

Orientation Invariance and Geometric Primitives in Shape Recognition

MARTHA J. FARAH

University of Pennsylvania

ROBIN ROCHLIN

KAREN L. KLEIN

Carnegie Mellon University

Although it is generally assumed that vision is orientation invariant, that is, that shapes can be recognized regardless of viewing angle, there is little evidence that speaks directly to this issue, and what evidence there is fails to support orientation invariance. We propose an explanation for the previous results in terms of the kinds of shape primitives used by the visual system in achieving orientation invariance: Whereas contours are used at stages of vision that are not orientation invariant, surfaces and/or volumes are used at stages of vision that are orientation invariant. The stimuli in previously reported studies were wire forms, which can be represented only in terms of contour. In four experiments, testing both short-term and long-term memory for shape, we replicated the previous failures of orientation invariance using wire forms, but found relatively good or perfect orientation invariance with equivalently shaped surfaces.

How do we recognize that a given object, viewed from different perspectives so that it casts different images on our retinæ, is the same object? Answers to this question can be subdivided into two very broad and general groups: Those that posit visual mechanisms whereby the different images of the same object are assigned the same description by the visual system, and those that posit the need for postperceptual learning. The first group includes a variety of theories that are, themselves, generally contrasted as representing

The research reported here was carried out while Martha J. Farah was at Carnegie Mellon University. We thank Irvin Rock for his interest and encouragement, and his helpful comments on an earlier draft of this article. We also thank David Bennett, Irving Biederman, Michael Corballis, Stephen Palmer, Hal Pashler, and an anonymous reviewer for their comments and suggestions. This research was supported by ONR grant N00014-91-J1546, NIMH grant R01 MH48274, NIH career development award K04-NS01405, and grant 90-36 from the McDonnell-Pew Program in Cognitive Neuroscience.

Correspondence and requests for reprints should be sent to Martha J. Farah, Department of Psychology, University of Pennsylvania, 3815 Walnut St., Philadelphia, PA 19104-6196.

very different approaches to the problem of shape recognition: Theories of vision in which shapes are represented in object-centered coordinate systems, as well as theories in which shapes are represented in a viewer-centered coordinate system with normalization processes that enable different viewer-centered representations of the same shape to be transformed and matched. The relevant commonality among these very different theories is that the representation of shape is not necessarily tied to a particular viewing angle. The visual system can separate shape and orientation information, and thereby determine whether different images could have been projected by the same shape from different viewing angles. This can be accomplished either by eliminating information about viewing angle altogether, as in object-centered theories, or allowing the visual system to transform viewing angle while holding shape constant, as in theories with viewer-centered representations with transformations. In other words, these theories all assume that visual perception is *orientation invariant*.

In the second group are theories in which shape and orientation are represented integrally, and different views of a shape would therefore be represented and stored within the visual system as separately as views of different shapes. According to these latter theories, our ability to know that two different images both came from the same shape would depend on additional, postvisual learning, for example, by seeing the shape continuously turned to reveal that the different images derive from the same shape. Such systems would not be orientation invariant.

The difference between these two types of theory essentially comes to this: When you reencounter a shape from a new perspective, do you know that it is the same shape just by looking at it? If so, visual recognition is orientation invariant. Or, must you have had prior experience with both views in order to know that the same object gave rise to both? If so, then visual recognition is not orientation invariant. Note that the term "orientation invariant" has sometimes been used to refer specifically to object representations that do not contain information about orientation. The usage here is broader than this, in that it refers to the system property of allowing or not allowing generalization across orientations on the basis of one view.

Despite the fundamental nature of this issue, relatively little research has been directed towards it. For the most part, vision researchers have just assumed orientation invariance (e.g., Biederman, 1987; Marr, 1982; Ullman, 1989). However, what data are available seem to suggest the opposite (e.g., Rock, 1983). The goal of the research here was to provide additional evidence on the issue of orientation invariance, and to help reconcile the inconsistency between the widespread assumption of orientation invariance on the one hand, and the data showing absence of such independence on the other.

