Journal of Experimental Psychology:
Human Perception and Performance
1990, Vol. 16, No. 3, 665–674

0096-1523/90/$00.75

OBSERVATIONS

Visual Perception of Smoothly Curved Surfaces From Double-Projected Contour Patterns

James T. Todd
Brandeis University
and Center for Adaptive Systems, Boston University

Francene D. Reichel
Brandeis University

This research was designed to examine how human observers are able to perceive the 3-dimensional structure of smoothly curved surfaces from projected patterns of surface contours. Displays were generated by using a method of double projection that made it possible to cover a surface with a wide range of contour patterns of varying geometric structure and to eliminate systematic variations of image shading. The compelling impression of 3-dimensional form from these patterns provides strong evidence that the ability of observers to perceptually interpret surface contours is considerably less restrictive than would be reasonable to expect on the basis of existing computational models. Results suggest that the perceptual analysis of surface contours is able to exploit statistical regularities of contour structure over appropriately large regions of visual space.

Of the many different methods for pictorial depiction of smoothly curved surfaces, one of the most perceptually compelling involves regular patterns of surface contours. Consider, for example, the image presented in the left panel of Figure 1. At one level of analysis, this image consists of nothing more than a pattern of wavy lines in the picture plane, but at the level of perceptual experience, it appears as a smoothly curved surface in three-dimensional space. The depiction of curved surfaces using patterns of image contours is especially common in engineering drawing and in optical art. However, compared with other sources of pictorial information, it has received relatively little attention in the study of human vision.

One of the only systematic analyses of how patterns of image contours might be perceptually interpreted by human observers has been proposed by Stevens (1981, 1986). Stevens argued that the optical projections of surface contours can only provide information about surface shape if the physical contours in three-dimensional space have some restricted relation with the surface on which they lie. Within the context of human vision, he proposed more specifically that surface contours are assumed perceptually to be lines of principal curvature (cf. Hoffman & Richards, 1984). According to Stevens, this assumption provides the necessary computational constraint to allow the determination of local surface shape, especially when contours are locally parallel to one another as in Figure 1. If parallel contours are assumed to be lines of curvature, then the depicted surface must be locally cylindrical. Stevens has shown that it is possible in that case to estimate the approximate surface orientation in each local neighborhood (see also Stevens & Brookes, 1987).

We now consider the psychological validity of this hypothesis. To demonstrate that human vision incorporates a line of curvature assumption in the analysis of surface contours, it would be necessary to show that this assumption provides sufficient constraint for the correct perception of smoothly curved surfaces. That is to say, when surface contours do indeed lie along lines of principal curvature, their optical projections should produce a compelling impression of three-dimensional form. Unfortunately, however, Stevens's (1981, 1986) efforts to confirm this prediction have been methodologically ambiguous. Although he has provided numerous examples that contoured images of parabolic surfaces can appear compellingly three-dimensional, these examples have all been confounded by the presence of image shading as an alternative artifactual source of information about the structure of the depicted surfaces.

In considering Stevens's analysis as a possible model of human perception, it is important to keep in mind that a strict line of curvature assumption is an extremely restrictive condition that is frequently violated in natural vision. Thus, to provide a convincing demonstration of the psychological validity of this assumption, it should also be necessary to show that the restriction is consistent with the perceptual limitations of actual human observers. Figure 1 would appear at first blush to provide an immediate countercase. The three-dimensional structure of the depicted surface is perceived quite clearly, even though the contours do not lie along lines of principal curvature. Stevens has attempted to address this

This research was supported in part by the Air Force Office of Scientific Research (AFOSR 89-0016) and by a joint grant from the National Science Foundation, the Office of Naval Research, and the Air Force Office of Scientific Research (BNS-898426). The participation of Francene D. Reichel was also supported under a National Science Foundation Graduate Fellowship

We are most grateful to James Cutting, Joe Lappin, Kent Stevens, and an anonymous reviewer for their comments on an earlier draft of the manuscript.

Correspondence concerning this article should be addressed to James T. Todd, Department of Psychology, Brown 125, Brandeis University, Waltham, Massachusetts 02254-9110.

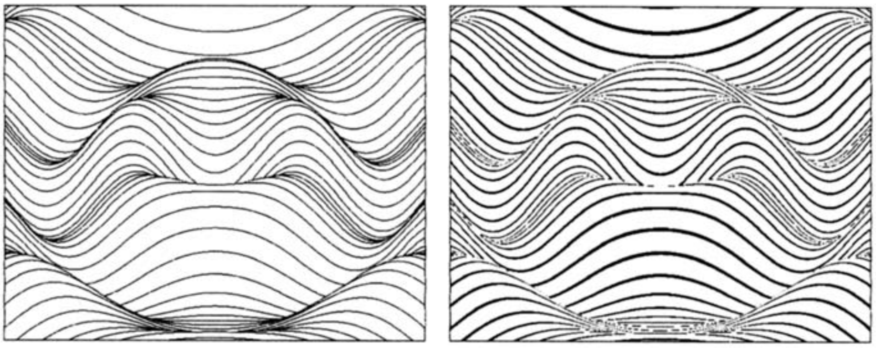

Figure 1. Two visual images depicting patterns of surface contours. (The pattern on the left contains continuous contours of uniform width and is similar to those used in previous investigations by Stevens [1981, 1986]. The pattern on the right contains discontinuous contours of varying width to eliminate any systematic variations of image shading.)

problem by postulating a process of local approximation: "Parallel correspondence is defined only locally, hence is applicable to surfaces that are not cylinders. If the contours are locally parallel, the surface may be approximated locally as a cylinder" (Stevens, 1981, p. 68). Whether this approach will work in practice without introducing large global distortions in an object's computed shape remains to be demonstrated; however, even if it is eventually shown to be computationally feasible, it would still be limited by the highly restrictive condition of contour parallelism. If such an approach is to be taken seriously as a plausible psychological model, it should be necessary to demonstrate that similar restrictions are applicable to actual human perception. Stevens has provided no such demonstrations.

