



The spatial region of integration for visual symmetry detection

Steven C. Dakin^{1*} and Andrew M. Herbert²

¹McGill Vision Research, Department of Ophthalmology, 687 Pine Avenue West H4-14, Montréal, Québec, Canada H3A 1A1 (scdakin@vision.mcgill.ca)

²École d'optométrie, Université de Montréal, C.P. 6128 Succursale Centre-Ville, Montréal, Québec, Canada H3C 3J7

Symmetry is a complex image property that is exploited by a sufficiently wide range of species to indicate that it is detected using simple visual mechanisms. These mechanisms rely on measurements made close to the axis of symmetry. We investigated the size and shape of this integration region (IR) by measuring human detection of spatially band-pass symmetrical patches embedded in noise. Resistance to disruption of symmetry (in the form of random phase noise) improves with increasing patch size, and then asymptotes when the embedded region fills the IR. The size of the IR is shown to vary in inverse proportion to spatial frequency; i.e. symmetry detection exhibits scale invariance. The IR is shown to have rigid dimensions, elongated in the direction of the axis of symmetry, with an aspect ratio of *ca.* 2:1. These results are consistent with a central role for spatial filtering in symmetry detection.

Keywords: symmetry; spatial frequency; filtering; scale selection

1. INTRODUCTION

Mirror symmetry is attractive. Many species rely upon it as an indicator of the fitness of potential mates: from swallows, with their predilection for symmetrical ornamentation (Møller 1992), to humans, who find faces with symmetrically distributed features attractive (Grammer & Thornhill 1994). Moreover, symmetry detection does not demand great computational resources; given the limitations of insect spatial vision, bees can detect symmetry (Horridge 1996), as manifested in their tendency to visit more symmetrical flowers (which yield greater amounts of nectar (Møller 1995)). More fundamentally, symmetry can signal the *presence* of objects and could direct visual attention. In short, symmetry is biologically significant and can be detected by simple visual processes.

Visual grouping is generally thought to be achieved by features mutually exciting receptive fields, a process modelled using spatial filters (e.g. Marr & Hildreth 1980; Watt 1988). How could symmetry be derived from basic visual mechanisms such as filters? The operation of a horizontal filter (followed by thresholding) on the symmetrical pattern shown in figure 1a is illustrated in figure 1b (the filter used is inset). Notice the clustering of 'blobs' around the axis. Such filtering, followed by measurement of the co-alignment of the position of blobs, is a biologically plausible scheme for pre-attentive measurement of local symmetry (Dakin & Watt 1994; for alternative filtering approaches see Osorio (1996), Bowns & Morgan (1993) and Gurnsey *et al.* (1998)). This scheme predicts human symmetry detection in the presence of various forms of disruption (Dakin & Watt 1994) and also

correctly predicts that information *orthogonal* to the axis is more resistant to the intrusion of noise than information *parallel* to the axis (Dakin & Hess 1997).

This model relies on correlations introduced by symmetry and, consequently, makes implicit assumptions about symmetrical feature pairs. For example, it assumes that the features comprising each pair will have the same sign of contrast. Contrast inversion, which reduces structured filter output to zero, is similarly disruptive to symmetry perception (figure 1c; Wenderoth 1996). This model also relies exclusively upon features that are sufficiently close to one another for their grouping to be made explicit by a single blob in the filter's response (e.g. figure 1b). Although symmetrical feature pairs are positioned at a range of separations, the model only uses features close to the axis. However, human observers seem to be similarly dependent on such features. Figure 1d shows a strip of symmetry embedded in noise. At exposure times that preclude detailed scrutiny of the image the entire pattern appears symmetrical, although only 16% of it is.

