Journal of Experimental Psychology:
Human Perception and Performance
2013, Vol. 39, No. 4, 1193–1204

0096-1523/13/$12.00 DOI: 10.1037/a0030837

Symmetry Facilitates Shape Constancy for Smoothly Curved 3D Objects

Young Lim Lee and Jeffrey A. Saunders
University of Hong Kong

We tested whether the presence of symmetry improves shape discrimination across changes in viewpoint and lighting for smoothly curved 3D objects. We constructed symmetric and asymmetric versions of random 3D shapes by manipulating their spherical harmonic representations. Matched objects had the same power spectra and appear highly similar except for the presence of symmetry. Observers discriminated sequentially presented pairs of either symmetric or asymmetric objects. Objects were presented in conditions that provided different 3D cues: shading only, stereo only, and combined shading and stereo. To control for 2D cues, standard and test objects had matched boundary contours and were rendered with different light sources. Test objects were also rotated in depth by variable amounts (0° to 60°). Across all viewpoint and 3D cue conditions, we found that shape discrimination for symmetric objects was better than for asymmetric objects. The symmetry benefit was not limited to monocular viewing or to conditions with large rotations in depth. In a second experiment, we blocked trials by viewpoint rotation to eliminate uncertainty in object orientation. This improved performance for asymmetric objects relative to symmetric objects, suggesting that symmetry contributes by providing a cue to object orientation. However, a symmetry advantage was still observed in all shape cue conditions, so this was not the sole source of benefit. Our results demonstrate that symmetry improves shape constancy for smooth 3D objects and suggest that one role of symmetry is to provide a reference orientation for an object.

Keywords: 3D shape, symmetry, stereo, shading

Perceiving an object to have a constant 3D shape across changes in viewpoint is a challenging problem because the image of an object can vary greatly depending on how it is viewed. A number of studies have found that rotation of a 3D object in depth makes it difficult to recognize the object (Bülthoff & Edelman, 1992; Rock & DiVita, 1987) or to discriminate from other objects with different 3D shapes (Lee & Saunders, 2011; Norman, Swindle, Jennings, Mullins, & Beers, 2009). Figure 1 shows an example of the same object viewed at two different orientations differing by 60° rotation in depth. Although the images produce a vivid 3D percept, it is difficult to perceive that it is the same object. Lee and Saunders (2011) tested 3D shape constancy for random objects like those shown in Figure 1 and observed significant viewpoint costs, even when 3D shape was specified by shading, specularities, and binocular cues. Increasing the information about 3D surface structure led to better shape discrimination in general, but did not reduce viewpoint costs.

Poor shape constancy for random objects like that shown in Figure 1 could be due to lack of structure in the objects rather than lack of information about 3D surface relief. Although Lee and Saunders (2011) observed poor shape constancy despite rich 3D cues, other studies have found that line drawings of objects can be sufficient for shape constancy across changes in viewpoint. Biederman and Gerhardstein (1993) found that observers could recognize line drawings of objects composed of simple volumetric primitives across 45° rotations in depth. Pizlo and Stevenson (1999) also observed good shape constancy for monocular line drawings of polyhedron objects. They found that observers could reliably discriminate symmetric polyhedra, or polyhedra with planar faces, across 90° rotations in depth. Performance for unstructured polyhedra was markedly worse. Chan, Stevenson, Li, and Pizlo (2006) tested similar line drawing stimuli under monocular and binocular viewing conditions, and found that binocular presentation provided little benefit for structured polyhedrons. These results suggest that structural properties of 3D objects, such as planarity of surface contours and symmetry, may be more important for reliable shape constancy than the amount of viewing information about 3D structure.

There are a number of ways that symmetry might contribute to 3D shape constancy. The presence of symmetry constrains the possible 3D interpretations of an image and thereby provides a 3D cue (Li, Pizlo, & Steinman, 2009; Pizlo, Sawada, Li, Kropatsch, & Steinman, 2010; Saunders & Knill, 2001; Sawada, 2010; Vetter, Poggio, & Bülthoff, 1994; Wagemans, 1993). Symmetry could constrain the interpretation of other 3D cues, resolving ambiguities and improving the accuracy of 3D shape estimation. Symmetry could also allow shape discrimination based on affine rather than metric structure (Vetter et al., 1994), and more efficient encoding of shape (Liu & Kersten, 2003).

Many 3D cues provide ambiguous information about metric 3D structure. Binocular disparities specify relative depths along a surface but require scaling by absolute distance to infer metric 3D structure.

This article was published Online First December 31, 2012.

Young Lim Lee and Jeffrey A. Saunders, Department of Psychology, University of Hong Kong, Hong Kong.

This work was supported by a grant from the Hong Kong Research Grants Council, GRF HKU-752010H. We thank Celina Wong and Xing Xing for assistance with data collection.

Correspondence concerning this article should be addressed to Jeffrey A. Saunders, Department of Psychology, The University of Hong Kong, 6/F, The Jockey Club Tower, Centennial Campus, Pokfulam Road, Hong Kong. E-mail: jsaun@hku.hk

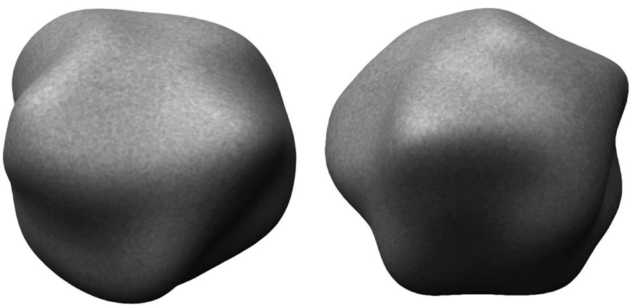

Figure 1. Images of the same 3D object from two different viewpoints. In the right image, the object has been rotated in depth around a vertical axis by 60°, relative to the orientation of the object in the left image. This amount of viewpoint change greatly impairs 3D shape discrimination, even when rich 3D cues are available (Lee & Saunders, 2011).

A number of studies have observed perceptual compression or expansion of depth in shape from stereo (e.g., Johnston, 1991; Norman, Todd, & Phillips, 1995; Todd & Norman, 2003). Similarly, shading indicates relative variations in 3D orientation but has a fundamental ambiguity when illumination is not known (Belhumeur, Kriegman, & Yuille, 1999). Some recent studies have found that judgments of local surface shape can be biased by changing the direction of illumination (Nefs, Koenderink, & Kappers, 2005; Caniard & Fleming, 2007). Such biases could potentially interfere with shape constancy across changes in viewpoint and illumination.

An additional constraint of symmetry could resolve the ambiguity of other 3D cues, thereby improving shape constancy. Saunders and Knill (2001) observed interactions between symmetry and stereo information for perception of 3D surface orientation, which could potentially improve perception of 3D shape from stereo. Li, Sawada, Shi, Kwon, and Pizlo (2011) observed veridical depth judgments from binocular views of symmetric polyhedra, which they suggest is due to symmetry and planarity constraints. Some recent computational models have shown that a symmetry assumption could potentially improve reconstruction of 3D shape from shading (Shimshoni, Moses, & Lindenbaum, 2000; Zhao & Chellappa, 2001).

For a symmetric object, recovery of 3D structure is not necessarily required for shape discrimination. Vetter et al. (1994) point out that the constrained set of 3D interpretations of a monocular image would generally be sufficient to distinguish the object from other objects with different shapes. Thus, even if symmetry did not directly interact with other 3D cues, it could still simplify the task of 3D shape discrimination.

Liu and Kersten (2003) suggested that 3D symmetric objects could be processed more efficiently than asymmetric ones in the human visual system. They showed that deformations of symmetric objects were more easily detected than deformations of an asymmetric object. They also found that a silhouette of a symmetric object already appears to be a 3D symmetric object, which suggests that symmetry has a role in representing 3D structure internally in the visual system.

The goal of the present study was to investigate the role of symmetry in 3D shape perception. We compared shape discrimination for symmetric and asymmetric objects that were closely matched in terms of complexity and variation, under conditions that provided rich 3D information even when symmetry is not present. We also systematically varied the amount of viewpoint change across views, and the presence of stereo and shading information.

Previous studies of the role of symmetry in shape discrimination have used artificial stimuli and restricted classes of objects: line drawings of random polyhedra (Pizlo & Stevenson, 1999; Chan et al., 2006), or randomly connected segments (Liu & Kersten, 2003). In the absence of symmetry, such stimuli provide little information about 3D structure. The role of symmetry for minimal stimuli like these may not be representative in the more natural situation of viewing solid objects with multiple sources of 3D information.

