3-D SHAPE REPRESENTATION BY CONTOURS

ISAAC; WEISS*

Massachusetts Institute of Technology-'

ABSTRACT

The question of ,3-D shape representation is studied on the fundamental and general level. The two aspects of the problem, (i) the reconstuction of a 3-D shape from a given set of contours, and (ii) finding "natural" coordinates on a given surface, are treated by the same theory. We first set a few basic principles that should guide any shape reconstruction mechanism, regardless of its physical implementation. Second, we propose a new mathematical procedure that complies with these principles and offers several advantages over the existing ad hoc treatments. Some general results are derived from this procedure, which conform very well with human visual perception.

I. Introduction

A major component of Image Understanding is associating 3-D shapes with contours First, given a set of contours, such as may be provided by some image processing device, one wants to infer the shape of the surface that they most likely describe. In the complementery problem, given the surface, one seeks its "natural parametrization . namely contours that will convey its essential characteristics in the most economical, yet reliable way One known mechanism that performs these tasks well is the human visual system A few drawn lines can create a surprisingly vivid and convincing impression of a 3-D shape [Barrow & Tannenbaum. 1981

In this report, we address this problem on the fundamental and general level of first, finding the principles that must guide any process of representing a surface by contours, regardless of its physical implementator!, and second, we propose a mathematical mechanism that can perform these tasks, conforming with our principles Nevertheless, as the eye is the most successful image processing system we have, we shall test the performance of our abstract procedure against its In the following we shall summarize the requirements that we impose on a surface reconstruction mechanism, and their relation to previous work

1) A mechanism should build a surface in accordance with the information it has available about it, like boundary or other contours on it. but it should not add extraneous information of its own It is qualitatively clear that information is closely related to the smoothness of the curve or the surface A straight line can be described by the coordinates of the end points onl\. while a more complicated shape will require more informaton Thus it is reasonable, and mathematically convenient, to associate minimal information with minimal curvature, and to assume that a surface reconstruction mechanism will look for a surface that will have a minimum overall curvature while fitting the given contours.

2) The reconstruction mechanism should be invariant to rotations and translation, i.e. if the input image is rotated or translated. the output should move by the same amount but not change its shape The principle of minimum "energy" assuming that the overall scalar curvature $\int k^2$ is minimal satisfies requirement (1) and (2). but not the following demands, which our proposed measure does

3) Dimensionlessness or scale invariance When the distance between an object and a viewer changes, the object's apparent size changes but not its shape Thus a reconstruction mechanism should yield the same output shape from an input image regardless of its size Thus the mathematical method of reconstruction must be dimensionless, or invariant to scale The minimum energy principle does *not* satisfy this important demand A simple remedy examined later is to multiply the energy by the total length of the contour, but this will not satisfy the rest of our requirements.The principle of maximum compactness, which seeks to maximize the ratio of the area to the square of the perimeter, is dimensonless, but it has only been successfully applied to closed planar curves

4) Handling of different scales of variations A variation in the tangent occuring over a small length results in a large curvature. A small bump, or "noise", will have considerable effect on the integral. A sharp corner will totally dominate it, with the integrals value being determined by the exact way the corner is formed, which should be immaterial for determining the shape. Thus, the energy principle is only applicable for very smooth curves while the compactness criterion is inherently quite insensitive to noise. The energy principle also tends to be insensitive to slowly changing, large scale features, and it will completely ignore straight sections of the curve Our procedure deals with both small- and large-scale variations quite successfully.

5) We would prefer the same mechanism to handle both aspects of the problem mentioned above, namely finding both the surface and a suitable set of coordinates on it While previously suggested mechanisms only attempt to find the surface, ours also lead to a "natural" parametrization of it.

This report is an abbreviation of a paper (to appear) which contains the mathematical details.

II. Smooth Contours

As we want to represent the amount of information contained in a curve by its curvature, the natural mathematical quantities to deal with are the derivatives of the tangent vector such as the curvature vector $\vec{k} = d\vec{t}/ds$ As we demand rotational invariance, we have to use a scalar product of these vectors This has led to the suggestion of the principle of minimum energy, by which one wants to minimize the integral $\int \vec{k}^2 ds$, with the integral taken over the whole curve, and to its 3-D generalizations. This may be implemented in several ways first, given the coordinates and tangents of two points on the curve, one can find the curve that will pass through them and will minimize this integral Second, given a closed boundary such as an ellipse, one can try to interpret it as a projected image of a surface in 3-D. The human eye will usually regard an ellipse as a slanted circle, and we would like an extremum principle to yield

"Now at the Center for Automation Research, University of Maryland College Park, MD 20742 the circle as an extremum over the set of all curves compatible with the projected ellipse. Limiting ourselves to planar curves (zero torsion), this means the set of ellipses (produced by different slant angles) with the same major axis as the apparent one.

The energy principle fails in this task as it does not extremize the circle over the set of ellipses. Even worse, as it is not dimensionless, the extremum will depend on the apparent size of the image, contrary to to our demands. We now propose a modification of the energy principle, which will work for smooth curves.

Letting Δs be the total length of the curve, we can define the dimensionless variable

$$l = \frac{s}{\Delta s}$$

where s is the length of the curve lying between a point on the curve and one of its ends. l is a measure of the relative position of a point on the curve, regardless of the curve's length. We define the "action", in analogy to the physical quantity, as the integral

$$A = \int_0^1 \frac{d\bar{t}}{dl} \cdot \frac{d\bar{t}}{dl} \, dl$$

This is equivalent to mutiplying the "energy" by Δs. As both l and $\bar{t}$ are dimensionless, so is A. (Unlike the physical action). Another advantage of this normalization is related to the way the extremum is found. In general, one can use the Euler-Lagrange equations to obtain a differential equation for $\bar{t}(l)$. But this requires that both the dependent and the independent variables will have fixed values at the end points. Our normalization provides that, as l always runs between 0 and 1, unlike s.

