

Detection of skewed symmetry

Tadamasa Sawada

Department of Psychological Sciences, Purdue University,
West Lafayette, IN, USA



Zygmunt Pizlo

Department of Psychological Sciences, Purdue University,
West Lafayette, IN, USA



This study examined the ability of human observers to discriminate between symmetric and asymmetric planar figures from perspective and orthographic images. The first experiment showed that the discrimination is reliable in the case of polygons, but not dotted patterns. The second experiment showed that the discrimination is facilitated when the projected symmetry axis or projected symmetry lines are known to the subject. A control experiment showed that the discrimination is more reliable with orthographic, than with perspective images. Based on these results, we formulated a computational model of symmetry detection. The model measures the asymmetry of the presented polygon based on its single orthographic or perspective image. Performance of the model is similar to that of the subjects.

Keywords: symmetry, skewed symmetry, planar figures, perspective projection, orthographic projection

Citation: Sawada, T., & Pizlo, Z. (2008). Detection of skewed symmetry. *Journal of Vision*, 8(5):14, 1–18, <http://journalofvision.org/8/5/14/>, doi:10.1167/8.5.14.

Introduction

An ability to detect symmetric objects is important because many natural and man-made objects in the real world are symmetric or approximately symmetric. More specifically, most of these symmetric objects are mirror (bilaterally) symmetric. There are also other types of symmetry: rotational and translational (Mach, 1906/1959). It has been argued that mirror symmetry is detected by humans more efficiently (shorter reaction times and smaller proportion of errors) than the other types of symmetry (Wagemans, 1997). In this paper, we will use “symmetry” to mean “mirror symmetry.”

The human visual system can reliably detect symmetry on the retina, especially when the axis of symmetry is vertical and the stimulus is projected to the center of the retina (e.g., Barlow & Reeves, 1979; Julesz, 1971). This result is attributed by some authors to the symmetric structure of the brain; the cortical activation in the left hemisphere, produced by the left half of a symmetric retinal stimulus, is compared, point by point, to the corresponding activation in the right hemisphere, which is produced by the other half of the retinal stimulus (Herbert & Humphrey, 1996; Julesz, 1971; Mach, 1906/1959). Obviously, this mechanism cannot be the only one for symmetry detection because human observers are able to detect symmetry when the axis is not vertical or the stimulus is not in the center of the retina. The human visual system is likely to use multiple mechanisms for detecting symmetry and one of them is tuned to vertical symmetry at the fovea. This would explain more reliable performance with such stimuli. The operation of multiple mechanisms can explain apparent

contradictions in the literature on human symmetry perception (for a discussion of this issue; see Herbert & Humphrey, 1996; Wagemans, 1997).

Note that a symmetric object “out there” produces a symmetric retinal image only for a small set of viewing directions, but human observers seem have little difficulty in determining whether a given retinal image was produced by a symmetric object. Empirical results, supporting this claim, exist for the case of 2D figures slanted relative to the observer. The asymmetric retinal image produced by a slanted symmetric 2D figure is called *skewed symmetry* (Kanade, 1981; Kanade & Kender, 1983). Originally, the term skewed symmetry was used exclusively in the context of an orthographic projection. An orthographic projection is an approximation to a perspective projection, which is the correct model for the formation of images in the eye or camera. This approximation is good when the object is small compared to the viewing distance. More precisely, an orthographic approximation to a perspective projection is good when the range in depth of the object or figure is small relative to the viewing distance. In practice, it is usually assumed that “small” means less than 10%. Despite the fact that retinal images of symmetric figures are almost never themselves symmetric, skewed symmetry received much less attention in the prior psychophysical research than the case of symmetric images. Before we discuss prior psychophysical research on skewed symmetry, we briefly review the relevant geometry.

The 3D orientation of a 2D figure is characterized by three angles: slant (σ), tilt (τ), and roll (ρ). Slant is the angle between the observer’s line of sight and the normal to the plane of the figure. Slant ranges between 0 and

90 deg. Tilt is the angle between the projection to the image plane of the normal of the plane of the figure and the x -axis of the image plane. Tilt specifies the axis of rotation, around which the plane is rotated in depth; the axis of rotation is in the image plane and it is orthogonal to tilt. Tilt ranges between 0 and 360 deg. Roll is the angle of rotation of a 2D figure about the normal to the plane of the figure. Psychophysical experiments on skewed symmetry are usually done by means of computer graphics. Namely, instead of using physical figures slanted in depth, the subject is presented with perspective images of the figures shown on a computer screen. If the observer's eye is placed at the center of perspective projection that was used to compute the perspective images, the retinal image in the observer's eye produced by the perspective image of a slanted figure is itself a perspective image of the slanted figure. Let the line of sight of the subject be parallel to the z -axis of the 3D Cartesian coordinate system. Let $z = 0$ be the plane of the computer screen, called an image plane, and the x - and y -axes of the 3D Cartesian coordinate system be used as the 2D coordinate system on the image plane. Let $(C_x, C_y, C_z)^T$ be the center of perspective projection ($C_z \neq 0$) (recall that $(C_x, C_y, C_z)^T$ is a column vector obtained by transposing (C_x, C_y, C_z)). When $\sigma = 0$ deg, the plane of the figure is parallel to the computer screen and is represented by the following equation $z = z_f$, where z_f is constant (z_f can be set to zero, without restricting generality). Let the 2D figure be represented by a set of points $(x_{2D}, y_{2D}, z_f)^T$. Note that all these points have the same z -coordinate when $\sigma = 0$ deg. When slant σ is not zero, the 3D coordinates of each point of the simulated 2D figure can be computed as follows:

$$\begin{pmatrix} x_{3D} \\ y_{3D} \\ z_{3D} \end{pmatrix} = \mathbf{R}_z(\tau) \cdot \mathbf{R}_y(\sigma) \cdot \mathbf{R}_z(-\tau) \cdot \mathbf{R}_z(\rho) \cdot \begin{pmatrix} x_{2D} \\ y_{2D} \\ z_f \end{pmatrix} \quad (1)$$

$$\begin{cases} \mathbf{R}_y(\theta) = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix} \\ \mathbf{R}_z(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{cases}$$

Now that the simulated 2D figure is slanted in the 3D space, its projection (orthographic or perspective) on the image plane will be computed. Consider perspective projection first. A perspective projection of a point $(x_{3D}, y_{3D}, z_{3D})^T$ in a 3D space to the 2D image plane (computer screen) can be computed as follows (see Figure 1):

$$\begin{cases} x_p = (x_{3D} - C_x) \cdot C_z / (C_z - z_{3D}) + C_x \\ y_p = (y_{3D} - C_y) \cdot C_z / (C_z - z_{3D}) + C_y \end{cases}, \quad (2)$$

where $(x_p, y_p)^T$ is a perspective image of $(x_{3D}, y_{3D}, z_{3D})^T$.

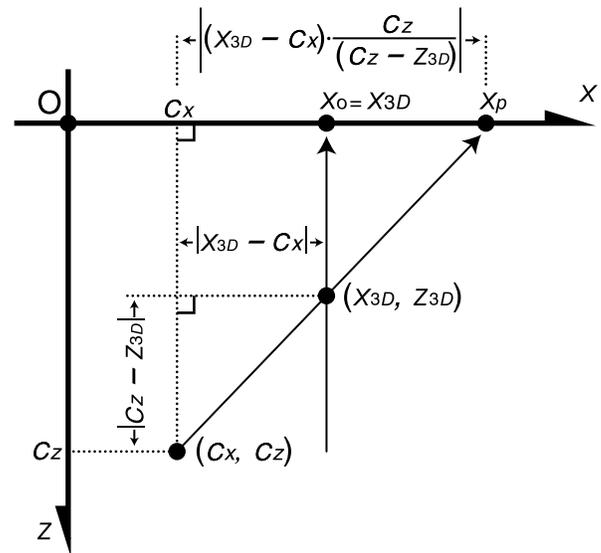


Figure 1. Top view of the viewing geometry.

Orthographic projection is different from perspective projection in that the projecting lines are all parallel to each other and perpendicular to the image plane. This can be modeled by moving the center of perspective projection to infinity ($C_z \rightarrow \infty$). In such a case, Equation 2 takes the following form:

$$\begin{cases} x_o = x_{3D} \\ y_o = y_{3D} \end{cases}, \quad (3)$$

where $(x_o, y_o)^T$ is an orthographic image of $(x_{3D}, y_{3D}, z_{3D})^T$.

It is important to emphasize that the observer's eye must be placed at the center of the perspective projection $(C_x, C_y, C_z)^T$ that was used to compute the perspective images. Only then will the retinal image in the observer's eye be a valid perspective image of the simulated 2D figure slanted in the 3D space. Otherwise, the retinal image will be a composition of two perspective projections, which is a *projective* not a perspective transformation of the simulated 3D figure (Coxeter, 1987; Pizlo, Rosenfeld, & Weiss, 1997a, 1997b; Wagemans, Lamote, & van Gool, 1997). If an orthographic approximation to perspective projection is used, then the position of the observer's eye is irrelevant, as long as the line of sight is orthogonal to the computer screen. In such a case, the retinal image will also be an orthographic transformation (up to size scaling) of the simulated 2D figure. Again, the orthographic approximation to perspective projection is good when the viewing distance is large as compared to the range in depth of the simulated 2D figure.

Next, consider the relation between skewed symmetry and symmetry for the case of orthographic and perspective projections. In a symmetric figure, the symmetry line

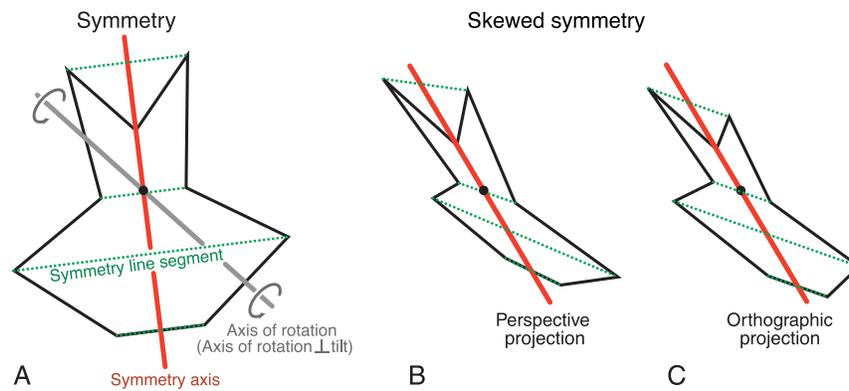


Figure 2. Perspective and orthographic projections of a symmetric polygon. Slant is 69 deg, and tilt is 228 deg.

segments, i.e., the line segments that connect symmetric points, are parallel to one another (see Figure 2A). The symmetry axis is perpendicular to the symmetry line segments and it bisects these segments. Some or all of these properties are changed in skewed symmetry. First, consider skewed symmetry produced by an orthographic projection (Figure 2C). In this case, projected symmetry line segments are still parallel, and their midpoints are on the projected symmetry axis. That is, the parallelism of symmetry line segments and collinearity of their midpoints are invariants of orthographic projection. However, the projected symmetry axis is not perpendicular to the projected symmetry lines. In skewed symmetry produced by a perspective projection, projected symmetry line segments are not parallel (Figure 2B). Instead, they intersect at a vanishing point. Furthermore, the projected symmetry axis is not a bisector of the symmetry line segments.

