Local Shape Approximation from Shading

Daphna Weinshall*
Department of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, ISRAEL
Internet: daphna@cs.huji.ac.il

Abstract

Shading can be used as an independent cue for exact shape recovery, or it can be used as a supplementary cue for shape interpolation between features, whose depth is known from other cues. Exact shape cannot be inferred from a local analysis of shading. However, for shape interpolation, a crude local approximation may be sufficient. This paper explores the limits of such local approximations that are easy to compute. In particular, the shape of shading is used to approximate the surface in areas of monotonic change of intensity. This analysis is accompanied by a method to compute the direction of a single point light source from the shading on occluding contours. A qualitative classification of shape near shading singularities is also discussed.

1 Introduction

Shading can be used as an independent cue for exact shape recovery, or it can be used as a supplementary cue for shape interpolation between features, whose depth is known from other cues (such as occluding contours or stereo).

Psychophysical evidence support the idea that the second role of shading, as a secondary cue for surface interpolation, may be more important to humans than its role as a primary cue, possibly because exact shape from shading is very difficult. One clue follows from the observation that humans perceive shaded images of Lambertian surfaces with mutual illumination (or interreflections) as far more realistic than in the absence of mutual illumination. Mutual illumination can severely degrade the accuracy of shape from shading algorithms (see [21] for an exception). On the other hand, mutual illumination can be used to detect other shape cues, such as occluding contours in shaded regions [11].

Human psychophysics also supports the idea that shading may not necessarily be used to compute an exact depth map, but rather an approximate surface that provides relative depth information more accurately than exact depth. One study shows that humans perceive relative depth from shaded images more reliably than absolute depth [27]. Other studies show that the ability of humans to infer shape from shading is weak[1] [6].

*This work was done at MIT, at the Center for Biological Information Processing.

[1] B. Cumming, personal communication

The problem of inferring shape from shading proved to be one of the most difficult in low-level vision. Even with the simplest Lambertian shading model, which is an idealization criticized as irrelevant since it does not necessarily approximate reality [11], and assuming constant albedo and a single point light source, significant ambiguity remains. For example, there exist concave, convex and saddle-like surfaces that appear the same from certain viewpoints[2]. Therefore most exact shape from shading algorithms involve the propagation of information from special locations on the surface, such as occluding contours or intensity singularities, where the surface is assumed known [32, 14, 12, 22, 26]. Techniques such as the method of characteristic strips were used to solve systems of differential equations with known initial conditions. Typically these techniques are computationally expensive, requiring thousands of iterations, and rarely is it possible to prove that they converge.

Because of these unattractive properties of global shape from shading algorithms, local shading analysis under some special circumstance was studied: spherical surfaces were discussed in [24, 18]; slanted surfaces, where the angle between the light source and the normal is sufficiently large, were discussed in [23]. Others have studied the behaviors of isophotes (lines of equal brightness) and their relationship to geometrical invariants of surfaces [17, 33, 4]. The present paper follows these directions in exploring approximations to shape using easy to compute local shading analysis (see also [10]).

The local shading approximation to shape may be useful for various limited purposes. Some examples are the following:

- If the task requires a a simple transformation, e.g. the prediction of the image of the surface illuminated from an infinitesimally different direction, the approximation may be sufficient. In that case, the transformation is very simple: the rotation of the shading surface by the amount of rotation of the light source.

- If an exact iterative shape from shading algorithm is to be performed, the shading approximation can give a better initial guess of the surface.

- The shading approximation gives relative depth of points on the surface in an unknown coordinate system, which may be sufficient for the computation of bumps and other surface features.

The rest of this paper is organized as follows: A qualitative classification of surfaces near some intensity singularities is discussed in section 2. An approximation of shape by the shading function itself is explored in section 3. Section 4 deals with the computation of the direction of the light source. Finally, the information that can be obtained on the surface directly from the shading, making use of the shading approximation discussed in section 3 and other cues, is discussed in section 5.

[2] e.g., under orthographic projection and when the light source is behind the viewer, the surfaces $z = x^2 + y^2$, $z = x^2 - y^2$, and $z = -x^2 - y^2$ all appear the same.

2 Geometrical properties of surfaces near global shading maxima:

Many shape from shading algorithms need, or can take advantage of, information on the shape near shading singularities (e.g. [13, 22]). A qualitative analysis of shading extrema was suggested by Koenderink & van Doorn [17], showing that these singularities often cling to parabolic points. In this section, a qualitative classification of the shape of surfaces near global shading maxima is obtained from the change in the appearance of isophotes in the image with a change of viewpoint.

Consider a Lambertian surface illuminated by a single light source (or a hemispherical sky), so that the irradiance function is proportional to $\vec{R} \cdot \vec{N}$, at a point where the normal to the surface is $\vec{N}$ and the light source direction (or the vertical direction) is $\vec{R}$. In most scenes, assuming it is not on the boundary of the surface, global shading maxima occur when $\vec{N}$ is parallel to $\vec{R}$. Choose such a point P. Let us use a coordinate system where the Z axis parallels $\vec{R}$, P defines the origin, and the X and Y axes are the directions of principal curvatures on the surface at P. With this selection, for $|\kappa_1| \leq |\kappa_2|$ the principal curvatures at P, the surface can be written as follows:

$$z(x,y) = \frac{\kappa_1 x^2 + \kappa_2 y^2}{2} + \epsilon, \quad \text{where} \quad \lim_{(x,y)\to(0,0)} \frac{\epsilon}{x^2+y^2} = 0. \tag{1}$$

If P is parabolic, namely $\kappa_1 = 0$, the isophotes near P are straight lines. If $|\kappa_1| = |\kappa_2|$, the isophotes are circles on the $X-Y$ plane (figure 1(a)), planar curves if P is elliptic ($\kappa_1 = \kappa_2$) and non-planar 3D curves if P is hyperbolic ($\kappa_1 = -\kappa_2$). Active exploration can be used to determine when the isophotes are planar curves (the parabolic and elliptic cases) and when they are not (the hyperbolic case) in the following way:

When the isophotes are parallel planar curves, they are projected obliquely to the image plane, and all of them are uniformly foreshortened. This affine transformation is responsible for contraction along one axis of the planar curves, as is illustrated in appendix A. Foreshortening preserves many curve features: straight lines are projected to straight lines, elliptical close contours to elliptical close contours, and inflection points to inflection points. Moreover, this projection is stable in the sense that as the camera (or the light source) moves, the curves uniformly get more or less contracted along a single axis.

The case is different for isophotes that are not planar curves, and therefore their projection to the image plane is not foreshortening. The further away the isophotes are from parallel planar curves, the less like simple single-axis contraction their projection to the image plane looks.

An example of the projection of a particular isophote on an elliptic surface with $\kappa_1 = \kappa_2$, and an example of the projection of the corresponding isophote on a hyperbola with $\kappa_1 = -\kappa_2$, are given in figure 1. As $|\kappa_1|$ gets further away from $|\kappa_2|$, the distinction between the hyperbolic and elliptic cases becomes less sharp. The projected isophotes still differ in the same way, however, as is shown in figure 2: in b, a hyperbolic surface where $\kappa_1 = -\frac{2}{3}\kappa_2$, the triangular shape of the isophotes (figure 2(a)) is maintained; in c, an elliptic surface where $\kappa_1 = \frac{2}{3}\kappa_2$, the isophotes look more like foreshortened ellipses.

This analysis suggests a heuristical method for the classification of surfaces near global shading maxima (that arise from the diffuse, rather than the specular, component of the illumination):

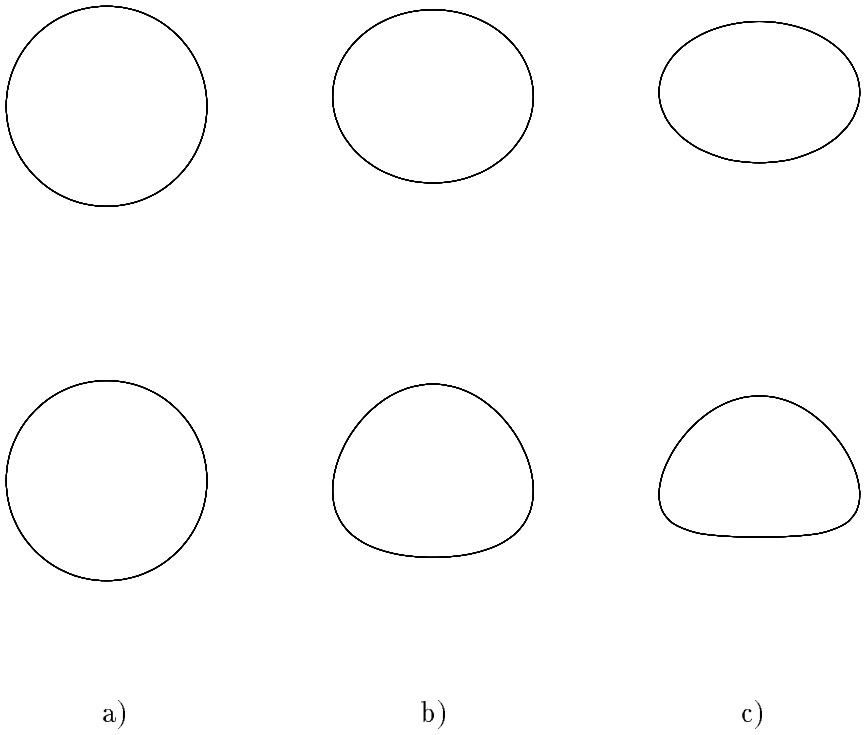

a) b) c)

Figure 1: The projection onto the image of an isophote on the surface $z = -(x^2 + y^2)$ (elliptic, above) and $z = x^2 - y^2$ (hyperbolic, below), illuminated by a single light source; a) viewing angle is parallel to $\vec{R}$, b) the angle between viewing angle and $\vec{R}$ is 30^o, c) the angle between viewing angle and $\vec{R}$ is 45^o.

