

Shape constancy from novel views

ZYGMUNT PIZLO and ADAM K. STEVENSON
Purdue University, West Lafayette, Indiana

Prior experiments on shape constancy from novel views are inconclusive: Some show that shapes of objects can be recognized reliably from novel views, whereas others show just the opposite. Our analysis of prior results suggests that shape constancy from novel views is reliable when the object has properties that constrain its shape: The object has volumetric primitives, it has surfaces, it is symmetrical, it is composed of geons, its contours are planar, and its images provide useful topological information about its three-dimensional structure. To test the role of some of these constraints, we performed a set of experiments. Solid shapes (polyhedra) were shown on a computer monitor by means of kinetic depth effect. Experiment 1 showed that shape constancy can be reliably achieved when a polyhedron is represented by its contours (most of the constraints are present), but not when it is represented by vertices or by a polygonal line connecting the vertices in a random order (all the constraints are absent). Experiments 2 and 3 tested the role of individual constraints. Results of these experiments show that shape constancy from novel views is reliable when the object has planar contours and when the shapes of the contours together with topological information about the relations among the contours constrain the possible interpretations of the shape. Symmetry of the object and the topological stability of its image also contribute to shape constancy.

Shape constancy refers to the fact that the percept of the shape of a given object remains constant despite changes in the shape of the retinal image. The retinal image may change because of changes in the orientation of the object relative to the observer (see Pizlo, 1994, for a review of prior research on shape constancy). Consider an example of shape constancy. When you walk around your car, your retinal image of the car changes. Yet, your percept of the shape of the car is constant. One can argue that achieving shape constancy in the case of your car may represent a rather trivial accomplishment: After all, you have seen your car from almost all viewing directions, and, therefore, the constancy of the percept of the car's shape may simply involve matching of the current view with all previously seen views. To prevent the observer from matching familiar views (or images), one should use unfamiliar objects and test shape constancy from unfamiliar (novel) views. There have been a number of studies performed during the last decade that tested whether shape constancy can be reliably achieved from novel views. We will first briefly review the most representative results.

Rock and DiVita (1987) used three-dimensional (3-D) wire objects as stimuli. Viewing was binocular with a

close viewing distance. The subjects were shown two objects from two viewing directions and were asked to decide whether these two objects had the same shape. When the viewing directions were identical, same shapes produced identical retinal images. When the viewing directions were different, the retinal images were different regardless of whether the shapes were or were not the same. Rock and DiVita found that, in the case of identical viewing directions, the subjects could reliably discriminate between same and different shapes. This result is not surprising because the subject's response could be simply based on the shape of the retinal image. When the viewing directions were different (45° apart), discrimination performance was at chance level. Results of this experiment suggest that there is no shape constancy from novel views.

Another set of experiments was performed by Farah, Rochlin, and Klein (1994). In one condition, they used objects that were obtained by bending oval disks of modeling clay. These surface objects resembled "curled potato chips." In the second condition, they used 3-D wire (contour) objects that were obtained by bending modeling strings to the shape of the surface boundary of the objects from the first condition. In the first two experiments, the viewing was binocular from close viewing distance. In the last two experiments, the stimuli were presented on a monoscopic computer screen. The angular difference between identical and novel views was in the range of 30°–60°. Farah et al. found that, in the case of wire objects, the accuracy of shape discrimination was quite poor for novel views and was much lower than that for familiar views (in one of the experiments, the performance for novel views was at chance level). In the case of surface objects, the accuracy of shape discrimination for novel and

This study was presented at the annual meetings of the Association for Research in Vision and Ophthalmology in 1997 and 1998. The authors thank the anonymous reviewers for their comments and suggestions, Moses Chan, Jan-Yuen Chen, Corrinne Lim, and Grant Wei who served as subjects, Filip J. Pizlo for his assistance with computer graphics, and Robert Melendez for his help in setting up Experiments 2 and 3. Correspondence should be addressed to Z. Pizlo, Department of Psychological Sciences, Purdue University, West Lafayette, IN 47907-1364 (e-mail: pizlo@psych.purdue.edu).

—Accepted by previous editor, Myron L. Braunstein

familiar views was about the same and well above chance level in the first two experiments, and the accuracy for novel views was moderately worse than that for familiar views in the last two experiments.

Next, consider the study of Bülthoff and Edelman (1992). They tested shape constancy using 3-D wire objects constructed from wire-like cylinders (tubes) and amoeba-like objects. The objects were displayed on a monitor by means of the kinetic depth effect (KDE). The subjects were familiarized with two sequences of views of the object being 75° apart and then were tested on a single view that differed from one of the familiar views by 15°–90°. Bülthoff and Edelman found that, regardless of the type of object, the accuracy of shape discrimination was quite poor for novel views and was much lower than that for familiar views.

In yet another study, Biederman and Gerhardstein (1993) used line drawings of 3-D objects composed of simple volumetric parts called *geons*. These line drawings were shown on a monoscopic computer monitor and did not contain any explicit surface features, such as shading, color, texture, or specular reflections. The subjects in their experiment could reliably achieve shape constancy when presented with two different views (45° apart) of an unfamiliar object.

