[18] Geometry from Specularities

Andrew Blake and Gavin Brelstaff

Department of Computer Science
University of Edinburgh, Edinburgh EH9 3JZ, Scotland, UK

Abstract

Relatively recently the problems for vision presented by specular reflections have begun to receive attention. In particular it has been noted that the monocular and stereoscopic appearance of specular reflections, and their dynamic behaviour, contain local shape information. Work reported here contributes to the understanding of specularity in several ways.

A new algorithm is described for accurate computation of horizontal and vertical stereoscopic disparities of specular points, relative to nearby surface points. Knowledge of such disparities is shown to restrict principal curvatures (with known light-source position) to lie on a hyperbolic constraint-curve. Monocular appearance of specularities is known also to constrain surface shape. We show that, at best, there remains a fourfold ambiguity of local surface curvature. In the case of a light source that is of unknown shape but known to be compact (in a precise sense), elongated specularities have geometrical significance. The "axis" of such a specularity, back-projected onto the surface tangent plane, approximates to a line of curvature. The approximation improves as the specularity becomes more elongated and the source more compact.

These ideas have been incorporated into an existing stereo vision system, and shown to work well with real and simulated images.

1 Introduction

Specular reflection represents both a problem and an opportunity in vision. It is a problem in that it disrupts processes such as edge-detection and stereoscopic matching, but an opportunity in that highlights or specularities are cues for surface geometry.

Clearly, to make any progress, it must be possible to detect specularities. Various processes have been proposed and tested. Some involve chromatic analysis [9,17] others achromatic analysis of intensity [25,8]. This paper addresses problems of geometric inference rather than of specularity detection but for the purposes of demonstrating a working system an achromatic specularity detector has been used. It builds on the ideas of Ullman's S-operator [25] and on the retinex process of Land and others [16,13,3], and will be the subject of a future paper.

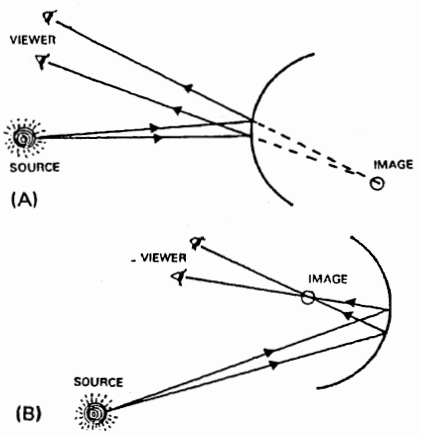

Figure 1: Specular stereo - the basic principle: specularities appear behind a convex mirror but in front of a concave one.

A discussion of the inference of shape from specularity is given by Koenderinck and van Doorn [15]. They elegantly expound the qualitative behaviour of specularities under viewer motion. Specularities travel freely in elliptic or hyperbolic regions, speeding up near parabolic lines, annihilating and being created, in pairs, on the parabolic lines. They travel most slowly in regions of high curvature and hence, for a given static viewer position, specularities are most likely to be found where curvature is high. More recently a number of quantitative analyses of specularity have emerged. Several involve active vision systems extending photometric stereo [26] to deal with and profit from specular reflection [7,14,20]. Other approaches are based on mathematical models of specular reflection, including specification of reflectance map, fixing various free parameters of the model by measure-

*Current addresses: A. Blake, Dept. Engineering Science, Parks Rd., Oxford. G. Brelstaff, IBM UKSC, Athelstan Ho., St Clement St., Winchester.