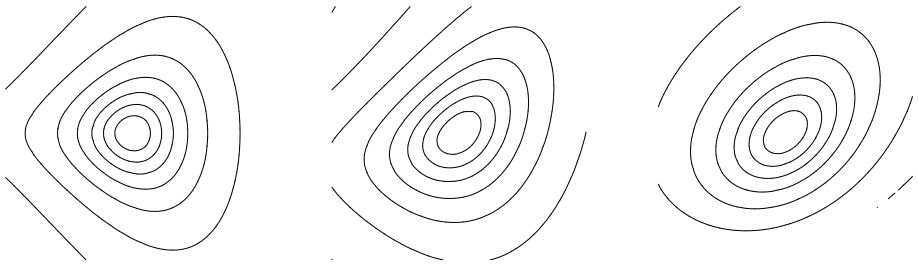

a) b) c)

Figure 2: The isophote map, projected onto the image, near a shading maximum. The viewing coordinate system is obtained by a rotation of 45^o around the Z axis, followed by a rotation of 10^o around the X. *a)* A generic hyperbolic surface $z = \frac{x^2-y^2}{2}$; *b)* a generic hyperbolic surface $z = \frac{x^2-\frac{2}{3}y^2}{2}$; *c)* a generic elliptic surface $z = -\frac{x^2+\frac{2}{3}y^2}{2}$.

- measure the projection of the isophotes in the image plane and their change with change of camera location;
- if the isophote projections are straight lines locally, the surface is parabolic (figure 3(a));
- if the isophote projections can be approximated by foreshortened circles, and their change can be approximated by contraction/expansion along a single axis, the surface is elliptic (figure 3(b));
- if the isophote projections can better be approximated by concentric triangles, the surface is hyperbolic (figure 3(c)).

This method requires oblique viewing and can be used more effectively with active exploration. It is assumed that $\kappa_1 = 0$ or $\kappa_1 = \kappa_2$. The classification becomes less reliable as κ_1 gets further away from κ_2.

3 A local approximation to shape from shading

In this section, the limits of the shape of the shading as an approximation to shape are studied. It is shown that if a linear transformation between the shapes of the shading and the surface is allowed, there exist families of surfaces for which equality holds. One such family is the locally-spherical surfaces, the only surfaces for which local shape from shading can be computed precisely [24]. In the case of such surfaces, and others, the shading is the shape, very little additional computation is required. For additional families of surfaces, the individual isophotes, or lines of equal brightness, are identical to the contour lines on the surface. The shading approximation to shape, proposed in this section, is the approximation of surface patches by members of these families.

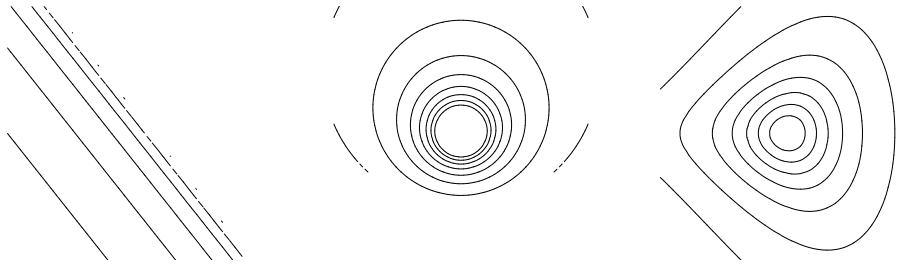

a) b) c)

Figure 3: Classification of surfaces near global shading maxima: *a)* parabolic, *b)* elliptic, *c)* hyperbolic.

A depth map of a surface can be represented by a contour map, a map of the lines of equal depth on the surface. We are interested in partial representations, where only the contour lines of the map are given, and a scalar function F is defined over the contour lines. When F is a monotonic function of the depth of the contour lines, this representation gives the relative depth between any two points on the surface. When F is a linear function of the depth, this representation gives the depth of the surface at each point up to a single scaling factor. The last two cases will now be discussed.

The question addressed in this section is when the isophotes (lines of equal brightness on the surface) are also the contour lines (lines of equal depth on the surface), and what is the function F that relates the intensity of the isophotes to the depth on the contours. Local stability analysis is used to detect cases where this shading approximation should not be used.

3.1 The 3D shading function:

Assume a Lambertian surface $z(x, y)$ with fixed albedo and a single light source or a uniform hemispherical sky (ambient illumination can also exist.) The reflectance function of a Lambertian surface depends only on the angle between the light source direction and the surface normal, ϵ in figure 4. The brightness of this surface at a point with normal $\vec{N}$ is:

$$I = \alpha \vec{R} \cdot \vec{N} + \beta \ , \tag{2}$$

for some constants α, β and a fixed direction $\vec{R}$, which is either parallel to the light source direction or vertical.

Since the surface is assumed to be Lambertian, the position of the viewer does not affect the shading at each point on the surface. Therefore it is possible to consider the shading as a scalar function defined on a surface in 3D, before computing the projection onto the image plane. Henceforth in this section, isophote map will refer to this function defined on a surface in 3D. The discussion of the projection to 2D is postponed to the next section.

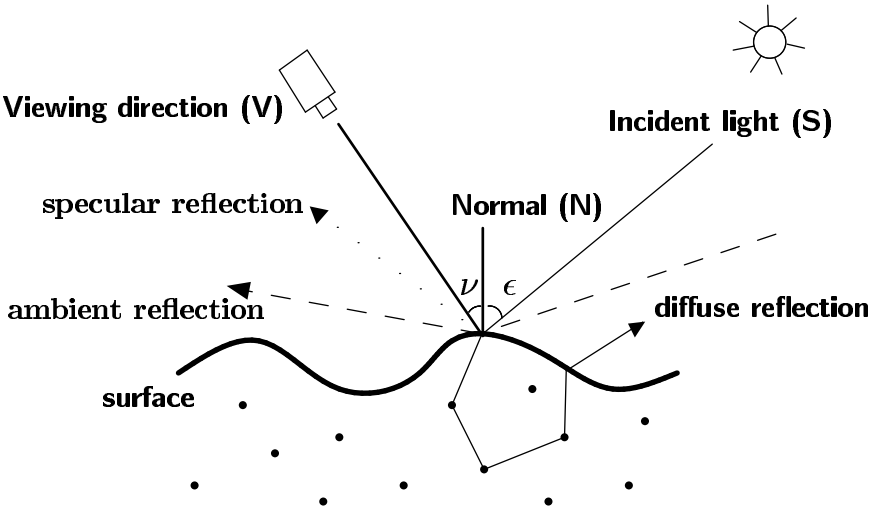

Figure 4: The reflectance function of a surface at a point is composed of three basic components: diffusive, specular, and ambient.

For the analysis of the isophote map, which is independent of the viewer position, a coordinate system is used where the depth Z is parallel to $\vec{R}$. This coordinate system is denoted $\Re$. With this selection of coordinate system, $\vec{R}$ in (2) is $(0,0,1)$. The shading at point $(x,y,z(x,y))$, with normal $(\frac{z_x}{\sqrt{1+z_x^2+z_y^2}}, \frac{z_y}{\sqrt{1+z_x^2+z_y^2}}, \frac{-1}{\sqrt{1+z_x^2+z_y^2}})$, is $I = \frac{1}{\sqrt{1+z_x^2+z_y^2}}$.