Results of these four studies appear quite inconsistent or even contradictory: Some of them show that shape constancy can be reliably achieved from novel views, whereas others show exactly the opposite. Since all these studies used similar procedures (i.e., comparison of shapes of two objects from different viewing directions), it is clear that the differences in performance are related to differences among the stimuli. Consider Biederman and Gerhardstein's (1993) experiment. Their experiment is of particular interest because they demonstrated shape constancy from novel views of solid objects under minimal conditions (i.e., in the absence of explicit surfaces and with monoscopic viewing of a stationary object). Their objects were composed of geons. Geons are simple solid objects whose shapes involve several constraints: volumetric primitives (such as 3-D junctions), symmetries (approximate), and planar or piecewise planar surface contours. We conjecture that it was the presence of these constraints that made the shape constancy task easy in Biederman and Gerhardstein's experiment. It is not clear, however, why performance was much poorer in Bülthoff and Edelman's (1992) study, since their stimuli seem to have contained some or all of those constraints. We hypothesize that there were still other constraints that existed in Biederman and Gerhardstein's stimuli, which were responsible for this difference in performance. For example, geons give rise to topologically stable retinal images. That is, a small change of viewing direction is not likely to produce large topological changes in the image. Furthermore, shapes of planar contours that are present in objects composed of geons are likely to constrain each other, and, in addition, they constrain the overall shape of the object. In other words, it is usually not possible to

change the shape of one contour of a geon without changing shapes of other contours that share edges with it. Such constraining among shapes of contours of an object seems to be weaker in the case of wire or amoeba-like objects. For example, in the case of Bülthoff and Edelman's wire object, one part of the object can be rotated around a vertex of an L junction, keeping the angle at this vertex constant. Such a rotation will produce an infinite number of differently shaped objects, even though all these objects have identically shaped planar parts of the contours (i.e., L parts).

If our conjecture about the importance of features that constrain the shape of an object is true, then it should be possible to demonstrate the harmful effect of elimination of individual constraints on performance in a shape constancy task. Below, we describe three experiments that tested such effects.

EXPERIMENT 1 Shape Constancy for Polyhedra, Vertices, and Polygonal Lines

Experiment 1 was designed to test the joint effect on shape constancy of several constraints described in the previous section. It was expected that if the objects contain a number of constraints (such as the volumetric 3-D objects of Biederman & Gerhardstein, 1993), the subjects would be able to achieve shape constancy reliably. If the objects lack all of these constraints (as in the case of Rock and DiVita's, 1987, 3-D wire objects), the subjects should perform at chance level.

Method

Subjects. Four subjects (including the 2 authors) were tested. The authors (Z.P. and A.S.) received much more practice before running the experiment than did the other 2 subjects. Subject G.W. was naive with respect to the hypotheses being tested.

Stimuli. All stimuli¹ involved 16 points that were generated as vertices of a polyhedron. The positions of the vertices were generated randomly in 3-D, subject to the following constraints: The faces of the object were planar, the object had one plane of symmetry, the "front" part of the object was a box that was smaller than the "back" part of the object, and, finally, the two boxes had a pair of faces that were coplanar.

Three types of stimuli were used in which the object was represented by (1) edges of the *polyhedron* (Figure 1 shows an example of this type of a stimulus. In this figure, the hidden edges have been removed to facilitate visualization of the object; however, in the experiment, all edges were visible), (2) *vertices* of the polyhedron without any contours, or (3) a *polygonal line* connecting the vertices in a random order. Note that the stimuli in all three conditions had the same overall 3-D structure involving 16 vertices of a polyhedron and differing only in the manner in which the vertices were connected (if at all).

The three types of objects contained different constraints (if any). In the polyhedron condition, the contours were planar, volumetric primitives were present (trihedral angles), the object had a plane of symmetry, the shapes of individual faces constrained each other because of the presence of useful topological information about which contours share edges and/or vertices, and the images of the polyhedra were topologically stable (i.e., small changes in the 3-D orientation of the object did not lead to large topological changes in the

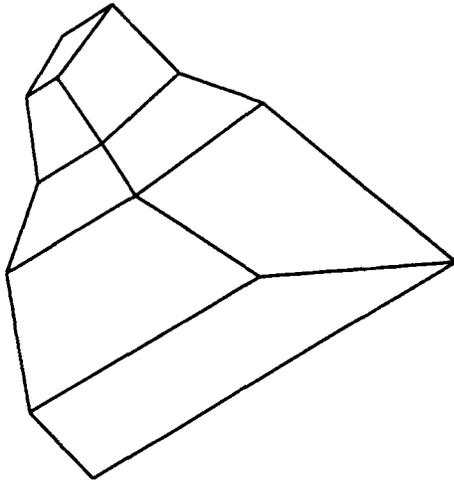


Figure 1. An example of an object used in the experiment (polyhedron condition). In this figure, the hidden edges were removed for better visualization. In the experiment, all edges were visible.

image). In the vertices condition, only the symmetry constraint was present. One could argue that in this condition the images were also topologically stable, except that the topological information present in the images was quite trivial—each image was a mere set of 16 points. In the polygonal line condition, no constraints were present.

