

Detecting mirror-symmetry of a volumetric shape from its single 2D image

Tadamasa Sawada

Department of Psychological Sciences, Purdue University, West Lafayette, IN, USA
tsawada@psych.purdue.edu
<http://psyclops.psych.purdue.edu/~tsawada/>

Zygmunt Pizlo

Department of Psychological Sciences, Purdue University, West Lafayette, IN, USA
pizlo@psych.purdue.edu
<http://psyclops.psych.purdue.edu/~pizlo/>

Abstract

We present a new computational model for verifying whether a 3D shape is mirror-symmetric based on its single 2D image. First, a psychophysical experiment which tested human performance in detection of 3D symmetry is described. These psychophysical results led to the formulation of a new algorithm for symmetry detection. The algorithm first recovers the 3D shape using a priori constraints (symmetry, planarity of contours and 3D compactness) and then evaluates the degree of symmetry of the 3D shape. Reliable discrimination by the algorithm between symmetric and asymmetric 3D shapes involves two measures: similarity of the two halves of a 3D shape and compactness of the 3D shape. Performance of this algorithm is highly correlated with that of the subjects. We conclude that this algorithm is a plausible model of the mechanisms used by the human visual system.

1. Introduction

Many, if not most objects in the real world are symmetric or at least approximately symmetric. More specifically, most symmetric objects are mirror, rather than rotational or translational symmetric. Hence, detecting mirror symmetry is important. In this paper, we will use “symmetry” to mean “mirror-symmetry”. Our everyday life experience suggests that human observers have little difficulty identifying which 3D objects are symmetric.

Figure 1 illustrates this fact. It is probably clear to the reader that the 3D shape in (a), but not in (b), (c), or (d) is mirror symmetric. Note that all 2D images in Figure 1 are asymmetric. It follows that the perceptual decision about the 3D symmetry of a shape cannot be based on evaluating symmetry of the 2D image. Instead, it requires recovering and evaluating the 3D shape itself. The recovery of a 3D shape from a single 2D image is an ill-posed inverse problem. Specifically, the family of possible 3D interpretations produced from a single 2D image is infinitely large. In order to recover a unique and correct 3D shape, the visual system has to impose *a priori* constraints on this family. In this paper, we examine how human observers detect 3D symmetry from a single 2D orthographic image. Based on results of this experiment we

formulated a new computational model of this ability. The model recovers a unique 3D shape by using the following constraints: symmetry of the 3D shape, planarity of its contours and 3D compactness. After the 3D shape is recovered, its asymmetry is evaluated by comparing the similarity of the two (symmetric) halves of the shape, as well as by measuring 3D compactness of the shape.

This paper is organized as follows. First, several observations about the relation between a 3D symmetric shape and its single 2D image are briefly described. Next, an overview of research on symmetry perception in human and computer vision is presented. Then, a psychophysical experiment on human 3D symmetry discrimination is described. Based on the psychophysical results, a new computational model for detecting symmetry is presented and the performance of the model is compared to the performance of human subjects. In the Appendix, a new computational model for recovering a 3D shape from a single 2D image is described. This model uses symmetry as one of the constraint, but it can recover asymmetric shapes, as well.

1.1. Properties of a 2D image of a 3D symmetric shape

Consider a 3D symmetric shape. The line segments, which connect pairs of symmetric points, are parallel to one another. These segments are called symmetry line segments. They all are perpendicular to the symmetry plane and their midpoints are on this plane.

Consider now an orthographic projection of a 3D symmetric shape on the 2D image. Parallelism of lines and midpoints of line segments are invariant under orthographic projection. It follows that images of symmetry line segments are parallel to one another. However, the midpoints of the projected symmetry lines segments are not collinear, unless these segments are coplanar in 3D space. Interestingly, if projected symmetry line segments are parallel to one another in the 2D orthographic image, then this image is consistent with a 3D mirror symmetric interpretation. This statement follows from the method of construction of a one-parameter family of 3D symmetric shapes (see Appendix). Thus, the images in Figures 1 c and d are consistent with 3D symmetric interpretations even though they are not perceived as such. This paper tries to

uncover the mechanisms that are used by the human visual system to discriminate between symmetric and asymmetric shapes.

3D mirror symmetry is a powerful constraint that can be used in recovering the 3D shape from its 2D image. The best way to understand the nature of this constraint is to consider the concept of a virtual image. A virtual image is computed from the given (real) 2D image. The virtual image is an image of the same 3D shape when viewed from a different viewing direction [17]. As a result, the problem of recovering a 3D symmetric shape from one 2D image is transformed to an easier (more constrained) problem of recovering a 3D shape from two 2D images. Two orthographic images, however, are still not sufficient for a unique recovery. This means that additional constraints have to be used (see below).