Most of the prior research on human perception of skewed symmetry involved dotted or textured stimuli. Attneave (1982), based on his informal observations of dotted stimulus, claimed that skewed symmetry cannot be detected reliably by the human visual system, but this claim was later shown to be inaccurate. Wagemans (1992, 1993) tested symmetry detection for both dotted patterns and contours in the case of symmetry and skewed symmetry produced by an orthographic projection. He found that performance was more reliable with contours than with dots and more reliable with symmetric than with skewed symmetric images. The superiority of contours did not receive much attention in the past, but contours of objects are at least as important as surface texture in representing the shapes of objects. In subsequent studies, Wagemans, Van Gool, and d'Ydewalle (1991, 1992) tested the role of symmetry lines and symmetry axis in detection of skewed symmetry of dotted patterns. They used skewed symmetry produced by orthographical projection and showed that skewed symmetry is detected more reliably when the orientation of the projected symmetry lines is horizontal and known to the subject, or when the

orientation of the projected symmetry axis is vertical and known to the subject. van der Vloed, Csathó, and van der Helm (2005) used perspective rather than orthographic projection. However, because the actual viewing distance (185 cm) was different from the simulated distance (10 cm) that was used to compute the perspective images on a computer screen, the retinal images in the subject's eye were actually projective, not perspective transformations of the symmetric figures. They found that performance in symmetry detection experiment was negatively correlated with slant: the larger the slant, the worse the performance. It is not clear as to whether this result applies to the case of skewed symmetry produced by perspective projection because the retinal images in that study were projective not perspective transformations of the symmetric patterns. Locher and Smets (1992) used real objects to test detection of symmetry for figures slanted in depth. Because real objects were used, the retinal images of the slanted figures were always perspective images of the figures. They showed that subjects' performance did not depend on slant. In particular, the performance for slanted figures was as good as for figures with slant zero. This result suggests that detection of skewed symmetry is as easy as detection of symmetry. However, the generality of their result is unclear because the axis of symmetry and the axis of rotation always coincided, and the viewing distance (95 cm) was large compared to the size of the figures (6.5 cm). Under these conditions, the perspective projection is essentially identical to orthographic projection (up to size scaling), and more importantly, skewed symmetric figures in the retinal image were themselves approximately symmetric. As a result, the lack of the effect of slant was to be expected because for all slants the subject was faced with the test of detecting symmetry, rather than skewed symmetry.

The purpose of our study was to provide further tests of human performance in detecting planar (2D) skewed symmetric figures and to develop a computational model of this perceptual ability. Specifically, we tried to replicate known results using somewhat different experimental

procedure and add new results that are needed to formulate a model and then to test it. In particular, (i) we replicated previous experiments that used either orthographic or projective images. We used perspective images, which represent the actual transformation from the 3D scene to the observer's retina; (ii) we used short exposure duration with controlled eye fixation. If the eye fixation is not controlled and the observer is allowed to make eye movements, the computational model of perception should include a model of the oculomotor strategy, which tends to be idiosyncratic; (iii) we directly compared performance using perspective and orthographic projection in order to verify which invariants are used in skewed symmetry detection; finally (iv) we tested the subjects' reconstruction of symmetry. Namely, we asked a subject to produce a symmetric polygon that is equivalent to the percept produced by a skewed symmetric polygon on the retina. Results of this last experiment are used in the model to "undo" the distortions of symmetry that were produced by (can be attributed to) projection from a polygon "out there" to the retinal image. This way, when our computational model measures asymmetry of a polygon, it measures only this part of asymmetry that cannot be attributed to the projection. It follows that our model (and the subjects) are detecting symmetry of a distal stimulus or skewed symmetry of a proximal stimulus (the retinal image). The model was tested using the same stimuli as those used in our psychophysical experiments. The simulation results allow us to conclude that the model is a plausible explanation of the underlying perceptual mechanisms.

Experiment 1: Symmetry detections for dotted stimuli and polygons

Method

Subjects

Four subjects (including the two authors TS and ZP) were tested. All subjects had prior experience as subjects in psychophysical experiments. TS and ZP received extensive practice before being tested. OK and YL were naive about the purpose of the experiment. All subjects had normal or corrected-to-normal vision.

Apparatus

The stimuli were shown on an LCD monitor with 1280×1024 resolution and 60 Hz refresh rate. The subject viewed the monitor with the right eye from a distance of 40 cm in a dark room. The subject wore an eye-patch on his left eye. A chin rest and forehead rest was used to support the subject's head.

Stimuli

Polygons were generated using a method similar to that used by Pizlo and Salach-Golyska (1995) in their Experiment 1. Vertices of the polygons were generated using polar co-ordinates r and θ . Each polygon had 9 to 12 vertices. For an asymmetric polygon, the radius r of each vertex was random in the range between 3.5 and 14.1 cm, and its orientation θ was random in the range between 0 and 360 deg. The vertices were connected by a polygonal line in the order of increasing value of θ . By doing this, we ensured that the polygon did not produce self intersections. For a symmetric polygon, half of the vertices were generated first in the same way as for an asymmetric polygon, but the orientations were restricted to the range between 0 and 180 deg. Next, mirror reflections of these points about the horizontal axis were generated and all vertices were connected as before. If the number of vertices was odd, one of the vertices was placed on the horizontal axis. If the number of vertices was even, none or two of them were placed on the horizontal axis. The origin of the polar coordinate system was placed at the center of the monitor. The dotted stimuli were generated the same way as the polygon stimuli, and the dots were placed at the vertices of the polygon. The stimulus occupied an area whose radius was at most 19.4 deg. We used large stimuli to make sure that perspective effects are clearly noticeable (see an example in Figure 2). For small stimuli, perspective projection becomes indistinguishable from orthographic projection. Note that if the fixation point were at the center of the figure, part of the stimulus would have been projected to the optic disk of the eye and thus remain invisible. To avoid this problem, the fixation point was shifted 9 deg to the right from the center of the monitor. The line of sight connecting the fixation point and the viewing eye was perpendicular to the monitor.

Perspective projection with a random tilt and random slant was used to produce the images. The center of rotation of the 2D figure was placed at the origin of polar coordinates used for generating the stimuli. Slants between 50 and 70 deg were used. Slants smaller than 50 deg were not used because they only produce slight distortions of symmetry. Slants greater than 70 deg were not used because the resulting perspective images are very narrow and many details are invisible. Roll depended on the condition. In the first condition, roll was equal to tilt. In this case the angle between the symmetry axis and the tilt was 0 deg. In such cases, the perspective image of a symmetric figure was almost exactly symmetric. It would have been exactly symmetric, if the plane defined by the center of perspective projection and the symmetry axis had been perpendicular to the computer monitor (image plane). This was not guaranteed in our experiment because the fixation point did not coincide with the center of the monitor. Note, however, the departures from perfect symmetry were very small in this case (see an example in

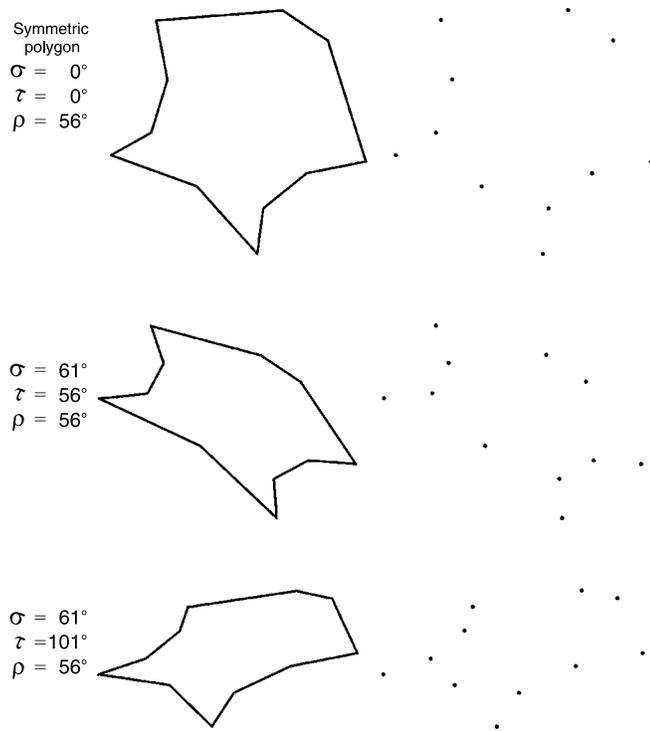


Figure 3. The polygon (left) and the corresponding dotted stimuli (right). The dots were placed at the vertices of the polygon. The symmetric polygon is shown on the top, and its perspective images are shown in the middle and on the bottom. Numbers on the left are values of slant (σ), tilt (τ), and roll (ρ). When tilt is equal to roll, the retinal image of a symmetric polygon is itself symmetric.

Figure 3). In the other two conditions, the perspective image of a symmetric figure was clearly asymmetric. In one of these two conditions, the angle between roll and tilt was ± 45 deg. As a result, the angle between the projected symmetry axis and the tilt was 45 deg. In such a case, the distortion of symmetry is maximal. In the third, random condition, roll was random in the range from 0 deg to 360 deg. As a result, the degree of distortion of symmetry was random, as well. The center of perspective projection was at the subject's right eye. As a result, the retinal image of the stimulus on the computer screen was a perspective image of the simulated 2D figure slanted in depth. Examples of the stimuli are shown in Figure 3.