We tested shape discrimination for smoothly curved 3D volumetric objects, rendered with shading and highlights and presented in stereo. Figure 2 shows an example of a matched set of symmetric and asymmetric shapes used in our experiments. Pairs of asymmetric shapes were constructed to have the same boundary contour but qualitatively different 3D structure. Previous studies have found that observers are capable of reliably discriminating such shapes using either monocular or binocular cues for modest changes in viewpoint (Norman, Todd, & Orban, 2004; Lee & Saunders, 2011). The asymmetric shapes were used to generate symmetric shapes that were highly similar except for the presence of symmetry (see Method section). The presence of symmetry could improve 3D shape perception for these stimuli, but is not necessary—previous results indicate that other cues are sufficient for high performance in some situations.

The presence of symmetry might be especially helpful for discriminating shapes across large changes in viewpoint. To the

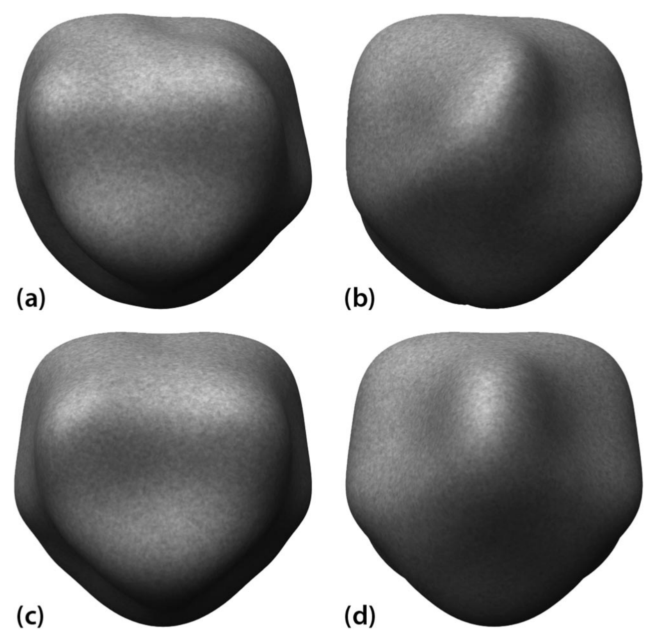

Figure 2. Example of a set of matched asymmetric and symmetric shapes. Two asymmetric shapes were constructed to have approximately the same boundary contour at a base viewing direction (a and b). The asymmetric shapes were represented as radial functions approximated by spherical harmonics. Symmetric shapes were generated from each asymmetric shape by rotating the phases of the spherical harmonic components to zero. These symmetric shapes were distorted in the x–y plane to make their boundary contours approximately the same, producing a matched pair of symmetric shapes (c and d).

extent that shape discrimination is impaired by view-dependent distortions in perceived 3D surface relief, such as compression along the line of sight, the benefit of additional 3D constraints would increase with viewpoint change. Shape discrimination for symmetric objects could also be based on an affine representation (Vetter et al., 1994), which would be robust to change in viewpoint. Previous studies of shape constancy by Pizlo and Stevenson (1999) and Chan et al. (2006) manipulated symmetry but tested only one amount of rotation in depth—90°. Consequently, one cannot determine whether the observed symmetry benefit is general or specific to the case of large viewpoint change. We measured shape discrimination across various amounts of rotation in depth for both symmetric and asymmetric objects to test whether symmetry reduces viewpoint dependence.

We also tested the relative benefit from stereo information for symmetric and asymmetric objects. Lee and Saunders (2011) observed a stereo benefit for asymmetric objects like those tested here, presented with rich monocular 3D cues. Stereo might be less advantageous when symmetry is present, because symmetric objects are easier to distinguish based on monocular images (Vetter et al., 1994). Chan et al. (2006) observed good shape discrimination performance for monocular line drawings of structured objects, and little benefit from stereo. Based on these results and theoretical grounds, Pizlo, Li, and Steinman (2008) argued that stereo information has little role in shape constancy for structured objects. On the other hand, cooperative interactions between symmetry and stereo have been observed for other type of 3D tasks (Saunders & Knill, 2001; Li et al., 2011).

For shaded objects, symmetry might contribute by resolving ambiguity in shading information. For example, symmetry might reduce illumination-dependent distortions such as that observed by Nefs et al. (2005). To test this possibility, we also included conditions with stereo information but no shading, shown in Figure 3. The stereo-only stimuli were binocular images of textured objects lit with only ambient illumination. If symmetry contributes to interpreting shading information, one would expect less benefit in these conditions than when shading is present.

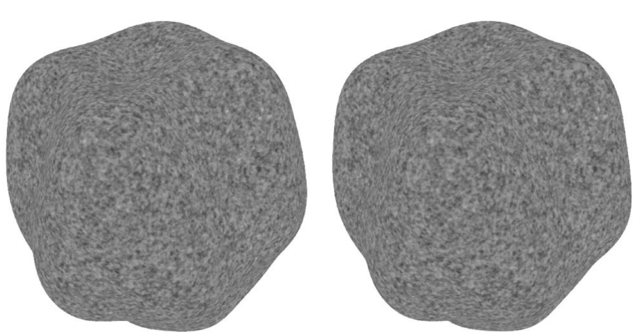

Figure 3. Example of a stereo-only stimuli. The two images are stereo views of a textured object with ambient illumination, with viewpoint rotation 15°. The right eye's image is shown on the left, and vice versa, to allow viewing with cross-fusion.

Experiment 1

Experiment 1 compared 3D shape discrimination for symmetric and asymmetric objects, for various changes in viewpoint, and under different monocular and binocular viewing conditions. Observers performed shape discrimination for sequentially presented objects, and the viewpoint of the test object was varied by 0°, ±15°, ±30° or ±60°. Objects were presented in three different shape cue conditions: shading only, stereo only, and combined shading and stereo. Based on the results of Lee and Saunders (2011), we expected significant viewpoint costs for the asymmetric objects and a consistent benefit from stereo. The presence of symmetry could potentially reduce viewpoint costs and reduce the benefit from stereo.

Method

Participants. Sixteen adults (six males and 10 females) at the University of Hong Kong participated in the experiment. All participants had normal or corrected-to-normal vision and passed a stereo acuity screening test. All participants were naïve as to the purpose of the study and were paid for participating. The procedures were approved by and conform to the standards of the Human Research Ethics Committee for Non-Clinical Faculties.

Apparatus and stimuli. Twelve sets of objects were used for the experiment and each set of objects included a pair of symmetric and a pair of asymmetric objects. Each object in a pair was constructed to have approximately the same boundary contour when viewed at a base orientation.

Random asymmetric 3D shapes were generated using a method similar to Lee and Saunders (2011). A sequence of 10 sinusoidal distortions was applied to a unit sphere, each in a different, randomly chosen direction. The distortion was of the form (x, y, z) → (x, y, z + 0.075 sin[1.6 x]), with the coordinate frame rotated to a different, random 3D orientation for each sinusoidal distortion. The distorted unit sphere was represented as a radial function and fit by spherical harmonics, with degrees up to $l = 16$. The shapes were smoothed by applying Gaussian blur to the radial function, with $\sigma = 4.8°$. Shapes were scaled to have an average radius of 10 cm, corresponding to about 9° to 10° of visual angle at the viewing distance. Shapes were further rotated in the x–y plane (around the line-of-sight axis) to be maximally symmetric, as measured by the average difference in radius for points with matching spherical coordinates. A large pool of asymmetric shapes was generated in this manner.

We then computed the projected boundary contour for each shape and selected pairs of shapes with similar contours. These pairs of 3D shapes were distorted around the line-of-sight axis to force their projected contours to be closely matched. To minimize artifacts, the boundary-matching distortion was forced to be smooth by applying a Gaussian blur with $\sigma = 4.8°$. We manually excluded some pairs of shapes for various reasons: Contour matching was not successful or produced visible artifacts, or a shape had an atypical amount of 3D variation.