The extremum problem is easily solved in the simple case of a curve with two boundary points. It is convenient to use a polar coordinate system, with the angle ϕ defined as the angle between the tangent $\bar{t}$ and (say) the x axis. We thus have:

$$t_x = \cos\phi \qquad t_y = \sin\phi$$

The action now reads

$$A = \int_0^1 \left(\frac{d\phi}{dl} \right)^2 dl$$

We shall now show that given two end points with the tangent there, the curve that passes through them is a circular arc. As a special case, extremizing over the set of ellipses with their extreme points fixed, will yield a circle. Our variational variable is now $\phi(l)$, with the boundary conditions $\phi(0) = \phi_1$ and $\phi(1) = \phi_2$, where ϕ_1, ϕ_2 are the inclinations of the curve at the end points. The EL equations yield

$$\frac{d^2\phi}{dl^2} = 0$$

with the unique solution

$$\phi = (\phi_2 - \phi_1)l + \phi_1$$

which is a circular arc.

III. Large- and Small-scale Variations

The simple minimum energy principle, our modification notwithstanding, suffers from a severe shortcoming. It cannot handle sharp turns in the curve. A sharp corner will completely dominate the curve, and the value of the integral will depend on the exact shape of the corner, with point-like edges leading to infinities. Even a small bump will have a considerable influence, and in fact, the smaller the bump, the greater its effect on the integral will be (keeping its shape similar). This makes the energy principle inapplicable for curves with edges, or noise, without elaborate filtering techniques. On the other hand, straight lines are ignored by this principle, regardless of how long they may be.

We propose a new kind of an extremum principle, one of whose advantages is essentially solving this small (and large) scale variation problem. (Another advantage will become apparent when we go over to 3-D). First, we parametrize the curve by a new coordinate α, running along the curve with its limits being 0 and 1. Then we define a new "action" A

$$A = \int_0^1 \left[\left(\frac{dt}{d\alpha} \right)^2 + \left(\frac{dl}{d\alpha} \right)^2 \right] d\alpha$$

Intuitively, one can regard α as a variable which is affected by two forces, generated by the fact that the extremizing drives α to follow l and $\bar{t}$, in analogy to an inertia effect. The second term in A pushes α to be as close as possible to l. In the absence of the first term (i.e. a straight line) α will coincide with l. This makes the first term as close as possible to the notion of curvature. On the other hand, the first term drives α to follow $\bar{t}$, (or ϕ) to some extent and concentrate in regions where $\Delta\bar{t}$ is large (such as a corner), so that $\Delta\bar{t}/\Delta\alpha$ does not get out of control.

As a "physical" analogy, one may think of a spring which can be stretched and bent unevenly along its length. The variable "α" will represent the amount of mass from a point on the spring to one end (as opposed to the spatial length l). The analogy is not perfect, though, as an actual spring usually does not have a dimensionless A. The spring is "ideal" in the sense that it, or part of it, can shrink to an infinitesimal length (like an ideal gas). This is what will happen in a corner. A *finite* "mass" $\Delta\alpha$ will concentrate in an infinitesimal corner, allowing the finite bending $\Delta\phi$ to spread over it.

In extremizing A, the unknowns are the functions $l(\alpha)$ and $\bar{t}(\alpha)$ (or equivalently $\phi(\alpha)$). Solving for them will give us both $\alpha(l)$ (the distribution of the "mass" along the mass) and $\bar{t}(l)$, which defines the curve. These unknown functions are not always completely independent, which may reduce the number of unknowns, or they may depend on a parameter such as the slant, that will become an unknown to be found by the extremization process. As we shall see, the first term will be rather dominant in sections of the curves having rapid curvature changes over short length, such as bumps, corners, or saw-teeth. The second term will be the major contributer in the large-scale features, having slow variations in curvature, such as the other boundaries of the saw. Thus we have obtained a "natural" way of treating curves with two very different scales of curvature change, without having to use arbitrarily preset filtering widths, as is commonly done.

A simple demonstration of the workings of this procedure can be made for a simple corner. We shall now examine the influence of a corner on our new A, as compared to its contribution to the energy.

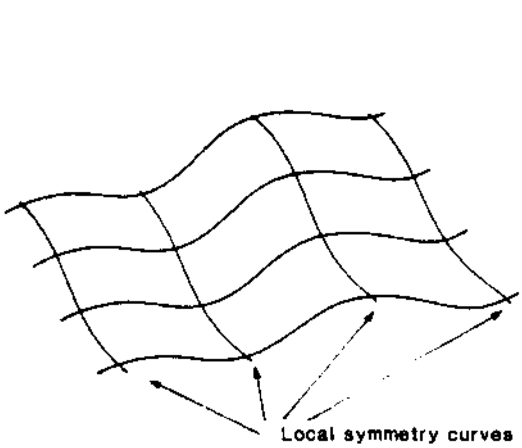

Figure 1. Local symmetry curves

We can build a corner from two "sticks" of length l, making an angle $\Delta\phi$ between themselves. We keep $\Delta\phi$ fixed, and we are only interested now in finding the distribution of the σ along this corner. This turnes out to be a quite simple extremum problem. The increments of σ along the sticks and on their joint is, respectively

$$\Delta\sigma_{stick} = \frac{l}{2l + \Delta\phi}, \qquad \Delta\sigma_{joint} = \frac{\Delta\phi}{2l + \Delta\phi}$$

As we noted before, the second quantity is indeed finite. (In fact, we derived this result by treating the joint as a small circular arc, with length l_{joint}, which tends to a point.) Substituting these quantities in A we obtain in the extremum.

$$A = (1 + \Delta\phi)^2$$

which is independent of the exact shape and size of the corner, only on the total turn $\Delta\phi$.