Take a region where the intensity changes monotonically. The isophotes are contour lines, and their values change in the same (or inverse) direction, if there exists a monotonic function $\Phi(z)$ such that $\Phi(z) = \frac{1}{\sqrt{1+z_x^2+z_y^2}}$. $\Phi(z)$ exists whenever $z_x^2 + z_y^2$ is some monotonic function of z. Thus the problem can be rephrased as follows: the shading depends monotonically on the depth when there exist a monotonic function $h(z)$ such that

$$z_x^2 + z_y^2 = h(z) \; . \tag{3}$$

There are (at least) two solutions to this differential equation: radially symmetric surfaces in $\Re$ and unidirectional surfaces. The isophotes of the first are concentric circles (figure 14), the isophotes of the second are straight lines (figure 17). If $h(z)$ is linear, the surfaces that satisfy equality in (3) are restricted to tori in $\Re$ and any sphere or cylinder. These results are discussed and illustrated in appendix B.

For a given surface patch, the shading approximation to the contour lines of the surface (monotonic h) or its unscaled depth (linear h) is equivalent to using a patch from a surface from these families to approximate the surface patch. In the following section, a method to determine when this approximate gives acceptable results is presented.

3.2 Stability analysis of the shading approximation

Assume oblique viewing, so that the projection of the isophotes to the image plane is foreshortened. Assume also orthographic projection. Consider first the case where isophotes are contour lines. Here the isophotes on the surface are planar curves lying on parallel planes $Z = \text{const}$ in some unknown Cartesian coordinate system. These parallel planar curves are projected obliquely to the image plane, and all of them are uniformly foreshortened. This foreshortening is responsible for contraction along one axis of the planar curve, as discussed in appendix A.

The case is different for isophotes that are not contour lines. Here the isophotes are not planar curves, and therefore their projection to the image plane is not foreshortening. We can therefore use the method of active exploration discussed in section 2 to detect the "bad" cases where the isophotes are not a good approximation to shape. By moving the camera, if the isophotes do not change by a uniform contraction and expansion along a single axis, if features such as inflection points appear and disappear, then the shading should not be used to approximate shape.

3.3 An approximate local shape from shading algorithm:

The contour map is approximated in regions where the brightness changes monotonically. Based on the above analysis, an algorithm to construct a shape representation is the following:

1. Find curves or points of brightness extrema;
2. compute a contour map in each region between these curves using the isophote map;
3. assign direction of depth change (whether the depth increases with increasing brightness or decreases);
4. Compute the relevant surface features.

In a minimalistic implementation of this algorithm, the depth on curves or points of brightness extrema that have been found in step 1 is either computed using other cues, such as motion or occluding contours, or set at an arbitrary constant, and the shading is then used to interpolate the depth between these features.

3.4 Discussion

The approximation to shape by shading discussed here is in some sense a generalization of Pentland's analysis of local shape from shading [24]. Pentland showed that only spherical points can be precisely recovered with local shading analysis. The analysis in section 3 shows that local shading can be used to give the shape of a richer family of surfaces, not limited only to spherical. It also gives an integrable consistent solution all over the surface even when the local approximation varies, e.g., when each surface patch is approximated by a different function. Finally, and most importantly, the shading approximation is computationally free, only the two parameters of the light source direction should be computed to obtain a viewer-centered depth map.

More specifically, the isophote map approximates the contour map in the coordinate system $\Re$, which is often different from the viewer-centered coordinate system. Thus typically the map of contour lines is foreshortened when projected into the image, as discussed above. This foreshortened map can be used to infer properties of the surface directly, as can be imagined from figure 5 (and will be discussed in section 5).

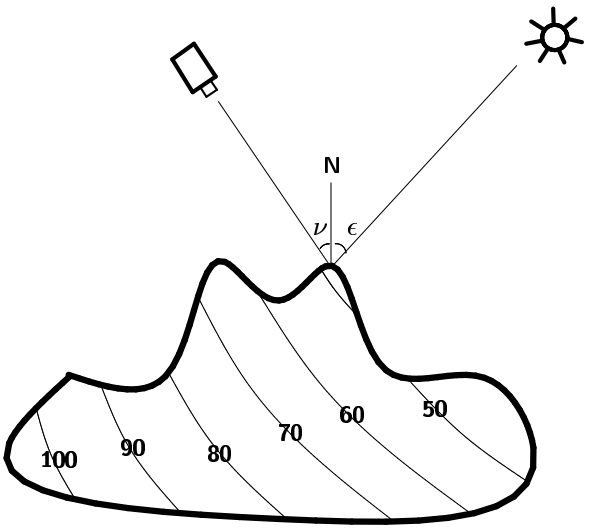

Figure 5: A Lambertian surface illuminated by a single point light source. In this case, the coordinate system in which the shading is the depth is the coordinate system of the light source.

4 Direction of illumination from occluding contours

If the surface is illuminated by a single point light source, the shading approximation discussed in section 3 is obtained in the coordinate system of the light source. If the exact depth map of the surface is needed rather than surface features, it may be necessary to compute the light source direction. In this section, direct computation of the light source direction from the shading on occluding contours in smoothly receding objects in the image is discussed (see also [5, 24, 25]).

Assume a single light source at a large distance from the surface. The source's position is defined by two angles (figure 6): tilt – the angle between the projection of the light direction on the image plane (denoted the $X-Y$ plane) and the X-axis, and slant – the angle between the light direction and the Z-axis (the viewing direction). In the following, the computation of these two angles from shading on occluding contours and self-shadow edges is discussed. This computation is mostly based on a general shading model with Lambertian, ambient and specular components.

4.1 Tilt computation

Zheng & Chellappa [34] have discussed the computation of the tilt of the illuminant from shading on occluding contours. Their scheme involves an integration of the shading information over the entire occluding contour. The following technique computes the tilt of the illuminant from local properties of the shading on occluding contours, and can be used when the contour is partially occluded.

The following computation assumes a single point light source and constant albedo. If more than one light source exists, the algorithm can be used with the self-shadow edges of each light source separately.

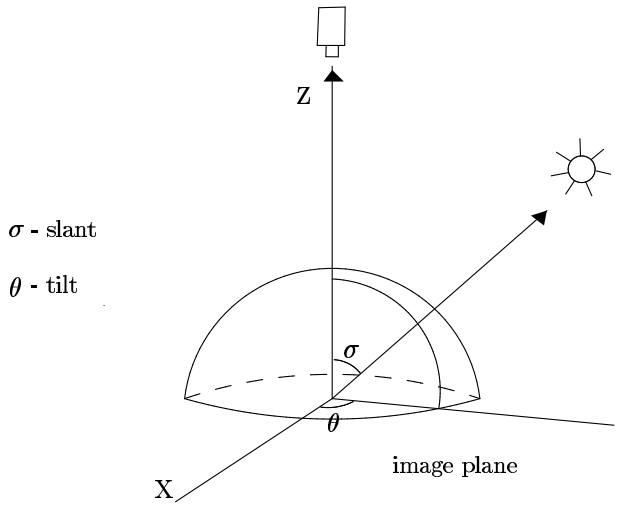

Figure 6: The tilt and slant of a point light source.

Tilt from shading on occluding contours:

Assume an occluding contour where the normal to the surface is perpendicular to the viewing direction (in the $X - Y$ plane).

Proposition 1 *The direction of the tilt of the illuminant is parallel or perpendicular to the occluding contour at points of shading extrema on the contour. This holds for points of extremum where the occluding contour does not have a singularity (such as a cusp).*

1. *At any point on the occluding contour where the minimum of shading occurs (must be the value of the ambient illumination, at the beginning of a self-shadow line), the angle of the tangent to the occluding contour (and the self-shadow line) is the angle of tilt of the illuminant.*

2. *At any point on the occluding contour where the maximum of shading occurs, the angle of the tangent to the occluding contour is perpendicular (in the image plane) to the angle of tilt of the illuminant.*

This proposition identifies a group of points on the occluding contour (at least 3, if all the occluding contour is visible) where the tangent to the occluding contour gives the angle of tilt of the illuminant. The tilt computation is therefore quite robust, since only one of these points should be visible (not occluded). Figure 7 shows an example of the occluding contour for a complex surface and the tangent to the occluding contour at the extrema of shading on the contour.

Tilt from self-shadow edges:

If both edges of self-shadow and cast-shadow are given, matching them can give the direction of illumination as well. As pointed out in [25], the angle of a line between a feature on a self-shadow line and its match on the cast-shadow line is the angle of tilt.