Three constraints for the perspective image of a polygon were used to make sure that all details of each polygon were visible: the length of the polygon's side had to be greater than 20 pixels (0.59 cm), the distance of each vertex from other vertices and sides had to be greater than 20 pixels, and the angle at the vertices could not be too close to 0 deg or 180 deg. Specifically, the angle had to be within the range 10–170 deg or 190–350 deg. If the perspective image of a polygon violated any of these constraints, another polygon was generated. The polygons and dots were drawn in white on a dark background with high contrast. The width of the polygon side was 4 pixels (0.12 cm), and the diameter of the dots was 8 pixels (0.24 cm).

Procedure

The method of signal detection was used. Each session consisted of 200 trials: 100 trials with symmetric stimuli and 100 trials with asymmetric stimuli, presented in a random order. There were 6 experimental conditions: two types of figures (polygon vs. dots) \times three conditions represented by the angle between the axis of symmetry and the tilt (0 deg vs. 45 deg vs. random). All conditions were blocked in each session. The subject ran two sessions for each condition. The order of sessions was randomized.

Each trial began with a fixation cross. After pressing the mouse button, the fixation cross disappeared and the stimulus was shown for 100 ms. The subject's task was to respond whether or not the stimulus on the computer monitor was skewed symmetric. That is, whether or not the stimulus was produced by a symmetric figure slanted relative to the subject. After each trial, the subject received feedback about the accuracy of the response. The subject's performance was evaluated by the discriminability measure d' and its standard error. The standard error was computed from two values of d' (recall that there were two sessions per each condition). Higher performance corresponds to higher values of d' . Chance performance is represented by $d' = 0$. Perfect performance is represented by $d' = \infty$. The subject ran a number of practice sessions to become familiar with the experiment and stimuli. Each session started with a block of 16 practice trials. After the practice trials, the subject was informed about the proportion of correct responses and was given an option to repeat it. This option was rarely exercised. Before each session, the subject was told which condition would be tested.

Results and discussion

Results of individual subjects are shown in Figure 4, and the averaged results are shown in Figure 5. The ordinate shows d' .¹ The results were analyzed using a two-way ANOVA within-subjects design: figure type (polygon vs. dots) \times angle between the symmetry axis and the tilt direction (0 deg vs. 45 deg vs. random).

It can be seen that the subject's performance was substantially higher in polygon than in dots conditions ($F_{1, 15} = 79.08$, $p < 0.001$). In fact, with dotted stimuli, performance was close to chance level for 45 deg and random angle conditions (d' is comparable to the standard errors). This means that subjects could not detect skewed symmetry in the case of perspective images of dotted patterns, although they could do this reliably in the case of polygons. The main effect of the angle between the axis of symmetry and the tilt was also significant ($F_{2, 15} = 59.42$, $p < 0.001$). Specifically, performance in 0 deg condition was substantially higher than that in 45 deg and random angle conditions. This means that symmetry on the retinal

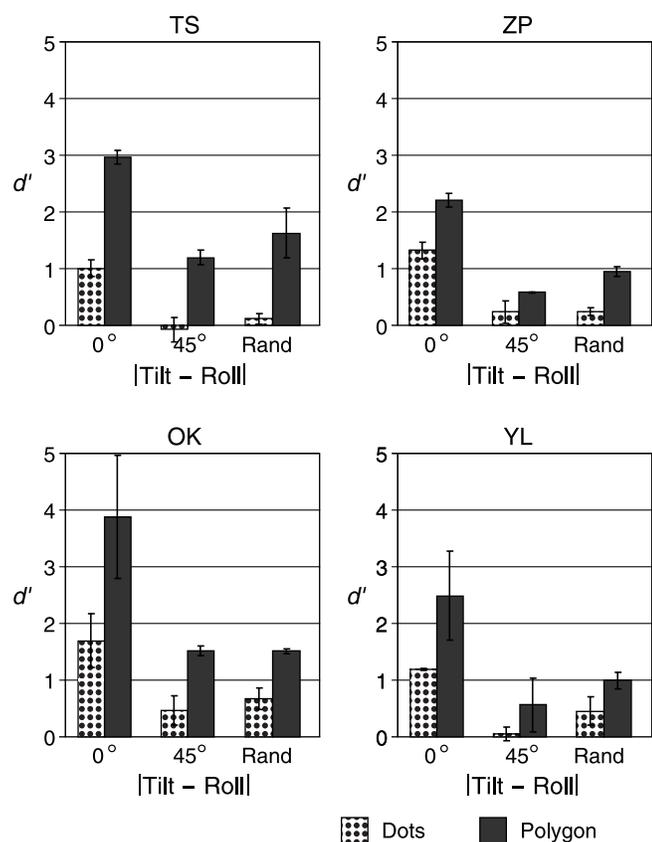


Figure 4. Results of individual subjects in [Experiment 1](#). The symbols indicate types of stimuli. Error bars represent the standard errors calculated from two sessions for each condition.

image is easier to detect than skewed symmetry. Performance in 45 deg condition was the worst. This is because symmetry is maximally distorted when the angle between the axis of symmetry and the tilt is 45 deg (Wagemans et al., 1991, 1992). Finally there was a significant interaction between the figure type and the angle between the symmetry axis and the tilt ($F_{2, 15} = 4.08, p < 0.05$). This interaction was due to the fact that the difference in performance between 0 deg and 45 deg conditions was large and significant in polygon condition (Tukey HSD: $p < 0.001$) but smaller and non-significant in dots condition ($p = 0.055$). This is most likely related to the floor effect; the performance in dots conditions was always quite poor.

Experiment 2: The effect of knowledge of the orientation of the projected symmetry axis, projected symmetry lines, and tilt

Results of [Experiment 1](#) show that the human visual system can reliably detect skewed symmetry from

perspective images, but only in the case of polygons. It is not clear whether these results can be explained by previous models because those models were designed for the case of (i) orthographic projection of (ii) dotted figures. The main idea behind all these models was related to the verification whether the midpoints of symmetry line segments are collinear (see the [Introduction](#)). According to these models, the visual system begins with finding the orientation of symmetry lines by trying a number of different correspondences among pairs of points. Then, it computes the midpoints and verifies their collinearity. This idea received support from Wagemans and his colleagues' (1991, 1992) experiments where they used orthographic images of dotted patterns and showed that performance in detecting skewed symmetry was higher when the projected symmetry lines or the projected symmetry axis had constant orientation and the orientation was known to the subject as compared to the case where the orientations of these lines were random. Indeed, according to the prior models, if the subject knows the orientation of symmetry lines, it is easier to find the right correspondences among pairs of symmetric points. Similarly, if the subject knows the orientation of the projected symmetry axis, then only one pair of skewed symmetric points is sufficient to determine all remaining pairs of skewed symmetric points. Recall that in perspective images of symmetric figures, the symmetry lines are not exactly parallel, and midpoints of symmetry line segments are not exactly collinear (see [Figure 2](#)). Does the beneficial effect of the known orientation of projected symmetry lines and symmetry axis generalize to perspective images of symmetric figures? More importantly, does this effect generalize to polygons? Recall that skewed symmetry is easy to detect when polygons but not dotted figures are used.

An alternative model of skewed symmetry detection can be derived from a theory of shape constancy, the “Perspective Invariants Theory” (Pizlo, 1994). Perspective invariants belong to the class of model based invariants (Rothwell, 1995) that were formulated for transformations

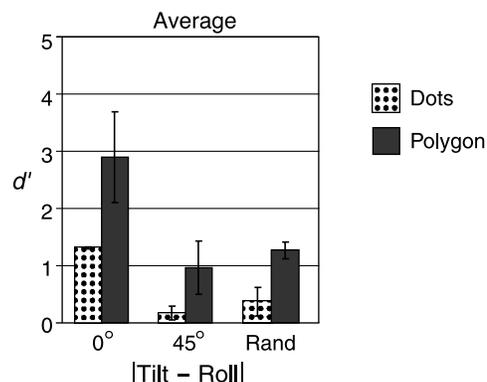


Figure 5. Averaged results from all 4 subjects in [Experiment 1](#). Error bars represent the standard errors calculated from four subjects.

that are not groups (Pizlo & Rosenfeld, 1992; Pizlo et al., 1997a, 1997b). Specifically, one can use perspective invariants to verify whether the two halves of a figure have identical shape. This way, symmetry detection task would become analogous to a shape constancy task. Note however, that perspective invariants require search for the correct tilt value because perspective invariants provide invariance in the case of slant but not tilt. In fact, Pizlo (1994) showed that if tilt is known to the subject, shape constancy performance improves. Does the beneficial effect of the known tilt generalize to skewed symmetry detection? If it does, it would suggest that skewed symmetry detection is based on the perspective invariants.

Method

Three subjects who participated in [Experiment 1](#) were tested (TS, ZP, and OK). The experimental method was the same as in [Experiment 1](#) except as indicated below.

There were six experimental conditions. In two conditions, the orientation of the projected symmetry axis was fixed and known to the subject. In the next two conditions, the orientation of the projected symmetry lines was fixed and known to the subject. In the remaining two conditions, the tilt direction was fixed and known to the subject. These three pairs of conditions represented three levels of a factor that will be called “known orientation.” The second factor was “direction”: main (vertical or horizontal) vs. diagonal (45 deg). In order to evaluate the magnitude of the effect of the known orientation, a random condition was included as well. This is the same as the random condition in [Experiment 1](#). Recall that in perspective projection, the projected symmetry lines are not parallel. Therefore, in the condition, in which the orientation of the projected symmetry lines was fixed, this orientation had to be represented by one of the symmetry lines. We chose the line going through the center of the screen and through the vanishing point of the projected symmetry lines.

The main direction for each “known orientation” condition was chosen based on previous results. Wagemans et al. (1991, 1992) tested a number of directions, including vertical, horizontal, and diagonal, of the projected symmetry axis and the projected symmetry lines under an orthographic projection. They showed that the benefit of the “known orientation” of the projected symmetry axis was maximal for the vertical direction, and the benefit of the “known orientation” of the projected symmetry lines was maximal for the horizontal direction. Because we are interested in measuring the maximal benefits of “known orientations,” therefore we used the vertical direction for the projected symmetry axis and the horizontal direction for projected symmetry lines. In the case of tilt, the vertical rather than the horizontal direction was chosen because the benefit of the “known orientation” of tilt was found, in a preliminary experiment, to be greater for the

vertical direction. Each subject ran 14 sessions in total: 7 conditions \times 2 replications.