In comparison, the (modified) energy principle will give for this corner.

$$A_{energy} = \frac{\Delta\phi^2}{\Delta l_{joint}}$$

which tend to infinity as the length of the corner goes to zero.

A bump can be regarded as a series of relatively sharp turns. Thus, given our expression for A at a corner, it is clear that its contribution is not dominating. With filtering, that will reduce the turns $\Delta\phi$ along the bump to small values, the influence of the bump will be quite negligible.

IV. Skew-symmetry

We now collect 4 corners to form a parallelogram, which is a skew-symmetric shape. We shall allow the $\Delta\phi$-s of the corners to vary, in accordance with with interpreting the paralleigram as a shape with a slant in 3-D. We want to extremize A in respect to these $\Delta\phi$-s, to find the slant. The problem is complicated slightly by the fact that there is a constraint, namely that the sum of the angle of the 4 corners is equal to 2π. this can be handles by the method of Lagrange multipliers, and the result is that all the corner angles are equal, namely equal to $\pi/2$. Hence, the parallelogram is interpreted by our extremum principle as a slanted rectangle, which is consistent with human perception. This result can be extended to general skew-symmetric shapes.

V. General Surfaces in 3-D

So far we have dealt with planar surfaces that can be described by a one-parameter boundary curve. We now turn to the general case of a (reasonably) arbitrary surface in 3-D. such a surface can be parametrized by two coordinates σ_1, σ_2. Our action A will now be a double integral

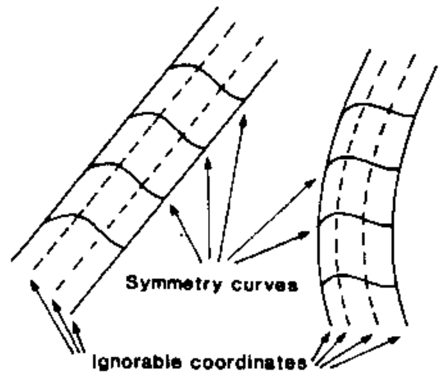

Figure 2. Ribbons with symmetry curves

$$A = \int_0^1 \int_0^1 \sum_i \left[\left(\frac{dt}{d\sigma_i} \right)^2 + \left(\frac{dl}{d\sigma_i} \right)^2 \right] d\sigma_1 d\sigma_2$$

where the summation is on $i = 1, 2$

Given a set of known contours, either on the boundary or otherwise. we construct the surface that best fits these contours as the one that extremzes the above integral. Moreover, the a, will now have a more tangible meaning then in the one-parameter case: they will be the natural coordinates parametrizing the surface. This is another advantage of the new procedure over the energy principle. It is interesting that the variables a, can serve the dual purpose of both handling bumps (and straight lines) and provide a set of 3-D natural coordinates.

it is easy to prove, for example, that a circular boundary will give rise to a sphere, complete with its longitudinal and latitudinal coordinates. Rather then do that. we shall state a general theorem about curves on surfaces, which will be very useful in finding natural coordinates as well as the surfaces themselves (including the sphere)

We first define a curve of local symmetry, T as one which divides the surface in two parts that are symmetric in the vicinity of the curve. Put another way, a reflection, say of the surface lying to right side of T will match the left side, near T. An example is the center line of any fold, ridge, valley or corrugation on a surface (fig 1), if this line is planar. Another example is the three symmetry lines on an ellipsoid (this is a global symmetry).

Theorem I A curve of local symmetry is planar

proof: Trivial. A reflection of the right side of T will not match the left side if T is not planar, (fig 1)

Theorem 2. A local network of coordinates, consisting of a local symmetry curve T and curves that are orthogonal to it, locally extremzes the action A

By "locally extremize" we mean that an infinitesimal variation of these curves, in the vicinity of P, will leave A unchanged

We shall discuss consequences of the last theorem. (The proof is ommited.) A very interesting class of surfaces to which our theorem is easily applicable is the one having an "ignorable coordinate", namely, its curvatures will not change as we move along this coordinate. As sub-classes one can mention (a) strips, or pipes, of arbitrary cross section, whether straight, circularly or helically bent, (fig 2), in which the ignorable coordinate is the length of the pipe, and (b) surfaces of revolution, in which the

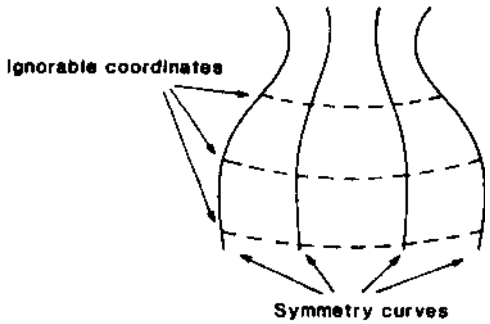

Figure 3. A surface of revolution

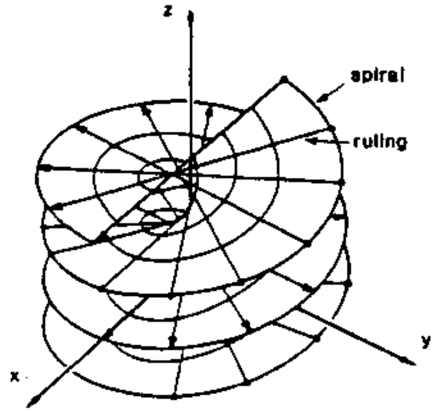

Figure 4. A helicoid

azimuth Φ is "ignorable" (fig 3) The coordinate *orthogonal* to the the ignorable one can be regarded as a local symmotry curve. as the surface stays the same on its sides Thus, this curve and the curves in the ignorable direction that intersect it. form a local network of curves that satisfy the condition- of theorem J . and thus it extremzes A