Several studies (Barlow & Reeves 1979; Jenkins 1982; Tyler *et al.* 1995; Labonté *et al.* 1995; Wenderoth 1995) have examined the latter phenomenon (first observed by Julesz 1975). Similarly, violations of symmetry are most salient near the axis (Bruce & Morgan 1975). Taken together, these results suggest that symmetry detection employs a limited spatial region of integration, a conclusion that is consistent with a central role for filtering operations. More detailed conclusions (e.g. the dimensions of the IR) are precluded by limitations of the stimuli used in previous studies (dot patterns). Dot patterns are not only spatially broad-band (and mask any systematic relationship between IR and spatial frequency), but they also over-represent high spatial frequencies. Furthermore, only one dimension of embedded symmetrical regions has been

*Author for correspondence.




Figure 1. (a) Spatially band-limited symmetrical texture. (b) Same texture convolved with low-frequency, horizontally orientated filter (inset), and 'thresholded' to remove values close to the mean grey level. Note clustering of resulting blobs around the axis, and the co-alignment of their centroids (crosses). (c) Symmetrical pattern where the contrast polarity of symmetrical features has been reversed across the axis. (d) Noise pattern containing a strip of symmetry around the axis; at brief presentation times, the entire pattern appears symmetrical.

studied systematically (the width, using embedded 'strips' (Jenkins 1982)). If local measurements of symmetry are to be useful (for segmentation, or to provide an object-based axis of representation), then the IR must be bound in two dimensions to avoid including irrelevant information falling outside the boundaries of the target object.

This paper examines the size and shape of the IR and, particularly, whether it scales with the spatial frequency of the pattern. A 'high-level' symmetry detector—operating on position and shape cues from an earlier form analysis stage—will show similar integration properties regardless of the spatial frequency content of the pattern. By contrast, a 'low-level' system, which is closely linked to spatial filtering, will show a strong dependence on pattern spatial frequency. Figure 2 illustrates the logic of the experiments reported. We measured detection of symmetrical patches embedded in noise backgrounds with varying degrees of disruption added to the symmetry. Resistance to this disruption ('noise resistance') will plateau when the patch covers the IR; by measuring this 'kneepoint' for circular and elongated patches we can estimate the IR. At first sight, this methodology might appear overly complex; why not just measure discrimination of embedded, undisrupted symmetrical patches from pure noise? The problem with this approach is that one would simply estimate the minimum detectable symmetrical patch size, which is largely determined by the physical properties of the patterns (i.e. the degree of undersampling introduced by restricting spatial frequency). Instead, we are attempting to measure the integration region by estimating the minimum patch size that gives optimal resistance to noise.

2. GENERAL METHODS

(a) Apparatus

A Macintosh 7500/100 computer generated stimuli and recorded subjects' responses. Stimuli were displayed (at 75 Hz) on a Nanao Flexscan 6600 monochrome monitor fitted with a video attenuator (ISR Instruments, Syracuse). Luminance levels were linearized using VideoToolbox routines (Pelli 1997). The screen was viewed binocularly at a distance of 98 cm and had a mean background luminance of 45 cd m⁻².

(b) Stimuli

Stimuli were 512 × 512 pixel, band-pass filtered noise textures containing patches with varying degrees of bilateral symmetry (examples are shown in figure 3). Images were generated from samples of Gaussian, randomly distributed luminance into which elliptical-shaped regions of symmetrical texture were inserted. Phase disruption, a uniform random offset added to the phase angle, and filtering were performed in the Fourier domain. Filters were idealized band-pass (i.e. sharp cut-off) with bandwidths of one octave. (Problems associated with idealized filters, particularly the appearance of 'ringing' artefacts, are limited to images containing significant phase-alignments across scale (e.g. edges) and are not visible in noise patterns.) Images were normalized to a root mean square contrast with $\sigma = 32$ grey levels.

Textures subtended 10 deg² and were presented in the centre of the display for 250 ms (abrupt onset and offset).