Symmetric shapes were generated from asymmetric shapes. We first represented each asymmetric shape as a radial functions on a sphere, expressed as a spherical harmonic expansion, with $l_{max} = 20$. The form of the radial function is

$$r(\theta, \phi) = \sum_{l=0}^{l_{max}} \sum_{m=-1}^{l} a_{lm} Y_{lm}(\theta, \phi)$$

where $\boldsymbol{r}(\theta, \phi\square)$ is radius, θ is azimuth, $\square$ is ϕ elevation, $\boldsymbol{Y_{lm}}(\theta, \phi\square)$ are the real spherical harmonic basis functions, and $\boldsymbol{a_{lm}}$ are the set of coefficients representing a particular object. The basis functions include pairs of cosine-type and sine-type functions, $\boldsymbol{Y_{lm}}$ and $\boldsymbol{Y_{l\text{-}m}}$, which are identical except for a phase rotation around the z-axis. If no sine-type components are present (all $\boldsymbol{a_{l\text{-}m}} = 0$), then the function is bilaterally symmetric. We generated a symmetric shape from an asymmetric shape by replacing pairs of coefficients ($\boldsymbol{a}_{\mathrm{lm}}$, $\boldsymbol{a}_{\mathrm{l-m}}$) with pairs that with the same total energy but no sine-type component, ($\sqrt{[\boldsymbol{a_{l-m}}^2 + \boldsymbol{a_{lm}}^2]}$, 0). This results in a bilaterally symmetric shape with the same spherical power spectra as the original asymmetric shape. For each pair of asymmetric shapes, a pair of symmetric shapes was created in this manner. The symmetric shapes in a pair were then distorted around the line of sight to force their boundary contours to match, in the same way as for pairs of asymmetric shapes. Figure 2 shows an example of matched pairs of symmetric and asymmetric objects.

Objects were simulated to have a combination of Lambertian and specular reflectance. We used a Phong model with a 5% ambient component, 70% Lambertian component, and 25% specular component with an exponent of 100. The Lambertian component was simulated to have a homogeneous surface texture, which modulated reflectance by between 60% to 100%. To create the texture for an object, we first generated a random set of 40,000 points that were uniformly distributed on the surface. These points were used as centers to form a Voronoi tiling of the surface and each tile was assigned a random reflectance. The resulting surface pattern was then approximated as a cube-map texture, with a resolution of 1024 × 1024 for each side. We repeated this procedure 10 times per object and averaged the results to get the final cube-map texture used to render the object.

There were three shape cue conditions: (a) shading only, (b) stereo only, and (c) combined shading and stereo cues. For the conditions with shading information, we simulated a diffuse point source light at an infinite distance. The illumination map was a Gaussian distribution, with width $\sigma = 30°$, centered around the light source direction. Four light source directions were used: (±.6209, .7399, .2588) and (±.2962, .8138, .5). The light source directions for standard and comparison images on a trial were always diagonally different to ensure that the task could not be performed based on 2D image similarity. For example, if the light source direction for standard image was (−.6209, .7399, .2588), the light source direction for comparison image would be (+.2962, .8138, .5). For the stereo-only condition, illumination was simulated to be ambient, with brightness equal to 60% of the maximum illumination in the shading conditions (see Figure 3). This brightness approximately matches the mean luminance contrast of the texture in the conditions with and without shading. Because the texture used in the stereo-only condition was fine and simulated illumination was ambient, the mean luminance in regions of the image was constant regardless of variations in texture.

A mirror stereoscope was used to present images. Observers viewed a pair of LCD monitors (Dell SP2208WFP) through two semisilvered mirrors positioned near the eyes and slanted 45° relative to the line of sight. The monitors had a visible region of 47 cm × 29.5 cm, with 1680 × 1050 resolution and a frame rate of 60 Hz. The monitors were positioned so that their virtual surfaces (viewed through the mirror) were frontal relative to the viewer and aligned at a distance of 60 cm. We measured interpupillary distance for each observer to compute the accurate stereo projections when rendering. In the shading-only condition, the same apparatus was used but observers wore an eye patch covering their nondominant eye. Objects were rendered so that their centers were at the distance of the screen, 60 cm from the observer.

Procedure. Observers performed a same–different 3D shape discrimination task. On each trial, a standard object was presented for 2 s, followed by a 600-ms random noise mask that covered the whole screen, and then a test object. Observers judged whether the test object was the same or different from the standard object. The test object was equally likely to be the same or different, and different objects had the same boundary contour at base viewpoint. The test object was presented either from same viewpoint or differed in rotation in depth by ±15°, ±30°, or ±60° relative to the standard object's viewpoint. Figure 4 shows examples of same and different pairs with different changes in viewpoint. The test object remained on the screen until observers' response, and observers had the option to repeat the standard and test sequence if desired. In stereo condition, a fixation point at the screen distance (60 cm) was shown prior to presenting each object.

For each shape cue condition, observers performed practice trials with feedback prior to the experimental session to become familiar with the stimulus and the task. During the practice session, one pair of symmetric objects and one pair of asymmetric objects were used, and these objects were not used in the experimental sessions. No feedback was given on experimental trials.

We divided the 12 sets of objects into two equal subsets. For half of the observers, we used six pairs of symmetric objects from the first subset and six pairs of asymmetric objects from the second subset, and for the other observers, we used asymmetric objects from the first subset and symmetric objects from the second subset. Thus, observers never saw both the symmetric and asymmetric versions of an object.

Different shape cue conditions were tested on separate days, with order counterbalanced across observers. In an experiment session, observers performed 28 practice trials and 336 experimental trials (12 object pairs × 7 viewpoint rotations × 2 same/different). Conditions were fully randomized within sessions. Trials were self-paced and observers were given breaks.

Within a session, each individual object was used as a test object on seven "same" trials and seven "different" trials, one time at each viewpoint rotation, and as a standard object in base viewing orientation (0°) on seven "different" trials. Over the course of an experiment, an object was presented a total of 24 times at base orientation and 18 times at other combinations of viewpoint rotation and shape cues.

Analysis. We conducted a hierarchical Bayesian analysis rather than a standard ANOVA. This analysis avoids the problem of inflated Type I error when doing multiple comparisons and has other advantages (Kruschke, 2010; Shiffrin, Lee, Kim, & Wagenmakers, 2008).

To model responses from a given subject and condition, we followed Lee (2008). The number of hits and false positives in a condition were assumed to be binomially distributed:

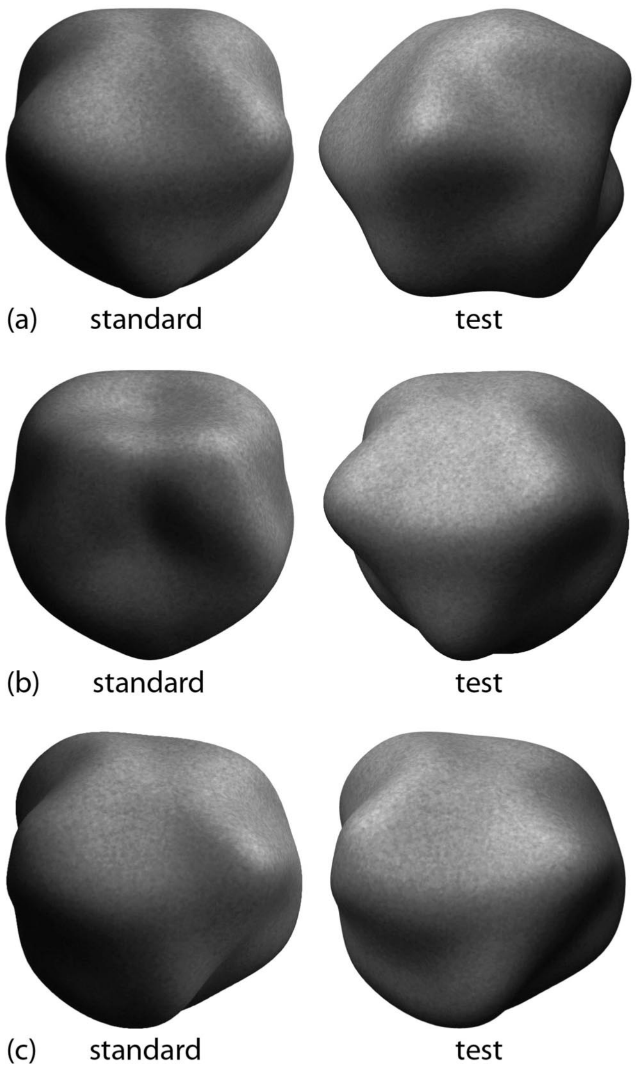

Figure 4. Standard and test objects for sample trials. (a) "Same" trial with symmetric object and 60° viewpoint rotation. (b) "Different" trial with symmetric object and 15° viewpoint rotation. (c) "Same" trial with asymmetric object and 0° viewpoint rotation. Note that for "different" trials, the standard object and the same boundary contour as the base view of the test object, so a direct comparison of 2D boundary contours would not be informative for the task. To further discourage a 2D strategy, different light source directions were used for the standard and test objects.

