The Visual Interpretation of Surface Contours

Kent A. Stevens
M.I.T. Artificial Intelligence Laboratory,
545 Technology Square, Cambridge, MA 02139, U.S.A.

ABSTRACT
This article examines the computational problems underlying the 3-D interpretation of surface contours. A surface contour is the image of a curve across a physical surface, such as the edge of a shadow cast across a surface, a gloss contour, wrinkle, seam, or pigmentation marking. Surface contours by and large are not as restricted as occluding contours and therefore pose a more difficult interpretation problem. Nonetheless, we are adept at perceiving a definite 3-D surface from even simple line drawings (e.g., graphical depictions of continuous functions of two variables). The solution of a specific surface shape comes by assuming that the physical curves are particularly restricted in their geometric relationship to the underlying surface. These geometric restrictions are examined.

1. Introduction

Of the means available to the visual system for determining the shape of a surface, stereopsis and motion predominate. Shading, when the illumination is directional, and texture gradients, when the surface is visibly textured, are also important. There are two more sources of shape information: boundary contours and surface contours. The contours that outline the boundary of a surface constrain the surface shape interior to the boundary, when the surface is smooth. Contours that lie across the surface are also useful, and they are the subject of this article.

Consider the common practice of mathematicians and engineers to graphically depict a continuous function of two variables, $z = f(x, y)$, as a surface seen from an oblique viewpoint. The technique is to project the curves that result from holding one parameter constant (for various values) while continuously varying the other parameter. An example is the sine function in Fig. 1 whose 3-D shape is readily apparent. The depiction of surfaces by contours is so familiar to us that we must pause to realize that it entails a significant problem of visual interpretation. Observe that a valid (and in fact, the correct interpretation) of Fig. 1 is that the surface is planar—it is the page of this journal on

Artificial Intelligence 17 (1981) 47–73
0004-3702/81/0000–0000/$02.50

which the undulating curves are printed. We do not easily take that interpretation: instead we see an undulating surface in 3-D. The widespread use of surface contours for conveying surface shapes—in mathematical texts, commercial drawings, and even cartoons—shows that their visual interpretation is definite and consistent. In short, there is a natural way of graphically conveying the shape of a surface by depicting how certain curves would lie across it. When drawn accordingly, we infer the intended shape.

Just as Fig. 1 has two radically different 3-D interpretations, undulating curves on a flat sheet or curves across an undulating surface, there is also an infinity of intermediate interpretations, at least theoretically. There are no physical laws that force any particular surface interpretation. Therefore human vision must incorporate particular constraints on the interpretation. This article examines what those constraints are. (For further discussion on the role of constraints in vision see [11, 13, 17].)

A word is needed regarding the 'restrictions' which we impose on the physical geometry. A surface contour is the image of a physical curve Γ across a surface Σ. This Γ may be the locus of some pigmentation change—a stripe on a zebra, say. When seen from a particular viewpoint, the physical curve in 3-D projects into a curve C in the 2-D image. The only evidence of surface is indirect, in the way Γ projects to C. Therefore, inferring something about the shape of Σ from C is possible only if the relation between Γ and Σ is restricted. This is a crucial point: we can constrain the surface shape only if the physical curve has some restricted relationship with the surface on which it lies.

The central problem is therefore cast as one of discovering the underlying

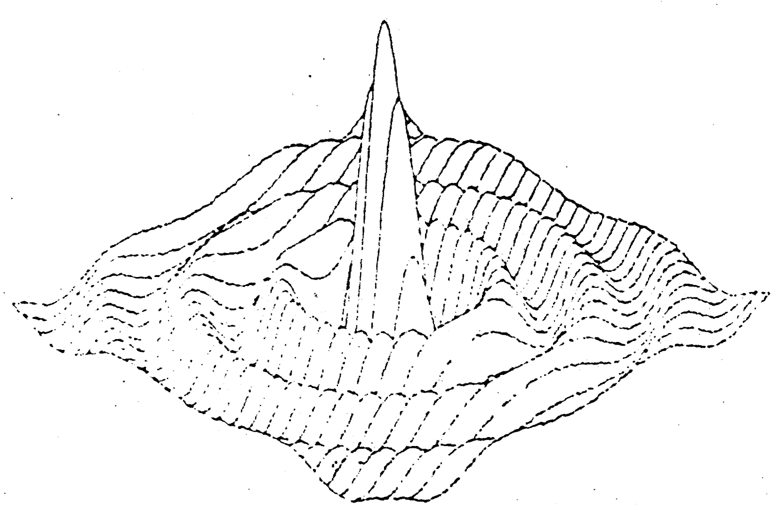

FIG. 1. It is commonplace to graphically depict a continuous function of two variables as a surface seen from an oblique viewpoint merely by a set of curves. The 3-D surface shape is immediately apparent.

geometric constraints. A traditional method for investigation would be to start with the basic physics of the situation, as Horn [6] successfully did with regard to shape from shading. The approach is to understand what physical laws govern the image formation process. In studying surface contours, the perspective geometry carries over from other investigations, but the problem breaks soon after into various, seemingly independent cases, according to what particular physical event causes the surface contour. Consider some possible causes of surface contours: an illumination change such as the edge of a cast shadow (tree shadow cast across fallen snow), a reflectance change such as a pigmentation marking (zebra stripe, stripe along the length of a spider plant leaf) or various sorts of surface features (circular joint between sections of a bamboo stalk, wrinkles on skin, fracture crack across a rock). Each case would be examined in order to uncover geometric restrictions on how the physical curves are produced in nature.

This approach, suggestive of the work of D'Arcy Thompson [16], seems intractable in its complexity. Certain statements may be made, e.g., that wrinkles due to compression and stripes on plants are usually close approximations to lines of curvature. But for every physical curve that has some neat geometric restriction, one may find another curve that is unrestricted, which meanders across the surface arbitrarily. Thus it would be infeasible to start with the basic physics—we would not know which physical situations to model.[1] But nonetheless, it is clear that certain assumptions about the *geometry* of the curves must be made if we are to use them to infer surface shape. So rather than study the basic physics of surface contours, we will study the basic geometric reasoning. This will give us some insight into the sorts of geometric constraints that are needed, and with that insight we may examine where those constraints arise in nature. The strategy that we will pursue is (i) to consider the various types of shape information that are plausible and useful, (ii) to determine the minimum geometric restrictions that are sufficient properties of various types of real physical curves.

We will see that there is a broad range in specificity of the 3-D shape information that might be inferred, and that the likelihood of the interpretation being correct decreases with increasing specificity. In other words, the more precisely we wish to determine the true 3-D shape, the less likely we will be successful in general. There is, therefore, a progression from weak to strong information about the shape of the surface, and associated with each type of information are restrictions that must be met in order to extract that information from the image. Rather than examine a range of possible shape descriptors, why not simply concentrate on the best that can feasibly be

[1] Stated another way, we do not know what physical interpretation we make regarding the curves in Fig. 1, if any. We might be assuming that the lines are pigmentation markings on the surface, or perhaps thin shadows cast across the surface, as if from a picket fence. Still other proposals may be made—this question probably cannot be settled by introspection.

computed from the image? The reason is that we do not yet know enough to judge what is the best, and furthermore, we do not know what specific shape information the human visual system extracts when viewing, say, Fig. 1.

While surface contours have not been studied psychophysically, the following is easily verified: we readily derive from a line drawing, such as Fig. 1, a *qualitative* appreciation for the shape of the depicted surface—e.g. where it undulates, where it is planar. We can also judge its relative orientation with some confidence, but we have little sense for its distance or scale. Reflecting on this, what can we infer from an ability to judge some 3-D property? It does not immediately follow that that property is explicitly represented. Surface orientation, however, is arguably represented explicitly by the human visual system [10, 12]. But what about the more qualitative descriptions of surface shape—are they also explicitly represented in the visual system? It is premature to say: the possibility is raised here primarily to emphasize that surface orientation is not the only possibility for local surface representation.