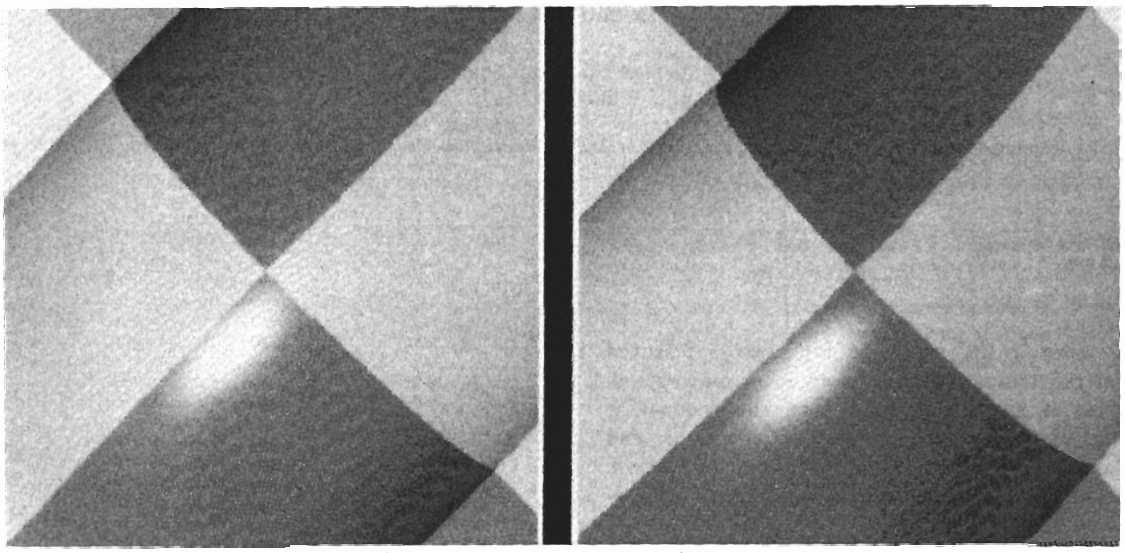

Figure 2: An example of a specularity whose "motion" relative to surface features is oblique - vertical relative disparity is not zero.

ments from individual specularities [1, 11,12]. In view of the difficulty of achieving adequate photometric models for specularity and of fixing their parameters from image data [11] it seems reasonable to try to restrict modelling to simple ray optics and the law of reflection. Koenderinck and van Doorn's work is in this vein, and other more quantitative models have been investigated [21,24,2]. In this paper we extend that theme and illustrate it by incorporating analysis of specularity in a stereoscopic vision system.

2 Stereoscopic analysis of specular reflection

The basic principle of "specular stereo" is illustrated in figure 1. According to the simple physics of curved mirrors, a specularity will appear behind a glossy, convex surface but (generally) in front of a concave one. Here this simple idea is expanded. For example, how does a specularity appear in a hyperbolic surface? Whether it appears behind or in front depends on the orientation of the surface. Even on elliptic, non-umbilic surfaces, astigmatic effects produce apparent depth variations as orientation is changed. In fact the notion of apparent depth is ill-defined here - what we actually observe are horizontal and vertical relative disparities (relative to disparities of surface features). Specularities, unlike physical surface features, need not satisfy the "epipolar" constraint [18] so vertical disparity may be non-zero. This is illustrated by the example of figure 2, in which relative displacement of the specularity in the right image (relative to the left) is oblique. Both horizontal and vertical disparities vary as the orientation of the stereo baseline changes relative to lines of curvature on the surface.

Analysis of the stereoscopic viewing geometry will establish the relationship between surface shape and measured disparities. It is helpful to consider two different kinds of analysis. The first is approximate, a linear system "driven" by the interocular separation, with disparities as its output. The characteristics of the linear system depend on surface geometry (curvature and orientation). This analysis appeared in earlier work [2] and is useful for characterising degeneracies - special alignments at which geometric inference will fail. An exact analysis is more convenient for computation as well as being more accurate and a method has been developed for accurate computation of relative disparities, together with error bounds.

2.1 Viewing geometry

The geometry for stereoscopic viewing of a specularity is shown in figure 3. Vectors $\mathbf{d}$, the stereo baseline, is assumed known, as is $\mathbf{S}$, the position of the light source [1]. The directions $\hat{\mathbf{V}}, \hat{\mathbf{W}}$ of vectors $\mathbf{V}, \mathbf{W}$ are given by the measured positions of the specularities in the left and right images. The projection of displacement vector $\mathbf{r}$ onto the left (say) image plane forms the observed relative disparity vector.

A few equations suffice to describe this geometrical arrangement. First, there are three cycles amongst the vectors:

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

[1] The assumption of known light source position is of course a strong one. However it can be computed from just one stereoscopic observation of a highlight on an object of known shape.

