

THE LINE OF CURVATURE CONSTRAINT AND THE INTERPRETATION OF 3-D SHAPE FROM PARALLEL SURFACE CONTOURS

Kent A. Stevens

Department of Computer and Information Science
University of Oregon
Eugene, Oregon 97403

ABSTRACT

Human vision is adept at inferring surface shape from a projection of curves lying across a surface, particularly from parallel undulating curves. Since the image projection loses 3-D information, the interpretation must be constrained by certain a priori assumptions. A theory of constraint on this problem [Stevens 1981a] proposes three assumptions, namely that neither the viewpoint nor the placement of the physical curves on the surface is misleading (two forms of general position), and that the physical curves are lines of curvature across the surface. It follows that parallel contours in an image correspond to parallel curves across an approximately cylindrical surface, and since the lines of curvature on a cylinder are geodesic and planar, one may reconstruct the image projection of an orthogonal net covering the surface. There is sufficient information available to estimate surface orientation at certain strongly constrained points, then to determine, given the geodesic and planarity properties, the surface orientation elsewhere by a simple extrapolation method. An implementation of this computation predicts local surface orientation that closely matches the orientation we perceive, thereby supporting the theory that our visual interpretation of surface shape from parallel contours embodies the line of curvature constraint.

I INTRODUCTION

The parallel curves in figure 1 are interpreted in 3-D as lying across an undulating




Figure 1.

* This work, conducted at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology, was supported in part by the Advanced Research Projects Agency under ONR contract N00014-80-C-0505 and in part by the AFOSR and NSF grant 79-23110MCS.

surface. (Ideally one should examine these figures monocularly, perpendicular to the line of sight, with an occluding mask that reduces the evidence that they are merely drawn on the printed page, and for several seconds to allow the 3-D shape to develop fully.) Our perception of surface shape in this figure is so immediate and effortless that we must be reminded that the image could correspond to any of an infinite set of possible configurations, including the literally correct 2-D interpretation as lines of ink on the flat printed page (which is surprisingly difficult for us to hold). The fact that we choose a definite 3-D interpretation (subjects make precise judgments of local surface orientation and different observers make similar judgments) suggests that the visual system strongly constrains the 3-D interpretation.

Observe that the undulating curves in figure 1 appear to be lines of curvature on a cylinder (a developable surface having parallel rulings -- see, e.g. [Hilbert & Cohn-Vossen 1952]). That is, the curves appear geodesic, planar, and oriented perpendicular to the ridges and troughs of the surface. The same interpretation can be had with even simpler configurations, such as in figure 2, where the curved contour (line of curvature) appears to intersect the straight line (ruling) at a right angle. Note again that the curved contour appears geodesic and planar, regardless whether seen as a physical edge (e.g. of a sheet of paper) or as a line across the surface. It turns out that considerable constraint on the 3-D interpretation




Figure 2.

is afforded by assuming that the curves are lines of curvature seen in general position, particularly for the case of parallel curves in an image.

A theory of the underlying constraints is given in [Stevens 1981a], which can account for these qualitative observations. We will review the theory and report quantitative support provided by an implementation that embodies these constraints. The implementation predicts 3-D shape that corresponds rather closely to the subjective impression of shape seen in these images (the detailed psychophysical correspondence is currently being examined). While the implementation is not meant as a perceptual model but only a means for examining the 3-D shape predicted by the constraints, it does have relevance to computer vision as a method for estimating local surface orientation from certain image configurations.

Prior to discussing the theory and examining the implementation it will be useful to comment on certain points that one might raise on first examination of these figures.

1. A shading effect is apparent in figure 1 (especially noticeable when blurred). Shading likely enhances the apparent surface impression, but it cannot be the sole cause of the three-dimensionality, since we perceive 3-D shape where shading is negligible (see figures 2, 4, and the surfaces suggested by dotted lines and subjective contours in [Stevens 1982, figure 2]).