Furthermore, the fact that we have a definite impression of surface orientation does not necessarily mean that we compute it *directly* from the contours in the image. A 'direct' computation would derive the given output without intervening representations. The input would be the set of curves in the image, the output would be a surface orientation map. An alternative 'indirect' method would consist of two representations: the final representation would be a surface orientation map as before; the earlier representation would be a symbolic shape representation specifying, say, where the surface is planar, where it is singly curved, and where it is doubly curved. Associated with each representation would be a visual process that develops the representation. So in this case, the input to the first process would be the surface contours in the image and the output would be the qualitative shape representation. That information would then feed a subsequent process whose job it is to make the surface orientation map consistent with the symbolic shape information and other constraints such as smoothness and various boundary conditions.

In summary, there are several possible representations of local surface shape incorporated in human vision. We do not yet know whether the more qualitative representations are incorporated in human vision at this level of processing, but we should examine how they would be developed from surface contour information.

2. Describing Surface Shape

This section examines a range of shape descriptors, but in doing so it introduces a number of concepts from differential geometry.[2] So we first list the

[2] The reader is referred to Hilbert and Cohn-Vossen [4] for an excellent discussion of differential geometry.

descriptors, in order to avoid losing track of them in the middle of all terminology.

(i) *Planar* versus *curved*. This is the most primitive qualitative statement about local surface shape we will consider.

(ii) *Developable* (and *cylindrical*) versus *locally convex* versus *hyperbolic*. Developable surfaces are single curved and correspond to smooth twisting and foldings of a paper sheet. (A cylinder is a further restriction to an untwisted paper sheet.) The hyperbolid (saddle shaped) and locally convex surfaces, on the other hand, are doubly curved. This three-way distinction is complete and is captured by the sign (zero, positive, negative) of the local Gaussian curvature.

(iii) *Local surface orientation* (*slant-tilt*). While (i) and (ii) are qualitative and describe the surface in a manner independent of the viewpoint, the local surface orientation is quantitative and describes the local shape relative to a particular viewpoint. Slant and tilt is a useful and natural formalism for this task.

Probably the most conservative statement about surface shape is simply to distinguish whether the surface is *curved* or *planar*.[3] While this attributes special status to planarity, the distinction is probably biologically important. In lay terms, 'planar' is synonymous with 'flat', and the distinction 'flat' versus 'curved' is highly intuitive and seemingly primitive. The planar/curved distinction tells us little about the surface; nonetheless, the knowledge that certain regions are planar would be useful to a smooth surface interpolation of the sort just discussed. We will not spend time considering whether planar/curved is made explicit in human vision—it is simply a starting point for a range of shape descriptors.

A stronger statement about shape than (i) is to describe *how* the curvature of the surface varies locally. If a patch of surface is non-planar, then a path drawn across it will be a curve in space whose curvature depends on the direction of the path across the surface, in general.[4] This gives us the basis for describing the way surface curvature varies locally. It will be treated formally momentarily, first let us visualize this notion with several examples. To begin with, consider the various directions in which one could proceed away from a point on a (right circular) cylinder. In the direction parallel to the cylinder axis the path is a straight line, in the perpendicular direction the path is circular, and in intermediate directions the paths are elliptical. If the surface patch is saddle-shaped, however, some paths would arc downward while others would arc upward. Finally, on a round pebble every path would arc in the same way. More formally, by 'path' we mean a curve defined by the intersection of a plane with a patch of surface, where the plane is oriented so that it contains the

[3] A planar surface patch has zero normal curvature in every direction at a point. I will use 'curved' to mean 'non-planar'.

[4] If the curvature is constant in all directions the surface patch is either spherical or planar.

normal to the surface at a given point and cuts the surface in some direction. The curvature of that curve measures the *normal curvature* at a point on a surface taken in the given direction. As the examples show, normal curvature depends on the type of surface and on the direction in which the surface is cut. A fact of fundamental importance to us is that the directions of maximum and minimum normal curvature, the so-called *principal directions*, are always mutually perpendicular at any point on a smooth surface. For example, the two principal directions on a cylinder are parallel and perpendicular to the axis of the cylinder.

The curvature in either principal direction is termed *principal curvature*. It is a signed quantity, and the previous examples illustrated three combinations of the two principal curvatures. On the saddle surface the principal curvatures have opposite sign, on the pebble they have the same sign, and on the cylinder one of the principal curves vanishes to zero.[5] These possibilities correspond to a three-way distinction which seems primitive and natural, as does the distinction between planar and curved. A curved surface is either singly curved (like a curtain or a gently folded sheet of paper) or doubly curved and either saddle-shaped or convex.

It is important to note that we seek only a qualitative description of shape here, and not any quantitative measure of surface curvature. Therefore it is sufficient to only consider the sign of the two principal curvatures. It so happens that the *Gaussian curvature* is convenient in this regard. Gaussian curvature κ is the product of the two principal curvatures κ_1 and κ_2 (the curvatures measured in the two principal directions). So if the two principal curvatures have the same sign, the Gaussian curvature will be positive; the surface at that point is called *elliptic*. A surface patch consisting entirely of elliptic points is termed *locally convex*, or simply *convex*, an example being a round pebble.[6] When the two principal curvatures have opposite sign their product κ is negative—the point is called *hyperbolic* and the surface is locally saddle shaped. We will call a surface patch *hyperbolic* as well, if it consists entirely of hyperbolic points. When either (or both) principal curvatures vanish, their product, the Gaussian curvature κ, is zero. A surface patch of zero Gaussian curvature is called *developable*.

Developable surfaces, or 'singly curved' surfaces, correspond to the smooth bendings and twistings of a piece of paper that are possible without tearing. If the paper is allowed to bend but not twist, one has the special class of surfaces called *cylinders*. To give a more precise definition it will be necessary to introduce the notion of *line of curvature*, a curve which follows one of the two principal directions. A line of *greatest* curvature, for instance, follows the direction of greatest normal curvature across the surface. (Since the two

[5] The fourth possibility is that both curvatures vanish, i.e. the surface is planar.

[6] We will not distinguish convex from concave—that distinction can only be made relative to a viewpoint.

principal directions are perpendicular at each point on a smooth surface, the lines of greatest and least curvature form an orthogonal net across the surface.) In an extended region where κ is zero, at least one of the principal curvatures vanishes, hence at least one of the lines of curvature is a straight line. We can now define a cylinder as a surface for which the lines of least curvature are parallel straight lines. (This is illustrated in Fig. 8a.) Cylinders have several important properties that we will exploit later. Furthermore, one may locally approximate real surfaces by cylinders, so long as the principal curvatures are very different in magnitude.

Some illustrations may help. The plane in Fig. 2a and the cylinder in Fig. 2b both have zero Gaussian curvature. The convex surface in Fig. 2c has positive curvature and the saddle surface in Fig. 2d has negative curvature.

Note that much of the surface in Fig. 1 may be approximated locally by cylindrical patches. This idea of local cylinder approximations to arbitrary doubly curved surfaces is powerful, but it is successful only if the two principal curvatures are very different in magnitude.