As we can see from the figures these curves are exactly the-ones that our visual intuition would expect as natural so that we are justified in calling the curves extremizing 4 the 'natural coordinates'

A further consequence of theorem 2 is the ability to make predictions about surfaces when they are not known in advance, without having to actually extremize A using the EL equations It is clear that generally, the greater the number of extremal local sections a surface has, the more "extremal" the global 2-dimensional integral A will be (This would be clearer if this "externum" were a "minimum", which it usually is, but as we have not examined here the second derivatives, we shall stick to the term "extremunY'.) Thus, a surface that extremizes A will contain as many symmetry curves as are compatible with the initial data As a consequnce, we will tend to obtain surfaces containing planar, spherical or cylindrical parts, exact or approximate, rather than bumpy ones As a particular example, we can use this general result to conclude, without having to solve the EL equations, that a circular boundary will give rise to a sphere. This is because a sphere is the most symmetric surface, thus having the greatest amount of symmetry curves More generally, boundary contours such as of fig. 3 will give rise to a surface of revolution, as this surface consists of a collection of local networks (the meridians and parallels) which fits those contours, and similarly for other surfaces

It is reasonable to assume, that when the conditions of the theorem are satisfied only approximately, e.g. when the curve is only approximately symmetric, a similar parametrization will still take place, with the natural coordinates approximately following the quasi-symmetry curves Thus our theory is applicable to shapes like generalized cones as long as the flutings vary slowly on the lenth-scale of the cone's radius This also is consistent with human perception.

VI. Relation to Other Work

We have already noted the intrinsic flaws of the energy principle, and its derivatives (such as Barrow and Tenenbaum's), as surface reconstruction mechanisms The compactness measure of Brady and Yuille [1984] does not suffer from these deficiencies It has been applied successfully to closed planar curves, but it is hard to see how it can be made a general theory It should be best viewed, perhaps, as a good measure of global symmetry.

Brady & Asada 1984 have studied local symmetries in planar shapes Strictly speaking, our defintion of "local" applies to an infinitesimal vicinity around a curve In this sense, every straight line on a plane is a local symmetry curve, except at it edges Brady examined larger vicinities that extend to the nearby boundaries of the curve so we shall elevate the type of symmetry he treated to a "regional" symmetry The fact that Brady's curves are also natural coordinates indicate that the notion of symmetry is of fundamental importance at all spatial scales (global, regional, and local) of shape representation

In differential geometry terms our local symmetry curves are planar geodesic ones. Brady *et al* 1984 have shown that planar geodesics are usually "natural" coordinates However, not all natural coordinates are planar geodesics For example, the parallels on a surface of revolution, and the spirals on a helicoid (fig 4), are not planar geodesics The question why these look natural was left open, as they did not seem to fit any consistent rule In our theory, these curves are natural coordinates by virtue of their being ***orthogonal at each point of their length to planar geodesics***, namely the meridians and the rulings, respectively Moreover, unlike the previous works, our results come out as a part of a general theory of shape representaton derived from sound first principles

Acknowledgements

The author thanks Michael Brady of MIT's AI Lab for carefully reading earlier drafts of this paper and making very valuable comments and suggestions.

References

Barrow, H. G. and J. M. Tenenbaum, [1981], "Interpreting line drawings as three-dimensional surfaces' , *Artif Jntrli,* 17. 75-117

Brady. Michael, J. Ponce, A. Yuille. H. Asada, [1984]. "Describing surfaces' To be published,

Brady. Michael, and H. Asada, (1984], "Smooth local symmetries and their implementation", A. I. Memo No. 757.

Brady. Michael, and A. Yuille, [1984], "An extremum principle for shape from contour", *IEEE Patt Anal & Mach. Int.,* PAMI-6

SPECULAR STEREO

A Blake

Computer Science Department,
Edinburgh University,
Scotland.

ABSTRACT

A glossy highlight, viewed stereoscopically, can provide information about surface shape. For example, highlights appear to lie behind convex surfaces but in front of concave ones.

A highlight is a distorted, reflected image of a light source. A ray equation is developed to predict the stereo disparities generated when a point source of light is reflected in a smooth, curved surface. This equation can be inverted to infer surface curvature properties from observed stereo disparities of the highlight. To obtain full information about surface curvature in the neighbourhood of the highlight, stereo with two different baselines or stereo with motion parallax is required

The same ray equation can also be used to predict the monocular appearance of a distributed source. A circular source, for instance, may produce an elliptical specular patch in an image, and the dimensions of the ellipse help to determine surface shape.

1 INTRODUCTION

When the reflectance of a surface has a specular as well as a diffuse component, the viewer may see highlights. Highlights can give extra information about surface shape Ikeuchi [8] uses photometric stereo with specular surfaces to determine surface orientation Beck [I] notes that stereo vision might be able to perceive highlights on a convex surface as lying beneath the surface Grirnson [7] incorporated lambertian and specular components of reflectance into stereo But he found that the computation of surface orientation could be numerically unstable.

