

The perception of doubly curved surfaces from intersecting contours

RICHARD B. IVRY and ASHER COHEN
University of Oregon, Eugene, Oregon

The perception of a three-dimensional (3-D) surface shape can be inferred from a 2-D image of two intersecting curves. Three experiments are reported in which we examined a possible method for determining the surface shape as a function of the geometry at the point of intersection. The method involves a two-step process in which the tangents to the two curves determine a skewed Cartesian coordinate system. The angle of the quadrant containing two arcs, the double arc quadrant (DAQ), is then examined. Experiment 1 showed that the surface is perceived as hyperbolic when the DAQ is acute and as locally convex when the DAQ is obtuse. Experiments 2 and 3 showed that even when the DAQ is 90° , the underlying 3-D shape may be unambiguously judged as either hyperbolic or locally convex. It is suggested that the viewer may use an extrapolation process in order to differentiate between these potentially ambiguous stimulus configurations.

One of the many sources of information from which the visual system can determine the shape of a surface is provided by surface contours (see Marr, 1982). A surface contour is the image of a physical curve across a surface. Figure 1 is an example of how a picture composed of surface contours is perceived as the image of a three-dimensional (3-D) surface. However, this drawing could be perceived as the projection of a 2-D object. In fact, there are an infinite number of objects that could be represented by this image. That one 3-D interpretation consistently dominates demonstrates that certain constraints are incorporated within the visual system. Certain geometric assumptions must be made in order to proceed from the curves projected on the 2-D image plane to a description of the distal surface shape.

The surface seen in Figure 1 is a cylindrical surface.¹ The intersection configuration in Figure 2a also seems to define a local patch of a cylinder. The surface patches in Figures 2b and 2c seem doubly rather than singly curved—Figure 2b suggests an ellipsoid and Figure 2c a saddle surface. Why does Figure 2a suggest a singly curved surface, in contrast to Figures 2b and 2c, and how do we differentiate between the latter two, radically different, surface shapes simply on the basis of the geometry of the intersecting curves?

Stevens (1981, 1983, 1986) has developed a theory of the underlying constraints on this interpretation problem. To infer the geometry of the surface, three critical assumptions must be made. First, the surface contours are assumed to be seen from a representative (generic) viewpoint. The general position of this viewpoint allows one to infer that straightness, collinearity, and parallelism in




Figure 1. One 3-D interpretation is consistently seen when this composition of curves is viewed.

the image curves hold for their corresponding physical curves. Second, the surface geometry in the vicinity of each physical contour is assumed to be similar, so that if shifted slightly its image projection would appear substantially the same. These two assumptions imply that no discontinuities, such as unforeseen intersections, troughs, or ridges, would stand out from one vantage point but not from another.

The third and strongest assumption of the theory is that each surface contour is a line of curvature across the surface; that is, the curve follows the principal directions on the surface. At any point on an arbitrary, smooth surface there are two principal directions. One of the principal directions defines the path of greatest curvature, while the other defines the path of least curvature. Moreover,




Figure 2. In (a), the intersection suggests a cylindrical surface patch. In (b) and (c), the two curves intersect in a manner that suggests a doubly curved surface.

This work was supported by Office of Naval Research Contract N00014-84-K-0533. The authors wish to express their appreciation to Kent Stevens for his many comments and criticisms. Jacob Beck, Fred Attneave, Norma Graham, and Richard Koch also provided valuable comments on this manuscript. The authors' address is: Department of Psychology, University of Oregon, Eugene, OR 97403.

the two paths are perpendicular to each other at each point on the surface. The principal lines of curvature correspond to these paths. For example, on the cylinder in Figure 2a, one of the principal curvatures is zero—the surface is not curved along one axis of the cylinder. Perpendicular to this is the other principal line in which the curvature of the cylinder is greatest. Any two intersecting image curves assumed to correspond to lines of curvature therefore meet at a right angle on the surface.

Given these three assumptions, what kinds of descriptions can be extracted from the image of surface contours? First, we are now in a position to make the simple inference that the image of the contour will be curved if and only if the surface is curved. It is possible that a contour on a curved surface will appear as a straight line from a specific point of view. But, as soon as the object is viewed from another point, the contour will no longer be straight. Thus, the assumption of general position is violated in these special cases.

In addition to this basic inference, a qualitative distinction among the three distinct types of surfaces depicted in Figure 2 can now be made. If either of the principal lines of curvature is a straight line, then the perception is of a cylinder. Figure 2a demonstrates an example of this type of surface. If neither of the principal lines of curvature is a straight line, then the surface is doubly curved. However, a further distinction can be made which divides doubly curved surfaces into two classes: surfaces that are locally convex, such as a sphere, and surfaces that are hyperbolic, such as a saddle. These two types of surfaces can easily be differentiated when the contours on a 3-D object are examined. To do this, we must examine the relationship of the surface with respect to the normal plane of the surface at the point of intersection. The normal plane is the plane that contains all of the tangents to the surface at the point of intersection. If all of the surface lies on one side of the normal plane, then the surface is locally convex.² If the surface includes regions that are on both sides of the normal plane, then the surface is hyperbolic.³

The perceptual problem in which we must interpret the 2-D representation of these surfaces is more difficult. Since we can clearly identify cylindrical surfaces by noting that one of the contours will always be a straight line in the image plane, we will not pursue this problem further. Doubly curved surfaces, however, present an ambiguous situation. Can we differentiate between surfaces that are locally convex and those that are hyperbolic merely from the information provided by the 2-D projection of two surface contours?