The sign of the Gaussian curvature (whether κ is positive, negative, or zero) provides a weak but useful characterization of surface shape in the immediate vicinity of a point. An artibrary surface may have some regions of positive and some regions of negative Gaussian curvature, with necessarily intermediate points of zero curvature. The description of the visible surfaces of an object in this manner would be potentially useful for visual recognition, for the description is object-centered [12]. Another representation of shape would consist of merely noting those places of zero versus non-zero Gaussian curvature. There

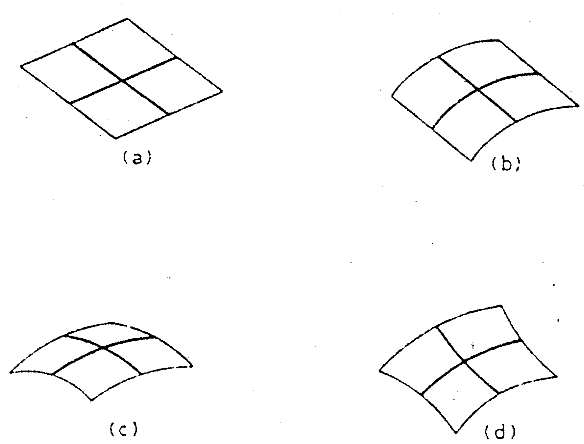

FIG. 2. Examples of surface patches of zero, positive, and negative Gaussian curvature. The planar surface in a and the cylinder in b both have zero Gaussian curvature. The surface in c is locally convex (positive curvature) and in d the surface is hyperbolic (negative curvature).

are two reasons for proposing this representation. First, an image often gives positive evidence for distinguishing patches of zero from non-zero Gaussian curvature (we will discuss some later, see also [2] regarding occluding contours [19] on shading). Second, Bruss [3] has shown that patches of negative and positive Gaussian curvature cannot be distinguished locally on the basis of the shading. Therefore an early representation of surface shape might only differentiate zero from non-zero Gaussian curvature, for that is often the most that can be stated conservatively (see 'the principal of least commitment' regarding visual processing discussed in [8]).

So far the shape descriptions are qualitative. To be quantitative about surface shape (e.g. to describe the magnitude of the Gaussian curvature, not just the sign) requires information relative to the particular viewpoint. The difficulty this would entail is equivalent to determining the relative surface orientation.

Surface orientation may be described in terms of the *slant* and *tilt* introduced in [15], which specify how much (slant) and which way (tilt) the tangent plane of the surface is inclined with respect to the image plane. Tilt is measured as an orientation in the image plane, and can be thought of as the orientation to which the surface normal would project. Slant is an angle, ranging from zero (where the surface is parallel to the image plane) to 90° (where the surface is completely foreshortened). Their relation to the more familiar Cartesian coordinates p and q of gradient space [6, 7] is:

$$\text{slant } \sigma = \tan^{-1}(p^2 + q^2)^{1/2} \quad \text{and} \quad \text{tilt } \tau = \tan^{-1}(q/p).$$

The slant-tilt form will be used in Section 3.3, where we consider the restrictions that allow us to derive surface orientation (see also [5, 14, 18]).

The various shape descriptors just given should be justified, for one might equally propose descriptors such as 'spherical', 'egg shaped', 'hour-glass', and so forth. The descriptors based on local curvature, however, have several advantages. First, each level of description is complete: every smooth patch of surface is either planar or curved, has either positive, negative, or zero Gaussian curvature, and every smooth visible point has a definable surface orientation. Second, they are local and (except for surface orientation) they are qualitative, while such descriptors as 'spherical' are more global and require quantitative knowledge of the curvature. Third, and most important, these descriptors are feasibly computed from contours in the image without requiring multiple views, prior knowledge of the surfaces, or other sources of information.

3. Geometric Restrictions

We now take each level of shape descriptor and examine what is required in order to infer that information from the surface contours in the image. It is important to keep in mind that the following discussion is only geometric with the goal of finding least restrictions. The justification or the relevance of these restrictions is a separate matter which will be examined subsequently. To give an example, we will find that restricting the physical curve to be a line of greatest curvature is useful. But when, in reality, are physical curves so restricted? In fact they often are, but for now we will merely seek geometric restrictions that are useful. And perhaps contrary to intuition, differential geometry has shown that there are few possibilities open to us.

We must assume that the surface in the vicinity of a curve (where we have no information) is 'well behaved'—that there are no invisible troughs or undulations. We rule out, for instance, the possibility the curves in Fig. 2a lie on two intersecting ridges. The property of the surface being 'well behaved' is captured by assuming the placement of the physical curves on the surface is not critical—that if displaced slightly relative to the surface they would appear qualitatively the same in the image.[7] This assumption is analogous to so-called *general position* which, as usually considered, means the viewpoint is not misleading—that the image is taken from a representative viewpoint. That assumption leads to many specific consequences: a curve that is straight in the image is straight in 3-D, lines that are collinear or parallinear or parallel in the image are collinear or parallel in 3-D, and so forth. Analogously, if the physical curves are assumed to lie in general position on the surface, we assume that a given curve of a given type is representative of those in its immediate vicinity. This also leads to a number of specific consequences, e.g., if two lines of curvature are parallel we assume that nearby lines of curvature are also parallel.[8] In practice this form of general position dovetails with the conventional form involving viewpoint (see Section 3.2.2). Superficially, general position of contour placement amounts to an expectation for surface smoothness, but more than that, it allows us to infer that any geometric property which holds along a particular curve (across a given surface) also holds in the vicinity of that curve.

3.1. Planar/curved

How can one determine whether or not a surface is curved when the only evidence is a curve C in the image, the projection of some physical curve Γ on some surface Σ? Observe that this cannot be determined simply on the basis of whether C is straight or curved—C being straight does not imply the surface is planar (or that it is cylindrical, or even that it is developable). Likewise, the mere fact that C is curved does not mean the surface is curved. The physical

[7] We must be careful with concepts such as 'displaced slightly' and 'qualitatively the same'. For now, these notions should seem intuitive; they are used more rigorously in Section 3.2.2.

[8] This implies that the underlying surface is cylindrical (see Section 3.2.2), a surprisingly strong consequence.

curve might be drawn on a planar surface, as are the figures on the pages of this journal.

We consider a single contour, for there are instances where that is sufficient. If we are to use the curvature of C to infer curvature of the surface Σ there are two hurdles: (i) inferring whether the physical curve Γ is curved or straight on the basis of the curvature of C, and (ii) inferring whether the surface is curved or planar from curvature of the physical curve Γ.

The first inference is straightforward, for in general we have that C is curved if and only if Γ is curved. The only exception is where C is a straight line but Γ is curved and planar and its curvature is hidden from view because the plane containing the curve is foreshortened into a line. This misleading situation is avoided by restricting the viewpoint to be in general position.

The second inference is more difficult. To use the curvature of Γ to tell us whether the surface is curved requires that we restrict the relation between the physical curve and the surface. To understand the restrictions that must be imposed, it is important to see when curvature (or lack of curvature) of Γ falsely implies curvature (or lack of curvature) of Σ. We have two cases to consider.

First suppose Γ is curved but Σ is actually planar. Then none of the curvature of Γ is normal curvature. In other words, the principal normal of Γ lies everywhere in the plane of Σ. In that special situation Γ is called *asymptotic*.[9] The other case is where Γ is straight but Σ is curved. When can this occur? (Since one cannot imbed a straight line on a surface of positive Gaussian curvature, we need only consider cases involving surfaces of either negative or zero Gaussian curvature.) There are special doubly curved surfaces (the hyperboloid of one sheet and the hyperbolic paraboloid) on which one may place a straight line. But of greater interest to us is the common occurrence of straight lines on developable surfaces (those of zero Gaussian curvature). The lines of least curvature on a developable surface are straight lines (see Figs. 2b and 8a), so this case would mislead us because the curve C would be straight despite the fact that the surface Σ is curved.