$$\mathrm{H} \sim \mathrm{Binomial}(h,\, N_{\mathrm{same}}),\ F \sim \mathrm{Binomial}(f,\, N_{\mathrm{diff}})$$

where H and F are the observed number of hits and false-positives, h and f are underlying hit and false-positive rates, and N_{same} and N_{diff} are the total number of "same" and "different" trials. Hit rate and false-positive rate were assumed to be related to sensitivity (d') and response bias (b) in a standard way:

$$d' = \Phi(h) - \Phi(f),\ b = -[1/2]\,(\Phi(h) + \Phi(f))$$

where Φ(p) is the inverse of the cumulative normal distribution. Each condition was assumed to have a separate mean sensitivity and response bias ($dcond_{\mathrm{ij}}$), and individual subjects were assumed to have idiosyncratic differences in sensitivity and response bias that varied depending on shape cue and viewpoint rotation but not symmetry ($dindiv_{\mathrm{jk}}$). The sensitivity and response bias for a condition and subject were

$$d_{\mathrm{ijk}} = dcond_{\mathrm{ij}} + dindiv_{\mathrm{jk}},\ b_{\mathrm{ijk}} = dcond_{\mathrm{ij}} + dindiv_{\mathrm{jk}}$$

where i is symmetric versus asymmetric, j is shape cue and viewpoint rotation condition, and k is individual subjects. For the condition mean parameters, we used a widely spread normal distribution as a prior, with a variance of 100 units of sensitivity/bias. Individual difference parameters were assumed to be normally distributed around zero, with an unknown variance (σ^2) that was fit during analysis. We assumed a broad, decreasing prior on σ: a gamma distribution with a mode of 0.25, a standard deviation of 1.0, and mean of 1.13. We tested other broadly distributed priors and found near-identical results.

Samples from the estimated posterior probability were computed using JAGS, an open-source software package for performing Gibbs sampling (Plummer, 2003). We computed a Monte Carlo Markov Chain (MCMC) with 100,000 samples, after a burn-in period of 10,000 samples, and then thinned by selecting every 10th sample to produce a final sample of size 10,000. Diagnostic measures indicated that the chain converged, and there was no indication of autocorrelation in the thinned samples.

In this approach, the distributions of differences or other contrasts are used to evaluate whether effects are statistically reliable. For example, to test whether two conditions have different means, one can compute the difference $dcond_1 - dcond_2$ for each sample, and then test whether this distribution is almost entirely positive or entirely negative. Specifically, we used the 95% highest density interval (HDI) of a distribution to characterize the range of credible values (Kruschke, 2010). If the HDI of a difference lies entirely above zero, then there is a probability of 95% or more that the difference is positive, given the data and model assumptions. Similarly, if the HDI is entirely below zero, it is strong evidence for a negative effect. When the HDI of a difference overlaps zero, the evidence for a difference is inconclusive. To test for the main effect of a variable in a factorial design, one can apply the same approach but marginalize over the other variables before computing differences.

Results

Figure 5 shows mean sensitivity and response bias measures as a function of viewpoint rotation for each shape cue and object type. These values correspond to the means of the d' and b parameters for conditions in the Bayesian model, averaged across samples of the posterior distribution. Error bars depict the standard deviation of the distributions, which is comparable with the standard error of the mean. To check the fits, we also computed sensitivity and response bias separately for individual subjects and computed the mean and standard errors, and found no qualitative difference in results.

Sensitivity was higher overall for symmetric objects compared with asymmetric objects. Otherwise, the pattern of results was similar. For both types of objects, viewpoint rotation decreased sensitivity and changed response bias, and stereo provided an overall benefit. To test the reliability of these effects, we computed contrasts from the posterior samples and assessed their distributions.

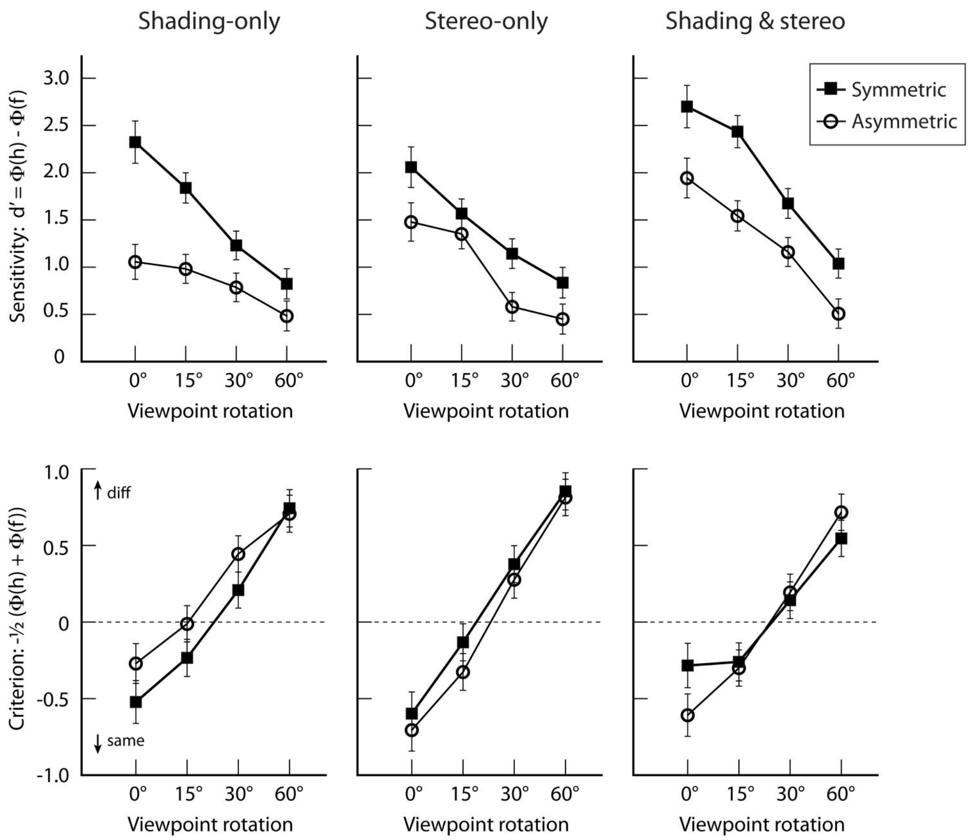

Figure 5. Mean sensitivity (top row) and response bias (bottom row) measures as a function of viewpoint rotation in Experiment 1. The three graphs on each row correspond to the different shape cue conditions: shading only (left), stereo only (middle), and combined shading and stereo (right). The two curves on each graph plot results for symmetric objects (solid squares) and asymmetric objects (open circles). Condition means were estimated with a Bayesian analysis (see Method section). Error bars depict ± 1 *SD* of the estimated posterior probability distributions, which is comparable with the standard error of the mean.

Symmetry effects. Figure 6 shows measures of the effect of symmetry on sensitivity for the different combinations of shape cue and viewpoint rotation. For each sample, we computed the difference in sensitivity between matched symmetric and asymmetric conditions, $d_{\text{sym}} - d'_{\text{asym}}$. The graphs plot the estimated posterior distribution of these differences. The 95% HDIs are depicted as shaded regions. This interval can be interpreted as the range of differences that is plausible, given the data and model assumptions. In almost all conditions, the HDIs are well above zero. The only exception was stereo-only condition, with 15° viewpoint rotation, which showed a trend in the same direction. There was also an interaction between symmetry and viewpoint for the shading-only condition: The symmetry effect decreased with amount of viewpoint rotation. There was no evidence for any other interactions.