The stimuli were shown on a computer monitor by means of the KDE. Each experimental trial consisted of two different sequences of images that represented either the *same* or *different* shapes. By “same” shapes we mean the same Euclidean structure up to size scaling. Consider first the trial with same shapes. Each sequence simulated rotation of the object around the vertical axis by an angle of 19° (20 images in steps of 1°). At the beginning of the first sequence, the object was put at a random 3-D orientation relative to the frame of reference in which it was generated. At the end of the sequence, the object disappeared. Before the second sequence started, the object was rotated by 71° around the vertical axis. A small random rotation (up to $\pm 20^\circ$) around the other two axes, size scaling (up to $\pm 50\%$ of the object size), and horizontal/vertical translations (up to $\pm 10\%$ of the image size) were added in order to minimize some 2-D artifactual cues. For example, if the size of the object was not randomized, then an identical size of the images shown in the first and second sequences would be a reliable cue in the case of same trials. Similarly, without the random rotations about the two horizontal axes, the subject could base his judgments on the distribution of image features along the vertical axis. Note that since the two sequences were separated by approximately 90° , they always contained different images regardless of whether the shape shown in the second sequence was the same as or different from the shape shown in the first sequence.

In a trial with different shapes, the second sequence involved a new, randomly generated object. This different object was rotated relative to the frame of reference, translated and scaled identically as in the trial with same shapes. In the polygonal line condition, in which a polygonal line connected the vertices in a random order, the order was the same in both sequences of a given trial, but it changed between trials. The stimuli were clearly visible: The luminance of the stimulus was 50 cd/m^2 , and the luminance of the background was less than 0.01 cd/m^2 .

Procedure. A signal detection experiment with confidence rating was used. Each session consisted of an equal number of same and different trials, presented in a random order. Viewing was monocular from a distance of 79 cm. The subject’s task was to dis-

criminate whether the two sequences contained the same or different shapes. After each trial, the subject received feedback about the accuracy of the response. Stimulus type was constant within a given session, and the order of the three sessions was random. For the authors (Subjects Z.P. and A.S.), each session consisted of 600 trials; for the other 2 subjects (Subjects M.C. and G.W.), each session consisted of 300 trials. The subject’s performance was evaluated by the discriminability measure d' and its standard error (Macmillan & Creelman, 1991).

Results and Discussion

In all three conditions, the subjects reported that they clearly saw the stimuli as three-dimensional. This means that the KDE was an effective depth cue in all three conditions. However, this clear perception of three-dimensionality did not lead to equally good performance in the shape discrimination task. Results are shown in Figure 2. The ordinate shows d' . Higher d' represents better (more accurate) performance. The polyhedron condition led to much better performance than did the other two conditions (the differences in d' were much larger than the standard errors), and performance in the case of the polygonal line was close to chance level. These results clearly support our hypothesis that shape constancy can be reliably achieved from novel views when the objects contain a number of constraints, and shape constancy may not be reliably achieved when all the constraints are absent.²

Note that Subject A.S. (one of the authors) performed better in the vertices condition than in polygonal line condition, in contrast with the other 3 subjects, who performed equally poorly in these two conditions. One possible reason is that A.S. received very extensive practice in viewing the stimuli due to his work in designing them. This explanation is not very plausible, however, because we did not observe any improvement in performance from one practice session to another in the case of Subjects A.S. and Z.P. (the other two subjects ran only one practice session). An alternative explanation is that A.S. could determine how to connect (group) the vertices to form the two boxes so that he could visualize the polyhedron. In fact, this is what A.S. reported. Z.P., on the other hand, reported that he could hardly visualize the polyhedron in the vertices condition. To test these two possible explanations, Z.P. ran a control experiment (one session with 300 trials) using the vertices condition. In this control experiment, the eight vertices representing the smaller box were shown in red, whereas the remaining vertices were shown in white (as in the main experiment). Obviously, color did not affect the 3-D structure of the object, but it provided a cue for the correct grouping (connecting) the vertices in order to visualize the polyhedron. If it was not grouping, but simply extensive practice, that led to higher performance of A.S. in the vertices condition, Z.P.’s performance in this control experiment should have been the same as that in the main experiment.

Z.P. achieved $d' = 0.95$ ($SD = 0.13$) in this control experiment. This performance was similar to that of A.S. in the vertices condition in the main experiment. This result