In summary, the two cases where the deduction 'C curved iff Σ curved' would fail are (i) when Γ is curved and asymptotic, and (ii) when Γ is straight and a line of least curvature. Both cases occur sufficiently often in real scenes that they must be seriously regarded. The first case arises, for instance, in water lapping on a beach, wood grain on a table top, mottle shadows on the ground, pigmentation markings on a relatively planar surface, and so forth. The second case arises in shading contours and glossy reflections on cylindrical surfaces (the 'terminator' or 'self-shadow' edge is straight and parallel to the axis of the cylinder, as is any specular reflection along the cylinder) and virtually all man-made objects have linear markings (seams, pigmentation edges, or what-

[9] The asymptotic case is discussed again in Section 4.1.

ever) which are rulings. The only situation which we may disregard as improbable concerns the rulings on a saddle surface as noted earlier.

The minimal restrictions necessary for distinguishing planar versus curved may be stated as follows: if Γ is curved, it cannot be asymptotic, and if straight, it cannot be a line of least curvature. We can accomplish the same thing by restricting Γ to a line of greatest curvature, but that is a much stronger restriction. Importantly, we will have independent motivation for this restriction momentarily. The geometric reasoning for distinguishing planar versus curved is summarized below. Observe that while we only need '$\Rightarrow$' for our purposes, we actually have the stronger biconditional '$\Leftrightarrow$'.

---

given a curve C, the geometric restriction

Γ is neither a line of least curvature nor an asymptotic curve

or the stronger restriction:

that Γ is a line of greatest curvature

plus general position (of viewpoint) allows the inference

C curved $\Leftrightarrow \Gamma$ curved $\Leftrightarrow \Sigma$ curved.

3.2. Sign of Gaussian curvature

We now examine geometric constraints that would allow one to determine whether a patch of surface has zero, positive, or negative Gaussian curvature. In order to reason about Gaussian curvature one clearly needs more than a single curve. The two cases that we will consider are intersecting curves and parallel curves.

3.2.1. *Intersecting contours*

We start with the intersection, illustrated in Fig. 2. Note that these four cases are exhaustive: either we have that both contours are straight (Fig. 2a) or one is straight and the other curved (Fig. 2b) or they are both curved with the two senses of relative curvature in the image (Figs. 2c and 2d).

To infer the sign of the Gaussian curvature the basic problem is to determine the signs of the two principal curvatures. The only input, remember, is the intersecting pair of curves C_1 and C_2 in the image. The corresponding physical curves Γ_1 and Γ_2 must therefore have a known relationship to the lines of curvature at that point. To illustrate this difficulty, Fig. 3 shows a cylinder with two choices of intersecting curves across it. In Fig. 3a the two curves are lines of curvature, and since the surface is developable, one of the lines of curvature is a straight line. In Fig. 3b both curves lie at an angle to the principal directions, and consequently both are curved. In Figs. 3c and 3d the outlines of the surfaces are removed, showing only the two curves. Note that in Fig. 3d the

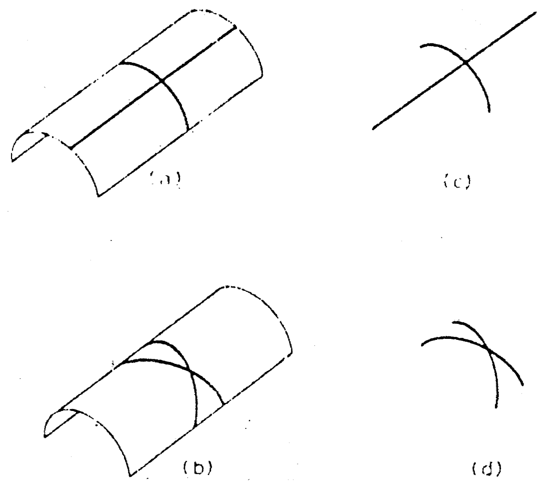

FIG. 3. The cylinder in a has two intersecting curves across it that are lines of curvature. Note that the straight line is a line of least curvature. The same cylindrical surface is shown in b, but this time the two intersecting curves do not lie in the principal directions, and therefore both are curved. In c and d the outline of the cylinder is removed, leaving only the intersecting curves in each case. The interpretation of c is the same as before, but d appears locally convex (like a sphere).

curves appear to lie on a convex surface. There is no information to indicate the special angular relation of the curves to the principal directions, and hence the relation between their curvature and the principal curvatures cannot be deduced. Given the goal of deducing the sign of the Gaussian curvature at an intersection of two curves in an image, there is apparently no feasible alternative to assuming that the corresponding physical curves are lines of curvature.

With the restriction of both physical curves to being lines of curvature, one may successfully deduce the sign of the Gaussian curvature by examining the curvature of the image curves at the intersection. In Fig. 2a the two lines are straight in the image, hence straight in 3-D, hence the surface is planar. In Fig. 2b one of the lines is straight, hence the surface has zero Gaussian curvature (it is developable, but we cannot further infer the surface is a cylinder—it may twist in space like a ribbon). In Fig. 2c the two lines have curvature of the same sign,[10] therefore the physical curves have the same sign[11] as well, and the

[10] There are many ways one may compare the sign of contour curvature in these intersection configurations. One approach is to proceed away from the intersection on two arcs, one from each curve, and compare how their normals rotate (whether they rotate clockwise or counterclockwise in the image). Note that one may attend to either of two pairs of arcs: those that bound the obtuse angle and those that bound the acute angle. If we take the arcs that define the acute angle of intersection and proceed on each away from the intersection, we will say that the curves have the same curvature if both normals rotate the same way.

[11] The sign of contour curvature in 3-D is with reference to the surface, as is customary. The two curves have the same sign of curvature if their principal normals are on the same side of the surface.

surface has positive Gaussian curvature (is locally convex). Finally, in Fig. 2d the two lines have opposite curvature and the surface has negative Gaussian curvature (is hyberbolic). The geometric reasoning is summarized below.

given intersecting contours C_1 and C_2, the geometric restriction
that Γ_1 and Γ_2 are lines of curvature
plus general position (of viewpoint) allows the inference

C_1 or C_2 straight	$\Leftrightarrow$ Γ_1 or Γ_2 straight	$\Leftrightarrow$ Σ developable
C_1 and C_2 have same curvature sign	$\Leftrightarrow$ Γ_1 and Γ_2 have same curvature sign	$\Leftrightarrow$ Σ locally convex
C_1 and C_2 have opposite curvature sign	$\Leftrightarrow$ Γ_1 and Γ_2 have opposite curvature sign	$\Leftrightarrow$ Σ hyperbolic

3.2.2. *Parallel contours*

Now we will examine parallel contours.[12] If two curves are parallel in the image and in general position relative to the viewer, the corresponding physical curves are also parallel in 3-D. What then can this tell us of the Gaussian curvature? Strictly speaking, the surface may have arbitrary shape[13] between the two parallel curves, so there are no inevitable consequences of parallelism. But 'if the surface is 'well behaved', as discussed earlier, then a strong restriction on the shape ensues. Specifically, if the placement of the physical curves on the surface is not critical (that if displaced slightly they would remain parallel), the surface is a *cylinder*. That is to say, a cylinder is the only surface in which one may embed parallel curves, in general. Note that one may find special situations that violate one or the other form of general position—either non-parallel curves on some non-cylindrical surface which look parallel from a particular viewpoint, or parallel curves on some non-cylindrical surface which are critically placed on the surface—but these are arguably improbable occurrences in nature. Hence these two forms of general position, of viewpoint and of placement of the curves on the surface, together allow one to infer the surface is a cylinder wherever the contours are parallel.

A word of caution is needed, however, regarding the practical definition of 'parallel'. For example, two lines of latitude on a globe are close to parallel, especially if they are closely spaced—and yet the surface is convex, not

[12] Two arbitrary curves are parallel if one may be superimposed onto the other by merely a translation.