Viewpoint effects. In all shape cue conditions, and for both symmetric and asymmetric objects, sensitivity decreased with viewpoint rotation. To evaluate this effect statistically, we computed the slope of the linear effect of viewpoint on sensitivity for each posterior sample. Figure 7 shows distributions of the viewpoint effect for symmetric and asymmetric objects in the three shape cue conditions. The HDIs are below zero in all cases, providing evidence for a significant negative slope. For the shading-only condition, the viewpoint effect was larger for symmetric objects than asymmetric objects (top rows). For the other conditions, the viewpoint effects were comparable for symmetric and asymmetric objects.

Viewpoint also had an effect on the response bias, corresponding to a bias toward judging objects as "same" when viewpoint change was small, and a bias toward judging objects as "different" when viewpoint change was larger. We evaluated the reliability of this effect by computing the linear effect of viewpoint for the response bias parameters of the posterior samples. The HDIs of these distributions were entirely negative for both symmetric and asymmetric objects, and for all shape cue conditions.

Stereo benefit. Stereo improved performance for both symmetric and asymmetric objects. To evaluate this effect, we computed the difference in sensitivity between the shading-only condition and the combined shading and stereo conditions, $d'_{\text{comb}} - d'_{\text{shading}}$, for all posterior samples. Figure 8 shows the distributions of the stereo effect for symmetric and asymmetric objects. The top two graphs show results averaged across viewpoint rotation, corresponding to a main effect of stereo. The HDIs for symmetric and asymmetric objects are entirely above zero, indicating a stereo benefit, and they overlap each other, indicating

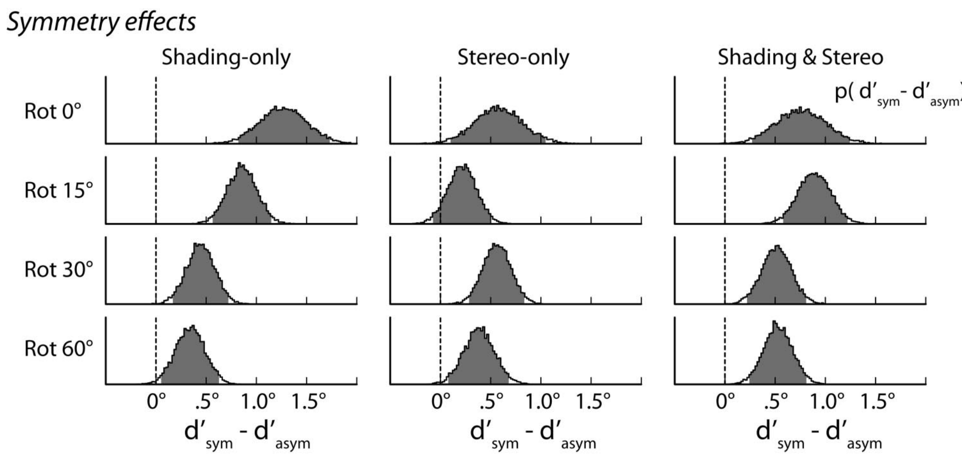

Figure 6. Effects of symmetry on sensitivity (d') in Experiment 1. The graphs plot estimated posterior probability distributions for the difference between asymmetric and symmetric conditions, $d'_{\mathrm{sym}} - d'_{\mathrm{asym}}$, for the different combinations of shape cue and viewpoint rotation. The shaded area under each curve depicts the 95% highest density interval (HDI) of the distribution. In almost all cases, the HDI lies entirely above zero, indicating a symmetry benefit.

that the effect was of comparable magnitude. The other graphs show the distribution of stereo effect for separate viewpoint rotation conditions. These distributions are more spread and the HDIs overlap zero in some conditions. Thus, although there is strong evidence for an overall stereo benefit in the case of both symmetric and asymmetric objects, we cannot determine from the present data whether the benefit was present for all viewpoint rotation conditions or for only a subset.

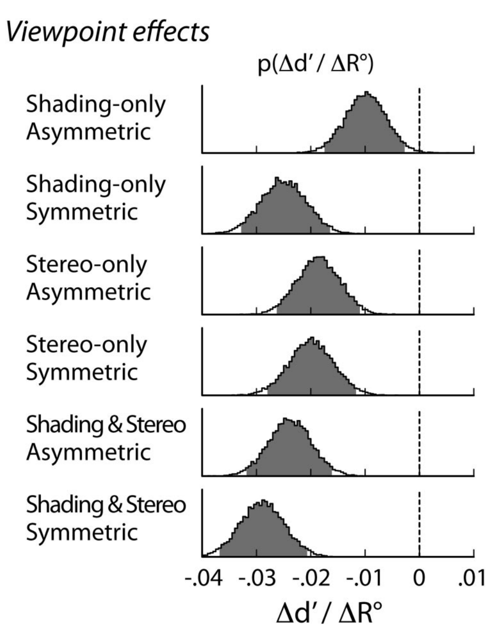

Figure 7. Linear effects of viewpoint rotation on sensitivity (d') in Experiment 1. The graphs plot the estimated distributions of the slope of d' as a function of viewpoint rotation for each shape cue and symmetry condition. The shaded area depicts the 95% highest density interval (HDI). In all cases, the HDI lies entirely below zero, indicating a significant viewpoint effect.

Discussion

Symmetry provided a general benefit for 3D shape discrimination. Sensitivity was higher for symmetric objects than for closely matched asymmetric shapes in all shape cue conditions, including the full cue

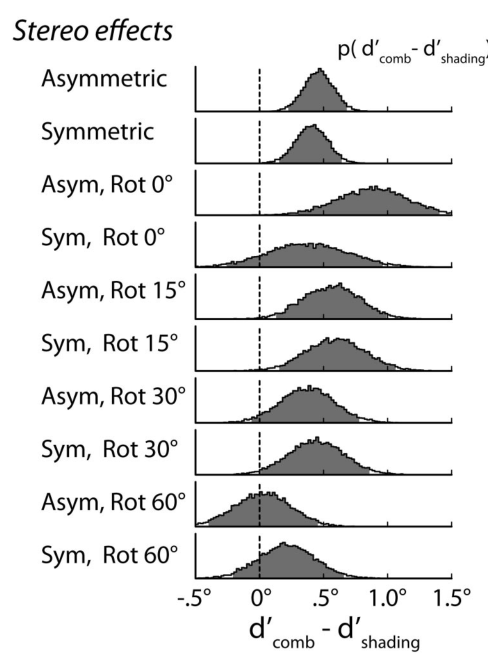

Figure 8. Effect of stereo on sensitivity (d') in Experiment 1. The graphs plot estimated distribution for the difference between the combined shading and stereo condition and the shading-only condition. Shaded areas depict 95% highest density interval (HDI). The top two graphs show overall stereo effect for asymmetric and symmetric objects after marginalizing over viewpoint rotation, and the other graphs show results for separate viewpoint rotation conditions.

condition that provided both shading and stereo information, and a symmetry benefit was observed both with and without change in viewpoint. Thus, the benefit from symmetry was not restricted to the situation of monocular viewing or to particular viewpoint conditions.

Symmetric and asymmetric objects produced a similar pattern of results across viewpoint change and shape cue conditions. For both types of shapes, rotation in depth produced a marked reduction in performance and binocular information provided an overall benefit. The results for asymmetric objects replicate the results of Lee and Saunders (2011), which tested similar objects and the same shape cue and viewpoint conditions. There were only a few differences in results for symmetric objects. Viewpoint change had a larger effect for symmetric objects in the shading-only condition, and there was not reliable evidence for a stereo benefit for symmetric objects in the 0° viewpoint rotation condition. Otherwise, the pattern of results was very similar and the primary difference was overall sensitivity.

Observers showed a viewpoint-dependent response bias for both types of objects, which was also found by Lee and Saunders (2011). Observers were biased toward judging objects to be the same when there was no change in viewpoint, and biased toward judging objects as different with a large change in viewpoint. The "same" bias may be due to the fact that pairs of objects had matching boundary contours when viewed at the base (0°) viewpoint and therefore had similar overall appearance. The "different" bias indicates that observers had difficult recognizing objects as being the same after a large change in viewpoint.

If symmetry contributed as an additional cue to 3D shape, one might expect symmetry to be especially beneficial for large changes in viewpoint. Contrary to this prediction, however, we observed that the effect of symmetry on sensitivity was as large for 0° viewpoint change as for 60° viewpoint change. In the case of the shading-only condition, the symmetry effect was actually larger in the condition without viewpoint rotation. Thus, although symmetry provided an overall benefit for shape discrimination, it did little to improve the cost of change in viewpoint.