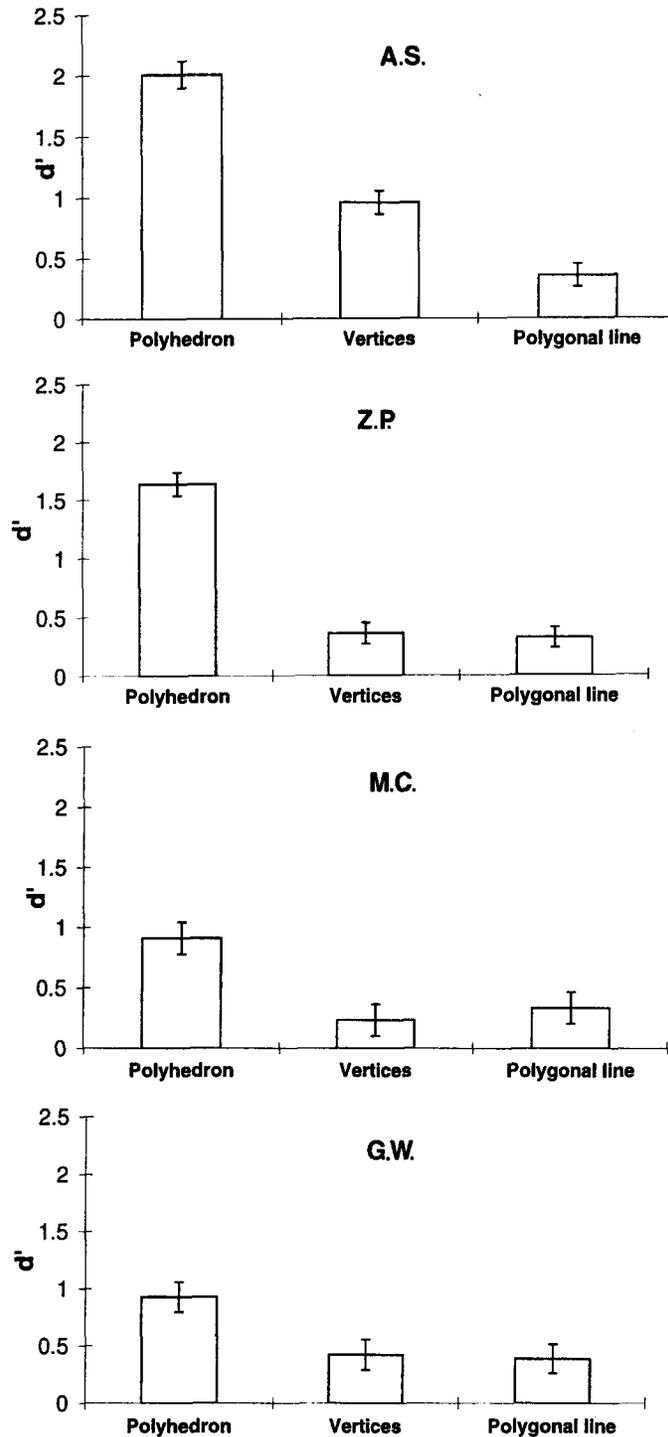


Figure 2. Results of Experiment 1. The ordinate shows d' . The error bars represent ± 1 standard deviation of d' . Subjects M.C. and G.W. received much less practice before running the experiment. G.W. was naive with respect to the hypotheses being tested.

suggests that it is the ability to group the vertices, rather than extensive practice, that improves performance in the vertices condition.

To summarize results of Experiment 1: We established that the presence of a number of features that constrain

the shape of an object, such as planar contours, volumetric primitives, symmetry, and topological information about which contours share edges, is crucial in achieving shape constancy. The next step is to test the individual effects of those various constraints to determine whether

all of them are equally important. These tests were performed in the next two experiments. These experiments used essentially the same method as that in Experiment 1.

EXPERIMENT 2

The Effect of Planarity and Symmetry on Shape Constancy

Method

Subjects. Four subjects (including the authors) were tested. Z.P. received more practice relative to the other 3 subjects. Subjects J.C. and C.L. were naive with respect to the hypotheses being tested.

Stimuli. Four types of stimuli were used: (1) *polyhedron*, which was the same as the stimulus in the polyhedron condition of Experiment 1; (2) *partially nonplanar symmetric polyhedron*, in which the planarity constraint of the contours of faces was removed (Note that the polyhedron still had a plane of symmetry, similar to the stimulus from the polyhedron condition. It is easy to show that if a quadruple of points has a plane of symmetry, such that the points are not on this plane, this quadruple is coplanar.³ Thus, even though the planarity constraint was not explicitly used in generating the stimulus in this condition, the symmetry constraint implied planarity of some contours, which is why this stimulus was called "partially nonplanar"); (3) *planar asymmetric polyhedron*, which was obtained by using only one half of the polyhedron from the polyhedron condition (more specifically, one of the two symmetric parts); and (4) *nonplanar asymmetric polyhedron*, in which the contours were nonplanar and the polyhedron was asymmetric.

In the partially nonplanar symmetric condition, the only constraint that was removed, relative to the polyhedron, was that not all contours were planar. In the planar asymmetric condition, the symmetry constraint was removed. In the nonplanar asymmetric condition, both planarity and symmetry constraints were removed.

Procedure. All subjects ran four sessions (in a random order); each session consisted of 600 trials.

Results and Discussion

Again, the best performance was in the polyhedron condition (see Figure 3). Removing the planarity constraint in some of the contours harmed performance slightly in all 4 subjects. Removing the symmetry constraint led to a small deterioration of performance in 3 subjects and to a larger deterioration in 1 subject (A.S.). Removing both the planarity and symmetry constraints led to performance that was worse by more than a factor of 2 in all 4 subjects, and this difference was much larger than the standard errors. Since the symmetry constraint implies planarity of some contours (see note 3), and since symmetry itself does not seem to affect performance much, we conclude that the deterioration of performance from the polyhedron condition to the nonplanar asymmetric condition was related mainly to the complete removal of the planarity constraint. Figure 4, which shows average d' 's from all 4 subjects, illustrates the overall trends in the results. The averages were computed by weighing d' 's by their variances (Bevington, 1969).