(c) Procedure

A two-interval, two-alternative forced-choice paradigm was used. Observers (the authors and a naive subject, H.Y.W.) were presented with two patterns (each preceded by a centrally located fixation marker) and judged in which interval the symmetrical region was present. In all experiments the threshold level of phase disruption was estimated as a function of the size of the embedded patch. A method of constant stimuli was used to sample representative amounts of phase disruption from 0° to 240° in steps of 30°. A pattern with 240° of imposed phase disruption is indistinguishable from pure noise and was therefore selected as the maximum level of noise to avoid 'steps' in measured psychometric functions. Runs consisted of 144 trials and were not interleaved. A cumulative Gaussian was fit to derived psychometric functions, and the standard deviations of these fits (corresponding to the 84% correct point) are the data reported for each patch size.

3. EXPERIMENT 1. SCALING OF INTEGRATION REGION WITH SPATIAL FREQUENCY

The first experiment looked at how the region of integration for symmetry detection depends upon the spatial frequency of the pattern. Subjects' discrimination of embedded symmetry from noise was measured with a variety of region sizes at spatial frequencies of 1.13, 2.26 and 4.52 cycles per degree (c.p.d.). Figure 3 shows example stimuli at two noise levels and two spatial frequencies.

(a) Results

Results for three subjects for experiment 1 are shown in figure 4. Each data point represents the level of phase disruption leading to 84% correct discrimination of




Figure 2. The logic of the experiments. (a) We assume symmetry detection operates within an integration region (IR) of limited spatial extent and (b) make an estimate of its longest dimension by measuring detection of embedded circular regions of symmetry as a function of the degree to which the symmetry is disrupted (phase noise). Subjects are able to withstand more phase noise (the accelerating portion of the graph) as the size of the symmetrical patch increases until the patch completely covers the IR. Beyond this point (indicated by the dashed arrow) subjects' ability to resist the intrusion of noise remains constant (the plateau shown in the graph). The patch radius at which this plateau occurs is therefore an indication of the size of the IR. (c) Experiment 2 refines our estimate of the IR by fixing the width or the height of the embedded region and measuring noise resistance as a function of the other dimension.

symmetry from noise, as a function of the size of an embedded patch. Notice that resistance to noise increases steadily with region size and then plateaus. Maximum resistance to noise peaks at around $180\text{--}190^\circ$ for all subjects at all spatial frequencies, although A.H. shows slightly higher noise resistance with high-frequency patterns (his data are fit separately in figure 4c). However, these data are basically consistent with results from Dakin & Hess (1997), who report no advantage for symmetry detection at any particular spatial frequency. Comparing data between conditions, curves are shifted versions of one another (see summary in figure 4d). The region size at which performance plateaus varies in inverse proportion to the spatial frequency of the pattern.

This experiment clearly indicates that the spatial frequency content of a symmetrical pattern determines the region over which subjects integrate information, which in turn suggests that symmetry detection is a 'low-level' visual task closely linked to the operation of spatial filters.

4. EXPERIMENT 2. THE ASPECT RATIO OF THE INTEGRATION REGION

Because circular embedded regions were used in experiment 1, the IR estimates given could be the result of the IR being limited either in its dimensions or its area. Assuming the former, results from experiment 1 effectively measured the *maximum* dimension of the IR. This point is illustrated in figure 2, as is the logic of experiment 2. By using elongated embedded patches, with one dimension fixed at the minimum patch size producing plateaued performance in the previous experiment, we can determine the direction of elongation (if any) of the IR. Moreover, elongation of patches allows us to test the fixed-area hypothesis. If the dimensions of the IR are actually flexible with respect to task demands, with IR area limiting performance, elongation of the embedded region along any dimension will have a similar effect.

All patterns in this condition had a peak spatial frequency of 2.26 c.p.d. We fixed either the width or the height of embedded patches at 3.0 degrees (the kneepoint




Figure 3. Examples of the stimuli used in (a–d) experiment 1 and (e–h) experiment 2. (a–d) Three degree-wide patches embedded in noise at (a,c) 1.13 and (b,d) 4.56 c.p.d. Adding 90° phase disruption (c,d) is more disruptive to perception of structure in the low-frequency pattern. (e–h) Embedded patches of (e,g) 3.0° × 1.5° and (f,h) 1.5° × 3.0°, with a peak spatial frequency of 2.26 c.p.d. (g,h) illustrate that noise is more disruptive to stimuli containing a horizontally elongated patch. (The dashed lines indicating the location of the patches did not appear during the experiment.)

patch radius from experiment 1) and varied the other dimension between 0.5 and 10 degrees.