In an effort to perform a more rigorous test of the psychological validity of Stevens's analysis, the research described in this article was designed to investigate images of curved surfaces on which contours could be arranged in a wide variety of geometric relations. To achieve this goal, however, we first had to develop an appropriately flexible method for generating different types of contour patterns and for controlling the presence of image shading as a potentially confounding source of information.

The Method of Double Projection

The usual procedure for creating images such as the one shown in Figure 1 is to adopt a Cartesian coordinate system (X, Y, Z), where X and Y are the horizontal and vertical axes of the picture plane, and Z is the depth axis parallel to the line of sight. The depicted surface is defined initially using a mathematical equation of the form $Z = f(X, Y)$. The individual contours on this surface are defined with the same equation by letting X vary continuously for fixed values of Y. These contours are then rotated in depth about the horizontal axis and projected back into the picture plane with their occluded portions removed. The problem with this procedure for our present purposes is that it is inherently structured to produce locally parallel contours. In order to investigate the psychological importance of this constraint, it is also necessary to generate images in which it is systematically violated.

To help overcome this difficulty, it is particularly useful to reconceptualize the stimulus generation procedure described earlier as a form of double projection (cf. Cutting, 1986, 1987; Gibson, 1966): A pattern P_1 of horizontal lines is projected onto a surface S, which is projected again into the picture plane to produce a new pattern P_2 (see Figure 2). For the purposes of mathematical analysis, the pattern P_1 can be defined as a scalar field in Cartesian two-space:

$$P_1 = [X,\ Y,\ g(X,\ Y)].$$

The surface S can be defined similarly as

$$S = [X,\ Y,\ f(X,\ Y)].$$

To rotate the surface about a horizontal axis through an angle θ with respect to the picture plane, the following transformation is involved:

$$h(X,\ Y,\ \theta) = Y \cos\theta + f(X,\ Y)\sin\theta.$$

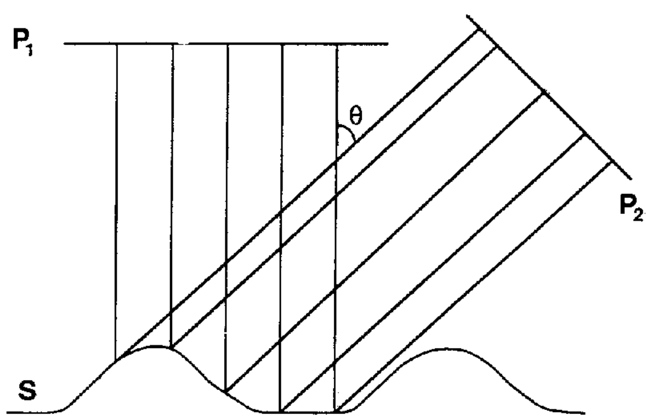

Figure 2. A schematic outline of the method of double projection: A pattern P_1 is projected onto a surface S, which is projected again onto the picture plane to produce a new pattern P_2. (The planes containing P_1 and P_2 are separated by an angle θ.)

These three equations uniquely determine the structure of P_2, so that

$$P_2 = [X, h(X, Y, \theta), g(X, Y)].$$

Thus, to generate a double-projected stimulus pattern, we need only specify the angle θ and the functions $f(X, Y)$ and $g(X, Y)$. For the pattern shown in Figure 1, these were defined as follows:

$$\theta = 40°$$

$$f(X, Y) = 25 \cos(\pi Y/75 + \pi/2) - 75 \cos(\pi \sqrt{X^2 + Y^2}/125)$$

$$g(X, Y) = \begin{cases} 1, & \text{if } \mathrm{mod}(Y, 10) = 0 \\ 0, & \text{if otherwise.} \end{cases}$$

It should also be noted, however, that the same basic procedure can easily be modified by changing $g(X, Y)$ to produce a wide variety of contour patterns for any given surface.

Our discussion thus far has considered the process of double projection as a continuous mapping, but that is not the case for most applications involving the use of computer graphics. If P_1 and P_2 are bit-mapped images composed of discrete picture elements (pixels), then the resulting stimulus pattern can be noticeably distorted because of the inherent rounding error of image rasterization. It is important to keep in mind that a pixel is not a point. Each pixel encompasses a finite area of image structure that can expand or contract as a result of projective transformations (cf. Todd & Mingolla, 1984).

A particularly simple algorithm for double-projecting bit-mapped images is provided in the Appendix as a FORTRAN subroutine. The subroutine receives a pattern array P_1 as input and returns a double-projected image of that pattern in the array P_2. The dimensions of these arrays must be appropriately scaled by the constants W and V, such that W is half the horizontal width of the display screen (in pixels) and V is its vertical height. It is also critical to orient the surface defined by the function h so that it is drawn from back to front. This will guarantee that all surface regions that are hidden by occlusion will be automatically overwritten.