It seems clear from Figures 2b and 2c that we can. The contours in Figure 2b are readily perceived as lying on a locally convex object, whereas those in Figure 2c lead to the perception of a hyperbolic surface. This is further demonstrated in Figure 3, in which the apparent surface shape varies from locally convex to hyperbolic as the stimuli are observed from left to right. Given that we are




Figure 3. A progression from locally convex to hyperbolic curvature produced by rotating one curve about its intersection point with a second, fixed curve.

skillful perceivers, the problem then becomes a computational question.

Stevens (1981) has shown that the curvature sign of the two intersecting curves may be compared in terms of the rotation of the tangents as one proceeds away from the intersection. Consider the two arcs that bound an acute angle (e.g., those with arrows in Figures 4a and 4b). If these arcs correspond to lines of curvature on a locally convex surface, the tangents in the image plane rotate in the same direction, either clockwise or counterclockwise. For a hyperbolic surface they rotate in opposite directions. An alternative method, which we have used in the following experiments to describe the local geometry of the stimulus intersection in 2-D, involves a simple construction on the intersecting curves. Construct a skewed Cartesian coordinate frame with four quadrants by extending the tangents at the point of intersection (Figure 4). Each curve is thus decomposed into two arcs that fall into different quadrants. Assuming neither curve is a straight line, one of these quadrants will contain two of the arcs, one from each curve. Two other quadrants will each contain only one arc, and the fourth quadrant will be empty. (Note that the empty quadrant will be diagonally opposite the quadrant containing two arcs.) We will refer to the quadrant containing two arcs as the double-arc quadrant (DAQ). In terms of the 3-D interpretation of a stimulus intersection, if the angle forming the DAQ is obtuse (Figure 4a), the 3-D surface would be convex, and if it is acute, it is hyperbolic (Figure 4b). It should be noted that both Stevens's (1981) rotation method and the DAQ method yield identical results. We introduce the DAQ al-




Figure 4. Two methods that distinguish between locally convex and hyperbolic surfaces. See text for explanation.

ternative for two reasons: (1) We find it to be the pedagogically simpler, and (2) the fact that there are alternative, equally valid geometric descriptions emphasizes the point that the actual computation is undetermined.

Both methods predict that the 2-D intersection should allow the viewer to identify the 3-D surface geometry. An experiment, reported below, tested that prediction. Additional experiments explored the relationships between contour curvature and apparent surface shape.

EXPERIMENT 1

The first experiment was designed to investigate two separate aspects of the theory. First, we sought to empirically demonstrate that the geometry of the intersection in 2-D governs the apparent surface shape in 3-D, according to the rule described above. To test this hypothesis, subjects were presented with stimuli composed of two intersecting contours and asked to judge if these contours appeared to be projected from a hyperbolic, saddle-shaped surface or from a locally convex, roughly spherical surface.

The second aspect of Experiment 1 was more exploratory in form. We expected that the major determiner of the apparent surface geometry was the way the two contours intersected in 2-D, that is, whether they appeared to have the same or opposite signs of curvature in 3-D. We recognized that when the curves are nonorthogonal, the impression of either a locally convex (Figure 3a) or hyperbolic (Figure 3e) surface is clear and unambiguous, but when the angle is near 90° (Figure 3c), the surface is more labile, with either 3-D interpretation achievable. Thus, an important issue was the determination of what governed the 3-D interpretation in these latter cases. In this exploration, we manipulated the relative curvatures and angle of the DAQ in both the ambiguous and unambiguous cases. In pilot studies subjects appeared confident in their qualitative judgments of not only 3-D shape, but also the spatial arrangements of the contours. The impression of a pair of crossing contours lying on an invisible surface is heightened when the curves are viewed in darkened conditions. To determine whether this impression was consistent across observers, we asked the subjects in Experiment 1 not only to report the 3-D shapes of the various stimulus configurations, but also to extend one of the contours beyond its actual termination point. Consistent patterns of extrapolation would reflect more constraints on the 3-D interpretation than that provided by the line of curvature assumption.

Method

Subjects

Seven subjects from the University of Oregon Cognitive Laboratory subject pool participated in the experiment. They were paid \$4 for the 1-h experiment.

Apparatus and Stimuli

All of the stimulus configurations were generated on a Symbolics 3600 Lisp Machine and presented on its high-resolution




Figure 5. Examples of stimuli (a-f) used in Experiment 1. In (g) and (h), the curves are extended to the extrapolation line.

monochrome display. The primary stimuli were chosen to form a 2 (curvature of the base contour) $\times$ 2 (curvature of the intersecting contour) $\times$ 4 (angle of the DAQ) factorial design. Some examples are shown in Figure 5 (a-f). So that the angle of the DAQ could be manipulated, the intersection consisted of a base contour and an intersecting contour. The base contour was always presented such that a line connecting its two endpoints was horizontal. The first two factors involved the degree of curvature of either contour. The base and intersecting contours were either a 90° segment of an ellipse whose radii were in a 3:2 ratio or a 90° segment of a circle. The visual angle from the chord to the highest point of each segment was 0.26°, and the chord connecting the endpoints subtended a visual angle of 1.94° for the ellipse segment and 1.31° for the circle segment. Thus, the segment taken from the circle involved greater curvature than did the ellipse segment. The third factor, the angle of the DAQ, was manipulated by rotating the intersecting contour about the point of intersection. This involved four levels—55°, 125°, and two instances of 90°, with the intersecting contour facing to the left in one instance and to the right in the other. The position of the intersection was always at the midpoint of the two contours; its position on the display screen was randomized. The experimental set also included two stimuli consisting of only a base contour presented without an intersecting contour, for each level of curvature. These stimuli were included to provide a baseline measure of the extrapolation process. Thus, the complete set contained 18 stimuli.

