MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 676

August 1982

IMPLEMENTATION OF A THEORY FOR INFERRING SURFACE SHAPE FROM CONTOURS

Kent A. Stevens[1]

ABSTRACT. Human vision is adept at inferring the shape of a surface from the image of curves lying across the surface. The strongest impression of 3-D shape derives from parallel (but not necessarily equally spaced) contours. In [Stevens 1981a] the computational problem of inferring 3-D shape from image configurations is examined, and a theory is given for how the visual system constrains the problem by certain assumptions. The assumptions are three: that neither the viewpoint nor the placement of the physical curves on the surface is misleading, and that the physical curves are lines of curvature across the surface. These assumptions imply that parallel image contours correspond to parallel curves lying across an approximately cylindrical surface. Moreover, lines of curvature on a cylinder are geodesic and planar. These properties provide strong constraint on the local surface orientation. We describe a computational method embodying these geometric constraints that is able to determine the surface orientation even in places where locally it is very weakly constrained, by extrapolating from places where it is strongly constrained. This computation has been implemented, and predicts local surface orientation that closely matches the apparent orientation. Experiments with the implementation support the theory that our visual interpretation of surface shape from contour assumes the image contours correspond to lines of curvature.

Acknowledgements: This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by the AFOSR and NSF grant 79-23110MCS.

---

1. Now at the Department of Computer and Information Science, University of Oregon, Eugene, Oregon 97403.

1. Introduction

The human visual system is adept at inferring the shape of a surface from an image consisting only of contours presumed to correspond to physical curves lying on the surface. In figure 1, for example, we perceive a rippled surface. The contours that comprise this figure are interpreted as arrayed in 3-D on a smooth but otherwise invisible surface.

Several questions are posed by figure 1. Most immediately, why do we take a 3-D interpretation instead of the (literal) interpretation of figure 1 as sinusoids on the plane of the printed page? The 2-D interpretation is almost impossible to hold, the tendency to see depth in the figure is so strong. Moreover, our 3-D interpretation is of a *particular* surface shape. Except for the expected depth reversal in figure 1 (which changes the apparent orientation of the surface), different observers see much the same rippled surface. What then determines the specific 3-D interpretation that we take? Clearly we incorporate geometric constraints strong enough to force a particular 3-D shape. What are they?

This article starts with a review of a theory [Stevens 1981a] regarding the specific geometric constraints underlying the interpretation of surface contours. As we shall see, a few constraints are sufficient to allow the determination of local surface shape from parallel[2] contours and certain related image configurations. Can we find evidence that our visual interpretation incorporates these particular geometric constraints? The intent of this article is largely empirical, and centers on experiments that compare human performance with the 3-D shape computed by an implementation of the theory. The observed similarity supports the theory.

2. Review of a theory of constraints on surface contour

2.1 Surface contours and their physical counterparts

A given contour in a natural image corresponds to perhaps an object edge, a crease, a shadow, or some surface

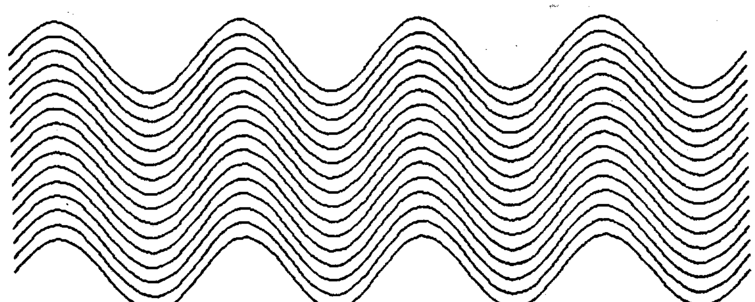

Figure 1. A set of parallel contours that are interpreted in 3-D as lying on a smooth cylindrical surface. The surface shape is vivid and one can judge with confidence the local surface orientation. This suggests that specific geometric constraints are imposed on the 3-D interpretation. Assuming that the contours are lines of greatest curvature on the surface plus weaker ancillary assumptions provides sufficient constraint on the problem of inferring local surface orientation from such configurations.

2. Recall that two arbitrary curves are parallel if one can be superimposed upon the other simply by a translation.

pigmentation marking. Each image contour has a specific physical cause. In this treatment, however, we will consider a surface contour only geometrically, as a curve in 2-D that corresponds to some curve in 3-D, without regard for what the physical curve may be. Of course, knowledge of the physical curve would be expected to simplify the interpretation task, but human vision does not rely on knowing this information: we infer a definite 3-D shape from purely abstract configurations of curves, as figure 1 demonstrates. (In the absence of information about the corresponding physical curves, we might take some default assumptions about them, but it is difficult to surmise *a priori* what those assumptions are.)

Not only do we infer surface shape from abstract lines, but the lines need not even be continuous in the image. In figure 2*b* the contours are defined by dots, and in figure 2*c* the contours are "subjective contours" defined by line terminations. The three surfaces seen in figure 2 are very similar. Thus we conclude that surface shape is inferred from a rather general representation of contour.

2.2 Terminology

The subsequent discussions largely concern the geometric properties of a particular type of contour (the line of curvature) lying across a particular class of surface (the cylinder). A brief discussion of these concepts is given below (see also *e.g.* [Hilbert & Cohn-Vossen 1952]).

Consider a smooth patch of surface and a plane that intersects the surface perpendicularly. That is, the surface normal lies in the plane. The intersection of the surface and the plane defines a curve whose curvature is called normal curvature. Suppose the plane is rotated, with the surface normal the axis of rotation. Depending on the surface in question, the normal curvature at the given point will vary between a minimum and maximum, called the principal curvatures. The tangent to the curve when the curvature is

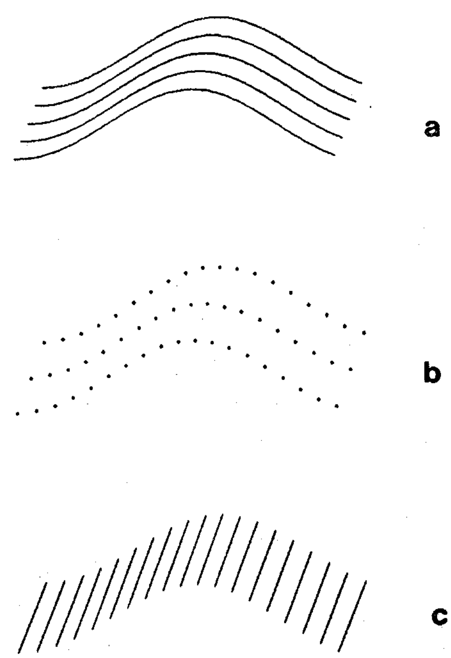

Figure 2. In *a* a ribbon-like surface is suggested merely by parallel lines. The same surface is suggested by dotted lines in *b*, and in *c* by line terminations. Various processes of contour construction precede the interpretation of surface shape from contour.

either maximum or minimum will point in one or the other of two *principal directions*, and for points on smooth surfaces the principal directions are mutually perpendicular. Thus each point on a smooth surface has two distinguished directions (except where the surface patch is planar or spherical, in which case the normal curvature is constant in all directions; such points are called umbilic).

The *lines of curvature* are curves across a given surface that follow one or the other principal directions. Specifically, the tangent at each point along a line of *greatest* (or *least*) curvature aligns with the direction of greatest (or least) normal curvature. Since the two principal directions, where defined, are mutually perpendicular, the lines of greatest and least curvature form an orthogonal net.

Of particular interest to us will be surfaces for which the lines of *least* curvature are straight lines, the so-called singly-curved or developable surfaces. Such surfaces correspond to the configurations that one can make with a sheet of paper, allowing bending and twisting but no tearing or creasing. On a singly-curved surface the tangent to a line of greatest curvature points in the direction in which the surface bends most rapidly, and for a line of least curvature the tangent points in the direction of zero normal curvature.

A *cylinder* is a restriction on the singly-curved surfaces wherein twisting is disallowed. That means the lines of least curvature (also called *rulings*) are all parallel. The cylinder corresponds to a rolled or rippled sheet of paper or a hanging curtain. The lines of greatest curvature, lying perpendicular to the direction of the ripples, are sometimes called *cross sections*.

The lines of curvature on a cylinder have two important properties: they are planar and geodesic. Being geodesic, the principal normal to the line of greatest curvature is identically the surface normal at that point. And since the curve is planar, the surface normal along that path is restricted to rotating in the plane containing the curve (the *osculating plane*).

2.3 Surface contours are distinct from texture contours and carry different 3-D information

An important distinction exists between the contours we will examine here and those that comprise image texture. While both correspond to physical curves across surfaces in 3-D and both contribute surface shape information, they differ in the way they "encode" that information.

Surface texture is foreshortened when the surface is slanted relative to the line of sight. The amount and direction of the foreshortening carries local information about the amount and direction of the surface slant. Thus local surface orientation can be estimated from texture foreshortening, provided the foreshortening can be quantified in the image, and provided certain physical restrictions in the physical texture are met. Early proposals for quantifying texture foreshortening (reviewed in [Flock 1964]) were in terms of the height-to-width ratio, which varies with the cosine of the slant angle in the ideal case of circular and flush-lying surface features. Recently, Witkin [1981] has shown that surface orientation can be estimated from local statistics of texture contour. A representative example of where contour foreshortening carries information about surface shape is given by the mottled pattern of sunlight and shadows cast on the ground below a tree. Instead of examining the height-to-width ratio of the individual patterns of light and dark, the method examines the contours that follow the boundaries between light and dark. The contour curvature varies in a statistical but systematic manner with tangent direction and surface orientation [Witkin 1981]. Hence from measurements of these statistics the local surface orientation can be estimated. The accuracy of the estimate depends on several factors. First, the physical texture must lie flush (as do mottled shadows on the smooth ground) so that they foreshorten according to slant angle in the expected manner (see [Stevens 1981b] for further discussion). Second, the curvature statistics of the physical texture must be isotropic, otherwise the texture anisotropy would resemble foreshortening and misleadingly appear slanted. Third, since the method is statistical, the surface must be roughly planar in each region over which the statistics are gathered, otherwise surface curvature would confound the foreshortening interpretation. Witkin [1981] demonstrates this computational method on slanted planar outlines and observes that if applied along the contour of an ellipse, the method estimates the spatial orientation of the corresponding circular disk.