2. When these figures reverse in depth, the perceived surface may change its appearance radically. For example, figure 1 rotated 90 deg will reverse readily, and with each reversal the troughs become ridges and vice versa (rotating the figure defeats our tendency to assume that distance increases towards the top of an image). The global changes in surface shape may be attributed, however, to local ambiguity in the direction of tilt: the apparent tilt reverses by 180 deg while the slant remains constant. (Slant is the angle, 0-90 deg, between the line of sight and the surface normal; tilt is the direction, 0-360 deg, of the gradient of distance from the viewer to the surface, and also the direction of the projected surface normal [Stevens 1979, 1983a; Marr 1982].) Tilt direction cannot be expected to be determined in these orthographic projections. It is presumably disambiguated by subsequent visual processes, as part of the problem of deriving a consistent depth interpretation from local surface orientation information. It is well known that image orientation is important to this process, as rotating figure 1 demonstrates.

3. The surface contours discussed here are distinct from contours of foreshortened surface texture (e.g. [Kender 1980; Witkin 1981]). Consider, e.g. the mottled pattern of sunlight and shadows cast on the ground below a tree, which project to foreshortened texture contours in an image. Provided the surface texture is isotropic, the image texture projected from a slanted surface is anisotropic, and the direction and magnitude of the anisotropy (measured, e.g., in terms of contour

curvature as a function of tangent direction [Witkin 1981]) indicates the tilt and slant of the surface in each locality. In addition to isotropy, the physical curves must lie flush on a surface that is approximately planar in each vicinity, otherwise the contour curvature in the image texture would be confounded by surface curvature [Stevens 1981b]. So on the one hand, texture contours are apparently interpreted as asymptotic: all of their curvature lies in the tangent plane. On the other hand, surface contours are apparently interpreted as geodesic: all of their curvature lies normal to the tangent plane. (In other words, all of the curvature of a surface contour is attributed to surface curvature, while none of the curvature of a texture contour is.)

4. Parallelism between contours in the image may be regarded as a potential triggering criterion for this analysis, analogous to bilateral symmetry between contours potentially triggering their interpretation as silhouette (or occluding) contours of smooth generalized cones [Marr 1977]. Note that for the contours we study the corresponding 3-D curves may or may not correspond to physical surface edges (examine the contours in figures 1 and 2). When seen as a physical edge, the surface orientation varies across the contour discontinuously, i.e., the edge is sharp, in contrast to the contours along the silhouette of smooth objects.

5. It is useful to consider three distinct categories of contour: texture contours [Witkin 1981], occluding contours [Marr 1977], and surface contours [Stevens 1981a]. Each category of contour poses a different computational problem and each is expected to be solved by substantially independent perceptual processes. But since the three categories are often present simultaneously in a natural image there is need for integrating the surface shape information contributed by each process, presumably within the 2 1/2-D Sketch [Marr 1978]. The implementation that we shall discuss does not address the integration issue; it is meant only as an empirical demonstration of the plausibility of the constraints.

II GEOMETRIC CONSTRAINTS ON THE INTERPRETATION OF SURFACE CONTOURS

The definiteness of our perception of surface shape from these contours suggests that we effectively constrain the infinite range of possible 3-D surfaces down to a unique perceived surface (and its reflection, as discussed in point 2 above) by applying certain specific geometric constraints. The important question is precisely what are the constraints.

There are many possible starting points for pondering this question. One is to characterize the problem as one of smooth surface interpolation, where perhaps a computation starts with a plane parallel to the image that is continuously deformed in 3-D until it becomes the solution surface when certain criteria are satisfied, such as minimizing quadratic variation. Consideration of the details of the computation of this sort, however, does not

directly provide insight into the constraints that allow the solution. I refer here to the distinction between a computational theory (an explanation of what is accomplished) and a description of the algorithms that implement it (how is it accomplished) that Marr and Poggio [1977] discuss. We therefore defer discussion of the computation and consider the constraints themselves.

Surface contours, such as those in figure 1, are assumed to be the projection of physical contours that lie across a continuous surface. If we assume general position of viewpoint we may conclude that parallel contours in an image correspond to parallel physical contours. Consider now the relation between the parallel contours and the underlying surface. Parallelism weakly constrains how the surface may lie between the contours, particularly assuming the surface is opaque. Stronger constraint comes from assuming that the placement of the contours is generic, that their location on the surface does not conspire with the underlying surface geometry to mislead the viewer. This general position argument implies that the underlying surface is a cylinder, since cylinders are the only surfaces on which parallel contours may lie, in general. The third assumption, that the contours are lines of curvature, implies that they are also planar and geodesic, as consequences of being lines of curvature on cylinders. These constraints allow the estimation of local surface orientation in the following way.