Experiment 3 was designed to test the effect of constraining among shapes of the individual contours. In the case of the polyhedron stimulus, the information about shapes of the planar contours, plus the topological information about which contours share vertices and/or edges, greatly constrains the possible reconstructions of the

polyhedra (the ambiguities of the reconstruction involve such cases as reflections of the object or its parts with respect to any plane containing an individual face).⁴ To test the role of the topological constraints, we presented only three faces of the polyhedron (three quadrilaterals) in such a way that each of two pairs of the faces shared only one vertex. As a result, the shapes of the three contours and this weak topological constraint determined an infinite number of possible interpretations of the object: One can rotate each contour about the common vertex or change its shape arbitrarily without affecting the shapes of the other two contours. If the joint constraining among shapes of the contours is critical to shape constancy, then this stimulus should lead to deterioration of performance, relative to the polyhedron stimulus.

EXPERIMENT 3

The Role of Topological Constraints

Method

Subjects. The 4 subjects who participated in Experiment 2 were tested again in Experiment 3. Again, J.C. and C.L. were naive about our hypotheses.

Stimuli. Four types of stimuli were used: (1) *polyhedron*, which was the same as the stimulus in the polyhedron condition of Experiments 1 and 2; (2) *three quadrilaterals* of the polyhedron, in which each of two pairs of the quadrilaterals shared one vertex (thus, the topological constraints were very weak); (3) *three triangles*, in which each of the triangles was one part of each quadrilateral that was used in the three-quadrilaterals condition (here, the topological constraints were equally weak as in the previous condition. Using triangles rather than quadrilaterals allowed additional testing of the planarity constraint. Although triangles are planar, the planarity constraint is weaker relative to quadrilaterals, since any three, but not any four, points are coplanar); and (4) *polygonal line* connecting the seven vertices that were used in the previous condition (specifically, the polygonal line stimulus was obtained by removing one side from each of the three triangles).

The polygonal line condition was added as a control, after testing the first three conditions. Since the performance in the three-triangles condition was better than that in the polygonal line condition in Experiment 1, we wanted to examine the nature of those differences. The polygonal line in this experiment was different from the polygonal line in Experiment 1 in that it was simpler (involved only 7, rather than 16, vertices). Furthermore, the polygonal line in this experiment led to fewer intersections in the image, and, thus, the image was topologically more stable. Finally, note that the three-triangles condition was different from the polygonal line condition in that the latter stimulus did not contain volumetric primitives (3-D junctions).

Procedure. All subjects ran the first three conditions (in a random order), and the polygonal line condition was added as the last condition.

Results and Discussion

Results of individual subjects are shown in Figure 5, and the averaged results are shown in Figure 6. Again, the best performance was in the polyhedron condition. Removing the topological constraints harmed the performance substantially in all 4 subjects: The discriminability in the three-quadrilaterals condition and in the three-triangles condition was much lower than that in the polyhedron condition (these differences were large rela-

