Pergamon

PII: S0042-6989(97)00031-X

Vision Res., Vol. 37, No. 20, pp. 2915–2930, 1997

Printed in Great Britain
0042-6989/97 $17.00 + 0.00

The Spatial Mechanisms Mediating Symmetry Perception

STEVEN C. DAKIN,*† ROBERT F. HESS*

Received 29 October 1996; in revised form 6 January 1997

This paper examines the role of spatial frequency and orientation tuned channels in the perception of visual symmetry. Subjects discriminated between band-pass filtered, white noise textures that either did or did not contain vertical bilateral symmetry (VBS, i.e., around a vertical midline) as a function of the spatial phase disruption imposed on the images. Resistance to phase noise is largely scale-invariant for isotropically filtered images, but horizontally filtered images are consistently more noise-resistant than vertical. However, when stimuli are rotated through 90 deg (*horizontal* bilateral symmetry, HBS) performance is better with *vertically* filtered images suggesting a general advantage for orientations orthogonal to the axis of symmetry. At these orientations symmetry may be signaled *directly* by clusters of features along the axis. Our data further suggest that the established disadvantage for HBS may be attributable to an over-reliance on the output of horizontal filters. We compare models which exploit feature clustering around the axis by measuring the co-alignment in the output of oriented filters. Models using filters oriented orthogonal to the axis of symmetry predict the psychophysical performance for isotropic patterns and for patterns filtered orthogonal to the axis. For patterns filtered parallel to the axis, our data suggest that visual attention may play a role. © 1997 Elsevier Science Ltd

Symmetry Filtering Orientation Masking Phase

INTRODUCTION

Given the sizable psychophysical literature exploring the human perception of visual symmetries (for a review see Wagemans, 1995) there is relatively little work linking established visual mechanisms to human performance. 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 (e.g. Palmer & Hemenway, 1978; Wagemans *et al.*, 1993; Zabrodsky & Algom, 1994). The most pervasive example of this is the modeling of human symmetry perception using cross-correlation around the axis (e.g. Barlow & Reeves, 1979; Pintsov, 1989). Jenkins (1983) points out that because experimental stimuli for the psychophysical investigation of symmetry detection use mirror reflection of elements it has been assumed that observers must perform some type of reverse mapping in order to detect structure. As well as requiring a complex physiological system dedicated solely to the processing of symmetry, this view is largely unsupported by the psychophysical literature.

Symmetry detection is highly resistant to perturbation of the position of elements comprising a pattern (Barlow & Reeves, 1979). This suggests that symmetry detection operates at a low spatial resolution, as do findings that the orientation of elements has little effect on symmetry detection (Koeppl, 1993; Locher & Wagemans, 1993). However, when symmetrical pairs are composed of opposite contrast elements (Zhang, 1991; Carlin, 1996; Wenderoth, 1996b) or when symmetrical elements are too distant from one another; as in the case when symmetry is disrupted only around the axis (Barlow & Reeves, 1979; Jenkins, 1982; Tyler *et al.*, 1995) symmetry detection breaks down. These results implicate mechanisms operating at a coarse spatial scale that are sensitive to both the relative contrast and the proximity of components. A natural candidate for this mechanism, given the body of work examining their role in other forms of grouping (e.g. Zucker, 1982), are spatial filters. The question is how one might measure symmetry from the output of filters?

Symmetry detection using spatial filters

Dakin & Watt (1994) observed that when images containing vertical bilateral symmetry (VBS) are spatially filtered and half-wave rectified, resultant "blobs" cluster around the axis of symmetry [Fig. 1(b)]. A measure of the degree of co-alignment of the blob

*McGill Vision Research, Department of Ophthalmology, 687 Pine Avenue West, Montreal, Quebec H3A 1A1, Canada.
†To whom all correspondence should be addressed [*Email*: scdakin@astra.vision.mcgill.ca].

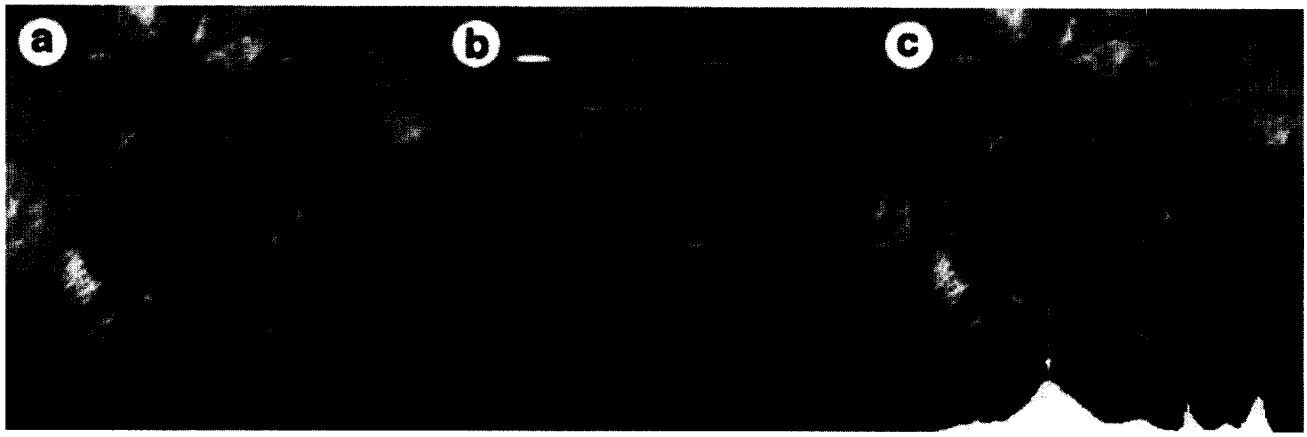

FIGURE 1. Symmetry detection using spatial filters. (a) A symmetrical image, processed with (b) a horizontal Difference-of-Gaussian filter, followed by half-wave rectification. (c) The histogram along the bottom of the image shows the alignment of blobs extracted from (b) lying along a particular image column. Notice that the peak, marked by the arrow, indicates the axis of symmetry.

centroids was proposed as a symmetry metric. The predictions of four symmetry detection models—using combinations of either isotropic or oriented filtering (with filters oriented orthogonal to the axis), and either the co-alignment metric or a bilateral correlation measure—were used to simulate data from various symmetry detection tasks. The tasks used were symmetry detection in the presence of uncorrelated elements and positional jitter (Barlow & Reeves, 1979), and the location of symmetrical regions embedded in noise (Jenkins, 1983). For VBS, a model using the co-alignment measure and horizontal Difference-of-Gaussian (DoG) filtering was shown to produce good agreement with data from all conditions.

Recently Osorio (1996) has proposed a model of symmetry detection using Gabor filters in quadrature phase (e.g. Daugman, 1985). The output from such filters oriented parallel to the axis of symmetry is squared. Areas in the image of maximum symmetric filter output and minimum anti-symmetric filter output indicate that local spatial harmonics are 90 deg/270 deg out of phase. Areas of local phase alignment (i.e., that have harmonics 90 deg/270 deg out of phase at two spatial scales) indicate the axis of symmetry. This model will detect element correlations that lay within a filter's receptive field i.e., extremely close to the axis of symmetry. Whilst it is well established that the area around the axis of symmetry is important to the perception of symmetry, it is equally clear that we are capable of detecting symmetry in patterns that do not contain such correlations, albeit at a reduced level of performance (Barlow & Reeves, 1979; Jenkins, 1982; Tyler *et al.*, 1995). This model also produces a clear and testable prediction: that, in the Fourier domain, it is information parallel to the axis of symmetry that will determine the percept of symmetry. The relative activity of two Gabors, 90 deg out of phase and oriented orthogonal to the axis of symmetry, cannot signal symmetry.

In summary, two models proposed for interpreting the output of spatial filtering produce contradictory hypotheses regarding the orientation of filters used to detect symmetry. This paper addresses the issue directly by considering the effect of spatial filtering on symmetry perception. Specifically, we consider what spatial frequencies and orientations are used in detecting symmetry, and whether this depends on the orientation of the axis of symmetry.

In order to investigate the relative importance of information at different scales and orientations we employed a selective phase randomization technique (Victor & Conte, 1996). Images were Fourier transformed and, at each point in their phase/amplitude spectra, a random offset added to the phase whilst maintaining equal power. The magnitude of this offset determines the coherence of phase information and, because this determines the positioning of local features, the strength of the symmetry percept. This technique is similar to a simultaneous luminance masking paradigm, where band-pass filtered white noise (which will have random phase) is added to a band-pass filtered image (Parish & Sperling, 1991). The advantage of the phase randomization technique is that it makes no assumptions about the summation of the mask and noise luminances, or the effect of disrupting the local power. Only phase information is altered.

GENERAL METHODS

Subjects

The first author and one naïve subject served as observers in all experiments. Both are corrected-to-normal myopes, and SCD has a small (<0.5 D) corrected astigmatism. Sufficient practice was undertaken for observers to reach asymptotic performance before threshold measurement began.

Apparatus

The stimuli were generated and presented using a Macintosh 7500 microcomputer which also recorded subjects' responses. Stimuli were displayed on a Nanao

Flexscan 6500 monochrome monitor, with a frame refresh rate of 75 Hz. Luminance levels were linearized using a look-up table derived using programs from Denis Pelli's VideoToolbox package, from which display routines were also derived. The screen was viewed binocularly at a distance of 95 cm and had a mean background luminance of 23.6 cd/m^2.

Stimuli

Noise patterns were generated from 256 pixel square white noise textures with an initially uniform random distribution of luminances. Symmetrical images were generated by mirror-reflecting half of such a noise pattern around (unless stated otherwise) a vertical axis centrally located in the pattern. Images were then Fast Fourier Transformed (using the FFT routine described in Press *et al.*, 1992), partially phase randomized, filtered, back-transformed and finally contrast normalized.

Phase randomization consisted of the addition of a uniform random offset to the phase value while maintaining local power. The degree of phase randomization imposed was the independent variable in all conditions. Filtering was performed in the Fourier domain and used idealized band-pass and orientation-limited filters (i.e., sharp cut-off). All filters has spatial band-widths of one octave and when orientation information was to be limited it was clipped in the range ±10 deg around horizontal or vertical. The reasons for selection of only horizontal and vertical filtering are given in Experiment 2. After FFT back-transformation, images were normalized to a root mean square (RMS) contrast with standard deviation (SD) = 32 gray levels. Note that by equating RMS contrast across conditions, one is also matching for spectral density (Brady & Field, 1995).

All textures were presented in the center of the display and subtended 5 deg square. Individual pixels were 1.17 arc min square.

Procedure

In each experiment a two-alternative forced-choice (2AFC) procedure was used. Subjects were presented with one texture for 100 msec, followed by a 1.0 sec delay, followed by a second texture for 100 msec. Before each trial a fixation mark was presented in the center of the screen. One randomly selected interval contained a symmetrical texture and the other a noise texture. Subjects were asked to "judge which interval contained the most symmetrical image" and to indicate their decision by depressing one of two keys on the computer keyboard.