Here a computation is proposed that is less ambitious than Crimson's, in that it. attempts to determine only local surface geometry, at specular points But it avoids relying on precise assumptions a b o u t surface reflectance

Instead, the only assumption is that a specularity can be detected in an image, and its position measured For instance a method like that of Ulltnan [15] could be used Thereafter specularities are matched in the same way as features in conventional stereo [6,9.10]. The disparity of a stereo-matched specular point is then compared with the disparity of any nearby surface features

The basic principle of the surface shape estimation relies on the properties of curved mirrors (fig 1).

To interpret, specular stereo, both horizontal and vertical disparities are used. Ideally, three non-collinear eyes are needed to obtain full information about local curvature. Alternatively, parallax from a known vertical motion of a viewer, combined with conventional stereo geometry, is just as good If only a static, stereo view is available then this still yields partial information This could be combined

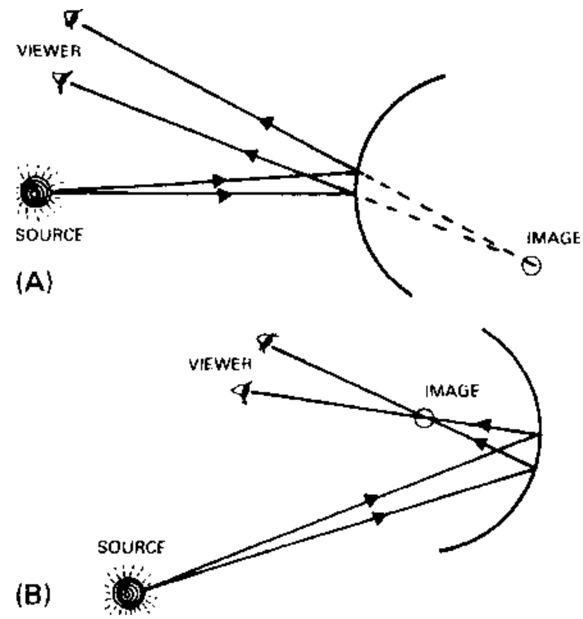

Figure 1: *Viewing geometry. In a convex mirror* (a) *the image of a distant point source appears behind the mirror surface In a concave one* (6) *the image may appear in front Study of the ordinary domestic soup spoon should confirm this*

with a *priori* knowledge or measurements from other sources (stereo, shape-from shading or specular reflection of a distributed source) to fully determine local surface shape

Finally, observe that the path of a light ray from source to viewer can be reversed Analysis developed to show the effect of moving the viewer also serves for movement, of the source The resulting equation is used to predict the appearance in an image of a distributed source under specular reflection in a curved surface. For instance a circular source generally produces an elliptical specularity in the image. The orientation and length of its major and minor axes, in principle, determine local surface shape.

2 IMAGING EQUATIONS

Equations are given to describe the process of formation of images of specular reflections Details of derivations are given in [3] These predict the dependence of observed stereo disparities on surface and viewing geometry Certain assumptions about the geometries are made, for the sake of mathematical simplicity Then the equations are inverted so that, given viewing geometry and disparities, local surface geometry can be inferred

2.1 Viewing, surface and reflection geometry

The stereo viewing geometry is shown in fig 2.

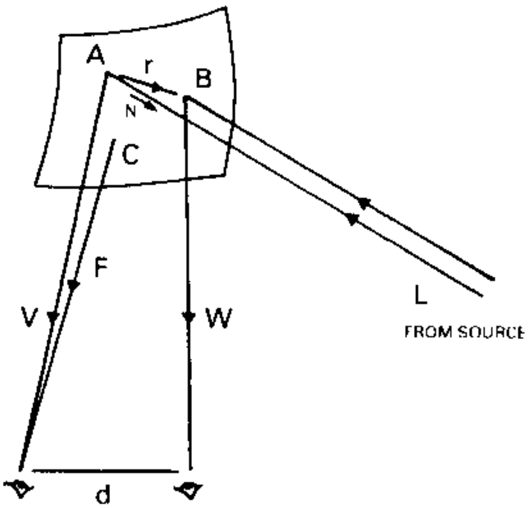

Figure 2: ***Stereo viewing geometry.*** *Illumination comes from a distant point source, in direction* $\mathbf{L}$. *Rays to left and right eyes lie along vectors* $\mathbf{V},\mathbf{W}$, *and strike the surface at points* A,B *respectively. Surface normals at* A,B *are* $\mathbf{N},\mathbf{N}'$ *respectively. The vector from* A *to* B *is* $\mathbf{r}$, *and the baseline lies along vector* $\mathbf{d}$. *A surface feature is assumed to be present nearby, at* C, *with position vector* $\mathbf{F}$ *relative to the left eye.*

It is assumed that the curved surface is locally well approximated by terms up to 2nd order in a Taylor series (see eq (4)). The vectors $\mathbf{V},\mathbf{d},-\mathbf{W},-\mathbf{r}$ form a closed loop, so that

$$\mathbf{V}+\mathbf{d}-\mathbf{W}-\mathbf{r}=0. \qquad (1)$$

A coordinate frame is chosen with origin at A, with $\mathbf{N}=(0,0,1)$, and with $\mathbf{L},\mathbf{V}$ lying in the x-z plane, so that

$$\mathbf{V}=(V\sin\sigma,0,V\cos\sigma),\ \mathbf{L}=(-\sin\sigma,0,\cos\sigma) \qquad (2)$$

where σ is the slant of the tangent plane at A. Its tilt direction lies in the x-z plane. Note that if viewing geometry and light source direction $\mathbf{L}$ are known (and the latter could be obtained as in [11]) then surface slant and tilt are known: the surface normal lies in the plane of $\mathbf{V},\mathbf{L}$ and bisects them.

