



Perception of 3D surface orientation from skew symmetry

Jeffrey A. Saunders *, David C. Knill

Center for Visual Science, University of Rochester, Meliora Hall 274, Rochester, NY 14627, USA

Received 29 June 2001

Abstract

In this paper, we investigate how symmetry can be used to perceive 3D surface orientation. When a symmetric planar object is viewed from an angle, the projected contour has skew symmetry, which provides partial information about the 3D orientation of the object. For a given skew symmetry, this information can be characterized by a constraint curve of possible slant/tilt combinations that are consistent with a mirror-symmetric interpretation. These constraint curves move around when an object is rotated within a plane, and depend on what we will term the *spin* of the object: the angle between its axis of symmetry and the direction of tilt. To test the influence of symmetry constraint curves, we presented subjects with stereo images of symmetric objects that varied in spin, and had them perform an orientation-matching task. We found that the judgments showed biases that depended on the spin of the objects. Since other sources of information depend only on slant and tilt, not on spin, the biases imply that skew symmetry contributed to subjects' judgments. In a second experiment, we introduced conflicts between stereo and symmetry cues, and found that the spin-dependent biases can be modulated by selectively changing stereo slant. We propose an explanation of these results involving the optimal integration of stereo and skew symmetry, and present a Bayesian model that can account for the pattern of biases. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: 3D surface orientation; Skew symmetry; Symmetric objects

1. Introduction

Symmetric objects form a large subset of the objects that we encounter in our environment, both natural and artificial. As early as the 19th century, Mach noted the importance of symmetry to visual perception (Mach, 1897). Since then, symmetry has been established as a central organizing principle for perceptual object representations in humans (Kohler, 1938; Attneave, 1954; Perkins, 1972, 1976; Leyton, 1992; Wagemans, 1995): observers are very efficient at detecting symmetry in random patterns (Barlow & Reeves, 1979), they are biased to interpret slightly asymmetric figures as symmetric (King, Meyer, Tangney, & Biederman, 1976), and they show improved recognition performance (e.g. viewpoint independence) for symmetric objects (Wagemans, 1992, 1993; Vetter & Poggio, 1994; Liu, Knill, & Kersten, 1995).

Fig. 1 illustrates another role that symmetry can play in perception: to support slant perception. Symmetric, planar figures project to approximately skew symmetric image contours when slanted away from the image plane (the approximation is exact for orthography). The amount and direction of skew provide potentially powerful cues to the three-dimensional orientation of the figure. A number of psychophysical experiments have tested whether or not the human visual system uses the information provided by skew symmetry for slant perception. Wagemans (1993) and McBeath, Schiano, and Tversky (1997) have independently shown that subjects are much more likely to perceive skew symmetric figures than asymmetric figures to be slanted out of the image plane in depth. While these results clearly demonstrate the perceptual salience of skew symmetry, they do not indicate *how* the visual system uses skew information.

The experiments described here were designed to quantify the influence of skew symmetry on the perception of surface orientation. The results of the experiments motivated us to develop a Bayesian model that

* Corresponding author.

E-mail address: saunders@cvs.rochester.edu (J.A. Saunders).

characterizes how the visual system integrates skew information with other cues to generate perceptual estimates of surface orientation. The model accurately accounts for the systematic biases that we find in subjects' judgments of surface orientation for stereoscopically presented images of symmetric figures. Moreover, it predicts a particular pattern of non-linear interactions between figural information and information from other cues such as stereo, when the cues are placed in conflict. Results of an experiment in which we systematically manipulated the degree of conflict between stereo and skew symmetry information bear out the predictions of the model.

1.1. Information content of skew

In this section, we describe the information about surface orientation provided by skew symmetry. Perspective projection maps bilaterally symmetric, planar figures into approximately skew-symmetric contours in the image. We will therefore refer to the projected symmetry axes of a figure as the *skew axes* of its contour.¹ The orientations of the two skew axes are a function of the orientation of the plane in which the symmetric figure lies (its slant and tilt) and the orientation of the figure within that plane (its spin). These three components of a three-dimensional orientation are illustrated in Fig. 2. Because the orientations of the two skew axes are functions of three unknowns, the axes do not uniquely determine the three-dimensional

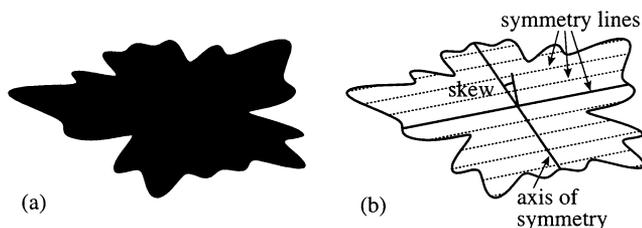


Fig. 1. (a) A skew symmetric contour can be seen as the image of a symmetric object viewed from an angle. (b) The skew symmetry is defined by an axis of symmetry and symmetry lines that implicitly connect matching parts of the contour. We will refer to these as the *skew axes* of the figure. The *skew* is the angular deviation from a right angle intersection. A skew symmetric contour approximates a perspective view of a symmetric figure.

¹ Under orthographic projection, the projected contour of a symmetric figure is exactly skew symmetric, and the axes of the skew symmetry correspond to the projected axes of symmetry of the 3D figure. This approximates perspective projection for small viewing angles and low to intermediate surface slants. Under perspective projection, the same mathematical constraint describes the relationship between the projected axes of symmetry at any point along the center line of the figure and the local slant and tilt of the surface relative to the line of sight to that point. We will take up a discussion of perspective later in Sections 2.3 and 3.4 (see also Appendix B).

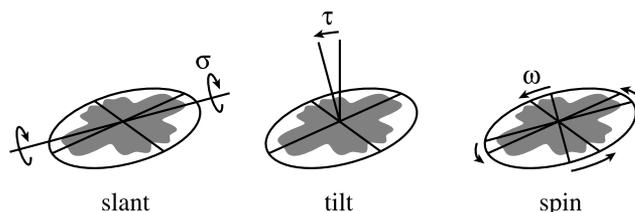


Fig. 2. Three angles that define a 3D orientation: slant (σ), tilt (τ) and spin (ω). Slant and tilt specify the surface orientation, and spin describes the orientation within the plane. Slant (σ) is the angle between the surface normal and the line of sight (slant of zero is frontal). Tilt (τ) is the orientation in the image plane of the surface gradient relative to vertical (tilt of zero corresponds to rotation around horizontal axis). Spin (ω) is defined as the angle between the symmetry axis and the back-projected tilt vector (spin of zero corresponds to a figure aligned with its tilt direction).

orientation of the figure. Rather, they define a constraint curve in two-dimensional slant/tilt, as illustrated in Fig. 3 (see Kanade, 1981; Stevens, 1981, and Appendix A for a derivation of the constraint curves).

The constraint curve for a given skew symmetric figure defines the set of surface orientations that are consistent with a symmetric interpretation of the figure. The visual system must rely on additional constraints on the shape of the three-dimensional figure or auxiliary cues like stereo to uniquely determine the orientation of a skew symmetric pattern in the image.

The constraint curves shown in Fig. 3 demonstrate that the ambiguity in the orientation of a symmetric figure changes significantly with the spin of the figure. When the symmetry axis is aligned with the direction of tilt (spin = 0°), as in Fig. 3a, the constraint curve is composed of two vertical lines: one at the true tilt and another 90° away (slant is indeterminate). When the spin is 45°, as in Fig. 3c, the constraint curve is symmetric around the tilt direction and concave up. Intermediate slants (Fig. 3b and d) give rise to constraint curves that are asymmetric around the figure's true slant and tilt.

While the spin of a figure strongly impacts the structure of the information provided by skew, it should have little effect on the information provided by a cue like stereo (assuming a reasonable distribution of tangents in the projected contour). We therefore quantified the influence of skew information by measuring judgments of surface orientation as a function of spin for stereoscopically presented images of symmetric figures.

2. Experiment 1

Pilot studies using monocular projections of symmetric figures revealed systematic biases in subjects' estimates of surface tilt as a function of a figure's spin (see Fig. 4). Experiment 1 is designed to characterize more fully how perception of surface orientation depends on

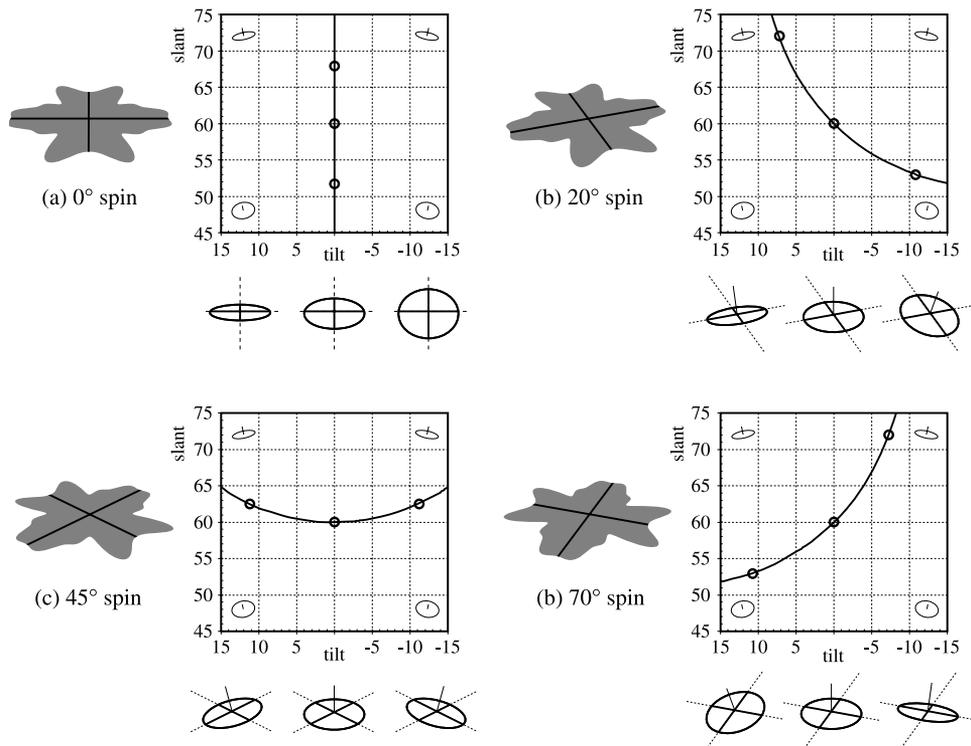


Fig. 3. Contours with different spins, and depictions of some possible mirror symmetric interpretations. The graph for each contour plots combinations of slant and tilt that admit symmetric interpretations of the skew symmetry. For any slant–tilt combination on a constraint curve, back-projecting the contour onto a plane would map the skew axes onto a pair of perpendicular lines in 3D, which define a mirror symmetric 3D interpretation. The circular figures below the graphs depict some sample interpretations, corresponding to the marked points on the constraint curves. In each case, the solid lines inside the circular figures would be the projections of perpendicular axes in 3D. Because there are different interpretations with the same projected axes, the orientations of the skew axes provide only ambiguous information about slant and tilt. The specific ambiguity varies for the different contours, depending on spin.

the spin of symmetric figures when viewed binocularly. The stereo information in the stimulus images was entirely consistent with the information provided by skew symmetry and was constant across changes in figural spin. Any biases that appear in subjects’ orientation settings should therefore be attributable to the influence of figural information and should help to elucidate the computational principles used by the visual system to interpret skew symmetry.

2.1. Methods

2.1.1. Subjects

Thirteen subjects participated in the experiment. Subjects were undergraduates at the University of Rochester and were naïve to the purposes of the experiment. All subjects had normal or corrected-to-normal vision.

2.1.2. Stimuli

Stimuli were stereoscopic images of planar, symmetric figures that were slanted relative to the frontal image plane. We used a weighted sum of cosine functions in polar coordinates to generate random figures,

$$r(\theta) = \sum_k \frac{w_k}{f_k} \cos(f_k \theta + \phi_k) \tag{1}$$

where r is the radial distance away from the origin, θ is the polar angle, f_k and ϕ_k are the frequency and phase of a particular cosine function, and w_k denotes random weights assigned to the cosine function with frequency f_k . To force the figures to be bilaterally symmetric, the phases of the cosine functions (ϕ_k) were constrained to be 0° or 90° .

