mirror symmetry: new findings and new model. *Perception (Supplement)*, 27, 58a.
- Rainville, S. J. M., & Kingdom, F. A. A. (1999a). The detection of mirror symmetry over a continuum of contrast polarities. *Investigative Ophthalmology and Visual Science*, 40(4), S359.
- Rainville, S. J. M., & Kingdom, F. A. A. (1999b). Spatial-scale contribution to the detection of mirror symmetry in fractal noise. *Journal of the Optical Society of America A*, 16(9), 2112–2123.
- Rainville, S. J. M., & Kingdom, F. A. A. (2000). The detection of mirror symmetry in spatially-redundant stimuli. *Investigative Ophthalmology*, 41 (4), S218.
- Rainville, S. J. M., & Kingdom, F. A. A. The spatial-frequency and orientation tuning of mirror-symmetry detection in sparse stimuli (in preparation).
- Rainville, S. J. M., & Kingdom, F. A. A. The perception of mirror symmetry is density invariant, not scale invariant. *Vision Research* (under revision).
- Swaddle, J. P., & Cuthill, I. C. (1994). Preference for symmetric males by female zebra finches. *Nature*, 367(6459), 165–166.
- Tapiovaara, M. (1990). Ideal observer and absolute efficiency in detecting mirror symmetry in random images. *Journal of the Optical Society of America*, 7(12), 2245–2253.
- Tyler, C. W., & Hardage, L. (1996). Mirror symmetry detection: predominance of second-order pattern processing throughout the visual field. In C. W. Tyler, *Human symmetry perception and its computational analysis*. Utrecht, The Netherlands: VSP.
- Tyler, C. W., Hardage, L., & Miller, R. T. (1995). Multiple mechanisms for the detection of mirror symmetry. Special issue: the perception of symmetry: II. Empirical aspects. *Spatial Vision*, 9(1), 79–100.
- van der Zwan, R., Leo, E., Joung, W., Latimer, C., & Wenderoth, P. (1998). Evidence that both area VI and extrastriate visual cortex contribute to symmetry perception. *Current Biology*, 8, 889–892.
- Victor, J. D., & Conte, M. M. (1996). The role of high-order phase correlations in texture processing. *Vision Research*, 36(11), 1615–1631.
- Wagemans, J. (1993). Skewed symmetry: a nonaccidental property used to perceive visual forms. *Journal of Experimental Psychology: Human Perception & Performance*, 19(2), 364–380.
- Wagemans, J., Van Gool, L., & Van Horebeek, J. (1991). Orientation selective channels in symmetry detection: effects of cooperation and attention. In B. Blum, *Channels in the visual nervous system: neurophysiology, psychophysics and models* (pp. 425–445). London: Freund Publishing House.
- Wenderoth, P. (1994). The salience of vertical symmetry. *Perception*, 23(2), 221–236.
- Wenderoth, P. (1996). The effects of the contrast polarity of dot-pair partners on the detection of bilateral symmetry. *Perception*, 25(7), 757–771.
- Wilson, H., McFarlane, D., & Phillips, G. (1983). Spatial frequency tuning of orientation selective units estimated by oblique masking. *Vision Research*, 23, 873–882.
- Wilson, H. R., & Gelb, D. J. (1983). Modified line element theory for spatial frequency and width discrimination. *Journal of the Optical Society of America A*, 1, 124–131.
- Zabrodsky, H., & Algom, D. (1994). Continuous symmetry: a model for human figural perception. Special issue: the perception of symmetry: part I. Theoretical aspects. *Spatial Vision*, 8(4), 455–467.
- Zhang, L. (1991). *Symmetry perception in human vision*. Unpublished Ph.D dissertation, University of Trieste, Trieste.