1.2. Symmetry in human vision

It is known that symmetry of a 3D shape plays an important role in human vision. Kontsevich [5] suggested that perceived 3D shape produced by a single 2D image is biased toward a 3D symmetric interpretation. He illustrated this suggestion by showing that a line drawing of a polyhedral object tends to be perceived as symmetric when the image is consistent with such an interpretation. He also showed that a nearly symmetric object tends to be perceived as symmetric even if the 2D image is not consistent with a perfectly symmetric 3D interpretation. Pizlo et al. [10] showed that symmetry prior is more important in perceptual interpretation of 3D shapes than binocular disparity. Pizlo and Stevenson [11], and Chan et al. [1] showed that human performance in shape constancy experiment is better for symmetric, than for asymmetric shapes.

1.3. Symmetry in computer vision

If two orthographic images of a 3D shape are given, a one-parameter family of 3D possible shapes can be recovered [3, 4]. Considering the fact that an orthographic image of a 3D symmetric shape allows one to produce a virtual image of the same 3D shape (see above), it follows that a single 2D orthographic image of a 3D symmetric shape also leads to a one-parameter family of possible 3D symmetric shapes.¹ If a 3D shape has two or more mirror symmetries, a unique 3D shape can be recovered [17]. If 3D data is available (e.g. from range sensors, motion or binocular disparity), the symmetry of a 3D shape can be evaluated and measured in a number of ways (e.g. [7, 8, 15, 19]). When only one 2D image is available (the case considered in our study), there are algorithms that can use

¹ If a perspective, rather than an orthographic image of a symmetric 3D shape is given, then a unique 3D interpretation can be computed (e.g. [3, 12, 9]).

topology of the 2D contours to determine pairs of symmetric points and features in the image, without actually recovering the 3D symmetric shape [18, 20].

It is trivially true that any 2D image of a 3D symmetric shape is consistent with infinitely many asymmetric 3D shapes. It is less trivial, but also true, that there are asymmetric 3D shapes such that each of their images is consistent with infinitely many 3D symmetric shapes. Images of such shapes are shown in Figures 1c and d. Interestingly, even though 3D symmetric interpretations are possible in these two cases, human observers do not perceive 3D symmetric shapes. Instead, they perceive asymmetric 3D shapes, which agrees with shapes that were used to produce these images. Formal psychophysical experiments testing this ability, as well as a computational model that can produce equally or even better performance, are described in this paper.

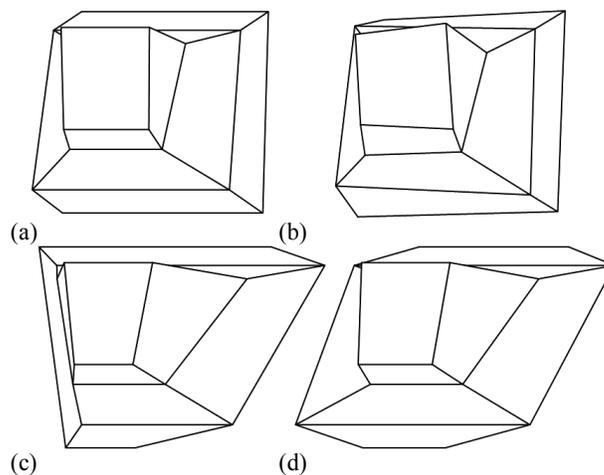


Figure 1. Orthographic images of (a) a symmetric polyhedron, (b) an asymmetric polyhedron generated by distorting the symmetric polyhedron shown in (a) randomly, (c-d) asymmetric polyhedra generated by distorting the symmetric polyhedron shown in (a) in such a way that the “symmetry line segments” remained parallel and faces remained planar.

2. Psychophysical experiment on 3D symmetry detection

2.1. Stimuli

Abstract symmetric and asymmetric polyhedra were used. Each symmetric shape consisted of a small and large box connected to each other. All faces of a symmetric polyhedron were planar (Figure 1a). Two types of asymmetric polyhedra were generated from symmetric ones. The first type (Type-A) was generated by distorting a symmetric polyhedron randomly (Figure 1b). The second type (Type-B) was generated by distorting a symmetric polyhedron in such a way that the symmetry line segments of the symmetric polyhedron remained parallel and the

faces remained planar (Figure 1c and d).

Orthographic projection was used to produce 2D images. Hidden edges were removed. Each polyhedron was randomly oriented in 3D space subject to two constraints: at least one vertex of each pair of symmetric vertices, and at least six symmetric pairs of vertices had to be visible. These constraints allow the recovery of the entire polyhedron (both the visible and the hidden parts) by our model (see Section 3.1 and Appendix). The stimuli were shown on an LCD monitor. The subject viewed the monitor with the right eye.

2.2. Procedure

The method of signal detection was used. Each session consisted of 200 trials: 100 trials with symmetric polyhedra and 100 trials with asymmetric polyhedra, presented in a random order. There were 8 experimental conditions: two types of asymmetric polyhedra (Type-A vs. -B) \times four levels of distortion for generating asymmetric polyhedra (L1 – L4). The levels corresponded to the extent by which the vertices have been moved. L1 corresponded to quite small distortion, whereas L4 corresponded to quite large distortion. The subject ran two sessions for each condition. The order of sessions was randomized.