The right panel of Figure 1 shows a double-projected image that was generated using this procedure from a pattern P_1 containing regularly spaced horizontal contours. Note in the figure that double projection causes the contours to expand and contract. Expansion occurs when one pixel in P_1 projects to multiple pixels in P_2. Contraction occurs conversely when multiple pixels in P_1 project to a single pixel in P_2. The color of the P_2 pixel in that case is determined by the last P_1 pixel that is mapped onto it. Thus, as is evident in the figure, small gaps can sometimes occur when a contour is overwritten by part of the background. One important consequence of these many-to-one and one-to-many mappings is that the pattern produced in P_2 can have systematic variations of texture density. It is also important to note, however, that this particular method of double projection does not produce systematic variations of pattern shading. Although the black and white regions of a pattern can expand or contract as they are mapped from P_1 to P_2, their relative proportions in each local neighborhood remain perfectly invariant.

There are several possible variations of this basic procedure of double projection that can easily be adapted to meet the demands of specific experimental applications. For example, it is especially useful in research on stereopsis or structure-from-motion to eliminate all possible information—including texture density cues—from each individual image of a multi-image display (e.g., see Reichel & Todd, 1990; Sperling, Landy, Dosher, & Perkins, 1989). This can be accomplished quite simply by replacing line 60 in the Appendix with

```
P2(X,YP) = mod(P1(X,Y)/(2**(YP - YO)), 2).
```

If this modification is used with a pattern array P_1 containing random integers, then the double-projected image that is output to P_2 will contain a homogeneous pattern of random noise, in which each individual pixel has an independent 50% probability of being black or white. By altering the parameter θ in the function h, it is possible to create multiple images of the same surface with a complete point-to-point correspondence, except for small amounts of noise that are automatically added or subtracted to compensate for the expansion and contraction of texture owing to the effects of double projection. An example stereogram that was generated with this procedure is shown in Figure 3.

It is also possible, if desired, to create patterns of double-projected contours that are smoothly continuous (i.e., without gaps) and to eliminate their variations in width. This requires two modifications to the basic procedure described in the Appendix. The first case involves many-to-one mappings, in which several adjacent pixels in P_1 project to a single pixel in P_2; the P_2 pixel should be colored black if a contour passes through any of its corresponding pixels in P_1 so as to eliminate the presence of gaps in regions that are highly foreshortened. To eliminate variations of contour width, it is necessary to control the effects of one-to-many mappings. Consider, for example, a single straight line segment in P_1 with a minimum possible width of 1 pixel. Figure 4 shows a possible mapping of this line segment, in which each pixel in P_1 projects to multiple pixels in P_2. Note in the figure that the transformed contour forms a bounded area that has large variations in width. It is possible to eliminate these variations as shown in the figure by only coloring those pixels at the upper (or lower) boundary of a contour's projected area. A typical example that was generated using these modifications is shown in the left panel of Figure 1.

When used with patterns of regularly spaced parallel contours in P_1, this latter modification to the method of double projection is formally equivalent to the procedures used for depicting surfaces in engineering drawing and in previous demonstrations of Stevens (1981, 1986). It is important to keep in mind, however, that (a) surface contours of uniform width do not generally occur in natural images, and (b) the imposition of such uniformity in experimental displays inevitably produces systematic variations of image shading that can serve as a potentially confounding source of information in addition to the geometric structure of the contours themselves (e.g., see Figure 1).

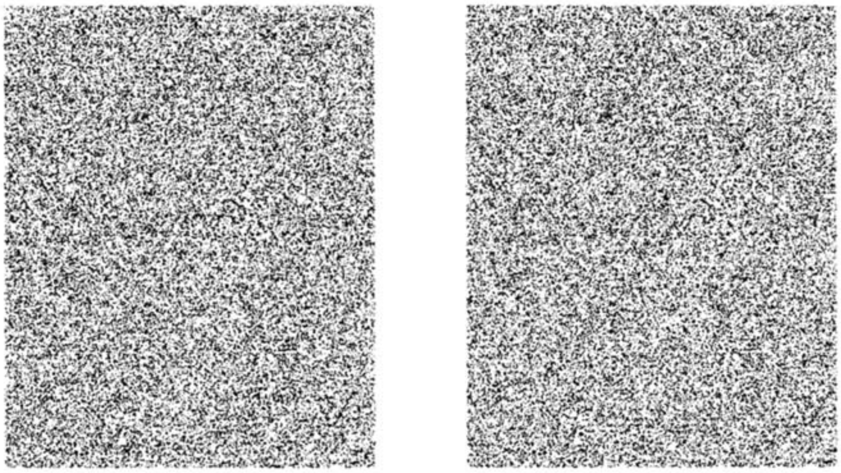

Figure 3. A random-dot stereogram of a smoothly curved surface. (Each image was generated using a modified method of double projection to eliminate all variations of texture density.)

The Visual Perception of Surface Contours

The method of double projection and its various modifications provide especially useful tools for exploring the underlying computational constraints on the perception of shape from surface contours. It is possible, with these techniques, to create a wide variety of contour patterns for any given surface by manipulating the specific geometric properties of P_1. In the discussion that follows, we present a large number of double-projected patterns in an effort to reveal some of the relevant boundary conditions for their perceptual interpretation. To facilitate a direct comparison of these figures, the depicted surfaces will all be identical to the one shown in Figure 1, and they will all be equated in terms of relative contour density. We also examine the effects of shading by generating each pattern both with and without the modified procedure described earlier for eliminating variations of contour width.