Phase noise was added to both noise and symmetrical textures and could be varied in the range 0 deg (unaltered phase) to 360 deg (complete phase randomization). A method of constant stimuli was used to sample representative phase noise levels along the psychometric function. Each block consisted of 288 trials. This consisted of 32 presentations at nine stimulus levels, corresponding to phase disruption from 0 to 240 deg in steps of 30 deg. This range was selected because pilot studies indicated that phase randomization outside that range produced images indistinguishable from pure noise.

EXPERIMENT 1: SYMMETRY DISCRIMINATION IN ISOTROPIC TEXTURES

As far as we are aware there have been no previous psychophysical investigations of the perception of symmetry in band-pass filtered textures. Julesz & Chang (1979) used demonstrations to show that if a horizontally and a vertically symmetric texture are filtered and added together, then perception of both symmetries is possible if the spatial frequency sensitivities of the filters differ by two octaves. They claimed that low spatial frequencies have "a stronger perceptual weight" than high and that such channels probably precede symmetry detection.

The first experiment examined the perception of symmetry in patterns containing information at a limited band of spatial scales, but at all orientations. Patterns were processed using idealized octave-wide band-pass filters centerd around 8, 16 and 32 cycles per image (corresponding to 1.6, 3.2 and 6.4 c/deg). Examples of symmetry and noise patterns are shown in Fig. 2. Subjectively at least, this figure suggests that it will be harder to detect symmetry in high spatial-frequency filtered images (Julesz & Chang, 1979).

Note that because isotropic band-pass filters remove information, their operation on partially phase-randomized symmetrical textures will change the degree of symmetry from the point of view of cross-correlation. One is effectively presenting a smaller sample of phase pairings in the Fourier domain. To ensure that this reduction in the number of phase-samples would not swamp any relative differences between scales, we ran the experiment described using a cross-correlator as an ideal discriminator and simulation results are presented as solid lines alongside data from observers in Fig. 3. The cross-correlator simply generated a correlation measure of all pixel-values on one side of the axis of symmetry, with pixel-values falling at mirror-reversed positions on the other side of the axis. The shape of these functions demonstrates that textures are equally discriminable in terms of their cross-correlation statistics at phase disruptions less than approximately 180 deg.

Results

Figure 3 shows the resistance to phase noise as a function of the peak spatial frequency of the pattern for the two subjects tested. It is evident that phase disruption below 90 deg has little effect on discrimination, but beyond this level performance rapidly decreases, approaching chance by about 180 deg. Both subjects show similar resistance to phase noise as spatial frequency increases. Such a lack of effect of scale is not surprising given that there was no uncertainty regarding the location of the axis within the texture. Subjects reported that they attended most closely to an area around the likely location of the axis of symmetry, and based decisions on feature pairings from that region.

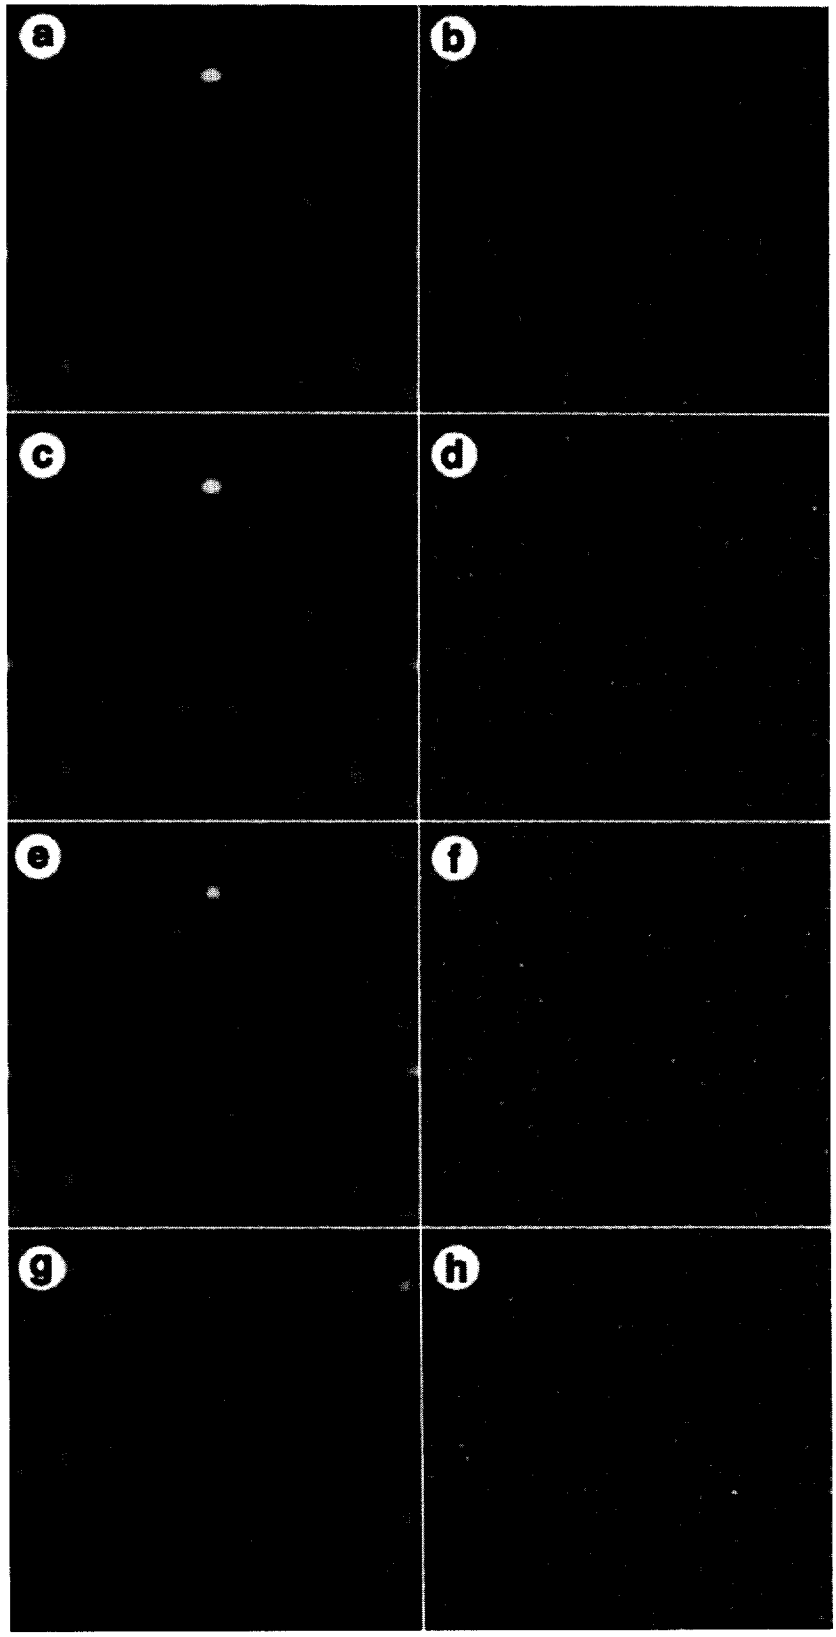

FIGURE 2. Examples of the stimuli used in Experiment 1. Patterns are band-pass limited to scales one octave around (a, c, e, g) 8 cycles per image and (b, d, f, h) 32 cycles per image. From top to bottom patterns have been phase-randomized to varying degrees: (a, b) 0 deg; (c, d) 60 deg; (e, f) 120 deg; (g, h) 360 deg.

Predictions from a cross-correlation mechanism are also presented in Fig. 3. Considering the cross-correlator as a discriminator, note that it grossly overestimates performance. At levels of phase noise around 180 deg, where the subject is approaching chance performance, this model still predicts performance at 95%. This is

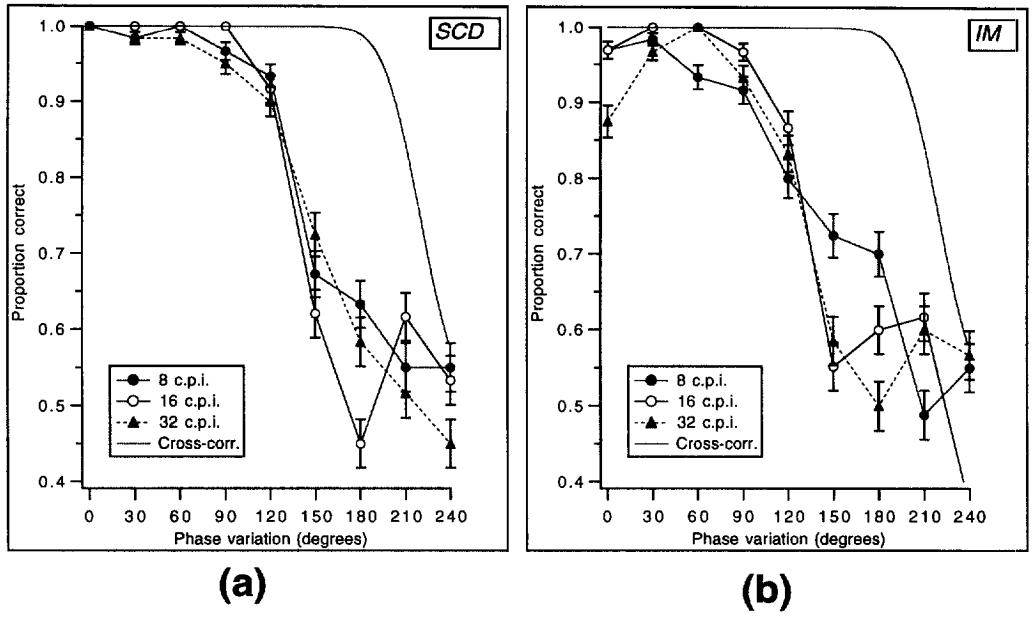

FIGURE 3. Symmetry discrimination performance, in the presence of phase disruption, for textures isotropically filtered at three spatial scales. The solid line is the prediction of a cross-correlator.

consistent with findings reported in Dakin & Watt (1994) that addition of an isotropic filtering pre-processing stage to a cross-correlator does not bring the model into better registration with subjects' performance on a number of symmetry discrimination tasks. In order to accommodate the data shown one would have to propose a plausible way of introducing noise into the cross-correlation process. Our data suggest that whatever this mechanism is, it will have to introduce variability equivalent to about 90 deg of phase variation to predict performance with these textures.