It is assumed that some feature at point C on the surface, is available near to the specular points A,B (fig 2) and that stereo is able to establish the position of C. Its position vector $\mathbf{F}$ is used to estimate V, the length of the vector $\mathbf{V}$. Assuming that C is not too far away from A, so that C lies, approximately, in the tangent plane at A,

$(\mathbf{V}-\mathbf{F}).\mathbf{N}=0$ so that

$$V\cos\sigma=\mathbf{F}.\mathbf{N} \qquad (3)$$

Since the choice of coordinate frame ensures that gradients vanish $(\partial z/\partial x=\partial z/\partial y=0)$, the surface, in the neighbourhood of A, is described by

$$z(x,y)=(1/2)\mathbf{x}.(H\mathbf{x})+O(|\mathbf{x}|^3) \qquad (4)$$

where $\mathbf{r}=(x,y,z)$, $\mathbf{x}=(x,y)$ and H is the (symmetric) hessian matrix [4] of the surface. Note that $\mathbf{r},\mathbf{d}$ etc are 3-dimensional vectors but $\mathbf{x}$ is a 2-dimensional vector, in the xy-plane. Similarly H is a 2×2 matrix operating on $\mathbf{x}$.

The law of reflection at A is that

$$2(\mathbf{V}.\mathbf{N})\mathbf{N}-\mathbf{V}\ \|\ \mathbf{L}. \qquad (5)$$

where $\|$ denotes "is parallel to". Similarly for the other eye, at B,

$$2(\mathbf{W}.\mathbf{N}')\mathbf{N}'-\mathbf{W}\ \|\ \mathbf{L}. \qquad (6)$$

Combining (1) (5) (6) gives (see [3] for details):

$$M H\mathbf{x}=\mathbf{x}+\mathbf{w} \qquad (7)$$

where $w_x=-d_x+d_z\tan\sigma$, $w_y=-d_y$ and

$$M=\begin{pmatrix}2(V\sec\sigma+d_z+d_x\tan\sigma) & 2d_y\tan\sigma\\ 0 & 2(V\cos\sigma+d_z)\end{pmatrix} \qquad (8)$$

The approximations used above hold good provided $|\delta\mathbf{N}|\ll\cos\sigma$ and $|\mathbf{x}|\ll V\cos\sigma$. This means that surface slant σ must not be close to 90°, and that both vergence angle and (angular) disparity should be small. It can be shown that these conditions will usually be satisfied when the stereo baseline is short, so that $|\mathbf{d}|\ll V\cos\sigma$.

To solve equation (7), we note also that

$$\det(M)=4(V\sec\sigma+d_z+d_x\tan\sigma)(V\cos\sigma+d_z)$$

so that, provided surface slant σ is not near 90° as above, and provided $|\mathbf{d}|<(1/2)V$ (baseline length less than half viewing distance), then $\det(M)\neq 0$. In that case, equation (7) can be inverted to give

$$H\mathbf{x}=\mathbf{v} \text{ where } \mathbf{v}=M^{-1}(\mathbf{w}+\mathbf{x}) \qquad (9)$$

which, in general, can be expected to impose 2 constraints on the 3 variables of H. If something is known already about H - say, that the surface is locally cylindrical at A - then it might be possible to determine H completely.

2.2 Disparity measurement

The equations just derived require the vector $\mathbf{x}$ to be determined from disparity measurements, as shown in fig 3.

Having obtained from the stereo images the angular disparity $\boldsymbol{\delta}$ of the specularity as in fig 3, it can be "back projected" onto the surface to obtain the length $\mathbf{x}$. The assumption that $|\delta\mathbf{N}|$ is small is used again to obtain

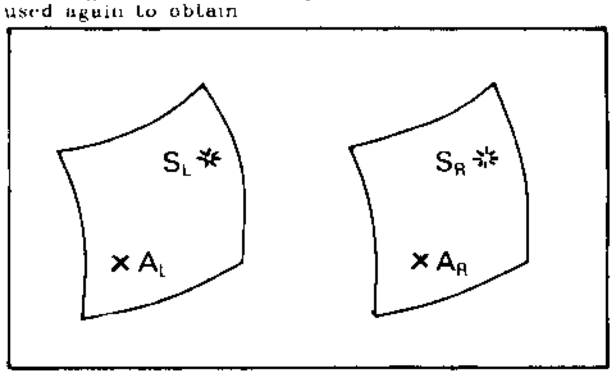

Figure 3: ***Disparity measurements.*** *The specular point is imaged at angular positions* S_L,S_R *in the left and right images respectively. A nearby surface feature is imaged at* A_L,A_R *and provides a disparity reference point on the surface. From these measured positions in the image,* $\boldsymbol{\delta}=(S_R A_R)-(S_L A_R)$ *is computed - the difference between the angular disparities of the 2 points. Then* $\mathbf{x}=(x,y)\approx V(\delta_x\sec\sigma,\delta_y)$

$$\mathbf{x}=(x,y)=VP\boldsymbol{\delta}, \text{ where} \qquad (10)$$

$$P=\begin{pmatrix}\sec\sigma & 0\\ 0 & 1\end{pmatrix}$$

2.3 Focusing effects

Equation (7) predicts that imaging of a specularity can become degenerate (fig 4). The equation can be rewritten as

(MH-l)x - *w* *(11)*

and the condition for degeneracy is that det(MH-/)-0.

The focusing effect may produce either a line or a blob in the image:

1 If the rank of *(MH-1)* is 1 the specularity will appear in the image as a line. Or else, there may be no solution for *x* in (7), and nothing of the specularity will be visible. Stevens [13] observes that, with infinitely distant source and viewer, any line specularity must lie on a plane curve in the surface a special cose of the rank 1 focusing effect.