PREVIOUS STUDIES

Although there is a large literature on the effects of orientation differences on the processing of visual patterns, most of it is not actually relevant to the issue of orientation invariance in visual recognition. For example, one research tradition has examined the effects of noncanonical viewpoints on the speed and accuracy of recognition of familiar objects (e.g., Corballis, 1988; Jolicoeur, 1985; Warrington, 1982). Those studies, which were designed for purposes other than assessing whether visual recognition has orientation invariance in the sense defined previously, might nevertheless seem relevant at first glance. Given that subjects are shown objects in unusual orientations and succeed in identifying them, doesn't this imply orientation invariance? It does not, because the stimuli in those studies were familiar, and subjects might have already learned to associate separate and distinct visual representations of the different views from prior experience.

Research on mental rotation of unfamiliar shapes might also appear to be relevant. For example, Shepard and Metzler (1971) showed that subjects can match different views of three-dimensional shapes even after changes of perspective in depth. However, these experiments tested simultaneous matching ability, not recognition. There are many judgments that we can make about visual stimuli when they are present—and thus supporting early visual representations with large information capacities—that we cannot make once a stimulus has been removed and we must rely on the long-term visual memory representations used in object recognition. For example, we can see whether a random dot pattern has a particular subpattern in it when the pattern is present but generally cannot make such judgments once the pattern has been removed. Perhaps the ability to associate different views of novel objects is also dependent on having the two stimuli simultaneously present.

Cooper (1975) and, more recently, Tarr and Pinker (1990) taught subjects to recognize novel patterns in one particular orientation and then required that the patterns be recognized after differing amounts of picture-plane rotation. Although their prime interest was in the issue of frames of reference and the role of mental rotation in this task, their data are also relevant to the issue of whether visual recognition has orientation invariance. The fact that subjects could perform these tasks with high accuracy implies that visual recognition is orientation invariant, at least for differences in picture-plane orientation.

Unfortunately, there are fundamental differences between the way an image changes under picture-plane rotation and depth rotation, which prevent us from assuming that the results of Cooper (1975) and Tarr and Pinker (1990) showing orientation invariance for picture-plane rotations will generalize to depth rotations. Unlike picture-plane rotation, when an object is

rotated in depth there is invariably foreshortening, and sometimes self-occlusion as well. Furthermore, picture-plane rotation is not a very "ecologically valid" kind of orientation change. Although it is the kind of change easiest to implement with tachistoscopes or simple computer graphics, it is rarely seen in the real world as viewers and objects move the respect to one another. The likelihood that a random change in angular position would occur only about the z-axis is very small. For these reasons, it is essential to ask about orientation invariance over depth rotations as well as picture-plane rotations.

Rock and colleagues were the first to address this critical issue in a series of elegant studies over the past decade (e.g., Rock & Di Vita, 1987; Rock, Di Vita, & Barbeito, 1981). Their results were surprising in that they seem to show a lack of orientation invariance for the visual recognition of shapes that had been rotated in depth. In those studies, subjects were shown bent wire forms in an incidental learning task. After viewing the forms from one perspective, subjects were later shown the same forms from both the same and different perspectives, along with new forms, and asked to recognize which had been seen previously. Subjects performed dramatically worse with new than with old perspectives, and in some cases were no more likely to recognize a form from a new perspective than they were to falsely recognize an entirely new form. More recently, Bultoff and Edelman (1992) reported similar experiments in which subjects were familiarized with particular angular, straight-edged wire forms and then tested for recognition in a forced-choice format with the target item presented from the same perspective as before or from a new perspective. Like Rock and colleagues, Bulthoff and Edelman found that subjects' representations of the stimuli were highly viewpoint-dependent.

The advantage of using wire forms in these experiments is to avoid stimulus self-occlusion, and thereby give the orientation-invariance hypothesis a better chance of being proved correct. This makes the results just summarized all the more striking. However, it is possible that the wire forms are somehow special. In this article we will offer an account of how wire forms might be special, in a way that has consequences for orientation invariance.