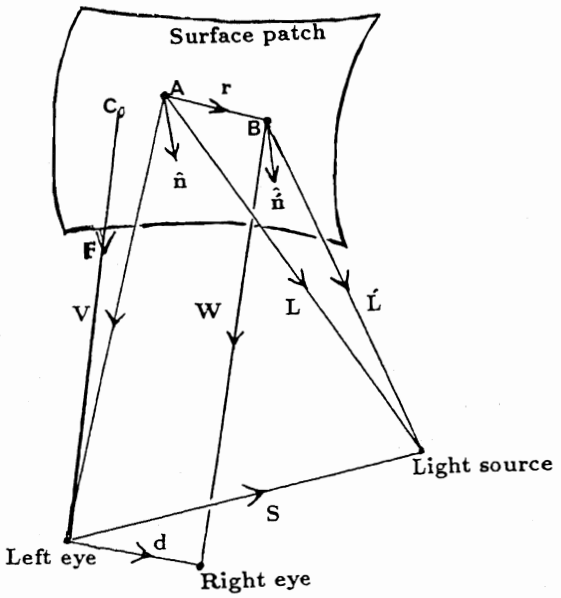

Figure 3: A smooth patch of surface is illuminated by a point source along vector **L**. Light striking point A is specularly reflected along vector **V** into the left eye. Similarly, light incident at B traverses **W** into the right eye. Surface normals at A and B are $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$ respectively. Vector **r** separates A and B. The stereo base-line lies along vector **d**. A surface marking lies nearby A at C.

$$\mathbf{V} - \mathbf{L} - \mathbf{S} = 0 \tag{2}$$

and

$$\mathbf{L} - \mathbf{L}' - \mathbf{r} = 0. \tag{3}$$

The physical law of reflection is expressed in the following equations:

$$\hat{\mathbf{n}} = \frac{\hat{\mathbf{V}} + \hat{\mathbf{L}}}{|\hat{\mathbf{V}} + \hat{\mathbf{L}}|}. \tag{4}$$

$$\hat{\mathbf{n}}' = \frac{\hat{\mathbf{W}} + \hat{\mathbf{L}}'}{|\hat{\mathbf{W}} + \hat{\mathbf{L}}'|}. \tag{5}$$

Finally, the vector **F** is computed from conventional stereoscopic viewing of the surface reference mark C, and it is assumed to lie in the tangent plane to the surface at A, so that

$$(\mathbf{F} - \mathbf{V}).\mathbf{n} = 0. \tag{6}$$

Of course C does not lie *exactly* in the tangent plane, but the error is small provided the surface reference C is not too far away from the ray intersection point A (figure 4).

2.2 Computation of surface depth

We need to know surface depth $V = |\mathbf{V}|$ but, so far, know only the depth of a nearby non-specular "reference point" on the surface. In fact V can be computed from equation

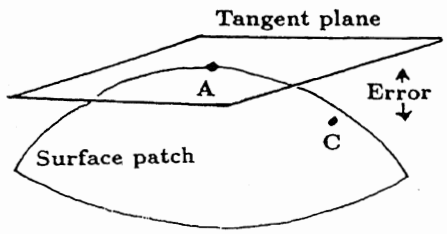

Figure 4: Point C lies on the surface at a little distance from A. Along this distance the surface curves gently out of the tangent plane. So the "tangent plane assumption" is only an approximation.

(6), given that $\hat{\mathbf{V}} = \mathbf{V}/V$ is known from the position of the specularity in the left image, by the iterative algorithm given in figure 5. This algorithm has not been

$i{=}0$, $|\mathbf{V}|^{(i)}{=}|\mathbf{F}|$.
repeat
$i = i+1$,
$\mathbf{V}^{(i)} = |\mathbf{V}|^{(i)}\hat{\mathbf{V}}$,
$\mathbf{L}^{(i)} = \mathbf{V}^{(i)} + \mathbf{S}$,
$\hat{\mathbf{n}}^{(i)} = \left(\hat{\mathbf{V}}+\hat{\mathbf{L}}^{(i)}\right) / |\hat{\mathbf{V}}+\hat{\mathbf{L}}^{(i)}|$,
$|\mathbf{V}|^{(i)} = \left(\hat{\mathbf{n}}^{(i)}.\mathbf{F}\right) / \left(\hat{\mathbf{n}}^{(i)}.\hat{\mathbf{V}}\right)$.
while $\left||\mathbf{V}|^{(i)}-|\mathbf{V}|^{(i-1)}\right| > 0.01|\mathbf{F}|$.
$\mathbf{V}{=}\mathbf{V}^{(i)}$, $\hat{\mathbf{n}}{=}\hat{\mathbf{n}}^{(i)}$, $\mathbf{L}{=}\mathbf{L}^{(i)}$.

Figure 5: Iterative algorithm for computing **V**, the position at which the reflected ray strikes the specular surface.

proved to converge, but seems well-behaved in practice.

2.3 Local surface shape

It is clear that, having computed **V** as above, **n** can be obtained from (4) and (2). In other words, knowing the light source position, and observing a nearby surface reference point C suffices to compute surface orientation at the specular point A.