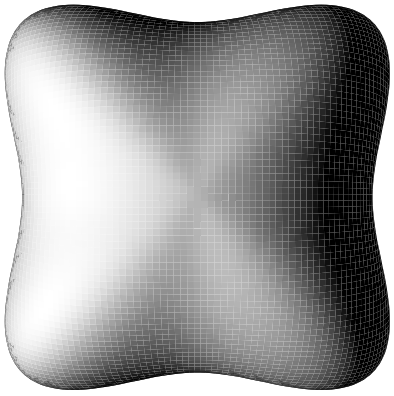

a)

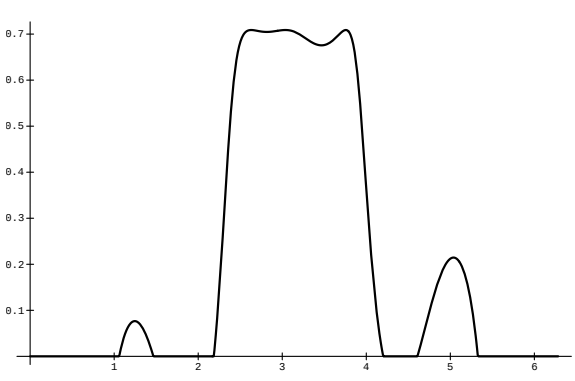

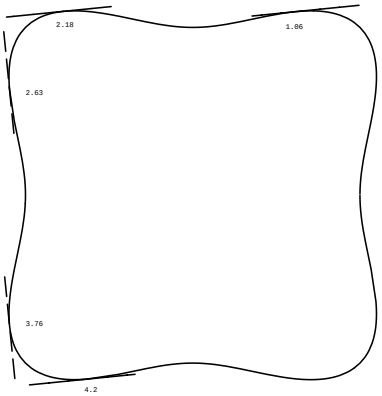

b) c)

Figure 7: a) A bird's eye view of a surface illuminated from direction $(-1, -0.1, 1)$. b) The shading on the occluding contour plotted as a function of polar angle (in radians) around the center of the object. c) The occluding contour of the surface with its tangent plotted at 3 points of shading minima (solid line) and 2 points of shading maxima (dashed line). Each point is identified by its polar angle (as plotted in b).

4.2 Slant computation

Slant from shading on occluding contours:

The occluding contour in the neighborhood of a point of global maximum of shading gives the direction of tilt of the illuminant. If the albedo of the surface and the intensity of the illuminant are known, the shading at such a point, which does not have to be unique, gives the angle of slant of the illuminant. If the albedo and light intensity are not known, the ratio between the global maximum of shading on the occluding contour to the global maximum of shading on the surface gives the angle of slant according to the following formula:

$$\cos\sigma = \frac{\max\limits_{\partial\Omega} I}{\max\limits_{\Omega} I} \tag{4}$$

for brightness I, slant σ, surface of an object Ω, and the occluding contour $\partial\Omega$.

This formula assumes a Lambertian surface. It may be used by humans to judge the slant of the light source, as discussed by Reichel & Todd in [27]. In general, $\min\limits_{\partial\Omega} I$ should be subtracted from both numerator and denominator to eliminate the component due to ambient illumination. If the surface is also specular, the denominator becomes $\max\limits_{\Omega\prime} I$ where $\Omega\prime$ is the surface area not including the regions of specular reflection (the use of $\Omega\prime$ may lead to an underestimation).

4.3 Results with a real image

Figure 8 shows an example, an image of a gourd, for which the direction of illumination was computed from the shading on the occluding contour. First, regions of maximal and minimal intensity on the occluding contour were identified (figure 8(c),(d)). The self-shadow edge in figure 8(b) does not actually intersect the occluding contour: the intersection was computed by locating the points on the occluding contour whose intensity was the closest to the mean intensity on the self-shadow edge (which was almost constant). The tangent to each of the edge segments in figure 8(c),(d) was computed, giving an estimate of -104^o to the tilt direction of the illuminant. The angle of slant was computed from (4) to be 75^o.

5 Geometrical properties of surfaces

Geometrical information on surfaces that can be obtained directly from the shading approximation will be discussed in this section. The ability of humans to judge shape information without knowledge of the light source direction was demonstrated and discussed by Mingolla and Todd in [19]. Humans also seem to perform better in tasks that require relative depth, which can be readily obtained from the shading approximation, than absolute depth [27]. Relative depth will not be discussed further, although the existence of ordering of image points (in any coordinate system) can be used to locate bumps and other important surface features.

Surfaces of objects can sometimes be concisely described as a collection of simpler parts, each of which described by a few parameters (e.g. generalized cylinders [2]; see also [16] and [9]). Classifying regions according to the sign of their Gaussian curvature, namely as elliptic (convex/concave),

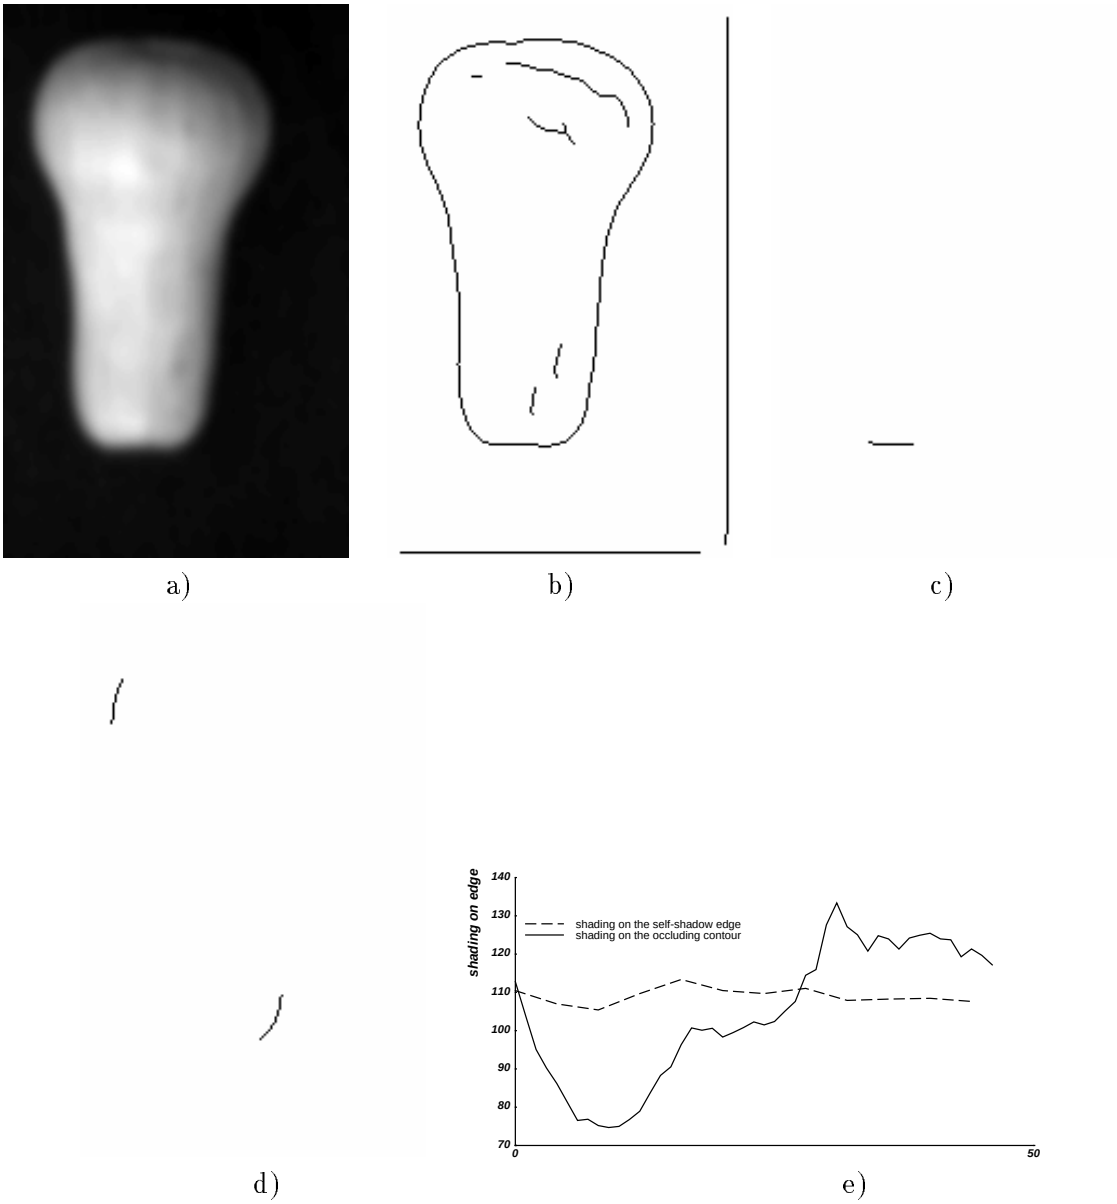

Figure 8: a) A gourd image; b) the edges of the image computed using Canny's algorithm [7]; c) a piece of the occluding contour where the intensity is maximal; d) two pieces of the occluding contour where the intensity is minimal (the beginning of a self-shadow edge); e) the (smoothed) shading profiles of the occluding contour and the self-shadow edge. The scale and absolute values of the units on the abscissa are unrelated for the two edges. The points of interest on the occluding contour are the maximum (giving the edge element in c) and the intersection of the contour line with the self-shadow edge (giving the edge elements in d).

planar, cylindrical, or hyperbolic, provides one important intrinsic surface feature (see also [1], [28], [30], [20], [3] and [29]). With this classification of parts as areas of the same sign of Gaussian curvature, part boundaries within an object are located on parabolic lines. The parts produced by this segmentation are often qualitatively similar to the parts produced by a generalized-cylinders-based scheme. I will now discuss cues in the contour map to the Gaussian curvature of the surface.