Procedure

The subjects were seated in a dimly lit room 6 ft from the display screen. At the beginning of the session, they were shown two 3-D clay figures. One was saddle-shaped and served as a prototype of a hyperbolic surface; the other was a hemisphere selected to be prototypical of a locally convex surface. Two intersecting contours had been painted on the clay models. The subjects were told to imagine that the contours were fluorescent. The objects were then freely rotated by the experimenter to demonstrate that the projection of the contours varied as a function of viewpoint. The subjects were then told that they would be presented with pairs of contours on the display console. They were first to try to imagine the underlying surface. Once this was accomplished, they were to indicate by pressing one of three buttons on a mouse (an interactive input device) whether the surface they had imaged was most similar to the saddle, the hemisphere, or neither.

After the subjects had responded, a dot appeared on the far right side of the display. The experimenter demonstrated that the x-coordinate of the dot could be changed by movement of the mouse. The y-coordinate was fixed at 32.4° of visual angle below the endpoints of the base contour. The dashed lines in Figure 5 show the relative distance between the contours and the extrapolation line. The subjects were then instructed to imagine how the base contour would continue if it were extended on the hypothetical surface. They were instructed to move the dot to the x-coordinate at which a projec-

tion of the right end of the base contour would cross the fixed y-coordinate. The dot became stationary when the subject pressed any of the buttons at a chosen point. Clicking the right button erased that point and allowed the subject to recommence the extrapolation process. Clicking the left button recorded the extrapolation; 2 sec later, the next stimulus appeared. The subjects were asked not to make shape judgments of the single contour stimuli, but to proceed directly to the extrapolation task.

Each subject performed the appropriate tasks six times on each of the 18 stimulus configurations. The experimental session was divided into two blocks, each consisting of 54 trials.

Results and Discussion

The first aspect of the data we examined concerned the subjects' judgments of surface shape when presented with only two contours. Forty-two responses (7 subjects $\times$ 6 responses) had been accumulated for each stimulus configuration in which the subjects had judged the surface to be hyperbolic, convex, or neither. As can be seen in the left half of Table 1, the basic prediction was convincingly supported by the results. Stimuli in which the DAQ was acute were judged to be more similar to the saddle-shaped clay object in 98.2% of the trials, and when the DAQ was obtuse the convex object was likewise selected on 98.2% of the trials. This outcome held despite variations in the degree of curvature of either the base or the intersecting contour.

Considerably less unanimity was found in the judgments of the stimuli in which the angle of the DAQ was 90°. These results are presented in the right half of Table 1. The stimuli that opened leftward were judged 22.6% of the time as hyperbolic, 48.2% of the time as convex, and ambiguous on 29.2% of the trials. For the stimuli that opened rightward, the percentages of 45.2%, 30.1% and 23.8% were obtained for judgments of hyperbolic, convex, and ambiguous, respectively. The near-perfect consistency that was found with stimuli in which the DAQ was either 55° or 125° was no longer evident across subjects or, in many instances, within an individual subject's six responses on the same stimulus configuration.

Table 1
Percent of Saddle-Shaped (S) and Locally Convex (LC)
Responses in Experiment 1

Intersecting Size		Judgment	DAQ			
			55°		90°	
					Open Left	Open Right
Base 1:1						
1:1	S	100.0	0.0	28.6	50.0	
	LC	0.0	97.6	42.8	28.6	
3:2	S	100.0	0.0	26.2	45.2	
	LC	0.0	100.0	50.0	31.0	
Base 3:2						
1:1	S	95.2	2.4	21.4	50.0	
	LC	2.4	97.6	35.7	19.0	
3:2	S	97.6	0.0	14.3	35.6	
	LC	0.0	97.6	64.3	45.2	

Note—Ambiguous responses are not included.

Table 2
Mean Angular Shift From Baseline Extrapolation
in Extrapolation Phase of Experiment 1

Intersecting Size	DAQ							
	55°				90°			
	Mean	SD	Mean	SD	Open Left	Open Right	Open Left	Open Right
Base 1:1								
1:1	-0.55	6.6	1.10	7.7	0.55	5.5	0.55	7.2
3:2	-0.55	6.6	1.10	7.7	-1.65	7.2	-1.10	7.2
Base 3:2								
1:1	-0.55	7.7	0.55	7.7	0.00	7.7	0.55	8.8
3:2	-1.65	7.7	1.10	7.7	-2.20	7.7	0.55	8.8

Note—Data given in minutes of arc. Negative numbers indicate that the mean shift was to the left of baseline extrapolation; positive numbers indicate that the mean shift was to the right.

