MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo 822 January, 1985

Describing Surfaces

Michael Brady
Jean Ponce
Alan Yuille
and
Haruo Asada

Abstract: This paper continues our work on visual representations of three-dimensional surfaces [Brady and Yuille 1984b]. The theoretical component of our work is a study of classes of surface curves as a source of constraint on the surface on which they lie, and as a basis for describing it. We analyze bounding contours, surface intersections, lines of curvature, and asymptotes. Our experimental work investigates whether the information suggested by our theoretical study can be computed reliably and efficiently. We demonstrate algorithms that compute lines of curvature of a (Gaussian smoothed) surface; determine planar patches and umbilic regions; extract axes of surfaces of revolution and tube surfaces. We report preliminary results on adapting the curvature primal sketch algorithms of Asada and Brady [1984] to detect and describe surface intersections.

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, the Office of Naval Research under contract number N00014-77-C-0389, and the System Development Foundation. This work was done while Haruo Asada was a visiting scientist at MIT on leave from Toshiba Corporation, Japan, and while Jean Ponce was a visiting scientist on leave from INRIA, Paris, France.

1. Introduction

Recent work in Image Understanding [Ballard and Brown 1982, Brady 1982, Marr 1982] has centered on the development of modules that compute three-dimensional depth, or depth gradients. Such modules include: stereo [Grimson 1981, 1984, Mayhew 1983]; shape from shading [Ikeuchi and Horn 1981, 1984] and photometric stereo [Woodham 1980, 1981]; shape from contour [Brady and Yuille 1984a, Witkin 1981, Barnard 1983, Barnard and Pentland 1983]; shape from motion [Bruss and Horn 1983, Ullman 1978]; and shape from texture [Vilnrotter, Nevatia, and Price 1981]. In applying vision to robotics, range finding and structured light have been investigated as techniques for recovering depth directly [Agin 1972, Holland, Rossol, and Ward 1979, Bolles, Horaud, and Hannah 1984, Faugeras et. al. 1982, 1984, Porter and Mundy 1982, 1984, Tsuji and Asada 1984] Although the work referred to above is largely experimental, it is clear that robust, efficient, practical three-dimensional vision systems will soon be available.

Several authors have suggested that the output of these "shape from" processes is a representation(s) that makes explicit the surface depth $z(x, y)$ or the local surface normal $\mathbf{n}(x, y)$. Barrow and Tenenbaum [1978] call such a representation an "intrinsic image", which emphasises that the representation has the same format as an image. Just as applications of two-dimensional vision depend upon the development of rich representations of shape (for example, Brady and Asada [1984]), computed from an image, so will applications of three-dimensional vision. This paper is concerned with the geometrical basis of such a representation. In particular, a set of curves are isolated that lie upon the surface and enjoy a global property, for example being planar. The structure of the representation is currently under development.

The representation that we are developing is based on the concepts of differential geometry, principally because it provides a hierarchy of increasingly stringent surface descriptions. A surface may simply be doubly curved, but it may be ruled, even developable, even conical. Our aim is to find the most stringent descriptors for portions of a surface. If, for example, there is a region of umbilic points, that indicates that part of the surface is spherical, then it is made explicit, as is the center of the corresponding sphere (Figure 1). If there is a portion of the surface that is part of a surface of revolution, it is described as such, and the axis is determined (see Figures 1 and 19).

Similarly, if there is a line of curvature or an asymptote that is planar or whose associated curvature (principal curvature or geodesic curvature respectively) is constant, then it is made explicit. For example, the asymptote that marks the smooth join of the bulb and the stem of the lightbulb in Figure 1, as well as the surface intersections marked on the oil bottle in Figure 21, are noted in the representation. We may associate a description with a curve that is a surface intersection; but only if it has an important property such as being planar. For example, a slice of a cylinder taken oblique to the axis of the cylinder produces a planar curve of intersection. On the other hand, the intersection of two cylinders is not a planar curve.

Figure 1 illustrates the representation we are aiming at. The stem of the lightbulb is determined to be cylindrical, because it is ruled and because it is a surface of revolution. We can compute the axis of the stem. The bulb is determined to be a portion of a sphere, because it is a connected region of umbilic points. The center of the sphere

can be computed. Similarly, the center of the sphere that forms the threaded end can be determined. The stem is smoothly joined to the bulb. Moreover, the axis of the cylindrical stem passes through the centers of the spheres defined by the bulb and threaded end. This distinguishes the diameters of each that are collinear with the stem axis, showing that the lightbulb is a surface of revolution. The geometrical primitives can be computed by the algorithms described below. We are working on the inference engine.

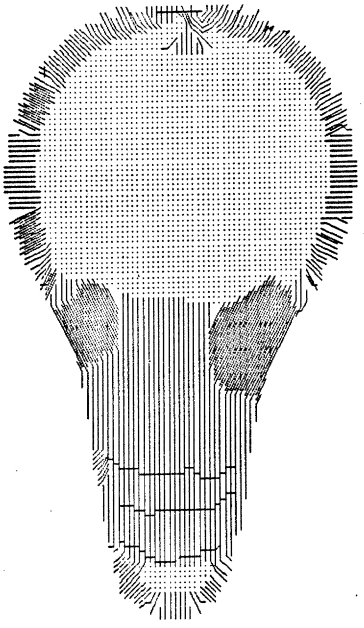

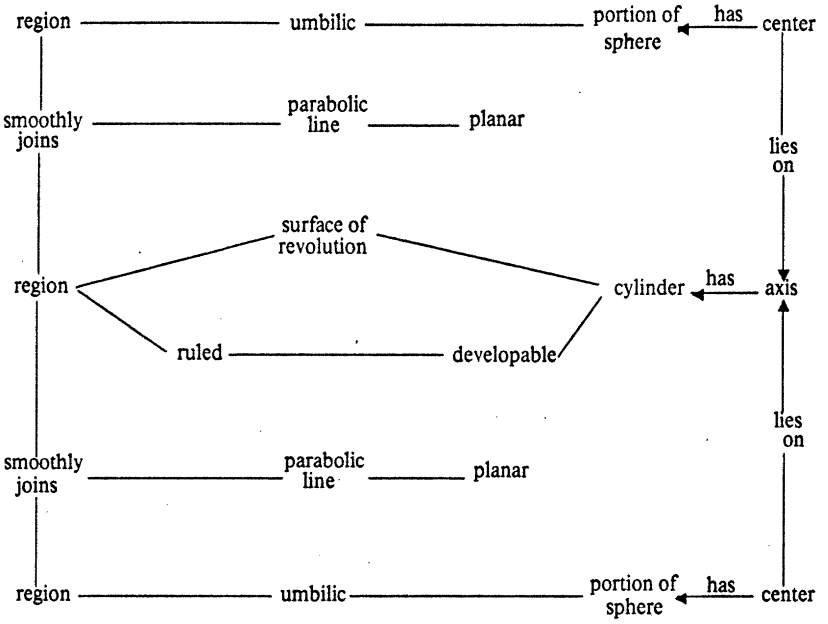

Figure 1. The representation of a lightbulb. a. The dotted region consists of umbilic points, indicating that the bulb is spherical. The parallel lines are the meridians of the cylindrical stem. The parallels, which are also rulings, are not shown. The significant surface changes are shown. They indicate the smooth join between the bulb and the stem and between the stem and the threaded end. b. The representation that we are working towards for the lightbulb.

In the next two sections we report recent theoretical developments and computational experiments. The theoretical component of our work is a study of classes of surface curves as a source of constraint on the surface on which they lie, and as a basis for describing it. We analyze bounding contours, surface intersections, lines of curvature, and asymptotes. Our experimental work investigates whether the information suggested by our theoretical study can be computed reliably and efficiently. We demonstrate algorithms that compute lines of curvature of a (Gaussian smoothed) surface; determine planar patches and umbilic regions; extract axes of surfaces of revolution and tube surfaces. We report preliminary results on adapting the curvature primal sketch algorithms of Asada and Brady [1984] to detect and describe surface intersections. The experimental results appear to contradict statements by a number of authors to the effect that second differential quantities such as lines of curvature cannot be computed reliably.

2. Surface curves

2.1. Introduction

A surprising amount can be learned about a surface from certain curves that lie upon it. Our theoretical work is an investigation of classes of surface curves as a source of constraint on the geometry of the surface and as a basis for description of the surface. The curves that we have studied are: bounding contours [Koenderinck and van Doorn 1982, Koenderinck 1984, Marr 1977, Barrow and Tenenbaum 1978, 1981, Binford 1982]; lines of curvature [Stevens 1982, 1983, Brady and Yuille 1984b]; asymptotes [Brady and Yuille 1984b]; and surface intersections [Binford 1982].

We begin by discussing (extremal) bounding contours of a surface, where the surface normal turns smoothly away from the viewer. First, we recall that the surface normal can be determined from its (orthographic) projection. Then we present a new proof of a recent result due to Koenderinck [1984] which states that the sign of the Gaussian curvature of the surface at points along the boundary curve is the same as the sign of the curvature of the projection of the boundary curve. We extend the result and show that, in general, points on the projection of an extremal boundary where the curvature is zero typically correspond to points on the surface that are locally flat (that is, where both principal curvatures are zero), or the boundary is locally aligned with a surface line of curvature whose curvature is zero. We suggest that these results may provide an alternative to line labelling for interpreting line drawings of curved objects.

The bounding contours of a surface are either extremal or mark discontinuities of some order C_n. Depth discontinuities of type C_0 typically occur at occluding boundaries. Most surface intersections are of type C_1. Smooth joins are C_2 discontinuities that can only be perceived if the curvature in a direction orthogonal to the join changes substantially. Asada and Brady [1984] discuss the analogous situation for planar curves. A Theorem of Joachimsthal [Weatherburn 1927, page 68] implies that surfaces rarely intersect along their lines of curvature. For example, the intersection of the two cylinders in Figure 2 is *not* a line of curvature of either surface.

The problem of detecting, localizing, and describing the discontinuities of a surface is analogous to computing the Primal Sketch representation for images [Marr 1976, Haralick et. al. 1983] and the Curvature Primal Sketch representation of the significant changes of curvature along planar curves [Asada and Brady 1984]. In Section 2.3, we show that Joachimsthal's Theorem can, under certain circumstances, be undone by Gaussian smoothing. More precisely, we shall prove that although two surfaces rarely intersect along a line of curvature, the intersection generates a line of curvature of the surface that results from smoothing with a Gaussian distribution:

Theorem Let $f(x, y, z)$ be a surface that is the cross product of a planar curve and a straight line. The lines of curvature of the convolution of f with a Gaussian distribution are in the plane of the curve and parallel to the generating line.

The restriction that the straight line be orthogonal to the planar curve is severe and implies that the surface is cylindrical. We are investigating ways to weaken the assumption as part of our development of the Surface Primal Sketch.

In Sections 2.4 and 2.5 we turn our attention to curves that lie in the interior of the visible portion of the surface, though they may intersect its boundaries. At each point of the surface, we define a set of *intrinsic directions* $\theta_i(x, y)$ in which the surface

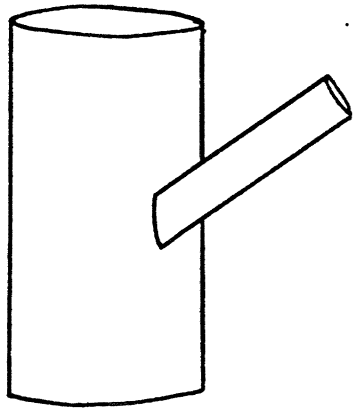

Figure 2. The curve that is the intersection of the two cylinders is not a line of curvature of either cylinder

change locally appears to be intrinsically important, and we associate a descriptor $\delta_i(x, y)$ with each such direction. The set of intrinsic directions that we have investigated so far consists of the directions of principal curvature, the asymptotic directions (directions in which the normal curvature is zero), and the parabolic directions, across which the sign of the Gaussian curvature changes. The corresponding descriptors $\delta_i(x, y)$ are the principal curvatures, and, for the asymptotes, the geodesic curvature. Currently, no descriptor is associated with parabolic directions. We may find it necessary to investigate additional intrinsically important directions in due course, though the directions of principal curvature, and asymptotic and parabolic directions suffice for a broad class of analytic surfaces that includes surfaces of revolution, ruled and developable surfaces, and generalized cones.