EXPERIMENT 2: SYMMETRY DISCRIMINATION IN ORIENTED TEXTURES

Given a lack of advantage for any particular spatial frequency, we asked if there would be any advantage for certain orientations of information in the pattern. The examination of orientation information in symmetrical textures introduces certain problems for stimulus generation. Consider filtering with an even-symmetric spatial filter with a symmetric orientation pass-band of limited extent (e.g. oriented DoG: Phillips & Wilson, 1983; Wilson & Gelb, 1984). This operation will have differential effects on the degree of symmetry in the pattern according to the orientation selected. For a vertically or horizontally bilaterally symmetric texture, only filtering with mechanisms centered at horizontal or vertical will not drastically reduce the degree of symmetry in the pattern. For that reason we assume, in the case of even-symmetric oriented filters, that filters at orientations other than horizontal or vertical cannot be responsible *in isolation* for signaling the presence of symmetry. Output from similar filters at other orientations could certainly be combined to calculate the degree of symmetry. Alternatively, this combination could occur through the use of filters with multiple preferred orientations, where those orientations are consistent with symmetry (e.g., 30 deg/150 deg, 45 deg/135 deg). Filter phase also differentially affects some orientations (for odd-symmetric filters, the degree of symmetry in horizontally filtered patterns will be unaffected, but will be disrupted in vertically filtered patterns). In order to allow comparison across orientation we considered only filtering operations that matched the degree of (cross-correlational) symmetry across filter orientation: i.e., only horizontal and vertical even-symmetric filters. Figure 4 shows examples of the filtered textures used in this experiment. All methods were identical to those used for Experiment 1. Orientations were not mixed within a single block of trials.

Results

Figure 5 shows the performance of two subjects on the symmetry discrimination task using horizontally and vertically oriented filtered patterns. Figure 5 also shows results from the previous isotropic condition, and the results of a simulation using identical experimental procedures with a cross-correlator as discriminator. The long dashed line in Fig. 5(a) shows the cross-correlating discrimination using isotropically filtered patterns. Notice that there is a small predicted deficit for oriented compared with isotropically filtered patterns (because they present a smaller sample of phase information). Discriminability is matched for the horizontally and vertically filtered cases.

Human data are considerably different. Figure 5 shows that subjects are significantly more sensitive to the introduction of phase noise with vertically filtered patterns at all spatial scales. The greater resistance to phase noise shown for horizontally filtered patterns produces performance similar to the task using isotropi-

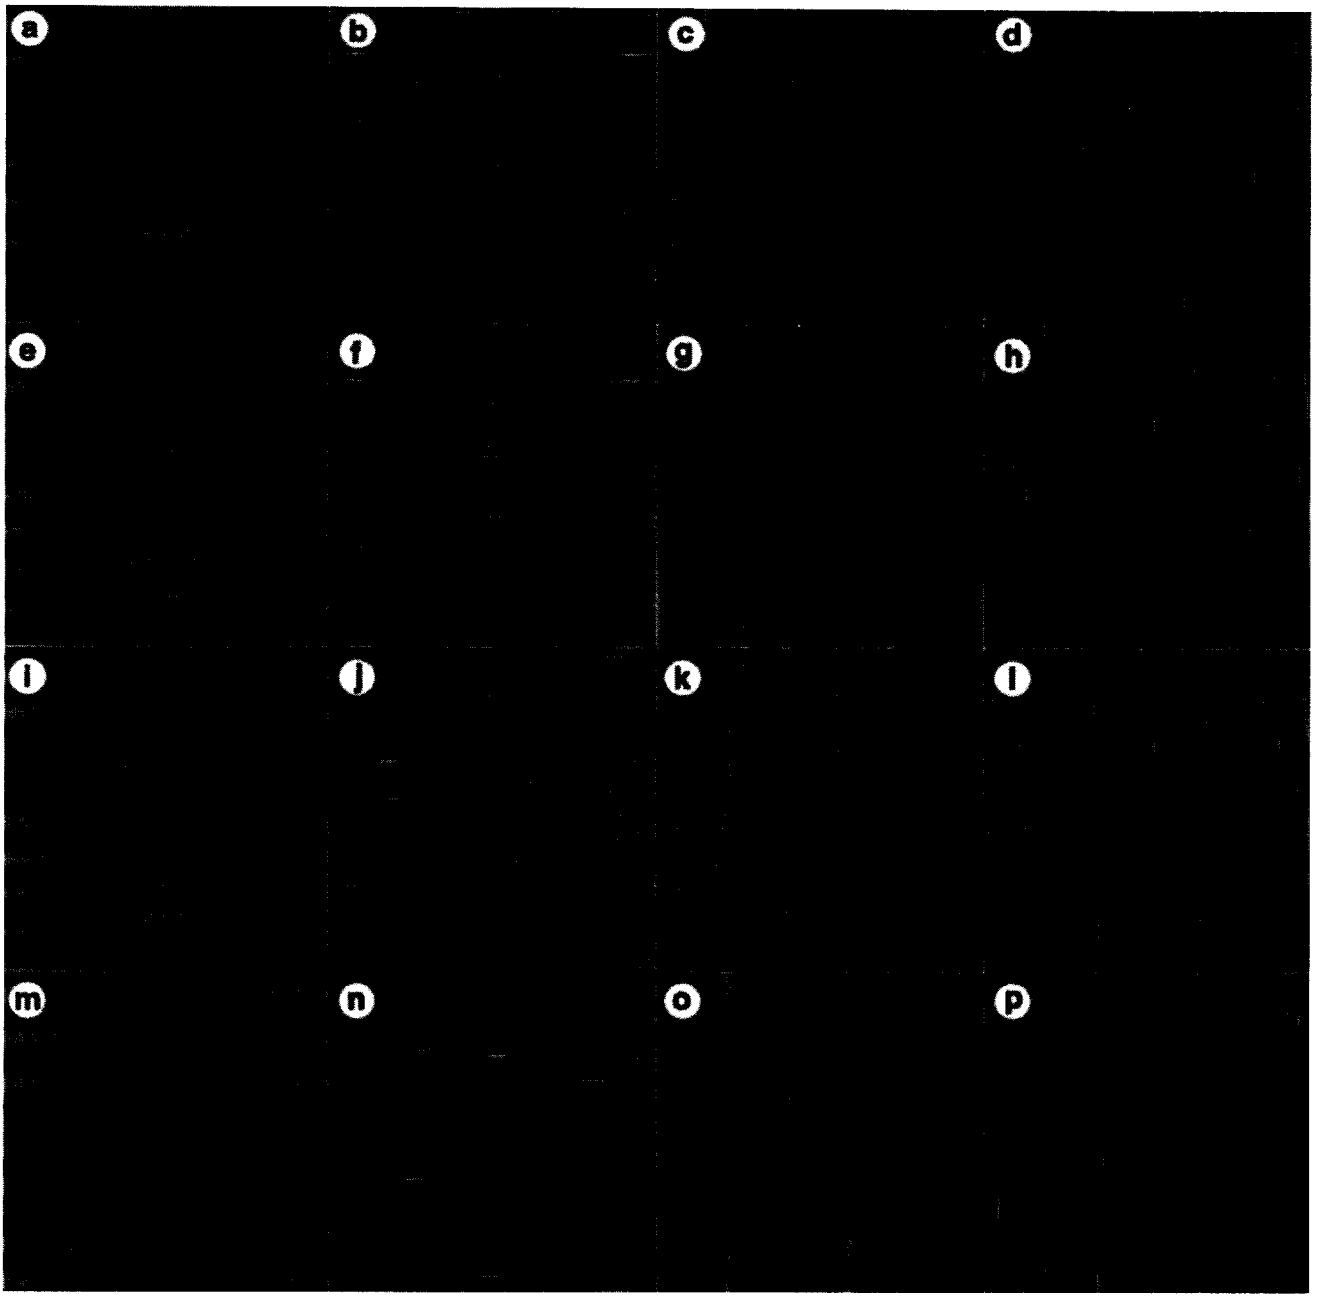

FIGURE 4. Examples of the stimuli used in Experiment 2. Patterns are limited to scales in a band 1 octave wide around (columns 1 and 3) 8 cycles per image and (columns 2 and 4) 32 cycles per image. From top to bottom, patterns have been phase-randomized to varying degrees: (a–d) 0 deg; (e–h) 60 deg; (i–l) 120 deg; (m–p) 360 deg.

cally filtered images (shown by the short-dashed line). This is interesting because of the aforementioned reduction in the sample of phase values that anisotropically filtered patterns make available. The contrary finding, that we are as good with orientationally band-limited stimuli, suggests that the operation of orientationally band-limited mechanisms effectively limits performance on this task. Furthermore, in the isotropically filtered condition, vertical information must be being actively "switched out" or subjects would be unable to perform as well as they do.

Notice that the magnitude of the advantage for horizontally filtered textures increases with spatial frequency. At high spatial frequencies subjects' performance with vertically filtered patterns with no added phase noise is not at ceiling, and collapses much faster than for other orientations.

In summary, this experiment has demonstrated an unexpected advantage for horizontal over vertical information in the detection of vertical bilateral symmetry. This is directly contrary to the predictions from models using the phase identity of spatial filters to measure local symmetry (Osorio, 1996), but is consistent with a model measuring feature co-alignment in the output of filters (Dakin & Watt, 1994). It is also particularly problematic for models based on cross-correlation, which is an inherently isotropic operation.