Results and discussion

Results of each subject are shown in [Figure 6](#), and the averaged results are shown in [Figure 7](#). The ordinate shows d' . The results were analyzed using a two-way ANOVA within-subjects design (the random condition was not included in this analysis). The interaction was not significant ($p = 0.425$), but both main effects were ($p < 0.001$). In the case of the “direction” factor, performance in the main direction (horizontal or vertical) was higher than that in the diagonal direction. To evaluate the effect of the “known orientation” factor, a posteriori test (Tukey HSD) was applied. Performance in the two conditions where tilt direction was fixed was worse than performance in the two conditions where the direction of the projected symmetry axis was fixed ($p < 0.005$) and where the direction of the projected symmetry lines was fixed ($p < 0.001$). To compare the performance in these conditions to that in the random condition, a one-way ANOVA ($F_{6, 12} = 28.90$, $p < 0.001$) followed by a posteriori test (Tukey HSD) was applied. Performance in the random condition was not significantly different from that where the tilt direction was diagonal ($p > 0.999$), but it was worse than the other five conditions ($p < 0.05$).

Our results on the effect of knowledge of the orientation of the projected symmetry axis and projected symmetry lines are similar to the results of Wagemans et al. (1991, 1992) despite the fact that they used orthographic images of dotted stimuli, and we used perspective images of polygons. This similarity suggests that there is one underlying perceptual mechanism for detecting skewed symmetry. Recall that Wagemans (and others) proposed that the detection of skewed symmetry involves two invariant features of symmetry: parallelism of projected symmetry lines and collinearity of midpoints of projected symmetry line segments. Recall that these two features are invariant under orthographic, but not perspective projection. Under perspective projection, they are only approximately invariant. This raises two possibilities. First, the visual system uses invariants of perspective projection, and orthographic projection is treated as an approximation of perspective projection. Second, the visual system uses invariants of orthographic projection, and perspective projection is treated as an approximation to orthographic projection. A number of studies demonstrated that the visual system uses the rules of perspective projection in shape perception (Kaiser, 1967; Pizlo & Salach-Golyska, 1995; Yang & Kubovy, 1999). Only one study suggested that the visual system uses the rules of orthographic, rather than perspective projection in shape perception (Hagen & Elliot, 1976). Results of our [Experiment 2](#) provided indirect evidence that detection of skewed symmetry does not involve perspective invariants. Recall

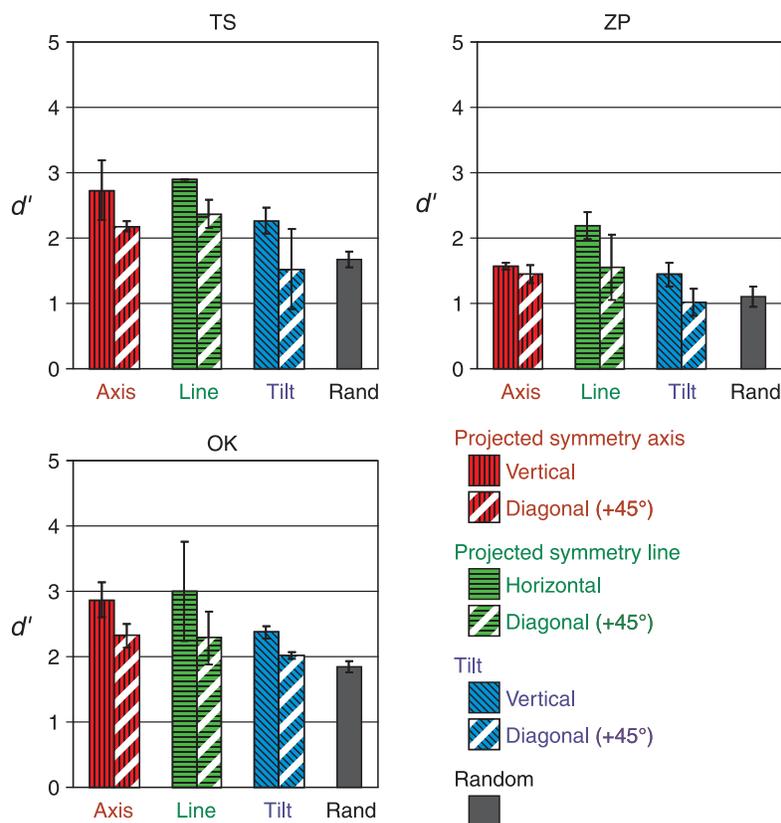


Figure 6. Results of individual subjects in Experiment 2. The symbols indicate “known orientation” and “direction.” Error bars represent the standard errors calculated from two sessions.

that the use of perspective invariants requires search for the tilt value. It follows that if tilt is known, detection should be easier. But the benefit of knowing tilt direction was small and present only in the case of vertical direction. This fact suggests that detection of skewed symmetry uses invariants of orthographic projection. If this is indeed the case, detection of skewed symmetry from orthographic images should be easier than from perspective images. This was tested in a control experiment.

Control experiment: Perspective vs. orthographic projection

To test whether the visual system uses the rules of perspective or orthographic projection in detecting skewed symmetry, the control experiment involved two types of stimuli. Some stimuli were perspective images of slanted 2D polygons. Others were orthographic images of slanted 2D polygons. When the image on a computer screen is a perspective image of a simulated 2D figure slanted in depth, then the retinal image in the observer’s eye is a valid (correct) image of the slanted 2D figure (assuming that the observer’s eye is at the center of perspective projection that was used to compute the perspective image). When the image on a computer screen is an

orthographic image of a simulated 2D figure slanted in depth, then the retinal image in the observer’s eye is *not* a valid image of the slanted 2D figure. So, from the point of view of geometrical optics, perspective projection is *better* than orthographic projection. But from the point of view of perceptual mechanisms of symmetry detection, this may or may not be true. Specifically, if the visual system uses properties of orthographic projection, then orthographic images are “better” than perspective images.

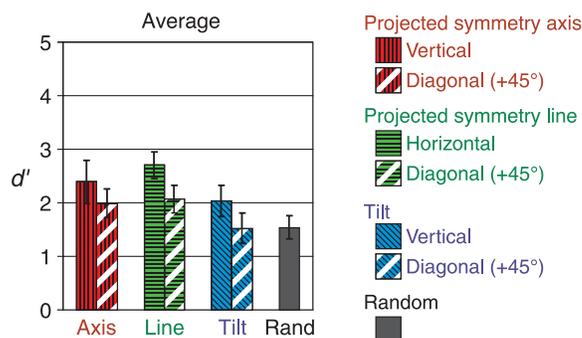


Figure 7. Averaged results from all 3 subjects in Experiment 2. Error bars represent the standard errors calculated from three subjects.

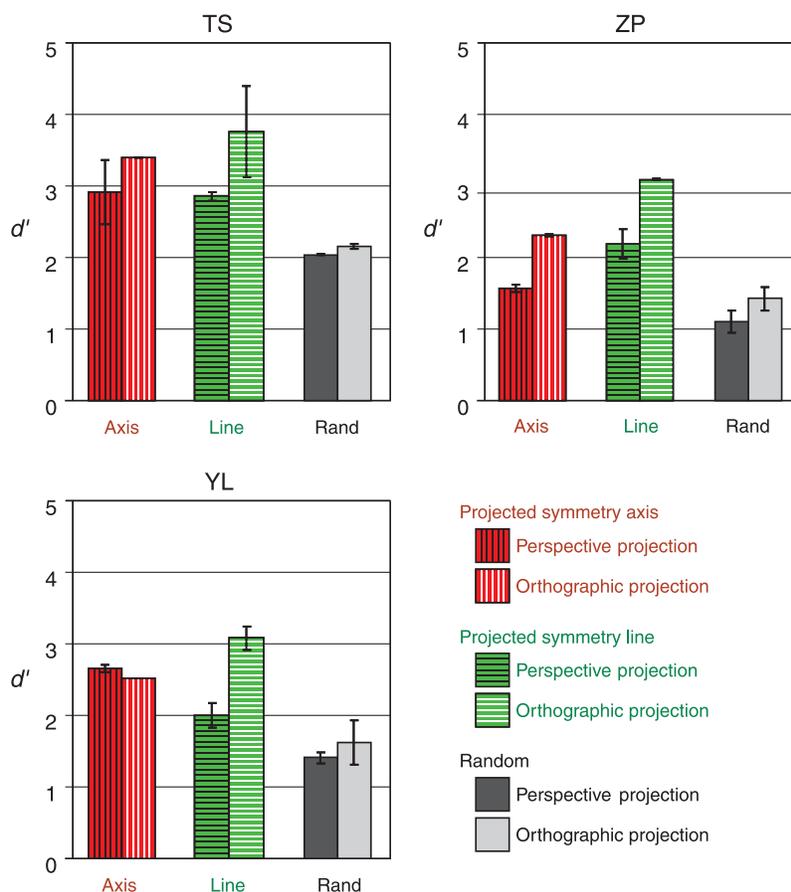


Figure 8. Results of individual subjects in control experiment, which compared skewed symmetry detection from perspective vs. orthographic images. The symbols indicate “known orientation” and type of projection. Error bars represent the standard errors calculated from two sessions.

Perspective images are only approximations. If, on the other hand, the visual system uses properties of perspective projection, then perspective images are “better” and orthographic images are only approximations. It follows that if the visual system uses the rules of perspective projection in detection of skewed symmetry, performance should be higher with perspective images. If the visual system uses the rules of orthographic projection in detection of skewed symmetry, performance should be higher with orthographic images.

Performance in detecting skewed symmetry under the orthographic and perspective projection was tested in three of the seven conditions that were used in Experiment 2: The direction of the projected symmetry axis is vertical, the direction of the projected symmetry lines is horizontal, and all directions are random. In total, there were six experimental conditions: two types of projection (perspective vs. orthographic) \times three “known orientations” (vertical projected symmetry axis vs. horizontal projected symmetry lines vs. random). Each condition was repeated twice in a random order.