The directions and descriptors are *local* statements about the surface. For example, the directions of principal curvature give the locally flattest coordinate system intrinsic to the surface. We aim to describe the *larger scale* structure of the surface. Whereas local surface structure is thoroughly discussed in differential geometry, larger scale structure is not. To determine the larger scale structure of the surface, we link the local directions to form smooth curves. These curves are the lines of curvature, asymptotes, and parabolic lines of the surface, and they are discussed in Sections 2.4 and 2.5. We propose that these smooth curves are only made explicit when they satisfy additional constraints. Currently, these constraints are that either (i) the descriptors δ_i in the directions θ_i are (nearly) constant along the curve, or (ii) the curve is (nearly) planar. The need for constraints on the surface curves is illustrated in Figure 3, which shows the lines of curvature on an ellipsoid. Only three of the lines of curvature are planar, and they are the intersections of the symmetry planes of the ellipsoid with the surface. The surface is effectively described by these curves. We show how choosing a set of surface curves in this manner can automatically suggest a "natural parameterization" of a surface.

A surface curve is a geodesic if and only if its geodesic curvature is zero. Geodesics

form a mathematically important class of curves. However, as we discuss below, there are situations in which, perversely, there are too many geodesics or too few. For the moment, we note that requiring planarity is less severe than requiring that the curve be a geodesic. A geodesic is planar if and only if it is a line of curvature. It follows that if the geodesic curvature is zero along a line of curvature then it is planar. However, planar lines of curvature are not necessarily geodesics, hence do not have zero geodesic curvature.

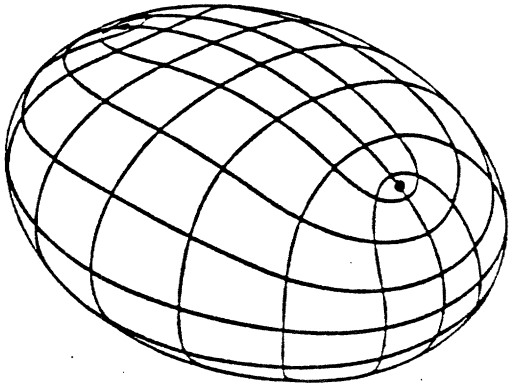

Figure 3. The lines of curvature of an ellipsoid. The only lines of curvature that are planar are the intersections of the surface with its planes of symmetry. Other lines of curvature do not seem to convey important information about the shape of the surface.

Section 2.4 is mostly given over to proving a theorem about generalized cones, a theorem that relates surface curves to a volumetric representation proposed by Marr [1977]. Marr considered generalized cones [Brooks 1981, Brooks and Binford 1980] with straight axes. He suggested that such a generalized cone is effectively represented by (i) those cross sections, called *skeletons*, for which the expansion function attains an extreme value; and (ii) the tracings, called *flutings*, for which the cross-section function attains an extremum. (A tracing is the space curve formed by a point of the cross section contour as the cross section is drawn along the axis.) Figure 4 illustrates these terms. We prove the following Theorem:

Theorem If the axis of a generalized cone is planar, and the eccentricity of the cone is zero, then (i) a cross section is a line of curvature if and only if the cross section is a skeleton; and (ii) a tracing is a line of curvature if the generalized cone is a tube surface (the expansion function is a constant), or the tracing is a fluting.

The cone's eccentricity at a point is the angle between the tangent to the axis curve and the normal to the cross-section through that point. It follows that the flutings and skeletons of a (planar axis) generalized cone are implied by both consideration of surface

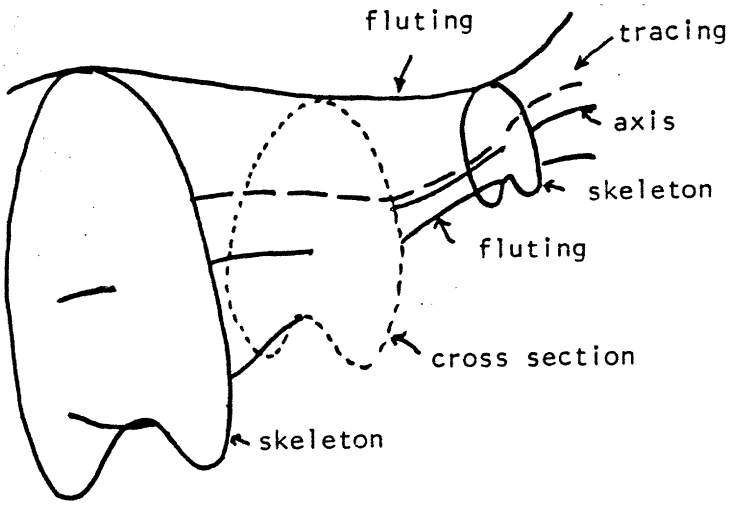

Figure 4. A generalized cylinder is formed by drawing a uniformly-expanding planar cross-section along a space curve.

curves and by Marr's considerations of volumetric representations. Lines of curvature that satisfy a large scale constraint, such as being planar, provide a natural basis for describing a surface. As is illustrated by the ellipsoid shown in Figure 2, and generalized cones more generally, there may only be a small number of lines of curvature that satisfy the large scale constraints. Coupled to a suitable scheme for surface interpolation (for example, Terzopoulos [1984]), they may provide a natural parameterization of the surface. Brady and Yuille [1984b] call such a parameterization a *curvature patch* representation.

As we shall show in Section 2.5, however, lines of curvature are a poor basis for describing many surfaces. In particular, ruled surfaces are seldom usefully described in terms of their lines of curvature, which typically fail to satisfy the large scale constraints. In such cases, the rulings are a better basis for description. The rulings of a developable surface are asymptotes, which are curves along which the normal curvature is zero. They are a useful basis for describing ruled surfaces. They are discussed in Section 2.4.

Section 3 reports some initial experiments on developing a *Surface Primal Sketch* representation of significant surface changes. Significant surface changes are found as follows. First, the surface is filtered with a Gaussian distribution and its lines of curvature are found. These are projected into a plane and input to Asada and Brady's Curvature Primal Sketch program. This isolates the significant curvature changes along the lines of curvature and describes each such change symbolically.

The final class of curves that we investigate are the parabolic curves. These are the smooth loci of parabolic points [Hilbert and Cohn Vossen 1952, pages 197 – 202]. Koenderinck's Theorem (re-proved in Section 2.2) suggests that parabolic points on extremal contours can be found. Figure 5 is reproduced from Hilbert and Cohn-Vossen [1952, page 197]. It shows the parabolic lines on the bust of the Apollo Belvidere. The footnote on page 198 of Hilbert and Cohn-Vossen has a familiar ring: "F. Klein used the parabolic

curves for a peculiar investigation. To test his hypothesis that the artistic beauty of the face was based on certain mathematical relations, he had all the parabolic curves marked out on the Apollo Belvidere, a statue renowned for the high degree of classical beauty portrayed in its features. But the curves did not possess a particularly simple form, nor did they follow any general law that could be discerned".

Figure 5. Bust of the Apollo Belvidere with superimposed parabolic lines. (Reproduced from [Hilbert and Cohn-Vossen 1952, page 197])

Consider, however, Figure 6a, which shows the computed regions of positive (white) and negative (black) Gaussian curvature on the surface of a lightbulb. The parabolic curves are planar and they mark smooth joins between significant parts, such as the bulb and stem. This suggests that, just like directions of curvature, a parabolic curve needs to satisfy additional large scale constraints in order to be made explicit by the visual system. The cross product (Figure 6b) of a straight line and a planar curve that is the smooth join of two curves whose (planar) curvatures have opposite signs (a smooth join in the terminology of Asada and Brady [1984]) is a simple surface that has a planar parabolic line. Asada and Brady point out that a smooth join is only perceivable if the difference between the curvatures of the flanking curves is sufficiently different. Few of the parabolic curves on the Apollo Belvidere have extended planar components. Those that do, such as the line on the forehead, are crossed by curves that do not have perceivable smooth joins. Two exceptions are the smooth join between the mouth and chin, and the smooth join between the nostril and the rest of the nose.

In general, it does not appear that the Gaussian curvature is a useful descriptor for a surface. The reason seems to be connected with the fact that the Gaussian curvature is a divergence expression [Courant and Hilbert 1953, p. 196]. This means that its Euler-Lagrange equation vanishes. A divergence term can be added to any variational problem, for example surface interpolation [Grimson 1981, Terzopoulos 1984, Brady and

Horn 1983], without changing the resulting surface. Representations based on Gaussian curvature typically suppress spatial structure.

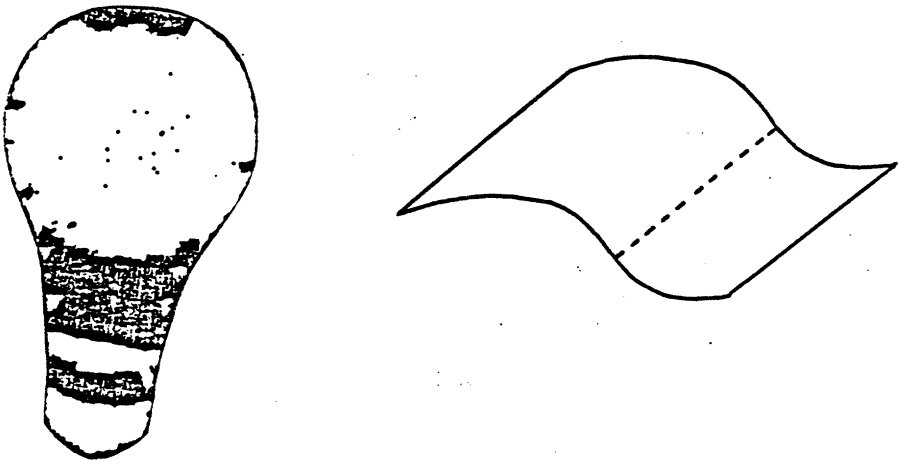

Figure 6. a. The parabolic curves on a lightbulb are the transitions from elliptic regions (white) to hyperbolic regions (black). b. The cross product of a planar curve with a smooth join and a straight line.

We omit analysis of certain classes of surface curve, particularly geodesics. There are at once too many geodesics, yet, in many cases, too few. There are too many since there is a geodesic through every point on a surface in every direction. The geodesics on a cylinder, for example, are all the helices $(R\cos\theta, R\sin\theta, k\theta)$, where R is the radius of the cylinder. On the other hand, only those meridians of a surface of revolution that are extrema of the surface width (skeletons, if it is thought of as a generalized cone) are geodesics. We have, moreover, proposed that to be made explicit, a surface curve must satisfy certain strong constraints. In particular, we would require a geodesic to be planar. However, if a geodesic is planar then it is a line of curvature [Weatherburn 1927, page 103].