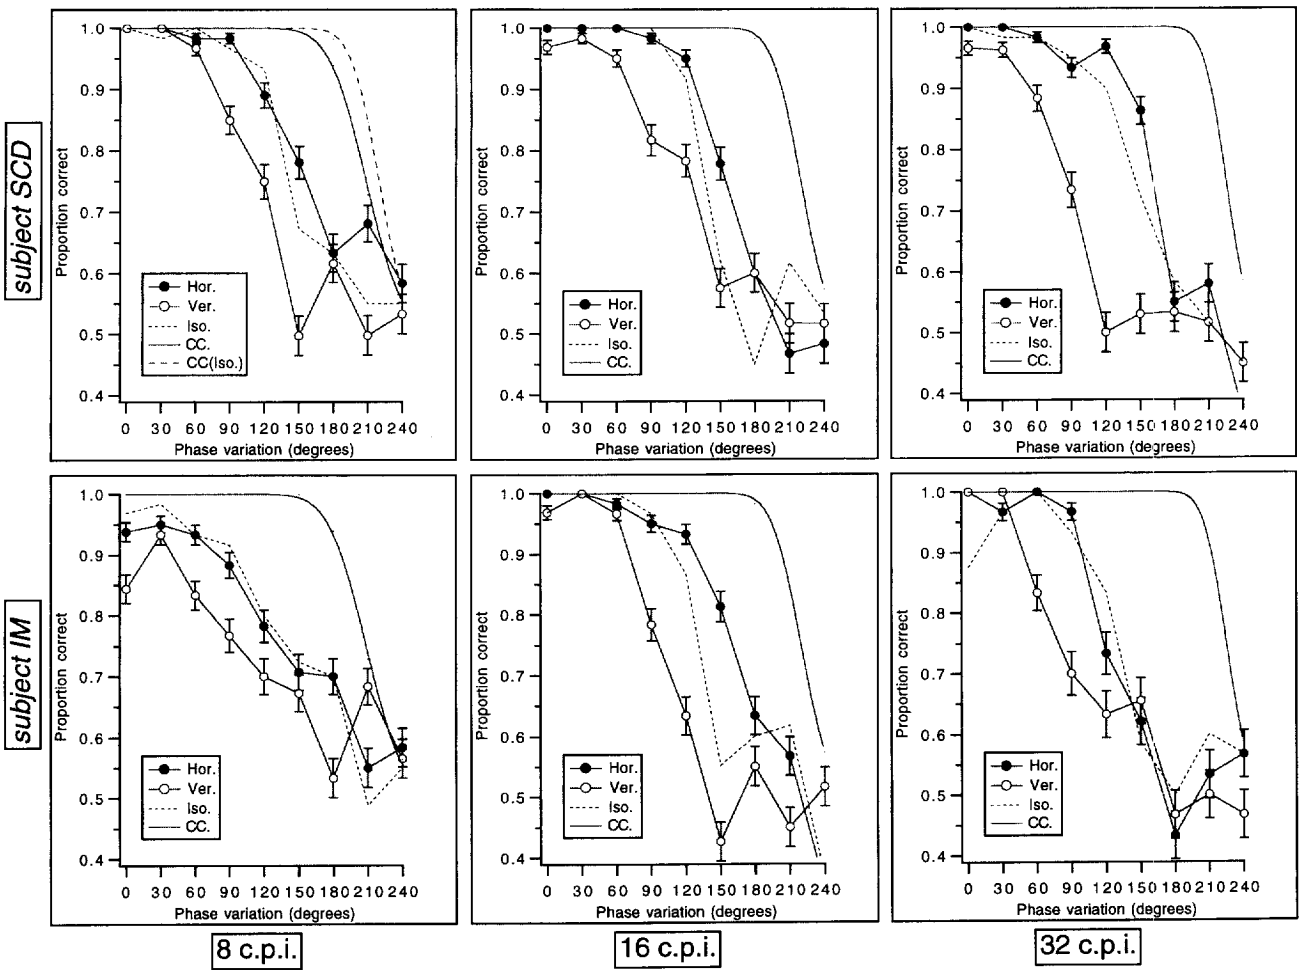

FIGURE 5. Effect of phase disruption on symmetry discrimination for textures filtered at three spatial scales, and retaining horizontal (solid symbols) or vertical information (open symbols). Subjects can consistently stand more phase disruption of horizontally than vertically filtered textures. For reference, the line composed of short dashes shows the performance of the subject with an isotropically filtered texture at the same spatial scale, and the solid line shows the prediction of a cross-correlator. The long-dashed line in (a) shows the prediction of a cross-correlator operating on isotropic patterns at 8 c.p.i.

EXPERIMENT 3: THE EFFECT OF AXIS ORIENTATION ON SYMMETRY DISCRIMINATION

The precision with which we can perform many tasks is dependent on stimulus orientation, and in particular there appears to be an advantage when stimuli are presented around horizontal and vertical, compared with oblique orientations. Orientation discrimination, for example, is better for horizontal and vertical stimuli (Caelli *et al.*, 1983; Heeley & Buchanan-Smith, 1990; Heeley & Timney, 1988; Orban *et al.*, 1984; Regan & Price, 1986) and furthermore, there is a small advantage for horizontal over vertical information (Heeley & Buchanan-Smith, 1990). If the advantage shown in the last experiment is specific to horizontal information, it might be attributable to the properties of the underlying channel. Alternatively, if it were to generalize to orientations orthogonal to the axis of symmetry, it is more likely to be related to the functional requirements of symmetry detection. In order to address this issue, we repeated Experiment 2 using patterns containing symmetry around a horizontal axis.

A number of experimental studies have investigated the effect of axis orientation on the perception of bilateral symmetry. On the whole, results indicate an advantage for vertical bilateral symmetry in terms of both speed (Corballis *et al.*, 1971; Palmer & Hemenway, 1978; Pashler, 1990; Royer, 1981), accuracy (Barlow & Reeves, 1979; Royer, 1981; Wagemans *et al.*, 1992; Wenderoth, 1994, 1996a), and contribution to the percept of symmetry in patterns with multiple axes (Fisher & Fracasso, 1987; Rock & Leaman, 1963). However, contrary findings (Fisher & Bornstein, 1982; Jenkins, 1983) and the importance of subjects' prior knowledge of axis orientation (Wenderoth, 1994) argue against the cause of this bias being at the neural level, but possibly as a consequence of visual attention.

Methods

The method used was identical to Experiments 1 and 2, except that symmetrical stimuli were generated using a reflection around a horizontal axis.

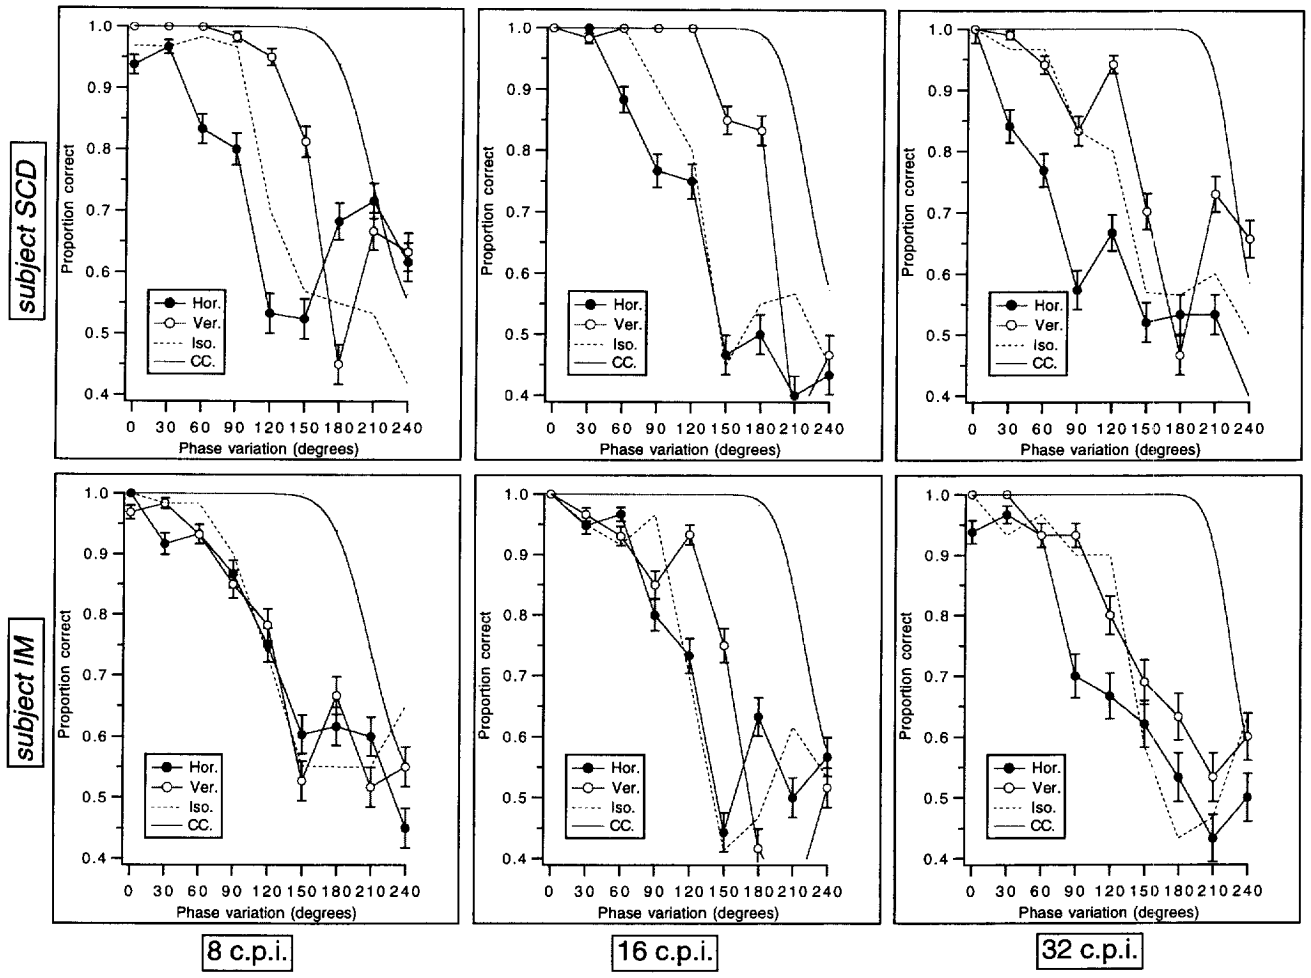

FIGURE 6. Discrimination of horizontally bilateral symmetry at three spatial scales, and retaining horizontal or vertical information. Notice that subjects now show greater resistance to phase disruption with vertically rather than horizontally filtered textures. This suggests the orientational advantage described is relative to the orientation of the axis of symmetry.

Results

The results shown in Fig. 6 indicate that in this condition, subjects could stand more phase disruption with vertically than horizontally filtered patterns. This confirms the hypothesis that the advantage shown in Experiment 2 is not solely for horizontal information but is for information orthogonal to the axis of symmetry. To highlight this trend in the data, Fig. 7(b) collapses data across spatial scale and subjects. The filled and open symbols show data from conditions where filter orientation was, respectively, orthogonal or parallel to the axis of symmetry. The advantage for the orthogonal case is clear.

Figure 6 also illustrates that for HBS the advantage for the orientation orthogonal to the axis (henceforth, referred to as the orthogonal orientation) extends not only to orientations parallel to the axis (henceforth, referred to as the parallel orientation), but in some conditions at least, over isotropically filtered textures. For one subject, performance is as poor on isotropic as on horizontal textures, while for the other performance falls somewhere between the horizontal and vertical conditions. Again, this is surprising given that the isotropic textures contain information at all orientations. Figure 7(a) shows that this is due to generally poorer performance with isotropically filtered textures for HBS than VBS.

Figure 7(b) shows that when horizontal information is removed from a pattern containing HBS (the vertically filtered case) performance improves and approaches data from the horizontally filtered VBS condition. This shows that poorer performance in either the vertically or isotropically filtered HBS conditions cannot be attributable to any absolute differences in the output of channels at different orientation. Instead, it suggests that when horizontal information is present, whether in an HBS or a VBS pattern, it cannot be ignored and will mask information from other orientations. Thus, there seems to be an inherent bias towards the use of horizontal information in the processing of symmetry.