Many current theories of vision, for example, those shown in Table 1, begin with a representation of the image in terms of contours, but progress on to representations whose primitive elements are surfaces and/or volumes. Whereas early vision is thought to be concerned with extracting contours, grouping them by Gestalt-like principles, and detecting nonaccidental relations among them, the higher levels of visual representation involved in recognition per se are not generally thought to be contour-based, but rather surface or volume-based. One way in which the wire forms are special is that they have no natural interpretation in terms of surfaces or volumes, but only contour. This might be significant from the point of view of achieving orientation invariance.

TABLE 1
Some Theories of Object Representation
and the Shape Primitives Associated with Different Stages of Perceptions

Theory	Primitive	
	Contours	Surfaces and/or Volumes
Biederman (1987)	nonaccidental image properties	geons
Marr (1982)	primal sketch	2.5 D sketch, 3 D model
Pentland (1986)	(unspecified)	needle map, superquadrix

A mechanism designed to operate on representations whose primitive elements are surfaces and/or volumes would presumably not be able to operate on representations built from a different set of primitives. Therefore, it is possible that the absence of orientation invariance with wire forms results from the fact they have no interpretation in terms of surfaces or volumes, and hence cannot be represented by the parts of the visual system that are responsible for orientation invariance.

We tested this conjecture by making two sets of stimuli, one from loops of wirelike contour and one from clay disk surfaces, and comparing the degree to which each type of stimulus was recognizable after a depth rotation. In order to isolate the difference between contours and surfaces as the cause of any difference we might find, the two sets of stimuli were yoked with one another in the following way: Each surface shape was the result of interpolating a surface within a corresponding contour shape. Thus, one can think of the yoked pairs as surface and contour "versions" of the same shape. Figures 1 and 2 show a pair of such stimuli, separated and superimposed to show the relation between the contour boundary and interpolated surface. In the following two experiments, we compared the orientation invariance for these two types of stimuli.

EXPERIMENT 1

In this experiment, subjects' ability to recognize shapes after a 5-second retention interval was tested. By comparing recognition of shapes viewed in the same orientation and different orientations, we can assess orientation invariance in short-term memory for shapes. By comparing shape-orientation independence for the surfaces and contours, we can test the conjecture that surfaces may be mentally represented in a way that allows more orientation invariance than contours.

Methods

Materials

Eight differently shaped surfaces were created by bending 3×4 in. (7.62×10.16 cm) oval disks of modelling clay into gently curved three-




Figure 1a. An example of a contour stimulus.




Figure 1b. An example of a surface stimulus.



Figure 2. The stimuli from Figure 1, superimposed to demonstrate their similarity in shape.

dimensional shapes. In order to prevent subjects from recognizing these shapes by recognizing individual distinctive parts, we avoided shapes with distinct features such as sharp ridges or small, high-curvature turns in the surface. As shown in Figure 1a, the resulting shapes resembled curled potato chips. Eight contour shapes, shown in Figure 1b, were made by bending wax-covered modelling strings to the shape of the surface boundary. Figure 2 shows the way in which the contour and surface shapes were matched. Two copies of each surface and contour shape were made, so that the "standard" and "comparison" items would not be physically identical, and subjects could therefore not identify a shape on the basis of some small flaw in the clay or wax. Two additional pairs of shapes were constructed for practice trials. All shapes were mounted on paper plates, the bottoms of which were labelled with shape identity and orientation (i.e., which direction is frontwards) for the use of the experimenter.

There were a total of 64 trials in the experiment. For the 32 trials of either surfaces or contours, each shape served as the standard four times, always presented in a particular orientation: Twice followed by its twin as the comparison stimulus ("same" trials) and twice followed by a different shape as the comparison stimulus ("different" trials). The pairings of shapes for the different trials were the same for the surface and contour stimuli. All shapes occurred equally often as comparison stimulus. For one of the two same trials for each shape, the twin comparison stimulus was presented in the same orientation as the standard ("same orientation—same" trials) and for the other the twin comparison stimulus was rotated by 45° about the y-axis ("different orientation—same" trials).

The viewing apparatus consisted of two separate platforms. The lower one held the stimulus being displayed, and subjects rested their chins on the upper one, which was 9 in. (22.86 cm) higher. Stimuli were placed approximately 34 in. (86.36 cm) away from the subject, resulting in a viewing angle of approximately 24° . The apparatus was placed in a quiet hallway outside the lab room where the stimuli were arrayed so that subjects could not see the stimuli except for the brief presentations from particular orientations during the experiment.