We further constrained the figures to be twofold symmetric (to be symmetric around both horizontal

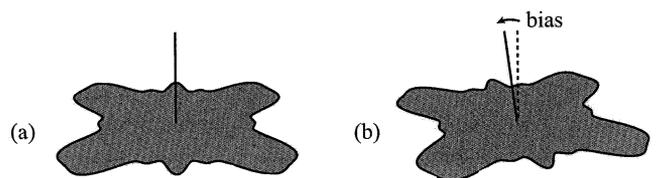


Fig. 4. Illustration of the bias observed in a pilot study. The left and right contours have same slant and tilt, but have different spins. The effect of spin on subjects’ tilt judgments is graphically depicted by the superimposed normal lines. For contours that are aligned with the direction of tilt (a), judgments of tilt were accurate, while for rotated contours (b), judgments of tilt tended to be biased in the direction of rotation.

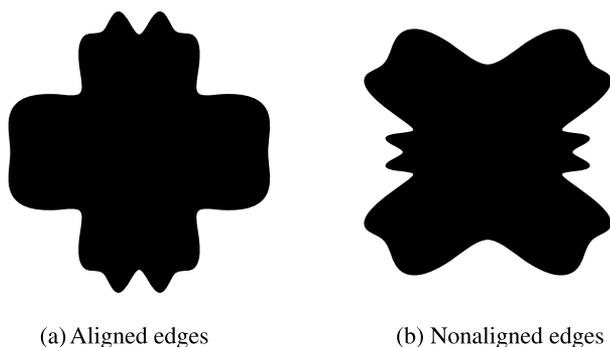


Fig. 5. Samples of the two types of twofold symmetric shapes used in the experiments. (a) For aligned-edge contours, as shown on the left, most edges are parallel to one of the two axes, which, in this case, is vertical and horizontal. (b) For nonaligned-edge contours, as shown on the right, few boundary edges are aligned with the symmetry axes.

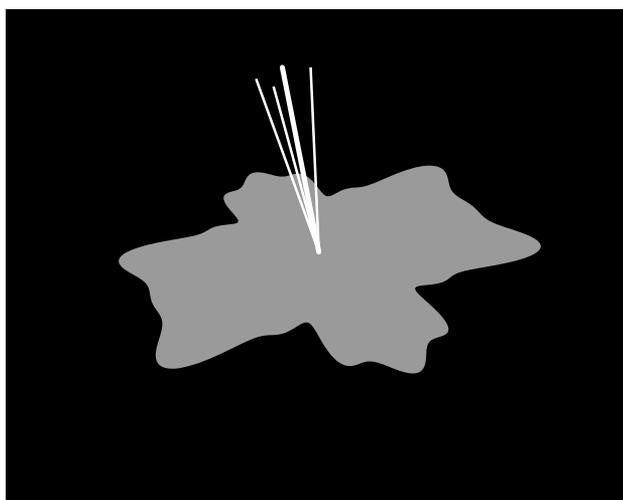


Fig. 6. The perpendicular line task. Subjects adjusted a probe line to appear perpendicular to the plane of the object. The line moved as if attached to the center of the figure.

and vertical axes), using only even numbered frequencies for the cosine functions in the series ($f_k \in \{2, 4, 6, 8, 10, 12, 14, 16\}$). Twofold symmetric figures were chosen for two reasons. First, it simplifies the analysis by eliminating the distinction between the direction of axes of symmetry and direction of symmetry lines. For a twofold symmetric figure, both directions that define a skew symmetry correspond to axes of symmetry, which insures that neither axis can hold a privileged perceptual position. Second, the presence of two directions of symmetry normalizes for the additional information present in a perspective projection. For figures with a single axis of symmetry, information from perspective distortion varies with spin, while for twofold symmetric figures, this information is approximately invariant to spin, and therefore would not be confounded with skew symmetry. This point will be made more clear later (see Appendix B).

We forced the figures to be isotropic by scaling the vertical coordinates of the contour so that the vertical and horizontal moments of inertia were equal (the diagonal elements of the moment tensor were constrained by symmetry to be zero). The procedure amounted to fitting an ellipse to each figure and re-scaling the figure so that the best-fitting ellipse was a circle. This guaranteed that, on average, the foreshortening of the projected contours specified a surface orientation equal to that specified by stereo. To minimize the distortions introduced by this re-scaling, randomly generated figures with moments of inertia that differed by more than 25% were discarded. Subjectively, this tended to eliminate shapes that were not very circular despite being isotropic by the matched moment of inertia criteria (for example, shapes with long thin protrusions).

To control for the possibility that subjects might be biased to align the probe line with the dominant tangent direction in the stimulus contours, we used two categories of random figures for experimental stimuli. The first had tangent direction histograms that had peaks at 0 and 90 degrees. The second had broader tangent direction histograms, selected to have no identifiable peaks at 0 and 90 degrees. These were drawn from the larger set of random figures generated using Eq. (1). Fig. 5 illustrates the two classes of stimuli. We will refer to the first class as *aligned-edge* figures (Fig. 5a), because their edges tend to align with the two symmetry axes. The second class has edges distributed primarily in off-axis directions. We will refer to these as *nonaligned-edge* figures (Fig. 5b). Twenty-six random contours of each type were generated for use in the experiment.

Stimuli were scaled so that when rendered to the computer monitor, they subtended a visual angle of approximately 12° across the longest axis of the figure (the axis perpendicular to the surface tilt). This was done by scaling them so that were they to have been shown in the fronto-parallel plane (0° slant), the best-fitting circle to the figures would have had a diameter of 12° .

Stimuli were rendered using a geometrically correct perspective projection of the figures to left and right eye views for a subject seated 50.8 cm (20 in) from the monitor. Interocular distance was measured for each subject to determine the accurate stereo perspective. To present separate images to left and right eyes, subjects viewed the computer monitor through liquid-crystal stereo glasses (Crystal Eyes). Left and right eye views were displayed at 120 Hz interlaced, giving an effective binocular refresh rate of 60 Hz. The monitor's resolution in stereo mode was 1280×512 pixels, with a viewing area of 38 cm horizontally by 30 cm vertically. Thus, for the viewing distance used, stimulus figures subtended approximately 360 pixels along their long axes. Stimulus figures were filled and colored red (blue and green guns were off) to take advantage of the relatively fast red

phosphor in the monitor and avoid ghosting effects in the stereo display.

2.1.3. Procedure

On each trial of the experiment, subjects used the computer mouse to adjust a stereoscopically displayed probe line to appear perpendicular to the plane of the figure in the display (Fig. 6). The line was centered at the center of mass of the three-dimensional stimulus figure. We mapped the x - y position of the mouse to the upper hemisphere of a unit sphere, so that movements of the mouse appeared to rotate the tip of the probe line over the surface of a sphere centered on the base of the probe.

The initial orientation of the probe was randomly chosen from an annulus of the normal directions on the unit sphere, which was centered on the true orientation of the stimulus figure, and had an inner radius of 30° and an outer radius of 50° . This guaranteed that the initial orientation of the probe was well away from the perceived perpendicular to the figure, forcing subjects to actively adjust the probe on every trial. The task was easily conveyed to subjects, and they were allowed 10–20 practice trials to familiarize themselves with the task and the interface.

Subjects performed the experiment in two separate hour-long sessions, scheduled on different days. In each session, subjects performed two blocks, one with aligned-edge contours and the other with nonaligned-edge contours, with order counterbalanced across sessions and subjects. Within each block, symmetric contours were presented at two slants (45° and 60°), 13 tilts (between -30° and $+30^\circ$ in 5° steps) and six spins (0° , 15° , 30° , 45° , 60° and 75°). Spin was defined relative to stimulus tilt, so that 0° represented a stimulus with a symmetry axis aligned with stimulus tilt. The range of spins used in the experiment spanned the range of possible spins for twofold symmetric figures, since spins of ω , $\omega - 90^\circ$, and $\omega + 90^\circ$ are equivalent for such figures. Each slant/tilt/spin combination was presented once within a block. Within an aligned- or nonaligned-edge block, 26 different randomly generated figures were used, with one assigned to each slant and tilt combination. The figures were shown once at each of the six spins over the course of a block, making 156 trials per block. We were interested in subjects' orientation settings relative to a frame of reference aligned with a figure's tilt; thus, each subject effectively repeated each combination of the three experimental factors (figure type, slant and spin) 13 times per block.

2.1.4. Apparatus

Images were generated on an SGI Indigo 2 Extreme and presented in stereo using Crystal Eyes shutter glasses that oscillated at 120 Hz (60 Hz for each eye's view). Subjects viewed the monitor in the dark from a chin rest 50.8 cm away from the screen, through

a matte black occluder that covered the edges of the monitor.

2.2. Results

Subjects' orientation settings consisted of a set of normal vectors indicating the direction of the probe line in three dimensions. In order to average each subject's orientation settings across different tilts, we first rotated the measured normal vectors into coordinate frames aligned with the simulated stimulus tilts. The resulting normal vectors represented subjects' orientation estimates relative to the true stimulus tilt. We averaged the rotated normal vectors within each experimental condition by first computing their arithmetic mean and then re-normalizing, giving a unit vector that represented the average orientation setting relative to stimulus tilt. We computed slant and tilt settings for each subject from the average normal vectors and performed statistical tests on these values.

Fig. 7 plots the resulting mean slant and tilt settings as a function of spin for the each of the slant and figure type conditions (45° and 60° slant, aligned-edge and nonaligned-edge figures). Orientation settings varied as a function of spin, forming a circular pattern in slant-tilt space. A useful way to view the results is to treat 60° and 75° spins as equivalent to -30° and -15° , respectively (the equivalence follows from the fact that the figures have two axes of symmetry). With this transformation, we see that positive spins of 15° and 30° led to positive biases in subjects' tilt settings relative to the mean response, while negative spins led to negative biases. The biases are qualitatively similar to those depicted in Fig. 4; they can be described as being toward the symmetry axis nearest the true tilt of the figures. Slant settings also varied systematically with spin, decreasing from a peak at 0° spin to a minimum at 45° spin. The pattern is most pronounced for the 60° slant stimuli (Fig. 7a and b), but is also evident for the 45° slant stimuli (Fig. 7c and d).

We performed independent ANOVAs on the orientation settings of the 60° and 45° slant stimuli. For the 60° stimuli, the ANOVAs revealed a main effect of spin ($F(11,263) = 15.31$, $P < 0.001$), no effect of figure type ($F(3,263) = 1.883$, $P = 0.1327$), and no interaction between figure type and spin ($F(9,263) = 0.9976$, $P = 0.4423$). The results for 45° stimuli were similar: a main effect of spin ($F(11,263) = 4.804$, $P < 0.001$), but no effect of type ($F(3,263) = 0.2127$, $P = 0.8876$), nor any interaction between type and spin ($F(9,263) = 0.6437$, $P = 0.7591$). The absence of interactions indicates that the pattern of bias as a function of spin was not different for the aligned-edge and nonaligned-edge figures.

The same pattern of biases was observed for both 60° and 45° slant conditions, but the pattern was more pronounced for the 60° slant stimuli. To further test the effect of spin in the 60° slant case, we did pairwise

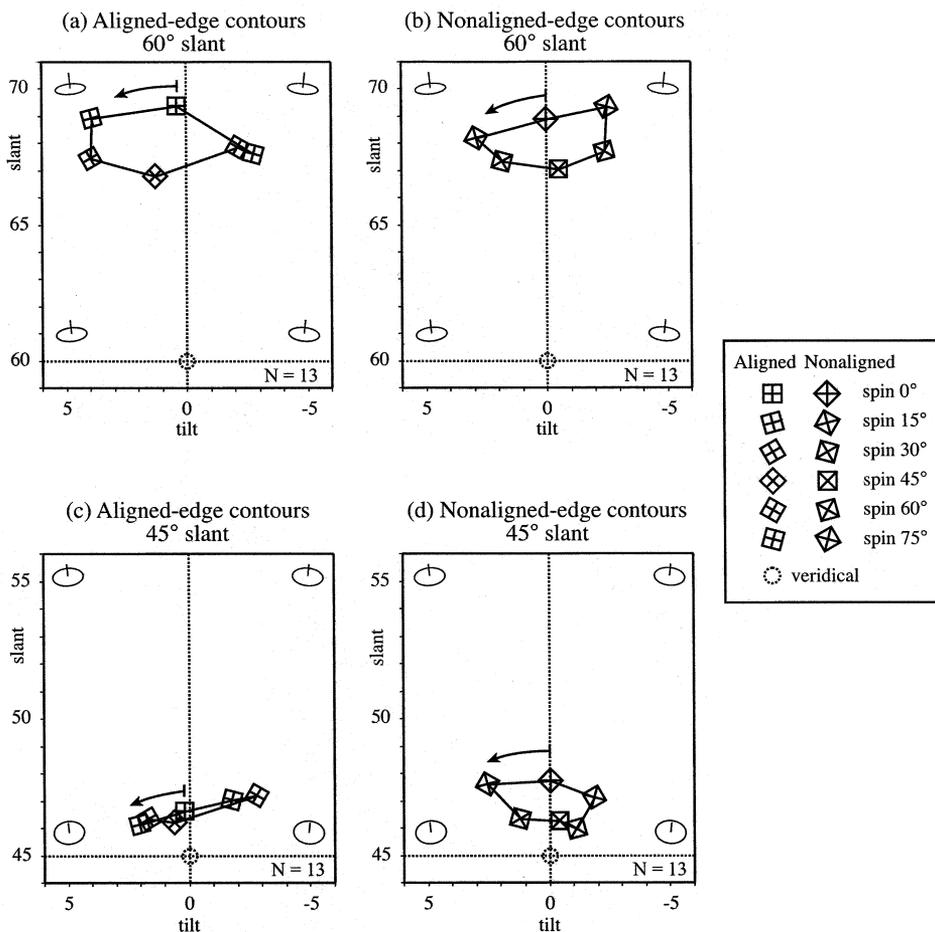


Fig. 7. Results of experiment 1. The four plots show mean orientations of the probe line in slant–tilt space for objects with 60° slant (top, a and b) and objects with 45° slant (bottom, c and d). The left two plots are the results for aligned-edge contours (a and c), and the plots on the right are for the nonaligned-edge contours (b and d). Points within a graph correspond to conditions with different spins, and are represented by squares rotated by the spin angle. The arrows show the direction of bias as spin is increased from zero. Note that the x -axis has been reversed to be consistent with the convention for signed rotation (a positive, counterclockwise rotational bias would correspond to a leftward shift in probe settings).