Some method for inferring surface orientation from contour foreshortening similar to [Witkin 1981] is likely incorporated in human vision, and is applied to image configurations in which the contours appear foreshortened on a planar surface, provided there is not better evidence to the contrary. Figure 3 is an example: the pattern resembles a sheet of plywood or waves lapping at a beach. That is not the only 3-D interpretation available in this pattern, however. Like figure 1, figure 3 may also be seen as an undulating surface, but because of the amplitude of the curves, the surface dips dramatically from crest to trough.

Figure 3 is rotated by 90°, relative to figure 1, to bias its interpretation as a planar surface over that of an undulating surface.[3] While we strongly prefer the planar interpretation in figure 3, one can also interpret it as a deeply convoluted surface (as noted, view it so that the peaks and troughs would be oriented as in figure 1). As the amplitude of the sinsuoid pattern increases the planar-surface interpretation tends to dominate over the undulating-surface interpretation, however there is a rather broad intermediate range for which the two interpretations are strongly rivalrous. In this range the interpretation taken depends strongly on the orientation of the figure.

The two cases have a neat geometric interpretation in terms of the physical curves. In the planar-surface case the physical curves are essentially asymptotic lines; in the undulating-surface case they are geodesic (as will be discussed). Asymptotic lines have zero normal curvature, all of their curvature lies in the tangent plane. In contrast, geodesic lines have zero curvature in the tangent plane, all of their curvature is normal curvature. Thus the two interpretations will be appreciated as being the two extremes of a continuum. We will use the terms "texture contours" and "surface contours" to distinguish the two cases.

2.4 Surface shape is described by local surface orientation

It is proposed that surface orientation is the primary form of 3-D shape information derived from contours such as those in figure 1. Several factors support this proposal: it is feasibly computed, sufficient for deriving more global shape descriptions, and the apparent surface orientation is closely predicted by the theory of

3. It is well known that in most natural scenes distance increases as one scans from the bottom to the top of the visual field, a simple consequence of our conventional viewpoint of objects on the horizontal ground. In terms of surface orientation, this tendency is reflected in a bias to choosing an upward pointing surface normal. For illustration, consider viewing an ellipse with horizontal major axis: if interpreted as a circular disk on the ground the normal would point vertically upward. But if the ellipse's major axis is vertical, so that the normal would point to either the left or right, I have observed no apparent bias to one interpretation over the other. The result is that when the apparent surface normal points vertically the depth reversals are less frequent than when it points horizontally. This biases can be demonstrated for a wide range of stimulus surfaces, and is used to advantage in figure 3. The practical consequence is that the planar-surface interpretation in figure 3 is strongest when the apparent surface is horizontal. Observe that by rotating figure 3 by 90° the undulating-surface interpretation is easier to achieve and the planar-surface interpretation, when achieved, is less stable because of depth reversals. Similarly, rotating figure 1 so that the normals point roughly horizontally, rather than vertically, results in more frequent depth reversals.

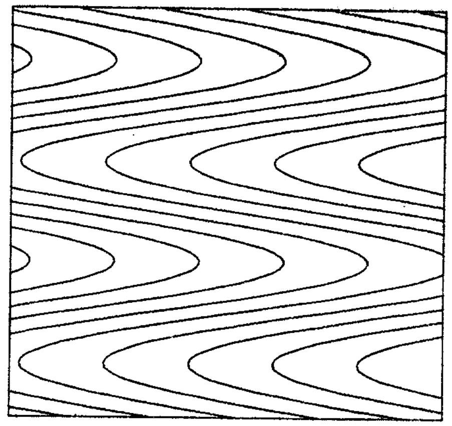

Figure 3. The curves tend to be interpreted as lying on a planar surface, although they may also be seen to lie across a dramatically undulating surface. The relative strength of the two spatial interpretations varies with the orientation of the figure (see text).

inferring surface orientation from contours. The computational feasibility will be reviewed momentarily; simply stated, local surface orientation is rather immediately encoded in the image while distance information is not. Surface orientation also provides a sufficient basis from which other surface shape descriptions can be subsequently derived such as distance (up to an additive constant and an overall multiplicative scalar), surface curvature, and topographical features such as ridges and troughs. The close match between apparent surface orientation and that predicted by the theory will be demonstrated later.

Surface orientation, having two degrees of freedom, can be regarded as a vector quantity having magnitude and direction. The magnitude component is the familiar variable *slant*, the angle σ between the line of sight and the local surface normal. Slant varies over a range from 0° (where the surface patch is perpendicular to the line of sight) to 90°. The other degree of freedom concerns the direction of slant, called *tilt* [Stevens 1979]. The tilt τ is the direction in which the surface normal would project in the image, and also corresponds to the direction of the gradient of distance from the viewer to the surface. Slant and tilt will be used to quantify orientation for our purposes.

The apparent surfaces suggested by contours in these figures are subject to the familiar "depth" reversals usually associated with the Necker cube. These reversals are expected, and correspond straightforwardly to ambiguity in the tilt *direction*. That is, the orientation of the surface can be recovered only to within a reflection in depth about the image plane, or equivalently, to within a 180° reversal in the direction the surface normal would point locally. Each choice of direction indicates the direction of the gradient of distance. While the subjective impressions of depth and of surface shape may change dramatically with these reversals, it amounts to only a reversal in the tilt component, the slant angle is unchanged. (It is worthwhile examining the tilt ambiguity in the various illustrations of this article, in order to convince oneself that when the apparent depth reverses, the tilt component of the surface orientation for any given patch reverses direction.) In light of this ambiguity, we expect to be able to compute surface tilt only to within a 180°

reversal.

2.5 The computational problem

The perceptual task we accomplish so effortlessly when viewing these figures can be described as a computational problem (in the sense of [Marr & Poggio 1977; Marr 1982]) as follows. We are presented with a set of contours $\{C\}$ in an image, assumed to correspond to a set of physical curves $\{\Gamma\}$ across a surface Σ. The task is to determine the local surface orientation of the visible patches of Σ, where surface orientation relative to the viewer is quantitifed by slant and tilt.

In studying how we accomplish this task, our approach will be to concentrate first on the case where the contours are parallel in the image, as in figure 1, then show how this problem extends naturally to certain other configurations.

Our tendency to take one specific 3-D interpretation in viewing figure 1 is so strong, we must be reminded that the problem is highly underconstrained. There are infinitely many 3-D surfaces that consistent with the given 2-D projection, including the surface of the page on which it is printed. What constrains the particular interpretation that we take? We have no expectations of viewing any specific surface shape, and no expectations for viewpoint relative to the surface. The central theoretical issue that emerges, then, is discovering what constraints are incorporated in the method adopted by the human visual system in "solving" this problem of 3-D perception.

Observe that figure 1 evokes the impression of viewing a continuous but transparent, or invisible, surface. One has a strong sense for the surface shape in the spaces between the contours as well as along them. We will be concerned primarily with determining orientation only for those surface points along a given physical curve Γ. In the intervening regions (where there is no surface information available) the orientation is presumably determined by interpolation. In the case of cylindrical surfaces, which are important here, the interpolation will be seen to be trivial. The more general case of interpolating doubly-curved surfaces is beyond the scope of this article.

2.6 Constraining the 3-D interpretation

In [Stevens 1981a] it is proposed that our interpretation is constrained primarily by assuming that the contours correspond to lines of curvature on the surface. That strong assumption, plus two assumptions to the effect that neither the viewpoint nor the particular placement of the physical curves on the surface is misleading, leads to an effective method for determining surface orientation from certain configurations.

2.6.1 General position of viewpoint: The surface Σ and the physical curves across Σ are assumed to be seen from a representative viewpoint, *i.e.* the image projection that results from this viewpoint is similar to the which would result from a slightly different viewpoint. By assuming this, we may exclude the following degenerate cases of projection: *i)* a planar curve projecting as a straight line, *ii)* two non-collinear lines projecting as collinear, and *iii)* two non-parallel curves projecting as parallel. Thus if a curve is straight in the 2-D image it is straight in 3-D, and likewise for two collinear or parallel lines in 2-D and their 3-D counterparts.

The *a priori* probability of encountering any of the above three degeneracies is small. It is particularly difficult for multiple, non-parallel curves appear parallel in the image. As the complexity of the projected curves increases it becomes increasingly improbable that two different and thus non-parallel curves would be placed fortuitously relative to each other and relative to the viewer so that their projections would differ only by a translation in the image. The other two degenerate cases, we recall, are: *i)* a straight line in the image that is actually the projection of a planar curve, and *ii)* two collinear lines in the image that are the projection of two lines that lie in the same plane but are non-collinear. Both are coincidental for the same reason: the line of sight happens to lie in the plane containing the planar curve or the two lines.

2.6.2 General position of contour: The surface geometry in the vicinity of the physical curve is assumed to be similar to that directly under the curve. Consider figure 1, where the surface between contours is invisible. An interpretation that we do *not* take is that the surface varies substantially between the contours. Instead, our interpretation is that the particular placement of the curves on the surface is not critical or fortuitous, so that if they were shifted laterally the image would appear substantially the same.

This assumption, plus the earlier assumption that the viewpoint is representative, allows one to conclude that in the vicinity of parallel physical contours the underlying surface patches are approximately cylinderical (see below). General position of contour is related to Grimson's [1981] observation that significant changes in surface geometry are usually reflected by intensity changes in the image. That is, if the surface deviates significantly from cylindricality between a given pair of parallel contours, it would usually be apparent in the image. Thus in the absence of evidence to the contrary, cylindricality can be assumed. This assumption seems *a priori* reasonable in the restricted context of parallel curves: in most instances where parallel curves lie on a surface, the surface in the vicinity is cylindrical.

2.6.3 The line of curvature assumption: The physical curve Γ is assumed to be a line of curvature across Σ, *i.e.*, the tangent to Γ everywhere coincides with one of the principal directions on the surface. In the case of parallel contours, since the surface in the vicinity is a cylinder, it follows immediately that Γ is also planar and geodesic.