Procedure

Subjects were tested individually. On each trial of the experiment, subjects were shown a standard stimulus for 3 seconds, followed by a 5-second interval, followed by a comparison stimulus, which remained present until subjects made their response. Timing was approximate, based on the experimenter mentally counting. Subjects closed their eyes during the stimulus placement and removal, and stimuli were transferred back and forth between the lab and the hallway in a box with high sides to prevent exposure of the stimuli for additional time or at other orientations. Between trials the experimenter

repaired to the lab room to record the subject's response, select the stimuli for the next trial, and check their orientation.

Half of the subjects performed the task with surfaces first, and half with contours first. Within these two blocks of trials, the shapes were presented in a fixed pseudorandom order, same for both surfaces and contours. The order was created by concatenating four repetitions of the eight shapes in different random order each time, and distributing same orientation—same, different orientation—same, and different conditions among these trials such that no more than four correct same or different responses could occur in a row, and equal numbers of same orientation—same, different orientation—same, and different trials occurred in the first and second halves of the trial order.

Subjects were instructed to try to remember the first shape that was shown on each trial so that they could decide whether the second shape was the same or different, regardless of its orientation. They were given five practice trials, with feedback that included showing the shapes side by side and turning them.

Subjects

Sixteen Carnegie Mellon University undergraduates participated for course credit; all had normal or corrected-to-normal vision.

Results and Discussion

The proportion of correct responses in each condition is shown in Figure 3. Consistent with previous results, subjects showed poor orientation invariance with the contour shapes: Whereas 89.1% of these stimuli were recognized when presented again at the same orientation, only 68.7% were recognized when the orientation was changed. In contrast, subjects showed virtually perfect orientation invariance with the surface shapes, recognizing 77.3% at the same orientation and 75.8% when the orientation was changed. The predicted pattern of results was found to be reliable using a matched-pairs Wilcoxon test comparing the difference in the size of the performance decrement with rotation for contours and surfaces, $T=28$, $N=16$, $p<.025$. This pattern was predicted by the hypothesis that orientation invariance in visual recognition involves visual representations whose primitive elements are surfaces rather than contours. In addition, as shown in Figure 3, subjects correctly rejected 89.8% of the different contour shapes and 73.4% of the different surface shapes.

In terms of their overall difficulty of recognition, the surface shapes were, if anything, harder to recognize than the contour shapes. However, orientation invariance is virtually perfect for surface shapes, and is poor for contour shapes. This is consistent with the hypothesis that the parts of the visual system that normally accomplish orientation invariance represent surfaces rather than contours.

Experiment 1




Figure 3. Performance of subjects in Experiment 1: The percentage of correct recognition for same shapes at the same orientation and same shapes at a different orientation, and percentage of correct rejection of different shapes, for contour and surface versions of the shapes.

EXPERIMENT 2

The previous experiment involved "recognition" of shapes over a 5-second interval. Such short-term memory for shape might involve different representations from the ones required for recognition of a shape on the basis of long-term memory. In this experiment, we assessed the orientation invariance of long-term memory recognition for the same contour and surface shapes used in Experiment 1.

Methods

Materials

The shapes and viewing apparatus from Experiment 1 were used in this study.

Procedure

As before, subjects were tested individually, with half the subjects performing the task with surfaces first and half with contours first. In the learning phase of the experiment, subjects were told that they would be required to learn the shapes of four stimuli in order to be able to recognize them later. The same four stimuli were shown to the subject one at a time repeatedly, first for 30 seconds each, then for 20 seconds each, and finally for 10 seconds each. Immediately following the learning phase, subjects were tested with 32 trials of test items, consisting of four presentations of each of the eight shapes. Two of the four presentations of each shape from the study set were in the same orientation as they were seen during the learning phase (same orientation—same) and two were rotated by 45° (different orientation—same). The order of trials was created by concatenating four random orderings of the eight stimulus shapes and distributing same orientation—same, different orientation—same, and different conditions among these trials such that no more than four correct same or different responses could occur in a row, and equal numbers of same orientation—same, different orientation—same, and different trials occurred in the first and second halves of the trial order. Subjects could inspect the test stimuli for as long as they wanted, although they were encouraged to answer within a few seconds and virtually always did so.