This is entirely consistent with, and provides an explanation for, reported effects of axis orientation. Previous experiments have typically used broadband stimuli composed of dots which contain equal information at both horizontal and vertical. The general deficit for horizontal bilateral symmetry, confirmed in this experiment for isotropically filtered patterns, appears to be due

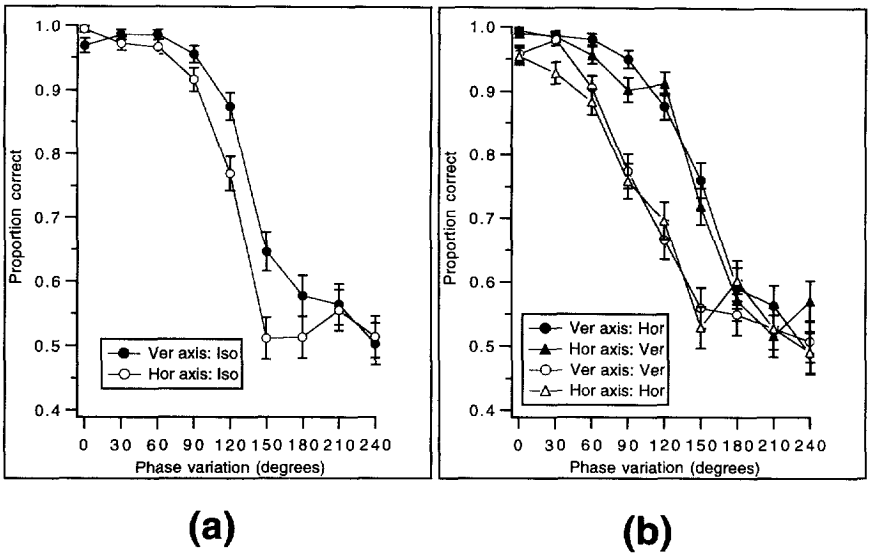

FIGURE 7. Summary of differences between VBS and HBS conditions. Data are averaged across scale and subjects. The effect of axis orientation is shown for (a) isotropically; and (b) horizontally and vertically filtered textures. (a) There is an advantage for VBS over HBS for isotropically filtered patterns. (b) The orientation of filtering in relation to the axis of symmetry determines performance. Performance is better with textures filtered orthogonal to the axis of symmetry, and worse for textures filtered parallel to the axis of symmetry, regardless of the axis of symmetry.

to an inherent bias towards using the output of the horizontal filters. This may arise as a consequence of a preponderance of HBS over other symmetries in the natural environment. The attentional and learning components of the effect (Wenderoth, 1994) suggest that given either enough time, or certain visual diets, observers may be able to ignore the output of horizontal mechanisms and thereby eradicate the effect.

MODELING OF SYMMETRY DISCRIMINATION

The experiments described are mutually consistent with respect to the influence of different oriented channels in the percept of symmetry. What has yet to be considered is *why* information at the orientation orthogonal to the axis of symmetry is more reliable than the orientation parallel to it. A model using an isotropic mechanism such as a cross-correlation would obviously not predict these orientational effects. We now describe two models of symmetry discrimination, using co-alignment in the output of oriented filters, that do.

Dakin & Watt (1994) propose the use of horizontal, linear filtering which, while successful at predicting performance in a variety of tasks using broadband stimuli, will not derive useful structure from vertically filtered VBS patterns. We therefore consider two variants of the model. The first, referred to as the "quasi-linear model", consists of both horizontal and vertical filtering mechanisms, followed by feature alignment calculation. The second, the "non-linear model", incorporates an early half-wave rectifying non-linearity prior to filtering. These two models are illustrated in Fig. 8. We first explain the operation of the horizontal filtering component of the quasi-linear model (the left-hand stream in Fig. 8), and then describe the two ways it was adapted to deal with information at other orientations.

Images are first convolved with an elongated DoG filter, composed of a DoG in the y-direction multiplied by a gaussian in the x-direction:

$$f(x,y,\sigma) = \left(e^{-y^2} - \frac{1}{2.23}e^{-y^2/2(2.23\sigma)^2}\right)e^{-x^2/2(3\sigma)^2} \quad (1)$$

where σ is the space constant of the filter, and the ratio of positive and negative parts are as derived by Wilson and colleagues (Phillips & Wilson, 1983; Wilson & Gelb, 1984). The exact form of this filter is not critical; any oriented, band-pass mechanism would suffice.

The filtered image is thresholded, by removing gray levels within one standard deviation of the mean, and the resultant "blobs" converted into a symbolic representation using Watt's image description scheme (Watt, 1991). Specifically a description of blob number i was of the form:

$$(cx_i, cy_i, \mu_i, \lambda_i, \theta_i) \quad (2)$$

where (cx_i,cy_i) is the centroid, μ_i the mass, λ_i the length and θ_i the orientation of the blob. Such a description is reminiscent of Marr's primal sketch (Marr, 1976, 1982).

The final stage is the measurement of feature alignment. Consider measuring the alignment of N_x blobs which are intersected by column x and which have a total mass M. This may be done using:

$$A(x) = \frac{1}{M}\sum_{i=1}^{N_x}\mu_i \exp\left[-\frac{(x-cx_i)^2}{2\lambda_i^2}\right] \quad (3)$$

which inversely weights the deviation of x from the blob's x-centroid (cx_i), but directly weights by the length and mass of the blob (λ_i and μ_i, respectively).

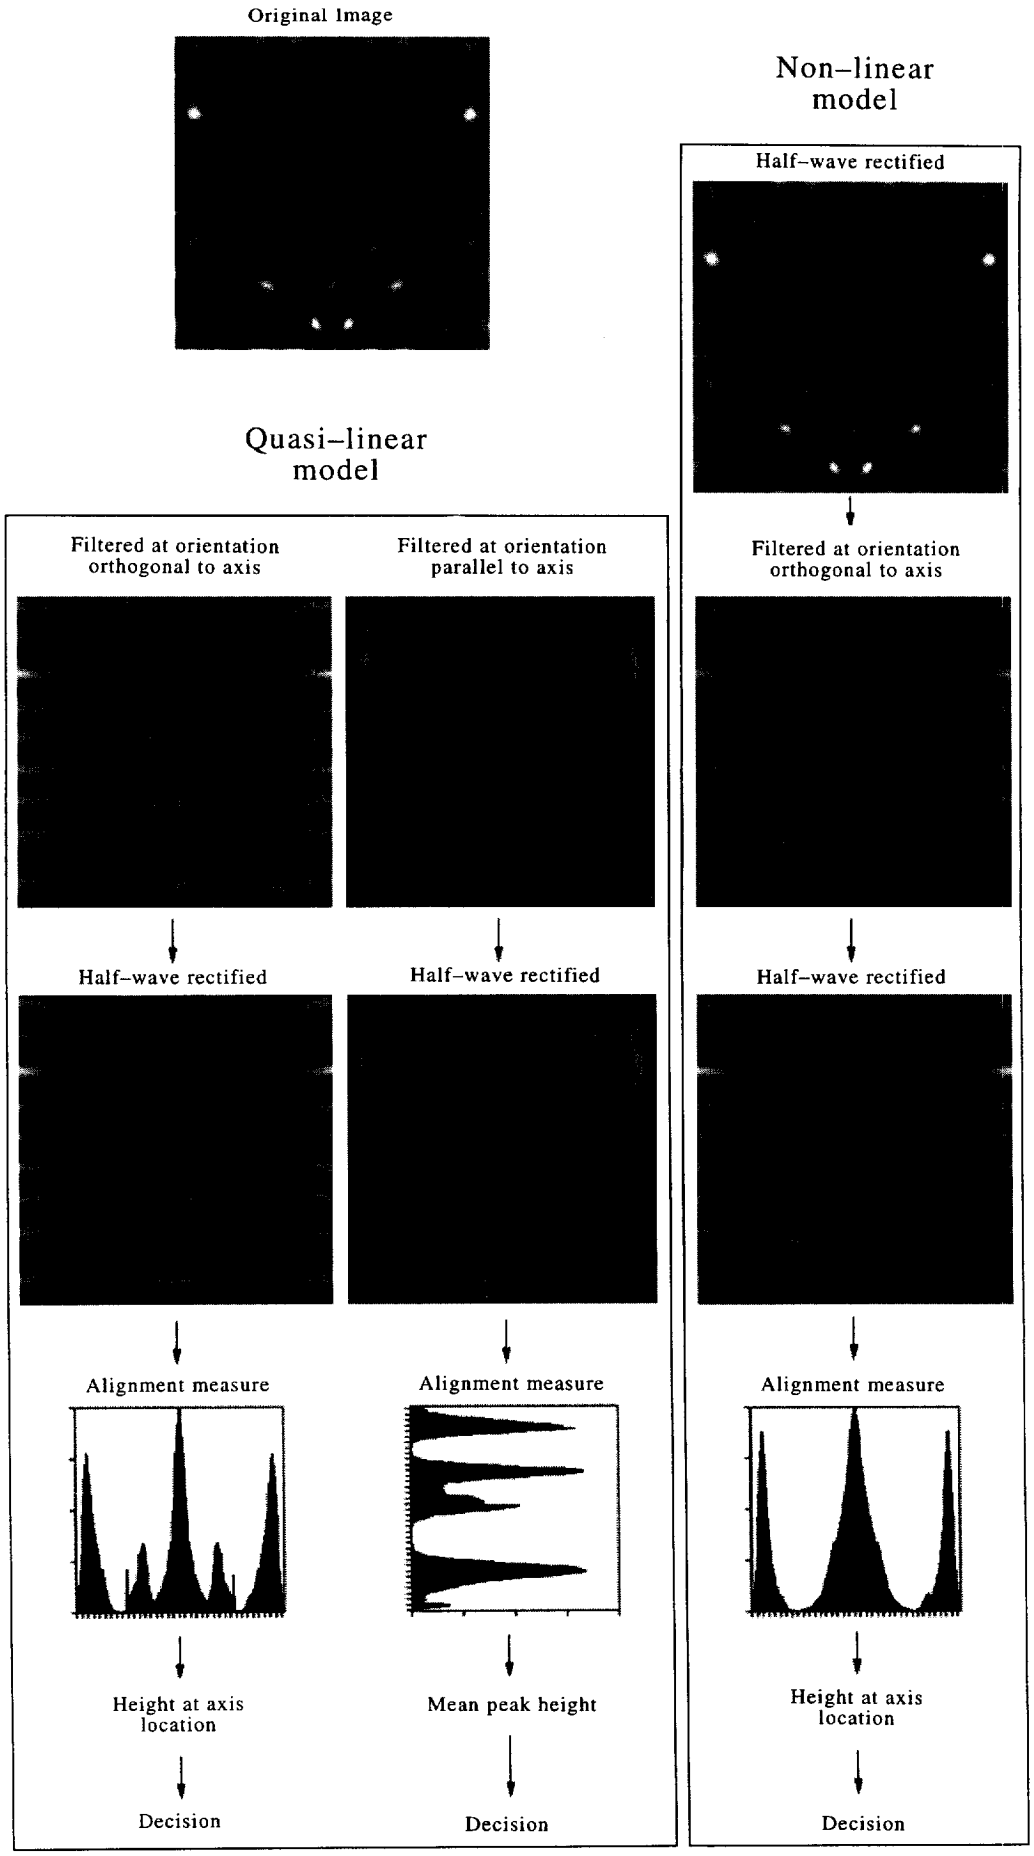

FIGURE 8. *Caption on facing page.*

TABLE 1. Chi-squares for the goodness-of-fit of the quasi-linear and non-linear models to data from all conditions

Model	Subject	Axis	I8	H8	V8	I16	H16	V16	I32	H32	V32
QL	SCD	V	0.023	0.025	0.596	0.225	0.114	0.155	0.048	0.143	0.155
QL	IM	V	0.031	0.069	0.408	0.246	0.126	0.213	0.189	0.270	0.186
QL	SCD	H	0.139	0.443	0.129	0.355	0.223	0.172	0.204	0.088	0.183
QL	IM	H	0.135	0.574	0.187	0.469	0.193	0.379	0.234	0.282	0.117
NL	SCD	V	0.017	0.032	0.116	0.169	0.030	0.148	0.073	0.143	0.045
NL	IM	V	0.029	0.015	0.056	0.198	0.058	0.069	0.191	0.291	0.116
NL	SCD	H	0.118	0.147	0.127	0.260	0.106	0.141	0.189	0.140	0.202
NL	IM	H	0.129	0.058	0.081	0.333	0.139	0.230	0.316	0.210	0.131

Abbreviations: QL, quasi-linear model; NL, non-linear model; V, vertical; H, horizontal. 'I8' refers to isotropically filtered patterns at 8 c.p.i., "V16" to vertically filtered patterns at 16 c.p.i., etc.