comparisons between conditions with matched positive and negative spins, and between conditions with 0° and 45° spin. Comparing figures with different directions of spin, there were significant differences in tilt settings (+15 vs. –15: $t(55) = 64.5$, $P < 0.001$; +30 vs. –30: $t(55) = 36.0$, $P < 0.001$), but no differences in slant settings (+15 vs. –15: $t(55) = 0.01$, $P = 0.906$ n.s.; +30 vs. –30: $t(55) = 0.72$, $P = 0.399$ n.s.). For the 0° spin and 45° spin conditions, there was a significant difference between slant components but not tilt components (slant: $t(55) = 12.2$, $P < 0.001$; tilt: $t(55) = 0.02$, $P = 0.894$ n.s.). These comparisons statistically confirm two salient features of the circular pattern of biases: different tilt biases for positive and negative spins ($\pm 15^\circ$ and $\pm 30^\circ$), and different slant biases for the two spins that produce skew axes that are symmetric relative to the tilt direction (0° and 45°).

For figures presented at a slant of 60°, subjects' slant settings showed a significant overall positive bias. The average slant setting was 68° for these conditions, with

a standard deviation of 4.6°. A trend toward overestimation is also present for figures with 45° slant (mean slant setting 47°, S.D. 5.7°). The data do not allow us to determine whether the positive bias reflects an overestimate of figure slant or an underestimate of probe line slant (or both).

The means reported above are of probe judgments that have been normalized for tilt direction. We do not have enough data to analyze the effect of spin separately for each tilt used, but it is possible to analyze the overall effect of tilt by averaging absolute tilt settings across spin conditions. We found that there was an overall tendency for judgments to be biased away from vertical. The bias was approximately linear, with tilts being overestimated by a constant factor of about 30%. For stimuli with the maximum tilts tested, $\pm 30^\circ$, the bias was about 9°. The presence of an overall bias away from vertical would add variability to the mean data, but would not be expected to contribute to differences between spin conditions.

2.3. Discussion

The results show that the alignment of a figure's symmetry axis within a plane (its spin) affects judgments of its slant and tilt. One possibility is that the influence of figure alignment is a consequence of how skew symmetry varies with spin. We chose to manipulate spin because it varies the amount of skew while leaving foreshortening and stereo information intact. If skew symmetry does play a role in the perception of surface orientation, the effect of figure alignment on subjects' judgments could be mediated by its effect on skew.

Before concluding this, we must rule out factors other than skew symmetry that also vary with spin. One source of stimulus variability that could potentially lead to the effect of spin was the distribution of tangent directions in stimulus figures. For individual exemplars of the figures, the distribution of tangents was not uniform. Consequently, the distribution of tangents for a projected contour was also not uniform, and varied with spin. The covariation of orientation estimates with spin could have arisen from an indirect effect of dominant edge orientation on stereo estimates of surface orientation or from a bias to match the probe line to the dominant orientation edge orientation in a stimulus. The fact that subjects' orientation settings did not differ significantly between the aligned-edge and nonaligned-edge stimuli, however, argues strongly against this interpretation. Were the distribution of edge orientations to have affected subjects' judgments, we would have expected a significant effect of edge type on subjects' orientation estimates.

Another factor that might influence judgments is the additional information provided by perspective distortion of a contour. In an orthographic projection, symmetry lines in an image are parallel, while in a perspective projection, symmetry lines converge to the horizon and therefore provide a cue to surface orientation. Over the extent of the projected contour, this effect would be expected to be a weak cue. Nevertheless, any information that it provides might vary with spin. We controlled for the possible role of perspective by using twofold symmetric contours, rather than figures with one axis of symmetry. Twofold symmetric contours have two axes of symmetry and hence two sets of symmetry lines. When viewed in perspective, the two sets of symmetry lines produce two complementary sets of perspective gradients. The combined information from the two sets of projected symmetry lines is approximately invariant to spin (see Appendix B) and would therefore be unlikely to induce spin-dependent biases.

Thus, the controls built in to the stimulus set implicate skew symmetry as the likely causal factor in the observed biases. The question remains as to why differ-

ent figural spins should lead to different biases. The method by which we generated the stimuli—geometrically correct binocular renderings of symmetric figures—offers no direct insight into why subjects' orientation settings should vary lawfully with the spin of a symmetric figure. The true orientation of the figure is geometrically consistent both with the stereo information provided in the displays and with the interpretation of the figure as mirror symmetric. The constraint on slant and tilt imposed by symmetry does change with spin, as illustrated earlier, but under conditions of consistent information, one would not expect the addition of symmetry information to result in biases.

Had there been a discrepancy between the information provided by skew symmetry and stereo, biases could arise due to the interaction between conflicting cues. Although stimuli were generated to be geometrically correct, the presence of measurement errors or perceptual biases independent of skew symmetry could effectively introduce conflicts between cues. We propose that an implicit cue conflict underlies the effect of figural spin observed in the data. By this explanation, the spin-dependent biases result from a process that integrates skew symmetry information, which depends on spin, with biased information from other sources that do not depend on spin. Empirical data on perceptual biases in estimates of relative depth from stereo suggest that in the near viewing distance used in the experiment (50 cm), subjects would overestimate slant from stereo.² This results from a bias to interpret stereoscopically presented stimuli as more elongated in depth at viewing distances less than 1–1.5 m (Johnston, 1991). The perceptual bias in slant-from-stereo estimates would effectively create a conflict between stereo and skew symmetry information.

Fig. 8 illustrates how a stereo slant bias could produce spin-dependent tilt biases when combined with the information provided by skew symmetry. In the figures, we represent the information provided by stereo by a point in slant-tilt space and the information provided by skew symmetry as a constraint curve, the latter representing the set of 3-D orientations for which the skewed symmetry axes in the image would back-project to orthogonal axes on a surface. For conditions that produced tilt biases (Fig. 8a and b), the constraint curves are sloped through the true orientation of the stimulus figure. An optimal integration process would

² Empirically measured perceptual distortions of space show underconstancy in perceived depth as viewing distance is increased (Foley, 1980), particularly in the absence of any visual context. For near viewing distances, this would correspond to a stretching of perceptual space in the depth dimension. Johnston (1991) observed perceptual stretching of depth for objects closer than 1–2 m; that is, subjects overestimate relative depth for objects closer than 1–2 m. The empirically measured depth stretching would lead to an overestimate of slant in the present experimental viewing conditions.

choose a compromise solution that lay somewhere between the orientation suggested by stereo (presumably biased) and the constraint curve defined by the skew axes of the figure. The resulting combined estimate would be biased upward in the direction of the sloped constraint curves, producing shifts in tilt as well as slant. The qualitative difference in estimated tilt between positive and negative spins is in the same direction as the tilt biases observed in the results

Fig. 8c and d illustrate how a cue integration process could explain the effect of spin on subjects' slant estimates as well. The constraint curve is vertical for 0° spin (Fig. 8c) but is approximately horizontal near the true orientation for 45° spin (Fig. 8d). Consequently, if an estimate of slant from stereo was biased, the symmetry axes in the 0° case would provide no contradictory information, while the axes in the 45° case would provide conflicting information that could be used to correct the biased estimate. If the two constraint curves were combined with an overestimate of slant, one

would expect that figures with 45° spin would be perceived to have lower slant, as is observed in the data. In the next section, we will elaborate this explanation into a statistically optimal model of cue combination.

3. Bayesian model

In the previous section, we discussed the possibility that the observed slant and tilt biases could be due to an interaction between the constraints imposed by a symmetric interpretation of the figures and a biased estimate of surface orientation from stereo (and perhaps foreshortening). The support for this explanation is that the directions of the observed biases were consistent with choosing orientation estimates that are intermediate between a biased estimate of slant from stereo and the constraint curves defined by figural skew. Under this interpretation, the results depend not only on whether skew symmetry is used, but also on how skew

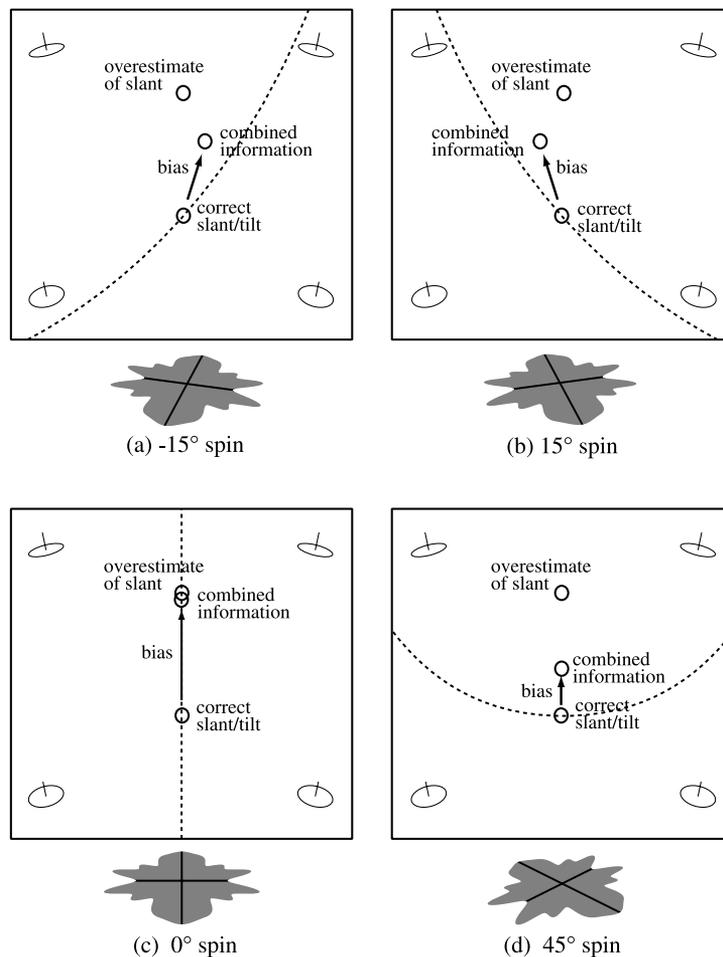


Fig. 8. Hypothetical effect of combining an overestimate of slant from stereo with information from skew symmetry for contours with different spins. When constraint curves are sloped in the neighborhood of the stereo-specified slant (a and b), integration with a biased slant estimate should produce tilt biases as well as slant biases. In the 0° case (c), contour symmetry provides no information about slant, so any bias in perception of slant from stereo would directly translate to bias in judgments. In the 45° case, skew symmetry does constrain the possible slant of the figure, so this information could reduce the effect of any bias in slant from stereo.

symmetry is combined with other sources of information. A relatively strong model of cue integration would be required, since one would have to assume that the process of cue integration is sensitive to the different geometric ambiguity of contours with different spins.

In this section, we develop a Bayesian model of cue integration that qualitatively accounts for the empirically measured pattern of orientation biases. We begin by describing the Bayesian approach to cue integration. We then apply the formulation to the problem of integrating skew symmetry information with stereo information.