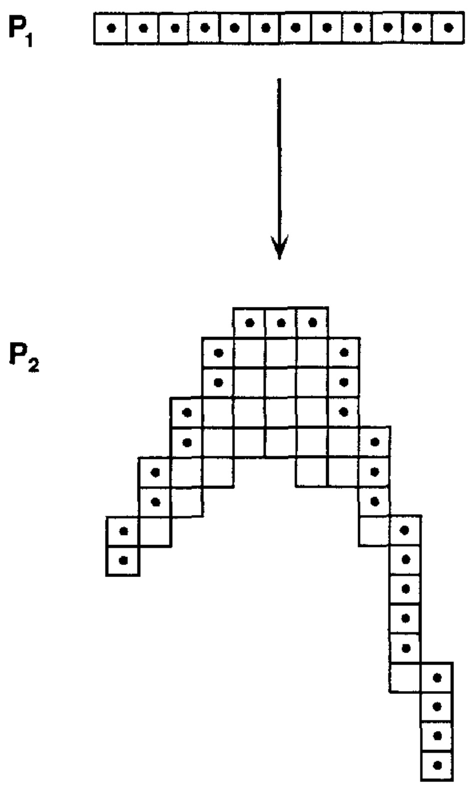

Figure 4. A single straight line segment in P_1 composed of discrete pixels, and its possible mapping in P_2. (In this particular example, each pixel in P_1 projects to multiple pixels in P_2, thus producing a bounded area with large variations in width. It is possible, if desired, to eliminate these variations by only coloring those pixels at the upper or lower boundaries of a contour's projected area. As represented here, the open square regions depict uncolored pixels, whereas regions containing a dot depict colored pixels.)

We begin by considering the perceptual significance of contour parallelism. On the basis of the analysis of Stevens (1981, 1986) described earlier, it would be reasonable to expect that the removal of contour parallelism in an image should have a highly detrimental effect on the perceptual appearance of three-dimensional form; however, our observations do not confirm this prediction. Consider, for example, the surfaces shown in the upper portion of Figure 5. These images were generated (both with and without shading) from a pattern P_1 in which lines were positioned at random with random orientations over a 45° range. Because both of these images appear compellingly three-dimensional, it seems reasonable to conclude that the perception of shape from surface contours need not necessarily depend on implicit assumptions about contour parallelism or the uniformity of contour spacing. There are limits, however, to how much variation of contour orientation can be tolerated. The lower portion of Figure 5 shows a pair of double-projected images for which the lines in P_1 were oriented at random over a full 180° range. Note in this case that the contours do not evoke a strong impression of a smoothly curved surface, either with or without shading. The patterns appear as little more than random configurations of wavy lines in the picture plane.

Another important property of these displays that could potentially influence their perceptual interpretation involves the length of the depicted contours. The upper portion of

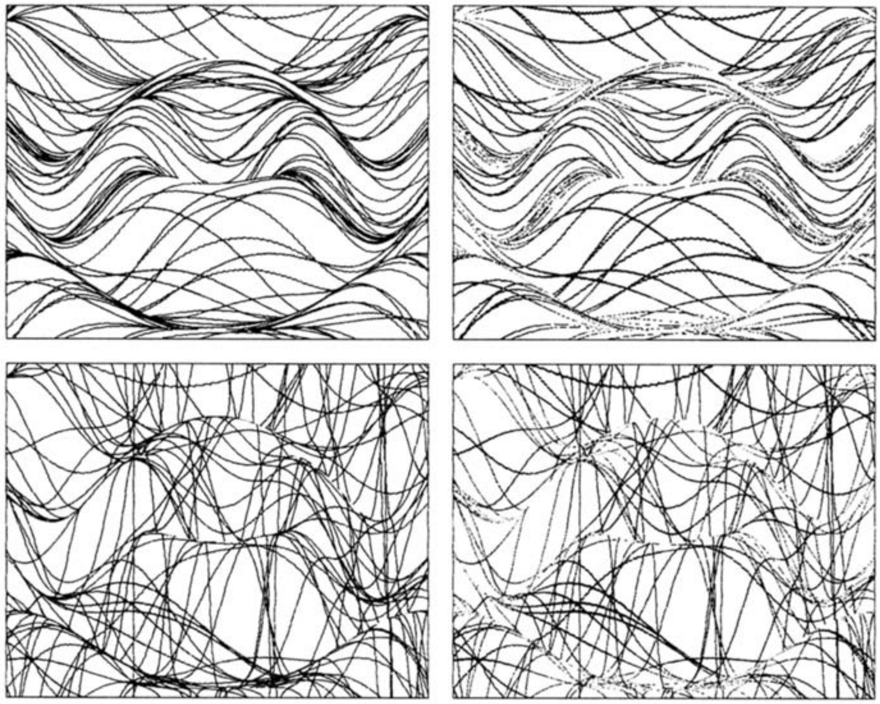

Figure 5. Four double-projected images of a smoothly curved surface for evaluating the perceptual significance of contour parallelism. (In generating each display, the pattern P_1 was composed of randomly positioned lines with random orientations. The range of orientations was 45° for the two upper images and 180° for the two lower images. The patterns on the right are identical to those on the left, except that they were generated with an unmodified double projection to eliminate any systematic variations of image shading.)