Given an alignment measure from each x-location, one has then to select a representative statistic to represent overall symmetry. If axis location is unknown then, because the horizontal filter output produces one large cluster around the axis of symmetry, a peak measure will suffice. However, in the experiments reported (as in most previous work) subjects had prior knowledge of the location of the axis. Consequently a better measure is simply to use the alignment measure at the x-location of the axis.

The basis of the scheme described so far is linear horizontal filtering which will not respond to vertically filtered textures. For this reason, two variants are considered. The second part of the quasi-linear model used vertical filtering, and similar post-filtering processing. The alignment measure is now measured in the y-direction:

$$A(y) = \frac{1}{M}\sum_{i=1}^{N_y} \mu_i \exp\left[-\frac{(y - cy_i)^2}{2\lambda_i^2}\right] \qquad (4)$$

where N_y blobs intersect row y, and all other parameters are identical to those specified for Eq. (3). There are two differences in the way vertical information, compared with horizontal information, is processed by the quasi-linear model. Firstly, because alignment is now measured along lines running perpendicular to the axis of symmetry, one would expect multiple feature clusters (Fig. 8, lowermost middle plot). A peak measure does not characterize this distribution well, and for that reason we chose to use the mean peak-height instead. The second difference is in the way models are cued as to where the axis of symmetry is. For horizontal filtering, the subjects' lack of uncertainty about axis location is simply encoded by measuring alignment at the correct x-location. For vertical filtering this knowledge is incorporated by a windowing parameter, which limits the area around the axis about which information is used. There is a trade-off, however. If the window is too narrow, insufficient information will be present to establish if the clusters extracted are due to symmetry or chance alignments in the noise pattern. Consequently the vertical component quasi-linear model was run with a range of window sizes.

The non-linear model uses a simpler way of measuring alignment from vertically filtered textures. Application of a non-linearity prior to filtering will introduce information at other orientations. The non-linear model is therefore identical to the horizontal component of the quasi-linear model except that DoG filtering was preceded by half-wave rectification:

$$R(x,y) = \begin{cases} I(x,y) & \text{if } |I(x,y)| > S \\ 0 & \text{otherwise} \end{cases} \qquad (5)$$

where $I(x,y)$ is the image intensity at (x,y), S is the standard deviation of all gray levels, and $R(x,y)$ is the half-wave rectified result. This form of non-linearity is both computationally simple and biologically plausible, and has been cited as a plausible non-linearity in models of "second-order" texture perception (e.g. Sutter *et al.*, 1995).

Simulation methods

Sixteen samples of stimulus and noise were generated at each stimulus level tested. Images generated at spatial frequency σ (the matched scale), were filtered with five mechanisms whose sensitivities peaked at 0.25σ, 0.5σ, σ, 2σ and 4σ. A set of symmetry measures from each model and at each spatial scale was generated and used to calculate the probability of correct discrimination of stimulus from noise at a particular level of phase disruption.

These raw probabilities were fit by a cumulative Gaussian function of the form:

$$P(x + \delta x) = \int_{-\infty}^{x+\delta x} \exp\left(\frac{-(a-\mu)^2}{2\sigma^2}\right) da \qquad (6)$$

FIGURE 8 (*opposite*). Two models for measuring visual symmetry. Both employ horizontal DoG filtering, followed by thresholding and calculation of the alignment at the axis location. However, the non-linear model half-wave rectifies input prior to filtering. The full quasi-linear model incorporates vertical filtering, followed by thresholding and measurement of the mean peak blob alignment across the image.

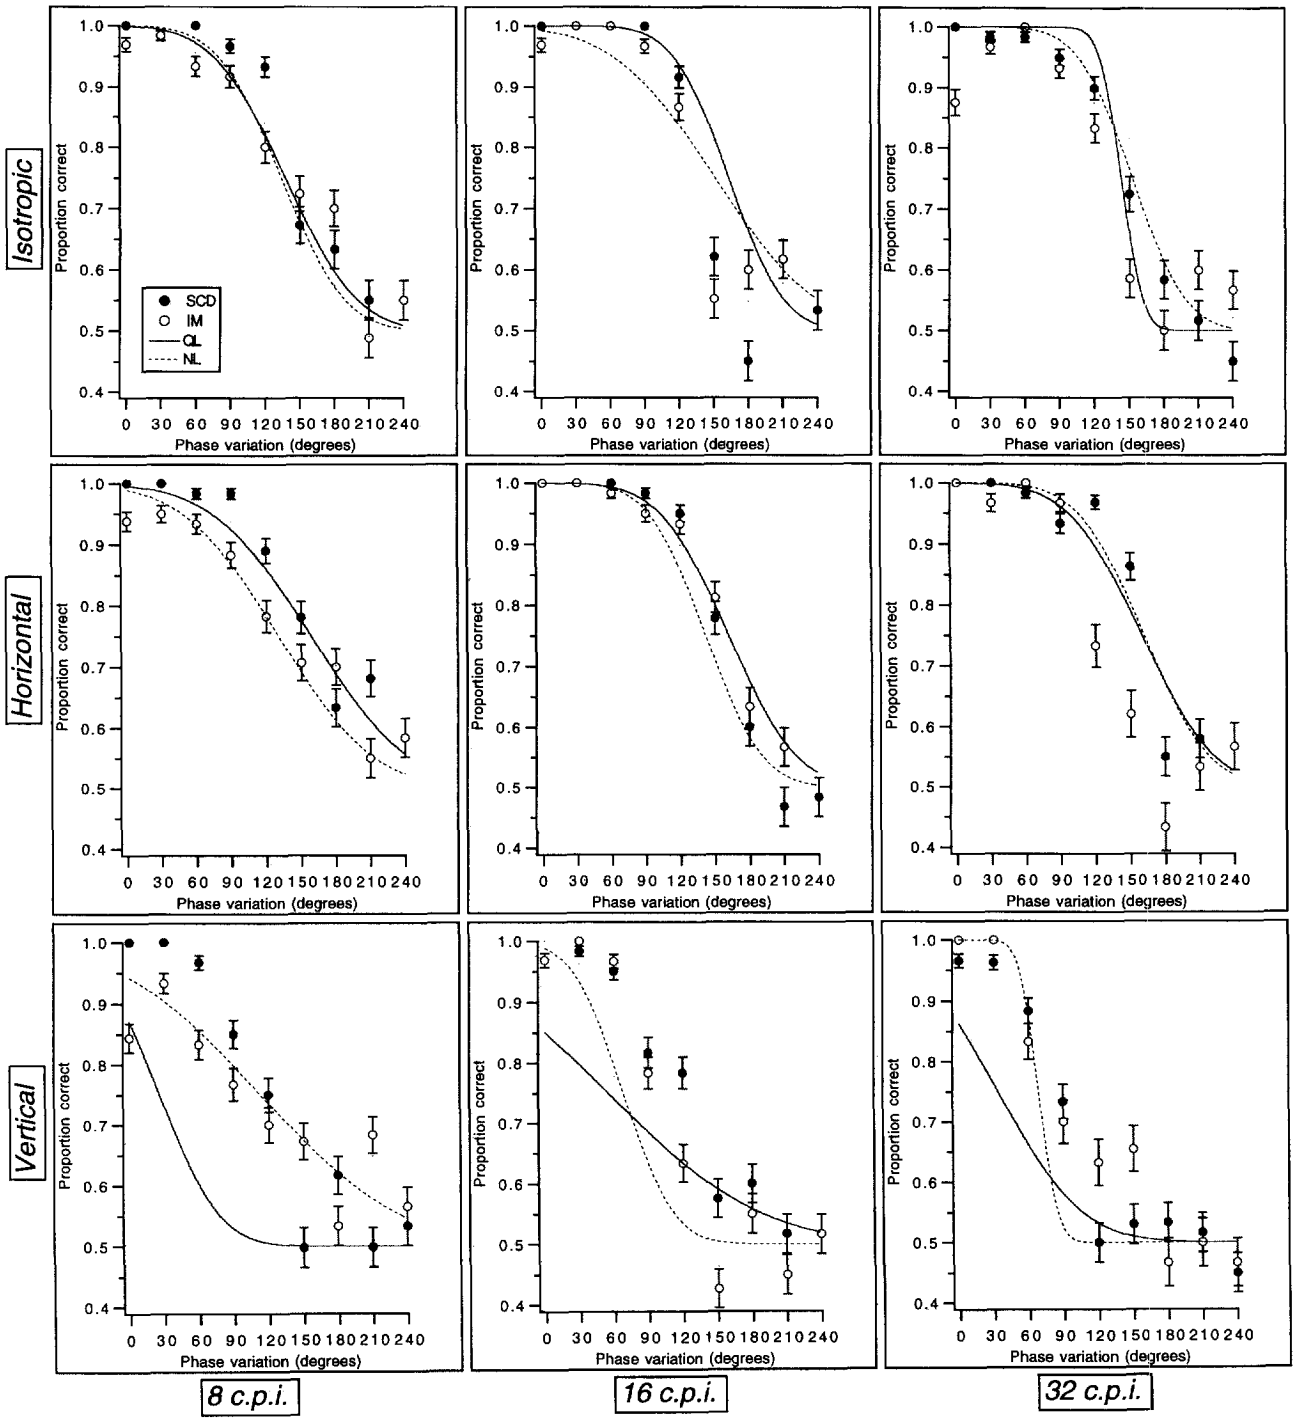

FIGURE 9. Predictions from the two models compared with psychophysical data from the vertical axis of symmetry condition. The two subjects' data are presented as the filled and open symbols; predictions from the quasi-linear (QL) and non-linear (NL) models as the solid and broken lines, respectively. Both models match human data well for isotropically and horizontally filtered textures. For the vertically filtered patterns, the non-linear model suffices at low spatial frequencies but, along with the quasi-linear model, is a poorer fit at high spatial frequencies.

Simulation results

Goodness of fit was assessed using a chi-square fit weighted by the standard error of subjects' data at each stimulus level. For the quasi-linear model the best fit was consistently provided by the filter whose sensitivity was matched to the spatial frequency at which the input texture was generated. For the non-linear model the best fit was achieved using a filter either at, or one octave below, this scale. Chi-square values for the fit of both models to data from all conditions are shown in Table 1.