3.1. Bayesian cue integration

Given two cues about a scalar surface property X (like slant), the easiest way to combine the information provided by the cues is to compute a weighted sum of the estimates of X derived independently from each cue. We represent this in the simple linear equation

$$\hat{X} = w_1 \hat{X}_1 + w_2 \hat{X}_2 \quad (2)$$

where \hat{X} is the combined estimate of X derived from the two cues, \hat{X}_1 and \hat{X}_2 are estimates of X derived from each cue independently, and w_1 and w_2 are the weighting factors for the two cues. Clearly, the weights, w_1 and w_2 , should depend on the relative reliability of the two cues, with the more reliable cue given a larger weight than the less reliable cue. If the values of \hat{X}_1 and \hat{X}_2 are Gaussian-distributed with unbiased means, the optimal estimate (i.e. estimate with lowest variance) of X is given by weighting the cues in inverse proportion to their variances (Landy, Maloney, Johnston, & Young, 1995).

Linear cue combination is a special case of an optimal strategy for combining cues. In order to model more general optimal cue integration strategies, we need to pose the problem in a full probabilistic, or Bayesian, framework (Clark & Yuille, 1990). The fundamental problem of cue integration then amounts to formulating a posterior conditional probability distribution of a surface property X , $p(X|I_1, I_2, I_3, \dots, I_n)$, that quantifies the information provided by the set of available image cues $I_1, I_2, I_3, \dots, I_n$. The posterior distribution specifies the relative probability of different values of X , given the information provided by the cues. Using Bayes' rule, the posterior distribution expands to

$$p(X|I_1, I_2, \dots, I_n) = \frac{p(I_1, I_2, \dots, I_n|X) \cdot p(X)}{p(I_1, I_2, \dots, I_n)} \quad (3)$$

where $p(I_1, I_2, I_3, \dots, I_n|X)$ is the likelihood of seeing the data given X , and $p(X)$ is a prior distribution on X reflecting our a priori knowledge about X , for example, that depth varies smoothly across a surface. If we assume that the image cues are independent, we can write the posterior as the product of a series of likeli-

hood functions, one for each cue, and the prior distribution,

$$p(X|I_1, I_2, \dots, I_n) \approx p(X|I_1)p(X|I_2)\dots p(X|I_n) \quad (4)$$

(we have discarded the denominator for simplicity, since, for our purposes, it merely serves as a normalizing constant). If the individual likelihood functions are Gaussian, and the prior is Gaussian, then the posterior distribution is Gaussian with a mean given by

$$\hat{X} = \frac{1}{(1/\sigma_1^2 + 1/\sigma_2^2 \dots + 1/\sigma_n^2 + 1/\sigma_p^2)} \times \left[\frac{1}{\sigma_1^2} \hat{X}_1 + \frac{1}{\sigma_2^2} \hat{X}_2 \dots + \frac{1}{\sigma_n^2} \hat{X}_n + \frac{1}{\sigma_p^2} \hat{X}_p \right] \quad (5)$$

where \hat{X}_i is the mean of the likelihood function for the i th cue, \hat{X}_p is the mean of the prior distribution and σ_i denotes the standard deviations of the likelihood functions (or the prior). Thus, assuming that we select the mean as our estimate of X , the Gaussian estimation problem reduces to the simple linear model described earlier. Deviations from the Gaussian assumption may lead to optimal integration models that are non-linear and, in general, may require explicit computation of likelihood functions.

While, in practice, computing full likelihood functions can be difficult, the Bayesian formulation of cue integration is conceptually no more complex than the simple linear one. For independent cues, a full Bayesian analysis requires one to multiply entire distributions over the space of values for X , rather than reducing the information provided by individual cues about X to single values and combining them algebraically through addition. When likelihood functions are significantly non-Gaussian, this can lead to qualitatively different behavior in Bayes optimal and linear estimators. As we will show, this is the case for information provided by skew symmetry.

3.2. Modeling information from skew symmetry

In order to incorporate the information from skew symmetry into a Bayesian model, we must formulate a likelihood function for skew symmetry. The desired likelihood distribution can be thought of as a fuzzy constraint curve, which has some additional uncertainty due to measurement noise beyond the fundamental geometrical ambiguity of the symmetry constraint. With generic assumptions, this likelihood function can be directly computed.

We define skew symmetry by a pair of skew axes $\{\theta_1, \theta_2\}$, which, for a symmetric interpretation of a contour, would be the projections of perpendicular axes on the figure. While symmetric figures do not project to strictly skew symmetric figures under perspective projection, the approximation is fairly good for the small stimuli that we used. Moreover, our analysis only as-

sumes that the visual system can extract noisy estimates of the projected axes of symmetry of a symmetric figure under perspective projection, not that the projected figure itself is skew symmetric. Our goal is to compute the likelihood of a given pair of skew axes for any possible slant–tilt combination, $p(\theta_1, \theta_2 | S, T)$, where S and T are the slant and tilt of a figure in three dimensions.

If the slant, tilt and spin of a figure are known, the projected skew axes are geometrically specified (see Appendix A). Observers' estimates of the skew axes would be expected to vary randomly around the geometrically specified skew axes. Assuming Gaussian noise in axis measurements, we have as a generative model of the image data:

$$\begin{aligned}\theta_1 &= f(S, T, \omega) + \varepsilon_1 \\ \theta_2 &= f(S, T, \omega) + \varepsilon_2\end{aligned}\quad (6)$$

where $f(S, T, \omega)$ is the projective mapping of a figure's axes to the image, ω is the spin of the figure in its three-dimensional plane, and ε_1 and ε_2 are noise processes. We assume that the noise variables are independent, identically distributed Gaussian random variables ($\varepsilon_1, \varepsilon_2 \sim N(0, \sigma^2)$), though the particular form of the noise distribution does not significantly impact the behavior of the model. Given the noise model, the joint likelihood function for a pair of projected axes $\{\theta_1, \theta_2\}$ conditioned on slant, tilt and spin, is simply the product of two normal distributions centered around the geometrically specified skew axes:

$$\begin{aligned}p(\theta_1, \theta_2 | S, T, \omega) \\ \sim \exp\left(-\frac{(f(S, T, \omega) - \theta_1)^2}{2\sigma^2}\right) \\ \times \exp\left(-\frac{(f(S, T, \omega + 90^\circ) - \theta_2)^2}{2\sigma^2}\right).\end{aligned}\quad (7)$$

Integration of Eq. (7) over spin gives a likelihood function parameterized only by slant and tilt. Assuming that all spins are equally likely, we obtain

$$p(\theta_1, \theta_2 | S, T) = \frac{1}{2\pi} \int_0^{2\pi} p(\theta_1, \theta_2 | S, T, \omega) d\omega.\quad (8)$$

Because the integral in Eq. (8) is intractable, we compute it numerically. Fig. 9a shows examples of likelihood functions derived for symmetric figures with a slant of 60° and a tilt of 0° (vertical) with various spins. In the absence of other sources of information, these likelihood distributions specify the relative probability of different combinations of slant and tilt given the skew axes. As one would expect, the likelihood functions appear as blurred skew symmetry constraint curves. The amount of blurring depends on the standard deviation of the noise, which, for these plots, was chosen to be 5° . Other noise models result in similar likelihood functions; the form of the noise distribution determines the exact pattern of drop-off in likelihood as

one moves away from the deterministic constraint curve.

3.3. Combining skew symmetry and stereo

The optimal estimate or surface orientation from combined sources would maximize the posterior probability of slant and tilt given the image data from both these cues, $p(S, T | D_{\text{stereo}}, \theta_1, \theta_2)$, where D_{stereo} represents the disparity information provided by stereo. If stereo is assumed to be independent of skew symmetry, this joint probability can be computed from Eq. (4),

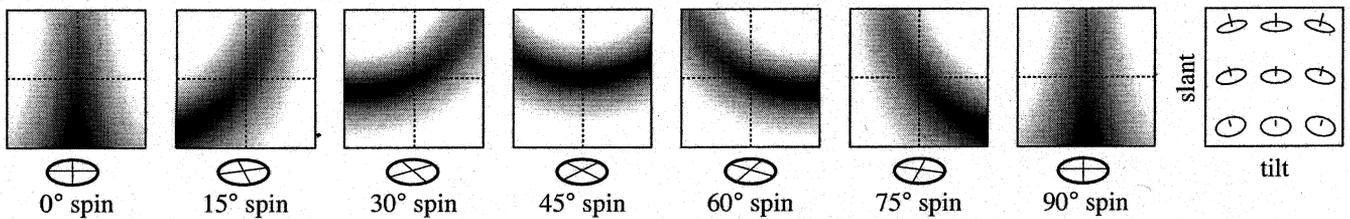
$$p(S, T | \theta_1, \theta_2, D_{\text{stereo}}) = p(\theta_1, \theta_2 | S, T) p(D_{\text{stereo}} | S, T) p(S, T)\quad (9)$$

where $p(\theta_1, \theta_2 | S, T)$ and $p(D_{\text{stereo}} | S, T)$ are the likelihood functions for the independent cues of skew symmetry and stereo, respectively, and $p(S, T)$ is the a priori likelihood of slant and tilt. For simplicity, we will assume the prior on slant and tilt to be uniform. In practice, this assumption is valid as long as the likelihood functions are reasonably compact in slant–tilt space.

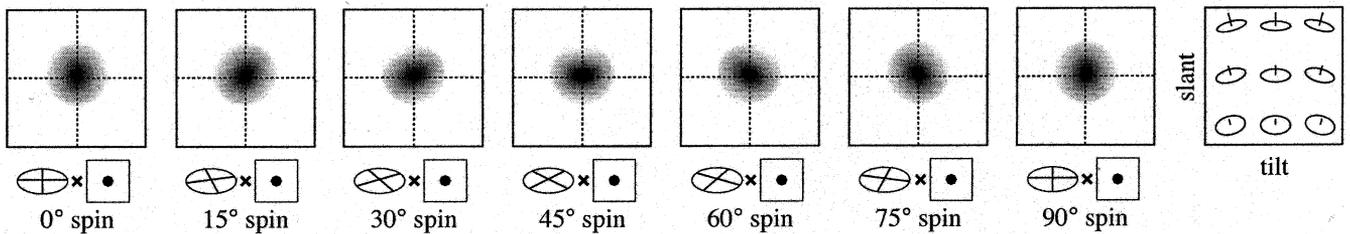
In the previous section, we derived likelihood functions for skew symmetry corresponding to different spins. Ideally, the likelihood functions for stereo would be based on a measurement model, as we have done in the case of skew symmetry. However, we do not have the data available to specify such a model. For purposes of making qualitative predictions about the interaction between cues, we will make the simplifying assumption that the likelihood function for stereo can be modeled as a Gaussian distribution in slant–tilt space. Because stereo specifies a slant–tilt combination, rather than a curve of possibilities, a distribution around a point is a natural generic representation. We further assume that information from stereo is independent of spin, which allows us to multiply the same stereo likelihood function with each of the spin-dependent likelihood functions from skew symmetry. What remains to specify are the parameters of the likelihood function for stereo: the mean and covariance matrices for the likelihood function. The mean value for slant reflects biases in perceptual estimates of surface orientation from stereo (we assume that for vertical tilts, subjects' estimates of tilt from stereo are unbiased). The variance and covariance terms reflect the uncertainty in the orientation specified by stereo information.