5.1 Geometrical properties of surfaces near inflection points:

5.1.1 Contour lines:

Let P be a point on the contour line $z = \text{const}$. Assume that the $X - Y$ plane is not tangent to the surface at P[3]. There exists a direction (which will be called Y without loss of generality) such that $z_y \neq 0$. In the neighborhood of P the contour line is some function $y(x)$, such that:

$$\begin{aligned} y'(x) &= -\frac{z_x}{z_y} \\ y''(x) &= -\frac{1}{z_y^3}(z_y^2 z_{xx} - 2z_x z_y z_{xy} + z_x^2 z_{yy}) \ . \end{aligned}$$

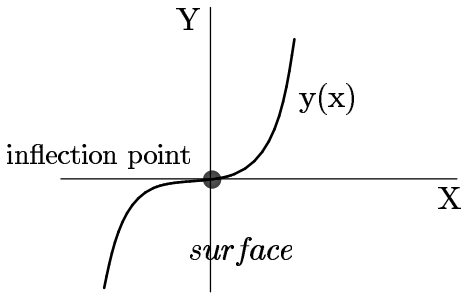

Figure 9: Inflection points in the contour lines can be seen from almost any viewpoint. They will correspond to inflections in the isophotes for surfaces whose isophotes are their contour lines.

In an inflection point $y''(x) = 0$, namely

$$z_y^2 z_{xx} - 2z_x z_y z_{xy} + z_x^2 z_{yy} = 0.$$

This equation has a solution if $z_{xy}^2 - z_{xx} z_{yy} \geq 0$, namely, the Gaussian curvature of the surface is not positive. In other words, the surface around an inflection point on a contour line is not elliptic, it is parabolic or hyperbolic.

The $X-$ and $Y-$ directions can be chosen so that $z_y \neq 0$ and $z_x = 0$ (figure 9). Then

$$y'(x) = 0, \quad y''(x) = -\frac{z_{xx}}{z_y}, \quad y'''(x) = -\frac{z_{xxx}}{z_y} \ . \tag{5}$$

[3] Inside a region where the intensity changes monotonically, it cannot be that $z_x = z_y = 0$. Thus the $X - Y$ plane is not the tangent plane to the surface at an internal point P.

P is an inflection point if and only if $y''(x) = 0$ and $y'''(x) \neq 0$, namely, $z_{xx} = 0$ and $Z_{xxx} \neq 0$. Thus z_{xx} changes sign at P whereas z_{yy} almost always does not. We can therefore conclude that any inflection point is one of the following:

- A hyperbolic point (less common);
- A parabolic point dividing an elliptic region from a hyperbolic region, and X is a principal direction (an example is given in figure 10).

Once again, the second case, where the inflection point is parabolic, can be identified by stability analysis. An inflection point in a hyperbolic region will move in all directions when the camera, or the light source, is moved a little around its location. An inflection point on a parabolic line will move on the parabolic line or disappear.

5.1.2 Isophotes are contour lines:

Let $(z(x, y)$ be the actual surface and $I(x, y) = f(z)$ the shading surface for some monotonic function f (i.e. $f'(z) > 0$). The Gaussian curvatures K of both surfaces are related as follows:

$$K_I = [f'(z)^2 \; K_z \; + \; f''(z)f'(z)(z_y^2 z_{xx} - 2z_x z_y z_{xy} + z_x^2 z_{yy})] \frac{(1 + I_x^2 + I_y^2)^2}{(1 + z_x^2 + z_y^2)^2} \; . \tag{6}$$

We can now conclude the following:

- From (6) it follows that near an inflection point, the sign of the Gaussian curvature of the shading surface $I(x, y)$ is identical to the sign of the Gaussian curvature of the surface $z(x, y)$. The shading can be used to determine whether the surface is locally parabolic or hyperbolic.
- From (6) it also follows that the sign of the Gaussian curvature of the shading surface $I(x, y)$ is identical to the sign of the Gaussian curvature of the surface $z(x, y)$ whenever $f''(z) = 0$, namely, when the shading is a linear function of the depth (cylinders and spheres).
- An inflection on an isophote is either a hyperbolic or parabolic point dividing an elliptic region from a hyperbolic one.

  In the reverse direction: on a parabolic point on an isophote, the X direction as define in (5) is a principal direction [33], namely, $Z_{yy} = 0$ or $Z_{xx} = 0$. If $Z_{xx} = 0$, P is an inflection point.

Finally, if $\vec{R}$ is not parallel to the image plane, the inflection point on the contour line (which is a planar curve) is projected to an inflection point on the projection of the contour line on the image. It can therefore be detected regardless of whether $\vec{R}$ is known.

5.1.3 An example:

Assume a close contour line. Since $z_y \neq 0$, any sign change of the curvature of the contour line (as determines by $y''(x)$) can only happen with sign changes of z_{xx} (note that the direction X changes continuously with P). A convex segment curves towards the direction of increase in depth on the surface, a concave segment curves away from it, and both curve towards the inside of the close

contour line when the Gaussian curvature does not change sign on the curve. (Thus close elliptic contour lines often resemble ellipses, see figure 18). Since an inflection point is either hyperbolic or it is separating a hyperbolic region from an elliptic one, contour line segments that have negative curvature relative to the inside of the contour line tend to be hyperbolic, and segments with positive curvature tend to be elliptic. This is always true for occluding contours, see [15].

Figure 10 illustrates the implication of these results. Figure 10(d) in particular shows that the parabolic lines intersect the contour lines at inflection points, the regions of the contour line with positive curvature relative to its inside are elliptic, and the regions with negative curvature relative to its inside are hyperbolic. Figure 10(e) shows that these qualitative relationship hold approximately for an isophote, as well as a contour line, though for this surface the isophotes and contour lines are not identical.

5.2 Geometrical properties of surfaces from isophotes:

For some surfaces the analysis of inflection points is not very useful, since the isophotes do not approximate the contour lines well, or there are no inflection points in regions of monotonic change of intensity. In such cases it is still possible to learn some shape properties from the isophotes. This section extends the work of Koenderink & van Doorn [17] and Yuille [33].

Yuille [33] has shown that at a parabolic point the isophote points along the line of curvature (or a principal direction) at the point (see also [31] for a similar result). He concluded that "typically the parabolic lines give rise to ridges, or valleys, in the image intensity". Analysis of surfaces of revolution shows that more often the isophotes are perpendicular, or almost perpendicular, to the parabolic line. In particular, in these cases, the parabolic line is perpendicular to any ridge or valley in the image intensity.

Consider the family of surfaces of revolution (generalized cylinders where the main planar axis is a straight line). For this family it is known [8] that the parallels (lines parallel to the generator curve) and meridians (lines defined by a given point on the generator curve as it sweeps around) are lines of curvature, and that parabolic lines are meridians. Thus, for surfaces of revolution, the isophotes at parabolic points are either parallel to the parabolic line or perpendicular to it.

The first case leads to parabolic lines being ridges or valleys in the image intensity as discussed by Yuille, but it is the rarer case of the two. An example of a torus is given in figure 11 (here the parabolic line is a ridge when the light source is behind the viewer). More often the tangents to the isophotes are perpendicular to the parabolic line and the isophotes have a local extremum or an inflection point. More specifically, the isophotes bend at the parabolic line (figure 12(a) and 12(b)) if the parabolic line separates a hyperbolic region from an elliptic one. If the parabolic line lies between two elliptic regions, the isophote has only an inflection point (figure 12(c)). A proof is given in appendix C.

This result suggests a heuristic for the detection of parabolic lines separating hyperbolic and elliptic regions: when all the isophotes bend and the line through the extrema of the bend is roughly perpendicular to the isophote, this line is a good candidate for a parabolic line.