So much for surface orientation - but what about local surface curvature? Curvature is expressed in terms of the Hessian matrix H. Choosing coordinates such that $\mathbf{x} = (x, y)$ is the restriction of $\mathbf{r} = (x, y, z)$ to the tangent plane at A:

$$z = \frac{1}{2}\mathbf{x}.(H\mathbf{x}) + O(|\mathbf{x}|^3). \tag{7}$$

(The choice of coordinates allows an arbitrary rotation in the tangent plane, which is conveniently fixed by choosing the (0,1,0) direction to be orthogonal to vector $\mathbf{V} - \mathbf{L}$.)

Differentiating (7), one can obtain

$$\delta\mathbf{n} = -H\mathbf{x} + O(|\mathbf{x}|^2). \tag{8}$$

where $\delta\mathbf{n}$ is the component of $\hat{\mathbf{n}}' - \hat{\mathbf{n}}$ lying in the tangent plane. Now $\hat{\mathbf{n}}$ is known already, so to compute $\delta\mathbf{n}$ we need $\hat{\mathbf{n}}'$ - which can be calculated from (5) if $\mathbf{W}$ is known. Observing that $\mathbf{r}.\mathbf{n} = 0$ and substituting this into (1) yields the following formula, in terms of measured quantities, for $|\mathbf{W}|$:

$$|\mathbf{W}| \approx \frac{(\mathbf{V}+\mathbf{d})\cdot\hat{\mathbf{n}}}{\widehat{\mathbf{W}}\cdot\hat{\mathbf{n}}}. \tag{9}$$

2.4 Graphical representation of geometric constraints

Measurement of $\delta\mathbf{n} = (\delta n_1, \delta n_2)$ imposes 2 independent constraints, via equation (8), on the components of the hessian H. But H is a symmetric matrix:

$$H = \begin{pmatrix} H_{xx} & H_{xy} \\ H_{xy} & H_{yy} \end{pmatrix}, \tag{10}$$

so it has 3 independent components. Clearly, further information is required to fix all 3. This can be obtained either by moving the stereo baseline or by monocular observation of the shape of the specularity. Both possibilities will be discussed in due course. In the meantime, it is natural to ask whether the 2 constraints already obtained represent *intrinsic* information about the surface - that is, do they constrain principal curvatures of the surface?

Brelstaff [4] has shown that there is indeed an intrinsic constraint. The principal curvatures κ_1, κ_2 are constrained to lie on a hyperbola. Equivalently, the corresponding principal radii of curvature r_1, r_2 lie on a hyperbola:

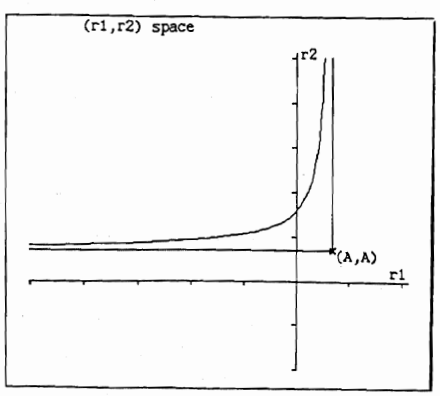

Figure 6: Stereo analysis constrains the principal radii of curvature $r_i = 1/\kappa_i$, $i = 1, 2$ to lie on the upper curve of a hyperbola.

$$-B^2 = (r_1 - A)(r_2 - A) \tag{11}$$

where

$$A = \frac{\delta\mathbf{n}.\mathbf{x}}{|\delta\mathbf{n}|^2} \text{ and } B = \sqrt{\frac{|\mathbf{x}|^2}{|\delta\mathbf{n}|^2} - A^2}, \tag{12}$$

as shown in figure 6. Without loss of generality, we can require $r_1 \leq r_2$, so that the constraint set includes only one curve of the hyperbola. For example, the family of surfaces allowed by the constraint in figure 6 is illustrated in figure 7.