2.2. Bounding Contours

Barrow and Tenenbaum [1981] propose that there are two distinct types of contour that bound a surface, which they call *extremal* and *discontinuity*. At an extremal boundary, the surface normal turns away smoothly from the viewer. A discontinuity boundary marks the abrupt termination of a smooth surface, perhaps to intersect another surface (see Barrow and Tenenbaum [1981, Figure 3.1]). Since a discontinuity boundary could be produced by the edge of a ribbon that is oriented in an arbitrary direction, it is difficult to infer a great deal about the flanking smooth surface solely from the image of a discontinuity boundary. Brady and Yuille [1984a], Barnard and Pentland [1983], and Barrow and Tenenbaum [1981] offer some suggestions.

Extremal boundaries are a rich source of information. First, suppose that a surface $f(x, y, z) = 0$ is orthographically projected onto a plane whose normal is $\mathbf{k}$. Suppose that the surface has an extremal boundary, then its boundary rim curve contains $\mathbf{k}$ and lies in the tangent plane of the surface. That is, $\mathbf{k} \cdot \nabla f = 0$, which can be written $\mathbf{k} \cdot \mathbf{n} = 0$, where $\mathbf{n}$ is the local surface normal.

Let the boundary curve $\mathbf{r}(s)$ lying on the surface have tangent vector $\mathbf{T}(s) = d\mathbf{r}/ds$. Now $\mathbf{n}$ is orthogonal to both $\mathbf{k}$ and $\mathbf{T}$, from which it follows that $\mathbf{n}$ is parallel to $\mathbf{k} \times \mathbf{T}$. The image of the boundary curve is

$$\mathbf{r}_P(s_P) = \mathbf{k} \times (\mathbf{r} \times \mathbf{k}),$$

whose tangent is

$$\begin{aligned} \mathbf{T}_P &= \frac{d\mathbf{r}_P}{ds_P} \\ &= \mathbf{k} \times (\mathbf{T} \times \mathbf{k}) \frac{ds}{ds_P} \end{aligned} \tag{1}.$$

That is, the image of the tangent to the boundary curve is parallel to the tangent to the image of the boundary curve. In fact, it is easy to see that

$$\frac{ds}{ds_P} = \{1 - (\mathbf{T} \cdot \mathbf{k})^2\}^{-\frac{1}{2}} \tag{2}.$$

Since the unit surface normal $\mathbf{n}$ is parallel to $\mathbf{k} \times \mathbf{T}$, it follows from Eqs. (1) and (2) that it is parallel to $\mathbf{k} \times \mathbf{T}_P$ and so the surface normal at an extremal boundary can be determined from the image if the (orthographic) viewing direction is known. This result is due to Barrow and Tenenbaum [1978].

An arbitrary curve parameterized by arclength s and lying on the surface satisfies (see, for example, [Millman and Parker 1977, page 103]):

$$\kappa(s)\mathbf{N}(s) = \kappa_n(s)\mathbf{n}(s) + \kappa_g(s)\mathbf{n} \times \mathbf{T}(s), \tag{3}$$

where $\mathbf{T}$ and $\mathbf{N}$ are the unit tangent and unit normal to the curve (in the moving trihedron of the curve); κ is the curvature $|d\mathbf{T}/ds|$ of the curve; κ_n is the normal curvature of the curve, defined as the curvature of the curve that is the intersection of the surface and the plane that contains $\mathbf{n}$ and $\mathbf{T}$; and κ_g is called the geodesic curvature (Figure 7).

From now on, we restrict attention to the extremal boundary curve. Since $\mathbf{k} \cdot \mathbf{n} = 0$, Eq. 3 implies

$$\kappa \mathbf{N} \cdot \mathbf{k} = \kappa_g[\mathbf{k}, \mathbf{n}, \mathbf{T}],$$

where $[\ldots]$ indicates the triple scalar product. Since $\mathbf{k} \cdot \mathbf{n} = 0$ (and $\mathbf{n} \cdot \mathbf{T} = 0$), the triple scalar product is zero if and only if $\mathbf{k}$ coincides with $\mathbf{T}$. Similarly, $\mathbf{k} \cdot \mathbf{N}$ is zero if and only if $\mathbf{k}$ lies in the rectifying plane spanned by the tangent $\mathbf{T}$ and the binormal $\mathbf{B}$. It follows from general position arguments that if $\mathbf{k}$ has a component out of the rectifying plane then:

$$\kappa = 0 \text{ if and only if } \kappa_g = 0. \tag{4}$$

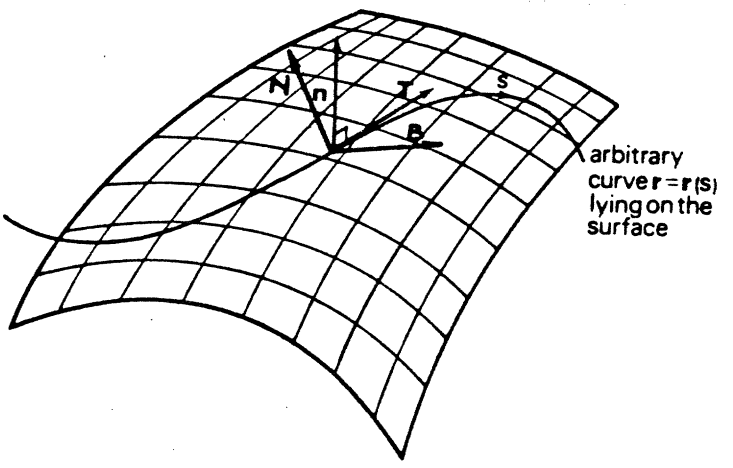

Figure 7. A curve lying on a surface. The vectors $\mathbf{T}, \mathbf{N}, \mathbf{B}$ are the moving trihedron of the curve, whose curvature is κ. The surface normal is $\mathbf{n}$.

Squaring Eq. 3 yields

$$\kappa^2 = \kappa_n^2 + \kappa_g^2. \tag{5}$$

It follows that if $\kappa = 0$ then $\kappa_n = \kappa_g = 0$, and if $\kappa_g = 0$, then by Eq. 4, in general, $\kappa = 0$, and $\kappa_n = 0$. It is easy to show (see Appendix A) that

$$\kappa_P^2 = \frac{\{1 - (\mathbf{T} \cdot \mathbf{k})^2 - (\mathbf{N} \cdot \mathbf{k})^2\}}{\{1 - (\mathbf{T} \cdot \mathbf{k})^2\}^3} \kappa^2. \tag{6}$$

where κ_P is the curvature of the projected curve $\mathbf{r}_P(s_P)$. It follows that, in general,

$$\kappa_P = 0 \text{ if and only if } \kappa = 0. \tag{7}$$

We can now present a new proof of a recent result due to Koenderinck [1984] that relates the curvature κ_P of the projection of the boundary curve to the Gaussian curvature κ_G of the surface. The key is to choose an appropriate parameterization of the surface at a point on the boundary curve (Figure 8). Define the radial curve at a point on the boundary to be the (normal) intersection of the surface with the plane that contains the surface normal $\mathbf{n}$ and the view vector $\mathbf{k}$. Let the radial curve be parameterized by u, and denote points along the radial curve by $\mathbf{r}_r(u)$. Using s and u to parameterize the surface in the neighborhood of a point on the boundary, we find the first and second fundamental forms of the surface:

$$\mathbf{G} = \begin{pmatrix} 1 & \mathbf{T}\cdot\mathbf{T}_r \\ \mathbf{T}\cdot\mathbf{T}_r & 1 \end{pmatrix}$$
$$\mathbf{D} = \begin{pmatrix} \kappa_r \mathbf{n}\cdot\mathbf{N}_r & \mathbf{n}\cdot\frac{\partial^2\mathbf{r}}{\partial s\partial u} \\ \mathbf{n}\cdot\frac{\partial^2\mathbf{r}}{\partial s\partial u} & \kappa\mathbf{n}\cdot\mathbf{N} \end{pmatrix}, \tag{8}$$

where κ_r is the curvature of the radial curve. In order for the boundary rim to be visible, κ_r must be positive.

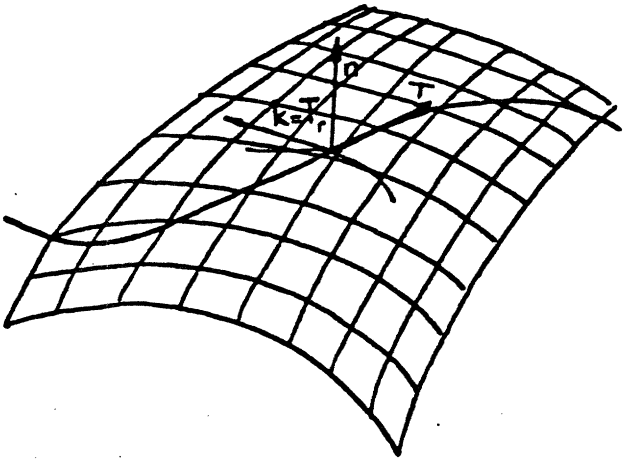

Figure 8. The rim and radial coordinate frame. The radial curve is the intersection of the surface and the plane that contains the view vector and the surface normal. The boundary curve defines the other parameter.

Note that, in general, $\mathbf{T}\cdot\mathbf{T}_r \neq 0$. Since the tangent $\mathbf{T}_r$ to the radial curve $\mathbf{r}_r$ lies in the plane spanned by both $\mathbf{k}$ and $\mathbf{n}$ (remembering that we are considering a boundary curve), and is in the surface tangent plane, we have

$$\mathbf{T}_r = \mathbf{k}, \text{ and } \mathbf{N}_r = \mathbf{n}. \tag{9}$$

It follows that $\mathbf{T}_r$ is constant along the boundary curve, so that

$$\begin{aligned} 0 &= \frac{\partial \mathbf{T}_r}{\partial s} \\ &= \frac{\partial^2 \mathbf{r}}{\partial u \partial s}. \end{aligned} \tag{10}$$

The second fundamental form reduces by Eqs. 9 and 10 to

$$\mathbf{D} = \begin{pmatrix} \kappa_r & 0 \\ 0 & \kappa\mathbf{n}\cdot\mathbf{N} \end{pmatrix}. \tag{11}$$

Since

$$\mathbf{n} = \frac{\mathbf{k} \times \mathbf{T}}{\{1 - (\mathbf{k} \cdot \mathbf{T})^2\}^{\frac{1}{2}}},$$

we find

$$\begin{aligned} \mathbf{n} \cdot \mathbf{N} &= \frac{[\mathbf{k}, \mathbf{T}, \mathbf{N}]}{\{1 - (\mathbf{k} \cdot \mathbf{T})^2\}^{\frac{1}{2}}}, \\ &= \frac{\{1 - (\mathbf{k} \cdot \mathbf{T})^2 - (\mathbf{k} \cdot \mathbf{N})^2\}^{\frac{1}{2}}}{\{1 - (\mathbf{k} \cdot \mathbf{T})^2\}^{\frac{1}{2}}}. \end{aligned}$$

The Gaussian curvature κ_G of the surface at points along extremal boundaries is given by [Faux and Pratt 1979, page 112]

$$\begin{aligned} \kappa_G &= \frac{|\mathbf{D}|}{|\mathbf{G}|} \\ &= \frac{\kappa_r \kappa \mathbf{n} \cdot \mathbf{N}}{1 - (\mathbf{T} \cdot \mathbf{k})^2} \\ &= \kappa_r \kappa \frac{\{1 - (\mathbf{k} \cdot \mathbf{T})^2 - (\mathbf{k} \cdot \mathbf{N})^2\}^{\frac{1}{2}}}{\{1 - (\mathbf{k} \cdot \mathbf{T})^2\}^{\frac{3}{2}}} \end{aligned} \tag{12}$$

It finally follows from Eq. 6 that

$$|\kappa_G| = |\kappa_r \kappa_P| \tag{13}.$$

Since κ_r is always positive, *the sign of the Gaussian curvature of the surface at points along the boundary curve is the same as the sign of the curvature of the projection of the boundary curve.* [Koenderinck 1984].