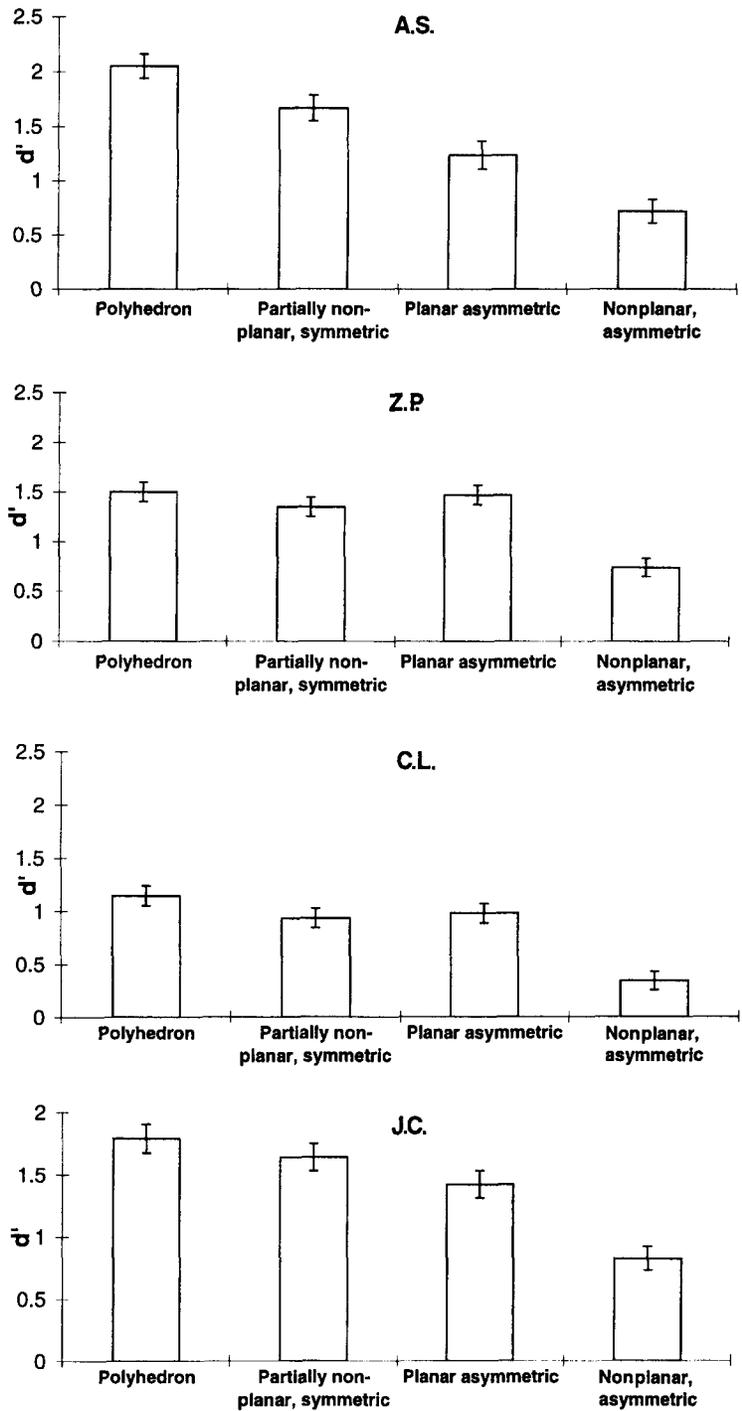


Figure 3. Results of individual subjects in Experiment 2. Subjects C.L. and J.C. were naive with respect to hypotheses being tested.

tive to standard errors). These results show that the topological constraints are quite important for shape constancy.

Performance in the three-triangles condition was substantially lower than that in the three-quadrilaterals condition in 3 subjects, and it was slightly higher in the case of Subject J.C. (in J.C., this difference was small relative

to the standard error). This result is consistent with our prediction based on the fact that a triplet of points is trivially planar and, therefore, the planarity constraint was weaker in the three-triangles condition than in the three-quadrilaterals condition.

Interestingly, even though the three-triangles condition did not contain many constraints, performance was

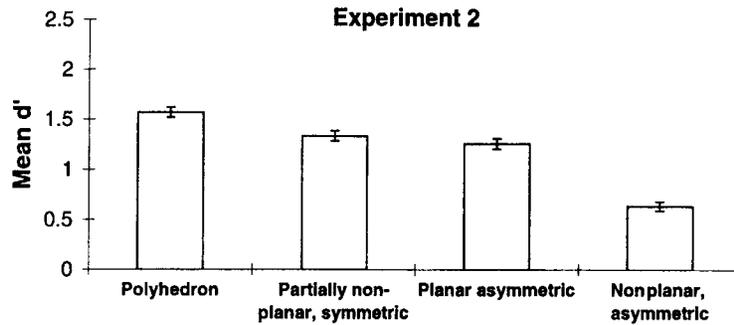


Figure 4. Averaged results from all 4 subjects in Experiment 2.

still somewhat better than that in the polygonal line condition in Experiment 1. To explain this difference, we ran a control session with the polygonal line stimulus (polygonal line condition). The performance in the polygonal line condition in this experiment was similar to that in the three-triangles condition (the differences were small and comparable to the standard errors). These two types of stimuli (three triangles and polygonal line) differed from the polygonal line in Experiment 1 in that they were simpler and they led to images that were topologically more stable. This means that simplicity and topological stability of the image(s) may contribute to shape constancy.

Finally, note that removing volumetric primitives (three triangles vs. polygonal line, in this experiment) did not seem to affect performance. This result suggests that shape constancy need not involve volumetric primitives.

SUMMARY AND CONCLUSIONS

In this study, we tested shape constancy from novel views for several types of stimuli. All stimuli had the same overall 3-D structure involving vertices of a polyhedron. The differences among the stimuli were related to the presence of some geometrical features constraining the shape of the object. Performance of the subjects in our experiments ranged from quite reliable ($d' \approx 2$) to almost chance level ($d' \approx 0.3$). One of the two factors having critical effect on performance was the presence of planar contours (Experiment 2). The other critical factor was related to the presence of topological features that, together with shapes of planar contours, constrain the possible interpretations of the polyhedra (Experiment 3). Symmetry and complexity of the 2-D image(s) also contribute to shape constancy, although to a much lesser extent. Finally, we did not find any effect of the presence of volumetric primitives like 3-D junctions. We would like to point out, however, that the constraints that were identified in our experiments do not necessarily form a complete list of the constraints that operate in the case of shape constancy. It is quite possible that the stimuli of Farah et al. (1994) or of Biederman and Gerhardstein (1993) contained some other features that contributed to shape constancy.