The *a priori* justification of assuming surface contours are lines of curvature is rather difficult. But as observed in [Stevens 1981a], it is interesting that the vast majority of curves across the surfaces of synthetic objects of all sorts *are* lines of curvature. The reasons for this are many, and include the following: most manufactured objects are composed of surfaces of revolution or cylinders, if they are not planar, and the markings placed on these surfaces of revolution and cylinders almost invariably follow the cross sections, meridians, and rulings -- all of which are intrinsically lines of curvature. Not only are markings across the surface usually lines of curvature, but the seams, creases, and edges are as well. Even in cases where the surface geometry is more complex, such as metal castings and plastic injection-moldings, the seams are usually planar, normal sections, and lines of curvature (although this is sometimes not as immediately obvious by inspection as are the meridians on a surface of revolution or the cross sections of a cylinder).

In the same vein, the contours that arise in biological forms are very often lines of curvature. The joints in a bamboo stalk are cross sections and the markings that run lengthwise along the stalk are rulings -- both of which are lines of curvature. Stripes and linear markings on vegetation for various reasons, are often aligned with the principal directions on the surface. In fact, wrinkles in skin accompanied with flexure at joints or compression due to underlying muscles tend to follow lines of curvature. A final observation should suffice: on specular surfaces, either inherently glossy or merely wet, the highlights either appear pointlike or extended. In the case of extended specular reflections, if they appear as straight in the image, the surface is locally a cylinder, and the path on the surface corresponding to the specularity is a line of least curvature (see [Stevens 1981a]).

The difficulty of this sort of argument, of course, is that while one may discover many physical curves that are lines of curvature, another might find many other curves that are not. The *a priori* justification for assuming that physical curves are lines of curvature does not follow by such argument. In the synthetic world that argument is rather compelling, but not so in the natural world. The likelihood of being correct in assuming a given image contour corresponds to a line of curvature is significantly improved, however, if one explicitly excludes from consideration certain classes of image contours that are likely not to be lines of curvature. One class of contours to disregard are those that exhibit systematic evidence of foreshortening, such as do the texture contours discussed earlier. That is, the characteristic evidence of foreshortening associated with the projection of random curves across approximately planar surfaces can be used both as a "triggering condition" for estimating surface orientation from foreshortening, and simultaneously, as a condition for excluding them from consideration as lines of curvature.

2.6.4 Discussion: The line of reasoning just given can be easily summarized as follows. Suppose one is presented with parallel contours in an image, as in figure 1. They almost certainly correspond to parallel contours in 3-D by general position of viewpoint. The chance is negligible of viewing two differently shaped (thus non-parallel) contours from a perspective which causes their projections to be identical and to differ only by a translation in the image plane. Moreover, without evidence to the contrary, one can assume that the physical placement of the contours does not conspire with the underlying surface geometry to mislead the viewer. The precise placement of the contours on the surface is assumed not to be critical; they could be displaced and appear very similar. It follows that the surface is a cylinder in the vicinity of the physical contours. (Intuitively, that means that one can slide one curve along a straight line *on the surface*, without dips or rises, until it superimposes over the other.) Parallel image contours therefore imply cylindricality. Now, since the contours are also assumed to correspond to lines of curvature across the surface, it follows that they are planar sections and geodesics. The parallel curves correspond to parallel planar, normal sections of the cylinder.

Note that we start with only three assumptions: *i)* general position of viewpoint, *ii)* general position of contour, and *iii)* the line of curvature assumption. These assumptions logically entail the geometric constraints of planarity, geodesic, perpendicular intersection, and cylindricality. Thus, if the three assumptions are incorporated in human vision, the geometric constraints follow as consequences and need not be independently motivated. Next we discuss how the constraints allow surface orientation to be determined.

2.7 Inferring surface orientation

With reference to figure 4, consider the case of multiple parallel surface contours. We assume

general position of viewpoint (1)

general position of contour (2)

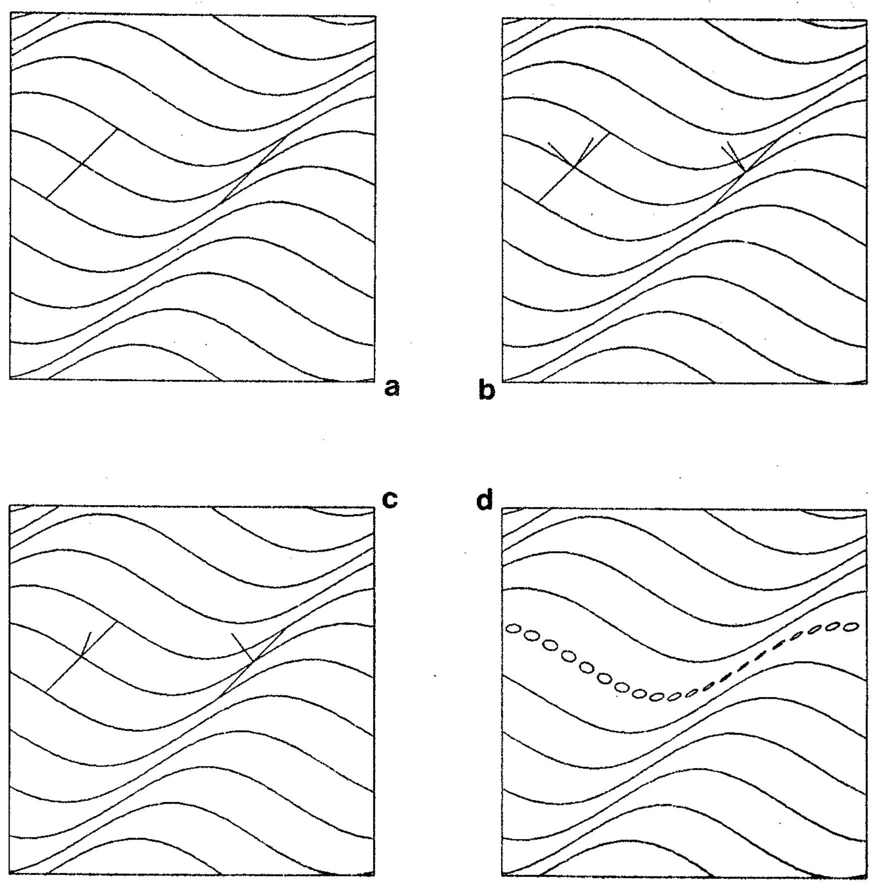

Figure 4. Parallel surface contours (*a*) are interpreted as parallel in 3-D and as lying on a cylindrical surface. At two points along one contour the rulings are shown to intersect. Note that the rulings are parallel straight lines. Each intersection is the oblique projection of a right angle on the surface, and this fact constrains the range of possible orientations for the normal, as shown in *b*. Note that where the angle of intersection is large the surface orientation is strongly constrained, and the bisector becomes an increasingly good estimate of the surface tilt. Provided the surface orientation is solved at one point along the curve it can be solved everwhere, see text.

Γ is a line of (greatest) curvature (3)

where Γ is the physical curve corresponding to a given contour C in the image. We will represent the physical curve Γ parametrically by $x_\Gamma = x_\Gamma(s)$ and define

$t_\Gamma(s)$ as the tangent vector of Γ

$n_\Gamma(s)$ as the principal normal vector of Γ

$\nu(s)$ as the normal to Σ at points along Γ.

Consider two contours C_i and C_j corresponding to physical curves Γ_i and Γ_j across the surface Σ. From (1) we have

$$C_i \parallel C_j \Rightarrow \Gamma_i \parallel \Gamma_j \quad (4)$$

and (2) and (4) yields

Σ cylindrical. (5)

That is, parallel curves imply a cylindrical surface, provided general position of both viewpoint and contour placement. From (3) and (5) it also follows that

Γ is geodesic (6)

$$\Gamma \text{ is planar, lying in some plane } \Pi \quad (7)$$

It follows, therefore, that the surface normal ν lies in the plane Π and is perpendicular to Γ. As the next step in deriving ν we examine the straight lines on the cylinder, the rulings.

The rulings, recall, are lines of least curvature on a cylinder, and since Γ is a line of greatest curvature, all rulings are perpendicular to Γ. Moreover, on a cylinder all rulings are parallel, as noted earlier. Hence, if we refer to a ruling as Λ, we have that

$$\Lambda \text{ is straight} \quad (8)$$

$$\Lambda \perp \Gamma \quad (9)$$

$$\Lambda_i \parallel \Lambda_j \quad \forall\, i, j. \quad (10)$$

Since all the rulings on a cylinder are straight and parallel, they share the same tangent vector t_Λ.

The tangent t_Λ to the ruling Λ is perpendicular to Π, the plane containing Γ. That is

$$t_\Lambda \perp \Pi \quad (11)$$

To see this, observe that Γ is geodesic (6) and planar (7) hence the plane Π is perpendicular to the tangent plane to the surface Σ. Since the ruling Λ lies in the tangent plane and is perpendicular to Γ, it is also perpendicular to Π. (Consequently the tangent to the ruling is equivalently the binormal to Γ, and the tangent, principal normal, and binormal to Γ are mutually orthogonal at each point along the curve.)

Thus we now may relate the surface normal $\nu(s)$ along the curve Γ in a simple manner to the tangents to the curve and the ruling

$$\nu(s) = n_\Gamma(s) = t_\Lambda(s) \times t_\Gamma(s). \quad (12)$$

Thus if the path of the curve Γ were determined in 3-D, the surface orientation along that path could be computed straightforwardly. In fact determining the path of Γ in 3-D reduces to the problem of determining the plane Π containing Γ (since Γ is simply the orthographic projection of C onto Π). Provided Π can be determined, the surface normal is computable either from the principal normal to Γ, or from the cross product relation in (12).

One possible strategy for solving for the surface orientation, therefore, would involve inferring the spatial orientation of the plane containing Γ. We shall show later, however, that the human visual system probably does not use this strategy. One that is more likely used is described by the following.