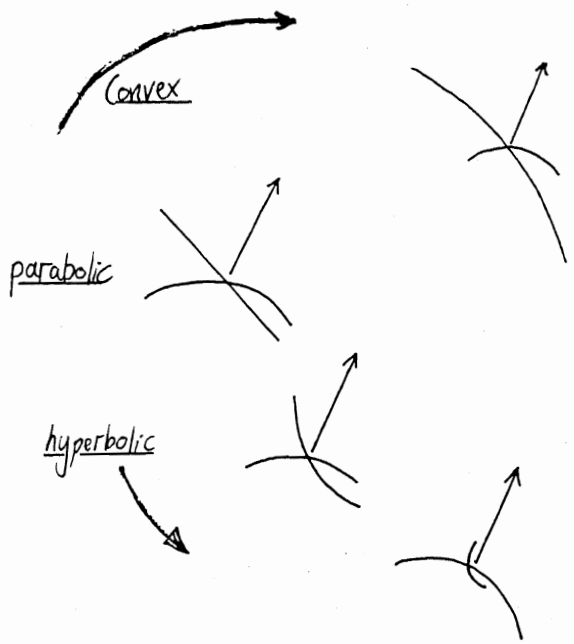

Figure 7: Specular stereo - constraints on local shape. A typical constraint, illustrated algebraically in figure 6, admits a one-parameter family of interpretations. Generally either concave or convex interpretations are excluded - in this case concave ones.

Note that a concave interpretation is excluded - the surface must be either convex or hyperbolic in this case. Generally a concave or convex interpretation is excluded according as the sign of A is positive or negative respectively - that is, according to the sign of $\delta\mathbf{n}.\mathbf{x}$. Recently a similar condition has been obtained by Zisserman et al. [27] but with the additional advantage of being independent of light source position. In that case, the discriminant is simply the scalar product of the projections onto the image plane of $\mathbf{x}$ (i.e. the relative disparity vector) and the baseline $\mathbf{d}$.

2.5 Test for umbilic points

Spheres are an interesting special case for specular stereo. Since $H = \kappa I$ on the surface of a sphere, we can see from

equation (8) that $\delta\mathbf{n}$ and $\mathbf{x}$ are parallel vectors or, in practice:

$$|\delta\mathbf{n} \times \mathbf{x}| \ll |\delta\mathbf{n}||\mathbf{x}|. \quad (13)$$

If the parallelism test is passed the point *may* be umbilic - but of course this is not guaranteed. For instance it could be on any surface patch specially oriented so that one line of curvature lies (locally) in the plane of the incident and reflected rays. So the test can be used to eliminate umbilic interpretations.

3 Specular stereo as a linear system

The previous section explained how constraints on geometry can be inferred from stereoscopic observation of a specularity and a nearby non-specular "reference" point. Further analysis can be applied, linearising the relationship between displacement vector $\mathbf{x}$ and baseline vector $\mathbf{d}$. Whilst this offers no particular improvements in convenience of computation of shape constraints, it affords insights into geometric degeneracies, and clarifies the relationship with conventional stereoscopic disparity.

Simplifying previous analysis [2], the linear system can expressed as

$$2V(MH - \kappa_{VL}I)\mathbf{x} = \mathbf{w}, \quad (14)$$

where

$$\mathbf{w} = (-d_1 + d_3 \tan\sigma, -d_2)^T, \quad (15)$$

$$M = \begin{pmatrix} \sec\sigma & 0 \\ 0 & \cos\sigma \end{pmatrix}, \quad (16)$$

σ is surface slant at A, and

$$\kappa_{VL} = \frac{1}{2}\left(\frac{1}{V} + \frac{1}{L}\right) \quad (17)$$

- a term familiar from the elementary ray-optics of lenses. It can be thought of as an "apparent curvature" induced on a viewed plane, owing to the finite distances from the plane to source and viewer. This linear approximation is valid whenever the baseline is relatively short, that is, when

$$|\mathbf{d}| \ll |\mathbf{V}| \cos\sigma, \quad (18)$$

and provided the surface does not focus incoming rays to a point or line close to the centre of projection (see discussion of degeneracy below).

3.1 Horizontal and vertical disparity

One insight that equation (14) affords is that the specular stereo constraints on H depend solely on imaging geometry ($\kappa_{VL}, V, M, \mathbf{w}$) and on the displacement vector $\mathbf{x}$. So the only *measured* quantity involved is $\mathbf{x}$, which is intimately related to the relative stereoscopic disparity δ, as follows:

$$\mathbf{x} = VP\delta \quad (19)$$

where

$$P = \begin{pmatrix} \sec\sigma & 0 \\ 0 & 1 \end{pmatrix}. \quad (20)$$

The two-component vector δ represents the horizontal and vertical disparities of the specularity, *relative* to the disparities of the nearby surface reference point. The conventional view of stereoscopic vision is that useful depth information is encoded entirely in horizontal disparities. Vertical disparity is fixed by epipolar geometry, so its possible influence is limited to calibration of view geometry and, in human vision, an associated illusory distortion, the induced effect [19].