[13] Actually, some restriction on the shape is imposed if the surface is opaque. Since the two curves are visible along their length, at no point can the intervening surface lie between the viewer and the physical curves. But this restriction is quite weak in terms of constraining the surface shape.

cylindrical. The cylinder deduction is valid only when the curves are precisely parallel. (And unfortunately, the deduction does not degrade gracefully: in the example just given even though the contours would appear roughly parallel, the surface would not be roughly cylindrical.) The geometric reasoning is summarized below.

---

given parallel contours C_1 and C_2 the geometric restrictions
that Γ_1 and Γ_2 are in general position on the surface and are *not* asymptotic
plus general position (of viewpoint) allows the inference
$C_1 \| C_2 \Leftrightarrow \Gamma_1 \| \Gamma_2 \Leftrightarrow \Sigma$ cylindrical.

3.3. Surface orientation

We have examined some geometric restrictions for determining whether a surface is planar versus curved, and for determining the sign of the Gaussian curvature. The final step is to consider surface orientation. As before, we will consider two cases where we can gain information about a patch of the surface: intersections and parallel contours.

3.3.1. *Intersecting contours*

The intersection in Fig. 4a will be the focus of our attention for the moment. The crucial measurement that we will use is the angle of intersection β measured in the image (Fig. 4b). The normal to the surface at that point is

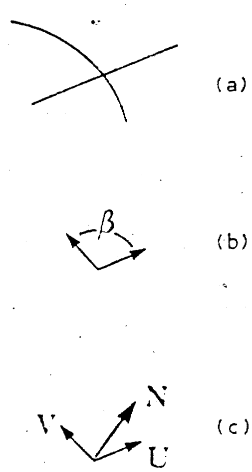

FIG. 4. The intersecting contours in a may be analyzed in terms of surface orientation given certain restrictions. The angle β shown in b is measured in the image plane. In c the three-dimensional vectors U and V lie on the surface, and N is the surface normal at their intersection.

shown in Fig. 4c. How might it be determined? We will find that the solution builds on the geometric restrictions established in the previous discussions.

Local surface orientation has two degrees of freedom, but the restriction of the physical curves to be lines of curvature reduces the degrees of freedom to one. To show this, we will derive a vector expression for the surface normal. Construct two vectors U and V that correspond to the tangents to the two physical curves at the intersection. For convenience they are constructed so as to project as unit vectors (Fig. 4c) and, without loss of generality, the x-axis is oriented with the projection of U. The angle β is measured between the projections of U and V.

$$U = \{1, 0, a\} \quad \text{and} \quad V = \{\cos\beta, \sin\beta, b\}.$$

The quantities a and b are the unknown components of U and V along the z-axis (perpendicular to the image plane). The surface normal N is the cross product:

$$N = \{-a\sin\beta, a\cos\beta - b, \sin\beta\} \tag{1}$$

and tilt τ and the slant σ are:

$$\tau = \tan^{-1}\frac{N_y}{N_x} \quad \text{and} \quad \sigma = \cos^{-1}\frac{N_z}{(N_x^2 + N_y^2 + N_z^2)^{1/2}} \tag{2}$$

where N_x, N_y, and N_z are the three components of the normal vector. Observe from (2) that the slant is the angle between the normal vector and the view vector and the tilt is the orientation to which the normal would project.

The expression for the normal in (1) carries two unknowns. This reflects the two degrees of freedom of surface orientation when no restrictions are imposed. Now, if the intersecting physical curves are lines of curvature, they are perpendicular at the intersection. Hence the dot product of U and V is zero, from which we have

$$b = -\frac{\cos\beta}{a}$$

which when substituted into (1) gives

$$N = \left\{-a\sin\beta, \cos\beta\frac{a^2+1}{a}, \sin\beta\right\} \tag{3}$$

and therefore the surface orientation is determined up to only one unknown, a.

Perpendicularity of the two physical curves therefore removes one degree of freedom of surface orientation. The constraint can be expressed geometrically in terms of slant and tilt, parameterized by the angle of intersection β in the image. For a given angle β to correspond to the projection of a perpendicular intersection, the intersection could only have certain orientations in 3-D

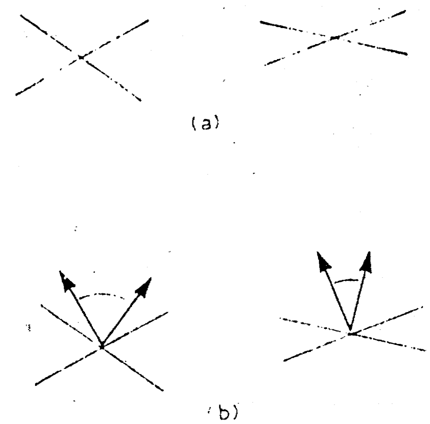

FIG. 5. A perpendicular intersection in 3-D is foreshortened in the image (a) in general. The projected angle of intersection in the image increasingly constrains the range of tilt as the angle approaches π.

relative to the viewer.[14] These limits are shown in Fig. 5. Observe that if β is near π, the foreshortening of the 3-D intersection must be large (Fig. 5a), but if β is near $\frac{1}{2}\pi$ one can say little about the surface orientation. This restriction has a simple geometric interpretation in terms of tilt, as shown in Fig. 5b. Suppose the tangents to the curves in the image plane have orientations τ_1 and τ_2 at their intersection. Then the surface tilt τ must lie within the perpendiculars to these tangents:

$$\tau_1 + \tfrac{1}{2}\pi \leq \tau \leq \tau_2 - \tfrac{1}{2}\pi.$$

Fig. 6 graphs this restriction and the corresponding restriction on slant, given β is a foreshortened right angle. For example, at $\beta = 135°$ the slant is $77.75 \pm 12.25°$, and the tilt is the bisector $\pm 22.5°$. And at $\beta = 150°$ the slant is $82.25 \pm 7.75°$, and the tilt is the bisector $\pm 15.0°$.

In conclusion, by restricting the physical curves to be lines of curvature, we can place bounds on the surface orientation at the intersection. For instance, the surface tilt at each intersection in Fig. 7a must lie within the limits shown in Fig. 7b. This restriction also places bounds on slant as a function of β.

Thus far we have only attended to the point of intersection of two surface contours, and have only utilized the angle of their intersection. Intuitively, it seems that we should be able to use the more global shape of the surface contour in order to determine surface orientation. But to do this requires substantially stronger geometric restrictions.

[14] We will ignore the reversals in the *direction* to which the normal would project due to depth reversals. This ambiguity is inherent in the orthographic projection and will not be considered here (see [15]).

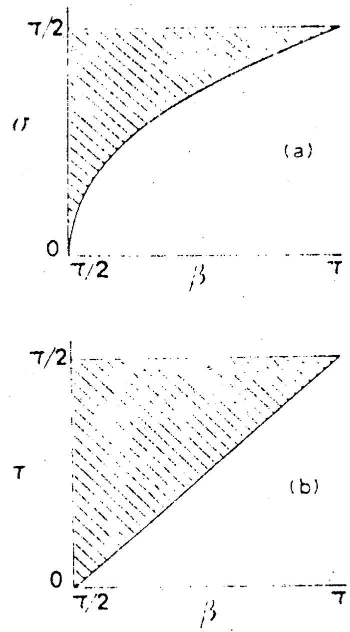

FIG. 6. The constraint on slant σ (shown in a) and on tilt τ (shown in b) is graphed as a function of the obtuse angle of intersection β, which ranges from $\frac{1}{2}\pi$ to π.