We can prove a slight extension to Koenderinck's result. If κ_P is zero, then the Gaussian curvature κ_G is zero. Also, recall from Eq. 6 that if κ_P is zero then (in general) the curvature κ of the boundary curve is also zero. It follows from Eq. 5 that the normal curvature κ_n of the boundary curve is also zero. Denote the principal curvatures at the surface point by κ_1, κ_2. The Gaussian curvature $\kappa_G = \kappa_1 \kappa_2$ is zero, and so at least one of κ_1, κ_2 is zero. But, by Euler's theorem,

$$\kappa_n = \kappa_1 \cos^2 \theta + \kappa_2 \sin^2 \theta,$$

and so $\kappa_1 = \kappa_2 = 0$ or θ is zero or $\pi/2$. That is, points on the projection of an extremal boundary where the curvature κ_P is zero typically correspond to surface points that are locally flat, that is, where both principal curvatures are zero, or the boundary is locally aligned with a surface line of curvature whose curvature is zero.

We have assumed that it is possible to determine which bounding contours are extremal and which mark discontinuities. This is a reasonable assumption in the case of dense surface data such as that used in the experiments reported in the next section. It is much more difficult in the case of line drawings such as that shown in Figure 9. Barrow and Tenenbaum [1981] propose that line labelling can suffice to make the distinction. We suggest that the results derived in this section hint at a more general approach that is

based on an analysis of the surfaces meeting at a corner. For example, if the smooth contour curves are all extremal then the Gaussian curvature would be positive along the curve with positive curvature and negative along the curves with negative curvature. This would imply that the surface changes the sign of its Gaussian curvature. But there are no surface markings or other evidence that it does. A more parsimonious assumption is that the surface has the same (positive) Gaussian curvature everywhere, and hence that the contour curves with negative curvature are discontinuities. This is, in fact, what is perceived.

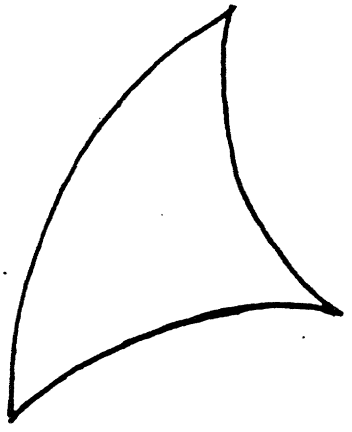

Figure 9. A line drawing that is perceived as the curved surface of a sail.

We can derive further results about the relationships between surface curves and the surface. For example, the sign of the normal curvature along the boundary rim can be determined from its projection [Yuille and Brady 1984]

2.3. Surface Intersections

In this Section, we prove the following Theorem, which suggests that we can find surface intersections at a given scale of Gaussian smoothing by inspecting the lines of curvature. We noted earlier that it undoes Joachimstal's Theorem.

Theorem The Gaussian convolution of a cylindrical surface is cylindrical. In more detail, let $f(x, y, z)$ be a surface that is the cross product of a planar curve and a straight line. The lines of curvature of the convolution of f with a Gaussian are in the plane of the curve and parallel to the generating line.

We begin with a Lemma, whose proof is simple and is omitted.

Lemma Let $\mathbf{r}(x, y) = x\mathbf{i} + y\mathbf{j} + f(x)\mathbf{k}$ be a cross product surface consisting of parallel instances of a curve $f(x)$ in the $x - z$ plane. The principal curvatures and directions of $\mathbf{r}(x, y)$ are zero in the y direction and the curvature of f in the $x - z$ plane.

Proof of Theorem. We assume that the surface has the form of the Lemma. Since the Gaussian is separable, the convolution of the surface is

$$G_\sigma(r) * f(x) = G_\sigma(y) * (G_\sigma(x) * f(x)).$$

By the derivative theorem for convolutions, it follows that

$$\{G_\sigma(r) * f(x)\}^{(n)} = G_\sigma(y) * \{G_\sigma(x) * f(x)\}^{(n)}.$$

Convolution with a constant is simply multiplication by a constant, and so the principal curvatures of $G_\sigma(r) * f(x)$ are the same as the principal curvatures of $G_\sigma(x) * f(x)$, which are given by the Lemma.

The Lemma can be extended straightforwardly to show that the Gaussian convolution of a surface of the form $z = f(x)(A + By)$ is a similar surface $z = (G_\sigma * f(x))(A + By)$ but the lines of curvature are not preserved.

2.4. Lines of Curvature

Brady and Yuille [1984b] argue that, in many cases, the lines of curvature give a natural parameterization of a surface. One practical advantage is that a computer-aided design (CAD) patch representation based on lines of curvature avoids problems of local flattening of the surface. Stevens [1982, 1983] has studied drawings consisting of a repeated pattern of "parallel" planar curves and the curved surfaces they suggest. He proposes that the given curves are often interpreted as lines of curvature of the perceived surface.

In Section 2.1 we noted that a line of curvature has to satisfy additional constraint if it is to be made explicit. For example, only the planar lines of curvature of the ellipsoid shown in Figure 2 are useful for describing the surface.

Consider a surface of revolution. Suppose that the axis is aligned with the $z-$ axis. The surface is formed by rotating the (one-parameter) curve $p(u)\mathbf{i} + z(u)\mathbf{k}$ about $\mathbf{k}$. The surface is

$$r(u, \theta) = p(u)\cos\theta\mathbf{i} + p(u)\sin\theta\mathbf{j} + z(u)\mathbf{k}.$$

The principal curvatures (see for example Millman and Parker [1977, p86]) are the meridians and the parallels, all of which are planar. In addition, the parallels are circular, so the curvature along any one of them is constant. The curvature along a parallel is $\mathbf{n}^T\mathbf{r}^*/p(u)$, where $\mathbf{r}^* = [\cos\theta \ \sin\theta \ 0]^T$. The foreshortening of the expected curvature $p(u)$ exemplifies Meusnier's theorem [Weatherburn 1927]. On the other hand, the asymptotes on a surface of revolution are, in general, complex space curves and the geodesic curvature is a complex function of position along the asymptote.

Surfaces of revolution are essentially one-dimensional in that their shape is completely determined by a planar curve. It is reasonable to ask whether lines of curvature are more generally useful. The theorem stated in Section 2.1 shows that they are. We first prove the theorem for the straight axis case originally studied by Marr [1977], relegating the more general case to Appendix B.

Theorem 2 If the axis of a generalized cone is straight, and the axis is normal to the cross-section, then (i) a cross-section is a line of curvature if and only if it is a skeleton; (ii) a tracing is a line of curvature if and only if it is a fluting.

Proof. To fix notation, we begin by analyzing a planar cross section (Figure 10). The curve is $(f(s), g(s))$, where s denotes arclength, and so its *radial distance* from the origin is given by

$$d^2(s) = f^2 + g^2. \tag{14}$$

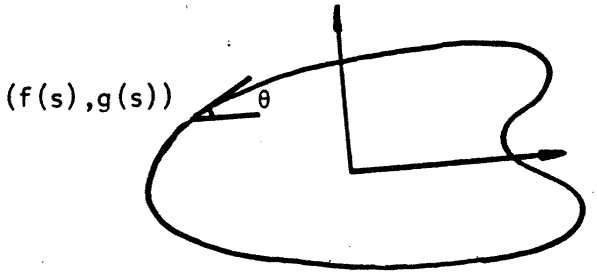

Figure 10. A planar cross-section curve

Differentiating Eq. 14 with respect to arc length gives

$$\dot{d} = \frac{f\dot{f} + g\dot{g}}{d}. \tag{15}$$

Thus the radial distance attains an extremum where the numerator of Eq. 15 is zero. Now consider the orientation θ of the tangent to the curve.

$$\begin{aligned} \tan\theta &= \frac{\dot{g}}{\dot{f}}, \\ \sec^2\theta\dot{\theta} &= \frac{\dot{f}\ddot{g} - \dot{g}\ddot{f}}{\dot{f}^2}, \\ \cos^2\theta &= \frac{\dot{f}^2}{\dot{f}^2 + \dot{g}^2}, \\ \dot{\theta} &= \frac{\dot{f}\ddot{g} - \dot{g}\ddot{f}}{l^2}, \end{aligned} \tag{16}$$

where $l^2(s) = \dot{f}^2 + \dot{g}^2$. Notice that $\dot{\theta}$ is the curvature of the cross section curve.

We are now ready to prove the Theorem. Suppose, without loss of generality, that the straight axis of the cone is $\mathbf{k}$. Let the cross-section be $f(s)\mathbf{i} + g(s)\mathbf{j}$, as above. Suppose the expansion function is $h(z)$, and assume that the eccentricity of the cone is zero, that is, the axis is normal to the cross-section. The generalized cone is

$$\mathbf{r}(s,z) = h(z)f(s)\mathbf{i} + h(z)g(s)\mathbf{j} + z\mathbf{k}.$$

To save on notation, we suppress parameters. We find

$$\begin{aligned} \frac{\partial \mathbf{r}}{\partial s} &= [h\dot{f}, h\dot{g}, 0]^T \\ \frac{\partial \mathbf{r}}{\partial z} &= [\dot{h}f, \dot{h}g, 1]^T \end{aligned},$$

(where $[\ldots]^T$ denotes the vector that is the transpose of the given row vector) and so the first fundamental form of the surface is (using the notation introduced above)

$$\mathbf{G} = \begin{pmatrix} h^2 l^2 & h\dot{h}d\dot{d} \\ h\dot{h}d\dot{d} & 1 + \dot{h}^2 d^2 \end{pmatrix}.$$

The surface unit normal $\mathbf{n}$ is parallel to $\mathbf{h} = [h\dot{g}, -h\dot{f}, h\dot{h}(\dot{f}g - f\dot{g})]^T$. The second fundamental form of the surface is

$$\mathbf{D} = \begin{pmatrix} \frac{1}{|\mathbf{h}|}h^2(\ddot{f}\dot{g} - \dot{f}\ddot{g}) & 0 \\ 0 & \frac{1}{|\mathbf{h}|}h\ddot{h}(f\dot{g} - \dot{f}g) \end{pmatrix}.$$

The principal directions of curvature are the eigenvectors of the matrix $\mathbf{G}^{-1}\mathbf{D}$, and this matrix is diagonal if and only if

$$\frac{dh}{dz}\frac{dd}{ds} = 0,$$

from which the result follows.