6 Discussion

The problem discussed in this paper was shape interpolation from local shading; ways to obtain a crude shape interpolation between distinct surface features were discussed. This approach is most

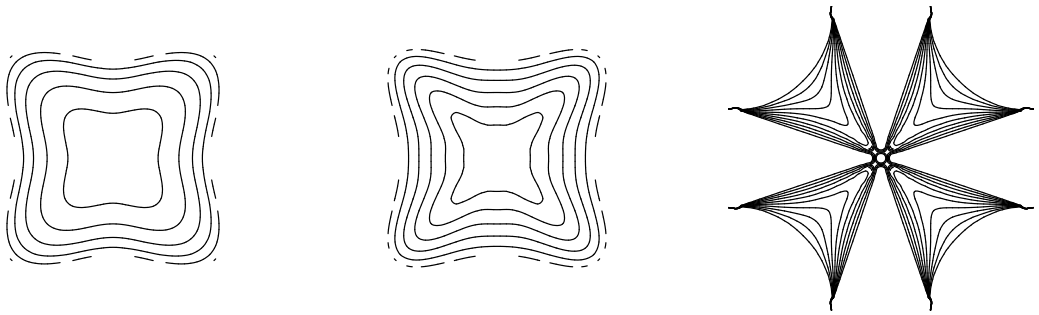

a) b) c)

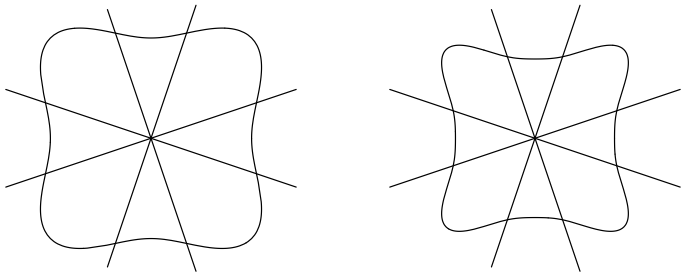

d) e)

Figure 10: For a starfish-like surface (shown in figure 7(a)): a) the contour map of the surface, b) the isophote map, c) elliptic regions marked by filling-in, d) the intersection of the parabolic lines and a contour line, e) the intersection of the parabolic lines and an isophote. The contour line in d) and the isophote in e) were chosen randomly and are unrelated.

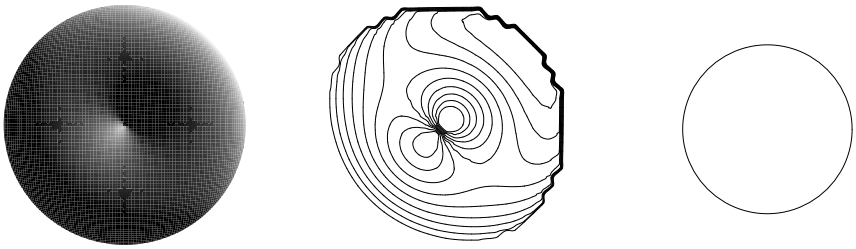

Figure 11: A torus illuminated from direction $(1,1,1)$, its isophote map, and its parabolic line.

useful when: *a)* the shading is used as a secondary cue, to interpolate a surface between features on it, whose depth is obtained using other cues (such as occluding contours or stereo); *b)* the surface interpolation need not be precise, and computation efficiency is preferred over accuracy.

The shape from shading problem, as defined here, is motivated by two observations on humans' perception of shape from shading. First, the importance of mutual illumination, a cue which complicates shape from shading but simplifies shape from occluding contours, to humans' perception of shaded images. Second, humans show higher accuracy in tasks that require only relative depth information rather than full depth information (see also section 1).

A solution to the problem of shape interpolation from shading gives a surface approximation which is computed using local information only. This approximation can be used to compute the appearance of the surface under infinitesimal changes of the light source direction, to provide an initial estimate for an exact and iterative shape from shading algorithm, or to compute qualitative surface features directly from the image.

Appendices:

A Foreshortening

When the isophotes are contour lines, they are planar 3D curves that are projected onto the image plane. This projection is described in figure 13.

Let Sr be the plane on which the isophote lie, and Im the image plane. Select directions X, Y in plane Sr so that X is parallel to the image plane Im, and Y perpendicular to X. These directions are projected to perpendicular directions X', Y' in Im. Let θ be the angle between axes Y and Y'. (X is parallel to X' by definition.) A curve on Sr of the form $y = f(x)$ is projected to a curve $y' = \cos\theta f(x')$. Thus the projected curve is foreshortened, namely, it undergoes uniform scaling along direction Y.

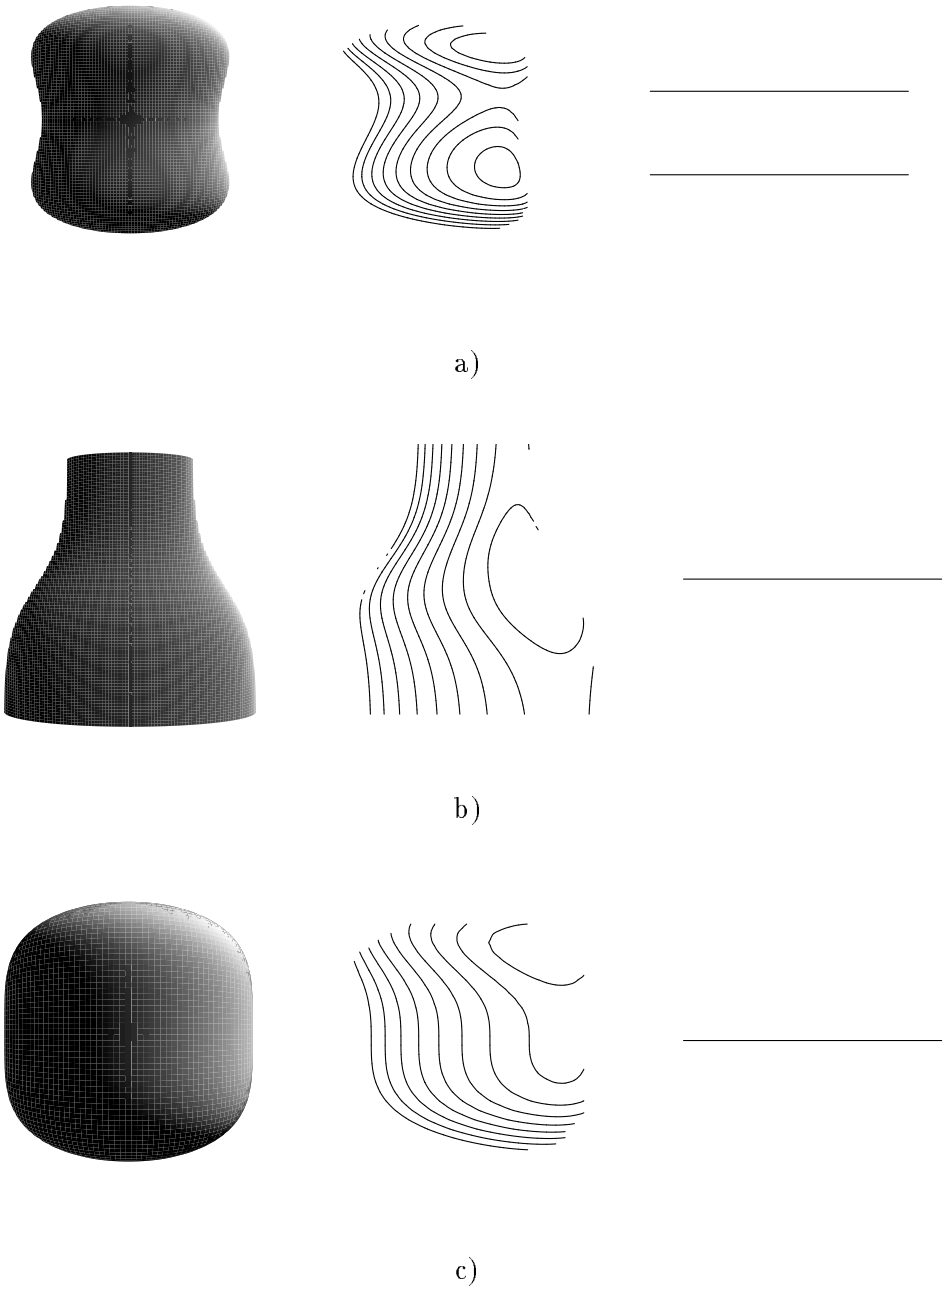

Figure 12: Each row, from left to right, shows a surface of revolution illuminated from direction $(1, 1, 1)$, its isophote map, and its parabolic lines. For illustration purposes, the object shown on the left of each row is actually illuminated from $(1, 1, 0.2)$.

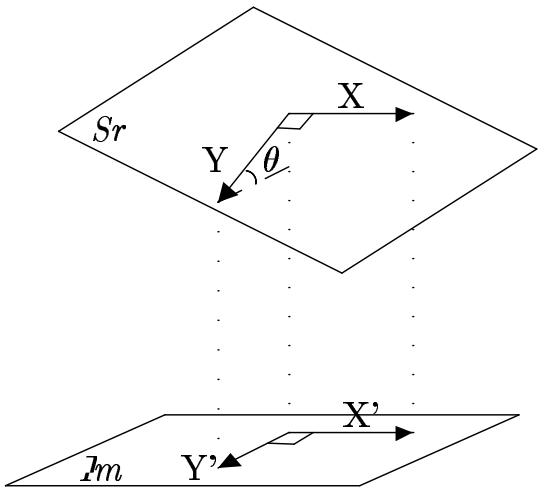

Figure 13: Illustration of foreshortening.