Results of each subject are shown in Figure 8, and the averaged results are shown in Figure 9.² It can be seen that the performance for orthographic projection was

better than that for perspective projection ($p < 0.005$), although the magnitude of this difference was rather small. Large difference was not expected simply because the perspective and orthographic images were similar to each other under the viewing conditions used here (an example is shown in Figure 2). The difference between perspective and orthographic images would have been larger, if either the viewing distance was smaller or the

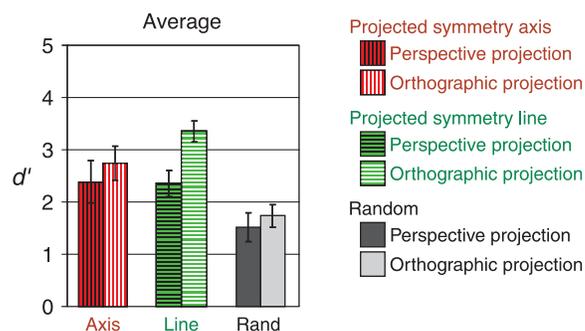


Figure 9. Averaged results from all three subjects in control experiment. Error bars represent the standard errors calculated from three subjects.

computer monitor was larger. The fact that orthographic images produced more reliable performance suggests that the visual system uses the rules of orthographic rather than perspective projection in detection of skewed symmetry. This result is important because it is perspective not orthographic projection that adequately describes the rules of image formation in the eye.

Next, both the effect of “known orientation” and the interaction are significant (known orientation: $p < 0.001$; interaction: $p < 0.05$). The effect of “known orientation” has already been demonstrated and discussed in [Experiment 2](#). The new result here is the interaction. The interaction effect seems to be produced by the fact that the effect of the type of projection is larger in the horizontal projected symmetry lines condition than in the other two “known orientation” conditions. To examine the interaction effect, a one-way ANOVA was applied to the difference of performance between the two types of projection (orthographic–perspective) for the three “known orientation” conditions ($F_{2, 6} = 6.50$, $p < 0.05$). A posteriori test (Tukey HSD) showed that the difference in the horizontal projected symmetry lines condition was significantly larger than that in the random condition ($p < 0.05$). These results suggest that the parallelism of the projected symmetry lines, as well as the orientation of the projected symmetry axis, is used by the visual system in the process of detection of skewed symmetry. Recall that the parallelism of symmetry lines is an invariant of orthographic but not perspective projection of a symmetric figure. Also, the projected symmetry axis connects the midpoints of the projected symmetry line segments under orthographic, but not under perspective projection.

Based on the results of the psychophysical experiments described above, as well on the results published by others, a computational model of skewed symmetry detection was formulated. This model is described, next.

Computational model of skewed symmetry detection

The model is intended to work with polygons, rather than dots because human subjects can reliably detect skewed symmetry in the case of polygons. Detection of skewed symmetry in the case of dots is substantially worse. In fact, it was at chance level in our experiment. Next, the model uses the rules of orthographic projection and perspective images are treated as approximations. The model is not only able to detect skewed symmetry. After skewed symmetry is detected, the model reconstructs the slanted 2D polygon. This model processes the polygon in the following way. First, the projected symmetry lines are detected using the parallelism of the projected symmetry line segments and collinearity of their midpoints. At the same time, the projected symmetry axis is computed using

the collinearity of midpoints of the projected symmetry line segments. Next, the symmetric polygon is reconstructed from the skewed symmetric polygon. Finally, symmetry of the reconstructed polygon is measured by comparing shapes of contours of the two halves of the reconstructed polygon.

Parallelism of the projected symmetry line segments and collinearity of their midpoints

A search is performed to find pairs of vertices which maximize a measure of (i) parallelism among the possible symmetry lines and (ii) collinearity of midpoints of symmetry line segments. The use of these two measures was motivated by Wagemans’s (1995) observation that in an orthographic image of a symmetric figure, symmetry line segments are parallel and their midpoints are collinear (see also Jenkins, 1983, who discussed these properties in the context of a symmetric image).

In this step of the model, a given image is assigned its “detected symmetry axis” and “detected symmetry lines.” Note that this assignment is performed for every image regardless as to whether or not it is an image of a symmetric polygon.

Assume that a polygon has n vertices. This polygon has n possible placements of the projected symmetry axis. If n is odd, each possible symmetry axis crosses $(n - 1) / 2$ symmetry lines including a side of the polygon. If n is even, each possible symmetry axis crosses $n / 2$ symmetry lines including two sides of the polygon or $(n - 2) / 2$ symmetry lines. Correspondence of pairs of symmetric vertices is uniquely specified for each possible symmetry axis: starting from an intersection of a possible symmetry axis and a polygon, the n th vertex in a clockwise direction should form a skewed symmetric pair with the n th vertex in a counterclockwise direction.³ (Note that in the case of dotted stimuli, an additional process for finding a skewed symmetric partner of each dot would be needed.) The model examines parallelism and collinearity of midpoints of the projected symmetry line segments for each possible symmetry axis. Then, it chooses the symmetry axis and the corresponding symmetry line segments that maximize the parallelism and the collinearity of midpoints.

Parallelism of the possible projected symmetry line segments and collinearity of their midpoints are measured for each possible projected symmetry axis in the following way. To evaluate the parallelism, an average orientation of possible projected symmetry line segments is computed, and the mean squared error of orientations from the average orientation is calculated for each possible symmetry axis. An inverse of this mean squared error of orientation is used as the measure of the parallelism. Then the midpoints of these possible projected symmetry line segments are computed, and a straight line is fitted to them by using the least squares method. This line is a possible projected symmetry axis. An inverse of the mean

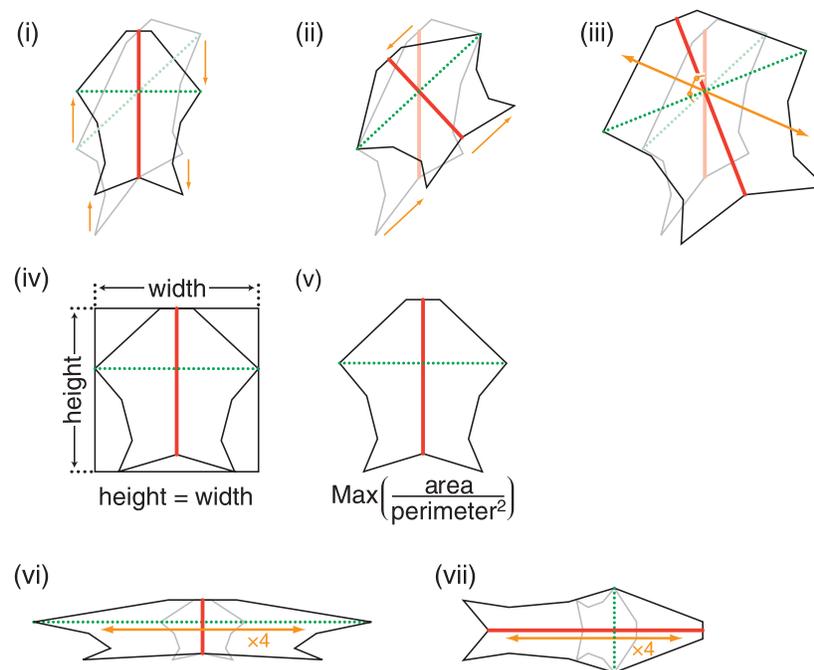


Figure 10. Seven ways of symmetry reconstruction with using the information about orientations of the projected symmetry axis and the projected symmetry lines. Symmetric figures are (i) reconstructed by shearing the figure along the projected symmetry axis, (ii) reconstructed by shearing the figure along the projected symmetry lines, (iii) reconstructed by stretching the figure along a bisector of an obtuse angle between the projected symmetry axis and the projected symmetry lines, (iv) chosen from the family of the symmetric figures by making the rectangle circumscribed on it a square, (v) chosen by maximizing its 2D compactness, (vi) obtained by stretching the shape reconstructed by model (ii) along the symmetry lines by a factor of four, and (vii) obtained by stretching the shape reconstructed by model (ii) along the symmetry axis by a factor of four.

squared error is used as the measure of the collinearity. The product of these two measures is computed for all possible orientations of the projected symmetry axis. The maximum of this product indicates the detected projected symmetry axis and the detected projected symmetry lines.

Reconstructing symmetry

Before the symmetry of a slanted figure is evaluated, the figure is reconstructed from the skewed symmetric image.

An orthographic image of a mirror-symmetric polygon determines a one-parameter family of possible symmetric interpretations (Kanade & Kender, 1983; Saunders & Knill, 2001). This family can be characterized by an aspect ratio, which is the ratio between the length of the longest symmetry line segment and the length of the intersection of the symmetry axis and the interior of the polygon. The question is which member of this family is actually perceived by an observer who is presented with a skewed symmetric image. To the best of our knowledge, this question has never been addressed in a psychophysical study before (although it was discussed by the computer vision community, e.g., Kanade & Kender, 1983).⁴ A symmetric polygon and its orthographic image are related by a 2D affine transformation. It follows that a

reconstruction of a symmetric figure must use a subset of a 2D affine transformation.

We consider seven different models (Figure 10). In the first model (model (i)), the polygon is sheared along the projected symmetry axis until the symmetry lines become perpendicular to the projected symmetry axis. In model (ii), the polygon is sheared along the projected symmetry lines, until the symmetry axis becomes perpendicular to the projected symmetry lines. In model (iii), the figure is stretched along a bisector of an obtuse angle between the projected symmetry axis and the projected symmetry lines, until they become perpendicular to each other. This model corresponds to the minimum slant interpretation: according to this interpretation, the observer recovers a symmetric figure by minimizing its slant (Kanade, 1981; Stevens, 1979). Model (iv) chooses a symmetric polygon for which a circumscribed rectangle is a square. Model (v) chooses a symmetric polygon whose 2D compactness, defined by the area divided by the perimeter squared, is maximal (Brady & Yuille, 1988; Hildebrandt & Tromba, 1996; Zusne, 1970). The sixth and the seventh models were simple modifications of model (ii). Namely, the shape reconstructed by model (vi) was obtained by stretching the shape reconstructed by model (ii) along the symmetry lines by a factor of four. The shape reconstructed by model (vii) was obtained by stretching

	R^2	Coefficient	Intercept
TS			
(i)	0.865	0.957 ± 0.038	0.403 ± 0.026
(ii)	0.954	1.005 ± 0.022	0.070 ± 0.018
(iii)	0.939	1.013 ± 0.026	0.207 ± 0.077
(iv)	0.207	1.592 ± 0.315	0.195 ± 0.084
(v)	0.249	0.922 ± 0.162	0.207 ± 0.077
(vi)	0.954	1.005 ± 0.022	-1.323 ± 0.042
(vii)	0.954	1.005 ± 0.022	1.463 ± 0.027
OK			
(i)	0.853	0.863 ± 0.036	0.125 ± 0.009
(ii)	0.887	0.958 ± 0.034	0.006 ± 0.010
(iii)	0.902	0.941 ± 0.031	0.065 ± 0.008
(iv)	0.112	1.188 ± 0.338	0.244 ± 0.068
(v)	0.107	0.554 ± 0.162	0.273 ± 0.064
(vi)	0.887	0.958 ± 0.034	-1.305 ± 0.064
(vii)	0.887	0.958 ± 0.034	1.343 ± 0.039

Table 1. Results of linear regression applied to the relation between the logarithm of the aspect ratio reconstructed by a subject and the logarithm of the aspect ratio reconstructed by each of the seven models.

the shape reconstructed by model (ii) along the symmetry axis by a factor of four. These last two models were added in order to provide a control for testing the models in shape discrimination (see [Simulation experiment](#) section).