A. Parallel correspondence

The lines in figure 1, by this hypothesis, are treated as lines of curvature across a cylinder. Moreover, the rulings that run perpendicular to these lines can be reconstructed by a straightforward geometric construction on the image, giving us the image of an orthogonal net lying across the surface. Since the rulings on a cylinder are parallel straight lines that project as parallel to the troughs and ridges in the figure and intersect successive curves at a constant angle, to reconstruct their projection it is sufficient to identify, for each point p on a contour, the point p' on the adjacent contour whose tangent is parallel to the tangent at p . The line pp' is the image of the ruling through p . Note that the constructed lines will themselves be parallel. Two such constructions are shown in figure 3a. This process may be regarded as establishing a parallel correspondence [Stevens 1981a] between adjacent contours. Note that contour parallelism detected by these means might serve to trigger the surface contour interpretation (see 4, above).

B. The perpendicularity between lines of curvature and rulings

The lines of least and greatest curvature on a surface are mutually perpendicular, hence the intersections in figure 3a are foreshortened images of right angles (examine the intersections in figures 2 and 4 as well). The perpendicularity




Figure 3.

places constraint on the range of possible surface orientations in these figures; both the slant and tilt are increasingly restricted as the right angle is increasingly foreshortened towards 180 deg in the image [Stevens 1981a]. The restriction on tilt can be determined by a simple construction in the image plane. At a point on a given image curve construct the perpendicular to the tangent and likewise the perpendicular to the ruling passing through that point, giving an angle less than 90 deg (figure 3b); the surface normal is then restricted so as to project within this range. Note that in figure 3b the surface orientation is least constrained at point 1 (tilt = 96.4 ± 38.6 deg) and most constrained at point 2 (tilt = 128.5 ± 6.5 deg).

Where the range of tilts is strongly constrained (and hence the slant is as well) the bisector is a good unbiased estimator of the tilt. For example at point 2 in figure 3c the bisector results in a tilt estimate of 128.5 deg, which implies a slant of 83.5 deg at that point. This estimate at point 2 allows surface orientation to be solved elsewhere along the contour where the orientation is less constrained. Since the lines of curvature on a cylinder are both geodesic and planar, the surface normal is constrained to rotate in the osculating plane, and once it is known at one point on the curve, it rotates in a predictable manner along the curve (see [Stevens 1981a, expressions 4 and 5]). Thus, based on the solution at point 2, these expressions predict tilt = 69.8 deg and slant = 47.4 deg at point 1. The surface normals at points 1 and 2 are depicted in figure 3c by directed line segments.

Figure 3d shows how the surface orientation may be solved by the same method at any point along the contour. Surface orientation is indicated in this figure by superimposing ellipses that correspond to circles lying on the "perceived"

surface.

Some final points should be made regarding tilt estimation by the bisector: (i) the estimate is increasingly accurate as the right angle is increasingly foreshortened towards 180 deg, (ii) it is an unbiased estimator, (iii) the bisector chosen as the tilt estimate corresponds to the least surface slant at that point [Stevens 1979], and (iv) while this method is applicable because lines of curvature form an orthogonal net across a surface, it is not central to the overall theory of constraint on the interpretation of parallel surface contours.

III SURFACE ORIENTATION PREDICTED BY THE THEORY

In figure 4a two simple figures are presented, each suggesting in 3-D a curve intersected at right angles by short line segments.

The predicted surface orientation was computed at the intersection points and indicated by foreshortened circles in figure 4b and by surface normals in figure 4c. The surface normal display in figure 4c is a sensitive probe of apparent tilt (particularly using tachistoscopic techniques [Stevens 1983b]). However indicating surface orientation by ellipses has two advantages over depicting surface normals. The first is the greater sensitivity that we have judging slant from the shape of an ellipse than from the projected length of a surface normal. (Observe that the differences in slant at the two intersections is much more apparent in figure 4b than in figure 4c.) The second advantage of using ellipses is that when the surface reverses in depth the ellipses reverse with it, while with normals, the reversal causes the line to appear to protrude behind the surface.