Figure 9 shows predictions from the models, operating at their optimal spatial scales, for discrimination of VBS from noise. It is clear that models using horizontal filtering, either with or without an early non-linearity,

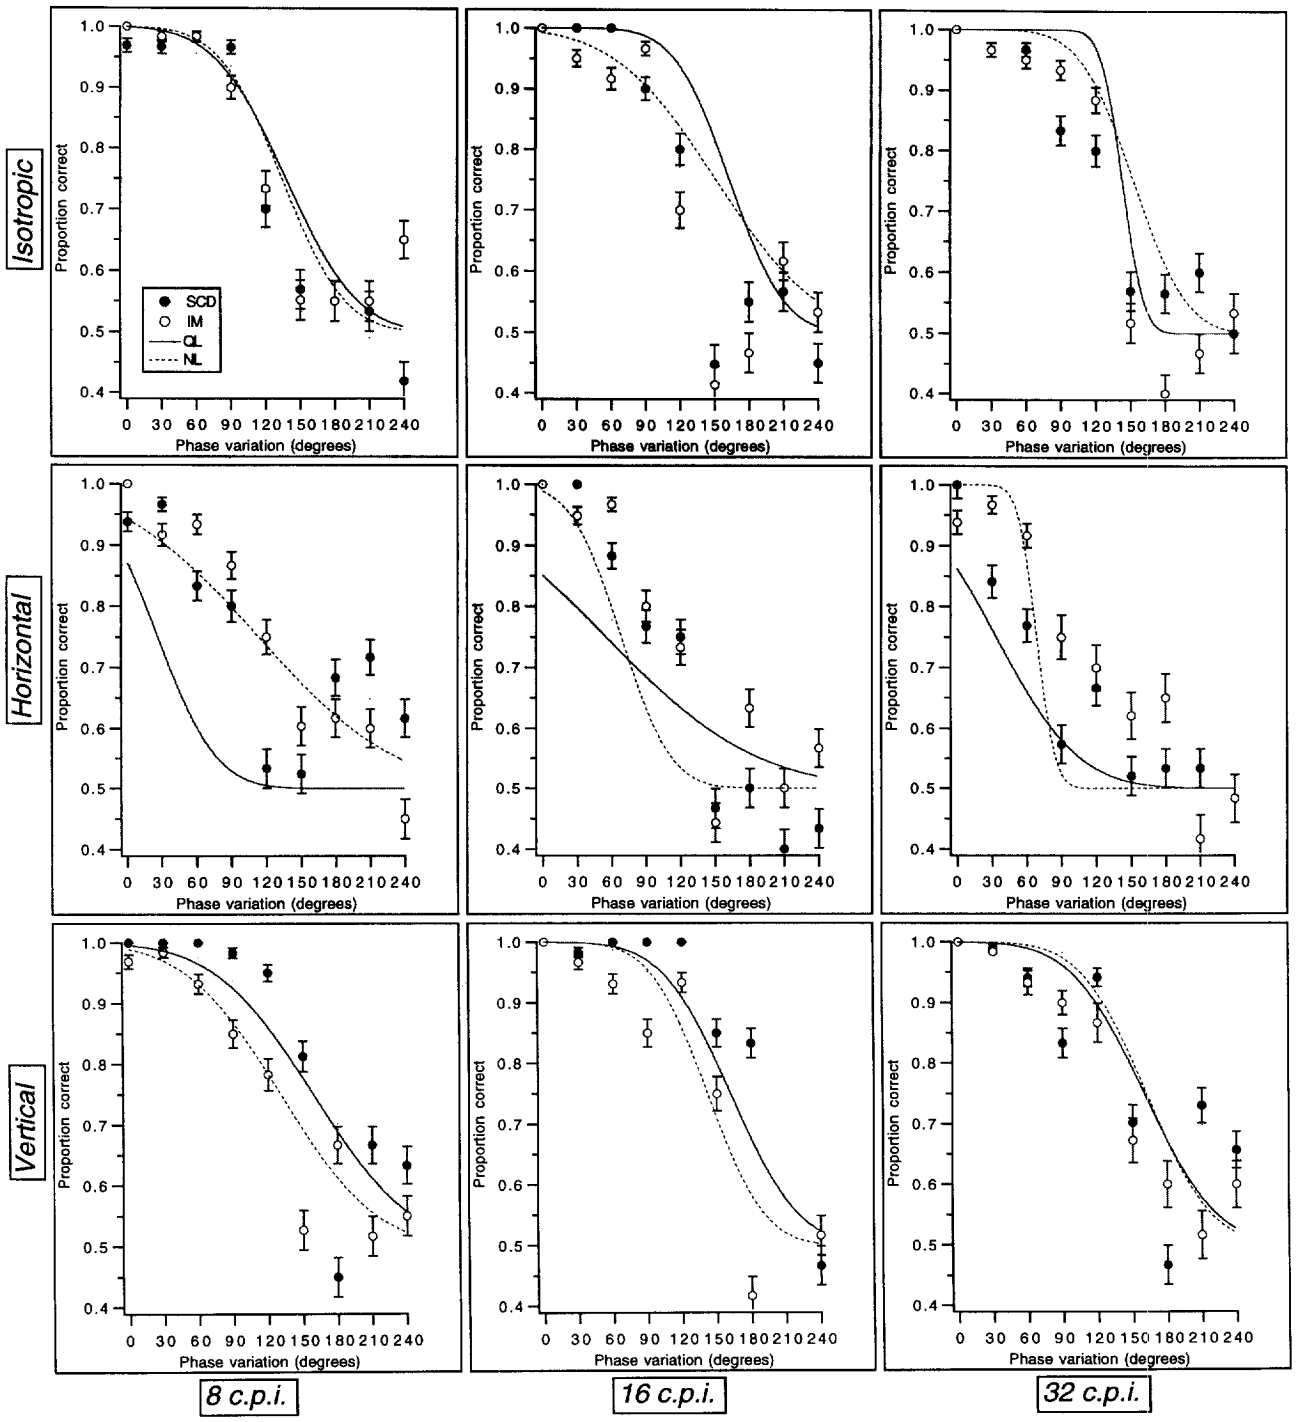

FIGURE 10. Comparison of the quasi-linear and non-linear models with human data from the horizontal axis of symmetry condition. Models match data well for isotropically and vertically filtered textures. Similarly to Fig. 9, the non-linear model provides an adequate fit to data when textures are filtered *parallel* to the axis of symmetry (i.e., horizontally) at low spatial frequencies but, along with the quasi-linear model, is a poorer fit at progressively higher spatial frequencies.

predict performance well in the case of horizontally and isotropically filtered textures. On the other hand, the vertical component of the quasi-linear model, even when window size is optimized (at a width of 75% of the image width), fails to achieve the subjects' level of performance with vertically filtered textures. The non-linear model fares better at low spatial frequencies, but still fails to achieve human levels of performance for many levels of phase disruption.

Predictions of the models for HBS are shown in Fig. 10. For isotropically and vertically filtered patterns, both linear and non-linear models predict human data well. However, for horizontally filtered patterns (which, being filtered *parallel* to the axis of symmetry, correspond to the vertically filtered condition shown in Fig. 9) fits are