An alternate explanation for the observed symmetry benefit is that symmetry helped to identify the change in orientation across views of an object. For our asymmetric, smoothly curved 3D objects, it might be hard to identify corresponding points across views. For symmetric objects, the symmetry plane potentially provides a salient reference frame, which could facilitate matching across views and improve shape discrimination. This possible explanation for a symmetry benefit is tested in Experiment 2.

Experiment 2

The presence of symmetry could potentially provide a cue to the orientation of an object, making it easier to compare features of an object across changes in viewpoint. In Experiment 2, we tested this possibility by blocking trials according to viewpoint change. If uncertainty about viewpoint change impaired shape discrimination for asymmetric objects in Experiment 1, removing this uncertainty by blocking should improve performance. An improvement would not necessarily be expected for symmetric objects, however, because presence of symmetry could help to identify viewpoint change even if randomized across trials. If this underlies the symmetry benefit observed in Experiment 1, then blocking would provide little benefit for symmetric objects compared with asymmetric objects. The symmetry benefit observed previously would therefore be reduced or eliminated. If symmetry contributes to shape discrimination in some other manner, such as by constraining the 3D interpretations of images, then a symmetry benefit would still be expected.

Method

Participants. Seventeen adults (five males and 12 females) at the University of Hong Kong participated in the experiment. None of them participated in Experiment 1. One participant was excluded from analysis based on unusually poor performance for 0° viewpoint change condition with full cues. Most participants are capable of good performance in this condition, so we suspected that this participant did not fully understand the task. All participants had normal or corrected-to-normal vision and passed a stereo acuity screening test. All were naïve as to the purpose of the study and were paid for participating. The procedures were approved by and conform to the standards of the Human Research Ethics Committee for Non-Clinical Faculties.

Apparatus and stimuli. Experiment 2 used the same objects, stimulus parameters, and viewing apparatus as the previous experiment. Objects were also presented in the same three shape cue conditions: shading only, stereo only, and combined stereo and shading.

Procedure. The procedure was the same as before, except that trials were blocked according to viewpoint rotation and we only tested positive rotations. Twelve pairs of symmetric and 12 pairs of asymmetric objects were randomly intermixed within blocks. Observers were told that the rotation was constant within blocks and were allowed 16 practice trials with feedback prior to the experimental blocks. The practice trials used different objects than the experimental blocks. No feedback was given on experimental blocks. Each experimental block consisted of 96 trials (48 total objects × 2 same/different trials) with a constant viewpoint rotation: 0°, 15°, 30°, or 60°. The order of blocks was randomized.

We applied the same Bayesian analysis to the data as in Experiment 1. The likelihood model and priors were identical. We used JAGS to compute a MCMC of 100,000 samples, which was thinned to 10,000 samples for the final estimates of condition means and posterior distributions.

Results

Figure 9 shows mean sensitivity and response bias measures from Experiment 2 as a function of viewpoint for the three shape cue conditions. As in the previous experiment, shape discrimination was better overall for symmetric objects, and stereo improved performance for both symmetric and asymmetric objects. However, blocking by viewpoint rotation also improved performance, especially for asymmetric objects.

Figure 10 shows the estimated effect of symmetry on sensitivity for both Experiment 1 (light gray) and Experiment 2 (dark gray), marginalized in different ways. Figure 10a shows the distribution of $d'_{\mathrm{sym}} - d'_{\mathrm{asym}}$ averaged across all shape cue and viewpoint rotation conditions. The HDI of this distribution for Experiment 2 is entirely above zero, indicating a significant symmetry benefit, and entirely below the HDI of Experiment 1, indicating that the overall benefit was smaller when conditions were blocked by viewpoint rotation. Figures 10b through 10d show distributions for separate shape cue conditions, averaged across viewpoint rotation conditions, and Figures 10e through 10h show distributions for

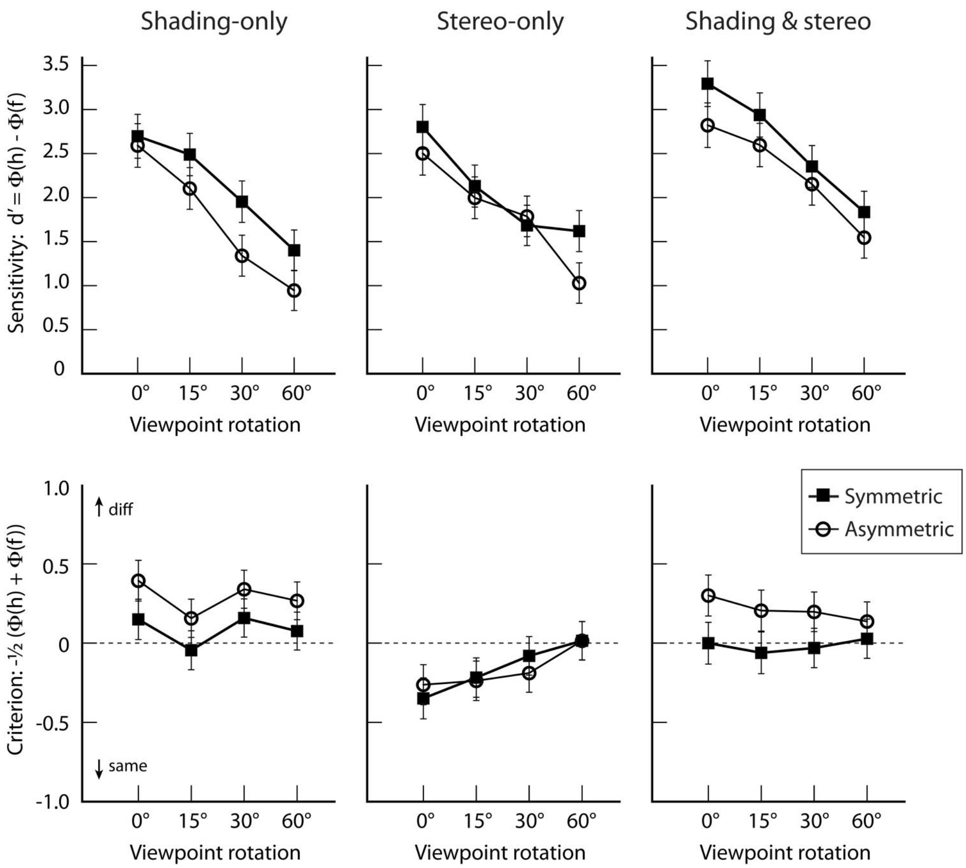

Figure 9. Mean sensitivity (top row) and response bias (bottom row) measures as a function of viewpoint rotation in Experiment 2. The three graphs on each row correspond to the different shape cue conditions: shading only (left), stereo only (middle), and combined shading and stereo (right). The two curves on each graph plot results for symmetric objects (solid squares) and asymmetric objects (open circles). Condition means were estimated with a Bayesian analysis (see Methods). Error bars depict ± 1 *SD* of the estimated posterior probability distributions, which is comparable with the standard error of the mean.

separate viewpoint rotations, averaged across shape cue conditions. In all cases and for both experiments, the difference in sensitivity was positive. For conditions with shading (b, d), and when viewpoint rotation was low (e, f), the symmetry effect was consistently smaller in Experiment 2. However, this interaction was not reliable in the other cases or when individual shape cue and viewpoint rotation conditions were analyzed. Thus, blocking by viewpoint in Experiment 2 reduced the benefit from symmetry, but we cannot determine whether this interaction occurred for all shape cue and viewpoint conditions or for only a subset.

Another effect of blocking was a change in the pattern of response bias. In Experiment 1, observers showed a strong bias toward "same" with no viewpoint change and a bias toward "different" with large viewpoint change. In Experiment 2, response bias showed little, if any, viewpoint dependence. To assess this interaction statistically, we computed the linear effect of viewpoint on response bias for all posterior samples and compared the distributions across experiments. In Experiment 1, the HDIs of the linear effect of response bias were entirely positive for all conditions. In Experiment 2, the HDI was above zero only for the stereo-only condition with symmetric objects, and in all conditions, the HDIs were entirely lower than in Experiment 1. By removing uncertainty about object orientation in Experiment 2, the viewpoint-dependent response biases observed in the previous experiment were reduced or eliminated.

Sensitivity decreased with viewpoint rotation in all conditions and there was no apparent interaction with symmetry or shape cue condition. Estimated distributions of the linear effect of viewpoint on sensitivity were entirely negative for all conditions (not shown), and overlapped for different symmetry and shape cue conditions.