Two surprising aspects of the results with the 90° stimuli should be noted. First, the majority of these trials (73.5%) produced responses of one of the two possible surface shapes. Although this outcome may reflect a bias on the subjects' part against choosing the ambiguous option, it also shows that factors other than just the DAQ are involved in the judgments. The second unexpected result in Experiment 1 is that when the intersecting contour opened leftward, the subjects tended to perceive convex surfaces, whereas when the intersecting contour opened rightward, the subjects generally saw hyperbolic surfaces. This unexpected bias might have been related to the fact that the extrapolation task was always performed on the right side of the stimulus figure.

As outlined in the introduction to Experiment 1, the extrapolation task was included as a more direct method of assessing how the various stimulus manipulations might affect perceived surface shape. The extrapolation data were analyzed by using each subject's extension of the single-contour stimuli as a baseline measure. The mean difference between a subject's six extrapolations of a given stimulus and the extrapolation of the matching single base contour was computed for each stimulus. Although this method may place unwarranted emphasis on the single-contour stimuli, it did provide a common reference for making comparisons between the different double-contour stimuli. The mean angular shift and standard deviation in the extrapolation task with respect to the appropriate baseline for each stimulus is given in Table 2. Negative numbers indicate a leftward shift from the baseline point, and positive numbers indicate a rightward shift. Two factors contribute to the finding that the mean angular shifts are quite small, especially in light of the much larger standard deviations. First, some subjects produced almost all negative shifts, whereas other subjects produced the opposite. These two patterns would have canceling effects on a mean score. Second, the extrapolation was restricted to a y-coordinate located only 32.4' of arc below the end of the base contour. Pilot work in which the extrapolation was carried out at a distance of 46.2' of arc had

produced much larger shifts. However, since some of the pilot subjects had complained that their extensions of the contour along the imagined surface had curved in such a direction so as to never reach the designated *y*-coordinate, we selected the shorter gap.

An alternative baseline measure, which is independent of the subject's judgments, is the points at which the target *y*-coordinate is crossed by the two basic contours when a larger segment of their respective figures is displayed, as is shown in Figures 5g and 5h. In fact, it had seemed reasonable to assume that the extrapolations performed on the single contours would yield this outcome. This turned out to be the case for the ellipse, for which both of the baseline measures produced identical results. However, the actual extension of the circle is located 2.75' of arc to the right of the subjects' extrapolation of the baseline circle segment and is displaced in this same direction in comparison to all of the test stimuli. Thus, the subjects were producing extrapolations that increased the curvature of this figure.

Although we were unable to identify any statistical differences in the extrapolation task using the mean scores, we did observe some interesting patterns by using matched comparisons. In one such comparison, we assigned the appropriate sign (+, -, or 0) for each acute-obtuse pair within each of the four base and intersecting contour arrangements. The sign was determined by subtracting the mean shift for the 55° stimulus from the mean shift for the 125° stimulus. For example, if a subject had produced a -1.0' shift for the 55° configuration and a +1.0' shift for the 125° configuration in which the two contours were segments of the 3:2 ellipse, then that pair received a sign of +. Mean shifts of -1' and -3' for a different 55° and 125° pair, respectively, would receive a sign of - since the mean score for the 55° was greater. This method indicated that stimulus configurations in which the DAQ was obtuse produced extrapolations that were more rightward than those on a corresponding stimulus in which the DAQ was acute [$t(6) = 2.38, p < .05$]. Since the extrapolation always involved the extension of the right endpoint of the base contour, this indicates that a contour along a surface that is perceived as convex is seen to be continuing to expand, whereas the same contour on a hyperbolic surface tends to be closing in on itself. In addition, in terms of both the mean scores and the sign test, the data suggest that this tendency becomes more exaggerated with a decrease in the curvature of the intersecting contour.

The overall mean extension of the 90° stimuli yielded a value of -0.34' of arc, which lay in between the 55° and 125° mean extrapolations. Surprisingly, the 90° stimuli that opened leftward and were generally judged convex, tended to be shifted more to the left than were the 90° stimuli that opened to the right and had been judged to be saddle-shaped. This is the opposite of what the extrapolation of the nonambiguous stimuli would have led us to expect. However, a sign test between the leftward- and rightward-facing pairs does not approach

significance [$t(6) = 0$]. In fact, an even number of differences between leftward/rightward pairs were less than and greater than 0. The difference between the mean shift of the two stimulus sets (-0.82' for leftward and 0.14' for rightward) is almost entirely attributable to the performance of 1 subject, who produced large shifts to the left when the intersecting contour faced left.

In summary, the judgment task of Experiment 1 provided strong empirical support for the basic prediction of the constraint theory. When the DAQ was acute, the perceived surface shape was hyperbolic, and when the DAQ was obtuse, the surface was seen as locally convex. The extrapolation task with these unambiguous stimuli showed that the local surface shape tends to influence the way in which the inferred surface is expected to continue. Furthermore, the results showed that when the two contours are perpendicular at the point of intersection, the perceived surface shape is not completely ambiguous. This last point demonstrates that, in addition to the angle of the intersection, there are other properties of intersecting contours that constrain the perceptual process. However, the extrapolation data from the 90° stimulus sets failed to reveal any consistent strategies concerning the extension of more ambiguous stimulus configurations. There are a number of possible explanations for this outcome. The subjects may not have used an extrapolation process as a means of disambiguating these stimuli. On the other hand, our present extrapolation task may have been too restrictive. As mentioned above, there is only a small angular difference between the contour endpoint and the *y*-coordinate at which the extrapolation must occur. In addition, the present task required the subjects to extend only one of four possible endpoints. An extrapolation task that requires the subject to consider more global aspects of the stimulus may prove to be more sensitive. We explored this possibility in Experiments 2 and 3.