Consider the intersection of the line of greatest curvature Γ and the line of least curvature, or ruling, Λ. In the image Γ projects as the contour C, and the ruling Λ projects as the straight line R. Figure 4*a* shows two rulings, which appear as parallel straight lines, each intersecting the contour C at some angle β. The rulings, while not visible in figure 1, can be reconstructed by a simple geometric construction that follows from (8-10), *i.e.* the rulings on the surface are parallel straight lines that intersect each line of greatest curvature at a constant (right) angle. Consequently the projection of a ruling would be a straight line that intersects successive image contours at a constant angle β, and all rulings would be parallel in orthographic projection. Thus it is sufficient to identify points on successive contours having parallel tangents, and to connect those points with straight lines that are themselves parallel. The correspondence between parallel lines can be determined by a simple local computation; see [Stevens 1981a] for more detail.

We impose a Cartesian coordinate frame whose x-y plane lies parallel to the image and whose x-axis is aligned with R. Construct tangent vectors to Γ and Λ scaled such that they project as unit vectors in the image plane:

$$t_\Gamma(s_0) = \{1, 0, a\}$$
$$t_\Lambda(s_0) = \{\cos\beta_0, \sin\beta_0, b\}$$

where β_0 is the angle of intersection at the given point. The quantities a and b are the unknown components along the z-axis (perpendicular to the image plane). The surface normal ν is their cross product:

$$\nu(s_0) = \{-a\sin\beta_0, a\cos\beta_0 - b, \sin\beta_0\} \quad (13)$$

and the tilt τ and slant σ are:

$$\tau = \tan^{-1}\frac{\nu_y}{\nu_x}; \quad \sigma = \cos^{-1}\frac{\nu_z}{(\nu_x^2 + \nu_y^2 + \nu_z^2)^{1/2}} \quad (14)$$

where ν_x, ν_y, and ν_z are the three components of the normal vector.

The expression for the normal in (13) carries two unknowns. This reflects the two degrees of freedom of surface orientation when no restrictions are imposed. Now, given (9),

$$t_\Gamma(s_0) \cdot t_\Lambda(s_0) = 0$$

from which we have

$$b = -\frac{\cos\beta_0}{a}.$$

Substituting b into (13) gives

$$\nu(s_0) = \{-a\sin\beta_0, \frac{a^2+1}{a}\cos\beta_0, \sin\beta_0\}. \quad (15)$$

and therefore the surface orientation is determined up to only one unknown, a. Thus perpendicularity between the lines of greatest and least curvature removes one degree of freedom of surface orientation.

The magnitude of this constraint depends on the angle of intersection β. For any angle of intersection the tilt is constrained to lie within the perpendiculars to R and C, and the slant is constrained accordingly. As β approaches 180° the tilt approaches the bisector of the angle and the slant approaches 90°. The restriction on the range of possible tilts becomes appreciable as β becomes large:

$\beta = 135°$:	$\sigma = 77.75 \pm 12.25°$	$\tau =$ the bisector $\pm 22.5°$.
$\beta = 150°$:	$\sigma = 82.25 \pm 7.75°$	$\tau =$ the bisector $\pm 15.0°$
$\beta = 170°$:	$\sigma = 87.49 \pm 2.51°$	$\tau =$ the bisector $\pm 5.0°$.

Hence the bisector of the angle of intersection becomes an increasingly good estimate of the surface tilt. (Incidentally, for any obtuse β the bisector choice for the tilt corresponds to the surface orientation with the least slant.)

Returning to figure 4*a*, observe that the angle of intersection varies along the curve; β is greatest at the intersection shown on the right (call that point 0) and least at the intersection on the left (point 1). Thus at point 0 the surface orientation is most strongly constrained by the perpendicularity. Specifically for figure 4, $\beta = 167.1°$ at point 0, and the tilt is $128.5 \pm 6.5°$. The bisector of the angle of intersection therefore is a good estimate of the local surface tilt at point 0.

At point 1, on the other hand, $\beta = 102.9°$ and the tilt is only constrained to $96.4 \pm 38.6°$. Nonetheless, having determined the orientation at point 0 will allow the tilt to be solved here and at all other points. (Intuitively, since the Γ is 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.)

Point 1 is defined as $\mathbf{x} = \mathbf{x}(s_1)$, and the tangents to the curve and the ruling are

$$t_r(s_1) = \{1, 0, a\}$$

$$t_A(s_1) = \{\cos\beta_1, \sin\beta_1, c\}$$

where β_1 is the angle of intersection at that point. Observe that since the rulings are parallel, $t_r(s_1) = t_r(s_0)$. The surface normal is, again, the cross product of the two tangent vectors:

$$\nu(s_1) = \{-a\sin\beta_1, a\cos\beta_1 - c, \sin\beta_1\} \tag{16}$$

where

$$c = -\frac{\cos\beta_2}{a}. \tag{17}$$

Note that the one remaining variable at this point, c, is related to a. But observe also that a is simply related to τ:

$$\tau = \tan^{-1}\frac{N_y}{N_x} = \tan^{-1}(-\cot\beta_1\frac{a^2+1}{a^2}).$$

Thus a can be solved in terms of the known tilt τ and the obtuse angle β_1

$$a = \left(\frac{-1}{\tan\tau\tan\beta_1 + 1}\right)^{1/2}. \tag{18}$$

Finally, substituting a into (17) and (16) completely determines the surface orientation.

Figure 4*c* shows the bisector estimate for the tilt at point 0 ($\tau = 128.5^\circ$, $\sigma = 83.5^\circ$), and the surface orientation solved for point 1 ($\tau = 69.8^\circ$, $\sigma = 47.4^\circ$). This was determined simply by the computation just described. As a first demonstration of the implementation, figure 4*d* shows how circles would appear if lying on the surface along the path of the curve, demonstrating that the orientation can be determined for all points along the curve.

To review, the surface orientation is strongly constrained in places where Γ intersects a ruling at a very large obtuse angle. Knowing the orientation at that point means that the orientation is known everywhere along the curve. This might be described as extrapolation of the solution from places where it is strongly constrained. But note that the solution at any point is locally determined given one global parameter (which we defined as the variable a).

Before examining further demonstrations of this method, it is instructive to examine, rather informally, whether our 3-D interpretations seem consistent with the underlying geometric constraints, namely, that the contour is planar, geodesic and a line of curvature, and that the surface is cylindrical.

3. Qualitative observations

3.1 Smooth curves generally appear planar in 3-D

A solitary curve in monocular presentation usually appears planar and in 3-D. The curve in figure 5*a*, for example, seems to lie in a different plane than the plane of the page. (It is suggested that this and subsequent figures be viewed monocularly with sufficient viewing time allowed for depth impressions to develop.) In fact it is quite difficult to perceive the torsion of a solitary twisting curve from a single orthographic projection -- figure 5*b* is the projection of a helix, although it is difficult to see it as such. The interpretation as a twisting space curve is easier to achieve if the curve projects as a self-intersecting image curve. The helix in figure 5*c*,

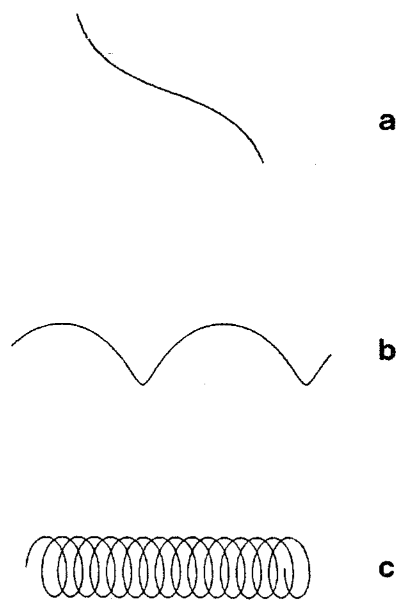

Figure 5. A smooth curve is usually seen as planar (*a*). (Observe this and subsequent figures monocularly, and allow sufficient viewing time for depth impressions to develop.) It is quite difficult to perceive the torsion of a solitary twisting curve from a single image unless, say, it is self-intersecting. Compare *b* and *c*, both of which are are helices.

for example, is more obviously twisting in space.[4]

3.2 Intersections are generally seen as right angle intersections in 3-D

A simple oblique intersection appears to be the image of a right-angle intersection in space (figure 6). The literal interpretation of either intersection in figure 6 as an oblique (non-right angle) intersection in the plane

Figure 6. Example of simple intersections that are strongly interpreted as right-angle intersections in 3-D.

4. Incidentally, planarity seems to be applied only piecewise along extended contours. An image curve without inflection points in curvature or discontinuities in tangent is usually seen as a single planar curve. But at such points the curve might appear to twist abruptly. As a result, the overall curve appears to be composed of piecewise-planar arcs in 3-D. In large part the spatial orientation of each individual arc seems to be independently determined. The reader is invited to draw various isolated curves to further examine the tendency for planarity and the instances where curves are seen as piecewise planar or to have torsion.

of the page is very difficult to achieve in monocular or even binocular presentation. This tendency to take a right-angle interpretation has been observed in various contexts, including illusions (see *e.g.* [Luckiesh 1965; Robinson 1972]), and in conjuction with 3-D interpretations of geometric figures [Attneave & Frost 1969; Attneave 1972; Perkins 1972; Finch 1975; Shepard 1979; Stevens 1982]. Further instances of perceived right-angles occur in other figures in this article.

3.3 Intersections generate apparent surfaces in 3-D

A simple curve having only a weak 3-D interpretation (figure 7*a*) suddenly provides a strong impression of a surface when intersected by a line segment. The apparent surface shape is substantially independent of whether the line intersects an endpoint (figure 7*b*, *c*, *f*, and *g*) or some interior point (figure 7*d* and *e*). Note further, by comparing figures 7*b*, *d* and *f* with figures 7*c*, *e* and *g*, that the curve can be interpreted either as a physical boundary or an interior marking without substantially affecting the apparent surface shape. In either case the surface appears to be cylindrical. (In figure 7 all intersecting lines have the same orientation; the intersection angle critically affects the apparent surface -- see later.)

3.4 Parallel curves appear to lie on cylindrical surfaces

Two arbitrary smooth curves that are parallel (*i.e.* differ only by a translation) are sufficient to elicit a 3-D interpretation. The curves are seen as the two edges of a ribbon-like surface (figure 8*a*). Note that additional parallel curves enhances the impression of a surface but the apparent surface shape is substantially unchanged (figure 8*b*). Note further that the spacing between curves is immaterial. In figure 9 the curves are precisely parallel but the spacing between curves is random. Under careful inspection figure 9 defines the same cylindrical surface as seen in figure 8*b*.