The problem of computing the surface orientation at a point along a curve on the basis of its 2-D projection is exceedingly underconstrained. It will be worthwhile discussing the problem informally for a moment. The three independent factors that enter into the problem are the viewpoint, the shape of the physical curve, and the behavior of the surface under the curve. Visualize the physical curve as a wire in 3-D and the surface along the curve as a thin ribbon that is glued to the wire. While the shape of the wire must be such that it projects to the given contour in the image, there are infinitely many bendings of the wire that would project identically. Furthermore, the wire may have

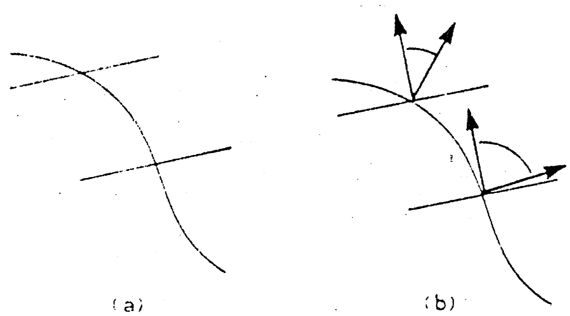

FIG. 7. The curve in a is intersected at two places. Assuming the intersections are perpendicular on the surface, the tilt at each point is constrained to lie somewhere within the bounds shown in b.

torsion, and the ribbon, which represents the strip of surface under the curve, may twist along the length of the wire arbitrarily.

Using the wire-and-ribbon analogy, let us rephrase the problem of solving surface orientation along a contour. One would start with an image of a wire (whose specific shape is not known *a priori*), and with information about how wires twist and curve in general and information about how ribbons are glued onto wires in general, determine how the particular ribbon would appear at each point along the wire from the particular viewpoint. It would amount to bending a wire appropriately then gluing the ribbon along its length in some manner, and holding the finished construction in some specific orientation in space relative to you. We discussed the bounds on slant and tilt afforded by a perpendicular intersection in 3-D. The intersection may be visualized as a line drawn across the ribbon so that it is perpendicular to the wire. Then, given an image which shows some angle β (Fig. 4b) one knows roughly how to hold the ribbon, but only at the point of intersection—the wire and ribbon would be free to bend and twist away from that point.

It is important to stress that the information about wires and ribbons must be specific enough to allow one to solve the problem, but general enough to be useful. We now will consider two geometric restrictions that together meet these criteria: that the wire lies in some plane and the ribbon is everywhere perpendicular to that plane. More formally, we restrict the physical curve Γ to be planar and geodesic. Our current goal is to show why these restrictions solve the problem; the motivation for proposing the geodesic and planar properties will be evident momentarily.

Suppose the physical curve Γ is planar, i.e., it lies in some plane Π. Then the problem of determining the 3-D shape of Γ reduces to determining the orientation of Π relative to the viewer—Γ is then simply the projection of the given image curve C back onto Π. A method for estimating the 3-D orientation of a planar curve (for example an extension of the method developed by Witkin [18]) may then be used to solve Γ. Now, given the 3-D shape of Γ, the fact that it is geodesic means the surface normal is identically the principal normal of Γ at each point. The geodesic and planar restrictions are clearly sufficient to solve the surface orientation along the curve. Now we should consider where they arise.

3.3.2. *Parallel contours*

Earlier we established that where contours in the image are parallel, the surface is locally a cylinder, subject to some assumptions of general position. This cylinder property, in conjunction with the restrictions of the physical curves to be lines of curvature (specifically lines of *greatest* curvature) will give us immediately the two geometric restrictions that we seek, because *lines of curvature across a cylinder are planar and geodesic*. The reasoning is summarized below.

for parallel contours C_1 and C_2, the geometric restrictions

that Γ_1 and Γ_2 are in general position on surface and are *not* asymptotic

plus general position (of viewpoint) allows the inference

$C_1 \| C_2 \Leftrightarrow \Gamma_1 \| \Gamma_2 \Leftrightarrow \Sigma$ cylindrical

and

Σ cylindrical and Γ line of curvature $\Rightarrow \Gamma$ planar and geodesic.

A complementary approach will now be considered, which deals geometrically with pairs of parallel contours.

The lines of least curvature, or *rulings* on a cylinder are parallel, straight lines, and are perpendicular to the lines of greatest curvature. Their projection would be parallel straight lines, but because of foreshortening they would no longer be perpendicular to the projected lines of greatest curvature (Fig. 8a). Nonetheless, a given ruling would intersect successive lines of greatest cur-

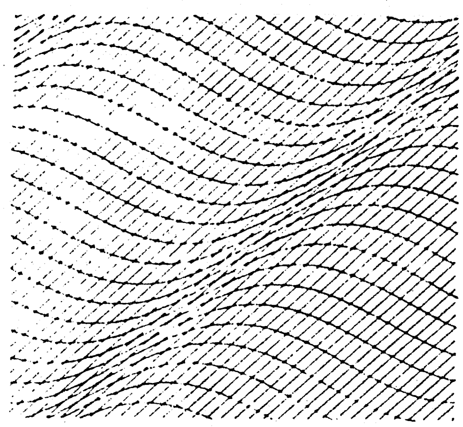

(a)

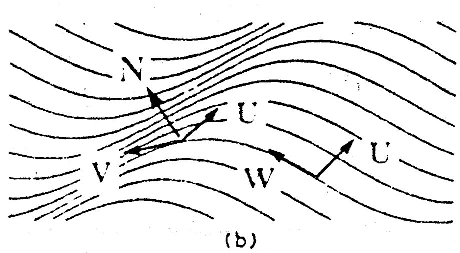

(b)

FIG. 8. In the orthographic projection of a cylinder (a) the lines of least curvature project as straight and parallel, and each intersects successive contours at a constant angle. In b the vectors at two intersections are shown. In the text it is shown how the constraint can be propagated along the curve.

vature at a constant angle. This fact allows us to reconstruct how the rulings would project in the image. We identify points on adjacent contours with parallel tangents, and connect those points with straight lines that are themselves parallel. This may be thought of as bringing points on adjacent contours into *parallel correspondence*. The line that connects corresponding points is the image of a ruling. The correspondence is unique in general. With reference to Fig. 9a, note that where the contours are straight the tangent to a point P on one contour would be parallel to various tangents on the adjacent contour, but only one choice would result in a correspondence line that is parallel to the correspondence lines elsewhere (Fig. 9b). With the rulings reconstructed in the image (as in Fig. 8a), we have at each intersection an angle β which corresponds to a foreshortened right angle, since lines of curvature are perpendicular. Thus we can place bounds on slant and tilt at each intersection, as discussed earlier. But we have another important dividend which stems from the fact that the rulings are perpendicular to the plane containing the lines of greatest curvature. To see this, consider the vector constructions in Fig. 8b.

The 3-D vector $\boldsymbol{U}$ is collinear with a given ruling, and $\boldsymbol{V}$, with a line of greatest curvature. Being perpendicular, the surface normal $\boldsymbol{N}$ at the intersection is their cross product. We wish to define the spatial orientation of the plane Π containing Γ. Since Γ is geodesic, the two vectors $\boldsymbol{V}$ and $\boldsymbol{N}$ define that

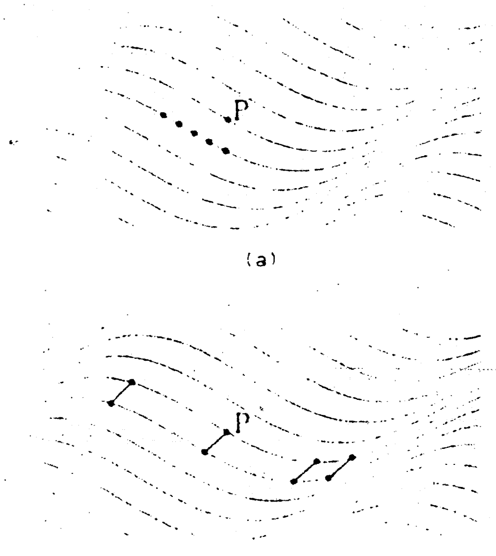

FIG. 9. Where the surface contours are straight, as in a, the tangent to a point P on one contour would be parallel to various tangents on the adjacent contour, however only one choice would result in a correspondence line that is parallel to the correspondence lines elsewhere.

plane. Furthermore, since U, V, and N are mutually orthogonal, U is the normal to the plane Π. This is important, because the orientation of the ruling immediately gives us the *tilt* of Π.