EXPERIMENT 2

We have hypothesized that through the process of extrapolation, the visual system may be able to resolve apparent ambiguities when presented with two intersecting contours. In the following two experiments, we employed a different method to investigate the extrapolation process and further extend our hypothesis that surface contours that intersect at 90° do not necessarily lead to an ambiguous perception.

The extrapolation task in Experiment 1 involved only a local extension of one of the contours. Although the data showed that the perceived shape might influence the imagined contour projection, no clear pattern emerged that allowed us to account for the differential perception of the 90° cases. That is, although some of the 90° stimuli were not judged ambiguously, the results of the extrapolation task did not correlate with the judgments. In Experiments 2 and 3, we exploited a different property of intersecting contours, namely that any set of contours that

intersect at one point will always intersect at a second point if each curve is extended until it connects back with itself. As opposed to the explicit intersection that is given in the image plane by the contours themselves, we have labeled this second intersection the *implicit intersection*. It is important to recognize that this implicit intersection is only a construct on the image plane and does not correspond to an actual, unseen point on the 3-D object. Although the explicit intersection does have a physical basis, in that it corresponds to a point on the surface which is shared by both contours, the implicit intersection is artificial in that it exists only in the 2-D projection. Note that if the contours were to continue to wrap around a spherical, locally convex object, there would be a second point of intersection. This second intersection does not necessarily correspond to the implicit intersection in the 2-D projection. In fact, such a correspondence would be quite rare. Moreover, there is no actual second intersection of the lines of principal curvature on a hyperbolic surface, even though an implicit intersection can always be projected from the 2-D image of such surfaces. Nonetheless, as we initially grappled with the extrapolation question, we observed that when the explicit intersection was either less than or greater than 90° , the angle of the DAQ at the projected implicit intersection was almost always in agreement. This is demonstrated in Figure 6. This observation led us to explore the possibility that one method for predicting how the surface shape of two perpendicular contours is perceived would be given by the angle of the DAQ at the implicit intersection. Support for this prediction would buttress the hypothesis that there is some consistent extrapolation involved in the perception of surface contours and that this extrapolation involves more global aspects of the stimulus.

Method

Subjects

Fourteen members of the Psychology Department served as subjects. Since all of the volunteers had participated in a pilot study in which they had been asked to draw projections of intersecting contours in which the DAQ was not 90° , they were familiar with the experimental procedure. All of the subjects performed in both the extrapolation task and the judgment task.

Extrapolation Task

Stimuli. Nine sets of curves (Figure 7), in which the angle of the DAQ at the explicit intersection was 90° , were hand drawn. In order to maintain uniformity, a compass was used to draw the curves, and thus they were all arcs of circles. One contour in each




Figure 6. The intersecting curves are extended (dashed lines) to a second point of intersection. The DAQ method makes the same prediction when applied at either intersection.




Figure 7. The stimuli used in Experiment 2. The letters indicate the degree of curvature (l = least, m = moderate, and g = greatest). Uppercase letters indicate curvature of the long segment and lowercase letters indicate curvature of the short segment.

set was a long arc of little (L), moderate (M), or great (G) curvature. The degree of curvature was manipulated by adjusting the radius of the compass. The second contour was a shorter segment of the three long arcs and is referred to as l, m, or g. The nine stimuli were constructed by pairing each long segment with each short one. Each stimulus set was placed on an individual sheet of 8.5×11 in. paper.

Procedure. The nine sheets of paper were shuffled into two different random orders and stapled together. The subjects were instructed to "continue the curves in the most consistent manner until each curve connects back upon itself." By prohibiting the subjects from drawing asymptotic curves, we hoped to insure that we would obtain a sufficient corpus of implicit intersections. The subjects were told that if an extrapolation extended beyond the edge of the paper, they should just resume drawing it at an imagined return point. The subjects were instructed to work through the sets in the order given and to not look back at a drawing once completed.

Judgment Task

Stimuli. The judgment task involved five sets of intersecting contours. Three of these stimuli were selected after the data from the extrapolation task were analyzed. Stimulus Ll was chosen because the DAQ at the implicit intersection was consistently drawn to be obtuse; Stimulus Lg was included because the DAQ at the implicit

intersection tended to be acute. The third stimulus, which contained an explicit intersection of 90° , was Stimulus Ml, which had yielded ambiguous results at the implicit intersection. The remaining two stimuli included one in which the explicit intersection was acute and one in which it was obtuse. These were also drawn with a compass and were selected by the experimenters as unambiguous examples of either a hyperbolic or saddle surface.