In order to test which of the seven models (if any) is used by human observers, we performed a control experiment (see below).

Psychophysical test: Reconstruction of symmetry from skewed symmetry

The stimuli were generated the same way as those in the previous experiments. The symmetric polygon slanted in depth was orthographically projected with random tilt to generate projected images in such a way that the projected symmetry axis was vertical. The line of sight connecting the center of the monitor, and the viewing eye was perpendicular to the monitor. Each subject ran 100 trials. In each trial, the test stimulus was skewed symmetric, and the response stimulus was a corresponding symmetric polygon. The subject's task was to adjust the aspect ratio of the response symmetric figure so that its shape matched the perceived symmetric shape produced by the skewed symmetric stimulus. The subjects could not see both of the stimuli at the same time. They could alternate between the test and the response stimuli as many times as they wanted. The interstimulus interval was 400 ms. This length of the interstimulus interval did not lead to the kinetic depth effect. The exposure duration was unlimited.

It is important to point out that this task is expected to be difficult to the subject because of the fact that there was

no unique correct answer in any given trial. As indicated above, any of the response stimuli in a given trial could actually produce the test stimulus. As a result, a large variability in the adjusted aspect ratio could be expected. Furthermore, the subject's adjustment could be affected by one or more response biases. Despite these inherent methodological weaknesses, we decided to perform this experiment in an attempt to shed some light on the mechanisms underlying human perception of symmetry.

Results were analyzed using linear regression applied to the logarithms of the aspect ratios ([Table 1](#)). The adjusted aspect ratio was compared to the aspect ratio of the symmetric figure reconstructed by each of the seven models defined in the previous section. It can be seen that the first three models provide substantially better prediction than the remaining four: the correlation coefficients are close to one, the slopes of the regression lines are close to one and the intercepts are close to zero. From the first three models, the best prediction seems to be provided by model (ii): the regression line was closest to the identity transformation. [Figure 11](#) shows scatterplots

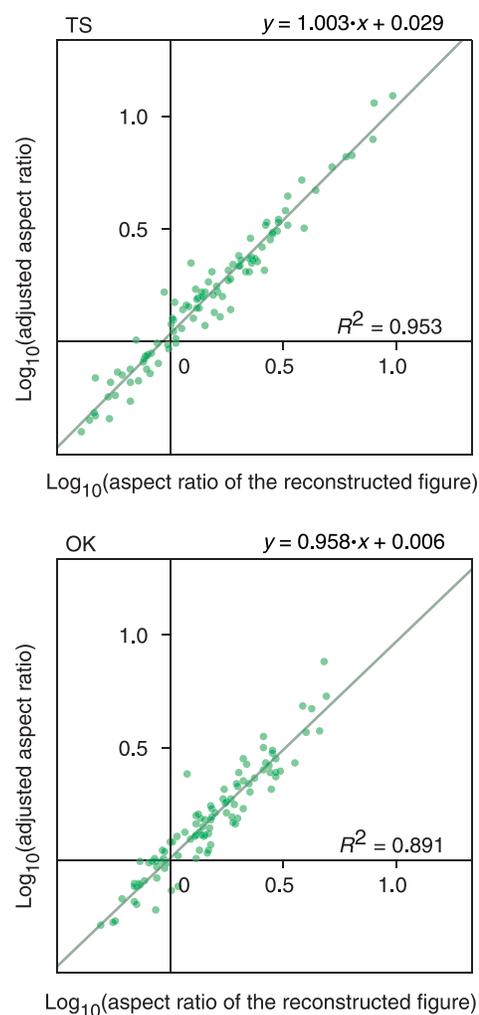


Figure 11. Scatterplots for each of the two subjects. The abscissa shows the aspect ratio reconstructed by model (ii).

illustrating the relation between the aspect ratio reconstructed by the subject and by model (ii). Therefore, model (ii) of symmetry reconstruction will be used in our model of symmetry detection.

A natural question arises as to whether the comparison between perceived shape and reconstructed shape depends on how the shape is measured. Recall that the family of symmetric shapes defined by a skewed symmetric image is represented by one parameter. Does it matter which parameter is used? The results presented in Table 1 used the aspect ratio defined in the previous section. We tried several other shape measures, including 2D compactness and aspect ratio of the circumscribed rectangle. All these other shape measures produced results very similar to those shown in Table 1. This means that the comparison between perceived shape and reconstructed shape is robust to the way the shape is actually measured.

Verification of symmetry of a polygon

To evaluate the degree of symmetry of the sheared polygon, the polygon is divided into pairs of vertices connected by symmetry line segments. The parts of the contour around the vertices of each pair are compared by using the Ψ function (Pizlo & Rosenfeld, 1992). The Ψ function is a contour-based shape descriptor; the independent variable of the Ψ function is the contour length measured from a given starting point, and the dependent variable is the contour orientation at that point measured relative to a given orientation. In this model, the Ψ function of a given vertex starts from one neighboring vertex and ends at the other neighboring vertex.

To compare the Ψ functions of each pair of vertices, the lengths of them are normalized to the averaged length of each pair. Then the absolute difference between these two Ψ functions is calculated for each pair. The differences for all pairs of vertices are summed and divided by the product of the maximum possible angle difference (180 deg) and the sum of the lengths of Ψ functions. The resulting value is the measure of asymmetry. The range of this measure is from 0 to 1; the bigger this value is, the less symmetric the polygon is. The decision as to whether or not the polygon is symmetric is based on a criterion chosen to make the hit rate and the correct rejection rate equal.

Simulation experiment

The model of detection of skewed symmetry was applied to the same type of skewed symmetric images as those used in our psychophysical experiments. Recall that our skewed symmetry detection model consists of three stages. First, the model estimates the orientation of the projected symmetry lines and the projected symmetry axis (see [Parallelism of the projected symmetry line segments and collinearity of their midpoints](#) section). Next, the

slanted polygon is reconstructed using one of the seven reconstruction models. Finally, the asymmetry of the polygon is measured by using the method described in the [Verification of symmetry of a polygon](#) section and thresholded in order to produce the answer “symmetric” vs. “asymmetric.” Note that each of the seven reconstruction models can be used in our detection model. As a result, we actually have seven detection models. We begin with the detection model that uses reconstruction model (ii). First, we assume that the projected symmetry lines and symmetry axis are found by performing exhaustive search. Recall that this exhaustive search involves n possible orientations, where n is the number of vertices in a polygon.

We generated symmetric and asymmetric polygons and then computed orthographic images using the same method as that used in our psychophysical experiments in the “random” condition. The same was done for perspective projection. The positions of the vertices of the projected polygon were randomly perturbed to simulate the visual noise. (If there was no visual noise, the measure of asymmetry would always be 0 for an orthographic image of symmetric polygon and performance of the model would be perfect.) To simulate the visual noise, zero mean, Gaussian noise was used (Chan, Pizlo, & Chelberg, 2006; Levi & Klein, 1990). The standard deviation of the noise was proportional to the eccentricity of the vertices to reflect the non-uniform distribution of cones in the retina. It is difficult to decide what the coefficient of this proportionality should be. Based on a number of psychophysical results related to line length discrimination, Chan et al. (2006) concluded that this coefficient should be 1–3%. This number corresponds to the Weber fraction in line length discrimination. Note, however, that Weber fraction of about 2% is observed for unlimited exposure durations. When exposure duration is shorter than 1 s, Weber fraction is elevated (Watt, 1987). In particular, Weber fraction for 100 ms is 2–3 times larger than for 1 s. Recall that our psychophysical experiments used 100 ms exposure duration. It follows that the visual noise of 6% is a reasonable guess. This is what we used in our simulations (for a review of related studies, see Pizlo, Rosenfeld, & Epelboim, 1995). We tried several other levels of noise, as well, and verified that 6% is a good choice. In particular, this level of noise produced a good fit to our psychophysical results in the condition where human performance was the best. For each type of projection, the model was tested in four sessions, each session consisting of images of 200 symmetric polygons and 200 asymmetric polygons.

Figure 12 shows frequency histograms of the asymmetry measure for symmetric and asymmetric polygons. The histogram for symmetrical polygons is substantially different from that for asymmetrical polygons in both orthographic projection ($d' = 3.09 \pm 0.060$) and perspective projection ($d' = 2.54 \pm 0.018$). The d' s were

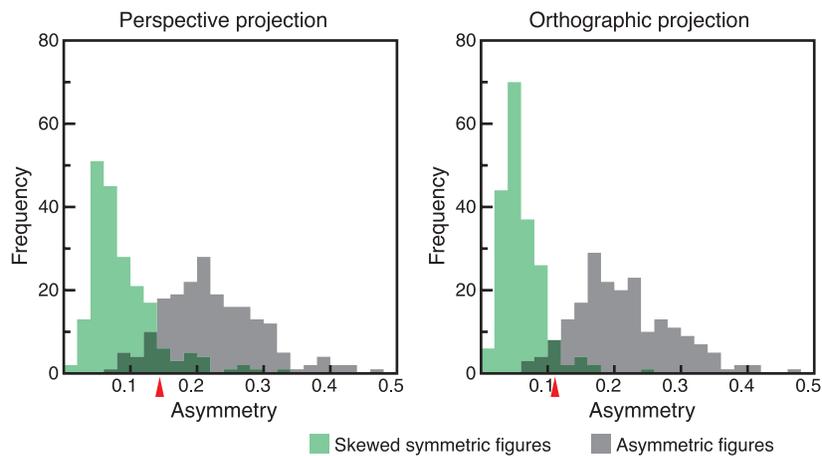


Figure 12. The histogram of the asymmetry measure produced by the proposed model for the orthographic and perspective projection. Red marks on the abscissa indicate the criteria.

calculated based on the criterion for which the hit rate was equal to the correct rejection rate. Note that performance for orthographic projection is slightly better than that for perspective projection.