Regardless of the exact values used for the free parameters, it is clear that if stereo information is unbiased, the result of combining the two cues will be a posterior peaked at the true orientation of a figure. Fig. 9b shows posterior distributions for different spins assuming unbiased stereo information, which were computed by multiplying a Gaussian centered around the true orientation (representing stereo information) with

(a) Likelihood distributions for skew symmetry



(b) Skew symmetry combined with accurate slant



(c) Skew symmetry combined with biased slant

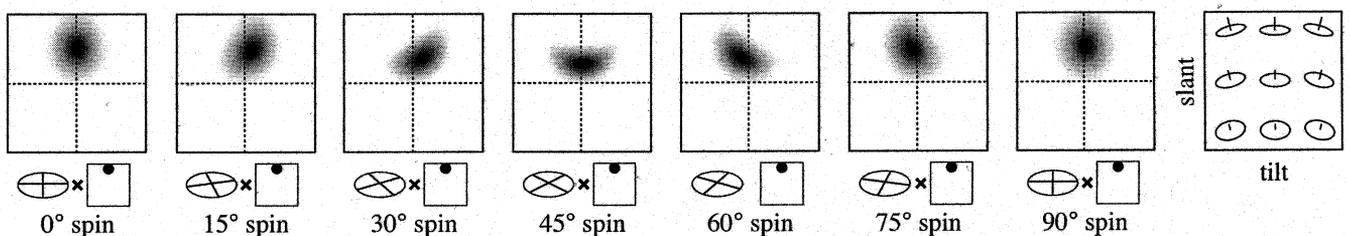


Fig. 9. Likelihood functions for skew symmetry alone, skew symmetry combined with unbiased stereo, and skew symmetry combined with biased stereo. The functions are all over slant/tilt space, with tilt represented along the x -axis and slant represented along the y -axis (center of the graph is veridical). (a) Likelihood functions from skew symmetry for different spins, which were derived geometrically assuming Gaussian noise in estimates of the orientations of skew axes. Note that the regions of high density fall along ridges corresponding to the skew-symmetry constraint curve for each spin. (b) Results of multiplying each of the likelihood functions from skew symmetry with a Gaussian distribution around the true slant and tilt, which represents unbiased stereo information. The peaks of the combined likelihood functions are centered around the correct slant and tilt, but have shapes that depend on spin. (c) Results of multiplying each of the likelihood functions from skew symmetry with a Gaussian distribution centered around the correct tilt but shifted upward by 10° , which represents biased stereo information. The peaks of these combined likelihood functions are shifted away from the true orientation, and the amount and direction of bias varies with spin.

each of the likelihood functions from skew symmetry shown in Fig. 9a. The peaks of all the distributions are at the true orientation; only the shapes differ, reflecting the differences in the skew symmetry likelihood function as a function of spin.

Fig. 9c shows the result of multiplying skew symmetry likelihood functions with biased stereo likelihood functions (slant from stereo is assumed to be overestimated by 10° , a figure consistent with existing data on perceptual distortions of depth at 50 cm viewing). The general effect of the stereo bias is to shift the peaks of the distributions upward from the true orientation, but the amount and direction of bias varies depending on spin. In effect, the information provided by the contour's symmetry axes 'pulls' the biased stereo interpretation towards the corresponding constraint curve. Since the curve shifts and distorts as a function of spin, the pulling effect of skew symmetry changes as a function of spin.

Fig. 10 shows the model results in a way that can be compared more easily to the human data. The graph plots the pattern of orientation estimates as a function of spin derived by taking the peaks of the posterior distributions shown in Fig. 9c. Given the assumed noise model and hypothesized overestimation of stereo, these would be the optimal estimates of surface orientation from a combination of skew symmetry and stereo information. The circular pattern of biases in model estimates is qualitatively similar to the pattern of biases observed in the human data. While these estimates were derived using fixed settings for the parameters of the component likelihood functions, the general pattern was invariant to the actual values of these parameters, as long as stereo had a positive bias in slant.

We have not attempted to fit the model parameters more quantitatively to subjects' data for several rea-

sions. First, the data reflect matches between perceived surface orientation and the perceived orientation of a probe, and not direct estimates of perceived surface orientation. The perceived orientation of the probe is undoubtedly biased (e.g. by the same perceptual depth stretching that we argue biases subjects' slant from stereo estimates), leading to biases in our measure of perceived surface orientation. We do not know how much of the measured biases are due to biased perception of probe orientation, so it would be difficult to infer any quantitative estimate of the perceived overestimation of the figure's slant. Second, the data show sizeable individual differences in the size of biasing effect of figural spin. This can result from individual differences in any of the model parameters. Independent measurements of these parameters for each subject are beyond the scope of the current study.

3.4. Additional contour cues

Our model thus far integrates information solely from stereo and skew symmetry, and neglects other contour cues. Clearly, the shape of a figure in the image provides more information about surface orientation than is contained in the projected axes of symmetry. For example, the overall foreshortening of a figure (represented by the orientation and aspect ratio of an ellipse fitted to the figure) provides information about surface orientation if one has some knowledge of the shape of the figure in three dimensions (e.g. that it is

isotropic). Perspective distortion of a projected contour provides another source of figural information about surface orientation. The symmetry lines in a perspective projection are not parallel, and the direction and rate of change provide a cue to surface orientation. Given that other cues are available, two questions arise. First, can the pattern of biases be predicted from other contour cues without including skew information from the projected axes of symmetry? Second, does the inclusion of additional information qualitatively change the results of the model?

The experimental stimuli were constructed so that both foreshortening and perspective information were approximately invariant to spin, so neither of these cues could by themselves account for the experimental results. In the case of foreshortening information, invariance to spin was a consequence of using isotropic figures. For views of a particular figure, there would be some idiosyncratic variations in foreshortening information as a function of spin, but one would not expect any consistent pattern across individual figure exemplars. To verify this, we computed best-fitting ellipses to the contours used in the experiment and found no overall correlation between spin and foreshortening information.

Invariance to spin for perspective information is a consequence of using twofold symmetric figures. The information provided by a single set of symmetry lines is strongly dependent on spin. However, perspective projections of twofold symmetric figures have two complementary sets of symmetry lines and the combined information is approximately invariant to spin (see Appendix B).

To test whether additional contour information changes the behavior of our model, we extended the model to include foreshortening and perspective information. Foreshortening can be assumed to be independent of stereo and skew symmetry and specifies a slant–tilt combination, so for simplicity, we assumed a Gaussian likelihood function for foreshortening, centered around the correct slant and tilt. Modeling the inclusion of perspective contour information is more complicated because one cannot assume independence from skew symmetry. Like skew symmetry, perspective contour information depends on the accuracy with which estimates of symmetry lines and their orientations can be extracted. To compute a joint likelihood for these cues, we generalized the noise model used for our analysis of skew symmetry. With an additional assumption (see Appendix B) of noisy measurements of perspective gradients, we computed likelihood functions for combined skew symmetry and perspective skew information for each of the spins used.

The extended model simulated the combination of biased stereo information with all three contour cues (foreshortening, skew symmetry and perspective). The

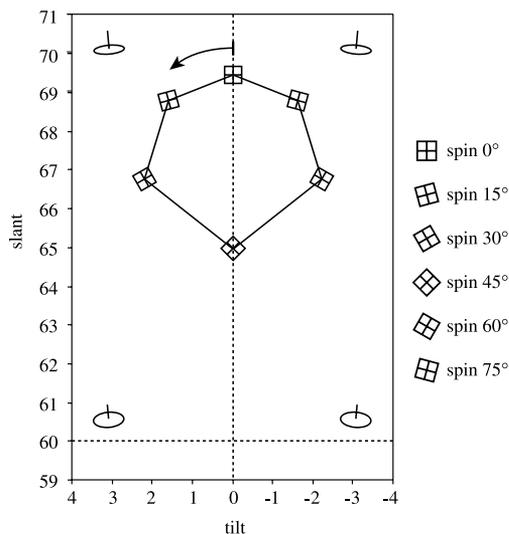


Fig. 10. Simulated orientation biases as a function of spin when skew-symmetry information is integrated with an overestimate of slant from stereo. The graph plots the peaks of the combined likelihood functions from Fig. 8, which would be optimal estimates given the likelihood functions for stereo and symmetry. The orientations of the graph markers indicate the spin condition. The model replicates the qualitative pattern of biases observed in the human data.

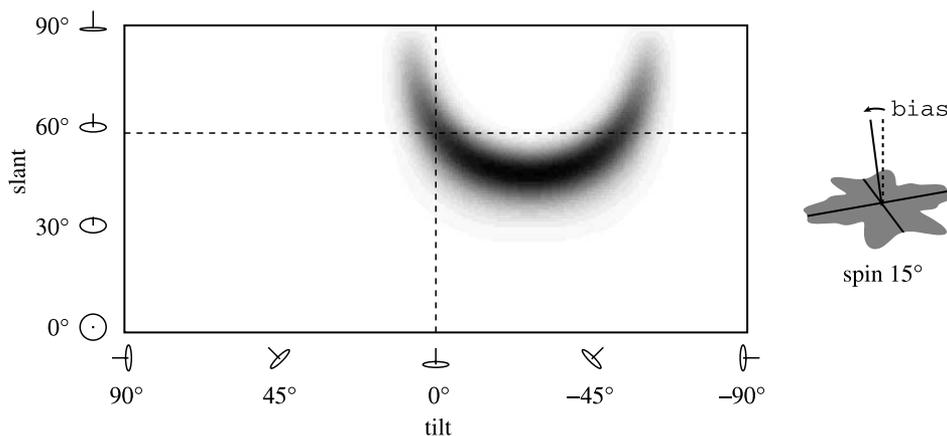


Fig. 11. Sample likelihood function from skew symmetry plotted over an expanded range of slants and tilts. This function was computed for a figure with 60° slant, vertical (0°) tilt, and 15° spin, assuming noise in measurements of skew axes in the projected image (see text for model details). Both the mean of the whole distribution and the minimum slant interpretation of skew symmetry are to the right of the true orientation, while the observed bias for this condition is to the left (positive tilt bias). Thus, neither of these sub-optimal interpretations could account for the orientation biases in the data.

resulting estimates showed the same circular pattern, but the magnitude of biases was smaller. The reduction in bias is not surprising, since the inclusion of additional consistent cues would generally be expected to reduce estimator uncertainty. If one built in some biases in the interpretation of other contour cues, the biases in combined estimates could be modulated. For example, overestimation of slant from foreshortening would have an effect similar to biased stereo and could exaggerate the effect of spin. However, as long as the net effect of stereo and other contour cues is a constant slant bias, the interaction with skew symmetry information would produce similar results. The additional simulation confirmed that for the stimuli we used, twofold symmetric isotropic contours viewed in perspective, the inclusion of additional contour cues does not change the qualitative pattern of results of an optimal cue integration model. Moreover, the circular pattern itself derives from the skew information, and other figure cues serve primarily to modulate the pattern.

3.5. Sub-optimal models

We have proposed that subjects' perception of surface orientation for stereoscopically viewed images of symmetric, planar figures results from a complex interaction between the constraint imposed by skew symmetry and information from stereo. Because the likelihood function for skew symmetry changes with the spin of a projected figure, the interaction between skew symmetry information and other cues like stereo changes with spin. Before accepting such an account, however, we must test whether simpler, sub-optimal models can account for the pattern orientation biases. The simple alternative to the Bayesian model is that the visual system first derives estimates of surface orientation

from figural information and stereo independently of one another and then combines these to arrive at a final estimate of surface orientation. In order to formulate such a model, we require a model of how the visual system might resolve the ambiguity in skew symmetry information. Two ways immediately present themselves. First, the visual system might select something like the mean orientation specified by skew symmetry. Second, the visual system might select the lowest slant solution (see, for example, Stevens, 1981).

Fig. 11 demonstrates why neither of these models could predict the measured biases in subjects' orientation matches. The figure shows an example likelihood function for a figure with a 15° spin. This spin leads to a positive tilt bias (to the left of vertical) in subjects' matches. The likelihood function, however, is shifted toward negative tilts (to the right of vertical). The average orientation specified by skew symmetry (the center of mass of the likelihood function) in this case has a large negative tilt. Similarly, the lowest slant consistent with the skew information has a negative tilt. While the bulk of the likelihood 'mass' is shifted towards negative tilts, subjects' orientation matches are biased in the opposite direction. This is true for all spins.

The most parsimonious explanation of the data is that skew symmetry influences judgments by means of its interaction with other cues. We demonstrated that a constant slant bias in other cues (i.e. stereo) would be sufficient to produce the pattern of biases in the results, assuming that cue combination is optimal.

3.6. Additional model predictions

In our model, spin-dependent biases are the result of integrating skew symmetry information with a biased

estimate of slant from stereo. The effect of the hypothesized slant bias is to create a conflict between symmetry and stereo cues, which, when optimally, combined leads to shifts in both slant and tilt of the combined estimate. If this model is true, a testable prediction is that biases would be modulated by stereo slant. If one increased stereo slant while keeping contour information constant, explicitly introducing a cue conflict, one would expect both the slant and tilt biases to be magnified. Conversely, by reducing stereo slant, one should be able to reduce the spin-dependent biases until they eventually go in the opposite direction.

Fig. 12 shows the model predictions when stereo slant is selectively changed to simulate various cue conflicts. The circular pattern of biases as a function of spin gets larger when stereo slant is increased and is reduced when stereo slant is decreased. When stereo information suggests a slant below the symmetry constraint curve, the biases go in the opposite direction. Another way to think of these predictions is in terms of the effect of stereo conflict for a fixed spin condition. In particular, varying stereo slant (but not tilt) should alter subjects' judgments of both slant and tilt. Most strikingly, the tilt bias can be made to reverse sign as a function of stereo slant.