The random spacing between the parallel contours, however, produces an illusion of shading, which in turn may purturb the impression of viewing a cylindrical surface. First consider figure 8*b*, where the line density varies across the pattern. There is a diagonal swath across the figure where the curves nearly touch, which appears shaded in parafoveal vision. The shading suggests surface curvature along that diagonal, which likely enhances the impression of surface shape. (Shading due to variations in line density probably contributes to the impression of three-dimensionality in these patterns, but as figure 2 demonstrates, shading is not wholly responsible.) Returning to figure 9, besides the shading along the diagonal, the random spacing of the parallel lines produces additional shading that suggests slight ripples running perpendicular to the major troughs and crests of the surface. When the surface shape is examined closely, however, these shading effects largely vanish.

Finally, consider figure 10, in which straight line segments intersect a given curve. If the lines segments are parallel the surface appears cylindrical (figure 10*a*). If they diverge slightly (figure 10*b*) we attribute the divergence to perspective. Thus the same surface is seen in each case, one in orthographic projection, the other in perspective (*i.e.* from close by). It is much more difficult to take the alternative interpretation that the surface is developable but twisted, as if a sheet of paper were twisted into a helicoid. (Incidentally, observe that the parallel rulings in figure 10*a* might actually appear to converge in 3-D.)

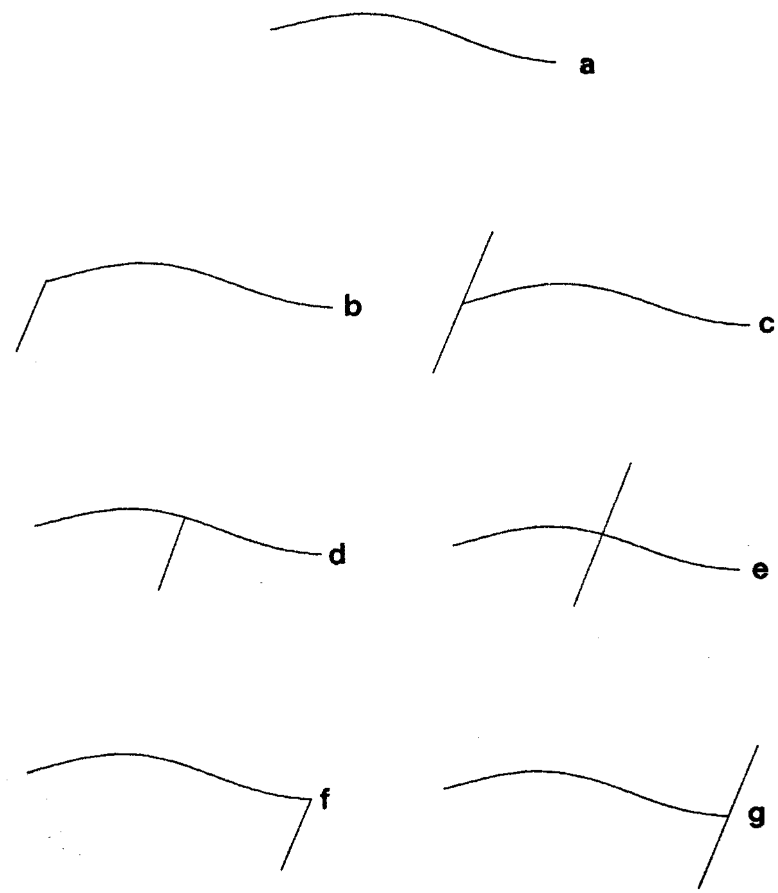

Figure 7. In *a* the simple curve has only a weak 3-D interpretation. But the addition of lines that intersect it, in *b*-*g*, cause us to see a surface that resembles a gently curved sheet of paper. To demonstrate that the curve can be interpreted as either a physical boundary or an interior marking, without changing the apparent surface shape, compare *b*-*c*, *d*-*e* and *f*-*g*. To demonstrate that the location of the intersection along the arc is not critical, compare *b*-*d*-*f*, and compare *c*-*e*-*g*.

3.5 Parallel curves appear to be lines of greatest curvature on cylindrical surfaces

Referring back to all preceding figures, the general impression is of curves that follow the greatest curvature across the surface. This is supported particularly in figures 4, 7, and 10, where the straight lines (that follow the direction of least curvature across the surface) appear to be perpendicular to the curves. Also supporting this is the following: the curves appear to be planar normal sections of cylinders and all planar normal sections of a cylinder are lines of curvature.

Moreover, consider the cylindrical surfaces depicted in figure 11. The curves appear to follow the direction of greatest curvature on the surface. Informally, they resemble flow lines over a waterfall, or sections of an airfoil. In each case they appear as planar normal sections, and as just mentioned, it therefore follows that they are lines of greatest curvature on the cylinder.

Figure 11 also demonstrates an important point regarding how the curves themselves are seen in 3-D. Each curve in figure 11 is an arc of an ellipse. Usually, an isolated arc of an ellipse assumes the 3-D orientation corresponding to an arc of a slanted circle, just as a complete ellipse is seen as a slanted circle in

a

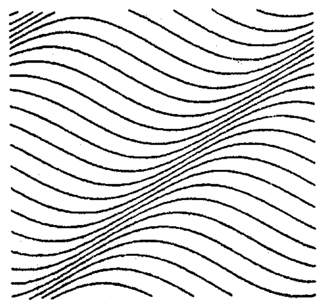

b

Figure 8. Two smooth curves that are parallel suggest a ribbon-like, cylindrical surface (*a*). The impression of a surface is enhanced but the apparent surface shape is unchanged when additional parallel curves are added (*b*).

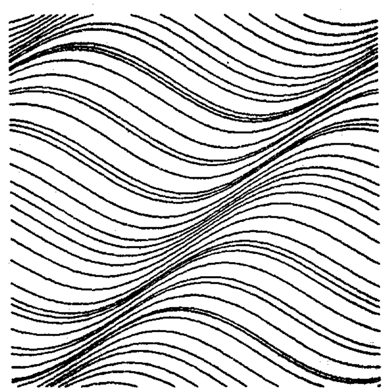

Figure 9. Randomly-spaced parallel contours generate the same apparent surface shape as seen in figure 8*b*, where the contours are spaced equally.

3-D.[5] (Hence an ellipse or a section of an ellipse appears as a section of a constant-curvature curve in 3-D.) But when the ellipse sections are arrayed in parallel fashion, as in figure 11, the dominant impression is of

5. To some extent this can be observed by scrutinizing the bottom-most ellipse of each surface in figure 11. But note that the other ellipses assume a different orientation in 3-D.

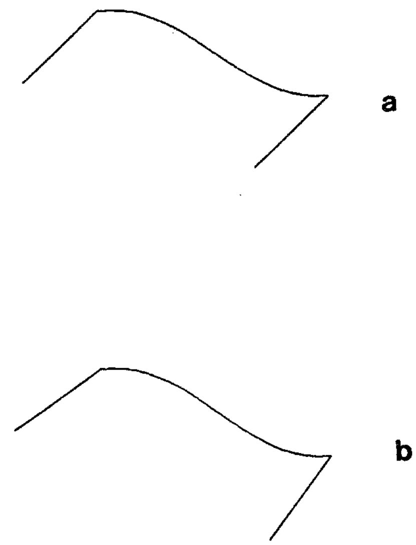

Figure 10. In *a*, where the intersecting lines are parallel, the surface appears cylindrical. In *b*, where they are not, the same surface is seen, but in perspective.

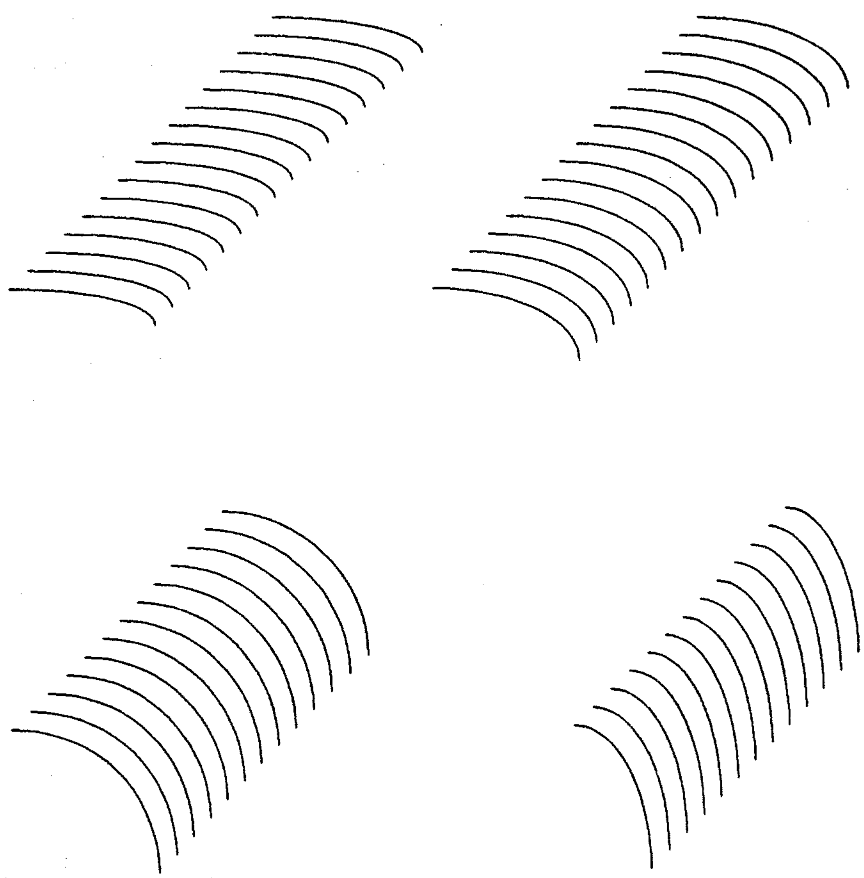

Figure 11. These cylindrical surfaces appear to be scored with lines of greatest curvature. Note that the curves appear planar and to be normal sections of the surface. It is significant that the 3-D curves do not appear to have constant curvature despite the fact that they are drawn as arcs of ellipses -- see text.

curves lying on a surface such that the normal curvature is *not* constant along the curve. The spatial orientation of the osculating plane of each curve is apparently not dominated by the shape of the individual projected curve. (This suggests that the human visual system does not first determine the orientation of the plane containing the curve, then use that to determine the surface orientation, as mentioned earlier.) The bottom left surface in figure 11 presents a particularly clear demonstration since the drawn curves are, in fact, arcs of circles. Rather than appear to lie in planes parallel to the printed page, they appear to be cross sections seen from an oblique viewpoint.