Stereo provided a benefit for both symmetric and asymmetric objects. Figure 11 shows the effect of stereo for symmetric and asymmetric objects, averaged across all viewpoint rotation conditions (top) and for separate viewpoint rotation conditions. As in Experiment 1, there is strong evidence for an overall stereo benefit for both symmetric and asymmetric objects, but we cannot determine whether this benefit is present for all viewpoint rotations. The magnitudes of the effects were comparable with the previous experiment.

Discussion

We found that blocking by viewpoint conditions improved shape discrimination, indicating that uncertainty about object ori-

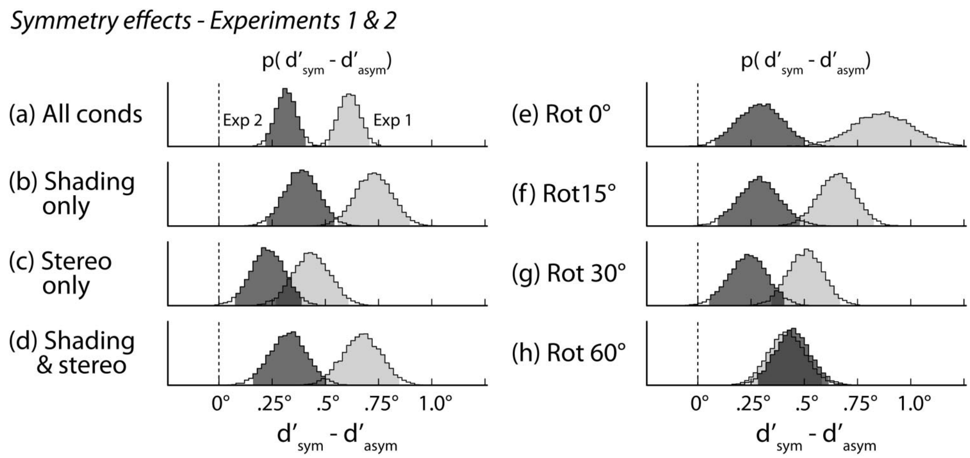

Figure 10. Effects of symmetry on sensitivity in Experiment 2 (dark gray) and Experiment 1 (light gray), shown as distributions of $d'_{\text{sym}} - d'_{\text{asym}}$ marginalized in different ways: (a) averaged across all conditions; (b-d) separate shading, stereo, and combined shape cue conditions, averaged across viewpoint rotations; (e-h) separate viewpoint rotation conditions, averaged across shape cue conditions. The shaded area under each curve depicts the 95% highest density interval (HDI) of the distribution. There is evidence for a symmetry benefit in all cases. The symmetry effect was smaller overall in Experiment 2. This interaction was observed for the shading and combined shape cue conditions (b, d) and for low viewpoint rotation (e, f).

entation did impair performance in the previous experiment. In Experiment 1, observers showed a bias toward judging same objects as “different” when there was a large change in viewpoint, and a bias toward judging different objects as “same” when viewpoint was constant. Blocking by viewpoint in Experiment 2 largely eliminated these biases. Sensitivity was also increased for both types of objects.

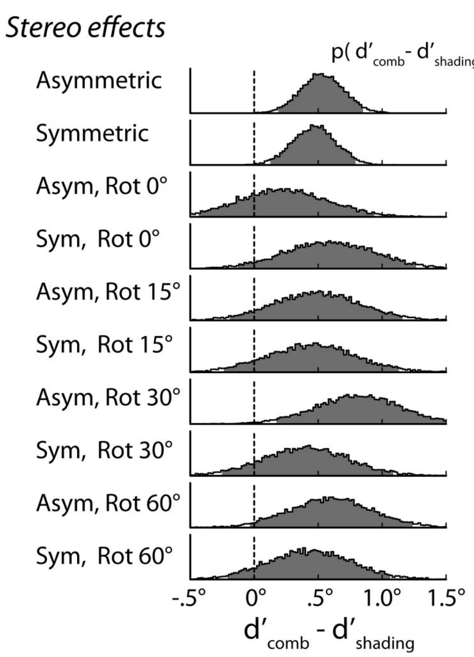

Figure 11. Effect of stereo on sensitivity (d') in Experiment 2. The graphs plot estimated distributions for the difference between the combined shading and stereo condition and the shading-only condition. Shaded areas depict 95% highest density interval (HDI). The top two graphs show overall stereo effect for asymmetric and symmetric objects after marginalizing over viewpoint rotation, and the other graphs show results for separate viewpoint rotation conditions.

The improvement in sensitivity across experiments was greater for asymmetric objects than symmetric objects, which suggests that some of the symmetry advantage observed in Experiment 1 is due to the use of symmetry as a cue to object orientation. Determining viewpoint change across two views of an unknown random object is potentially difficult (e.g., Figure 2), so knowledge of viewpoint change would be advantageous. For symmetric objects, the symmetry plane provides a salient reference frame for the object, so object orientation is better specified, even if not known in advance. Eliminating uncertainty about object orientation by blocking would therefore confer comparatively less advantage for symmetric objects than asymmetric objects, consistent with our results.

Although the symmetry effect was smaller than in the previous experiment, there remained a significant symmetry benefit across shape cue and viewpoint conditions. The symmetry benefit was present for both monocular and stereo conditions, and was not restricted to conditions with large viewpoint change. The other main findings of Experiment 1 were also replicated. Sensitivity decreased with viewpoint rotation in all conditions, and the rate was similar for symmetric and asymmetric objects. Stereo provided an improvement for both types of objects, though possibly not when viewpoint rotation is large. The main differences in results across experiments were the higher overall performance, the smaller symmetry effect on sensitivity, and the different pattern of response biases.

General Discussion

Symmetry Benefit

Our main goal was to test whether the presence of symmetry improves 3D shape discrimination across changes in viewpoint.

We simulated symmetric and asymmetric objects that were highly similar except for the presence of symmetry, and images of objects provided rich monocular and binocular information about 3D surface structure. We found that symmetry provided a consistent benefit for shape discrimination, both for small and large change in viewpoint.

Compared with previous studies of symmetry and 3D shape discrimination, our symmetric stimuli provided richer information about 3D surface structure. The previous studies of symmetry and 3D shape discrimination by Pizlo and Stevenson (1999) and Chan et al. (2006) used line drawings of polyhedra as stimuli. For unstructured polyhedra, these line drawing stimuli provide little, if any, monocular 3D information. In contrast, our monocular stimuli elicit a vivid 3D percept due to shading, specularities, and the occlusion contour. Our stimuli also likely provided better binocular 3D information. The objects were covered with texture, providing disparity information throughout the image, and the binocular occlusion contours could provide an additional cue. Thus, one contribution of our study is to demonstrate that symmetry improves shape discrimination even when rich information about 3D surface relief is present.

We varied change in viewpoint as well as symmetry and 3D shape cues, and found that rotation in depth impaired shape discrimination across all conditions. Although the presence of symmetry improved shape discrimination, symmetry did not significantly reduce viewpoint effects. Shape discrimination remained highly viewpoint dependent, even for symmetric objects presented binocularly and even when change in viewpoint could be anticipated in advance.

Role of Symmetry

The different results of Experiments 1 and 2 demonstrate that some of the difficulty in 3D shape constancy across viewpoints is due to uncertainty in viewpoint change, and that the presence of symmetry can reduce this difficulty. For a symmetric object, the symmetry plane provides a reference frame for the object, which could facilitate comparison across different views. For a random 3D object with no global structure, associating different views is potentially more difficult, adding an additional source of uncertainty. This could explain why removing uncertainty in viewpoint provided more benefit for asymmetric objects, resulting in a smaller symmetry effect in Experiment 2. Recent results by Egan, Todd, and Phillips (2012) are consistent with this interpretation. They tested the effect of adding a visible reference line to indicate the orientation of an object and found that this improved shape discrimination for asymmetric objects but not symmetric objects. Although we found that uncertainty about viewpoint affected performance for both types of objects, the interaction observed by Egan et al. (2012) is in the same direction as observed across our experiments.