This behavior can be contrasted to that of a linear model. Suppose that the interpretation of contour information in isolation was subject to spin-dependent

biases, ignoring the problem of why the disambiguated interpretations might be biased in the particular observed direction. As long as contour cues were given some weight when combined with stereo, spin-dependent biases would also be observed even for a linear model. However, if stereo slant were changed while leaving the figural information constant, one would not expect a significant modulation of the spin-dependent effect. Unless symmetry and stereo were for some reason assigned different weights depending on the amount of conflict, perturbing stereo slant would primarily affect the slant component of a linear combined estimate. Thus, the predictions shown in Fig. 12 provide a way to distinguish an optimal Bayesian cue integration strategy from a linear model, even if the a priori assumptions involved in the interpretation of contour information are not known.

4. Experiment 2

We further tested the Bayesian integration model by introducing explicit conflicts between stereo and symmetry and measuring the resulting changes in subjects' perceptual biases.

4.1. Method

4.1.1. Subjects

Eighteen subjects participated in the experiment. Subjects were undergraduates at the University of Rochester and were naïve to the purposes of the experiment. All subjects reported having normal or corrected to normal vision.

4.1.2. Stimuli

Stimuli were similar to those of Experiment 1, except that in most conditions, the slant specified by stereo conflicted with the information provided by skew symmetry. Four slant conflict conditions were used: -10° , -5° , 0° , and $+5^\circ$. The conflict was defined relative to a baseline condition, the 0° conflict condition, for which the stimuli were accurate binocular perspective views of planar symmetric figures. The stimuli in these baseline conditions were equivalent to those of Experiment 1. Stimuli in the $+5^\circ$ conflict condition had stereo information that suggested a slant 5° higher than the baseline, while stimuli in the -5° and -10° conditions had stereo slants 5° and 10° less than the baseline.

We kept the figural information presented in different conflict conditions as constant as possible by equating the cyclopean perspective projection of figures (projection to an imaginary point between the two eyes). Conflict stimuli generated in this way have a 3D interpretation as planar figures with the slant defined by stereo, but under this 3D interpretation, the figures

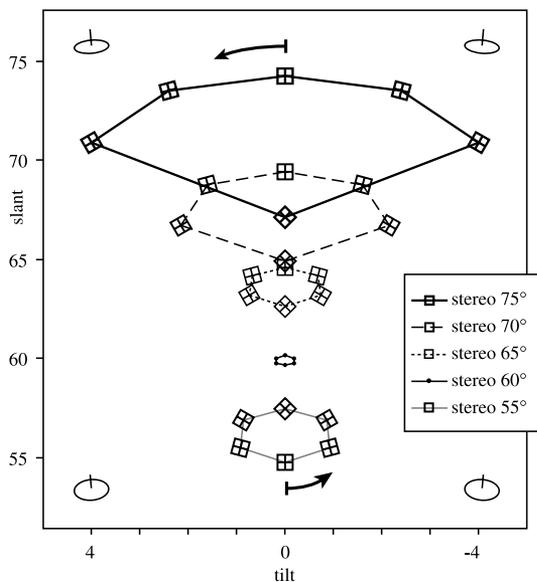


Fig. 12. Simulated biases as a function of spin when skew symmetry information is integrated with different amounts of conflicting stereo information. Each circular plot shows bias as a function of spin for a fixed stereo slant, which ranged from 55° to 75° . The arrows show the direction of bias as spin is increased from 0° . Changing stereo slant modulates the pattern of biases, affecting the tilt component of the optimal estimate in addition to the slant component. For negative stereo conflicts, the direction of bias for a given spin reverses direction.

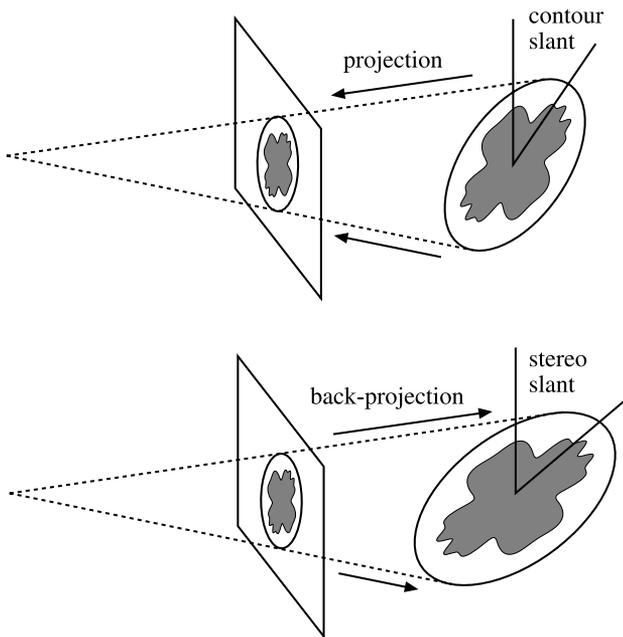


Fig. 13. Illustration of how the cue conflict stimuli in Experiment 2 were constructed. The cyclopean perspective view of a symmetric contour was computed for some slant, which we will term the contour slant. We then computed the 3D back-projection of the cyclopean view onto a plane with some other slant, which we will call the stereo slant. The final stimuli were simulated binocular views of the back-projected 3D figures. If the stereo slant and contour slant are the same, the resulting stimuli correspond to a binocular view of a mirror symmetric 3D figure. Otherwise, the stereo interpretation would be an approximately skew symmetric object in 3D.

would not be mirror-symmetric. We first determined the cyclopean projected view of a figure at the baseline slant (either 45° or 60°), then back-projected the cyclopean view onto a planar surface with the specified stereo slant (see Fig. 13). Note that when the cyclopean view was back-projected onto a surface with the same slant as the contour slant, the resulting binocular view was entirely consistent with a mirror symmetric object, as in Experiment 1. When the contour and stereo slants are different, the interpretation of the figure as indicated by stereo would be a skew symmetric 3D object rather than a mirror symmetric object.

The slant and tilt combinations tested in Experiment 2 were the same as in Experiment 1: two slants, 45° and 60° , and 13 tilts ranging from -30° to 30° in five degree steps. The stimuli included both aligned-edge and nonaligned-edge figures, presented in separate blocks. One hundred and sixty figures of each edge type were created. One figure from this larger set was randomly chosen for each slant and tilt combination, and this figure was presented at each of the spins tested. Different random samples of the 160 figures were used for each subject.

4.1.3. Procedure

The task was identical to that used in Experiment 1. Subjects adjusted a stereoscopically displayed probe line to appear perpendicular to the figure. The initial position of the probe line was chosen randomly between 30° and 50° away from the slant and tilt specified by stereo.

Subjects performed the experiment in two separate hour-long sessions, scheduled on different days. Each session consisted of a single block of either aligned-edge contours or nonaligned-edge contours, with the order of presentation counterbalanced across subjects. Within each block, symmetric contours were presented at two slants (45° and 60°), four conflict conditions (-10° , -5° , 0° , and $+5^\circ$), 13 tilts (between -30° and $+30^\circ$ in 5° steps) and six spins (0° , 15° , 30° , 45° , 60° and 75°). Each slant/tilt/spin combination was presented once within a block. Because judgments were normalized with respect to tilt direction for analysis, each subject effectively repeated each combination of the three experimental factors (figure type, slant and spin) 13 times per block.

4.1.4. Apparatus

The apparatus was identical to that of Experiment 1. Images were present in stereo on a computer monitor using shutter glasses oscillating at 120 Hz. Subjects viewed stimuli in the dark from 50.8 cm (20 inches) through a matte black occluding window.

4.2. Results

The orientation settings were analyzed as before, by first normalizing for tilt direction and then averaging across normalized normal vectors. The data from two of the 18 subjects were excluded from analysis because of excessive variability in their orientation settings. We speculate that the performance for these subjects could have resulted from some general problem with stereo vision, since we did not prescreen subjects for basic stereo ability, and stereo would be essential for perceiving the orientation of the probe line. Judgments from aligned-edge and nonaligned-edge figures were combined for the analysis. Fig. 14 shows mean normalized probe settings in slant–tilt space as a function of spin for each of the slant and conflict conditions.

Within each slant and conflict condition, the mean probe settings generally show a circular pattern of biases as a function of spin, as in experiment 1. The stimuli in the 0° conflict condition are equivalent to those of the previous experiment, so the biases in these conditions can be directly compared. Both the pattern and magnitude of biases in the no conflict condition are similar to those of Experiment 1, replicating the previous result. For conditions with positive stereo slant conflict, the magnitude of the circular pattern is larger

than in the no-conflict condition, while for negative stereo conflicts, the magnitude of the pattern is reduced. For the 45° slant stimuli with a negative 10° stereo slant conflict, the circular pattern is inverted.

The 45° and 60° slant conditions were considered separately, and independent ANOVAs were performed. For the 60° slant conditions, there was a main effect of spin ($F(11,735) = 17.03$, $P < 0.001$), a main effect of stereo slant conflict ($F(7,735) = 326.1$, $P < 0.001$), and an interaction between stereo conflict and spin ($F(29,735) = 2.687$, $P < 0.001$). For the 45° slant conditions, there was no main effect of spin ($F(11,735) = 1.751$, $P = 0.588$), but there was a significant interaction between spin and stereo conflict ($P(29,735) = 3.04$, $P < 0.001$) and a main effect of stereo slant conflict ($F(29,735) = 419.0$, $P < 0.001$). The main effects of stereo slant conflict for the 45° and 60° conditions indicate that the slant component of responses shifted when stereo slant was varied. The interactions between stereo conflict and spin reflect the fact that varying stereo slant also affected the amount of tilt bias. In the 45° slant condition, the interaction between stereo conflict and spin was sufficient to invert

the pattern of biases as a function of spin, which accounts for the absence of a significant overall effect of spin.

For both 60° and 45° slant stimuli, the effect of stereo conflicts was in the same direction relative to the consistent cues condition: exaggerating tilt biases for positive conflicts, and attenuating tilt biases for negative conflicts. We analyzed the effect of spin for each slant and conflict condition separately in independent ANOVAs, using a Bonferroni adjustment to compensate for multiple tests. Significant effects of spin were present for 60° slant stimuli with stereo conflicts of 5°, 0°, and -5° (5°: $F(11,159) = 16.7$, $P < 0.001$; 0°: $F(11,159) = 12.58$, $P < 0.001$; -5°: $F(11,159) = 2.959$, $P = 0.005$; -10°: $F(11,159) = 0.7404$, n.s.), and for 45° slant stimuli with stereo conflicts of 5°, 0°, and -10° (5°: $F(11,159) = 6.015$, $P < 0.001$; 0°: $F(11,159) = 4.919$, $P < 0.001$; -5°: $F(11,159) = 1.009$, n.s.; -10°: $F(11,159) = 2.562$, $P = 0.02$). There was an overall bias in slant settings toward overestimation, as in the previous experiment. The overestimation was larger for the 60° slant figures (mean error 11.2°, S.D. 5.7°) than for the 45° slant figures (mean error 4.2°, S.D. 7.0°).

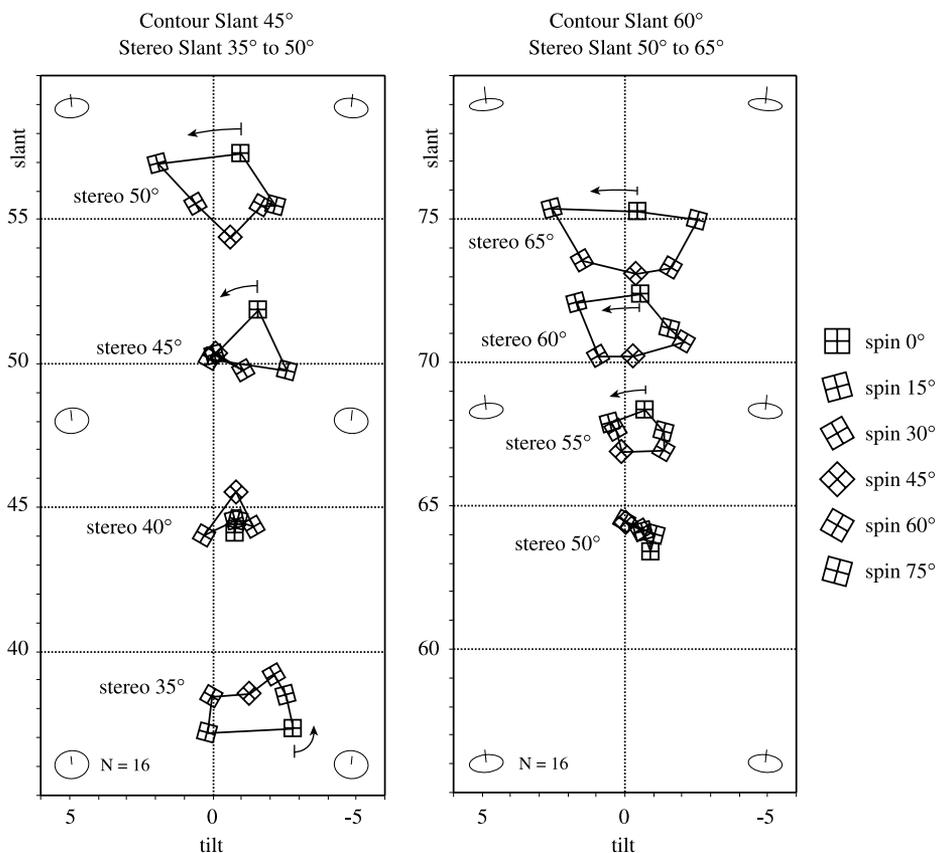


Fig. 14. Results of Experiment 2. The separate circular plots correspond to conditions with the same stereo slant but different spins. The arrows show the direction of change in bias as spin is increased from 0°.