Figure 4.

The implementation has been applied to a variety of contour configurations that suggest cylindrical surfaces, such as figures 5 and 6. Note that the surface orientation predicted by the theory corresponds rather closely with the apparent surface orientation in these figures.




Figure 5.




Figure 6.

IV DISCUSSION

The implementation supports the hypothesis that our visual interpretation of parallel surface contours is constrained by assuming that the image curves are projections of lines of curvature, and by assuming that both the viewpoint and the contour placement are representative. The cylindricity of the surface, the fact that the rulings can be reconstructed by determining parallel correspondence, and the fact that the rulings are perpendicular to the lines of curvature combine to constrain the surface shape. The constraint is of two sorts. First the perpendicularity property restricts the surface orientation, secondly the geometry of cylinders allows one to compute the surface orientation at any point along a line of curvature as long as it is known at one point.

The performance of the implementation also lends some support to the bisector method for estimating surface orientation at the point of greatest constraint due to perpendicularity. The use of the bisector, while not critical to the overall theory, seems consistent with our perception of local surface orientation in these figures.

The theoretical emphasis thus far has been on the simplest configurations of surface contours, namely, those that suggest singly-curved surfaces. In [Stevens 1981a] it is proposed that more complicated configurations of roughly parallel contours, which suggest doubly-curved surfaces, are analyzed locally by detection of local parallelism among the contours, and exploitation of the constraints that arise by assuming (falsely) that the surface is locally cylindrical. The doubly-curved surface in figure 7 seems amenable to such an interpretation.




Figure 7.

Finally, to return to questions of algorithm, I suggest that the line of curvature constraint, cylindricity in the case of parallel contours, and the restrictions caused by the orthogonality of the lines of greatest and least curvature are used primarily to estimate local surface orientation. This local information is then integrated with

other information contributed from shading, texture and so forth, in order to derive an at least locally consistent description of the scene. It is premature to propose a specific algorithm until we know how this integration proceeds.

ACKNOWLEDGEMENTS

The late David Marr made important contributions to this work. I also thank Whitman Richards, Shimon Ullman, Donald Hoffman and Eric Grimson for helpful discussions.

REFERENCES

- Hilbert, D. & Cohn-Vossen, S. 1952 Geometry and the imagination. New York: Chelsea.
- Kender, J.R. 1980 Shape from Texture. Ph.D. Thesis, Computer Science Department, Carnegie-Mellon University.
- Marr, D. 1977 Analysis of occluding contour. Proc. R. Soc. Lond. B. 197, 441-475.
- Marr, D. 1978 Representing visual information. Lectures on Mathematics in the Life Sciences 10, 101-180.
- Marr, D. 1982 Vision. San Francisco: W.H. Freeman.
- Marr, D. & Poggio, T. 1977 From understanding computation to understanding neural circuitry. Neurosci. Res. Prog. Bull. 15, 470-488.
- Stevens, K.A. 1979 Representing and analyzing surface orientation. In: Artificial Intelligence: an MIT perspective. P.H. Winston and R.H. Brown, eds. Cambridge: MIT Press.
- Stevens, K.A. 1981a The visual interpretation of surface contours. Artificial Intelligence 17, 47-73.
- Stevens, K.A. 1981b The information content of texture gradients. Biological Cybernetics 42, 95-105.
- Stevens, K.A. 1982 Implementation of a theory for inferring surface shape from contours. M.I.T. A.I. Lab Memo 676.
- Stevens, K.A. 1983a Slant-Tilt: the visual encoding of surface orientation. Biological Cybernetics 46, 183-195.
- Stevens, K.A. 1983b Surface tilt (the direction of slant): a neglected psychophysical variable. Perception and Psychophysics, 33, 241-250.
- Witkin, A.P. 1981 Recovering surface shape and orientation from texture. Artificial Intelligence 17, 17-45.