In each trial, the stimulus was shown for 500 ms. The subject responded whether the presented polyhedron was symmetric or not. After each trial, a feedback about the accuracy of the response was given. The performance of the subject was evaluated by the discriminability measure d' used in the signal detection theory. d' is computed from hit rate (symmetric polyhedra judged as symmetric) and correct rejection rate (asymmetric polyhedra judged as asymmetric) in each session [16]. Higher performance corresponds to higher values of d' ($d'=0$ - chance performance, $d'=\infty$ - perfect performance). In this study, perfect performance corresponds to d' about 5.

2.3. Results and Discussion

Results of one subject (the first author) are shown in Figure 2 (results of the other subject were very similar). The ordinate shows d' . The abscissa shows levels of distortion of the asymmetric polyhedra. The two curves indicate types of asymmetric polyhedra: black symbols Type-A, and gray symbols Type-B.

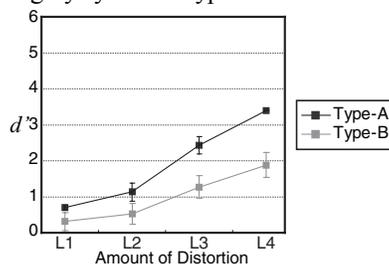


Figure 2. Results of the subject in the psychophysical experiment. The ordinate shows d' , and the abscissa shows levels of distortion of asymmetric polyhedra. Error bars represent standard errors.

As expected, performance improved with the level of distortion for both two types of asymmetric polyhedra. Also as expected, performance was better with Type-A asymmetric polyhedra. Recall that the Type-A asymmetric polyhedra were generated by distorting a symmetric polyhedron randomly. Random distortions removed the parallelism of symmetry line segments and planarity of contours. Parallelism of symmetry line segments and planarity of contours were preserved with the Type-B asymmetric polyhedra. The fact that performance with Type-A distortion was better is not surprising. What is surprising, however, is that the subject was able to perform above chance level with Type-B distortion. Recall that each image of an asymmetric polyhedron of this type is consistent with a symmetric interpretation. How then can the asymmetry of the 3D shape be detected? If all images in a session were consistent with symmetric polyhedra, how was the subject able to discriminate between images of symmetric vs. asymmetric 3D shapes? Figures 1c and d provides an illustration that this is possible. In fact, it is difficult to “see” that the images in Figures 1 c and d could be produced by 3D symmetric shapes. Clearly, the visual system uses a priori constraints that override, or compete with symmetry constraint in the process of recovery of 3D shapes. The nature of these constraints will be described next.

3. Computational model of detecting 3D symmetry

The model consists of two main stages. In the first, a polyhedron is recovered from a 2D image. In the second, the asymmetry of the recovered polyhedron is measured and compared to a criterion in order to decide whether or not the recovered 3D shape is symmetric. Performance of the model was evaluated using the same 3D simulated shapes and their images that were used in the psychophysical experiment. Note that our model does not perform image segmentation. Specifically, the model is provided with visible contours of a 3D shape, as well as with the information which vertices in the 3D shape are symmetric and which contours are planar. More precisely, the model was given two possibilities for the orientation of the symmetry plane. These two possibilities correspond to the vertical and horizontal planes of symmetry in the case of the objects shown in Figure 1. The model evaluated the symmetry of both interpretations and chose the more symmetric one.

Note that the 2D orthographic image is the only input data to the model (and to the human visual system). However, the 2D image is not the only information used by the model (and by the human visual system). The model

(and the human visual system) also uses a priori shape constraints. The constraints are used because 3D shape recovery from a single 2D image is underconstrained. So, even though the 2D image is the only input data for the discrimination between symmetric and asymmetric 3D shapes, the 2D image itself doesn't have enough information to perform this discrimination. Reliable discrimination can only be performed after the 3D shape is recovered, through the application of a priori constraints. In other words, the a priori constraints add information not only for the purpose of recovering the 3D shape, but also for the purpose of discrimination between two categories of 3D shapes (symmetric vs. asymmetric).

The computational details are explained below. Specifically, the second stage, in which the asymmetry of the recovered 3D shape is measured, is described in the next section. This second stage is essential for understanding how the model's performance was compared to that of the subject. The first stage, in which an approximately symmetric 3D shape is recovered from a single 2D image is described in the Appendix.