B Radially symmetric and unidirectional surfaces

I: Radially symmetric surfaces

Let $z(x, y)$ be radially symmetric in the coordinate system $\Re$, namely:

$$z(x, y) = z(r) \quad \text{for} \quad r = \sqrt{x^2 + y^2} \ .$$

Note that this surface, which is radially symmetric in $\Re$, is not necessarily radially symmetric in other coordinate systems.

To see why shading depends only on depth for these surfaces, note first (as can be readily verified) that:

$$z_x^2 + z_y^2 = z_r^2 \ .$$

Since $I(x, y) = I(r)$ is monotonic, z_r must be of the same sign and nonzero in the region. Thus there exists an inverse function $r(z)$. z_r is therefore a function of z only.

Figure 14 shows the contour map and isophote map of a volcano-like radially symmetric surface facing upwards. Note that the shape of the isophotes is identical to the shape of the contour lines, both concentric circles. This is the isophote shape for all radially symmetric surfaces. The spacing between the isophotes and the contour lines, which reflects the scaling of the maps, is different between the two maps.

Toruses and spheres: brightness is a linear function of depth

The sub-family of radially symmetric surfaces, whose brightness gives the depth up to a constant scaling factor, includes surfaces $z(r)$ for which the following equation has a solution:

$$I = az + b \tag{7}$$

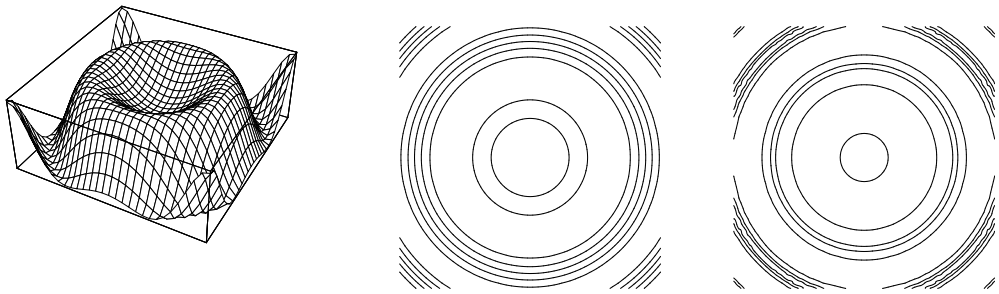

a) b) c)

Figure 14: a) An oblique view of a certain surface. b) The contour map of the surface in (a) illuminated by a uniform hemispherical sky, c) the isophote map.

for two constants a and b. The solution is (as can be readily verified):

$$z(r) = \sqrt{\lambda^2 - (r-\mu)^2} + \nu$$

for some constants λ, μ, ν. This surface is a torus when $\mu > 0$, and a sphere when $\mu = 0$. Note that a sphere is radially symmetric in any coordinate system, not only $\Re$.

Thus in a region that can be approximated by a patch on a sphere or a patch on a torus facing $\vec{R}$, the brightness itself gives almost all the information we can hope to get about z from a single image. Figure 15 shows the contour and isophote maps of a torus.

II: Unidirectional surfaces

A unidirectional surface is a surface that changes as a function of a single direction in the image, $z(x,y) = z(ax+by+c)$. The shading depends only on depth for these surfaces, with the following particular solution of (3):

$$x + cy = \int \sqrt{\frac{1+c^2}{h(z)}} dz - d \tag{8}$$

for two constants c, d. In this solution, which can be readily verified, z varies with a single direction in the image plane ($x + cy + d$). Figure 16 gives an example of a unidirectional surface.

a) b)

Figure 15: a) The contour map of a torus illuminated by a uniform hemispherical sky, b) the isophote map of the torus.

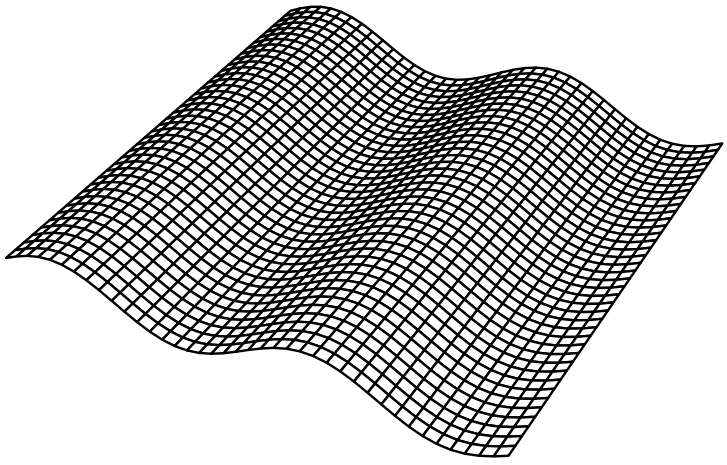

Figure 16: An oblique view of a unidirectional surface.

Cylinders: brightness is a linear function of depth

The sub-family of unidirectional surfaces, whose brightness gives a linear transformation of the depth, satisfies:

$$z(x,y) = \sqrt{\lambda^2 - \alpha((x + cy + d) - \mu)^2} + \nu$$

for some constants $\lambda, \alpha, \mu, \nu$. This surface is a cylinder, as illustrated in figure 17.

a) b)

Figure 17: a) The contour map of a cylinder illuminated by a uniform hemispherical sky, b) the isophote map of the cylinder.

III: General unidirectional surfaces

General Unidirectional surfaces are surfaces that can be described as $\tilde{z}(\tilde{x}, \tilde{y}) = \tilde{z}(\tilde{x} + a\tilde{y} + b)$ in some Cartesian coordinate system $\tilde{X}, \tilde{Y}, \tilde{Z}$ different from $\Re$. The isophote map of such surfaces is still composed of the same straight lines on the surface as characterized above. Thus the isophote map is a monotonic function of the depth $\tilde{Z}$.

IV: Other surfaces

For many surfaces $z_x^2 + z_y^2$ is not a function of z only. In many cases, the approximation of the surface by isophotes degrades gracefully as the following examples show. This issue is elaborated on in section 5, where the information in contour lines is discussed.

Ellipses:

Take an ellipse $z = c\sqrt{R^2 - (\frac{x}{a})^2 - (\frac{y}{b})^2}$ such that $a \neq b$ and the ellipse is not radially symmetric. The brightness is:

$$I = \frac{z}{c\sqrt{z^2 + \frac{x^2 c^4}{a^4} + \frac{y^2 c^4}{b^4}}}$$

$$= \frac{z}{c\sqrt{z^2 + \frac{c^2}{a^2}(c^2R^2 - z^2) + \frac{c^4y^2}{b^2}(\frac{1}{b^2} - \frac{1}{a^2})}} .$$

The only dependence on x, y is in the last term $\frac{c^4y^2}{b^2}(\frac{1}{b^2} - \frac{1}{a^2})$. Thus the isophote approximation deteriorates continuously as a gets further away from b.

Figure 18 shows an example of an ellipse. As can be seen, the isophotes approximation only amplifies the effect of $\frac{a}{b}$, the measure of anisotropy between the two principal directions of the ellipse. Ellipses behave "well" since they are an intermediate case between a sphere and a cylinder (in a certain parametric representation).

a) b)

Figure 18: a) The contour map of an ellipse, with $a = 1.5$, $b = 2$, $c = 1.25$, $R = 1$, illuminated by a uniform hemispherical sky, b) the isophote map of the ellipse.

General radially symmetric surfaces:

Figure 19 shows a rotated torus that does not face upwards (therefore it is not radially symmetric in $\Re$), its isophote and contour maps. In many regions the isophote map captures important aspects of the contour map (this is discussed in more detail in section 5), with distortions at increasing levels of severeness.

"bad" surfaces:

Finally, for some surfaces the isophote approximation is quite wrong, as shown in figure 20. This is an example of a hyperbolic saddle-like surface. It is not the case, though, that the isophote approximation fails for all hyperbolic regions, as the torus example shows.

C Differential properties of surfaces

Let us use the following coordinate system to describe a surface in the neighborhood of a point P. Let the tangent plane at P be the $X - Y$ plane, with the $X-$ and $Y-$ directions corresponding to the two principal directions at P. Let Z be the direction of the normal to the surface at P. Let the origin be at P. Let the surface be a function $z(x, y)$.