Figure 6 shows a pair of images that were generated using randomly positioned horizontal line segments with a length of only 10 pixels (cf. Zucker, 1985). The resulting patterns both with and without shading are immediately identified as a smoothly curved surface even by naive observers. Compare this, however, with the pair of images shown in the lower portion of Figure 6, for which the pattern P_1 was composed of randomly positioned dots. Consider first the right panel of this pair that was generated using an unmodified double projection and that contains no systematic variations of image shading. Although there are noticeable variations of dot density in this figure that could in principle provide useful information, this seems to be of no use whatsoever for the perceptual analysis of three-dimensional form (cf. Cutting & Millard, 1984; Todd & Akerstrom, 1987). It is possible to discern a vague impression of three-dimensional form when shading is added to the display (i.e., in the left panel), but the effect in that case is not overly compelling, and naive observers have a difficult time identifying this figure as a smoothly curved surface.

These observations suggest that systematic variations of contour orientation may be a critical source of information for the perceptual analysis of shape from surface contours. If this information is masked by excessive randomization as in Figure 5, or by excessive reductions in length as in Figure 6, then the perception of three-dimensional form is dramatically impaired.

One other property that is shared in common by all of the images presented thus far is that they have all been generated from patterns in P_1 composed entirely of linear contours. This is potentially of great theoretical importance. If the patterns in P_1 are restricted to linear contours, then any curvature in their double-projected images can only be due to the three-dimensional structure of the depicted surface. Is this a necessary condition for the perception of three-dimensional form in these displays? Our observations suggest it is not. Consider, for example, the two images presented in the lower portion of Figure 7. These images were generated using a pattern P_1 composed of "random-walk" contours, in which each successive step had a horizontal displacement of 1 pixel and a vertical displacement selected at random between ±2 pixels. Note in particular that the three-dimensional form of the surface is clearly visible either with or without shading. The upper panel of Figure 7 shows a similar pair of images that were generated with sinusoidal contours. Here the effects are somewhat more complex. When the display contains systematic variations of image shading, it produces an extremely compelling impression of a smoothly curved surface. When the shading is removed, however, the compellingness of the display is greatly attenuated.

In considering the properties of double-projected images, it is especially important to note that a pattern P_2 can contain two perceptually distinct categories of contours: *surface contours* that are in one-to-one correspondence with the contours

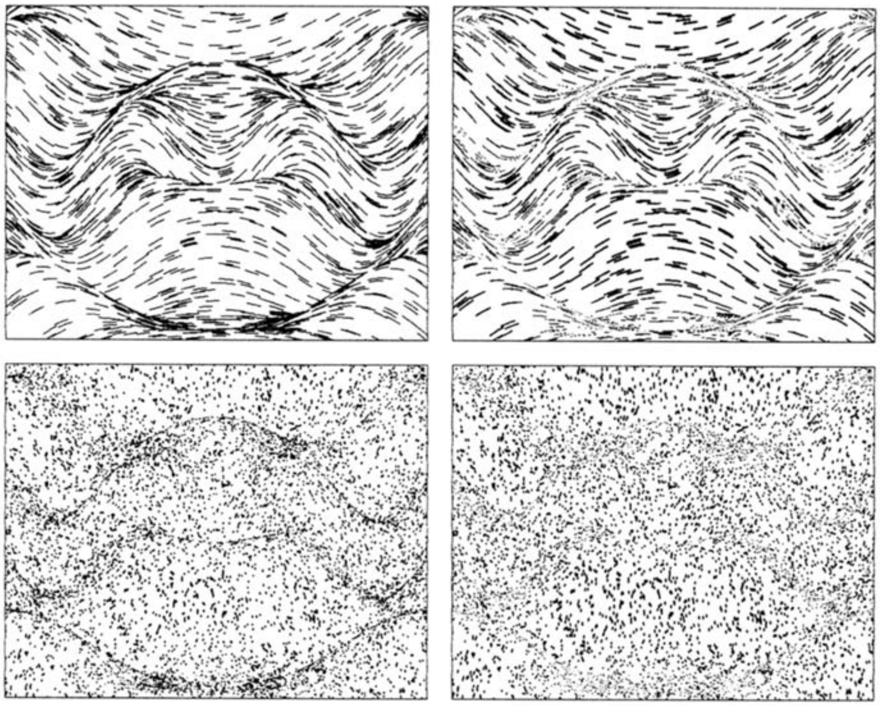

Figure 6. Four double-projected images of a smoothly curved surface for evaluating the perceptual significance of contour length. (The pattern P_1 was composed of randomly positioned horizontal line segments for the two upper images, and randomly positioned dots for the two lower images. The patterns on the right are identical to those on the left, except that they were generated with an unmodified double projection to eliminate any systematic variations of image shading.)

in P_1, and *occlusion contours* that are produced by the process of double projection when one part of a surface is hidden from view by another. One possible strategy for identifying occlusions in all of the images presented thus far is to search for the presence of contour T-junctions. This strategy will not work, however, for double-projected images such as those shown in Figure 8, for which the pattern P_1 was composed of randomly positioned Ts. Note in particular that the occlusion contours in these images are easily identified even though the majority of T-junctions are due to other factors.

It is also important to point out in this regard that the presence of occlusions is not a necessary condition for a pattern of image contours to be perceived as a smoothly curved surface. Consider, for example, the clear appearance of three-dimensional form for the pair of images in the upper portion of Figure 9. These images are identical to those shown in Figure 1, except that the depicted surface has a smaller slant in relation to the image plane so that there are no occlusions. (See also the similar demonstrations of Stevens, 1981, 1986). A pattern of occlusion contours is also not a sufficient condition for the perception of three-dimensional form. This can be demonstrated by examining the lower left panel of Figure 5. Note in particular that there are many contour terminations in this figure to specify the presence of occlusions, yet the pattern does not appear compellingly three-dimensional. Similarly, if the occlusion contours in these figures are presented in isolation (as in the lower left panel of Figure 9), they appear as nothing more than a two-dimensional pattern of curved lines.