3.1. Measure of asymmetry of the recovered polyhedron

Before the asymmetry of a 3D shape is evaluated, the shape is recovered. As pointed out above, recovery of a unique 3D shape from a single 2D image is underconstrained. In order to produce a unique 3D shape, one has to restrict the family of possible 3D interpretations, by using a priori constraints. Here we use the algorithm described by Li & Pizlo [6]. Specifically, given a 2D orthographic image of a symmetric 3D shape, the algorithm begins by producing a virtual image of this shape (see Appendix). Next, the algorithm constructs a one-parameter family of 3D symmetric shapes consistent with the given 2D image. Finally, 3D shape with maximal 3D compactness is selected as the recovered shape. 3D compactness is defined as V^2/S^3 , where V is the volume and S is the surface area of the 3D shape. In the case of opaque shapes, planarity of faces themselves [14] or planarity in conjunction with symmetry [9] can be used to recover the back part of the polyhedron. When the 3D shape is asymmetric, this algorithm must be modified. Specifically, the 2D image is first corrected so that it is consistent with a 3D symmetric shape. Then, the 3D shape is recovered as described above. Finally, the 3D symmetric shape is distorted in 3D, so that it becomes consistent with the given 2D image (see Appendix).

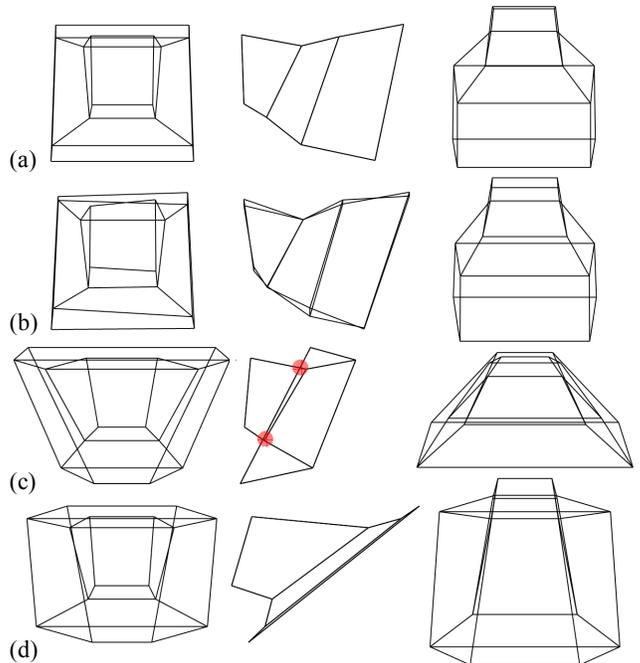


Figure 3. Recovered polyhedra from images in Figure 1. To give the reader a better intuition about the recovered 3D shapes, the polyhedra in this Figure are transparent, and three views (front, side and top) are shown. Red marks in (c) indicate the line segment which represents self-intersection of the object's surface.

Now, that the 3D shape is recovered, one can evaluate how asymmetric it is. This is done by comparing the two halves of the 3D shape. If the shape is perfectly symmetric, the two halves are identical. If the 3D shape is only approximately symmetric, the two halves will only be approximately identical. We use the sum of squared differences between the corresponding 2D angles, α_a and $\alpha_{counterpart(a)}$, of the polyhedron H , as a metric:

$$ap(H) = \sum_a (\alpha_a - \alpha_{counterpart(a)})^2, \quad (1)$$

The greater $ap(H)$ is, the more asymmetric H is. In the simulation experiment, a criterion was used, to which $ap(H)$ was compared, in order to decide between symmetric and asymmetric 3D shapes.

The measure $ap(H)$ defined in (1) can correctly detect asymmetry in the case of Type-A asymmetric 3D polyhedra that were used in our experiment (see above). However, it will not detect asymmetry in the case of Type-B asymmetric polyhedra. The reason is that every image of such a polyhedron is actually consistent with a 3D symmetric interpretation. As a result, our algorithm will recover a symmetric polyhedron. This means that we need another measure of 3D symmetry.

An examination of the 3D symmetric shapes recovered from images produced by Type-B asymmetric polyhedra shows that the recovered shapes have surfaces with self-intersection. An example of such a case is shown in

Figures 1c and 3c. Self-intersecting surfaces are not “valid” in the sense that some vertices and edges that should not be visible in the 2D image are actually visible. They should not be visible because the surfaces are assumed to be opaque. In other words, the recovered symmetric shapes with self-intersections are not consistent with the given images. It follows that 3D interpretations that are consistent with the given 2D image are not symmetric. There are also other cases of Type-B asymmetric polyhedra, in which self-intersection of surfaces does not occur, but the 3D shape is correctly perceived by observers as asymmetric. Such a case is shown in Figures 1d and 3d. An examination of the recovered 3D shape shows that it tends to be “thin” and have very low 3D compactness. This observation suggests that the human visual system “prefers” compact asymmetric 3D shapes over symmetric non-compact ones. It follows that 3D compactness is a more important prior than 3D symmetry. Interestingly, compactness can also be used to detect 3D recovered shapes that have self-intersecting surfaces because such shapes tend to have low compactness.² Indeed, using compactness alone allows one to detect most self-intersecting 3D shapes. To test this, we generated 800 Type-B asymmetric polyhedra and recovered 3D symmetric shapes from their 2D images. 106 of the recovered 3D shapes had self-intersecting surfaces and most of them (101) had very low compactness.