Subjects

Sixteen Carnegie Mellon University undergraduates participated for course credit; all had normal or corrected-to-normal vision. None had participated in Experiment 1.

Results and Discussion

The proportion of correct responses in each condition is shown in Figure 4. As expected on the basis of Rock et al's (1981) results and those of Experiment 1, subjects showed poor orientation invariance with the contour shapes:

Experiment 2




Figure 4. Performance of subjects in Experiment 2: The percentage of correct recognition for same shapes at the same orientation and same shapes at a different orientation, and percentage of correct rejection of different shapes, for contour and surface versions of the shapes.

Whereas 85.2% of the contour shapes were recognized when presented at their original orientations, only 64.8% were recognized when rotated. In contrast, subjects' recognition of the surfaces was virtually unaffected by orientation: When presented in their original orientations 76.6% of the surface shapes were recognized, and when rotated 74.2% were recognized. The reliability of the predicted pattern was again tested with a matched-pairs Wilcoxon test comparing the size of the performance decrement with rotation for contours and surfaces, $T=1$, $N=12$, $p < .005$. This was predicted by the hypothesis that orientation invariance in visual recognition involves visual representations whose primitive elements are surfaces rather than contours. Figure 4 also shows that subjects correctly rejected 87.9% of the different contour shapes and 78.1% of the different surface shapes.

As in the previous experiment, subjects showed orientation invariance for the shapes of surfaces, but not for contours that have no interpretation in terms of surfaces or volumes. Also, as in the previous experiment, this was true even though subjects found the task generally harder with surfaces: When the orientation of the shapes was the same, or when the two shapes were different, subjects performed better with the contours.

EXPERIMENT 3

This experiment was essentially a replication of Experiment 2 with three main differences: We varied the angle through which stimuli were rotated; we used a slightly larger range of stimulus shapes; and all of the stimuli were presented to subjects on videotape.

Methods

Materials

The shapes from the previous experiment were videotaped from an angle of approximately 45° above level for both the study and test trials. Because the slight irregularities in the shapes were invisible on videotape, we were able to use the same physical stimuli for both study and test, rather than using twin stimuli as in the previous two experiments. This allowed us to use a total of 10 shapes rather than the 8 used previously because we had originally made 10 shapes and had to eliminate 2 of them because the surface versions did not match their twins closely enough. Each of the 10 shapes was presented four times: twice in its standard orientation (i.e., the orientation in which it would be studied if it were in the study set) and twice at a different orientation. One of these different orientations was 30° from the standard orientation and the other was 60°. The shapes were displayed on a 19-in. (48.26) monitor, and ranged in width from 6–10 in. (15.24–25.40 cm), depending on the shape and viewing angle. Subjects' distance from the monitor was not controlled, but was approximately 2–3 ft (0.6096–0.9144 m).

Procedure

As before, half of the subjects performed the task with surfaces first, and half with contours first. In addition, half of each of these halves was trained with 5 of the shapes in the learning phase, and the other half were trained with the other 5. Subjects were tested either 1 or 2 at a time. The timing of stimulus presentations for the learning phase was the same as in the previous experiment. In the test trials, each stimulus was shown for 9 seconds, during which time the subject recorded his or her response on a response sheet. Subjects almost always responded before the end of the 9 seconds, but the experimenter would pause the videotape to allow additional time for response recording when needed. The order was created by concatenating four repetitions of the 10 shapes in different random order each time, and distributing same orientation—same, different orientation—same, and different conditions among these trials such that no more than four correct same or different responses could occur in a row, and equal numbers of same orientation—same, different orientation—same, and different trials occurred in the first and second halves of the trial order.

Subjects

Sixteen Carnegie Mellon University undergraduates participated for course credit; all had normal or corrected-to-normal vision. None had participated in Experiments 1 or 2.