Figure 12 demonstrates that parallel copies of virtually any curve can induce the impression of a cylindrical surface. Here, rather than precise sinusoids or arcs of ellipses, the repeating curve is a spline that was drawn through some points that were chosen arbitrarily. These parallel lines appear to lie on a surface that is precisely cylindrical (a straight line placed on the surface parallel to the troughs and ridges would touch the surface at all points along its length).[6]

Finally, observe in figure 13 that the global surface shape is not restricted to being strictly cylindrical. This surface appears doubly curved, and resembles concentric ripples on a pool. The curves are locally, but not globally, parallel. It is conjectured that doubly-curved surfaces such as that depicted by figure 13 are perceived by the same processes that derive the 3-D shape of cylindrical surfaces from parallel contours. That is, the method involves a local correspondence between parallel segments of contours (to estimate the angle β,

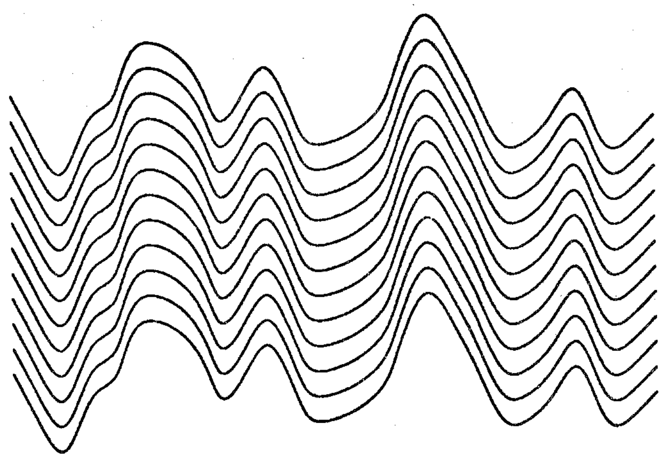

Figure 12. The irregular cylindrical surface is suggested by multiple parallel copies of a smooth spline drawn through arbitrarily-chosen points. The curves are interpreted as normal sections of the cylinder, and the ridges and troughs run perpendicular to the curves, thereby following the lines of least curvature or rulings of the surface.

6. Except along some ridges where the surface may appears to dip slightly between the contours. This effect, curiously, seems associated only with apparent ridges. Therefore when the the figure reverses in depth so that what were seen as (convex) ridges now appear as (concave) troughs, the surface along the trough appears precisely cylindrical. The depth reversal can be induced most easily by simply inverting the figure (since we tend to interpret distance as increasing as one scans from bottom to top on the figure, this biases the depth interpretation in these figures). Where the apparent surface undulates as one scans along a ridge, the contours seem to be interpreted as bounding (or silhouette) contours [Marr 1977] of slight peaks. The impression is enhanced by obscuring all but a single ridge and its flanking slopes on either side: each individual curve may appear to be the outline of a small conical hill. Again, when the figure is inverted, this strip looks like a culvert (*i.e.* cylindrical); it is quite difficult to interpret it as an inverted view of a row of small conical hills. This curious phenomenon seems independent of the tendency to interpret parallel contours as lying on a cylindrical surface.

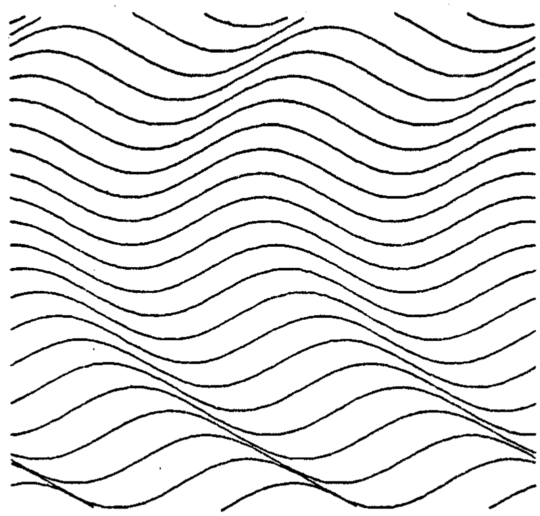

Figure 13. The doubly-curved surface resembling ripples on a pool arise from sinusoids having a gradually-varying phase shift. The curves are locally, but not globally, parallel. Arbitrary doubly-curved surfaces are amenable to a local approximation by cylinders provided the contours are sufficiently closely spaced.

etc., as discussed in section 2.7) and this method is tolerant of gradual variations in the correspondence across the surface. Thus, provided the contours are sufficiently closely spaced, arbitrary surfaces can be analyzed by this method.

4. Demonstration of the implementation

Informal observation suggests that we interpret parallel curves as planar, geodesic and as lines of curvature across cylindrical surfaces. These observations are consistent with the hypothesis that the human visual system incorporates assumptions (1-3). Stronger support is provided by a computer implementation of the method defined by (14-18). The surface orientation predicted by this method is in close agreement with the apparent surface orientation.

The implementation computes the surface normal at points along a curve from image configurations such as we have examined. The bisector of the angle β is used as an estimate of the surface tilt τ at the point where β is largest. The tilt estimate allows the variable a to be computed from (18). Given a, the normal at that point is completely specified, so the slant can also be determined from (14). (Note that the bisector solution is but one of several independent means for estimating the orientation at a point, from which a may be determined.) For any other point along the contour, the value of c may be computed from (17) given a and the measured value of b at that point. Given c at that point, the surface normal (16) is also determined, so that the tilt and slant at that point can also be computed from (14). The computation of surface orientation at any point is therefore a simple function of the local parameter β and the global parameter a.

The surface orientation computed by this method can be graphically depicted by superimposing a directed line segment that represents the local surface normal (figure 4*c*) or, say, by computing how a small circle lying on the surface would project onto the image plane, (figure 4*d*). The local surface orientation predicted by this method would seem in close correspondence with that perceived by the human observer if the superimposed line segments appear perpendicular to the surface or the superimposed ellipses appear to be circles that lie in the tangent plane of the surface. A demonstration of both schemes is given in figure 14. Figure 14*a* shows a curve intersected by parallel straight line segments. The implementation computed the surface normals shown in figure 14*b*, and the equivalent foreshortened circles in figure 14*c*. One may compare the apparent surface orientation in figure 14*a* with the predictions below it. (The figure should be viewed monocularly, oriented perpendicular to the line of sight, and studied for a few seconds in order to develop a strong impression of three-dimensionality.)

While the depiction of a surface normal by a short line segment is an excellent means for indicating local surface tilt, the slant, indicated by the length of the line segment, is much harder to judge from the displayed line. While we are sensitive to the length of the depicted surface normals (because they appear obviously incorrect when all are displayed as with equal length rather than foreshortened by the sine of the slant angle) it is difficult to make subtle judgments of slant from variations in the projected line lengths.[7] Furthermore, the line segments that represent surface normals may influence the apparent surface shape. In figure 14*c*, the addition of short line segments to the simple configuration of figure 14*a* brings a complicating factor. As noted earlier, we have a tendency to interpret intersecting straight lines segments as perpendicular and equal-length in 3-D. Hence the added line segments in figure 14*c* have a strong potential for influencing the spatial interpretation of the basic figure. (There seems to be less interaction when the depicted surface normal

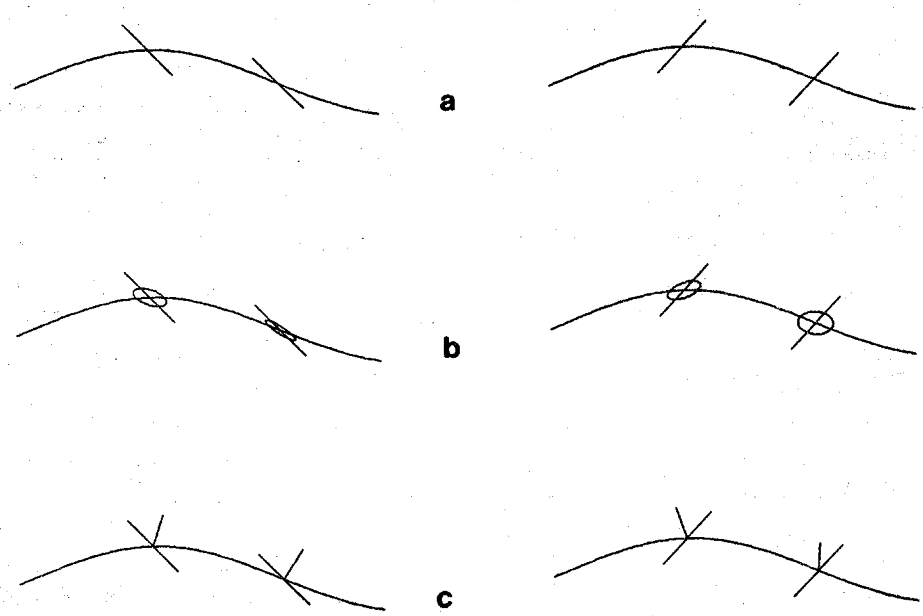

Figure 14. The configurations in *a* appear in 3-D. Compare the spatial orientation at each intersection point with that predicted by the implementation. Orientation is depicted in *b* by the projection of a unit surface normal, and in *c* by the projection of a circle lying flush on the surface.

7. When one considers what inferring local slant by this means would entail, this observation is not surprising. Note how little the line lengths vary in figure 14*c* compared to the differences in foreshortening to the corresponding ellipses in figure 14*b*.