In the simulations (described below) we tested two models of symmetry discrimination. One used 3D compactness only (model C), and the second used both $ap(H)$ as defined in (Eq 1) and 3D compactness (model A+C). Model A+C judges a polyhedron as symmetric only when the polyhedron is symmetric according to both $ap(H)$ and 3D compactness criteria.

3.2. Simulation experiment

The two models of symmetry discrimination (C and A+C) were applied to the images that were used in the psychophysical experiment. Each model was applied to

² The self intersection makes an upper bound of compactness smaller, as compared to objects without self intersection. Take an object O with a self intersection of its surface. Let the self intersection separate the surface into two surfaces with surface areas S_a and S_b . Compactness of this object is maximized when these two surfaces form two spheres that contact each other at the self intersection:

$$\begin{aligned} \max(C_o) &= \frac{(4\pi r_a^3/3 + 4\pi r_b^3/3)^2}{(S_a + S_b)^3} = \left(\frac{4\pi}{3}\right)^2 \cdot \frac{\left(\sqrt{S_a/4\pi^3} + \sqrt{S_b/4\pi^3}\right)^2}{(S_a + S_b)^3} \\ &= \frac{1}{36\pi} \left(\frac{S_a^{1.5} + S_b^{1.5}}{(S_a + S_b)^{1.5}}\right)^2 < \frac{1}{36\pi} \end{aligned}$$

where C_o is compactness of the object, and r_a and r_b are radii of spheres with surface areas S_a and S_b . Recall that $1/(36\pi)$ is the upper bound of 3D compactness (the upper bound corresponds to a single sphere). So, self intersection of the surface of the object makes an upper bound of its compactness smaller.

images of symmetric and asymmetric shapes for each of the eight experimental conditions, and computed asymmetry measure for all shapes. This led to two overlapping frequency distributions: one representing symmetric and the other representing asymmetric shapes (as in a signal detection model). When these two distributions overlap, perfect discrimination is impossible. By changing the criterion for classifying the shape as asymmetric, the proportion of hits and false alarms are changed. Furthermore, the discriminability d' is also changed to some degree. In the simulation results shown in Figure 4, the criteria for compactness and for $ap(H)$ were chosen in order to provide the best fit of the model to the subject's data.

Results of the model superimposed on the results of the subject are shown in Figure 4. The ordinate shows d' . The abscissa shows levels of distortion for the asymmetric polyhedron.

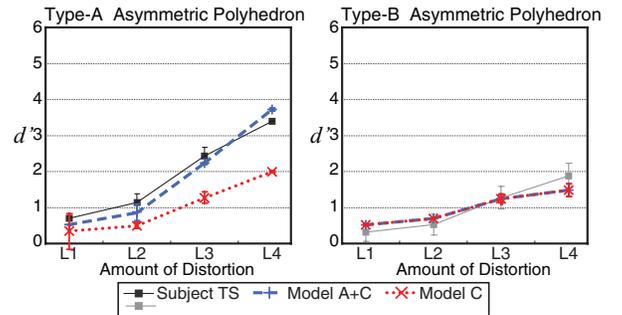


Figure 4. Results of the two models in the simulation experiment. The ordinate shows d' , and the abscissa shows the level of distortion for generating asymmetric polyhedra. Error bars represent standard errors.

Consider the discrimination between symmetric and Type-B asymmetric polyhedra (Figure 4, right panel). Recall that Type-B asymmetric polyhedra were produced by distorting symmetric polyhedra in such a way that the symmetry line segments remained parallel and the faces remained planar (see Figures 1c and d). As a result, there is always a 3D symmetric interpretation for any image of this type of asymmetric polyhedron. Hence, the asymmetry of Type-B asymmetric polyhedra can not be evaluated by $ap(H)$, but only by 3D compactness. It can be seen that the performance of Model C is quite similar to that of the subject. Obviously, Model A+C would have produced the same performance as Model C because the criterion $ap(H)$ does not contribute anything to the discrimination in this case.

Next, consider the discrimination between symmetric and Type-A asymmetric polyhedra (Figure 4, left panel). Recall that Type-A asymmetric polyhedra were produced by randomly distorting symmetric polyhedra (see Figure 1b). It can be seen that model A+C produced a good fit to the subject's results. Model C's performance is

systematically lower. This was to be expected. In this condition, many asymmetric polyhedra had high 3D compactness. Therefore, 3D compactness itself could not lead to reliable discrimination between symmetric and asymmetric polyhedra.

From these results, we can conclude that model A+C is a plausible model of symmetry discrimination. Its performance, as measured by d' is similar to that of the subject across all experimental conditions. Specifically, a human observer perceives a 3D shape as symmetric when a symmetric interpretation is possible and the recovered symmetric shape is compact with non self-intersecting surfaces.