Earlier we lamented that somehow the shape of the curve in the image should help solve the surface orientation. Now we see how to do this. The planarity and geodesic restrictions allow us to propagate the surface orientation *along* a curve, from places where it is strongly constrained to places where it is not. Consider the two intersections in Fig. 8b.

$$U = \{1, 0, a\}, \quad V = \{\cos\beta_1, \sin\beta_1, b\}, \quad \text{and} \quad W = \{\cos\beta_2, \sin\beta_2, c\}.$$

where

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

The normals at the two intersections are

$$N_1 = \left\{-a\sin\beta_1, \cos\beta_1 \frac{a^2+1}{a}, \sin\beta_1\right\} \quad \text{and}$$

$$N_2 = \left\{-a\sin\beta_2, \cos\beta_2 \frac{a^2+1}{a}, \sin\beta_2\right\}. \tag{4}$$

Observe that β_1 is large, and therefore the slant and tilt are strongly constrained at that point. We take the bisector of the range of tilt as a best estimate[15] of τ in order to solve for the unknown a. Since

$$\tau = \tan^{-1}\frac{N_y}{N_x} = \tan^{-1}\frac{\cos\beta_1 \dfrac{a^2+1}{a}}{-a\sin\beta_1}$$

we have that

$$a = \left(\frac{-1}{\tan\tau\tan\beta_1 + 1}\right)^{1/2}. \tag{5}$$

Now, for some other point where the surface orientation is not at strongly constrained, such as at the other intersection (where β_2 is smaller) we may substitute (5) and (4) to find that the normal N_2 is determined completely.

Note that the parallel correspondence process which reconstructs the rulings is local, and therefore might be applied to the images of surfaces that are not cylinders globally, such as Fig. 10 and even Fig. 1. This suggests a computation for generating local cylinder approximations to doubly-curved surfaces.

[15] Incidentally, for a given intersection the bisector choice results in the surface orientation with the least slant.

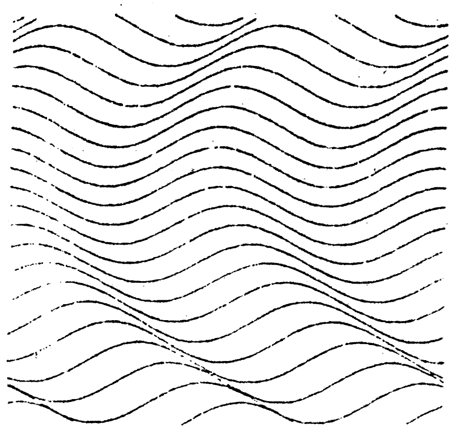

FIG. 10. Parallel correspondence is defined only locally, hence is applicable to surfaces that are not cylinders. If the contours are locally parallel, the surface may be approximated locally as cylindrical.

4. Discussion

4.1. The distinction between asymptotic curves and lines of curvature

Earlier when we were determining the restrictions necessary to distinguish planar from curved surfaces, the asymptotic curve emerged as a special case because its curvature is not related to surface curvature. Recall that an asymptotic curve follows the direction of zero normal curvature on the surface. It may therefore be regarded as the antithesis of a line of curvature, which follows the direction of extremal normal curvature. Asymptotic curves only exist of surfaces of negative or zero Gaussian curvature, i.e., hyperbolic and developable surfaces.

Asymptotic curves exist on hyperbolic surfaces because the principal curvatures have opposite sign and so the normal curvature must pass through zero in some direction between the two principal directions. In fact, an asymptotic curve's tangent always bisects the two principal directions; it is therefore a rather special curve across a saddle surface and will be disregarded because it is virtually non-existent in actual situations (and non-intuitive and hard to visualize as well).

Asymptotic curves also occur on developable surfaces—they correspond to rulings. But a degenerate case is when the surface is planar, for any curve across the plane is asymptotic trivially, since there is no normal curvature in any direction on a plane. It was this case that we explicitly excluded in our analysis of surface shape. Nonetheless, asymptotic curves on planar surfaces may be used to estimate the orientation of the planar patches. Witkin [18] describes a method for estimating the direction and magnitude of the *foreshortening* a planar curve has undergone in projecting into the image. From measurements of contour curvature and the distribution of tangent orientation in the image one may estimate the tilt and slant of the plane containing the corresponding 3-D curve. This approach seems useful, for instance, with mottled shadows on the ground or the spots on a dalmatian, and many other natural situations involving curves on relatively[16] planar surfaces.

I would suggest that there is a meaningful distinction to be made between the analysis of asymptotic curves described by Witkin and the analysis of lines of curvature introduced here. Fig. 11 shows an interesting case where the human visual system interprets the contours as asymptotic, i.e., as lying on a flat surface. There are probably geometric criteria which govern whether the contours are interpreted as lines of curvature or as asymptotic curves.

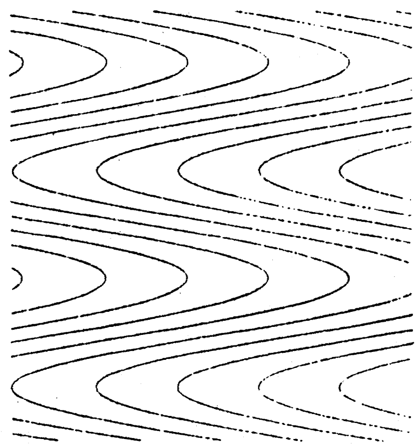

FIG. 11. The contours are interpreted as lying across a planar surface, i.e., as asymptotic curves.

4.2. Where do geodesics, lines of curvature, and planar curves occur in reality?

The previous discussions revealed three important geometric restrictions on physical curves: *line of curvature*, *geodesic*, and *planarity*. Where can we expect these restrictions to hold in actuality? A brief discussion will be given here.

The curves across the surfaces of virtually all synthetic objects are seldom arbitrary in the geometrical sense. In fact, they are almost always planar and geodesic lines of curvature. This is a consequence of the way we fabricate, label, and decorate objects. In particular, we tend to generate complicated surfaces piecewise out of developable surfaces (usually cylinders) and planes. The seam where these surfaces join is usually planar, and that plane is usually normal to the surface—hence the seam is both planar and geodesic. Moreover, for synthetic objects that are surfaces of revolution, the markings almost

[16] I.e. the curvature of the surface is small relative to the contour detail on the surface.

invariably follow the object's axis or are perpendicular to it. Importantly, lines on a surface of revolution that are parallel and perpendicular to the axis of rotation are lines of curvature and are planar. I am purposefully staying at a general level of discussion. The reader is invited to examine nearby objects to find specific examples. Without addressing the issue of nature versus nurture, it is clear that the fabricated objects which comprise our everyday environment have the geometric restrictions that we seek—the curves across their surfaces are almost invariably planar, geodesic, and lines of curvature.