One implication is that previous studies may have underestimated observers' ability to perceive 3D shape constancy. In previous studies of 3D shape discrimination across change in viewpoint, either the direction and angle of viewpoint rotation was randomized (Lee & Saunders, 2011; Norman, Bartholomew, & Burton, 2008; Norman et al., 2009) or the axis of viewpoint rotation was randomized (Chan et al., 2006). Our results and those of Egan et al. (2012) indicate that shape discrimination can be improved by knowledge or information specifying the orientation of an object, particularly if the object is random and unstructured.

We observed better shape discrimination for symmetric objects than asymmetric objects, even when uncertainty in viewpoint change was eliminated in Experiment 2, so this is not the sole contribution of symmetry. There are a number of other possible explanations for a symmetry benefit and our results do not strongly constrain the possibilities. Symmetry could potentially contribute by resolving depth ambiguity in stereo information. However, some benefit from symmetry was observed in conditions with no change in viewpoint, for which symmetry does not constrain the magnitude of depth relief. Symmetry could also contribute to resolving ambiguous shading information, which would be applicable, even with no change in viewpoint. However, a benefit from symmetry was also observed in the stereo-only condition, so this cannot be the sole explanation. Thus, the additional benefit from symmetry appears to be relatively general or due to a combination of factors. Further study would be required to distinguish possible mechanisms for this benefit.

Stereo Benefit

Consistent with Lee and Saunders (2011), we observed a benefit from stereo even when monocular information from shading and specularities are available. The present data do not clearly establish a stereo benefit at all separate viewpoint rotation conditions. This may simply reflect lack of power. The trend was positive in all cases, and the overall effect was clearly reliable for both types of objects.

Pizlo and colleagues have argued that stereo information has little role in shape perception for structured objects (Chan et al., 2006; Pizlo et al., 2008). Our results are not consistent with this view. If stereo contributes to shape perception primarily when an object is unstructured, one would expect more benefit from stereo for our asymmetric objects. However, we found that the benefit from stereo was similar for symmetric and asymmetric objects. There was no indication of a substantive interaction; both symmetry and stereo appeared to provide an overall benefit for shape discrimination.

Conclusions

We found that symmetry provides a benefit for discriminating random 3D shapes across a range of conditions: both monocular and binocular viewing, with and without shading information, and for various amounts of change in viewpoint. When change of viewpoint could be anticipated (Experiment 2), the benefit from symmetry was reduced but not eliminated. This suggests that symmetry provides a cue to object orientation and also contributes in a more general manner.

Although shape discrimination was better overall for symmetric objects, performance remained highly viewpoint dependent, and we observed a similar stereo advantage for both symmetric and asymmetric objects.

References

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision, 35,* 33–44. doi:10.1023/A:1008154927611

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1162–1182. doi:10.1037/0096-1523.19.6.1162

Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences of the United States of America, 89,* 60–64. doi:10.1073/pnas.89.1.60

Caniard, F., & Fleming, R. W. (2007). Distortion in 3D shape estimation with changes in illumination. In C. Wallraven, V. Sundstedt (eds.), *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization,* APGV 2007, Tübingen, Germany (pp. 90–105). doi:10.1145/1272582.1272602

Chan, M. W., Stevenson, A. K., Li, Y., & Pizlo, Z. (2006). Binocular shape constancy from novel views: The role of a prior constraints. *Perception & Psychophysics, 68,* 1124–1139. doi:10.3758/BF03193715

Egan, E., Todd, J., & Phillips, F. (2012, May). *The role of symmetry in 3D shape discrimination across changes in viewpoint*. Poster session presented at the Annual Meeting of Vision Sciences Society, Naples, FL.

Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research, 31,* 1351–1360. doi:10.1016/0042-6989(91)90056-B

Kruschke, J. K. (2010). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences, 14,* 293–300. doi:10.1016/j.tics.2010.05.001

Lee, M. D. (2008). Bayes SDT: Software for Bayesian inference with signal detection theory. *Behavior Research Methods, 40,* 450–456. doi:10.3758/BRM.40.2.450

Lee, Y., & Saunders, J. A. (2011). Stereo improves 3D shape discrimination even when rich monocular shape cues are available. *Journal of Vision, 11,* 1–12.

Li, Y., Pizlo, Z., & Steinman, M. R. (2009). A computational model that recovers the 3D shape of an object from a single 2D retinal representation. *Vision Research, 49,* 979–991. doi:10.1016/j.visres.2008.05.013

Li, Y., Sawada, T., Shi, Y., Kwon, T., & Pizlo, Z. (2011). A Bayesian model of binocular perception of 3D mirror symmetrical polyhedra. *Journal of Vision, 11,* 1–20. doi:10.1167/11.4.11

Liu, Z., & Kersten, D. (2003). Three-dimensional symmetric shapes are discriminated more efficiently than asymmetric ones. *Journal of the Optical Society of America, A, Optics, Image Science & Vision, 20,* 1331–1340. doi:10.1364/JOSAA.20.001331

Nefs, H. T., Koenderink, J. J., & Kappers, A. M. L. (2005). The influence of illumination direction on the pictorial reliefs of Lambertian surfaces. *Perception, 34,* 275–287. doi:10.1068/p5179

Norman, J. F., Bartholomew, A. N., & Burton, C. L. (2008). Aging preserves the ability to perceive 3D object shape from static but not deforming boundary contours. *Acta Psychologica, 129,* 198–207. doi:10.1016/j.actpsy.2008.06.002

Norman, J. F., Swindle, J. M., Jennings, R., Mullins, E. M., & Beers, A. M. (2009). Stereoscopic shape discrimination is well preserved across changes in object size. *Acta Psychologica, 131,* 129–135. doi:10.1016/j.actpsy.2009.03.009

Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychological Science, 15,* 565–570. doi:10.1111/j.0956-7976.2004.00720.x

Norman, J. F., Todd, J. T., & Phillips, F. (1995). The visual perception of surface orientation from multiple sources of optical information. *Perception & Psychophysics, 57,* 629–636. doi:10.3758/BF03213268

Pizlo, Z., Li, Y., & Steinman, R. M. (2008). Binocular disparity only comes into play when everything else fails; a finding with broader implications than one might suppose. *Spatial Vision, 21,* 495–508. doi:10.1163/156856808786451453

Pizlo, Z., Sawada, T., Li, Y., Kropatsch, W. G., & Steinman, R. M. (2010). New approach to the perception of 3D shape based on veridicality, complexity, symmetry and volume. *Vision Research, 50,* 1–11. doi:10.1016/j.visres.2009.09.024

Pizlo, Z., & Stevenson, A. K. (1999). Shape constancy from novel view. *Perception & Psychophysics, 61,* 1299–1307. doi:10.3758/BF03206181

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs Sampling. . In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* (DSC 2003). Vienna, Austria: Technische Universität Wien.

Rock, I., & DiVita, J. (1987). A case of viewer-centered object perception. *Cognitive Psychology, 19,* 280–293. doi:10.1016/0010-0285(87)90013-2

Saunders, J. A., & Knill, D. C. (2001). Perception of 3D surface orientation from skew symmetry. *Vision Research, 41,* 3163–3183. doi:10.1016/S0042-6989(01)00187-0

Sawada, T. (2010). Visual detection of symmetry of 3D shapes. *Journal of Vision, 10,* 4. doi:10.1167/10.6.4

Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science, 32,* 1248–1284. doi:10.1080/03640210802414826

Shimshoni, I., Moses, Y., & Lindenbaum, M. (2000). Shape reconstruction of 3D bilaterally symmetric surfaces. *International Journal of Computer Vision, 39,* 97–110. doi:10.1023/A:1008118909580

Todd, J. T., & Norman, J. F. (2003). The visual perception of 3-D shape from multiple cues: Are observers capable of perceiving metric structure? *Perception & Psychophysics, 65,* 31–47. doi:10.3758/BF03194781

Vetter, T., Poggio, T., & Bülthoff, H. H. (1994). The importance of symmetry and virtual views in three dimensional object recognition. *Current Biology, 4,* 18–23. doi:10.1016/S0960-9822(00)00004-X

Wagemans, J. (1993). Skewed symmetry: A nonaccidental property used to perceive visual forms. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 364–380. doi:10.1037/0096-1523.19.2.364

Zhao, W., & Chellappa, R. (2001). Symmetric shape-from-shading using self-ratio image. *International Journal of Computer Vision, 45,* 55–75. doi:10.1023/A:1012369907247

Received June 12, 2012
Revision received September 4, 2012
Accepted September 6, 2012 ■