Although abrupt discontinuities resulting from occlusion are neither necessary nor sufficient for a pattern of image contours to be perceived as a smoothly curved surface, there is considerable evidence to suggest that the presence or absence of these occlusions can have a substantial impact on observers' perceptions (e.g., see Reichel & Todd, 1990; Todd & Reichel, 1989; see also the related theoretical analyses of Beusmans, Hoffman, & Bennett, 1987; Hoffman & Richards, 1984; Koenderink, 1984; Koenderink & van Doorn, 1976, 1982; Richards, Koenderink, & Hoffman, 1987). This can be observed informally from the displays presented thus far by comparing the amount of perceived surface undulation for the patterns shown in Figure 1 with the comparable patterns without occlusions in the upper portion of Figure 9. An even more dramatic demonstration of the perceptual significance of occlusions can be obtained by examining the lower right panel of Figure 9. This pattern was generated with a procedure identical to the one described for Figure 1, except that the orientations of P_1 and P_2 were reversed (see Figure 2). Although the contours in this figure were projected at the same orientation in relation to the line of sight as in Figure 1, the surface was positioned parallel to the image plane so that there would be no occlusions. Note in particular that the

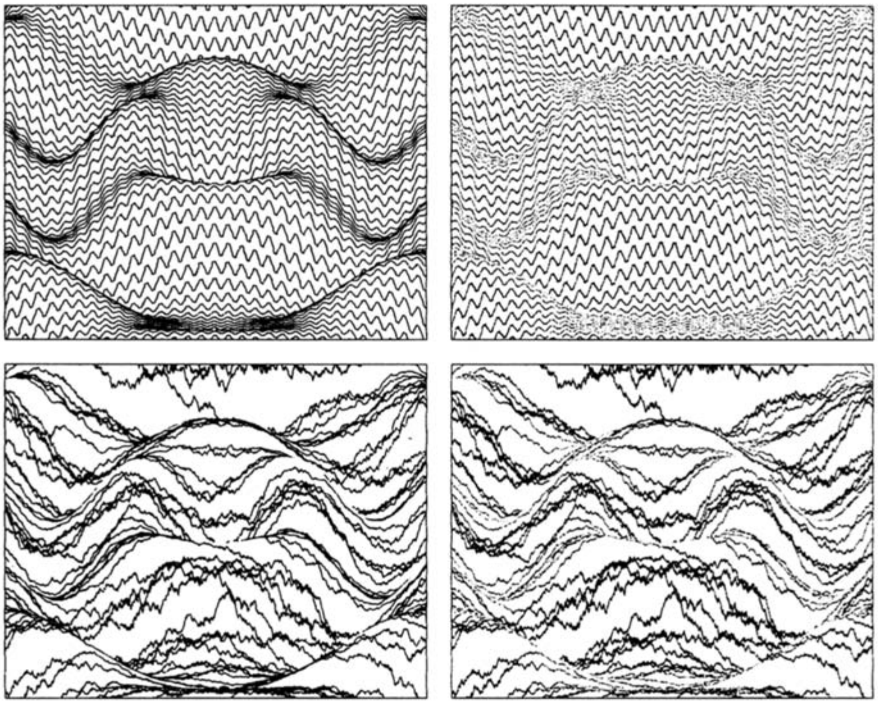

Figure 7. Four double-projected images of a smoothly curved surface for evaluating the perceptual significance of contour curvature. (The pattern P_1 was composed of sinusoidal contours for the two upper images, and random-walk contours for the two lower images. The patterns on the right are identical to those on the left, except that they were generated with an unmodified double projection to eliminate any systematic variations of image shading.)

resulting pattern bears no resemblance to a smoothly curved surface, even though it satisfies the condition of local parallelism required for Stevens's analysis.

Summary and Conclusions

The research described in this article has been designed to examine some of the limiting conditions for the perceptual interpretation of surface contours. Our working hypothesis at the start of this investigation was based on an analysis by Stevens (1981, 1986), in which he argued that the physical contours on a smoothly curved surface must be highly constrained with respect to the surface geometry in order for their optical projections to provide useful information for the visual perception of three-dimensional form. To investigate the psychological validity of this hypothesis, we first had to develop

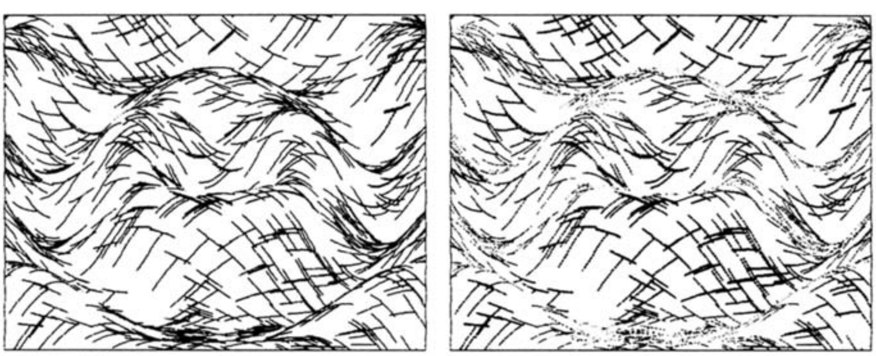

Figure 8. Two double-projected images of a smoothly curved surface, for which the pattern P_1 was composed of randomly positioned Ts. (The pattern on the right is identical to the one on the left, except that it was generated with an unmodified double projection to eliminate any systematic variations of image shading.)