One may also seek to find geometric restrictions on natural curves as well. Several interesting observations may be made, but the task is significantly more difficult, in general. As mentioned in the introduction, planar lines of curvature may be found in many biological forms: stripes on plant leaves, wrinkles, the joints on bamboo stalks (they are also geodesic). As also discussed earlier, it is difficult to weigh this evidence, for one may easily give counterexamples—curves that meander across the surface rather arbitrarily. The stripes on a zebra are one example.[17]

An interesting case of geometric restriction in nature is provided by the glossy reflections from surfaces that have a strong specular component in their reflectance functions. If a surface is specular, then given directional illumination and a favorable viewpoint, a specular reflection—either a point-like *highlight* or a *gloss contour*—will appear in the image. For this to occur some patch on the surface must be oriented such that the surface normal bisects the angle defined by the point light source, the given surface point, and the viewer. I will refer to this alignment as the *specularity condition*. If the surface is doubly curved, the specularity condition is met only at a point, if at all, and causes a point-like reflection in the image—a so-called highlight. Similarly, on developable surfaces a specular reflection would be point-like in general—the important exception being a cylinder. On a cylinder, the fact that the tangent plane does not twist along a ruling means that if the specularity condition is met at some point, it will also be met along the ruling passing through that point. Furthermore, the specularity condition is strong enough that if a given gloss contour C (produced by specular reflection) is straight, the surface is locally a cylinder:

$$C \text{ straight} \Leftrightarrow \Sigma \text{ cylindrical.}$$

While this is true in the case of orthographic projection with a distant viewer and distant point light source, in reality the viewer is often near the surface and the light source is neither distant nor a point source. The most important

[17] Incidentally, that may be one reason why they are effective in obscuring the true shape of the animal. It is well-known that one technique that nature seems to have adopted for protective coloration is to place high contrast contours across the animal which do not correspond to the underlying 3-D shape. In our terms, those contours would be non-planar, and would not be lines of curvature or geodesic.

exception concerns an extended light source, such as a bright window or a ceiling light panel. Instead of a tiny point or thin line of specular reflection, the reflection will be extended. Nevertheless, if the outline of the reflection is straight, the surface is cylindrical.

Can we infer anything about surface shape in other situations? For instance, if the gloss contour is an arc, rather than straight, what does that tell us? Generally it means the surface is doubly curved, and the curvature across the arc is much greater than the curvature along it (in other words, the surface may be approximated locally as a cylinder along the arc.)[18] To be more analytic would be quite difficult: the arc is the projection of some path along which the specularity condition is met, but if the viewer is relatively near the surface, and especially if the light source is also near by, that condition changes from point to point. Without knowledge of the viewer (and illumination) geometry, the specular reflection cannot be interpreted further. I suggest that images of specular surfaces are not feasibly analyzed by any analytic 'shape-from-shading' method. Rather, only the qualitative shape of the gloss contours are used in order provide rough information about the local Gaussian curvature, and only with the additional constraint afforded by the smooth boundaries can one feasibly compute local surface orientation. But it is clear that specular reflections can tell us not only something of the reflectance properties of the surface (that the surface is specular [1]), but also some qualitative information about the surface shape.

4.3. Summing up

This article has introduced an approach towards understanding how surface contours may be used as information about surface shape. Rather than start with the 'physics of the situation', which would be intractably difficult in its generality, we started with the basic geometric reasoning required to make 3-D shape assertions of various sorts. A range of assertions were studied: whether the surface is planar versus curved, the sign of the Gaussian curvature, and finally, the local surface orientation. Even the simple task of distinguishing planar from curved required some strong geometric restriction. The weakest restriction is that the physical curve is neither asymptotic nor a line of least curvature. A somewhat stronger restriction is that the curve is a line of greatest curvature. Next we saw that to conclude anything about the sign of Gaussian curvature, the physical curves must be lines of curvature. (Actually, the special case of parallel contours requires only that the corresponding physical curves be in general position on the surface and not be asymptotic—then parallelism implies the surface is a cylinder.) We then looked at the problem of computing

[18] One often resorts to physical movement relative to the surface in order to sort out the various reflections. Those reflections that stay fixed relative to the surface (or displace only slightly) correspond to places of high curvature, i.e., corners.

surface orientation and found that the perpendicularity of lines of curvature places useful constraint on the possible surface orientation at the intersection. Furthermore, it reduces the degrees of freedom of surface orientation to one. And provided that the curve is geodesic and planar (which is the case if the curves are parallel in general), we found that one may propagate the surface orientation from places where it is strongly constrained or known outright, to places where it is not. Thus parallelism emerged as very important in this analysis. It also emerged that the physical curve should be restricted to being planar, a line of curvature, and geodesic. Then we noted the strong distinction between lines of curvature and asymptotic curves. Finally we reflected momentarily on where these various restrictions occur, both in the man-made world, and in nature.

ACKNOWLEDGMENT

Primarily I must thank David Marr, who supervised the dissertation from which this article derives. Discussions with Shimon Ullman, Whitman Richards, Mike Riley, and Eric Grimson were most helpful. I also thank the two reviewers for their criticism. This article 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 N00014-75-C-0643 and in part by National Science Foundation Grant MCS77-07569.

REFERENCES

1. Beck, J., *Surface Color Perception* (Cornell University Press, Ithaca, 1972).
2. Brady, M., Finding the axis of an egg. *Proceeding Sixth International Joint Conference on Artificial Intelligence* (1979) 85–87.
3. Bruss, A., The image irradiance equation: its solution and application, Ph.D. Thesis (Electrical Engineering and Computer Science), M.I.T., forthcoming.
4. Hilbert, D. and Cohn-Vossen, S., *Geometry and the Imagination* (Chelsea, New York, 1952).
5. Hoffman, D.D., Inferring shape from motion, *Biological Cybernetics* (submitted for publication). Also available as *M.I.T. A.I. Lab. Memo* 592 (1980).
6. Horn, B.K.P., Obtaining shape from shading information, in: P.H. Winston, ed., *The Psychology of Computer Vision* (McGraw-Hill, New York, 1975).
7. Mackworth, A.K., Interpreting pictures of polyhedral scenes, *Artificial Intelligence* 4 (1973) 121–137.
8. Marr, D., Early processing of visual information, *Philos. Trans. Roy. Soc. B*, **275** (1976) 483–524.
9. Marr, D., Analysis of occluding contour, *Proc. R. Soc. Lond. B*, **197** (1977) 441–475. Also available as M.I.T. A.I. Lab. Memo 372.
10. Marr, D., Representing visual information, A.A.A.S. 143rd Annual Meeting, Symposium on: Some mathematical questions in biology, February 1977, published in *Lectures on Mathematics in the Life Sciences* 10 (1978) 101–180. Also available as M.I.T. A.I. Lab. Memo 415.
11. Marr, D., *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (W.H. Freeman, San Francisco, in press).
12. Marr, D. and Nishihara, K., Representation and recognition of the spatial organization of three-dimensional shapes, *Philos. Trans. Roy. Soc. B* 200 (1978) 269–294.

13. Marr, D. and Poggio, T., From understanding computation to understanding neural circuitry, *Neuroscience Research Progress Bulletin* **15** (1977) 470–488. Also available as M.I.T. A.I. Lab. Memo 357.

14. Prazdny, K., Egomotion and relative depth from optical flow, *Biological Cybernetics* **36** (1980) 87–102.

15. Stevens, K.A., Surface perception by local analysis of texture and contour, M.I.T. A.I. Lab. Technical Report 512 (1980).

16. Thompson, W. D'Arcy, *On Growth and Form* (Cambridge University Press, Cambridge, 1961).

17. Ullman, S., *The Interpretation of Visual Motion* (MIT Press, Cambridge, MA, 1979).

18. Witkin, A.P., Estimating shape from texture, *Artificial Intelligence* **17** (1981) 17–45 [this volume].

19. Woodham, R.J., Reflectance map techniques for analyzing surface defects in metal castings, *M.I.T. A.I. Lab. Technical Report* 457 (1978).

Received March 1980; revised version received September 1980