4. Summary and future work

This paper examined the problem of discrimination between 3D symmetric and asymmetric shapes based on a single 2D orthographic image of a randomly generated polyhedron. We reported a psychophysical experiment and formulated a computational model of this ability. To the best of our knowledge, this is the first such model. Performance of the model was similar to that of the subject suggesting that this model provides a possible explanation of the underlying perceptual mechanisms.

The model performs a 3D recovery of a symmetric or asymmetric 3D polyhedron. It is an elaboration of previous models that could recover only symmetric shapes. After the 3D shape is recovered, its symmetry is evaluated by comparing corresponding angles and evaluating 3D compactness. The fact that 3D compactness is a critical component of 3D shape recovery and 3D symmetry discrimination implies that this constraint is at least as important as 3D symmetry in perception of 3D shapes.

Our future work will address several aspects of the model. First, although our current model can detect whether the 3D shape is asymmetric, it will not always recover this asymmetric shape. This is the case with Type-B asymmetric polyhedra. The next step will be formulating a model that can explain what the subject actually perceives when looking at images like those in Figures 1c and d. Next, the model will be elaborated by adding an image segmentation stage. In the current version, the object's contours are given to the model, as is the possible position of the symmetry plane. In our future work, the model will be generalized so that it can detect vertices and edges, group contours and detect pairs of symmetric features. Finally, this model will be generalized to the case of smoothly curved objects, such as animal bodies.

5. Appendix: A new algorithm for recovering a 3D approximately symmetric polyhedron from a single 2D orthographic image

Let $z=0$ be the image plane, the x-axis point to the right and y-axis point up. Consider an orthographic image of a 3D polyhedron in the image plane xy . The polyhedron is assumed to be opaque. We begin with the analysis of those pairs of symmetric vertices that are both visible. It is assumed that the correspondence between 3D symmetric points is given. The algorithm consists of seven steps (operations). These steps will be described below.

5.1. Canonical orientation on the image plane.

First, the 2D image of the polyhedron is rotated in the image plane (in clockwise direction) so that the projected symmetry line segments become horizontal. There are two rotation angles that differ by 180 deg, which can accomplish this. The smaller of these two angles is used. This rotation is unique when all projected symmetry line segments are parallel. When these segments are not parallel (as is the case with an image of an asymmetric polyhedron), the rotation makes the average orientation of the projected symmetry line segments horizontal:

$$\mathbf{R}_{2D}(-\phi) \cdot \mathbf{p}_i = \mathbf{P}_i, \quad (\text{A1})$$

(A1) can be written explicitly as follows:

$$\begin{bmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{bmatrix} \cdot \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} X_i \\ Y_i \end{bmatrix}$$

where $\mathbf{p}_i = [x_i, y_i]^t$ and $\mathbf{P}_i = [X_i, Y_i]^t$ are positions of a projected vertex i before and after the rotation, and ϕ is the average orientation of the projected symmetry line segments. Let a symmetric counterpart of the vertex i be a vertex j ($\mathbf{p}_j = [x_j, y_j]^t$ and $\mathbf{P}_j = [X_j, Y_j]^t$).

5.2. Correction of the 2D image

When the projected symmetry lines are all parallel in the 2D image, this step is skipped. When they are not parallel, their orientations are changed (corrected) so that they become parallel. This way, the corrected image will be consistent with a 3D symmetric interpretation. Specifically, each projected symmetry line segment is made parallel to the x-axis by applying the following transformation:

$$\begin{aligned} \mathbf{P}'_i &= \begin{bmatrix} X'_i \\ Y'_i \end{bmatrix} = \begin{bmatrix} X_i \\ (Y_i + Y_j)/2 \end{bmatrix} \\ \mathbf{P}'_j &= \begin{bmatrix} X'_j \\ Y'_j \end{bmatrix} = \begin{bmatrix} X_j \\ (Y_i + Y_j)/2 \end{bmatrix} = \begin{bmatrix} X_j \\ Y'_i \end{bmatrix} \end{aligned} \quad (\text{A2})$$

where $\mathbf{P}'_i = [X'_i, Y'_i]^t$ and $\mathbf{P}'_j = [X'_j, Y'_j]^t$ are positions of projected vertices i and j after the correction.

Note that this transformation leads to the smallest change

of the two endpoints, in the least sum of squares sense. This corrected image is an orthographic image of a perfectly symmetric shape.