(a) Results

Figure 5 gives data from two subjects in this condition. The lack of coincidence between the fixed-width and fixed-height conditions indicates that the dimensions of the IR do not depend upon the type of stimuli being presented to the subject. Estimates from experiment 1 are therefore due to limits on the dimensions of the IR and not a limited *area* (measurements from experiment 2, with similar positions on the abscissa, have equal areas). Notice also that the effect of reducing the height of the embedded patch (open symbols) is similar to the effect of changing the radius of an embedded circular patch (solid line) at the same spatial frequency. This suggests that the limit on the IR measured in the previous experiment was

due to the *height* of the integration region. Changing the width of the patch produces quite different results (filled symbols). Subjects can withstand much greater reduction in the area of fixed-height patches than fixed-width patches before noise resistance collapses (the kneepoint of the fit to this condition is at approximately 1.5 degrees). These data are consistent with the IR having an aspect ratio of approximately 2:1. Figure 3f illustrates the kneepoint of the variable-width condition for this experiment so that the dashed region thus effectively indicates the integration region employed for this task.

Elongation of the IR in the direction of the axis might be seen as arguing for mechanisms for coding symmetry that are themselves orientated in the same direction as the axis (e.g. organized in a phase-categorization scheme (Osorio 1996)). This view is incorrect because a filter elongated in this manner must have a selectivity for lower spatial frequencies in this direction. This means the filter is effectively grouping information in the direction of the axis, whereas symmetry introduces structure/correlation that is *perpendicular* to the axis. A filter orientated in the axis direction will therefore group clusters arising in the noise by chance and will tend towards a zero response. Figure 1b illustrates that a system exploiting blob clustering in the output of filters elongated perpendicular to the axis would be expected to integrate over a region that is elongated in the direction of the axis in order to make an accurate estimate of blob co-alignment.

5. DISCUSSION

The striking visual impression gained from symmetry, coupled with its ecological significance, suggests that there might be highly specialized mechanisms for detecting and coding visual symmetry. The experiments presented here extend previous findings (e.g. Jenkins 1982; Dakin & Watt 1994) suggesting that this view is wrong. The limited and rigid dimensions of the spatial region of integration, when combined with the knowledge that symmetry embedded in noise is only detectable at or near fixation (Gurnsey *et al.* 1998), argue against a general symmetry detection system operating in parallel over large portions of the visual field.

Our finding that the region of integration for symmetry detection scales with the spatial frequency content of the pattern indicates, at least within the psychophysical domain, that symmetry perception is achieved using the type of low-level visual operations that are more usually associated with the processing of texture. Moreover, our estimates of the IR are very small. In general, the IR captures about seven cycles of a filter output in the *y*-direction and about 3.5 cycles in the *x*-direction. For a 10 deg² symmetrical pattern, at 2.26 c.p.d., this means that over 95% of the pattern is completely redundant. Symmetrical transformations introduce structure that is correlated across large distances, and it seems highly unlikely that a specialized symmetry detection system (performing some kind of back transformation) would ignore so much information. The alternative hypothesis is that humans are very poor at estimating the true symmetry of objects. Instead, we rely on correlated image structure that is little more than a by-product of symmetry, but that is readily signalled by the visual mechanisms we have available, i.e. spatial filters.