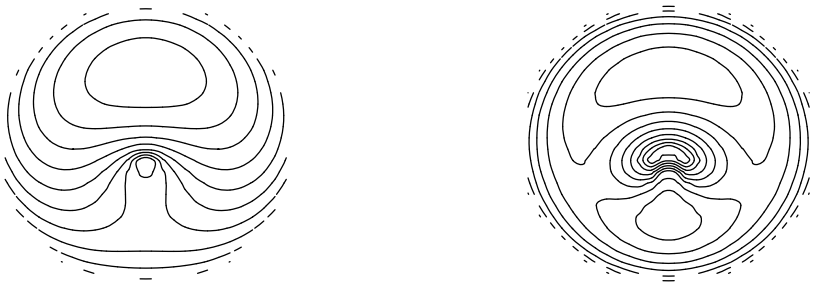

a) b)

Figure 19: a) The contour map of a rotated torus illuminated by a uniform hemispherical sky, b) the isophote map.

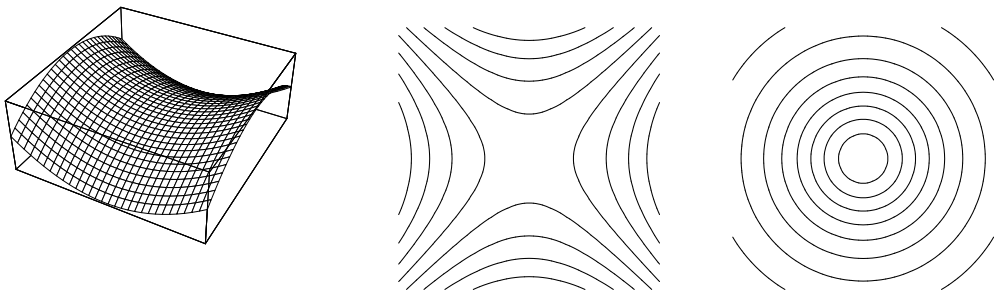

a) b) c)

Figure 20: a) An oblique view of a hyperbola. b) The contour map of the hyperbola illuminated by a uniform hemispherical sky, c) the isophote map.

With this selection, the depth function z and the first two derivatives z_x and z_y at the origin P are 0. z_{xx} and z_{yy} are the two principal curvatures λ_1 and λ_2, and $z_{xy} = 0$.

For $\Xi = \frac{1}{\sqrt{1+z_x^2+z_y^2}}$, an isophote curve on the surface is defined by

$$(-z_x\Xi, -z_y\Xi, \Xi) \cdot (s_1, s_2, s_3) = const$$

for a distant light source at (s_1, s_2, s_3). The tangent to the isophote lies in the $X - Y$ plane by definition, and its direction is given by:

$$\begin{aligned}(\tilde{x}, \tilde{y}) &= (s_1(z_{xy}\Xi + z_x\tfrac{\partial\Xi}{\partial y}) + s_2(z_{yy}\Xi + z_y\tfrac{\partial\Xi}{\partial y}) - s_3\tfrac{\partial\Xi}{\partial y}, \\ &\quad -(s_1(z_{xx}\Xi + z_x\tfrac{\partial\Xi}{\partial x}) + s_2(z_{yx}\Xi + z_y\tfrac{\partial\Xi}{\partial x}) - s_3\tfrac{\partial\Xi}{\partial x})) \\ &= (s_2\lambda_2, -s_1\lambda_1)\end{aligned}$$

namely,

$$\begin{pmatrix} \tilde{x} \\ \tilde{y} \end{pmatrix} = \begin{pmatrix} \lambda_2 & 0 \\ 0 & \lambda_1 \end{pmatrix} \cdot \begin{pmatrix} s_2 \\ -s_1 \end{pmatrix} . \tag{9}$$

From (9) it follows that at a parabolic point, where $\lambda_2 = 0$, the isophote points along the direction Y. Let the parabolic line point along the X direction (for a surface of revolution it must point along either X or Y). It follows that in the neighborhood of P the Y component of the isophote changes sign when the principal curvature λ_1 changes its sign at P (i.e. when the parabolic point is a transition between elliptic and hyperbolic regions) and the isophote bends. If λ_1 does not change sign (in the transition between two elliptic regions), the isophote has only an inflection point. Note that when $\lambda_1 = \lambda_2$, the isophotes are perpendicular to the projection of the light source on the $X - Y$ plane (which is not the image plane). On a self-shadow edge $s_3 = 0$, in which case the light source direction itself is perpendicular to the edge.

ACKNOWLEDGMENTS: I'm very grateful to W. Richards, A. Zisserman, T. Poggio, S. Ullman, E. Hildreth and A. Shashua for many helpful suggestions regarding the manuscript. I also thank T. Breuel for helping me in using his image processing package.

References

[1] P. J. Besl and R. C. Jain. Invariant surface characteristics for 3D object recognition in range images. *Computer Vision, Graphics, and Image Processing*, 33:33–80, 1986.

[2] T. O. Binford. Generalized cylinders representation. In S. C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 321–323. Wiley, New York, 1987.

[3] A. Blake, A. Zisserman, and G. Knowles. Surface descriptions from stereo and shading. In B. K. P. Horn and M. J. Brooks, editors, *Shape from shading*, pages 29–52. The MIT press, Cambridge, MA, 1989.

[4] M. J. Brady. Criteria for representations of shape. In J. Beck, B. Hope, and A. Rosenfeld, editors, *Human and machine vision*, pages 39–84. Academic Press, New York, 1983.

[5] M. J. Brooks and B. K. P. Horn. Shape and source from shading. In B. K. P. Horn and M. J. Brooks, editors, *Shape from Shading*, pages 53–68. MIT Press, Cambridge, MA, 1989.

[6] H. H. Bülthoff and H. A. Mallot. Interaction of depth modules: stereo and shading. *Journal of the Optical Society of America*, 5:1749–1758, 1988.

[7] J. F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–698, 1986.

[8] M. P. DoCarmo. *Differential geometry of curves and surfaces.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1976.

[9] G. J. Ettinger. Large hierarchical object recognition using libraries of parametrized model sub-parts. In *Proceedings of the 2nd International Conference on Computer Vision*, pages 32–41, December 1988.

[10] F. P. Ferrie and M. D. Levine. Where and why local shading analysis works. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):198–206, 1989.

[11] D. Forsyth and A. Zisserman. Mutual illumination. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition*, pages 466–473, San-Diego, CA, 1989.

[12] B. K. P. Horn. *Robot vision.* MIT Press, Cambridge, Mass., 1986.

[13] B. K. P. Horn. Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75, 1990.

[14] K. Ikeuchi and B. K. P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 15:141–184, 1981.

[15] J. J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.

[16] J. J. Koenderink and A. J. van Doorn. The internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32:211–217, 1979.

[17] J. J. Koenderink and A. J. van Doorn. Photometric invariants related to solid shape. *Optica Acta*, 27(7):981–996, 1980.

[18] C-H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. In B. K. P. Horn and M. J. Brooks, editors, *Shape from Shading*, pages 323–348. MIT Press, Cambridge, MA, 1989.

[19] E. Mingolla and J. T. Todd. Perception of solid shape from shading. *Biological Cybernetics*, 53:137–151, 1986.

[20] V. S. Nalwa. Representing oriented peicewise c2 surfaces. In *Proceedings of the 1st International Conference on Computer Vision*, pages 40–51, December 1987.

[21] S. K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *International Journal of Computer Vision*, 6(3):173–195, 1991.

[22] J. Oliensis. Uniqueness in shape from shading. *International Journal of Computer Vision*, 1990. in press.

[23] A. Pentland. Shape information from shading: a theory about human perception. In *Proceedings of the 2nd International Conference on Computer Vision*, pages 404–413, Tarpon Springs, FL, 1988. IEEE, Washington, DC.

[24] A. P. Pentland. Local shading analysis. In A. P. Pentland, editor, *From pixels to predicates*, pages 40–77. ablex, New Jersey, 1986.

[25] S. A. Shafer and T. Kanade. Using shadows in finding surface orientation. *Computer Vision, Graphics, and Image Processing*, 22:145–176, 1983.

[26] R. Szeliski. Fast shape from shading. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53:129–153, 1991.

[27] J. T. Todd and F. D. Reichel. Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, 96(4):643–657, 1989.

[28] B. C. Vemuri, A. Mitiche, and J. K. Aggrawal. Curvature-based representation of objects from range data. *image and vision computing*, 4:107–114, 1986.

[29] D. Weinshall. Direct computation of qualitative 3D shape and motion invariants. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(12):1236–1240, 1991.

[30] I. Weiss. Projective invariants of shapes. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition*, pages 291–297, June 1988.

[31] L. B. Wolff. A photometric invariant and shape constraints at parabolic points. In *Proceedings of the 3rd International Conference on Computer Vision*, pages 344–349, Osaka, Japan, 1990. IEEE, Washington, DC.

[32] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19:139–144, 1980.

[33] A. L. Yuille. Zero crossings on lines of curvature. *Computer Vision, Graphics, and Image Processing*, 45:68–87, 1989.

[34] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):680–702, 1991.