Procedure. The judgment phase was conducted a few days after the subjects had completed the extrapolation task. The clay models of hyperbolic and the convex surfaces (without any contours painted on them) were placed in front of the subjects. The subjects were then presented with the five stimuli, one at a time, and asked to judge if the 2-D stimulus looked more like a drawing of the saddle or more like one of the hemisphere. If the subject was unable to choose, an ambiguous response was recorded. The order of presentation was randomized for each subject. In addition, the direction of the page that contained the stimulus was also randomly varied.

Results and Discussion

Extrapolation Task

The two authors independently judged the angle of the DAQ depicted at the implicit intersection. Note that although we continue to refer to this intersection as implicit, the subjects' drawings had produced an actual intersection. Each judge rated the angle on a 5-point scale in which 1 = *very acute* and 5 = *very obtuse*. Implicit intersections that were within 5° of 90° received a rating of 3. When the two judgments did not agree, the average of the two ratings was recorded, thus creating half-steps on the rating scale. This somewhat subjective scoring procedure was employed because, due to some of the subjects' lack of artistic expertise, the tangents at the point of intersection were not always clearly defined. However, the raters did not have any a priori expectations that certain stimuli would yield consistent extrapolations whereas others would not. Any drawings in which the implicit intersections were located at a point off the data sheets or in which the subjects failed to follow directions were discarded.

The results of this rating procedure are presented in Table 3. It should first be noted that 51 % of the implicit in-

tersections that could be measured contained an angle of 90° at the DAQ. This is probably due to a flaw in our methodology in this experiment. We chose arcs of circles as a means of creating a balanced set of stimuli (three sizes crossed with three levels of curvature). However, as occurred to us after the experiment, any two circles that intersect at 90° at one point will also be perpendicular to each other at their second point of intersection. Many of the drawings verified this theorem. In addition, our conservative range for assigning a rating of 3 also increased this frequency.

More interesting, however, is the finding that, for a number of stimuli, the implicit intersection was consistently shifted away from 90° . If we consider only the extrapolations that received a rating of other than 3, five of the nine stimuli produced implicit intersections in which at least 75 % of the implicit intersections shifted in either one direction or the other. (The number of minimum data points required was set at five non- 90° observations.) Stimuli Ll, Gl, Lm, and Mg all produced implicit intersections in which the angle of the DAQ tended to be obtuse. Stimulus Lg was extrapolated so that the DAQ at the implicit intersection was acute. This consistency surfaced despite the difficulty some of the subjects experienced in drawing smooth projections. Note that any noise introduced by the motor system should work against an outcome of systematic, unidirectional shifts from 90° .

Judgment Task

The set of stimuli for the judgment task comprised five stimuli—two in which the explicit intersection was predicted to lead to an unambiguous perception and three in which the contours were perpendicular. However, the latter stimuli had yielded three different outcomes in the extrapolation task. The DAQ at the implicit intersection was consistently obtuse for Stimulus Ll, acute for Stimulus Lg, and ambiguous for Stimulus Ml. Thus, the predicted judgments for these three stimuli were convex, saddle, and ambiguous, respectively. In addition, the constraint theory clearly predicts that the stimulus in which the DAQ at the explicit intersection is obtuse should be matched with the convex hemisphere, whereas the remaining stimulus should be paired with the hyperbolic saddle.

As shown in Table 3, all of these predictions were supported by the data. The stimuli in which the DAQ at either the explicit or the implicit intersection was obtuse were judged to be contours on a convex surface. The two stimuli in which the angle was acute were judged to be projections from a hyperbolic surface. The fifth set of contours in which the DAQ at the explicit intersection was 90° and the implicit intersection tended to be similarly perpendicular was matched with the saddle 36 % of the time, the hemisphere 36 % of the time, and judged ambiguous by 28 % of the subjects.

The results clearly show that contours that contain an unambiguous intersection either at the explicit or at the implicit intersection are consistently matched with the object predicted by the constraint theory. Surprisingly, the results were equally convincing when the DAQ measure-

Table 3
Angle of Implicit Intersection and Judgment
of Surface Shape in Experiment 2

Stimulus	Rating		Percent Judgment		
	Mean	SD	Saddle	Ambiguous	Locally Convex
Ll	3.83	0.56	14.3	0.0	85.7
Lm	3.50	0.50			
Lg	2.61	0.59	64.3	14.3	21.4
Ml	2.89	0.52	35.7	28.6	35.7
Mm	2.82	0.70			
Mg	3.31	0.66			
Gl	3.36	0.46			
Gm	3.15	0.38			
Gg	2.92	0.40			
Control with explicit intersection greater than 90°			14.3	0.0	85.7
Control with explicit intersection less than 90°			78.6	0.0	21.4

Note—Uppercase letter refers to large segment having least (L), moderate (M), or greatest (G) curvature; lowercase letter refers to shorter segment with similar notation. See text for explanation of rating procedure.

ment had to be taken at the implicit intersection as when it could be unambiguously determined at the explicit intersection. This does not mean that we believe the subjects were actually calculating the implicit intersection. As stated earlier, the implicit intersection does not correspond to an actual point on the 3-D surface, and thus seems an unlikely goal of the computational process. However, the high correlation between the present extrapolation process and the subjects' judgments leads us to believe that the process of projecting the contours can be a useful heuristic in capturing the essence of how the surface is perceived. This implies that there must be some explicit cues given in the image that can lead to an unambiguous interpretation of the surface beyond those discussed in the basic constraint theory. Although we have not yet been able to pinpoint these cues, our (accidental) use of arcs of circles provides one hint. As noted in the Method section of this experiment, an accurate extension of the arcs should have produced an implicit intersection of 90° . However, the fact that the implicit intersections consistently differentiated between the two types of doubly-curved surfaces for some of the stimuli leads us to suspect that the extrapolation cues involve some kind of relationship between the two contours rather than properties of each of the curves by themselves.