Results and Discussion

The proportion of correct responses in each condition is shown in Figure 5. Once again, subjects showed poor orientation invariance with the contour shapes, recognizing 87.3% of them when presented at the same orientation at which they were studied, and only 56.9% of them when they were rotated. In contrast, subjects were less affected by rotation with the surface shapes, recognizing 90.0% of them when viewed from the original orientation, and 72.3% of them when they were rotated. As before, the predicted interaction between orientation (same or different) and stimulus type (surface or contour) was reliable by matched pairs Wilcoxon test, $T=21.5$, $N=14$, $p<.05$. In general, rotation by 30° had a smaller effect on recognition than rotation by 60° (74.4% vs. 55.6% respectively), $T=5$, $N=12$, $p<.01$. Although the difference between the two orientations appeared larger for contours than for surfaces, this interaction was not statistically significant. Turning to different trials, subjects correctly rejected 62.9% of the contour shapes and 67.8% of the surface shapes.

Although the same predicted interaction was observed in this experiment as in the previous ones, namely the greater effect of rotation on recognition on the contour than surface shapes, the recognition of both contour and surface shapes showed greater effects of orientation change in this experiment than in the previous two. This is presumably due to the loss of

Experiment 3




Figure 5. Performance of subjects in Experiment 3: The percentage of correct recognition for same shapes at the same orientation and same shapes at a different orientation, and percentage of correct rejection of different shapes, for contour and surface versions of the shapes.

binocular depth information when viewing stimuli on videotape. Nevertheless, despite the greater effects of rotation for all stimuli, the surface shapes retained relatively more orientation invariance than the contour shapes. We will return to the issue of the role of depth information in orientation invariance in the General Discussion.

EXPERIMENT 4

The purpose of Experiment 4 is to assess the degree to which the results of the previous three experiments are general across different shapes. Although cognitive psychologists usually test the reliability of their effects over different subjects, in research on pattern recognition we should be concerned about reliability over different patterns as well. Experiment 4 was designed to allow this by introducing a larger set of stimulus shapes.

Methods

Materials

Thirty new surface and contour forms were created using the same method described for Experiment 1. Only one copy of each stimulus was made because they were to be videotaped and, as noted earlier, the slight differences between the different copies of a single shape were not visible on videotape. Each stimulus appeared in three trials on the tape: once in a same orientation—same trial, once in a same orientation—different trial, and once in a different trial. For different orientation—same trials, the stimuli were rotated 50°. As before, the shapes were videotaped from an angle of 45°, and subtended approximately 6–10 in. (15.24–25.40 cm) on the monitor, which subjects viewed from a distance of approximately 2–3 ft (0.6096–0.9144 m).

Procedure

Because of the difficulty of committing a large number of shapes to long-term memory, this experiment tested short-term memory for the stimuli. On each trial, a stimulus was shown for 3 seconds, followed by a blank screen for 6 s, followed by the second shape for 4 s. The intertrial interval was 1 second. A pseudorandom order of trial was created by concatenating three repetitions of the 30 shapes in a different random order each time, and distributing same orientation—same, different orientation—same, and different conditions among these trials such that no more than four correct same or different responses could occur in a row, and equal numbers of same orientation—same, different orientation—same, and different trials occurred in the first and second halves of the trial order. Subjects were instructed to try to remember the first shape that was shown on each trial, so that they could decide whether the second shape was the same or different, regardless of its orientation. They were also told that there would be twice as many same

trials as different trials. They were given five practice trials with feedback before the experiment began. As usual, half of the subjects were tested with contours first and half were tested with surfaces first.

Subjects

Ten Carnegie Mellon University undergraduates participated for course credit; all had normal or corrected-to-normal vision. None had participated in Experiments 1, 2, or 3.