2.5. Asymptotes

We begin with an example that illustrates that lines of curvature are not always the best basis for describing a surface. Consider a helicoid of a single blade (Figure 11), which can be parameterized as follows

$$\mathbf{r}(l,\theta) = [l\cos\theta, l\sin\theta, m\theta]^T,$$

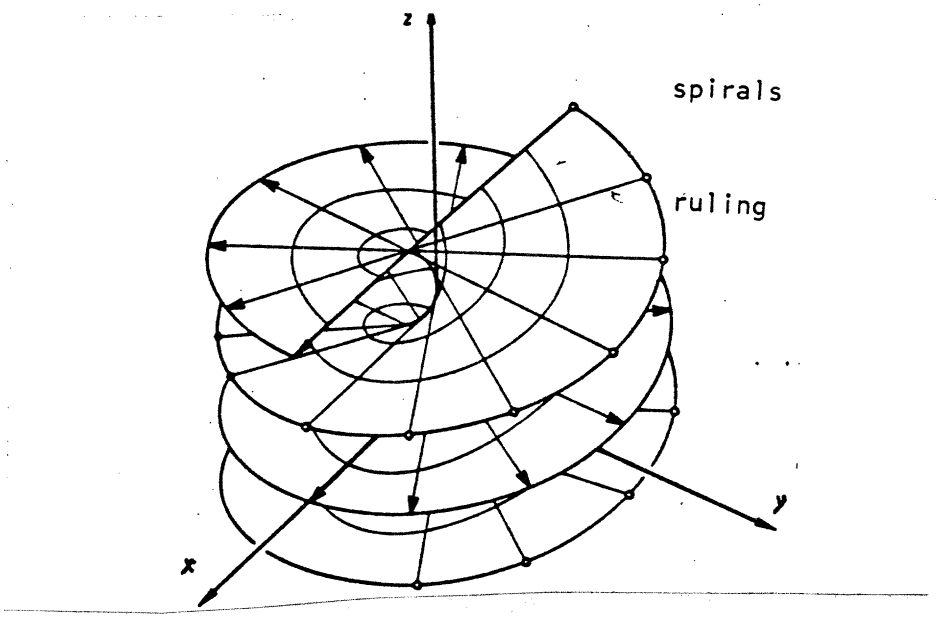

Figure 11. The helicoid of a single blade. (Reproduced from [do Carmo 1976, Figure 2-27 Page 94])

where m and l are assumed positive. Denoting $\sqrt{l^2+m^2}$ by d, and m/l by $\tan\psi$, we find

$$\mathbf{r}_l = [\cos\theta, \sin\theta, 0]^T$$
$$\mathbf{r}_\theta = [-l\sin\theta, l\cos\theta, m]^T$$
$$\mathbf{n} = [\sin\theta\sin\psi, -\cos\theta\sin\psi, \cos\psi]^T$$
$$\mathbf{G} = \begin{pmatrix} 1 & 0 \\ 0 & d^2 \end{pmatrix}$$
$$\mathbf{D} = \begin{pmatrix} 0 & -\frac{m}{d} \\ -\frac{m}{d} & 0 \end{pmatrix}$$

The principal directions of curvature are given by $\dot{l} = \pm d\dot{\theta}$, and the principal curvatures are $\pm m/d^2$. Since the parameter l varies in the principal directions, so does m, hence so does the principal curvature. It is easy to show that the lines of curvature are not planar.

Notice that the diagonal of $\mathbf{D}$ is zero. It follows that the normal curvature in the directions of the tangent vectors $\mathbf{r}_l$, $\mathbf{r}_\theta$ is zero. These are the asymptotes, and they correspond to the rulings and spirals that make up the helicoid. The geodesic curvature along the spirals is simply the curvature of the spiral considered as a curve. The geodesic curvature of the ruling is zero. (It is a curious fact that the lines of curvature cut the asymptotes at a constant angle of $\pi/4$.)

In general, ruled surfaces are poorly described by their lines of curvature. Examples such as the helicoid and the surface $z = kxy$ suggest that the asymptotes may be a better basis for description. The asymptotic direction can be found from Euler's theorem if the principal curvatures have opposite signs, that is, the Gaussian curvature is negative. Note that a given ruled surface may be generated by more than one set of rulings. Also, it is

possible for a ruled surface to admit a significantly different description. For example, the hyperboloid $x^2 + y^2 - z^2 = 1$ is both a ruled surface and a surface of revolution.

In general, a ruled surface can be parameterized in the form

$$\mathbf{r}(s,t) = \mathbf{u}(s) + t\mathbf{w}(s),$$

where $\mathbf{u}$ is a curve called the directrix, and $\mathbf{w}(s)$ is the set of rulings. The theory of ruled surfaces emphasizes a unique curve, called the line of striction, that lies in the surface and is orthogonal to the $\mathbf{w}'$. It is not clear what role, if any, it plays in perception.

The normal to a ruled surface is parallel to

$$\mathbf{u}' \times \mathbf{w} + t\mathbf{w}' \times \mathbf{w},$$

and varies with t along the ruling. The normal direction is constant along a ruling if and only if $\mathbf{n} \cdot \mathbf{u}'$ is constant, which is if and only if the triple scalar product $[\mathbf{u}', \mathbf{w}', \mathbf{w}]$ is zero. Along such rulings, the determinant of the second fundamental form $\mathbf{D}$ is zero, and this in turn implies that the Gaussian curvature κ_G is zero. If this condition holds for all rulings on a ruled surface, the surface is called developable and the Gaussian curvature is everywhere zero. Informally, a surface is developable if it can be rolled out flat onto a plane. For such surfaces the rulings are both asymptotes and lines of curvature. For developable surfaces, the descriptive bases of Sections 2.4 and 2.5 coincide.

This section has advanced a number of mathematical methods for isolating curves that embody important information about a surface. Now we turn to their computation.

3. Computational Experiments

In this Section we report on a number of computational experiments that investigate whether the surface curves and regions proposed by our theoretical analysis can be computed reliably and efficiently. The input to our programs are mostly (dense) depth maps obtained by the structured light systems at MIT [Brou 1984] and INRIA [Faugeras et. al. 1983]. Both systems are accurate to about 0.5 millimeters. The objects that we have worked with include: a bottle, an egg, a sphere, a styrofoam cup, a lightbulb, and a pen (all surfaces of revolution); a telephone handset (surface intersections and an approximately ruled surface); a coffee mug with a handle, a plastic container, a hammer, and a Renault part [Faugeras et. al. 1983] (complex surfaces with surface intersections). We have also conducted experiments with artificial data to which controlled amounts of noise have been added.

3.1. Gaussian Smoothing

Depth maps generated by structured light systems are noisy, as are image surfaces. In recent years, Gaussian smoothing filters have been extensively investigated for early processing of images. For example, Marr and Hildreth [1980] suggest Laplacian of a Gaussian filters $\nabla^2 G_\sigma$ for edge finding. These filters are closely approximated by difference-of-Gaussian (DOG) filters that can be efficiently implemented. Poggio and Torre [1984] and Canny [1983] have suggested directional edge finders whose first step is Gaussian smoothing.

Witkin [1983] has proposed *scale-space filtering* in which a (one-dimensional) signal is filtered at a variety of spatial scales to produce a hierarchical description. Witkin suggests that it is possible to automatically determine a discrete set of "natural scales" at which to describe a signal symbolically. Witkin's scale space representation is a ternary tree of zero crossings of G''_σ. He did not attempt to *interpret* the multiple descriptions in terms of primitive events. Asada and Brady [1984] have shown how scale space filtering can be used to generate a symbolic description of the significant curvature changes along a planar contour.

Yuille and Poggio [1983a] have provided some theoretical underpinning for the scale space representation. They have shown that the contour of zero crossings of second derivatives ("fingerprint") may preserve enough information to reconstruct the original signal to within a constant scale factor. They also show [Yuille and Poggio 1983b] that a Gaussian filter is essentially unique in having the property that zero crossings are not introduced as one moves to coarser scales.

In view of this background with images, the first processing stage of our program is Gaussian smoothing. Initially, we applied the Gaussian filter at every surface point $z(x, y)$. This is unsatisfactory as it smooths across the depth discontinuities that are the bounding contours of an object (Figure 12b). This is an advantage in edge detection; but a disadvantage in smoothing within a given region. In our experimental data, the depth discontinuities are sharp and easy to find before smoothing. It is difficult to prevent smoothing across discontinuities with a straightforward implementation of Gaussians using convolution masks. It is particularly difficult when the distance between opposite sides of a surface is roughly the same size as the footprint of the convolution mask as it typically is in our data.

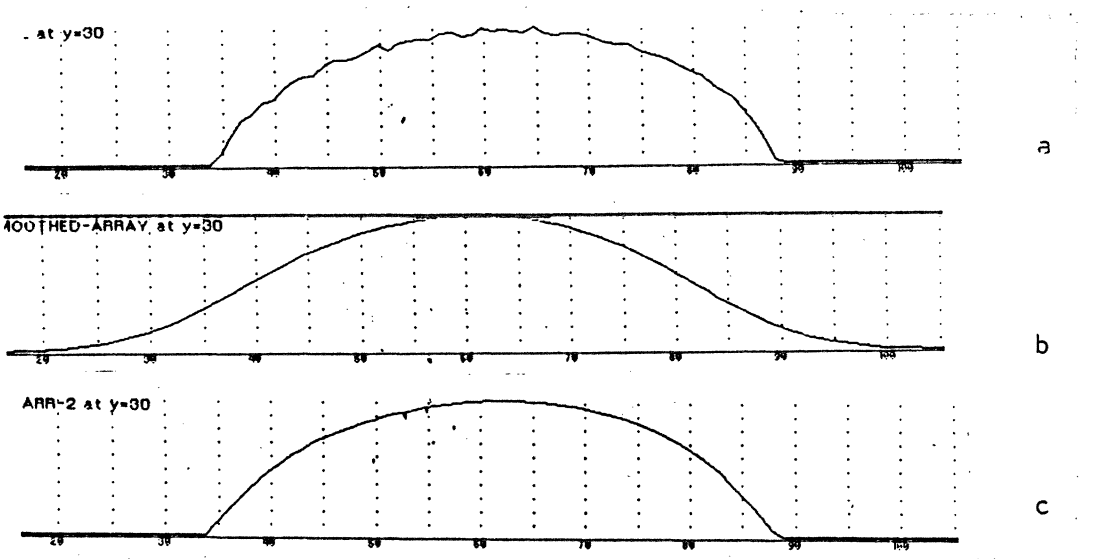

Figure 12. a. Raw data from a cross section of an oil bottle after scanning using the INRIA system. b. Smoothing across surface boundaries with a Gaussian mask that is applied everywhere. c. Gaussian smoothing using repeated averaging and computational molecules.

We appeal to the central-limit theorem and implement Gaussian filtering using repeated averaging with the 3×3 mask shown in Figure 13a (see Burt [1981, 1983]). Iterating n times approximately corresponds to filtering with a Gaussian whose standard deviation is proportional to $\sqrt{n}$. To prevent smoothing across previously marked depth discontinuities, we use the technique of computational molecules proposed by Terzopoulos [1983] in his application of finite-element techniques to surface reconstruction. In more detail, the mask shown in Figure 13a is viewed as the sum of the four molecules shown in Figure 13b. To apply the filter to a surface point, only those molecules that do not overlap the set of points marked as discontinuities are used. Figure 12c shows the result of Gaussian smoothing using repeated averaging and computational molecules.

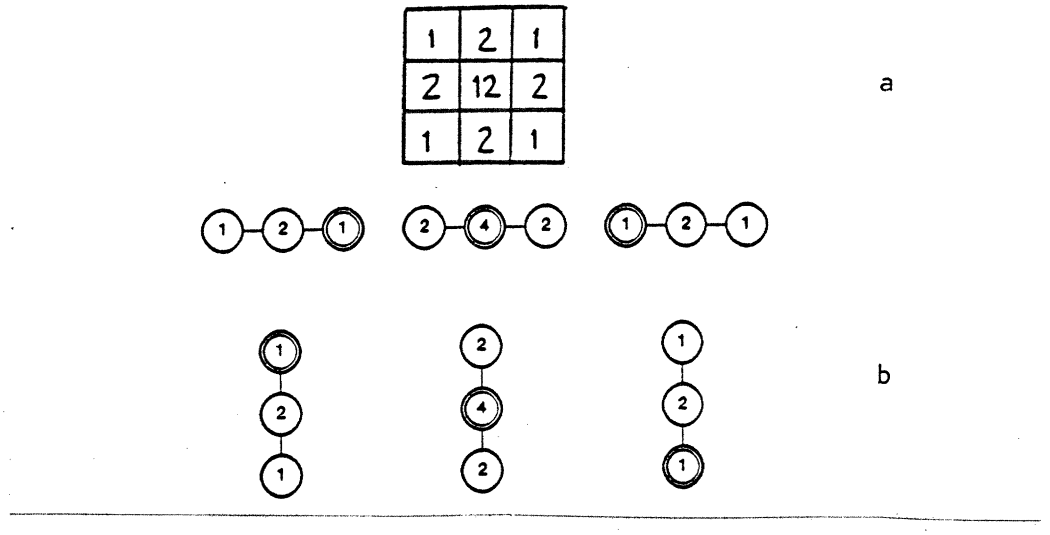

Figure 13. Gaussian smoothing. a. The 3×3 mask that is repeatedly applied to approximate Gaussian filtering. b. The computational molecules whose sum is the mask shown in a.