Consider now ways in which the shape constraints may be involved in perceptual mechanisms. One possibility is that the visual system forms the representations of objects, assuming that objects in our environment satisfy the constraints. This assumption could speed up the information processing, and, when it is violated, the percept would be inaccurate and unreliable. Although this way of using shape constraints is possible, it seems too restrictive. It implies, for example, that, when the contours of objects are nonplanar, the observer perceives them as planar. This does not seem to be the case (the reader is encouraged to view samples of our stimuli; see note 1). The observer does see the nonplanarity of the contours.

An alternative way of using constraints in shape constancy is related to regularization methods of solving inverse problems. Consider the relation between retinal image formation and visual perception. Retinal image formation is a transformation from a 3-D scene to the retinal image(s) according to the rules of optics. This is a direct (forward) problem. Having the retinal image(s), the observer is faced with the task of making inferences about the 3-D scene. Thus, the observer has to solve the inverse problem (Poggio, Torre, & Koch, 1985; Tikhonov & Goncharsky, 1987). It is known that many inverse problems in science are ill posed (the problem is well posed when it has a solution, the solution is unique, and it depends continuously on the initial data; otherwise, it is ill posed). In order to solve an ill-posed problem, the observer has to impose constraints on the family of possible solutions. Such a method of solving inverse problems is called *regularization* (Tikhonov & Goncharsky, 1987). The constraints, whose role in shape constancy was demonstrated in our experiments, could be involved in the operation of a regularizing mechanism in order to solve the (ill-posed) shape constancy problem.

The fact that shape perception does involve constraints that lead to a unique percept even if the number of possible interpretations is infinite has been well known. It was discussed for the first time by Gestalt psychologists in the context of their principles of organization (Koffka, 1935). Consider Figure 1. The retinal image produced by this figure could have been produced by an infinite number of 3-D structures, including asymmetrical ones whose

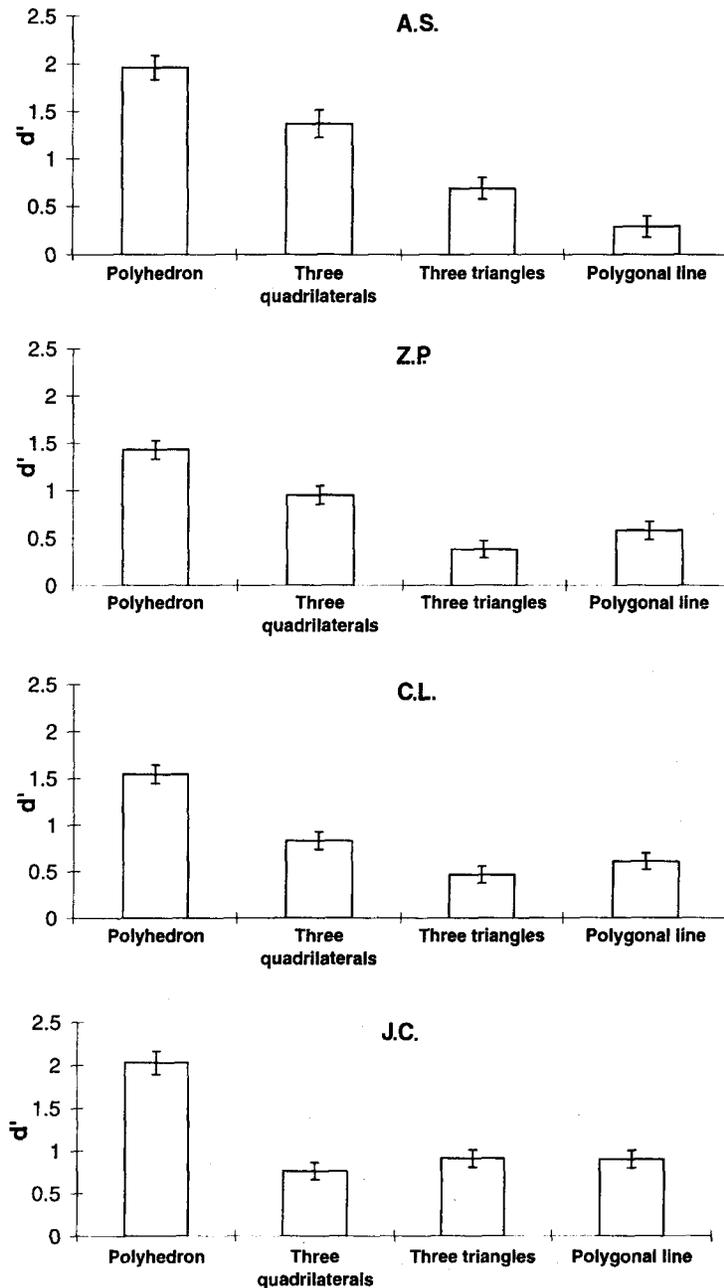


Figure 5. Results of individual subjects in Experiment 3. Subjects C.L. and J.C. were naive with respect to hypotheses being tested.

contours are not planar. Yet, the reader perceives a symmetric polyhedron with planar faces. This phenomenological observation has been a subject of a number of psychophysical experiments. These experiments showed that shape perception involves solving an optimization problem with the objective function containing measures of simplicity or compactness of the 3-D interpretation (Attneave & Frost, 1969; Hochberg & Brooks, 1960; Perkins, 1972; Pizlo & Scheessele, 1998). These simplicity and compactness constraints allowed a unique per-

ceptual interpretation to be obtained from a single image. Our experimental results presented in this paper extend these prior results by demonstrating the role of planar contours and topological features as possible regularizing constraints in shape-from-motion phenomenon.