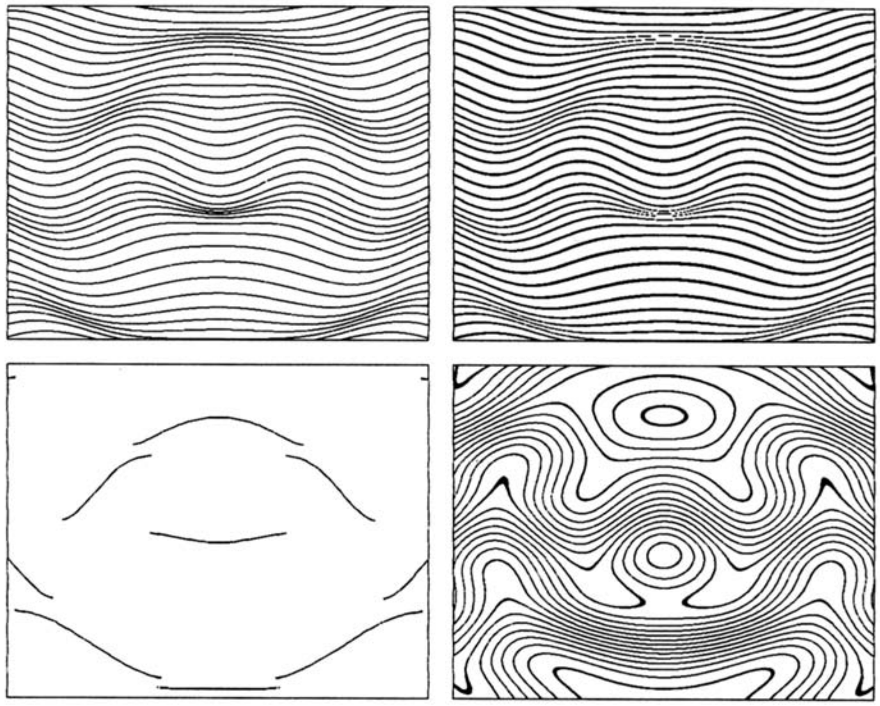

Figure 9. Four double-projected images of a smoothly curved surface for evaluating the perceptual significance of occlusion contours. (The lower left image shows the occlusion contours of Figure 1 presented in isolation. The two upper images show the surface in Figure 1 at a smaller slant, $\theta = 15°$, with no visible occlusions both with and without systematic variations of image shading. The lower right image shows the same surface oriented parallel to the picture plane. This image was generated using the same procedure as described for Figure 1, except that the orientations of P_1 and P_2 were reversed—see Figure 2.)

an appropriate method of stimulus generation that would allow us to cover a surface with a wide range of contour patterns of varying geometric structure. We then set out to investigate the effects of these patterns on actual human perception.

The results of this research have indicated that the ability of observers to perceptually interpret surface contours is surprisingly robust. The visual impression of a smoothly curved surface can be achieved over a wide range of conditions: To generate a perceptually compelling double-projected image, the contours in P_1 can be short or long, straight or curved, and have randomly varying positions and orientations. Perceptually compelling effects can also be obtained for patterns composed of discontinuous contours of varying width, in which there are no systematic variations of image shading. What, then, is the computational basis for the perceptual appearance of three-dimensional form in these figures? Our observations suggest that the perception cannot be due to an assumed restriction on how individual contours relate to the surface geometry. If that were the case, then the perceived three-dimensional shape of a surface would not be expected to remain invariant over large changes in its depicted pattern of contours, as was observed in this investigation.

A more plausible explanation of this finding is that the three-dimensional structure of a depicted surface is determined from the statistical distributions of contour orientation within each local neighborhood. For images of smoothly curved surfaces, the average orientation within each local region forms a two-dimensional vector field that is everywhere continuous except at points of occlusion. If it is assumed that the distribution of contour orientations on the physical surface in three-dimensional space is spatially isotropic, then any variations in their average projected orientations could be used to estimate the approximate surface curvature in each local neighborhood. It would be reasonable to expect, by using this analysis, that performance should deteriorate for any manipulation of the depicted pattern that masks changes in average orientation. One such manipulation would be to reduce the signal-to-noise ratio by increasing the variance of contour orientation as in Figure 5. Another is to reduce the length of the contours in such a way that their relative orientations are no longer detectable as in Figure 6.

It would be possible to compute the vector field of average contour orientations by using a neural network analysis similar to the boundary contour system developed by Grossberg and Mingolla (1985a, 1985b, 1987). The network could be constructed so that neighboring neurons tuned to parallel orientations would excite one another, whereas those tuned to orthogonal orientations would inhibit one another. To estimate surface curvature, it would then be necessary to measure how the average contour orientation varies from one region to the next, perhaps by using a method like the one developed by Zucker (1985).