3.2. Lines of Curvature

Suppose that the (Gaussian smoothed) surface is given in vector form as $\mathbf{z}(x, y)$. Its first and second derivatives can be estimated using the facet model [Haralick 1980]. We derived the finite difference operators shown in Figure 14 by least squares fitting a quadratic to a 3×3 facet of the surface. Although these operators are sensitive to noise in raw image or surface data, they perform well after Gaussian smoothing.

From these estimates we can compute the fundamental forms of the surface [Millman and Parker 1977, Faux and Pratt 1979]. The first fundamental form is:

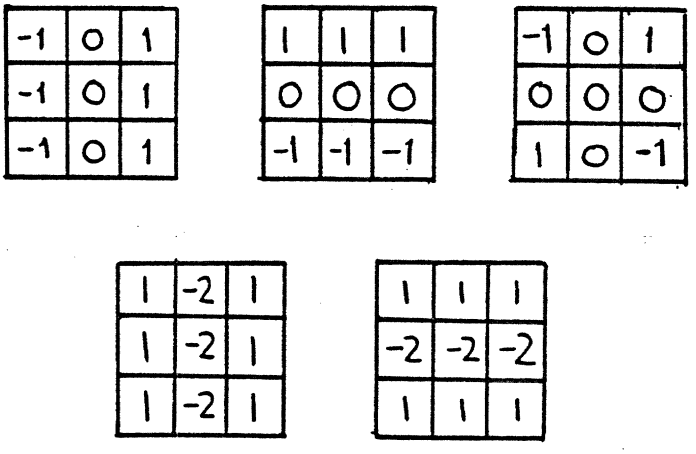

Figure 14. Operators used to estimate first and second derivatives of the Gaussian smoothed surface.

$$\begin{aligned}\mathbf{G} &= \begin{pmatrix} \mathbf{z}_x^2 & \mathbf{z}_x \cdot \mathbf{z}_y \\ \mathbf{z}_x \cdot \mathbf{z}_y & \mathbf{z}_y^2 \end{pmatrix} \\ &= \begin{pmatrix} 1+p^2 & pq \\ pq & 1+q^2 \end{pmatrix} \\ &= \begin{pmatrix} g_{11} & g_{12} \\ g_{12} & g_{22} \end{pmatrix}\end{aligned}$$

where p is the z component of $\mathbf{z}_x$ and q is the z component of $\mathbf{z}_y$. Similarly, the second fundamental form is:

$$\begin{aligned}\mathbf{D} &= \begin{pmatrix} \mathbf{n} \cdot \mathbf{z}_{xx} & \mathbf{n} \cdot \mathbf{z}_{xy} \\ \mathbf{n} \cdot \mathbf{z}_{xy} & \mathbf{n} \cdot \mathbf{z}_{yy} \end{pmatrix} \\ &= \begin{pmatrix} d_{11} & d_{12} \\ d_{12} & d_{22} \end{pmatrix}\end{aligned}$$

where $\mathbf{n}$ is the unit surface normal. The (tangent) principal directions are the solutions of the quadratic equation [Faux and Pratt 1979]:

$$(d_{12}g_{22} - d_{22}g_{12})(\frac{dy}{dx})^2 + (d_{11}g_{22} - d_{22}g_{11})\frac{dy}{dx} + (d_{11}g_{12} - d_{12}g_{11}) = 0.$$

Suppose that $\lambda = dy/dx$ is a tangent principal direction. Then the unit vector lying in the surface tangent plane in the principal direction is parallel to

$$\mathbf{t} = [1, \lambda, p + q\lambda]^t.$$

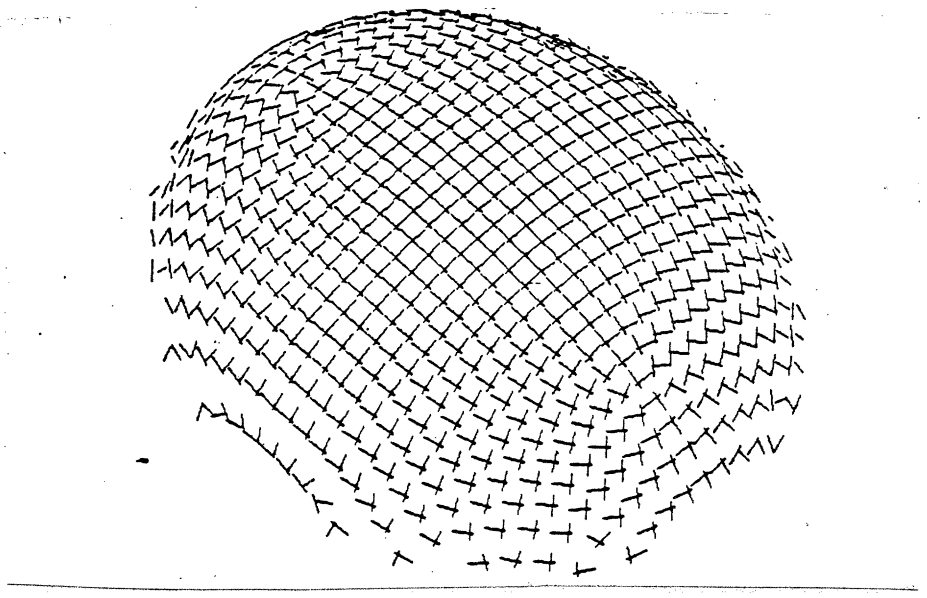

Figure 15. The orthographic projections of the principal directions of curvature for an ellipsoid.

By Euler's theorem, the surface tangent vectors $\mathbf{t}_{max}$ and $\mathbf{t}_{min}$ in the principal directions are orthogonal. Of course, their (orthographic) projections onto the plane $z = 0$ are not orthogonal, as Figure 15 shows.

The surface curvatures in the principal directions are the roots of the quadratic [Faux and Pratt 1979, page 112]:

$$|\mathbf{G}|\kappa^2 - (g_{11}d_{22} + g_{22}d_{11} - 2g_{12}d_{12})\kappa + |\mathbf{D}| = 0. \tag{17}$$

The curvature values can be used to extract planar patches of a surface and regions that consist entirely of umbilic points. To find planar patches, for example, we have adopted the simple approach of requiring the absolute values of both principal curvatures to be below a small threshold. Figure 16a shows the planar patches found by this technique for a Renault part used in experiments by Faugeras and his colleagues at INRIA. Faugeras et. al. [1982, 1984] have investigated a variety of techniques that exploit special properties of planes. Figure 16b shows the planes computed by the INRIA group for the Renault part. It is the opinion of Jean Ponce, who has worked both at INRIA and at MIT, that the two methods are equally effective.

A point is umbilic if its principal curvatures are equal. The principal curvatures define a best fitting ellipsoid in the local neighborhood of a surface point. In the case of an umbilic point, the ellipsoid is a sphere. To determine umbilic points, we simply require that

$$\frac{|\kappa_1| - |\kappa_2|}{\max(|\kappa_1|, |\kappa_2|)}$$

is less than a threshold, where κ_1 and κ_2 are the solutions to Eq. 17. Figure 1 shows a region of umbilic points found for the lightbulb. It is well-known (for example, do Carmo

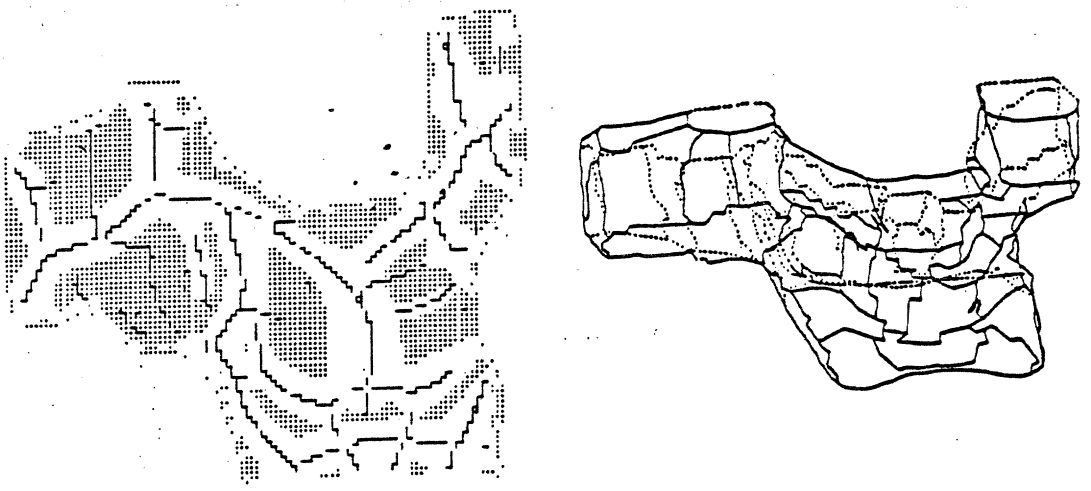

Figure 16. a. Planar patches on the Renault part found by simply requiring the absolute values of both principal curvatures to be below a threshold. b. The planar patches on the same part found by the INRIA group. (Reproduced from [Faugeras et. al. 1984, page 432])

[1976, page 147 proposition 4]) that if all the points of a connected surface are umbilic, then the surface is either contained in a sphere or in a plane. Since the curvatures are non-zero, the connected region that is the bulb of the lightbulb is a portion of a sphere.

In order to extract the lines of curvature, the directions at the individual surface points need to be linked. This is more difficult than linking zero crossings to form an edge for example since the principal directions form a dense set. Figure 17 shows that lines of curvature cannot be extracted by simply choosing the eight-connected neighbor with the nearest direction after projection onto $z = 0$.

Instead, linking is based on the vectors $\mathbf{t}_{max}(x, y)$ in the surface tangent plane. Initially, each point (x, y) can potentially link to its eight-connected neighbors. As the program proceeds, neighbors become inhibited as they are linked to other points. The program conducts a breadth-first search. At each iteration, the point (x, y) for which a *closeness* evaluation function is minimized is chosen. If its minimizing neighbor is part of a growing line of curvature, (x, y) extends it, otherwise a new line of curvature is started. In either case, neighbors that are not minimizing inhibit their link to (x, y). The closeness function that is currently used is the sum of three dot products involving principal vectors at neighboring points in the surface tangent plane:

$$\begin{aligned} c(x_1, y_1, x_2, y_2) = \mathbf{t}(x_1, y_1) \cdot \mathbf{t}(x_2, y_2) \\ + \mathbf{t}(x_1, y_1) \cdot \mathbf{r}(x_1, y_1, x_2, y_2), \\ + \mathbf{t}(x_2, y_2) \cdot \mathbf{r}(x_1, y_1, x_2, y_2) \end{aligned}$$

where $\mathbf{r}(x_1, y_1, x_2, y_2) = \mathbf{r}(x_1, y_1) - \mathbf{r}(x_2, y_2)$. Other evaluation functions could have used the curvature values as well as their directions, but we have not found this to be necessary. Figure 18 shows the lines of curvature found by the algorithm for a coffee cup and an

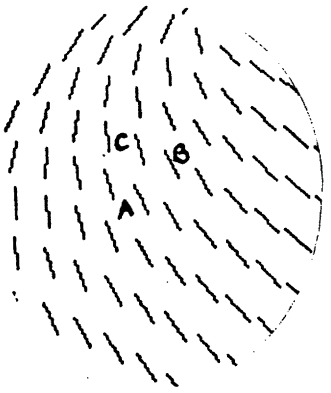

Figure 17. A swirling pattern of principal directions shows that lines of curvature cannot be extracted by simply choosing the eight-connected neighbor with the nearest direction after projection onto $z = 0$. The direction closest to the point marked A is at the point marked B, but C is more consistent with global judgments.

oil bottle. The program gives similarly good results on all our test objects. Note that because of discretization it is possible for all the neighbors of a point to become inhibited before it is selected. Hence some points may not lie on the lines of curvature computed by the program.