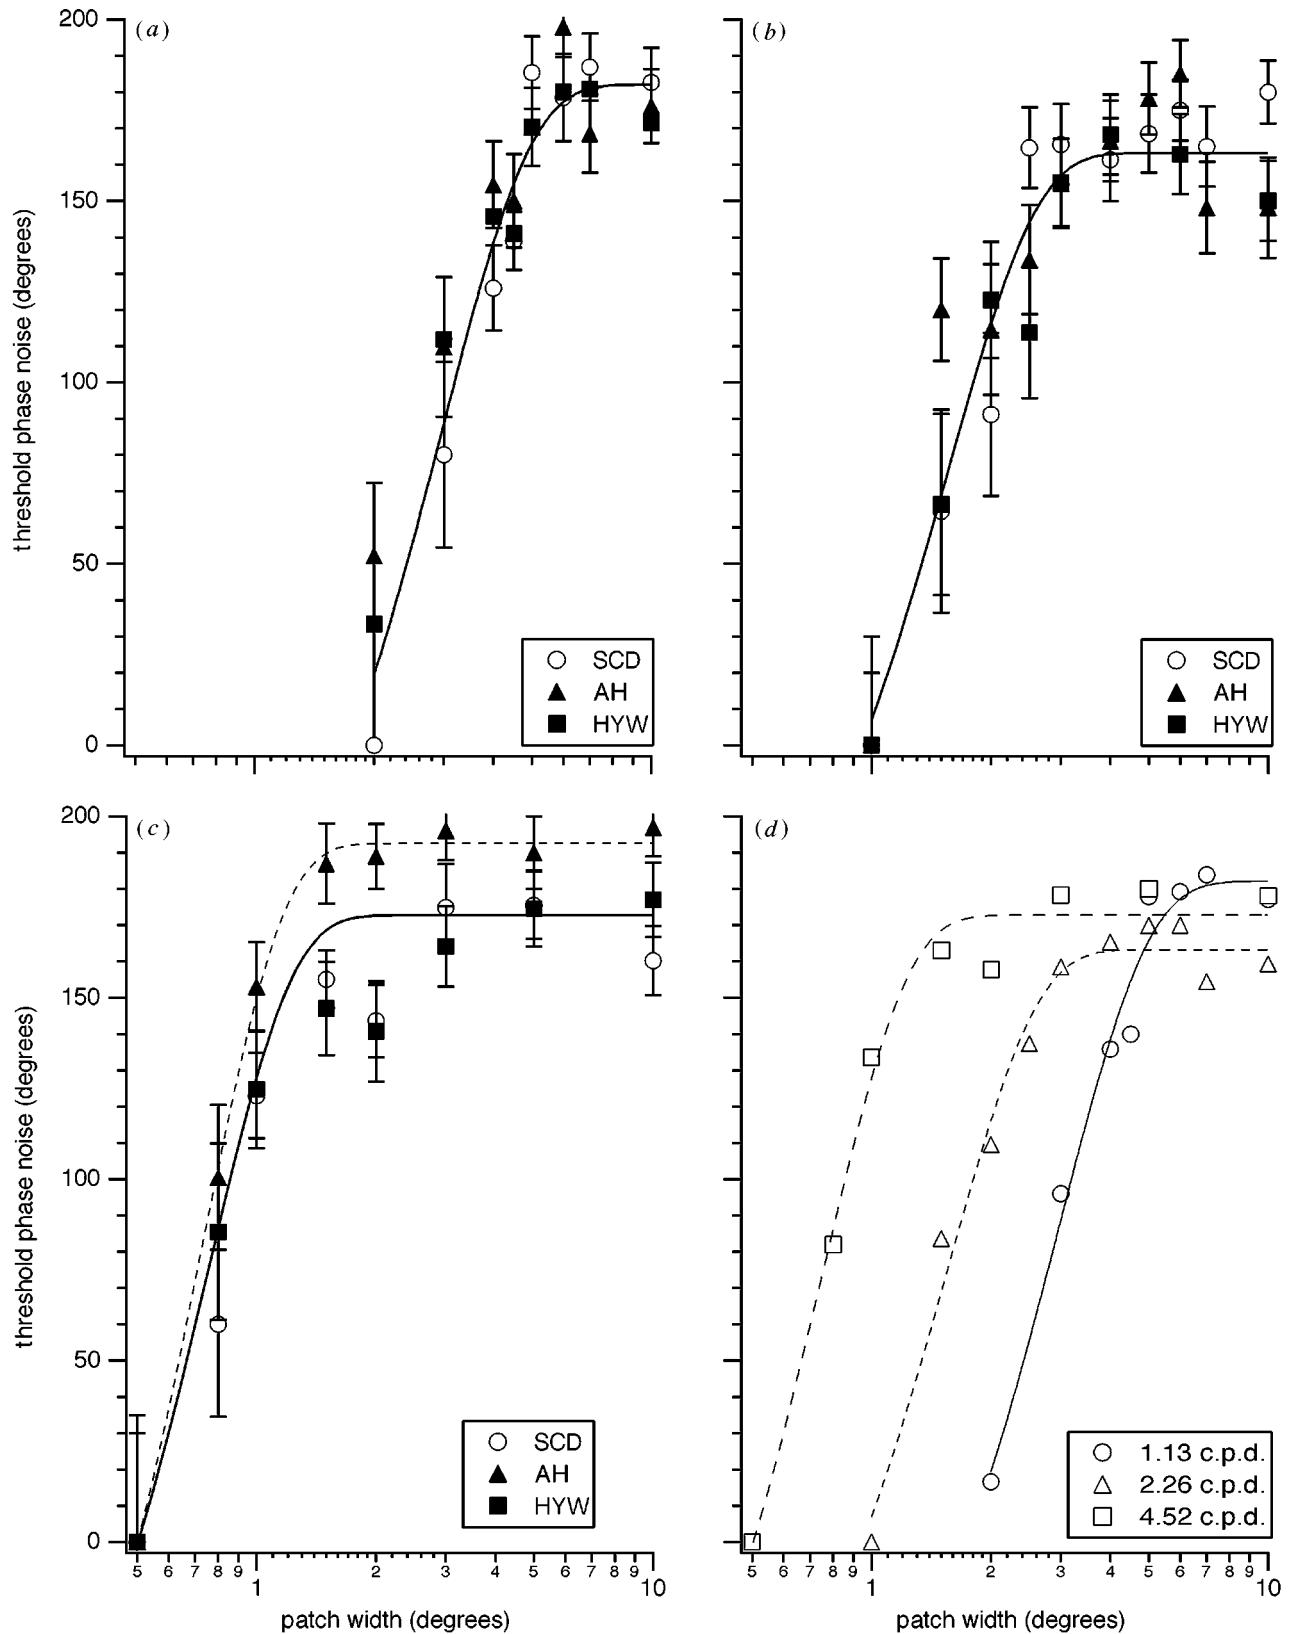


Figure 4. Results from experiment 1, for patterns with spatial frequencies of (a) 1.13 c.p.d., (b) 2.26 c.p.d. and (c) 4.52 c.p.d. Fits are the positive portion of an error function. Error bars show ± 1 s.e. Notice that the patch size at which noise resistance is maximized depends on the spatial frequency of the pattern. (d) A summary plot collapsing data across subjects.




Figure 5. Results from two subjects for experiment 2. Noise resistance as a function of patch height (open symbols; thin dashed fit) is compared with noise resistance as a function of patch width (filled symbols; thin solid fit). Notice that a greater reduction in patch width than height can be tolerated (data at the same position on the abscissa have the same area). These data are consistent with the IR being quite elongated (aspect ratio *ca.* 2:1) in the direction of the axis of symmetry.

However, this is not to say that filters are the only route to symmetry open to human vision. Given time, we can certainly detect structure in patterns that produce poor filter output (outline figures (e.g. Mach 1897); opposite contrast patterns, e.g. figure 1*d*) apparently by way of an explicit comparison of the shape and position of symmetrical features. Filter- and form-based symmetry processing could complement one another, with the former being used to rapidly signal the presence of symmetrical objects, and the latter being used to construct a full representation of shape. This two-stream approach is broadly consistent with the suggestions of many authors (e.g. Julesz 1975).