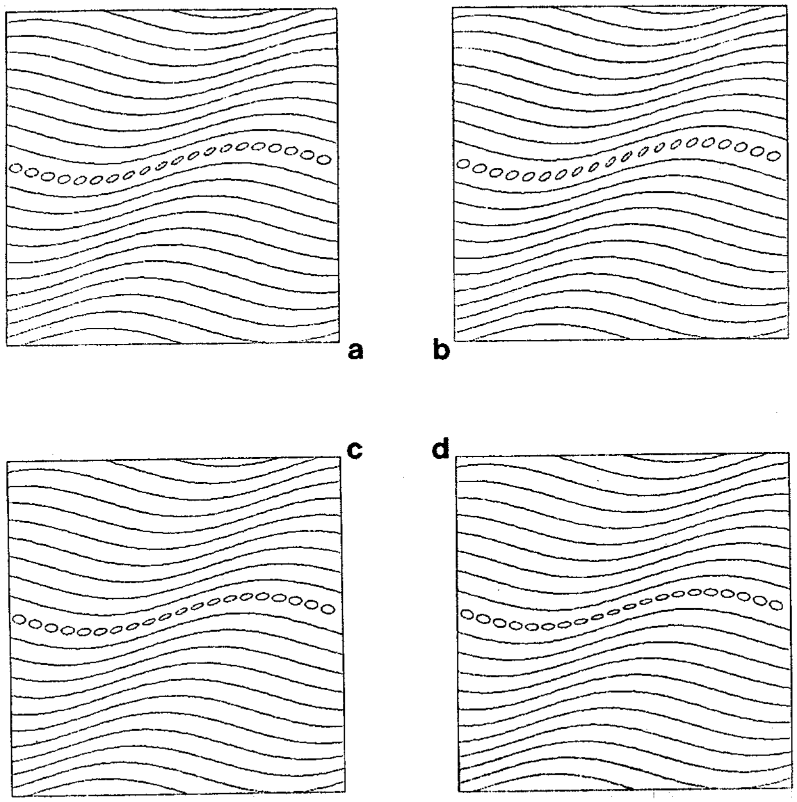

Figure 15. The shape and orientation of the ellipses in *a* are predicted by the theory and appear as circular disks, each lying in the tangent plane to the surface at that point. In *b* the ellipses are systematically rotated 10° counterclockwise relative to the corresponding disks in *a*. Similarly, in *c* and *d* the ellipses are rotated 10° and 20° clockwise. Observe that when the orientation of the ellipse is inappropriate the interpretation is either of a circular disk touching the surface at only one point on its perimeter, or of an ellipse drawn on the surface.

---

is presented only briefly while continuously viewing the stimulus figure [Stevens 1982].)

Consider now the use of ellipses to depict the surface orientation predicted by the implementation. An ellipse of given shape and orientation is very effective in suggesting a circular disk of the corresponding surface orientation. Superimposed on the image of a surface, an ellipse is interpreted as a disk that contacts the apparent surface, and if oriented appropriately, the disk appears to lie in the tangent plane to the surface at that point. The ellipse may be used as a probe of apparent surface orientation, since ellipses of inappropriate orientation or shape will not appear to lie on the surface. To demonstrate, figure 15*a* shows a row of ellipses, each of which appears, to most observers, to lie flush on the surface. In figure 15*b* the ellipses are systematically rotated 10° counterclockwise relative to the corresponding ellipses in figure 15*a*. Similarly, in figures 15*c* and 15*d* the ellipses are rotated 10° and 20° clockwise. Observe that when the orientation of the ellipse is inappropriate the interpretation is either of a circular disk touching the surface at only one point on its perimeter, or of an ellipse drawn on the surface. In figure 15 the apparent tilt of the disks was varied, holding slant constant. Figure 16 varies the slant while holding tilt constant. Figure 16*a* (which is identical to figure 15*a*) shows ellipses that appear to be circles on the surface, while the ellipses in figure 16*b* correspond to disks that are systematically slanted and additional 10°, and in figures 16*c* and 16*d* the disks are slanted by 10° and 20° less than the corresponding disks in 16*a*. Note that when the disks are inappropriately slanted

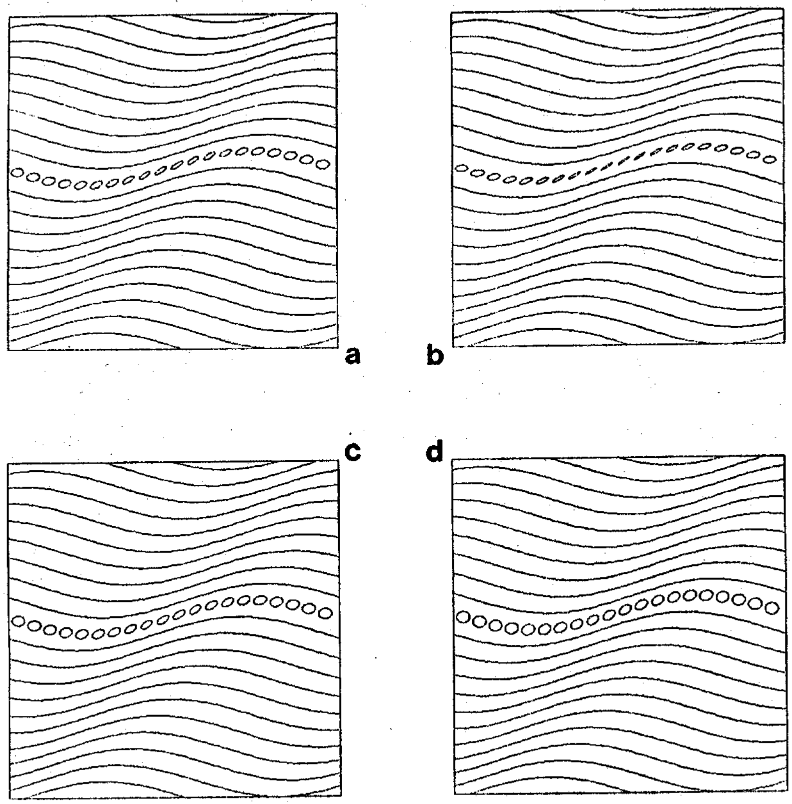

Figure 16. In *a* (which is identical to figure 15*a*) the ellipses appear to be circles on the surface. The ellipses in *b* correspond to disks that are systematically slanted and additional 10°, and in *c* and *d* the disks are slanted by 10° and 20° less than the corresponding disks in *a*. Note the apparent step-like discontinuity in the surface in *c* and *d* (in *a*, on the other hand, the surface appears smooth and continuous across the row of disks).

relative to the surface they tend to introduce an apparent step-like discontinuity in the surface (see figures 16*c* and 16*d*). In figure 16*a*, on the other hand, the surface appears smooth and continuous across the row of disks. As these figures demonstrate, we are sensitive to "errors" in the orientation and shape of the superimposed ellipses, which correspond to discrepancies of (somewhat less than) 10° in tilt and slant relative to the apparent orientation of the underlying surface.

The local surface orientation predicted by the theory is demonstrated, using ellipses, in figures 17, 18, and 19. Observe that in each instance the ellipses appear as disks lying on the surface, even when the figure reverses in depth. Figure 17 shows variations in apparent surface shape with changing the amplitude of the sinusoids. Likewise, in figure 18 the phase between adjacent sinusoids is varied. Finally, figure 19 shows the implementation applied to the cylinders of figure 11, with similar results.

5. Discussion

The implementation supports the hypothesis that our visual interpretation is constrained by treating the image curves as projections of lines of curvature assuming both the viewpoint and contour placement are representative. Moreover, the specific method of estimating surface tilt at the point of greatest constraint by the bisector, while not critical to the overall theory, seems well supported by the implementation. The

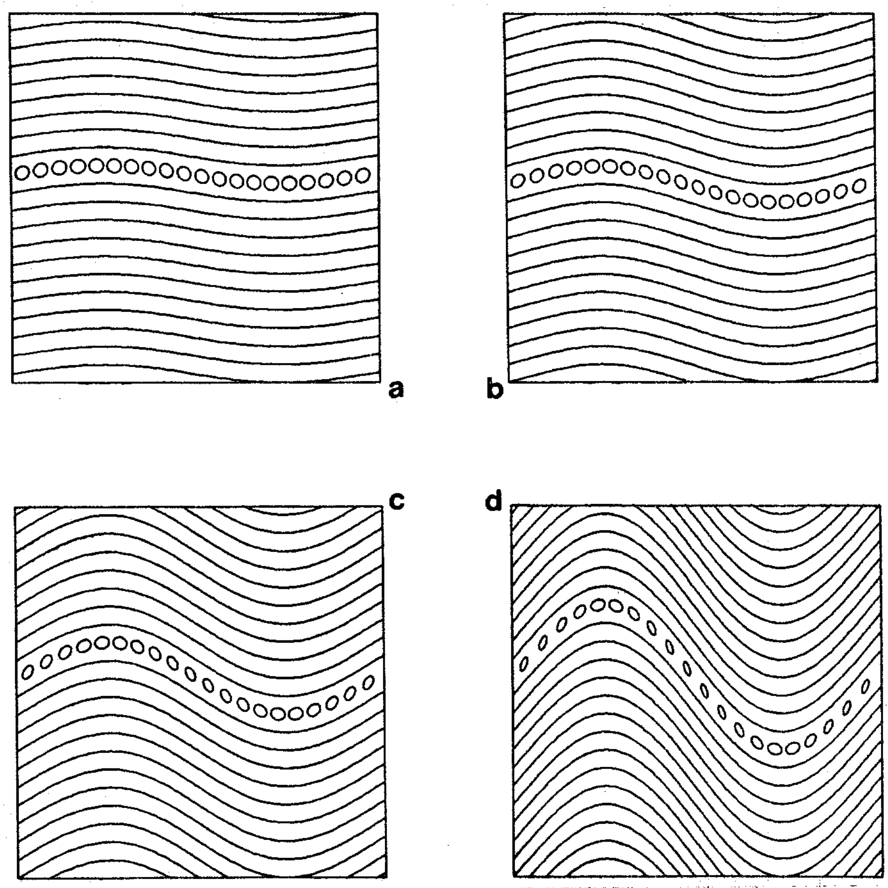

Figure 17. Sinusoid patterns of differing amplitude. In each case one sinusoid is replaced with a series of ellipses that appear to lie flush on the subjective surface. Note that they continue to appear appropriate when the figure suddenly reverses in depth or as one rotates the figure.

implementation was intended only to examine these two points; it was not otherwise meant as a model of the corresponding visual process in human vision.