5.3. Producing a virtual image

The method proposed by Vetter and Poggio [17] is applied to the corrected image. The virtual image of the symmetric 3D shape is generated by reflecting the corrected image with respect to the y-axis:

$$\begin{aligned} \mathbf{Q}_i' &= \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \cdot \mathbf{P}_i' = \begin{bmatrix} -X_i \\ Y_i' \end{bmatrix} \\ \mathbf{Q}_j' &= \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \cdot \mathbf{P}_j' = \begin{bmatrix} -X_j \\ Y_j' \end{bmatrix} \end{aligned} \quad (\text{A3})$$

where \mathbf{Q}_i' and \mathbf{Q}_j' are positions of projected vertices i and j in the virtual image. This virtual image is an image of the same 3D shape after a 3D rigid rotation of the shape around the y-axis. Let the 3D coordinates of the symmetric pair of vertices i and j of the real (corrected) image be $\mathbf{V}_i' = [X_i, Y_i', Z_i']^t$ and $\mathbf{V}_j' = [X_j, Y_j', Z_j']^t$. Note that x- and y-values of \mathbf{V}_i' and \mathbf{V}_j' are identical to those of \mathbf{P}_i' and \mathbf{P}_j' on an orthographic image. In the same way, let the 3D coordinates of the symmetric pair of vertices i and j of the virtual image be $\mathbf{U}_i' = [-X_i, Y_i', Z_i']^t$ and $\mathbf{U}_j' = [-X_j, Y_j', Z_j']^t$. Then, the vertex that corresponds to \mathbf{V}_i' after the 3D rigid rotation can be written as follows:

$$\mathbf{\Lambda}_i' = \mathbf{R}_{3D} \cdot \mathbf{V}_i' \text{ and } \mathbf{\Lambda}_i' = \mathbf{U}_j' \quad (\text{A4})$$

(A4) can be written explicitly as follows:

$$\begin{bmatrix} -X_j \\ Y_j' \\ Z_j' \end{bmatrix} = \mathbf{R}_{3D} \cdot \begin{bmatrix} X_i \\ Y_i' \\ Z_i' \end{bmatrix},$$

where \mathbf{R}_{3D} is a 3×3 rotation matrix, and $\mathbf{\Lambda}_i'$ is the 3D vertex i after the 3D rigid rotation.

The 3D rigid rotation has three parameters. Recall, however, \mathbf{R}_{3D} in Equation (A4) has only one parameter, the angle θ of rotation around the y-axis:

$$\mathbf{R}_{3D} = \mathbf{R}_y(\theta) = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}. \quad (\text{A5})$$

5.4. Recovering one-parameter family of symmetric polyhedra

From the first row of Equation (A4) we obtain:

$$-X_j = \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix}^t \cdot \begin{bmatrix} X_i \\ Z_i' \end{bmatrix}. \quad (\text{A6})$$

An equation for Z_i can be derived by combining Equation (A6) with (A1):

$$Z_i = \frac{\cos \phi (x_j + \cos \theta x_i) + \sin \phi (y_j + \cos \theta y_i)}{-\sin \theta}. \quad (\text{A7})$$

Hence, the vertex i of the recovered 3D symmetric shape can be written as follows:

$$\mathbf{V}_i' = \begin{bmatrix} X_i \\ Y_i' \\ \frac{\cos \phi (x_j + \cos \theta x_i) + \sin \phi (y_j + \cos \theta y_i)}{-\sin \theta} \end{bmatrix}. \quad (\text{A8})$$

It can be seen that \mathbf{V}_i' depends on one parameter, the angle θ .

5.5. Undoing the 2D correction in 3D space

When the projected symmetry lines are all parallel in the real 2D image, this step is skipped. When they are not parallel, the recovered 3D shape is distorted so that its image agrees with the given 2D real image:

$$\mathbf{V}_i'' = \mathbf{V}_i' + \mathbf{\Delta}_{3D}, \quad (\text{A9})$$

where $\mathbf{\Delta}_{3D}$ is a 3D distortion and \mathbf{V}_i' is position of vertex i after the distortion. Let the 3D coordinate of $\mathbf{\Delta}_{3D}$ be $[\Delta_x, \Delta_y, \Delta_z]^t$. From Equation (A2), $\mathbf{\Delta}_{3D} = [0, Y_i - Y_i', \Delta_z]^t$ and Δ_z can be arbitrary. Obviously, this distortion ($\mathbf{\Delta}_{3D}$) is minimized when $\Delta_z = 0$. Hence, the minimally distorted symmetric shape which is consistent with the real 2D image can be written as follows:

$$\begin{aligned} \mathbf{V}_i'' &= \mathbf{V}_i' + \min(\mathbf{\Delta}_{3D}) \\ &= \begin{bmatrix} X_i \\ Y_i \\ \frac{\cos \phi (x_j + \cos \theta x_i) + \sin \phi (y_j + \cos \theta y_i)}{-\sin \theta} \end{bmatrix}. \end{aligned} \quad (\text{A10})$$

Note that the transformation of x and y coordinates in Equation (A10) is an inverse transformation of that in Equation (A2).