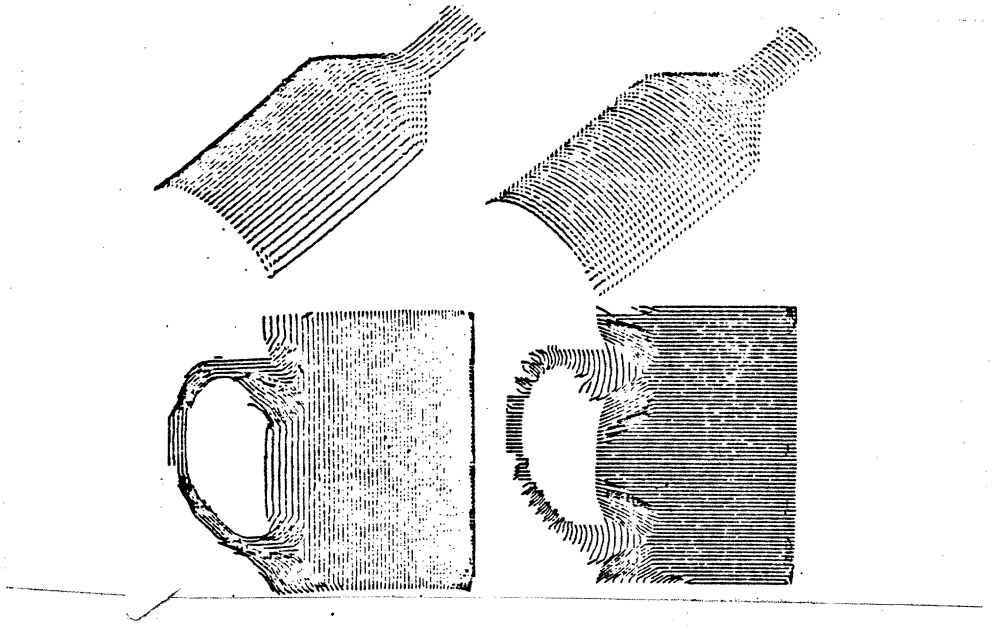

Figure 18. Linked lines of curvature found for an oil bottle and a coffee mug.

3.3. Using lines of curvature

Earlier, we showed that lines of curvature that are planar, or along which the principal curvature is constant, are important for describing surfaces. Given a linked list of surfaces points forming a line of curvature, we can determine the best fitting plane. If the set is $\{(x_i, y_i, a_i) | 1 \leq i \leq n\}$, then the least-squares fitting plane $ax + by + cz + d = 0$ is determined from a solution of

$$\begin{pmatrix} \text{var}\ (x) & \text{cov}\ (x,y) & \text{cov}\ (x,z) \\ \text{cov}\ (x,y) & \text{var}\ (y) & \text{cov}\ (y,z) \\ \text{cov}\ (x,z) & \text{cov}\ (y,z) & \text{var}\ (z) \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = 0,$$

where var (x) is the variance of the x_i's, and cov (x, y) the covariance of the x_i and y_i. We can determine whether a given population of points is planar by examining the condition number of the covariance matrix (compare Brady and Asada [1984, pages 341 – 342]). Similarly, we can compute the best fitting circle to a line of curvature and determine whether the population lies on that circle. Figure 19a shows the best fitting circles computed for the lines of curvature that are the parallels of a cup. Figure 19b shows the axis that is the locus of the centers of the circles in Figure 19a.

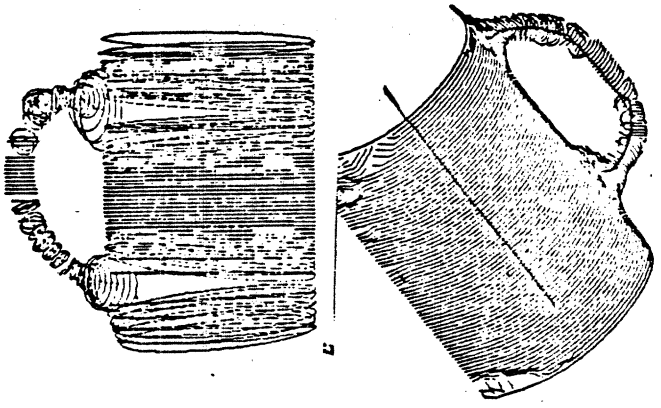

Figure 19. a. The best fitting circles to the parallels of the cup. b. the axis that is the locus of the centers of the best fitting circles shown in a.

We need to determine the significant discontinuities in a surface. The result would be a *Surface Primal Sketch* analogous to Marr's [1976] intensity change Primal Sketch (for image surfaces) and Asada and Brady's [1984] Curvature Primal Sketch for significant curvature changes along planar contours. In each case, the problem is to *detect* all significant changes, *localize* those changes as accurately as possible, and to *symbolically describe* the change. Yuille and Poggio [1983a, 1983b] have proved that, in principle,

scale space filtering enables a discontinuity to be accurately localized. Canny [1983] uses the smallest scale at which a given intensity change can be detected to most accurately localize it. Figure 20 shows the surface intersections of a telephone handset found by a program described below after the surface has first been smoothed at a variety of scales. The increasing localization of the surface intersection flanking the elongated portion of the surface can be clearly seen. Canny's [1983] claim that signal to noise increases proportional to the scale of the filter can also be seen. We are currently working to integrate these separate descriptions to yield a single description.

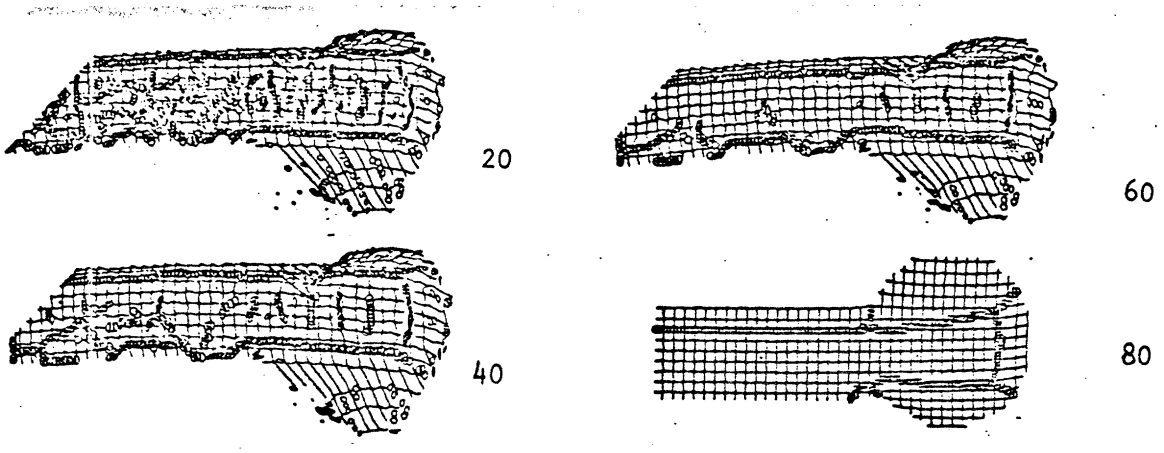

Figure 20. Significant surface discontinuities found by the curvature primal sketch program at multiple scales. The input to the program is the lines of curvature computed at each scale. The scales shown are 20, 40, 60, and 80

Earlier, we noted a theorem by Joachimsthal that shows that surfaces rarely intersect along their lines of curvature. We also showed that Gaussian smoothing overcomes this problem. So long as the curvature of the curve of intersection is small compared to the Gaussian filter, the lines of curvature of the smoothed surface lie parallel and perpendicular to the locus of curvature maxima of the smoothed surface. Asada and Brady [1983] filter a planar contour at multiple scales to detect, localize, and describe the significant changes in curvature. As an initial experiment, we have applied the curvature primal sketch program to lines of curvature after they have been projected into their best fitting plane. The results encourage us to extend the Asada and Brady program to two-dimensions. Figures 20 and 21 show some of the surface intersections found by this method.

4. Acknowledgements

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence

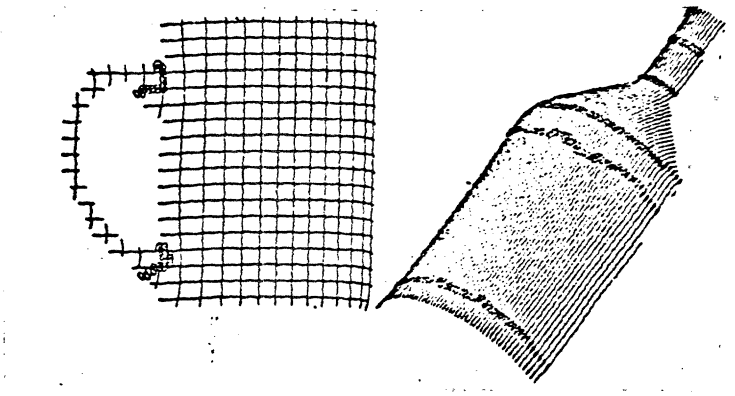

Figure 21. Surface intersections found by the curvature primal sketch program applied to lines of curvature after they have been projected into their best fitting plane.

research is provided in part by the the System Development Foundation, the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643, and the Office of Naval Research under contract number N00014-80-C-050. We thank several people who have commented on the ideas presented in this paper, particularly Brian Barsky, Tom Binford, Jean-Daniel Boissonat, Philippe Brou, Olivier Faugeras, Eric Grimson, Berthold Horn, Joe Mundy, Tommy Poggio, Demetri Terzopoulos, and Shimon Ullman. We are indebted to Olivier Faugeras and the INRIA group for giving us access to their ranging system. More especially, we appreciate the exchange of ideas, results, and personnel with INRIA. Steve Bagley, Margaret Fleck, and Eric Grimson made valuable comments on drafts of this paper.

5. Appendix A

We derive Equation 6.

$$\begin{aligned}\kappa_P^2 &= \left(\frac{d\mathbf{T}_P}{ds_P}\right)^2 \\ &= \left(\frac{ds}{ds_P}\right)^2\left(\frac{d}{ds}\frac{\mathbf{k}\times(\mathbf{T}\times\mathbf{k})}{\{1-(\mathbf{T}\cdot\mathbf{k})^2\}^{\frac{1}{2}}}\right)^2 \\ &= \left(\frac{ds}{ds_P}\right)^2\frac{\{1-(\mathbf{T}\cdot\mathbf{k})^2-(\mathbf{N}\cdot\mathbf{k})^2\}}{\{1-(\mathbf{T}\cdot\mathbf{k})^2\}^2}\kappa^2\end{aligned},$$

and the result follows from Eq. 2.

6. Appendix B: The Generalized Cone Theorem

We prove a more general form of the Theorem proved in Section 2.3. We relax the condition of the Theorem to planar axes. As before, we assume that the eccentricity of the generalized cone is zero, so that the tangent to the axis curve is normal to the cross section. Without this assumption, the proof becomes quite complex.