One may ask a question about ecological justification of the constraints that were found to operate in shape constancy. If constraints such as planarity of contours were not present in the objects from our natural environment, one would expect the percepts of shapes to be often

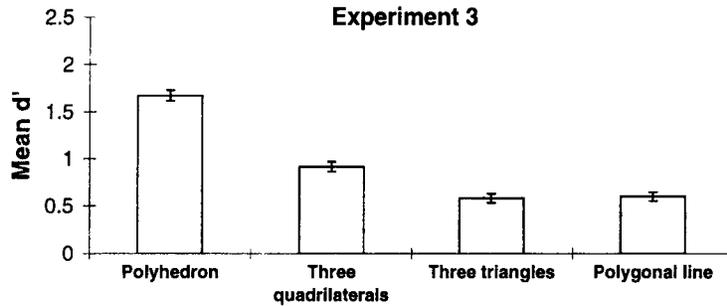


Figure 6. Averaged results from all 4 subjects in Experiment 3.

inaccurate and not constant. However, shape perception in our everyday life seems to be accurate and constant. This means that the natural environment is likely to include such constraints as (approximate) planarity of contours and topological features. This suggestion is consistent with some observations reported in the past (e.g., Stevens, 1981, 1986).

REFERENCES

- ATTNEAVE, F., & FROST, R. (1969). The determination of perceived tridimensional orientation by minimum criteria. *Perception & Psychophysics*, *6*, 391-396.
- BEVINGTON, P. R. (1969). *Data reduction and error analysis for the physical sciences*. New York: McGraw-Hill.
- BIEDERMAN, I., & GERHARDSTEIN, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception & Performance*, *19*, 1162-1182.
- BÜLTHOFF, H. H., & EDELMAN, S. (1992). Psychophysical support for a 2-D view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences*, *89*, 60-64.
- FARAH, M. J., ROCHLIN, R., & KLEIN, K. L. (1994). Orientation invariance and geometric primitives in shape recognition. *Cognitive Science*, *18*, 325-344.
- HOCHBERG, J., & BROOKS, V. (1960). The psychophysics of form: Reversible perspective drawings of spatial objects. *American Journal of Psychology*, *73*, 337-354.
- KOFFKA, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt Brace.
- MACMILLAN, N. A., & CREELMAN, C. D. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.
- PERKINS, D. N. (1972). Visual discrimination between rectangular and nonrectangular parallelepipeds. *Perception & Psychophysics*, *12*, 396-400.
- PIZLO, Z. (1994). A theory of shape constancy based on perspective invariants. *Vision Research*, *34*, 1637-1658.
- PIZLO, Z., & SCHEESSELE, M. R. (1998). Perception of 3D scenes from pictures. *Proceedings of the SPIE*, *3299*, 410-422.
- POGGIO, T., TORRE, V., & KOCH, C. (1985). Computational vision and regularization theory. *Nature*, *317*, 314-317.
- ROCK, I., & DiVITA, J. (1987). A case of viewer-centered object perception. *Cognitive Psychology*, *19*, 280-293.
- STEVENS, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence*, *17*, 47-73.
- STEVENS, K. A. (1986). Inferring shape from contours across surfaces. In A. P. Pentland (Ed.), *From pixels to predicates: Recent advances in computational and robotic vision* (pp. 93-110). Norwood, NJ: Ablex.
- TIKHONOV, A. N., & GONCHARSKY, A. V. (1987). *Ill-posed problems in the natural sciences*. Moscow: Mir.
- ULLMAN, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.

NOTES

1. To see examples of the stimuli, the reader is referred to the demo on our Website: <http://bigbird.psych.purdue.edu/shapedemo>. Note that this demo does not show the entire experimental trial, which consisted of two short sequences of motion. Instead, it shows the objects themselves undergoing 3-D rotation. The speed of the rotation can be adjusted by the viewer, and the duration of the rotation is under the viewer's control. Pressing the reset button generates another random example from a given experimental condition.

2. Note that the geometrical information that was present in the visual stimulus (20 images of a rotating rigid object) was sufficient, theoretically, for a complete reconstruction of the 3-D object in all conditions (e.g., see Ullman's, 1979, structure-from-motion algorithm). Thus, if the theoretical models based on geometrical analysis of images were an adequate explanation of human perception of shape from motion, the subjects' performance in our experiments would have been the same across all conditions. It was not.

3. PROOF Let point A be symmetric to B with respect to plane π , and let point C be symmetric to D with respect to the same plane π . It follows that line AB is orthogonal to π , and similarly, line CD is orthogonal to π . This implies that AB is parallel to CD. Hence, these lines—and, thus, all four points—are coplanar.

4. Note that the topological constraints that provide the information about which contours share vertices and/or edges are available in each 2-D image.

(Manuscript received May 29, 1997,
revision accepted for publication July 29, 1998)