A common geometry is shared by the configurations consisting of parallel curves (figure 1), a single curve intersected by parallel straight lines (figure 10*a*) and the rather "minimal" case of a single curve intersected by a single straight (figure 7). All visually suggest cylindrical surfaces, where the curves correspond to lines of greatest curvature and the straight lines correspond to rulings. We saw a geometric argument for why parallel curves imply that the surface is cylindrical in section 2.7. The same argument can be adapted to the case of a single curve intersected by parallel straight lines -- that geometry also implies cylindricality. The only curious case is why a single curve intersected by a straight line also looks like a cylinder. At most, if the line segment corresponds to a ruling, one could conclude that the surface is developable but not necessarily cylindrical. (Since there is only one ruling visible, one can determine from this configuration whether or not the surface normal twists along the curve.)

Let C be the image of a given line of curvature Γ and R that of a given ruling Λ. Since the surface is a cylinder, all rulings are mere translations of one another, and all lines of curvature are likewise translations of one another. Thus one can reconstruct how all rulings and lines of curvature would appear in the image by

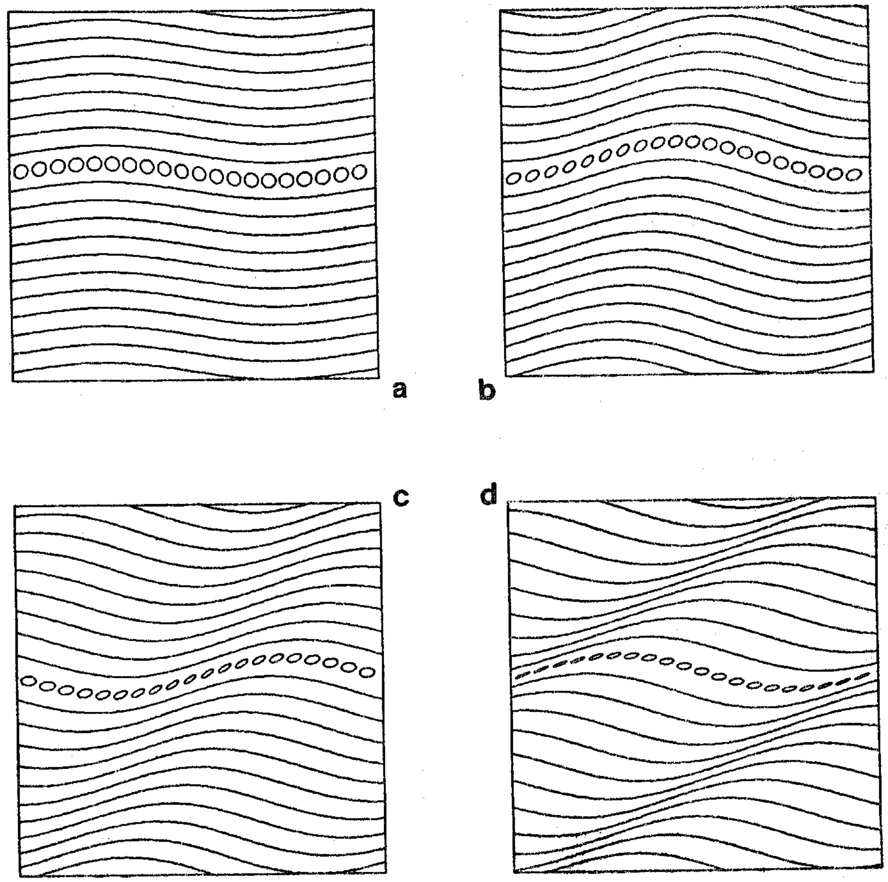

Figure 18. Sinusoid patterns of differing phase shift, constant amplitude. Note that the ellipses appear appropriate in each case, even when depth reverses and as the figure is rotated.

translating C in either direction along R, and by translating R in either direction along C. This would reconstruct the orthographic projection of an orthogonal net or grid across the surface. Thus although one starts with a simple configuration of a few parallel curves or a curve intersected by one or two straight lines, one can effectively reconstruct how a dense grid across the surface would appear. By this means the surface orientation can be determined for arbitrary points along a cylinder. The interpolation of surface orientation is conceivably applied to doubly-curved surfaces (such as figure 13) provided they may be approximated locally by cylinders. It remains to be determined how well a simple grid interpolation accounts for the apparent surface orientation in those cases.

To summarize, three assumptions (line of curvature and general position of viewpoint and contour) allow one to conclude that the underlying surface is cylindrical when presented with parallel contours and certain other configurations. It is possible to reconstruct the projection of an orthogonal net across the surface, and provided that the surface orientation is determined at one point on the surface it can be solved everywhere else. (The bisector is one of several methods that may be used to determine orientation at certain points in order to seed the computation.) Thus the computation of surface orientation from the image configurations we have examined is, at least in principle, amenable to a very simple local algorithm requiring but one global parameter.

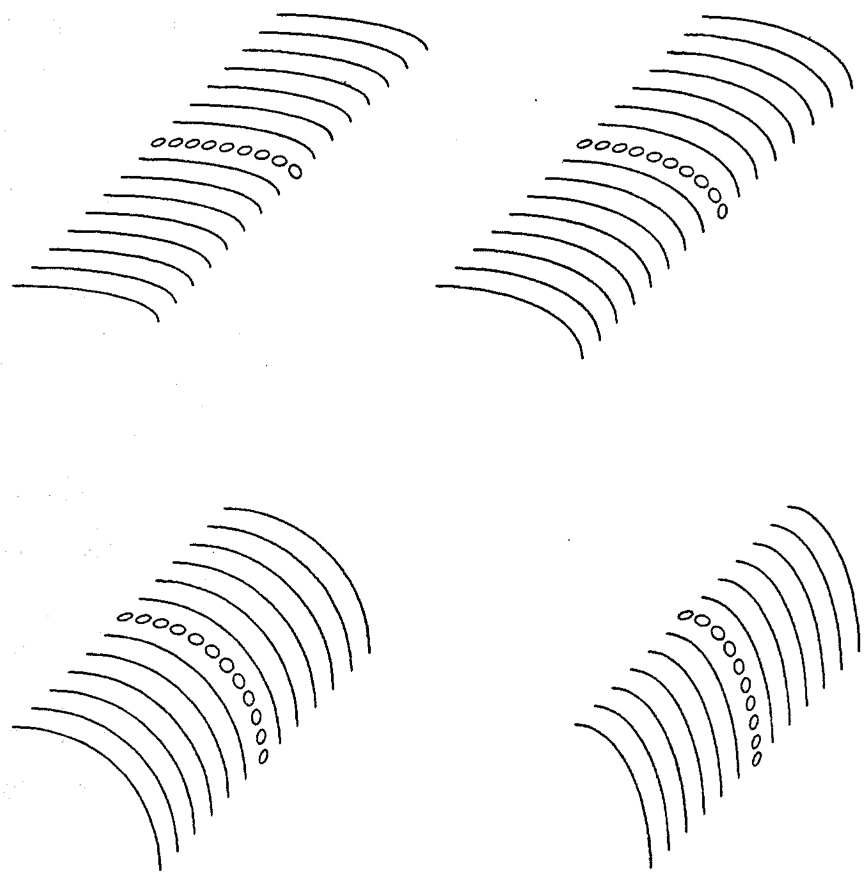

Figure 19. The cylinders in figure 11 are reproduced here with ellipses drawn by the implementation. Note that in each case they appear appropriate.

Acknowledgments

The late David Marr supervised the dissertation research from which this work derives; I am deeply indebted to him. I also thank Whitman Richards, Shimon Ullman, Donald Hoffman and Eric Grimson for helpful discussions. This work was conducted at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by the AFOSR and NSF grant 79-23110MCS.

References

Attneave, F. 1972 The representation of visual space. In *Coding processes in human memory.* A.W. Melton & E. Martin (eds). New York: John Wiley.

Attneave, F., Frost, R. 1969 The determination of perceived tridimensional orientation by minimum criteria. *Perception and Psychophysics* 6, 391-396.

Finch, D. 1975 Description of a new visual illusion: a drawing of a four-dimensional cube. *Perception* **4**, 111-113.

Flock, H.R. 1964 Three theoretical views of slant perception. *Psychological Bulletin* **2**, 110-121.

Grimson, W.E.L. 1981 *From images to surfaces.* Cambridge: MIT Press.

Hilbert, D. & Cohn-Vossen, S. 1952 *Geometry and the imagination.* New York: Chelsea.

Luckiesh, M. 1965 *Visual illusions: their causes, characteristics and applications.* New York: Dover.

Marr, D. & Poggio, T. 1977 From understanding computation to understanding neural circuitry. Neurosci. Res. Prog. Bull. **15**, 470-488.

Marr, D. 1977 Analysis of occluding contour. *Proc. R. Soc. Lond. B.* **197**, 441-475.

Marr, D. 1982 *Vision: a computational investigation into the human representation and processing of visual information.* San Franscisco: W.H. Freeman.

Perkins, D.N. 1972 Visual discrimination between rectangular and nonrectangular parallelopipeds. *Perception and Psychophysics* **12**, 396-400.

Robinson, J.O. 1972 *The psychology of visual illusions.* London: Hutchinson University Library.

Shepard, R.N. 1979 Psychophysical complementarity. In *Perceptual organization*, M. Kubovy & J.R. Pomerantz, eds. Hillsdale, N.J.: Lawrence Earlbaum Associates.

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 Surface tilt (the direction of slant): a neglected psychophysical variable. *Perception and Psychophysics*, in press.

Witkin, A.P. 1981 Recovering surface shape and orientation from texture. *Artificial Intelligence* **17**, 17-45.