Next, we repeated the experiment using the other six models of symmetry reconstruction. This test was done in order to verify whether symmetry detection depends on how the symmetry is reconstructed. The d' s for all seven models are shown in Figure 13. It can be seen that performance of the first five models is equally good, but that of the last two models is somewhat lower. These results show that evaluating asymmetry of a given shape does not depend strongly on how the shape is reconstructed. The rest of the simulations uses model (ii) of symmetry reconstruction.

Note that symmetry detection of our model is substantially higher than that of human subjects. This can be seen by comparing d' s of our model shown in Figure 13 (model (ii)) to d' s of our model shown in Figure 9, random condition. Performance of the model is about twice as high as that of the subjects. Recall that our model performs an exhaustive search for the projected symmetry lines and symmetry axis. It is reasonable to expect that subjects did not perform exhaustive search, considering the fact that exposure duration was short (100 ms). Indeed, fixing the orientation of symmetry lines or symmetry axis led to substantially better performance (see Figure 9). This result strongly suggests that without knowing these orientations, the visual system tries only a few of them. In order to evaluate the effect of the amount of search for the orientation of the projected symmetry lines on performance, we performed the next simulation experiment. We tested 10 conditions, corresponding to the number of possible orientations of the projected symmetry lines that were tried. This number ranged from 1 to 9, plus an exhaustive search. The actual orientations were chosen randomly, and the best was used to measure the asymmetry of a polygon.

Results are shown in Figure 14 for both orthographic and perspective projections. The ordinate shows d' , and the abscissa shows the number of orientations that were tried as possible orientations of projected symmetry lines. As expected, performance is better when the search involves more orientations. For eight orientations, the performance of the model ($d' = 1.35 \pm 0.040$ under the perspective projection, and $d' = 1.65 \pm 0.086$ under the orthographic projection) is closest to the average performance of human subjects in the random condition (see Figure 9).

Finally, we tested the effect of knowledge of the orientation of projected symmetry lines and projected symmetry axis on the model's performance. On each trial, the model was given the orientation of the projected symmetry lines (or symmetry axis). The model examined eight orientations of the projected symmetry lines

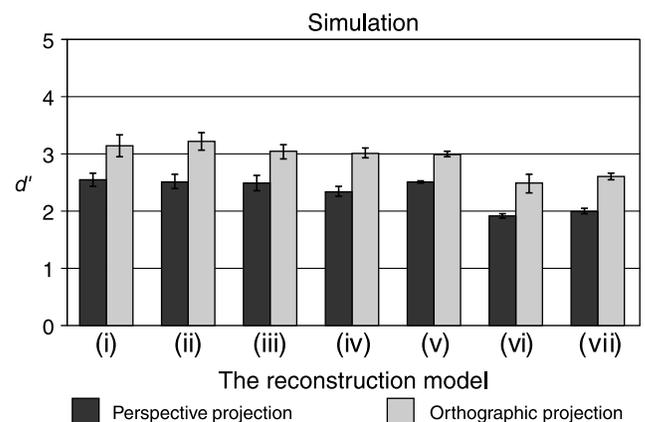


Figure 13. Detection performance for all seven models of symmetry reconstruction. Dark bars represent perspective projection, and light bars represent orthographic projection. Error bars represent the standard errors calculated from four replications (sessions).

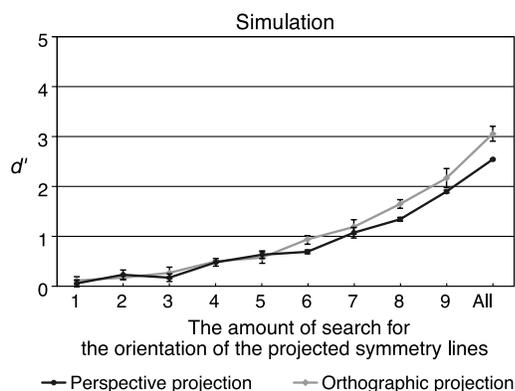


Figure 14. The effect of the amount of search for the orientation of the projected symmetry lines on performance. The abscissa shows the number of orientations that were tried.

(symmetry axis) from the range ± 45 deg around the given orientation. If there were more than eight orientations in this range, the eight closest to the given orientation were used. If there were fewer than eight, then only the orientations from this range were used. Again, symmetry detection used the reconstruction model (ii).

Results of the simulations are shown in Figure 15 for both perspective and orthographic projection. It can be seen that the performance of the model is very close to the performance of subjects shown in Figures 8 and 9. Specifically, performance for orthographic projection is slightly better than that for perspective projection, and the knowledge of orientation of the projected symmetry axis and the projected symmetry lines improves the performance. These results show that the proposed model can adequately simulate the psychophysical results.

Summary and discussion

In Experiment 1, performance in detecting symmetric and skewed symmetric figures was compared using dotted stimuli and polygons. Performance was substantially higher with symmetric, as compared to skewed symmetric figures. More importantly, performance was substantially higher with polygons as compared to dots. In Experiment 2, the effect of knowledge of the projected symmetry axis, projected symmetry lines, and tilt was examined. The results showed that the knowledge of the projected symmetry axis and projected symmetry lines improved performance, but that of tilt did not. In the Control experiment, performance in detecting skewed symmetry under the perspective and orthographic projection was compared, and the results showed that the performance under the orthographic projection was slightly better than that under the perspective projection. The effect of the type of projection was highest when the orientation of the projected symmetry lines was known to the subject.

Finally, reconstruction of a symmetric figure from the skewed symmetric figure was tested. The results showed that the perceived symmetric figure was very close to the figure produced from the skewed symmetric figure by shearing it along the projected symmetry lines.

Based on these results, a computational model of detecting skewed symmetry was proposed. The model applies to polygons and uses the rules of orthographic projection. The model begins with detecting the projected symmetry line segments. A search is performed to find a set of the projected symmetry line segments based on a criterion involving their parallelism and collinearity of their midpoints. This leads to the estimation of the projected symmetry axis. Next, the slanted polygon is reconstructed by applying an affine transformation. Finally, asymmetry of the reconstructed polygon is measured by comparing shapes of contours of the two halves of the reconstructed polygon. But note that the model does not compare the two halves in one step holistically. Instead, it compares shapes of parts of contours around corresponding vertices. The model was tested with noisy orthographic and perspective images to simulate the psychophysical results. The model provided good fit to the subjects' data.

Our results showed that both subjects' and model's performance in skewed symmetry detection are reliable with polygons. Recall that Attneave (1982) claimed that skewed symmetry cannot be detected reliably by the human visual system. In his demonstration, he used only dotted patterns. Clearly, Attneave's claim does not generalize to polygons. Why symmetry and skewed symmetry are detected much easier from polygons than from dotted stimuli? There are two main differences between polygons and dotted stimuli:

- the presence of the orientation information of contours in the case of polygons and
- the information about the order of vertices in the case of polygons.

Consider the role of orientation information. This was tested in the study of Locher and Wagemans (1993). They

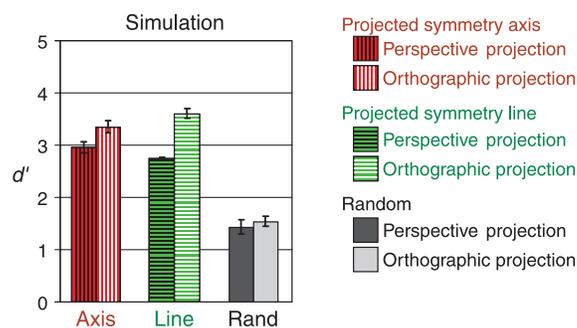


Figure 15. The effect of the knowledge of the orientation of projected symmetry lines and symmetry axis on the model's performance.

compared performance for two types of stimuli: dots vs. line segments in a symmetry detection task. The line segments stimulus was produced by replacing dots by short line segments. Note that the line segments were not connected, and thus they did not form any shapes. Performance in these two conditions was similar. This suggests that the mere presence of line segments is not critical in symmetry detection. In other words, local information about orientations of line segments does not add much, if anything to symmetry detection. The superiority of polygons over dotted stimuli in our experiments is probably related to more global aspects of the polygons. One such global feature is the information about the order of vertices. Recall that this order makes the search for symmetry lines and symmetry axis easier. The order information comes from the continuity of contours of polygons; the contour of a polygon is a continuous and closed line.

Note that most prior studies on symmetry detection with dotted patterns used large number of dots: hundreds and thousands of dots (Barlow & Reeves, 1979; Jenkins, 1983; Julesz, 1971). They demonstrated rather reliable performance with such stimuli. Recall that these studies tested detection of symmetric patterns, not skewed symmetric ones. Our study of skewed symmetry detection used not more than 12 dots (see our Experiment 1). Is it possible that the poor performance with dotted patterns in our experiment was an artifact related to the small number of dots used? This is quite unlikely. Our informal observations suggest that increasing the number of dots in a skewed symmetry stimulus does not make the detection of skewed symmetry easier. If anything, it seems to make it more difficult. This is illustrated in Figure 16. These dotted stimuli consist of 20, 50, and 100 dots. It is probably impossible to detect skewed symmetry with these stimuli.

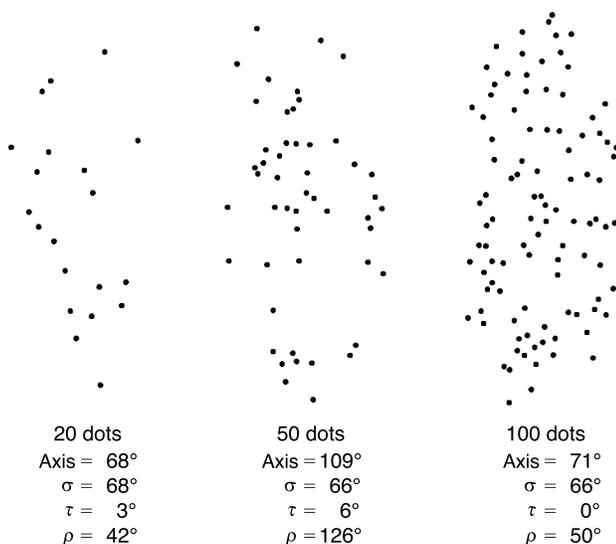


Figure 16. Three skewed symmetric dotted patterns. These patterns consist of 20, 50, and 100 dots. The numbers on the bottom indicate the orientation of the projected symmetry axis, slant (σ), tilt (τ), and roll (ρ).