Results

The proportion of correct responses in each condition is shown in Figure 6. Consistent with previous results, subjects showed poor orientation invariance with the contour shapes: Whereas 94% of these stimuli were recognized when presented again at the same orientation, only 49.3% were recognized when the orientation was changed. As predicted, subjects showed relatively more orientation invariance with the surface shapes, recognizing 91.7% at the same orientation and 72.7% when the orientation was changed. Eighty percent of contour shapes and 75.3% of surface shapes were correctly classified as different. The prediction of greater decrement with rotation of contours than surfaces was tested with a matched-pairs Wilcoxon test over stimuli and found to be reliable, $T=55.5$, $N=30$, $p<.005$. This demonstrates that the greater orientation invariance allowed by surfaces relative to contours is general over the types of stimuli used in these studies.

GENERAL DISCUSSION

In answer to the question of whether vision is orientation invariant, we would answer "well, yes and no." Vision does appear to be invariant over moderate changes in depth orientation for some kinds of stimuli, but not for others. Specifically, wire forms show poor orientation invariance, and clay surfaces show good orientation invariance when viewed binocularly over at least the ranges of orientation change used here. One way of interpreting this difference between wire forms and clay surfaces is in terms of the different kinds of shape primitives used by the visual system to represent stimuli at different stages of perception. At early stages of vision, which are not orientation invariant, the stimulus is represented in terms of contours. At later states of vision, where orientation invariance would be accomplished if it exists at all, the stimulus is represented in terms of surfaces and/or volumes. As wire forms cannot be represented in terms of these latter, higher order geometric primitives, they should not be expected to show orientation invariance.

Let us now consider some limitations of these conclusions. The most obvious is that we cannot comment on the range of orientations over which perception of surface shape will be invariant. Our goal was to address the

Experiment 4




Figure 6. Performance of subjects in Experiment 4: The percentage of correct recognition for same shapes at the same orientation and same shapes at a different orientation, and percentage of correct rejection of different shapes, for contour and surface versions of the shapes.

issue of the *existence* of orientation invariance, in the face of previous data that seemed to show its nonexistence. It is possible that depth rotations larger than the 45° change used in our first two experiments would reveal a decrement in recognition accuracy for binocularly viewed surface shapes. However, the finding of perfect orientation invariance for binocularly viewed surfaces after 45° rotations is not trivial; the same amount of rotation produced a failure of orientation invariance for wire shapes.

Another limitation of these results is that they do not explain *why* surfaces are better suited to the computation of an orientation-invariant representation than contours. Although our goal was not to address this question, and our experiments were therefore not designed to provide an answer to it, the unexpected difference between the results with binocularly viewed shapes and videotaped shapes may provide a clue.

The accurate representation of depth is a necessary precondition for orientation invariance. If a viewer cannot appreciate the equivalence of a given distance between two points on a shape when that distance occurs in depth, at one orientation of the shape, and in the picture plane, at the orthogonal orientation (as well as when the distance is comprised of both depth and picture-plane components at intermediate orientations), then the viewer will not be able to determine the equivalence of the shape over depth rotations. The decrement in orientation invariance in the videotape experiments, in which subjects were deprived of binocular depth cues, is consistent with this idea.

Of course, subjects in Experiments 1 and 2 viewed the wire forms binocularly and at a relatively close distance, which should have afforded them good depth vision. Furthermore, Rock, Wheeler, and Tudor (1989) tested their subjects' depth perception for wire forms in their experiment and found it to be quite accurate. Why, then, didn't the wire forms display orientation invariance in these experiments? Doesn't this imply that the representation of depth is not the critical factor in determining the degree of orientation invariance found with surfaces and contours?

The distinction between initial encoding of depth information and ability to retain it and operate upon it may be relevant here. Although subjects may have encoded depth information as accurately from the wire forms as from the clay surfaces in our first two experiments, the differential *redundancy* of depth information in a contour representation and a surface representation may make the latter more robust under the transformations and comparisons involved in the computation of orientation invariance. The depth of a given point on a smooth contour is roughly equivalent to the depth of the points on either side of it, and this redundancy would allow for some degree of noise reduction by the user of these local constraints. The depth of a given point on a smooth surface is constrained by the larger set of points surrounding it on all sides, and this may confer greater resiliency to noise introduced in the processing required for orientation invariance.