In addition to the effects of stereo and spin, the orientation curves show an overall shift to the right, toward negative tilts. Since positive and negative spins give rise to images that are symmetric around the vertical, one would expect tilt settings for negative spins to have been the negative of tilt settings for equivalent positive spins. One would further expect that the average tilt settings for 0° and 45° spins would have been zero, since the projected symmetry axes are symmetric around the vertical for these conditions. We performed a post-hoc test for an overall tilt bias by averaging subjects' tilt settings across positive and negative spins of equal magnitude and over different slant and figure type conditions. The composite average was small (tilt = -0.34) but significantly different from zero ($t(271) = -2.51$, $P = 0.013$). An ANOVA further revealed that the amount of tilt bias varied depending on conflict conditions ($F(3,112) = 15.56$, $P < 0.001$), with smaller stereo slants producing larger biases, but was not affected by spin magnitude ($F(1,112) = 0.68$, $P = 0.40$ n.s.). These results suggest that subjects suffer from a small, absolute tilt bias in their settings and that this tilt bias is larger when stereo slant is small. One possible cause for this bias would be a perceptual bias to represent the stimuli in a coordinate frame aligned with the left eye's view.

4.3. Discussion

Orientation judgments showed circular patterns of biases as a function of spin, as in the previous experiment. Increasing the slant suggested by stereo exaggerated subjects' biases, while decreasing the slant suggested by stereo shrunk the biases, as predicted by the Bayesian model. Most notably, subjects' orientation settings in the 45° slant baseline, -10° conflict condition showed an inverted pattern of biases. The inversion of biases is perhaps the strongest distinguishing prediction of the Bayesian model. Taken together, the results fit the predictions of the Bayesian model in all qualitative aspects. Accepting the model allows us to use the pattern of results to infer the degree of slant over-estimation from cues other than skew symmetry in the cue consistent stimuli (presumably dominated by stereo). In the 60° baseline condition, the circular pattern of orientation disappears in the -10° conflict condition, suggesting that subjects' bias in slant from stereo estimates was 10° . This is consistent with predictions from previous data on the human depth from stereo biases at near viewing distances. In the 45° baseline condition, the pattern of subjects' biases compresses to a point in the -5° conflict condition and inverts in the -10° conflict condition. This suggests a positive bias in slant from stereo estimates at that slant near 5° .

5. General discussion

5.1. Skew symmetry

The two experiments were designed to test whether skew symmetry contributes to the perception of surface orientation and to explore how this information is integrated with other cues. We manipulated the information provided by skew symmetry by varying the alignment of a figure within a three-dimensional plane (its spin) and found systematic biases in orientation judgments that depended on spin. The different directions of bias can be explained in terms of the nonlinear constraints imposed by a symmetric interpretation of a contour, which change when a figure is rotated within a plane. Specifically, the tilt biases observed in Experiment 1 were in a direction consistent with moving upward along the skew symmetry constraint curves, which led us to postulate that the effect was due to the nonlinear interaction of symmetry information and a biased estimate of slant from stereo. We developed a Bayesian model of cue integration for stereo and symmetry based on this idea and found that it could reproduce the qualitative pattern of biases. To further test this model, we introduced a range of slant conflicts between stereo and symmetry information in addition to varying spin. The effect of the cue conflicts was to increase or decrease the magnitude of the spin effect, and in one case invert the effect, consistent with the predictions of the Bayesian optimal model of cue integration.

The results clearly implicate symmetry as a strong constraint for perceiving the quantitative slant and tilt of an object, even in the presence of reliable stereo information. We can conclude further that at least part of the contribution of contour information specifically involves skew symmetry. The projected contour of a planar symmetric object provides other figural cues, including the aspect ratio or foreshortening of the projected contour, and the gradient of symmetry lines produced by perspective distortion. However, the stimuli were designed to control for the possible effects of these other cues. For twofold symmetric isotropic figures, both foreshortening and perspective cues are approximately invariant to spin, so neither of these cues could produce the spin-dependent biases observed in the data. The similar pattern of biases for aligned-edge and non-aligned-edge figures further rules out explanations based on the distributions of contour edge orientations. While it is likely that other figural cues do contribute to the perception of surface orientation, for the experimental stimuli, the only expected effect of other cues would be to attenuate biases due to skew symmetry, as demonstrated in our model simulations. Skew symmetry appears to be the only remaining factor that could

explain the presence of spin-dependent biases in these conditions.

Skew symmetry is available as a cue for 3D orientation whenever an observer views a surface that is planar and symmetric. Many architectural features are approximately planar and symmetric (windows, doors, tables), and volumetric objects often contain a planar symmetric face (particularly among man-made objects). Thus, this cue would be applicable for a large range of objects. Our geometric analysis does not apply to objects that are symmetric in three dimensions but are not planar and do not contain planar facets (objects with smoothly curved surfaces). In these cases, the bounding contours of objects are generally not skew-symmetric. While a symmetry constraint may play an important role in the perception of such objects, it is not mediated by skew as analyzed here.

5.2. Bayesian cue integration

The results reported here provide insight into how different sources of 3D information are combined. In an earlier section, we contrasted a linear cue integration model, which computes a weighted average between estimates from different sources, with a probabilistic Bayesian model, which makes full use of information provided by noisy or ambiguous cues. We noted that in some cases, these models converge, but that when substantial nonlinearities are present, the predictions of the two models differ. The use of skew symmetry as information for surface orientation provides an interesting test case for the study of cue integration. When considered as a constraint on slant and tilt, skew symmetry specifies a curve of possibilities, rather than a single estimate, so the constraint is both ambiguous and nonlinear. Consequently, linear and optimal models of integrating skew symmetry with other cues make different predictions regarding the effect of conflicting information.

Experiment 2 was a direct test of the contrasting predictions of a linear and optimal integration strategies. The results provide clear evidence in favor of the nonlinear, optimal model. Varying stereo slant affected both slant and tilt components of judgments, depending on the spin of a contour. In contrast to a linear model, a Bayes optimal model accounts for the qualitative results data with relatively few reasonable assumptions. Both the basic spin effect in Experiment 1 and the modulation of this effect in Experiment 2 can be explained as direct consequences of the orientation and shape of the skew-symmetry constraint curves. A crucial aspect of our cue integration model is that the information from skew symmetry is represented by a likelihood distribution rather than a single estimate, incorporating the full probabilistic information. The success of the model in explaining the biases suggests that the process of interpreting skew-symmetric contours is sensitive to the range

of ambiguity of symmetry information and is not separate from the interpretation of other sources of information.

The Bayesian model provides a functional model of performance. What we have shown is that subjects' biases in estimating the orientation of skew symmetric figures follow the qualitative behavior of a Bayes optimal observer. We emphasize, however, that the representation of figural information and the particular algorithm that we use to simulate the Bayes optimal observer need not correspond directly to the process used by the visual system. In particular, our model assumes a front-end process that extracts noisy estimates of a figure's skew symmetry axes. The estimated skew axes, along with stereo disparity information, serve as the inputs to the Bayesian estimator. An alternative formulation would be one in which orientation estimates are derived by simultaneously maximizing the fit of estimated surface orientation with the stereo disparity data and maximizing a measure of the symmetry of the inferred figure in three dimensions. This formulation provides an alternative path to formulating a Bayesian observer, one that uses all of the available figural information rather than compressing it into a representation only of the skew axes. Since the geometric basis of the latter formulation is the same as our formulation, the qualitative predictions of the model should be the same; that is, the qualitative predictions of Bayesian estimators that use a symmetry constraint will be the same regardless of the assumed input to the estimators. Thus, our data do not implicate a particular figural representation or estimation process, but rather constrain process models of skew symmetry interpretation to a family of Bayes optimal observers. More particularly, the data reveal a cooperative cue integration strategy, rather than a strategy in which the outputs of modular estimators of surface orientation are combined (what Maloney & Landy refer to as a modified weak fusion model (Landy et al., 1995)).

In this paper, we have focused on the contour information provided by skew symmetry. We are currently following up this work with studies of how other contour cues interact with stereo and with one another. Specifically, we are investigating the foreshortening and perspective distortion of symmetric contours as potential cues for surface orientation. Our model can be easily extended to incorporate information from these cues, which leads to additional testable predictions. For example, cue conflicts can be introduced by varying the anisotropy of symmetric figures, analogous to varying stereo in Experiment 2. An optimal model predicts similar modulation of biases. In general, the results reported here establish that skew symmetry contributes to perception of 3D orientation, and further suggest that the visual system makes full use of the ambiguous information provided by skew symmetry when integrating skew symmetry with other cues.

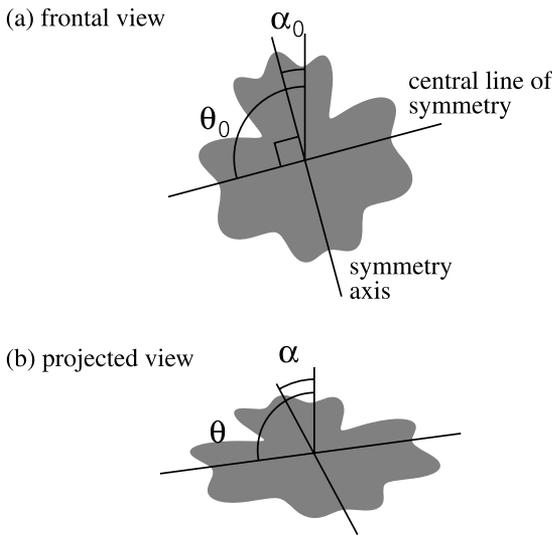


Fig. 15. (a) Frontal view of a mirror symmetric object. The axis of symmetry and lines of symmetry meet at right angles. (b) Same object viewed from an angle. The projected axis of symmetry and the projected lines of symmetry no longer meet at right angles.

Acknowledgements

We would like to thank Jeff Hall and Susan Damske for their work collecting data. This work was supported by NIH grant EY09383, and was begun while Jeff Saunders was a fellow on and NEI training grant to the University of Pennsylvania. Part of this work was presented at the annual meeting of ARVO in 1999.

Appendix A. Symmetry constraint

In this appendix, we describe more explicitly the constraint that an assumption of mirror symmetry imposes on the possible slant and tilt of a skew symmetric contour. Fig. 15 shows frontal and projected views of a symmetric contour, with labels representing the orientation of the symmetry axis and lines before (α_0 and θ_0) and after projection (α and θ). The angle between α_0 and the tilt direction is what we have been referring to as spin.

If one uses orthographic projection to approximate perspective projection, the effect of 3D viewing orientation is to compress the projected contour uniformly in the direction of tilt. This can be expressed mathematically by the following equation, which relates the axis and symmetry line orientation in a frontal view (α_0 and θ_0) to the corresponding orientation in the projected image (α and θ) when viewed with slant σ and tilt τ :

$$\begin{aligned} \tan(\alpha - \tau) &= \cos(\sigma)\tan(\alpha_0 - \tau) \\ \tan(\theta - \tau) &= \cos(\sigma)\tan(\theta_0 - \tau). \end{aligned} \tag{A1}$$

An assumption of mirror symmetry corresponds to the constraint that the symmetry lines and axis meet at right angles in a frontal view, which can be expressed as $\theta_0 = \alpha_0 + 90^\circ$. Incorporating this assumption, the equations can be written in terms of spin:

$$\begin{aligned} \tan(\alpha - \tau) &= \cos(\sigma)\tan(\omega) \\ \tan(\theta - \tau) &= \cos(\sigma)\tan(\omega + 90^\circ) \end{aligned} \tag{A2}$$

where ω is the spin of a contour. Using the fact that $\tan(\omega + 90^\circ) = -1/\tan(\omega)$, Eqs. (A1) and (A2) can be combined to yield a relation between the skew axes in an image and the possible slant/tilt combinations for a mirror symmetric interpretation:

$$\tan(\alpha - \tau)\tan(\theta - \tau) = -\cos^2(\sigma). \tag{A3}$$

For the case of orthographic projection, the relation in Eq. (A3) fully describes the constraint that an assumption of mirror symmetry imposes on the interpretation of a skew-symmetric contour. The same relationship applies to perspective projection when applied to the projected axis of symmetry and any given symmetry line, when the slant and tilt are specified relative to the line of sight to the point of intersection between the two lines. In particular, for the twofold symmetric figures used in the experiments, Eq. (A3) holds for the orientations of the projected axes of symmetry and the slant and tilt of the surface relative to a line of sight through the center of mass of the 3D figure.