EXPERIMENT 3

Experiment 3 was undertaken primarily to replicate the results of Experiment 2. In addition, we were interested in examining whether the correlation between the extrapolation and judgment tasks would persist with contour sets that were not composed of arcs of circles and in which the degree of curvature varied within the single curves. The latter manipulation seemed especially relevant, since a surface contour on a doubly curved surface might undergo radical changes in curvature in the image plane due to the orientation of the surface with respect to the viewer.

Method

Subjects

Thirteen members of the Psychology Department served as subjects. None of these volunteers had participated in our pilot studies or in Experiment 1 or Experiment 2.

Stimuli

Eight sets of contours (Figure 8) were constructed in which the DAQ at the explicit intersection was 90° . The sets were prepared by randomly pairing 16 curves that had been traced from either a French curve or the base of curved household objects. Note that, unlike that used in Experiment 2, this method did not allow us to create any sort of factorial design. However, we were primarily interested in exploring the generality of the results of the previous experiment, and therefore chose to employ a wide variety of curves. As before, each stimulus set was drawn on an individual sheet of 8.5×11 in. paper.

Following an analysis of the extrapolation data, three of the eight stimuli were selected for the judgment task. These included Stimulus 2, which was consistently drawn so as to have an acute implicit intersection, Stimulus 8, which yielded an obtuse implicit inter-




Figure 8. The stimuli used in Experiment 3.

section, and Stimulus 5, for which no clear pattern emerged at the DAQ of the implicit intersection. The two stimuli from Experiment 2 that were unambiguous at the explicit intersection were again included for control purposes in the judgment task.

Procedure

The procedure was identical to that of Experiment 2. All of the subjects first participated in the extrapolation task. Following completion of this part of the experiment and subsequent data analysis, the judgment task was conducted.

Results and Discussion

Extrapolation Task

The two authors employed the same procedure to record the angle of the DAQ at the extrapolated implicit intersections. The results of the ratings are presented in Table 4. The percentage of implicit intersections that were rated 3 (90°) was reduced to 28.8% as a result of using curves that were not arcs of circles. Thus, the number of instances in which the implicit intersection deviated from 90° was greatly increased. This allowed us to examine a larger data set to determine if the present method would lead to a consistent surface extrapolation across subjects. In fact, all of the stimuli produced at least six extrapolations that received a rating other than 3.

The results provide support for our basic hypothesis. Four of the eight stimuli led to implicit intersections in

Table 4
Angle of Implicit Intersection and Judgment
of Surface Shape in Experiment 3

Stimulus	Rating		Percent Judgment		
	Mean	SD	Saddle	Ambiguous	Locally Convex
1	3.32	0.75			
2	2.18	0.72	61.4	15.4	23.1
3	3.41	1.22			
4	3.68	1.01			
5	3.17	0.72	23.1	23.1	52.8
6	2.64	0.55			
7	3.29	0.99			
8	4.18	0.75	0.0	7.7	92.3
Control with explicit intersection greater than 90°			0.0	0.0	100.0
Control with explicit intersection less than 90°			92.3	0.0	7.7

Note—This experiment did not include an ordered set of stimuli. See text for explanation of rating procedure.

which at least 75% of the implicit intersections were shifted unidirectionally. Two of the stimuli (Nos. 4 and 8) produced implicit intersections in which the DAQ tended to be obtuse, and two of the stimuli (Nos. 2 and 6) led to implicit intersections that were generally acute. The results for the remaining four stimuli were less uniform.

Judgment Task

As mentioned in the Method section, Stimuli 2, 8, and 5 were selected for the judgment task. On the basis of the extrapolation data, we predicted that they would be judged as hyperbolic, convex, and ambiguous, respectively. It should be noted that the extrapolation task did not produce a clear candidate for the ambiguous prediction. Stimulus 5 was chosen over the other three possibilities because its range of ratings were most closely distributed about the midpoint of our scale. As in Experiment 2, two stimuli in which the contours were not perpendicular at the explicit intersection were included to provide baseline measures.

The results of the judgment task are shown in Table 4. The last two rows confirm the prediction of the constraint theory when the explicit intersection is either less than or greater than 90°. Furthermore, as was observed in Experiment 2, determining the DAQ at the implicit intersections appears to correspond well with subjects' shape judgments of contours that intersect at 90°. This outcome was most convincing for Stimulus 8, which was judged to be a projection of a convex surface by 12 of the 13 subjects. The results were less definitive for Stimulus 2 in that only 62% of the subjects selected the predicted saddle-shaped figure while the remaining 38% were approximately evenly divided between the hemisphere and the ambiguous response. In addition, Stimulus 5, which had produced the least consistency at the implicit intersection, was judged by a majority of the subjects as convex.