Our conjecture that the differential redundancy of depth information in contour-based and surface-based representations underlies the greater utility of surface-based representation for orientation invariance can also explain an observation made by Rock and DiVita (1987), which might have seemed incompatible with our findings. After presenting their experiments with wire forms, Rock and DiVita raised the issue of the generality of their findings for other types of stimuli. Although they did not carry out formal experiments with nonwire stimuli, they did construct two new stimuli objects and reported that those objects did not appear to them to be orientation invariant. They provided photographs of the objects from two perspectives for readers to verify this intuition for themselves. The two objects were a piece of crumpled paper and a complex-shaped mass of clay, both of which had complex, irregularly shaped surfaces. Similar unpublished findings were mentioned by Bultoff and Edelman (1992) with amoeba-like shapes. Interestingly, Hoffman and Richards (1985) pointed out that volumetric-primitive-based object-recognition schemes, such as those listed in Table 1, are poorly suited to objects with irregular crenulated surfaces.

The results here are consistent with the known physiology of vision. Cells in area V1 represent the visual field in terms of contours, that is, edges and lines at different spatial scales (e.g., Desimone, Schein, Moran, & Underleider, 1985). Cells in later visual areas, for example in inferotemporal cortex (IT), do not appear to code shape in terms of contours. For example, spatial frequency filtering of the image has little effect on the responses of at least some cells in this area (Rolls, 1984), and changes in patterns of shadow and light, which change the contours in the scene but not the surfaces and volumes, do not affect recognition if IT is intact but do impair recognition if IT has been ablated (Weiskrantz & Saunders, 1984). Consistent with the present conjecture about orientation invariance and shape primitives, representations in V1 are highly orientation-sensitive. A change in stimulus orientation will result in a completely new population of neurons becoming active. In contrast, at least some representations in IT show good orientation invariance. Some cells will maintain a response selective to shape over at least a 45° depth rotation (Desimone, Albright, Gross, & Bruce, 1984), and changes in stimulus orientation do not affect recognition if IT cortex is intact but do impair recognition if this region has been ablated (Weiskrantz & Saunders, 1984).

REFERENCES

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115-147.
- Bultoff, H.H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences*, *89*, 60-64.

- Cooper, L.A. (1975). Mental rotation of two-dimensional random shapes. *Cognitive Psychology*, 7, 20-43.
- Corballis, M.C. (1988). The recognition of disoriented objects. *Psychological Review*, 95, 115-123.
- Desimone, R., Albright, T.D., Gross, C.D., & Bruce, C. (1984). Stimulus-selective responses of inferior temporal neurons in the macaque. *Journal of Neuroscience*, 4, 2051-2062.
- Desimone, R., Schein, S.J., Moran, J., & Underleider, L.G. (1985). Contour, color and shape analysis beyond the striate cortex. *Vision Research*, 25, 441-452.
- Hoffman, D.D., & Richards, W.A. (1985). Parts of recognition. In S. Pinker (Ed.), *Visual cognition*. Cambridge, MA: MIT Press.
- Jolicoeur, P. (1985). The time to name disoriented natural objects. *Memory & Cognition*, 13, 289-303.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Rock, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
- Rock, I., & DiVita, J. (1987). A case of viewer-centered object perception. *Cognitive Psychology*, 19, 280-293.
- Rock, I., DiVita, J., & Barbeito, R. (1981). The effect on form perception of change of orientation in the third dimension. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 719-732.
- Rock, I., Wheeler, D., & Tudor, L. (1989). Can we imagine how objects look from other viewpoints? *Cognitive Psychology*, 21, 185-210.
- Rolls, E.T. (1984). Neurons in the temporal lobe and amygdala of the monkey with responses selective for faces. *Human Neurobiology*, 3, 209-222.
- Shepard, R.N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701-703.
- Tarr, M.J., & Pinker, S. (1990). Human object recognition uses a viewer-centered reference frame. *Psychological Science*, 1, 253-256.
- Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193-254.
- Warrington, E.K. (1982). Neuropsychological studies of object recognition. *Philosophical Transactions of the Royal Society (London)*, B298, 15-33.
- Weiskrantz, L., & Saunders, R.C. (1984). Impairments of visual object transforms in monkeys. *Brain*, 107, 1033-1072.