2 If the rank of *MH I* is zero, that is MH-/-0. then either the specular reflection is invisible as above, or it is focussed in 2 dimensions onto the imaging aperture, and appears in the image as a large bright blob

3 It can be shown that if the surface patch is convex, the effect cannot occur This corresponds to physical intuition Only concave mirrors focus distant light sources

2.4 A source at a finite distance

if the source is at a finite distance *L,* rather than infinite as assumed so far, the imaging equation (7) becomes

(MH (l+p)/)x = *w* (12)

The constant p = ($d_2$seeo + *V) /L* and clearly, as L->,

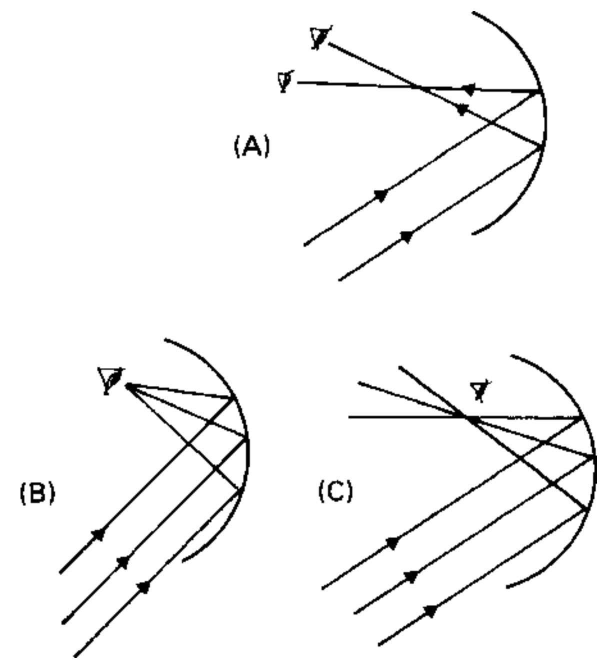

Figure 4: Degenerate imaging of a specularity. Normal imaging of a specularity (a) becomes *degenerate because of the focusing action of the curved surface. The surface may focus onto the viewing aperture (b) to produce a bright blob or line in the image, or away from the aperture (c) in which case nothing is visible in the image*

*p~*0* which gives the infinite source equation (7)

Suppose the infinite source computation of surface shape is performed, when in fact the source is at a finite distance *L,* how great is the resulting error[9] The answer is that the error in the curvature (along a given direction in the xy-plane) is of the order of ±√*VL.* The error is negligible if. assuming *a* not to be close to 90°, either

- the light source distance is large compared with the viewing distance: $L \gg V$, or
- the surface has *high curvature* for *both* principal curvatures, $\kappa_i \gg 1/L$, $i=1,2$

The first case is intuitively reasonable; if the light source is distant compared with the observer distance V, then equations for an infinite light source can be used with little error. What is perhaps less obvious is the second condition that, for highly curved objects, the source need *not* be further away than the observer.

2.5 Distributed sources

The mathematical model that has been used so far assumes a point source. In practice the source may be distributed, so that it subtends some non-zero solid angle, at the surface

Equation (12), for a source at a finite distance, is used but source and viewer positions are interchanged The light ray is reversed Vector $\mathbf{d}$ now represents the movement of *source* for a fixed *viewer* position. After some rearrangement, this yields a new equation, looking rather like (12) but with a factor V/L on the right hand side

$$(\mathbf{MH} - (1+p)\mathbf{I})\mathbf{x} = (V/L)\mathbf{w}. \qquad (13)$$

It would be most convenient to express the shape of the image specularity (using angular position in the image, $\boldsymbol{\delta}$) in terms of source distribution (using a new angular variable $\boldsymbol{\alpha}$). From (10) $\mathbf{x}=V\mathbf{P}\boldsymbol{\delta}$, and it is straightforward to show that $\mathbf{w}=L\mathbf{P}\boldsymbol{\alpha}$ So now

$$\mathbf{T}\boldsymbol{\delta} = \boldsymbol{\alpha} \text{ where } \mathbf{T}=\mathbf{P}^{-1}\mathbf{MHP} - (1+p)\mathbf{I}. \qquad (14)$$

What equation (14) says is that the viewer sees an image of the source that has undergone a linear transformation $\mathbf{T}^{-1}$. The effect of the transformation depends on surface shape. For a planar mirror for example, $\mathbf{H}=\mathbf{0}$ so that $\boldsymbol{\delta}=\boldsymbol{\alpha}/(1+p)$ an isotropic scaling that preserves the shape of the source Note that if the source is very distant, $p \approx 0$ and the scaling factor is unity.

If the angular dimensions of the source are known then, in principle, surface shape may be recovered completely by monocular observation. For a circular source with slant $\sigma=0$, the ellipse axes coincide with the principal curvature directions of the surface In general, when $\sigma \neq 0$, measuring the length and direction of ellipse axes enables $\mathbf{T}$ and hence $\mathbf{H}$ to be found from (14) Note that for a circular source, because of its symmetry, principal curvatures are determined only up to sign inversion (approximately)

3 INFERRING LOCAL SURFACE SHAPE

3.1 Locally cylindrical surface

On a surface that is *known* to be locally cylindrical, equation (7) is sufficient to recover both parameters of local surface shape. For instance, when the source is distributed, a strip shape image-specularity indicates that the surface may be cylindrical - or at least that one principal curvature may be much larger than the other (This can be deduced from (14))

The parameters to be determined are the direction of the cylinder axis *9* and the radius *R* Using (9):

$$\tan\theta = v_y/v_x$$

$$R = (x\cos^2\theta + y\sin\theta\cos\theta)/v_x$$

3.2 Spherical surface

The knowledge that the surface is locally spherical could be derived monocularly, from (14), as in the cylindrical case except that rather more must be known about the source for example, that it is circular

On a spherical surface there is only one parameter to specify - the radius of curvature *R,* and from (9).