GENERAL DISCUSSION

In this paper, we have proposed and tested a method that may differentiate between the two classes of doubly curved surfaces. Specifically, when the DAQ is acute, the surface is perceived as hyperbolic, and when the DAQ is obtuse, the surface is judged to be locally convex. This prediction was strongly supported by the results of Experiment 1. Experiments 2 and 3 showed that even when the explicit intersection is ambiguous (i.e., the contours are perpendicular), the perceived surface shape will still be unambiguous when the DAQ at the implicit intersection consistently deviates from 90°. In such situations, the DAQ method is now applied at the implicit intersection.

It should be pointed out that in all three experiments, subjects were forced to choose between two possible 3-D interpretations. One may argue that the subjects' responses reflect a judgmental rather than perceptual process. That is, instead of relying on their perceptual systems, the subjects may have used some knowledge of geometrical properties in order to make their decisions. However, two points strongly appear to argue against this possibility. First, most subjects do not possess such geometric knowledge. Second, subjects were always given the option of indicating that the stimulus produced an ambiguous percept in terms of its underlying 3-D shape.⁴

These experiments are consistent with the theory developed by Stevens (1981, 1983, 1986) concerning the perception of intersecting contours. The results show that the interpretation of 3-D surface shape is constrained, since the visual system makes certain assumptions about the surface contours. Most relevant to the present study is the assumption that these contours are interpreted as lines of principal curvature. Once this assumption has been made, the two classes of doubly curved surfaces can be differentiated by the 2-D geometry at the point of intersection. Support for this assumption was found in all three experiments when the intersection offered an unambiguous interpretation. It should be noted that at least one alternative explanation cannot be directly ruled out by these experiments. This is that the visual system develops algorithms that are specifically designed to operate at the 2-D level without constraints based on the underlying 3-D structure. However, as Marr (1982) has argued, it seems most reasonable to employ a working hypothesis in which the 2-D representations are processed with reference to the actual world.

We recognize that it cannot presently be stated what actual processes are involved in the perception of intersecting contours. Two computational candidates have been proposed for inferring the surface shape when the explicit intersection is either acute or obtuse. Stevens (1981) proposed that the curvature sign can be determined by comparing the rotation of the tangents as one proceeds away from the intersection. We have employed a different geometric computation based on the construction of the DAQ. In addition, we have introduced the idea of implicit inter-

sections to extend the scope of either method (i.e., when the contours are perpendicular). Although the rotation and DAQ methods yield equivalent results, it is unclear if either is involved in the actual computation. Furthermore, as stated earlier, we believe the idea of implicit intersections should not be regarded as a hypothetical algorithm. It is possible that the visual system attempts to determine implicit intersections, but it strikes us as an unlikely computation since there is no physical correspondence between a 2-D implicit intersection and the 3-D surface.

Nonetheless, the present results lead us to expect that the actual process used in disambiguating contours that intersect at 90° involves some form of extrapolation. We believe that the extrapolation tasks employed in the present study provide a useful first step. The extrapolation task in Experiment 1 showed that the inferred surface shape constrains the manner in which a single contour is perceived to continue along the surface. Experiments 2 and 3 demonstrated that the way in which the two contours interact as they continue away from the intersection can be predictive of how potentially ambiguous intersections are perceived. The implication of these findings is that the cues that allow some perpendicular contours to be unambiguously perceived are a function of more global aspects of the stimulus configurations than just the local information provided at the point of intersection.

REFERENCES

- HILBERT, D., & COHN-VOSSEN, S. (1952). *Geometry and the imagination*. New York: Chelsea.
- MARR, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- STEVENS, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence*, 17, 47-73.
- STEVENS, K. A. (1983). The line of curvature constraint and the interpretation of 3-D shape from parallel surface contours. *Eighth Annual International Joint Conference on Artificial Intelligence* (pp. 1057-1061). Los Altos, CA: William Kaufmann.
- STEVENS, K. A. (1986). Inferring shape from configurations of contours across surfaces. In A. P. Pentland (Ed.), *From pixels to predicates: Recent advances in computational vision*. Norwood, NJ: Ablex.

NOTES

1. Hilbert and Cohn-Vossen (1952) give an excellent discussion of the geometric concepts we will use.
2. The term "locally convex" refers to the intrinsic geometry of the surface independent of viewpoint. The usual convex/concave distinction depends on the viewpoint. Thus, looking at the inside of a sphere, one sees a locally convex surface.
3. The distinctions among cylindrical, locally convex, and hyperbolic surfaces is naturally captured by the sign of the Gaussian curvature. The Gaussian curvature at a point is the product of the two principal curvatures and is a signed quantity. If either of the principal lines of curvature is a straight line, then the Gaussian curvature is zero since the product will also be zero. The Gaussian curvature is positive for locally convex surfaces such as a sphere since the two principal curvatures have the same sign. By common sign, we mean that the two principals follow the same direction away from the normal plane. The common sign can be either positive or negative (an arbitrary distinction determined relative to the normal plane) since the product of either set will produce a positive sign. Hyperbolic surfaces, such as saddles, have negative Gaussian curvature since the two principals have opposite signs. Specifically, one principal moves in one direction from the normal plane, whereas the other principal moves in the opposite direction.
4. The authors are indebted to an anonymous reviewer for stressing the importance of this point.

(Manuscript received March 24, 1986;
revision accepted for publication November 28, 1986.)