One important aspect of the Grossberg and Mingolla (1985a, 1985b, 1987) model for computing the field of average contour orientations is that it involves networks of local mechanisms that interact globally. There are several phenomenological aspects of our displays that lead us to suspect that similar global interactions are of fundamental importance in the perceptual analysis of surface contours. Of particular relevance to this issue is the apparent tolerance of observers' perceptions to large random variations of contour orientation or curvature. This finding suggests that the perceptual analysis of these displays may involve a process of spatial integration in order to exploit statistical regularities in an appropriately large population of contours (see Todd & Akerstrom, 1987). It also is interesting to note in this regard that the possible extent of this spatial integration can be artificially limited by viewing the displays through a reduction tube. Indeed, if a small local region is viewed in isolation, it does not appear compellingly three-dimensional (cf. Todd & Reichel, 1989). Still other evidence suggests, moreover, that the necessary integration cannot be achieved through a simple process of low spatial frequency filtering. If the displays are optically blurred, their perceptual compellingness is dramatically reduced, especially those displays in which systematic variations of shading have been eliminated. When considered as a whole, these observations suggest a global integration of high spatial frequency structure similar to what is achieved by the Grossberg and Mingolla model.

Although we are in general agreement with Stevens (1981, 1986) that objects in three-dimensional space must somehow be constrained to reliably interpret their optical projections, we believe it is essential to consider seriously the functional limitations of any proposed constraint before it is accepted as a plausible foundation for human perception. This is, in our opinion, a fundamental issue that existing models have not adequately addressed. Many of the constraints that have been proposed in the literature for the computational analysis of various image properties (e.g., shading, texture, or motion) seem to have been adopted more for their mathematical convenience than for their ecological or psychological validity, and there is a growing body of psychophysical evidence that large systematic violations of these constraints often have little impact on observers' perceptions (see Mingolla & Todd, 1986; Todd, 1984, 1985; Todd & Akerstrom, 1987; Todd & Mingolla, 1983; Todd & Reichel, 1989). Stevens's (1981, 1986) suggestion that surface contours are perceptually assumed to be lines of principal curvature is a typical case in point in that its highly restrictive limitations seem to have little in common with the perceptual capabilities of actual human observers. Our approach, in contrast, is based on a more general assumption that the distributions of contour orientation on a surface are spatially isotropic, so that changes in the average orientation of their optical projections can be used to estimate the approximate surface curvature in each local neighborhood. Although this constraint is less computationally precise than Stevens's line of curvature assumption, we believe that it is much more consistent with the properties of human perception.

References

Beusmans, J. M. H., Hoffman, D. D., & Bennett, B. M. (1987). Description of solid shape and its inference from occluding contours. *Journal of the Optical Society of America A, 4,* 1155–1167.

Cutting, J. E. (1986). The shape and psychophysics of cinematic space. *Behavior Research Methods, Instruments, and Computers, 18,* 551–558.

Cutting, J. E. (1987). Rigidity in cinema seen from the front row, side aisle. *Journal of Experimental Psychology: Human Perception and Performance, 13,* 323–334.

Cutting, J. E., & Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General, 113,* 198–216.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Grossberg, S., & Mingolla, E. (1985a). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review, 92,* 173–211.

Grossberg, S., & Mingolla, E. (1985b). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics, 38,* 141–171.

Grossberg, S., & Mingolla, E. (1987). Neural dynamics of surface perception: Boundary webs, illuminants and shape from shading. *Computer Vision, Graphics, and Image Processing, 37,* 116–165.

Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18,* 65–96.

Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception, 13,* 321–330.

Koenderink, J. J., & van Doorn, A. J. (1976). The singularities of the visual mapping. *Biological Cybernetics, 24,* 51–59.

Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception, 11,* 129–137.

Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics, 53,* 137–151.

Reichel, F. D., & Todd, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 653–664.

Richards, W. A., Koenderink, J. J., & Hoffman, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A, 4,* 1168–1175.

Sperling, G., Landy, M. S., Dosher, B. H., & Perkins, M. E. (1989). Kinetic depth effect and identification of shape. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 826–840.

Stevens, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence, 17,* 47–74.

Stevens, K. A. (1986). Inferring shape from contours across surfaces. In A. P. Pentland (Ed.), *From pixels to predicates* (pp. 93–110). Norwood, NJ: Ablex.

Stevens, K. A., & Brookes, A. (1987). Probing depth in monocular images. *Biological Cybernetics, 56,* 355–366.

Todd, J. T. (1984). The perception of three-dimensional structure from rigid and nonrigid motion. *Perception & Psychophysics, 36,* 97–103.

Todd, J. T. (1985). The perception of structure from motion: Is projective correspondence of moving elements a necessary condition? *Journal of Experimental Psychology: Human Perception and Performance, 11,* 689–710.

Todd, J. T., & Akerstrom, R. A. (1987). Perception of three-dimensional form patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance, 13,* 242–255.

Todd, J. T., & Mingolla, E. (1983). The perception of surface curvature and direction of illumination form from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 583–595.

Todd, J. T., & Mingolla, E. (1984). Simulations of curved surfaces from patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 734–739.

Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review, 96,* 643–657.

Zucker, S. W. (1985). Tangent fields and the dimensionality of their support. *Computer Vision, Graphics, and Image Processing, 32,* 74–103.

Appendix

```
Subroutine Double-Project (P1, P2, θ)
Integer P1 (−W:W,−V:V), P2 (−W:W,−V:V), X,Y,YO,YN,YP
 10    Do 100 X = −W, W
 20          Do 90 Y = −V, V
 30            YN = h(X,Y,θ)
 40            If (Y.EQ.−V) YO = YN
 50            Do 70 YP = YO, YN
 60              P2(X,YP) = P1(X,Y)
 70              Continue
 80            YO = YN + 1
 90            Continue
100          Continue
       END
```

Received March 3, 1989
Revision received August 12, 1989
Accepted August 24, 1989 ■