Compared to a matching process, the filtering scheme has the virtue of simplicity: its component mechanisms are likely to be possessed by even the simplest of visual systems capable of symmetry detection (e.g. neurons with orientated receptive fields are present in the mid-brain of the honeybee (Yang & Maddess 1997)). It remains an open empirical question as to whether other species also rely on spatially limited visual correlations to signal symmetry.

S.C.D and A.M.H. were supported by Canadian MRC grants (MT 108-18, to R. F. Hess and MA 135-33 to J. Faubert, respectively). Thanks to Isabelle Mareschal and Stephane Rainville for

their insightful comments on this project, and to Hok Yean Wong for acting as a subject.

REFERENCES

- Barlow, H. B. & Reeves, B. C. 1979 The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Res.* **19**, 783–793.
- Bowns, L. & Morgan, M. 1993 Facial features and axis of symmetry extracted using natural orientation information. *Biol. Cybern.* **70**, 137–144.
- Bruce, V. & Morgan, M. 1975 Violations of symmetry and repetition in visual patterns. *Perception* **4**, 239–249.
- Dakin, S. C. & Hess, R. F. 1997 The spatial mechanisms mediating symmetry perception. *Vision Res.* **37**, 2915–2930.
- Dakin, S. C. & Watt, R. J. 1997 Detection of bilateral symmetry using spatial filters. *Spatial Vis.* **8**, 393–413.
- Grammer, K. & Thornhill, R. 1994 Human (*Homo sapiens*) facial attractiveness and sexual selection: the role of symmetry and averageness. *J. Comp. Psychol.* **108**, 233–242.
- Gurnsey, R., Herbert, A. M. & Kenemy, J. 1998 Bilateral symmetry embedded in noise is detected accurately only at fixation. *Vision Res.* (In the press.)
- Horridge, G. A. 1996 The honeybee (*Apis mellifera*) detects bilateral symmetry and discriminates its axis. *J. Insect Physiol.* **42**, 755–764.
- Jenkins, B. 1982 Redundancy in the perception of bilateral symmetry in dot textures. *Percept. Psychophys.* **32**, 443–448.
- Julesz, B. 1975 Experiments in the visual perception of texture. *Scient. Am.* **232**, 34–43.
- Labonté, F., Shapira, Y., Cohen, P. & Faubert, J. 1995 A model of global symmetry detection in dense images. *Spatial Vis.* **9**, 33–55.
- Mach, E. 1897 *Contributions to the analysis of the sensations*. La Salle, Illinois: Open Court.
- Marr, D. & Hildreth, E. 1980 Theory of edge detection. *Proc. R. Soc. Lond. B* **207**, 187–217.
- Møller, A. P. 1992 Female swallow preference for symmetrical male sexual ornaments. *Nature* **357**, 238–240.
- Møller, A. P. 1995 Bumblebee preference for symmetrical flowers. *Proc. Natn. Acad. Sci. USA* **92**, 2288–2292.
- Osorio, D. 1996 Symmetry detection by categorization of spatial phase. *Proc. R. Soc. Lond. B* **263**, 105–110.
- Pelli, D. D. 1997 The VideoToolbox software for visual psychophysics: transforming number into movies. *Spatial Vis.* **10**, 437–442.
- Tyler, C. W., Hardage, L. & Miller, R. 1995 Multiple mechanisms for the detection of visual symmetry. *Spatial Vis.* **9**, 79–100.
- Watt, R. 1988 *Visual processing: computational psychophysical and cognitive research*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Wenderoth, P. 1995 The role of pattern outline in bilateral symmetry detection with briefly flashed dot patterns. *Spatial Vis.* **9**, 57–77.
- Wenderoth, P. 1996 The effect of contrast polarity of dot-pair partners on the detection of bilateral symmetry. *Perception* **25**, 757–772.
- Yang, E. & Maddess, T. 1997 Orientation-sensitive neurons in the brain of the honeybee (*Apis mellifera*). *J. Insect Physiol.* **43**, 329–336.