To derive the constraint curves for the skew symmetry constraint (shown earlier in Fig. 3), we first used Eq. (A2) to compute the orientations of projected symmetry lines and symmetry axis, θ and α , and then used Eq. (A3) to find other slant/tilt pairs consistent with these skew axes. Eq. (A2) was also the basis of the generative model used to compute likelihood distributions for skew symmetry information. We assumed that image measurements of α and θ varied randomly around the orientations geometrically specified by Eq. (A2) given the slant, tilt and spin of a figure. Assuming a noise model (Gaussian) one can derive $p(\alpha, \theta | \sigma, \tau, \omega)$. The likelihood over slant and tilt is computed by integrating out the spin variable and applying Bayes' rule (see text).

Appendix B. Perspective information

In this appendix, we describe the information conveyed by the perspective distortion of a projected contour. Under orthographic projection, symmetry lines are parallel in the image and can be characterized by a single orientation. In a perspective projection, lines of symmetry are not parallel but rather converge to a point on the horizon. Thus, the orientations of projected symmetry lines are a function of position in an image, $\theta(x)$. For small regions, the function $\theta(x)$ can

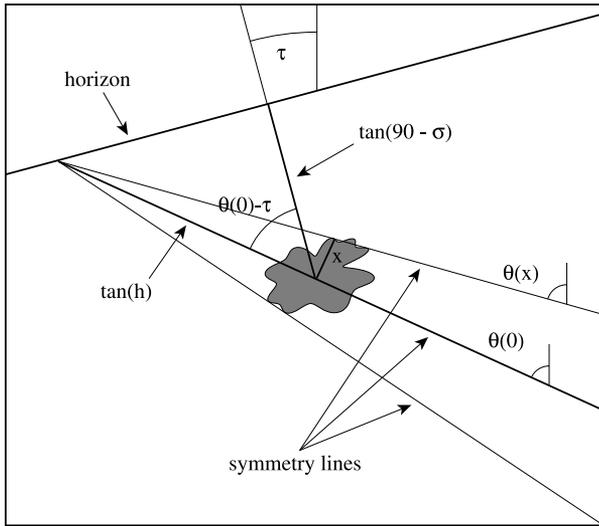


Fig. 16. Illustration of the geometry of symmetry lines in a perspective view. Projected symmetry lines converge to a point on the horizon. Their orientations vary as a function of position in the image, represented by $\theta(x)$. The distance to the horizon along the direction $\theta(0)$ is related to the slant (σ) of the figure and the difference between $\theta(0)$ and the tilt direction (τ).

be characterized by two parameters: the orientation of a central symmetry line, $\theta(0)$, and the rate of change of the orientation of symmetry lines, $d\theta/dx$ (Fig. 16). These two parameters can be directly computed for a figure given its slant, tilt (relative to the line of sight through the center of mass of a figure) and spin.

As noted previously, the constraint on skew axes for orthographic projection applies to the perspective projection of the symmetry axis and the central line of symmetry:

$$\begin{aligned} \tan(\theta(0) - \tau) &= \cos(\sigma)\tan(\omega + 90^\circ) \\ &= \cos(\sigma)/\tan(\omega) \end{aligned} \tag{A4}$$

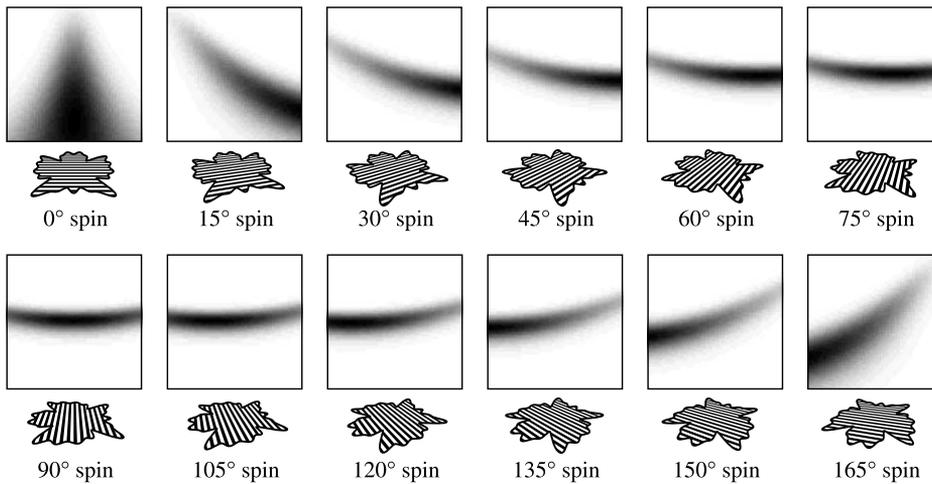
where σ is slant, τ is tilt, and ω is spin.

The rate of change in the orientation of symmetry lines, $d\theta/dx$, is inversely related to the distance to the horizon:

$$d\theta/dx = 1/\tan(h) \tag{A5}$$

where h is the difference in visual angle between the reference point and the horizon (Fig. 16). This distance is related to the slant and tilt of the figure:

(a) Symmetric



(b) Bisymmetric

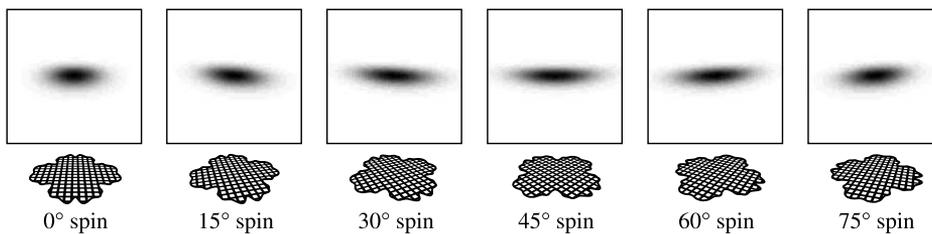


Fig. 17. Likelihood functions for perspective information provided by a single set of symmetry lines (a) and by a pair of symmetry lines (b) for a figure with 60° slant and vertical tilt, with varied spin. The functions in (a) would represent perspective information for a figure with a single axis of symmetry, while the functions in (b) represent perspective information for a bi-symmetric figure. The lines within the figures depict symmetry lines when viewed in perspective. The amount of perspective distortion has been exaggerated for illustration. The axes for the functions (not shown) are $[-25^\circ, 25^\circ]$ for tilt and $[35^\circ, 85^\circ]$ for slant. Note that the likelihood functions from a pair of symmetry lines (b) show little variation as a function of spin, in contrast to the dramatic differences for likelihood functions from a single set of symmetry lines (a).

$$\tan(h) \cos(\theta(0) - \tau) = \tan(90^\circ - \sigma). \quad (\text{A6})$$

Combining the previous equations and applying trigonometric identities yields an equation for $d\theta/dx$ solely in terms of slant, tilt and spin:

$$d\theta/dx = \frac{\tan(\sigma)\tan(\omega)}{\sqrt{\cos^2(\sigma) + \tan^2(\omega)}}. \quad (\text{A7})$$

From Eq. (A7), we can see that $d\theta/dx$ is zero when spin is 0° , corresponding to horizontal projected symmetry lines when the axis of symmetry is aligned with the direction of tilt, and is maximum when spin is 90° , corresponding to the case in which symmetry lines converge to the closest point on the horizon (e.g. receding lines in a one-point perspective painting).

To derive likelihood directions for perspective information, we assume noise in the measurements of $\theta(0)$ and $d\theta/dx$ to get $p(\theta(0), d\theta/dx|\sigma, \tau, \omega)$, and then integrate over the spin variable. Fig. 17a shows the likelihood distributions from perspective cues for different spins. The free parameters for this computation are the noise functions for $\theta(0)$ and $d\theta/dx$, which, for these simulations, we assumed to be Gaussians with widths 5° and 15° , respectively (this level of noise for $d\theta/dx$ corresponds to about a 2° slant discrimination threshold at 60° slant). The likelihood functions change dramatically as a function of spin, but cycle after 180° of spin rather than after 90° of spin in the case of skew-symmetry information.

For twofold symmetric figures, there are two sets of symmetry lines, which effectively have spins 90° apart. The perspective information conveyed by the two sets of symmetry lines is roughly complementary, because the constraint curves intersect transversely. Fig. 17b shows likelihood distributions for perspective information for twofold symmetric contours, using the same noise assumptions as in Fig. 17a. Although spin has a large effect on the information provided by a single set of symmetry lines, the combined information from a pair of symmetry lines is less affected. There is a difference in shape of the combined likelihood distributions, which can be made more or less pronounced by varying the amount of noise in $d\theta/dx$. In conditions in which perspective information is weak (e.g. when visual angle of projected contour is small), the uncertainty in $d\theta/dx$ would be high, resulting in perspective likelihood distributions that are approximately invariant to spin. In particular, there would be little ‘slope’ in the dense portions of the likelihood function.

The information provided by perspective distortion and skew symmetry depend on conceptually independent assumptions: in the former case that symmetry lines are parallel, and in the latter case that symmetry lines and the symmetry axis meet at right angles. However, when computing likelihood functions from combined sources, the two cues cannot be treated as independent, since both skew and perspective information rely on measurement

of the orientation of symmetry lines. Thus, to simulate the use of both symmetry cues, we integrated a generative model for both skew and perspective parameters, $P(\theta_1(0), \theta_2(0), d\theta_1/dx, d\theta_2/dx|\sigma, \tau, \omega)$.

Acknowledgements

We would like to thank Jeff Hall and Susan Danaske for their work collecting the data. This work was supported by NIH grant E709383, and was begun while Jeff Saunders was a fellow on an NEI training grant to the University of Pennsylvania. Part of this work was presented at the annual meeting of ARVO in 1999.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychology Reviews*, *61*, 183–193.
- 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.
- Clark, J. J., & Yuille, A. L. (1990). *Data fusion for sensory information processing systems*. Boston, MA: Kluwer Academic Press.
- Foley, J. M. (1980). Binocular distance perception. *Psychological Review*, *87*(5), 411–434.
- Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research*, *31*(7–8), 1351–1360.
- Kanade, T. (1981). Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, *17*, 409–460.
- King, M., Meyer, G. E., Tangney, J., & Biederman, I. (1976). Shape constancy and a perceptual bias towards symmetry. *Perception & Psychophysics*, *19*(2), 129–136.
- Kohler, W. (1938). *The place of value in a world of facts*. New York: Liveright.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, *35*(3), 389–412.
- Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.
- Liu, Z., Knill, D. C., & Kersten, D. (1995). Object classification for human and ideal observers. *Vision Research*, *35*(4), 549–568.
- McBeath, M. K., Schiano, D. J., & Tversky, B. (1997). Three-dimensional bilateral symmetry bias in judgments of figural identity and orientation. *Psychological Science*, *8*(3), 217–223.
- Mach, E. (1897). *Contribution to the analysis of the sensations* (C. M. Williams, Trans.). Chicago, IL: Open Court original work published in 1890.
- Perkins, D. N. (1972). Visual discrimination between rectangular and nonrectangular parallelepipeds. *Perception & Psychophysics*, *12*(5), 396–400.
- Perkins, D. N. (1976). How good a bet is good form? *Perception*, *5*(4), 393–406.
- Stevens, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence*, *17*, 47–73.
- Vetter, T., & Poggio, T. (1994). Symmetric 3D objects are an easy case for 2D object recognition. *Spatial Vision*, *8*(4), 443–453.
- Wagemans, J. (1992). Perceptual use of nonaccidental properties. *Canadian Journal of Psychology*, *46*(2), 236–279.
- 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. (1995). Detection of visual symmetries. *Spatial Vision*, *9*(1), 9–32.