$R = x/v_x = y/v_y$

· the second equality being available as a check for consistency of assumptions

3.3 Known orientation of principal axes

If the orientation of principal axes about the surface normal, is known then the complete local surface geometry can be obtained. Orientation could be derived monocularly (assuming source shape known) from (14).

Rotating coordinates about the z axis, a primed (') frame can be obtained in which H in (9) becomes diagonal:

$H'x' = v'$

Now, in general, the two diagonal components of H' can be obtained immediately. Experiments with computer generated images have obtained curvature to an accuracy of 10%.

3.4 General case

In the general case, the surface curvature at the specular point is described by 3 parameters, but the specular stereo measurements yields only 2 constraints. However two additional constraints - 4 in all - are available if a second baseline is used. The extra baseline could be derived either from a third sensor, suitably positioned, or from known motion of the viewer (parallax).

Suppose now that there are 2 baselines $d^{(i)}, i=1,2$ with corresponding $x^{(i)}, M^{(i)}, w^{(i)}, v^{(i)}$. Now equation (9), applied once for each baseline, gives

$$HX = V \quad \text{where} \tag{15}$$

$$X = \begin{pmatrix} x^{(1)} & x^{(2)} \\ y^{(1)} & y^{(2)} \end{pmatrix} \quad V = \begin{pmatrix} v_x^{(1)} & v_x^{(2)} \\ v_y^{(1)} & v_y^{(2)} \end{pmatrix}$$

and H can be recovered provided X is non-singular.

It appears to be impossible to suggest baselines that guarantee to generate a non-singular X, for *all* viewing geometries and surfaces. This is because $\det(X)$ depends on the surface and the viewing geometry, as well as on the baselines. This is probably best achieved (see [3]) by making the baselines $d^{(i)}$ fairly near orthogonal, and certainly nowhere near collinear.

The disparity measurements give 4 constraints. If H is the only unknown, it is now overdetermined. One could either

1. Test whether the H obtained from (15) is indeed symmetric as a check on validity of assumptions (for example, the validity of the local approximation of (4), over the range of movement of the specular point on the surface).
2. Use a least-squares error method to find the symmetric H that fits the data best. Then H is the solution of linear equations:

   $$HXX^T + XX^TH = VX^T + XV^T.$$

   The error measure $\|HX-V\|$, if it is too large, indicates that some assumptions were not valid.

4 CONCLUSION

Is specular stereo actually useful? We argue that it is. Of course the presence of specularities in the image cannot be guaranteed; specular stereo is not an autonomous process in the sense that conventional stereo is. Indeed specular stereo itself *relies* on conventional stereo to provide a disparity reference. In the case of a densely textured surface, conventional stereo with surface fitting [2,5,6,12,14) would be able to give an accurate estimate of surface shape. But for a smooth surface, stereo features may be relatively sparse, and fitting a surface to disparity measurements may be difficult and inaccurate. Then, provided at least one nearby surface feature is available as a disparity reference, specular stereo, together with monocular analysis of specularity, provides valuable surface shape information

Acknowledgement

This work was supported by SERC grant GR/D 1439.6 and by the University of Edinburgh. Thanks are due to G Brelstaff. A Zisserman and R Fisher for valuable discussion

REFERENCES

1 Beck, .1 (1972) *Surface color perception* Cornell University Press, Ithaca, U.S.
2 Blake, A. (1904) Reconstructing a visible surface *Proc AAAI conf* 1984, 23-26
3. Blake, A (1984) Inferring surface shape by specular stereo. Report CSR 179-84, Depi Computer Science, Edinburgh University
4. do Carmo, MP (1976) *Differential geometry of curves and surfaces.* Prentice Hall, Englewood cliffs, USA
5 Faugeras. 0 1) and Hebert, M (1983) A 3-D recognition and positioning algorithm using geometrical matching between primitive surfaces *JJCAI* 83. 996-1002
6 Grimson, W.E.I, (1982) *From images to surfaces* MIT Press, Cambridge, USA
7. Grimson, WEI.. (1982) Binocular shading and visual surface reconstruction *AI Lab Memo* 697, MIT. Cambridge, USA
8. Ikeuchi.K. (1981) Determining surface orientations of specular surfaces by using the photometric stereo method *IFFF trans PAMI.* 3, 6, 661 669
9. Marr, D and Poggio, T (1979) A computational theory of human stereo vision *Proc R. Soc Lond B.* 204, 301-328
10 Mayhew, JEW and Frisby. J.P. (1981). Towards a computational and psychophysical theory of stereopsis. *AI Journal,* 17, 349-385.
11 Pentland, A.P. (1984) Local shape analysis *IFFF trans PAMI,* March 1984. 170 187
12. Potmesil, M (1983). Generating models of solid objects by matching 3D surface segments *IJCAI* 83, 1089-1093
13. Stevens,K.A (1979) *Surface perception from local analysis of texture and contour* Ph.D. thesis, MIT, USA
14. Terzopoulos, D (1983) Multilevel computational processes for visual surface reconstruction, *Computer Vision Graphics and Imaye Processing,* 24, 52-96
15 Ullman. S (1976). On visual detection of light sources *Biol Cybernetics,* 21, 205 212