5.6. Applying planarity constraint to recover hidden vertices

If all vertices are visible, this step is skipped. Symmetric pairs whose one vertex is visible and the other is hidden are recovered by applying two constraints: planarity constraint for the visible vertex and symmetry constraint for the hidden counterpart [9]. In order to use a planarity constraint, at least three vertices of a face on which the visible vertex is located have to be recovered first. Assume that the face is planar, and the orientation of the face is known. The z-value of the visible vertex is obtained by computing an intersection of the face and the projection line emanating from the image of this vertex. The hidden counterpart is recovered by reflecting the visible vertex with respect to the symmetry plane of shape.

5.7. Applying the maximum compactness constraint

The maximally compact 3D shape is chosen from the one-parameter family, as the recovered shape.

Interestingly, this algorithm that was formulated for the case of 3D symmetric shapes can be applied to 2D shapes, as well. Specifically, the steps 1-5 of the algorithm work the same way in the case of 3D and 2D points simply because they are applied to one pair of symmetric vertices at a time. As a result, the algorithm produces a one-parameter family of symmetric shapes and the shapes are either 2D or 3D depending on whether all midpoints of the symmetry line segments are collinear or not. This seems to be the first such algorithm. Prior algorithms for recovering 3D symmetric shapes needed either four [17] or three [6] non-coplanar symmetric pairs of vertices. It follows that these prior algorithms cannot recover planar symmetric figures. After a one-parameter family of 2D symmetric shapes is computed, the human visual system chooses a symmetric shape that can be produced from the given image by shearing the 2D shape along the projected symmetry line segments [13].

Acknowledgment

This project was supported by the National Science Foundation (grant #0533968) and the US Department of Energy.

References

- [1] M. W. Chan, A. K. Stevenson, Y. Li, and Z. Pizlo. Binocular shape constancy from novel views: the role of a priori constraints. *Perception & Psychophysics*, 68(7):1124-1139, 2006.
- [2] A. R. J. François, G. G. Medioni, and R. Waupotitsch. Mirror symmetry \Rightarrow 2-view stereo geometry. *Image and Vision Computing*, 21(2):137-143, 2003.
- [3] T. S. Huang, and C. H. Lee. Motion and structure from orthographic projections. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(5):536-540, 1989.
- [4] J. J. Koenderink, and A. J. van Doorn. Affine structure from motion. *Journal of the Optical Society of America A*, 8(2):377-385, 1991.
- [5] L. L. Kontsevich. Symmetry as a depth cue. Human Symmetry Perception. In C. W. Tyler, editor, *Human Symmetry Perception and Its Computational Analysis*, 331-347, Lawrence Erlbaum Associates, 2002.
- [6] Y. Li, and Z. Pizlo. Recognition of shapes of 3D symmetric objects by using planarity and compactness constraints. *Proc. SPIE-IS&T Electronic Imaging*, 6499:64990B/1-10, 2007.
- [7] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3D geometry. *Proc. ACM SIGGRAPH*, 25(3):560-568, 2006.
- [8] D. O'Mara, and R. Owens. Measuring bilateral symmetry in digital images. *IEEE TENCON - Digital Signal Processing Applications*, 1:151-156, 1996.
- [9] H. Mitsumoto, and S. Tamura. 3-D reconstruction using mirror images based on a plane symmetry recovering method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):941-946, 1992.
- [10] Z. Pizlo, Y. Li, and G. Francis. A new look at binocular stereopsis. *Vision Research*, 45(17):2244-2255, 2005.
- [11] Z. Pizlo, and A. K. Stevenson. Shape constancy from novel views. *Perception & Psychophysics*, 61(7):1299-1307, 1999.
- [12] C. A. Rothwell. *Object Recognition through Invariant Indexing*. Oxford University Press, 1995.
- [13] T. Sawada, and Z. Pizlo. Detection of skewed symmetry. *Journal of Vision*, 2008. (In Press).
- [14] K. Sugihara. *Machine Interpretation of Line Drawings*. MIT Press, 1986.
- [15] C. Sun, and J. Sherrah. 3D symmetry detection using the extended gaussian image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):164-169, 1997.
- [16] W. P. Tanner, Jr. and J. A. Swets. A decision-making theory of visual detection. *Psychological Review*, 61(6):401-409, 1954.
- [17] T. Vetter, and T. Poggio. Symmetric 3D objects are an easy case for 2D object recognition. *Spatial Vision*, 8(4):443-453, 1994.
- [18] A. P. Vicent, P. C. Calleja, and R. R. Martin. Skewed mirror symmetry in the 3D reconstruction of polyhedral models. *Journal of Winter School on Computer Graphics*, 11(3):504-511, 2003.
- [19] H. Zabrodsky, and D. Weishall. Using bilateral symmetry to improve 3D reconstruction from image sequences. *Computer Vision and Image Understanding*, 67(1):48-57, 1997.
- [20] H. L. Zou, and Y. T. Lee. Skewed mirror symmetry detection from a 2D sketch of a 3D model. *Proc. the 3rd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, 69-76, 2005.