Theorem If the axis of a generalized cone is planar, and the eccentricity of the cone is zero, then (i) a cross section is a line of curvature if either the generalized cone is a surface of revolution or the cross section is an extremum; and (ii) a tracing is a line of curvature if the generalized cone is a *tube surface* (the expansion function is a constant), or the tracing is a fluting.

Proof. Denote the Frenet-Serret moving trihedron of the axis curve by $(\mathbf{T}, \mathbf{N}, \mathbf{B})$. Since the axis is planar it has zero torsion, which simplifies the Frenet-Serret formulae:

$$\frac{d\mathbf{T}}{dt} = \kappa\mathbf{N}, \text{ and } \frac{d\mathbf{N}}{dt} = -\kappa\mathbf{T} \qquad (B2.1)$$

Note that the binormal $\mathbf{B}$ is constant, being the unit normal to the plane of the axis. Since the eccentricity of the cone is assumed to be zero, the tangent $\mathbf{T}$ to the axis is normal to the cross-section, which has $\mathbf{N}$ and $\mathbf{B}$ as a basis. Let the axis of the cone be $\mathbf{a}(t)$, where t is arclength, and let the expansion function be $h(t)$. As in Section 2.3, suppose that the cross-section is given by $f(s)\mathbf{N} + g(s)\mathbf{B}$. The generalized cone is defined to be

$$\mathbf{r}(s,t) = \mathbf{a}(t) + h(t)\{f(s)\mathbf{N} + g(s)\mathbf{B}\}. \qquad (B2.2)$$

The analysis proceeds as in Section 2.3. The first fundamental form is

$$\mathbf{G} = \begin{pmatrix} h^2 l^2 & h\dot{h}d\dot{d} \\ h\dot{h}d\dot{d} & \dot{h}^2 d^2 + (1-\kappa h f)^2 \end{pmatrix}.$$

As expected, this reduces to the expression in Section 2.3 in the case $\kappa = 0$. The second fundamental form hints at the complexity of the most general case, when the axis is not restricted to lie in a plane:

$$\mathbf{D} = \begin{pmatrix} \frac{1}{|\mathbf{n}|}d_{11} & \frac{1}{|\mathbf{n}|}d_{12} \\ \frac{1}{|\mathbf{n}|}d_{12} & \frac{1}{|\mathbf{n}|}d_{22} \end{pmatrix},$$

where

$$\begin{aligned} d_{11} &= h^2(\dot{f}\ddot{g} - \ddot{f}\dot{g})(1-\kappa h f) \\ d_{12} &= -\kappa h^2 \dot{f}\dot{h}(\dot{f}g - f\dot{g}) \\ d_{22} &= h\ddot{h}(\dot{f}g - f\dot{g})(1-\kappa h f) \\ &\quad -(1-\kappa h f)^2 h\kappa\dot{g} \\ &\quad -fh\dot{h}(\dot{f}g - f\dot{g})(\dot{\kappa}h + 2\kappa\dot{h}) \end{aligned} \qquad (B2.3)$$

We establish the result as before.

7. References

Agin, G. J, [1972], Representation and description of curved objects, Stanford University, AIM-73.

Asada, Haruo, and Brady, Michael, [1984], The curvature primal sketch, MIT Artificial Intelligence Laboratory, AIM-758.

Ballard, D. H., and Brown, C. M, [1982], *Computer Vision*, Prentice-Hall.

Barnard, S, [1983], "Interpeting perspective images," *Artificial Intelligence*, **21 (4)** , 435 – 462.

Barnard, S., and Pentland, Alex P, [1983], Three-dimensional shape from line drawings, Proc. 8th Int. Jt. Conf. Artif. Intell., Karlsruehe, 1061 – 1063.

Barrow, H. G., and Tenenbaum, J. M, [1978], Recovering intrinsic scene characteristics from images, *Computer vision systems* Riseman, E. M., and Hanson, A. (eds). Academic, 3 – 26.

Barrow, H. G., and Tenenbaum, J. M, [1981], "Interpreting line drawings as three-dimensional surfaces," *Artif. Intell.*, **17** , 75 – 116.

Binford, T. O, [1982], "Inferring surfaces from images," *Artificial Intelligence*, **17** , 205 – 245.

Bolles, Robert C., Horaud, Patrice, and Hannah, Marsha Jo, [1984], 3DPO: a three-dimensional parts orientation system, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 413 – 424.

Brady, Michael, [1982], "Computational approaches to image understanding," *Computing surveys*, **14** , 3 - 71.

Brady, Michael, and Asada, Haruo, [1984], "Smoothed local symmetries and their implementation," *Int. J. Robotics Research*, **3 (3)** .

Brady Michael, and Horn B. K. P, [1983], "Rotationally symmetric operators for surface interpolation," *Computer Vision, Graphics, and Image Processing*, **22** , 70 - 95.

Brady, Michael, and Yuille, Alan, [1984a], "An extremum principle for shape from contour," *IEEE Patt. Anal. and Mach. Int.*, **PAMI-6** , 288 – 301.

Brady, Michael, and Alan Yuille, [1984b], Representing three-dimensional shape, Romansy Conf., Udine, Italy.

Brooks, R.A, [1981], "Symbolic Reasoning Among 3-D Models and 2D Images," *Artificial Intelligence*, **17** , 285-348.

Brooks, R. A., and Binford, T. O, [1980], Representing and reasoning about partially specified scenes, *Proc. Image Understanding workshop* Baumann Lee S. (ed)., Science Applications Inc., 95 - 103.

Brou, P, [1984], "Finding the orientation of objects in vector maps," *Int. J. Rob. Res.*, **3 (4)** .

Bruss, A., and Horn, B. K. P, [1983], "Passive navigation," *Comp. V. Graph. and Im. Proc.*, **21** , 3-20.

Burt, Peter J, [1981], "Fast filter transforms for image processing," *Comp. Graph. and Im. Proc.*, **16** , 20 - 51.

Burt, Peter J, [1983], "Fast algorithms for estimating local image properties," *Comp. Graph. and Im. Proc.*, **21** , 368 - 382.

Canny, John Francis, [1983], "Finding Edges and Lines in Images," *MIT Artificial Intelligence Lab. TR-720.*

Courant R., and Hilbert D, [1953], *Methods of mathematical physics*, John Wiley Interscience, New York.

Do Carmo, Manfredo P, [1976], *Differential Geometry of Curves and Surfaces*, Prentice-Hall, Englewood Cliffs, NJ.

Faugeras, O. D., *et. al*, [1982], Towards a flexible vision system, *Robot vision*, ed. Pugh, Alan, IFS, UK.

Faugeras, O. D., *et. al*, [1984], Object representation, identification, and positioning from range data, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 425 - 446.

Faux, I. D., and Pratt, M. J, [1979], *Computational Geometry for Design and Manufacture*, Ellis Horwood, Chichester.

Grimson, W. E. L, [1981], *From Images to Surfaces: a Computational Study of the Human Early Visual System*, MIT Press, Cambridge.

Grimson, W. E. L, [1984], Computational experiments with a feature based stereo algorithm, MIT Artificial Intelligence Laboratory, AIM-762.

Haralick, R. M, [1980], "Edge and region analysis for digital image data," *Comp. Graph. and Im. Proc.*, **12** , 60 - 73.

Haralick, R. M, Watson, Layne T., and Laffey, Thomas J, [1983], "The topographic primal sketch," *Int. J. Robotics Res.*, **2** , 50 - 71.

Hilbert, D., and Cohn Vossen, S., [1952], *Geometry and the Imagination*, Chelsea, New York.

Holland, Steven W., Rossol, L., and Ward, Mitchell R, [1979], CONSIGHT 1: a vision controlled robot system for transferring parts from belt conveyors, *Computer vision and sensor based robots* eds. Dodd, G. and Rossol, L. Plenum Press.

Ikeuchi, K., and Horn, B. K. P, [1981], "Numerical shape from shading and occluding boundaries," *Artificial Intelligence*, **17** , 141- 185.

Ikeuchi, Katsushi, Horn, B. K. P., et. al., [1984], Picking up an object from a pile of objects, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 139 - 162.

Koenderinck, Jan J, [1984], What tells us the contour about solid shape?, Dept. Medical and Physiol. Physics, Univ. Utrecht, Netherlands.

Koenderinck, Jan J., and van Doorn, Andrea J, [1982], "The shape of smooth objects and the way contours end," *Perception*, **11** , 129 - 137.

Marr, D, [1976], "Early processing of visual information," *Phil. Trans. Roy. Soc. London*, **B275** , 843 - 524.

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

Marr, D, [1982], *Vision*, Freeman.

Marr, D., and Hildreth, E. C., [1980], "Theory of edge detection," *Proc. Roy. Soc. London*, **B207** , 187 – 217.

Mayhew, J, [1983], Stereopsis, *Physical and Biological Processing of Images*, Braddick, O. J., and Sleigh, A. C.(eds.), Springer, New York, 204 – 216.

Millman, Richard S., and Parker, George D, [1977], *Elements of Differential Geometry*, Prentice-Hall.

Poggio, T., and Torre, V., [1984], Ill-posed problems and regularization analysis in early vision, MIT Artificial Intelligence Laboratory, AIM-773.

Porter, G., and Mundy, J, [1982], Non-contact profile sensor system for visual inspections, *Robot Vision*, Rosenfeld A. (ed.), Proc. SPIE, 67 – 76.

Porter, G. B., and Mundy, J. L, [1984], A model-driven visual inspection module, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 371 – 388.

Stevens, K. A, [1982], "The visual interpretation of surface contours," *Artificial Intelligence*, **17** , 47 – 73.

Stevens, K. A, [1983], The line of curvature constraint and the interpretation of 3-D shape from parallel surface contours, Proc. 8th Int. Jt. Conf. Artif. Intell., Karlsruehe, 1057 – 1061.

Terzopoulos, D, [1983], "The role of constraints and discontinuities in visible-surface reconstruction," *Proc. 7th Int. Jt. Conf. Artif. Intell., Karlsruehe*, 1073 - 1077.

Terzopoulos, D, [1984], The Computation of Visible Surface Representations, MIT, Artif. Intell. Lab..

Tsuji, S., and Asada, M, [1984], Understanding of three-dimensional motion in time-varying imagery, *First International Symposium on Robotics Research* Brady, Michael, and Paul Richard (eds). MIT Press, 465 – 474.

Ullman, S, [1978], *The Interpretation of Visual Motion*, MIT Press, Cambridge Mass..

Vilnrotter, F., Nevatia, R., and Price, K. E, [1981], Structural analysis of natural textures, *Proc. Image Understanding Workshop* Baumann L. S. (ed.), 61-68.

Weatherburn, C. E., [1927], *Differential Geometry of Three Dimensions*, Cambridge University Press, Cambridge, UK.

Witkin, A., [1981], "Recovering surface shape and orientation from texture," *Artificial Intelligence*, **17** , 17 – 47.

Witkin, A., [1983], "Scale-Space Filtering," *Proc. 7th Int. Jt. Conf. Artif. Intell., Karlsruehe*, 1019 - 1021.

Woodham, Robert J, [1980], "Photometric method for determining surface orientation from multiple images," *Opt. Eng.*, **19** , 139 – 144.

Woodham, Robert J., [1981], "Analysing images of curved surfaces," *Artif. Intell.*, **17** , 117 – 140.

Yuille, A. L., and Brady, Michael, [1984], Surface information from boundary projections, MIT Artificial Intelligence Lab (forthcoming).

Yuille, A.L. and Poggio, T, [1983a], "Fingerprints Theorems for Zero-Crossings," *MIT Artificial Intelligence Laboratory AIM-730.*

Yuille, A.L. and Poggio, T, [1983b], "Scaling theorems for zero crossings," *MIT Artificial Intelligence Laboratory AIM-722.*