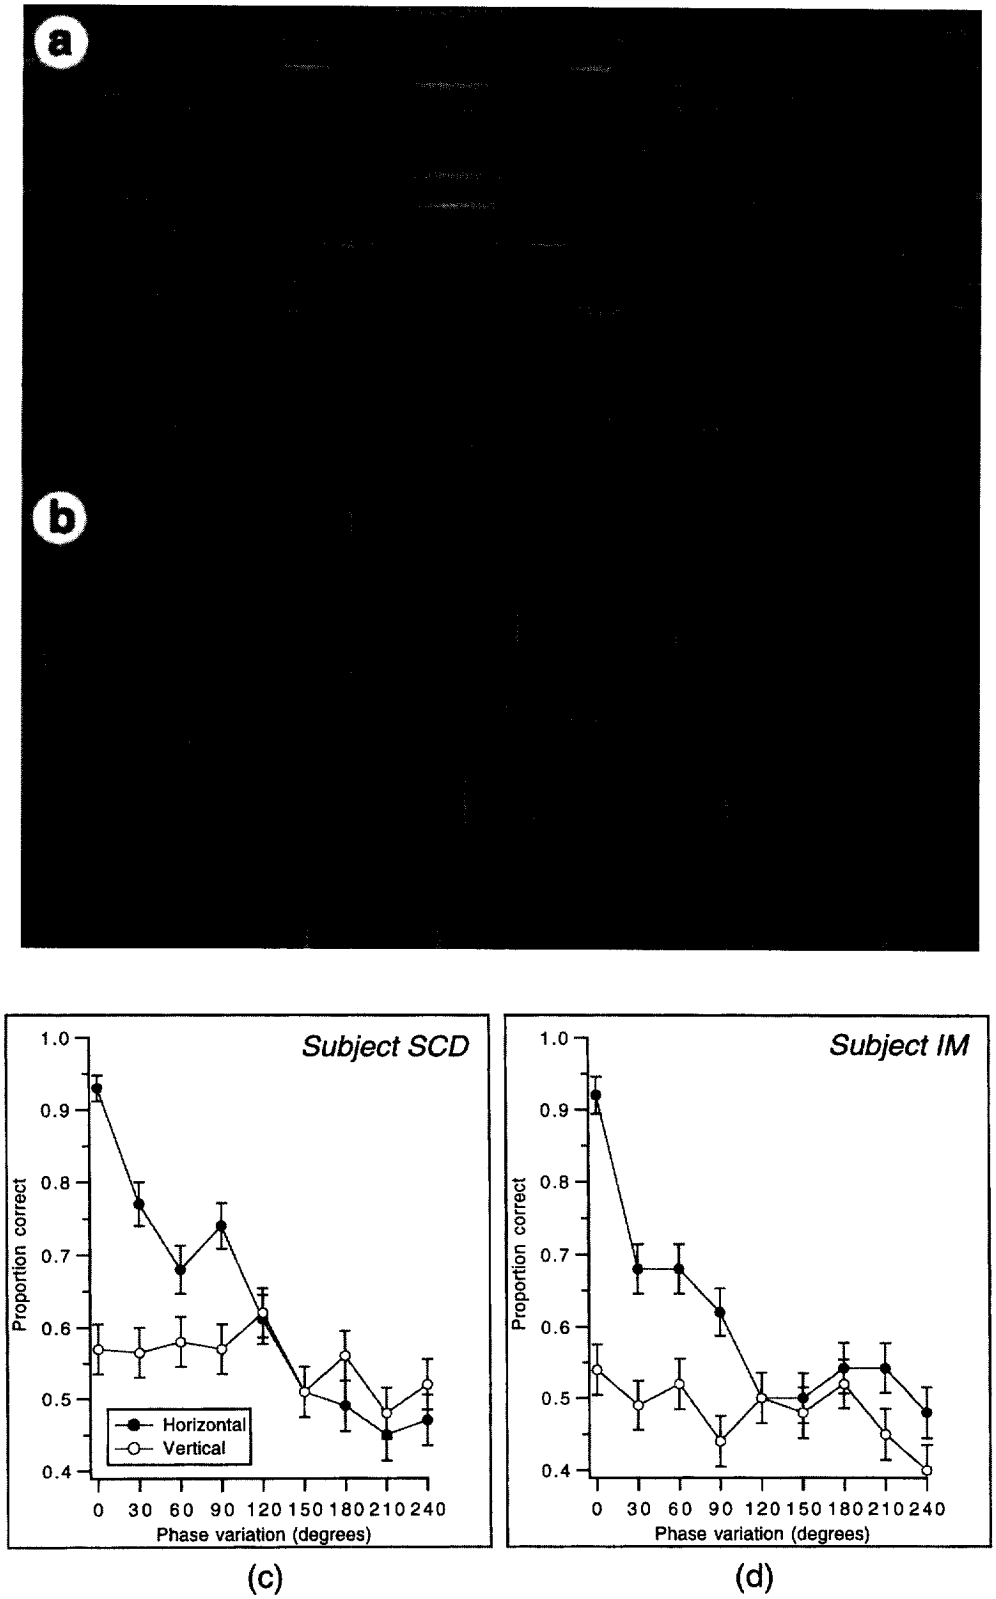

FIGURE 11. (a, b) Examples of the stimuli used to examine the effects of uncertainty about the location of the axis of symmetry. (c, d) Discrimination of symmetry from noise with axis-location uncertainty. Note that performance with vertically filtered patterns falls to chance, whereas for the horizontally filtered textures, discrimination is still possible in the presence of phase disruption. This suggests that visual attention is particularly important for the perception of symmetry in vertically filtered textures.

uniformly poorer. Again, at low spatial frequencies the non-linear model is adequate but fails to reach human performance at finer spatial scales.

To summarize, a model using filters oriented orthogonal to the axis of symmetry, either with or without a preceding non-linearity, adequately accounts for human

performance with both isotropically filtered patterns, and patterns filtered orthogonal to the axis of symmetry. Neither the non-linear nor the quasi-linear model can account fully for human performance with patterns filtered parallel to the axis of symmetry, at medium to high spatial frequencies.

GENERAL DISCUSSION

To summarize, this paper has three main findings:

- We have demonstrated a new phenomenon in human visual processing of symmetry: that detection performance for symmetry in patterns that are limited to orientations around the axis of symmetry is "poorer" (i.e., more vulnerable to the intrusion of noise) than when orientation is either not limited, or is limited to orientations orthogonal to the axis.

- Generally poorer performance with horizontal bilateral symmetry is attributable to masking by the output of mechanisms sensitive to orientations parallel to the axis of symmetry.

- A model measuring feature co-alignment, operating on the output of filters oriented orthogonal to the axis of symmetry, accounts for discrimination performance in the presence of phase disruption, for isotropic patterns, and patterns containing information only at orientations orthogonal to the axis of symmetry.

That both the quasi-linear (QL) and non-linear (NL) models perform poorly with patterns filtered parallel to the axis of symmetry could be for a number of reasons. For the QL model it is possible that the windowing parameter fails to convey knowledge of the axis of location information effectively, although a range of window sizes was tested. Alternatively, the mean-peak height may simply not be a good metric for measuring the alignment statistics of these patterns. A number of alternative statistics were attempted (the mean height, the overall peak, etc.) but no improvement in the fits was shown. With respect to the NL model, it may be that some aspect of the pre-filtering non-linearity affected the outcome. We re-ran the simulations using half-wave rectification maintaining separate positive and negative components, and there was no significant effect on the predictions of the model. Although it is not possible to exhaustively check every possibility for each stage of the models, we conclude that it is unlikely that the poorer performance of the models with images filtered at parallel orientations could be corrected with a trivial manipulation of any one component.

A second alternative is that attentional factors are more important for the parallel orientation case than for the orthogonal orientation filtered case. It is possible that, because the output of pre-attentive processes are unreliable at parallel orientations, their output is supplemented by more detailed matching information (such as shape, size, and orientation of elements around the axis). This explanation would suggest that, for VBS, placing additional load on attentional processing should reduce discrimination of symmetry in vertically filtered patterns more than with horizontally filtered patterns.

We tested whether attentional factors would differentially affect perception of symmetry in horizontally and vertically filtered textures, by rendering subjects uncertain as to the location of the axis of symmetry. Symmetrical patterns were 512 by 256 pixel images containing a 256 pixel square image embedded in noise [Fig. 11(a) shows examples]. The offset of the embedded image was randomly determined, so that subjects knew only that the axis of symmetry must lie in a 256 pixel-wide region around the center. Noise textures were identical but contained no embedded image. All images were presented for 100 msec. All other experimental details were identical to those described in Experiment 2, except that only one spatial frequency (16 c.p.i.) was tested.

Results shown in Fig. 11 indicate that not having prior knowledge of the axis location is disastrous for perception of symmetry in vertically filtered patterns, and subjects performed consistently near to chance. For horizontally filtered patterns the drop-off in performance is steeper than for the equivalent condition in Experiment 2, but subjects appear to be quite capable of performing the task at low levels of phase disruption. We conclude that attentional factors are more critical in perception of symmetry in patterns filtered at parallel orientations. This is a likely explanation both for the failure of the models to account for performance in this condition, and for subjects' generally poorer performance with patterns filtered at this orientation.

REFERENCES

Barlow, H. B. & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research, 19*, 783–793.

Brady, N. & Field, D. (1995). What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns. *Vision Research, 35*, 739–756.

Caelli, T., Brettel, H., Rentschler, I. & Hilz, R. (1983). Discrimination thresholds in the two-dimensional frequency domain. *Vision Research, 23*, 129–133.

Carlin, P. (1996). On symmetry in visual perception. Department of Psychology, University of Stirling, Stirling, Scotland.

Corballis, M. C., Miller, A. & Morgan, M. J. (1971). The role of left-right orientation in interhemispheric matching of visual information. *Perception and Psychophysics, 10*, 385–388.

Dakin, S. C. & Watt, R. J. (1994). Detection of bilateral symmetry using spatial filters. Special Issue: The perception of symmetry: Part I. Theoretical aspects. *Spatial Vision, 8*, 393–413.

Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial-frequency, and orientation optimized by two dimensional cortical filters. *Journal of the Optical Society of America A, 2*, 1160–1169.

Fisher, C. B. & Bornstein, M. H. (1982). Identification of symmetry: effects of stimulus orientation and head position. *Perception and Psychophysics, 32*, 443–448.

Fisher, C. B. & Fracasso, M. P. (1987). The Goldmeier effect in adults and children: environmental, retinal, and phenomenal influences on judgments of visual symmetry. *Perception, 16*, 29–39.

Heeley, D. W. & Buchanan-Smith, H. M. (1990). Recognition of stimulus orientation. *Vision Research, 30*, 1429–1437.

Heeley, D. W. & Timney, B. (1988). Meridional anisotropies or orientation discrimination for sine wave gratings. *Vision Research, 28*, 337–344.

Jenkins, B. (1982). Redundancy in the perception of bilateral symmetry in dot textures. *Perception and Psychophysics, 32*, 171–177.

Jenkins, B. (1983). Component processes in the perception of bilaterally symmetric dot textures. *Perception and Psychophysics, 34*, 433–440.

Julesz, B. & Chang, J. (1979). Symmetry perception and spatial-frequency channels. *Perception, 8*, 711–718.

Koeppl, U. (1993). Local orientation versus local position as determinants of perceived symmetry. *Perception, 22 Supplement*, 111.

Locher, P. J. & Wagemans, J. (1993). Effects of element type and spatial grouping on symmetry detection. *Perception, 22*, 565–587.

Marr, D. (1976). Early processing of visual information. *Proceedings of the Royal Society of London, B, 275*, 483–534.

Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.

Orban, G. A., Vandenbussche, E. & Vogels, R. (1984). Human orientation discrimination tested with long stimuli. *Vision Research, 24*

Osorio, D. (1996). Symmetry detection by categorization of spatial phase, a model. *Proceedings of the Royal Society of London, B, 263*, 105–110.

Palmer, S. E. & Hemenway, K. (1978). Orientation and symmetry: effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance, 4*, 691–702.

Parish, D. H. & Sperling, G. (1991). Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination. *Vision Research, 31*, 1399–1415.

Pashler, H. (1990). Coordinate frame for symmetry detection and object recognition. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 150–163.

Phillips, G. & Wilson, H. (1983). Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America, 62*, 226–232.

Pintsov, D. A. (1989). Invariant pattern recognition, symmetry, and Radon transforms. *Journal of the Optical Society of America, 6*, 1544–1554.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992). *Numerical recipes in C: the art of scientific computing*. Cambridge, U.K.: Cambridge University Press.

Regan, D. & Price, P. (1986). Periodicity of orientation discrimination and the unconfounding of visual information. *Vision Research, 26*

Rock, I. & Leaman, R. (1963). An experimental analysis of visual symmetry. *Acta Psychologica, 21*, 171–183.

Royer, F. L. (1981). Detection of symmetry. *Journal of Experimental Psychology: Human Perception and Performance, 7*, 1186–1210.

Sutter, A., Sperling, G. & Chubb, C. (1995). Measuring the spatial frequency selectivity of second-order texture mechanisms. *Vision Research, 35*, 915–924.

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*, 79–100.

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

Wagemans, J. (1995). Detection of visual symmetries. Special Issue: The perception of symmetry: II. Empirical aspects. *Spatial Vision, 9*, 9–32.

Wagemans, J., Van Gool, L. & d'Ydewalle, G. (1992). Orientational effects and component processes in symmetry detection. *Quarterly Journal of Experimental Psychology, 44A*, 475–508.

Wagemans, J., Van Gool, L., Swinnen, V. & Van Horebeek, J. (1993). Higher-order structure in regularity detection. *Vision Research, 33*, 1067–1088.

Watt, R. J. (1991). *Understanding vision*. London: Academic Press.

Wenderoth, P. (1994). The salience of vertical symmetry. *Perception, 23*, 221–236.

Wenderoth, P. (1996a) The effects of dot pattern parameters and constraints on the relative salience of vertical bilateral symmetry. *Vision Research, 36*, 2311–2320.

Wenderoth, P. (1996b) The effects of the contrast polarity of dot-pair partners on the detection of bilateral symmetry. *Perception, 25*, 757–772.

Wilson, H. & Gelb, D. (1984). Modified line-element theory for spatial-frequency and width discrimination. *Journal of the Optical Society of America, A1*, 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*, 455–467.

Zhang, L. (1991). Symmetry perception in human vision. In *Psychology*. University of Trieste, Trieste.

Zucker, S. (1982). Early orientation selection and grouping: evidence for type I and type II processes. McGill University, Montreal, Quebec.

Acknowledgements—We thank Roger Watt for valuable discussions of this work and particularly for suggesting the use of a non-linearity prior to filtering. Thanks also to Isabelle Mareschal for being a patient subject. This research was funded by MRC Grant #MT 108-18.