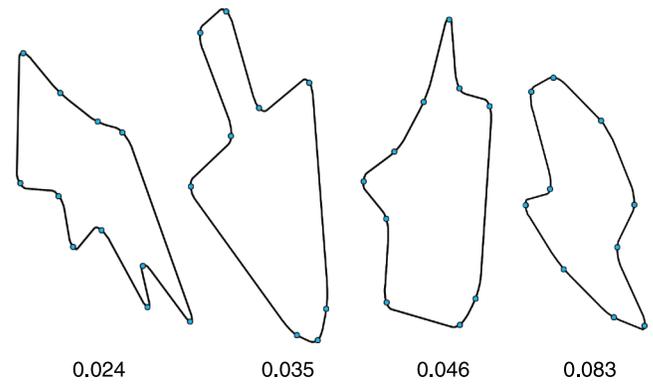


Figure 17. Examples of skewed symmetric smooth contours produced by an orthographic projection. Blue dots indicate positions of feature points that were detected based on the curvature of the contours. Numbers on the bottom are measures of asymmetry produced by our model.

Finally, consider possible generalizations of our model. The model was designed to detect symmetry of 2D (planar) polygons. Can this model be applied to 2D smoothly curved contours (shapes) such as those shown in Figure 17? These shapes were produced from polygons used in our psychophysical experiments by replacing vertices with curved arcs. Before the model was applied, characteristic points were detected. This was done by taking points of maximal (positive and negative) curvature of the contour (Attneave, 1954).⁵ To make the characteristic points stable, the curvature function was smoothed out by a local smoothing operator. The numbers below the shapes represent the measure of asymmetry produced by our model. When these numbers are compared to the criterion that our model used for discriminating between skewed symmetric and asymmetric polygons (Figure 12), it is seen that the shapes in Figure 17 would all be classified as skewed symmetric. Once we know that the model can be applied to smooth contours, the next question is whether the model can be applied to contours of 3D smooth surfaces and 3D volumetric objects. The answer is “probably yes,” as long as the contours are piece-wise planar. Testing human performance in detecting symmetry of 3D surfaces and 3D shapes, and generalizing the model to such stimuli will be addressed in our future work.

Acknowledgments

We are very grateful to Yunfeng Li for helpful discussions about the computational model. We also thank the two reviewers whose questions and suggestions led to a substantial improvement of this paper. This project was supported by the National Science Foundation (Grant # 0533968) and the U.S. Department of Energy.

Preliminary results of our experiments were published in Sawada and Pizlo (2007).

Commercial relationships: none.

Corresponding author: Tadamasa Sawada.

Email: tsawada@psych.purdue.edu.

Address: 703 3rd Street, West Lafayette, IN 47907-2081, USA.

Footnotes

¹We also plotted the results using overall proportion correct, rather than d' , as a dependent variable. These two dependent variables led to the same conclusions.

²Results of subject ZP in Experiment 2 were used as his results in the perspective condition in this experiment because he already reached an asymptotic performance. TS ran the perspective conditions again because his performance improved, somewhat, compared to Experiment 2. YL ran all conditions because he was not tested in Experiment 2.

³In real images of symmetric figures, spurious vertices may be present due to the imperfection of image segmentation. In such cases, the model should examine the neighboring vertices, from $n - i$ to $n + i$, as well.

⁴A related question about the types and order of affine transformations determining the shape percept was studied by Wagemans, Vanden Bossche, Segers, and d'Ydewalle (1994).

⁵Bitangent points (Rothwell, 1995) and inflection points (De Winter & Wagemans, 2004) are also possible characteristic points on smoothly curved contour.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193. [PubMed]
- Attneave, F. (1982). Prägnanz and soap bubble systems: A theoretical exploration. In J. Beck (Ed.), *Organization and representation in perception* (pp. 11–29). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Barlow, H. B., & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research*, *19*, 783–793. [PubMed]
- Brady, M., & Yuille, A. (1988). Inferring 3D orientation from 2D contour (An experimental principle). In W. Richards (Ed.), *Natural computation* (pp. 99–106). Cambridge: MIT Press.
- Chan, M. W., Pizlo, Z., & Chelberg, D. M. (2006). Binocular shape reconstruction: Psychological plausibility of the 8-point algorithm. *Computer Vision and Image Understanding*, *74*, 121–137.
- Coxeter, H. S. M. (1987). *Projective geometry* (second edition). New York: Springer.
- De Winter, J., & Wagemans, J. (2004). Contour-based object identification and segmentation: Stimuli, norms and data, and software tools. *Behavior Research Methods, Instruments, & Computers*, *36*, 604–624. [PubMed] [Article]
- Hagen, M. A., & Elliott, H. B. (1976). An investigation of the relationship between viewing condition and preference for true and modified linear perspective and adults. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 479–490. [PubMed]
- Herbert, A. M., & Humphrey, G. K. (1996). Bilateral symmetry detection: Testing a 'callosal' hypothesis. *Perception*, *25*, 463–480. [PubMed]
- Hildebrandt, S., & Tromba, A. (1996). *The parsimonious universe*. New York: Springer.
- Jenkins, B. (1983). Component processes in the perception of bilaterally symmetric dot textures. *Perception & Psychophysics*, *34*, 433–440. [PubMed]
- Julesz, B. (1971). *Foundation of cyclopean perception*. Chicago: University of Chicago Press.
- Kaiser, P. K. (1967). Perceived shape and its dependency on perceived slant. *Journal of Experimental Psychology*, *75*, 345–353. [PubMed]
- Kanade, T. (1981). Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, *17*, 409–460.
- Kanade, T., & Kender, J. R. (1983). Mapping image properties into shape constraints: Skewed symmetry, affine-transformable patterns, and the shape-from-texture paradigm. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 237–257). New York: Academic Press.
- Levi, D. M., & Klein, S. A. (1990). The role of separation and eccentricity in encoding position. *Vision Research*, *30*, 557–585. [PubMed]
- Locher, P. J., & Smets, G. (1992). The influence of stimulus dimensionality and viewing orientation on detection of symmetry in dot patterns. *Bulletin of the Psychonomic Society*, *30*, 43–46.
- Locher, P. J., & Wagemans, J. (1993). Effects of element type and spatial grouping on symmetry detection. *Perception*, *22*, 565–587. [PubMed]
- Mach, E. (1959). *The analysis of sensations and the relation of the physical to the psychical*. New York: Dover. (Original work published 1906).
- Pizlo, Z. (1994). A theory of shape constancy based on perspective invariants. *Vision Research*, *34*, 1637–1658. [PubMed]

- Pizlo, Z., & Rosenfeld, A. (1992). Recognition of planar shapes from perspective images using contour-based invariants. *CVGIP. Image Understanding*, *56*, 330–350.
- Pizlo, Z., Rosenfeld, A., & Epelboim, J. (1995). An exponential pyramid model of the time-course of size processing. *Vision Research*, *35*, 1089–1107. [PubMed]
- Pizlo, Z., Rosenfeld, A., & Weiss, I. (1997a). The geometry of visual space: About the incompatibility between science and mathematics. *Computer Vision and Image Understanding*, *65*, 425–433.
- Pizlo, Z., Rosenfeld, A., & Weiss, I. (1997b). Visual space: Mathematics, engineering, and science. *Computer Vision and Image Understanding*, *65*, 450–454.
- Pizlo, Z., & Salach-Golyska, M. (1995). 3-D shape perception. *Perception & Psychophysics*, *57*, 692–714. [PubMed]
- Rothwell, C. A. (1995). *Object recognition through invariant indexing*. Oxford: Oxford University Press.
- Saunders, J. A., & Knill, D. C. (2001). Perception of 3D surface orientation from skew symmetry. *Vision Research*, *41*, 3163–3183. [PubMed]
- Sawada, T., & Pizlo, Z. (2007). Symmetry detection in 3D scenes. *Proceedings of SPIE*, *6498*, 64980Y/1–12.
- Stevens, K. A. (1979). Representing and analyzing surface orientation. In P. H. Winston & R. H. Brown (Eds.), *Artificial intelligence: An MIT perspective, Volume 2* (pp. 101–125). Cambridge: MIT Press.
- van der Vloed, G., Csathó, A., & van der Helm, P. A. (2005). Symmetry and repetition in perspective. *Acta Psychologica*, *120*, 74–92. [PubMed]
- Wagemans, J. (1992). Perceptual use of nonaccidental properties. *Canadian Journal of Psychology*, *46*, 236–279. [PubMed]
- Wagemans, J. (1993). Skewed symmetry: A nonaccidental property used to perceive visual forms. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 364–380. [PubMed]
- Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision*, *9*, 9–32. [PubMed]
- Wagemans, J. (1997). Characteristics and models of human symmetry detection. *Trends in Cognitive Sciences*, *1*, 346–352.
- Wagemans, J., Lamote, C., & van Gool, L. (1997). Shape equivalence under perspective and projective transformations. *Psychonomic Bulletin & Review*, *4*, 248–253.
- Wagemans, J., Vanden Bossche, P., Segers, N., & d'Ydewalle, G. (1994). An affine group model and the perception of orthographically projected planar random polygons. *Journal of Mathematical Psychology*, *38*, 59–72.
- Wagemans, J., van Gool, L., & d'Ydewalle, G. (1991). Detection of symmetry in tachistoscopically presented dot patterns: Effects of multiple axes and skewing. *Perception & Psychophysics*, *50*, 413–427. [PubMed]
- Wagemans, J., Van Gool, L., & d'Ydewalle, G. (1992). Orientation effects and component processes in symmetry detection. *Quarterly Journal of Experimental Psychology*, *44A*, 475–508.
- Watt, R. J. (1987). Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *Journal of the Optical Society of America A, Optics and Image Science*, *4*, 2006–2021. [PubMed]
- Yang, T., & Kubovy, M. (1999). Weakening the robustness of perspective: Evidence for a modified theory of compensation in picture perception. *Perception & Psychophysics*, *61*, 456–467. [PubMed]
- Zusne, L. (1970). *Visual perception of form*. New York: Academic Press.