For specularities, this is not the case. Vertical disparity plays a strong role, in two ways. First, vertical disparities that violate epipolar geometry are *evidence* that specular reflection is occurring. Second, as we just saw, measured vertical disparity imposes an independent constraint on curvature, in addition to the one imposed by measurement of horizontal disparity. Both of these computational "truths" call for psychophysical investigation. Is either theory exploited in human vision?

3.2 Degeneracy

Inspection of the linearised imaging equation (14) reveals degeneracy when

$$\det(MH - \kappa_{VL}I) = 0.$$

What exactly is observed physically? First of all, this can happen only on non-convex surfaces (MH is negative definite on a convex surface) and even then only for special alignments - when the viewer collides with one or both "focal surfaces" [27][2]. A convex surface lies *in front* of its focal surfaces; the viewer cannot collide with a focal surface because the convex surface is in the way. When degeneracy does occur, stereoscopic analysis fails for the very simple reason that the specularity is visible only in one eye. Moreover, the focusing effect ensures that when it is visible, it is likely to be very bright.

3.3 Combining information from two baselines

A final result from the linearised view is that when two independent baselines are used, for instance when a stereo observer is in motion, the baselines should not be nearly parallel. If they are not, $\mathbf{H}$ can be recovered completely; if they are, computation of $\mathbf{H}$ is ill-conditioned. Details of the argument are given in [2].

[2]The specularity is focussed onto a line or onto a blob, according as the rank of $MH - \kappa_{VL}I$ is 1 or 0 respectively.

4 Monocular analysis of specular reflection

A simple theory of monocular analysis of specularity [2] is summarised here, and possible ambiguity of interpretation is explored. In the case of a circular source, assuming that the diameter of the source is known, there is a possible fourfold ambiguity of interpretation, corresponding roughly to independent inversion of each principal curvature, but generally accompanied by some rotation (about the surface normal) of the lines of curvature.

The geometrical arrangement for monocular viewing, with a distributed source, is shown in figure 8. It mirrors the earlier stereo geometry but with the baseline between stereo views replaced by virtual baselines between pairs of points on the source. This duality can be tapped mathematically to derive a linear mapping relating the position

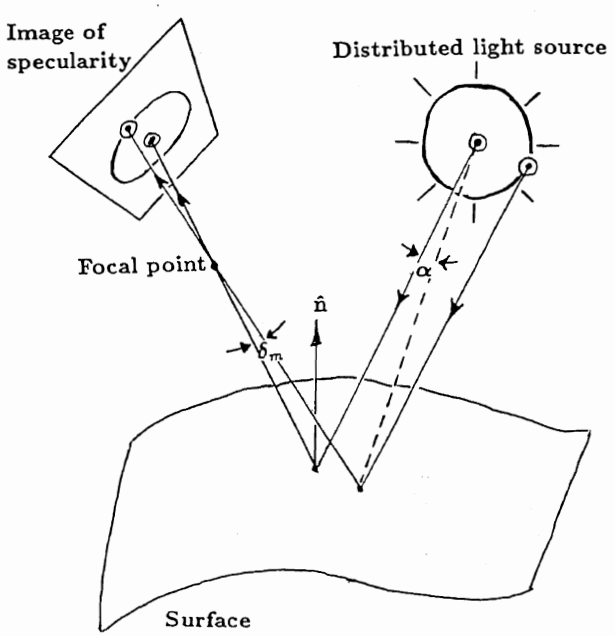

Figure 8: Monocular analysis: A distributed light source of known shape is reflected by a curved surface as a specular image region. Surface curvature information may be inferred by measuring the shape in the image. A point on the source, transforms to a point on the specularity. The angular positions of two such points are specified by α and δ_m, for source and specularity respectively.

δ_m of a specular point (in polar projection) to the angular position α on the source from which the illuminating ray came.

$$T\delta_m = \alpha \tag{21}$$

where

$$T = 2VP^{-1}MH^*P, \tag{22}$$

where

$$H^* = H - \kappa_{VL}M^{-1}, \tag{23}$$

and P, H, M, κ_{VL} are defined as previously. Since T is symmetric, the mapping is a linear scaling in two orthogonal directions.

Consider the case of a circular source, so that points on the outline of the source satisfy

$$\alpha^T\alpha = \rho^2. \tag{24}$$

The shape of the ellipse (on the polar projection) is obtained using the transformation (21) to give (using the fact that T is symmetric):

$$\delta_m^T T^2 \delta_m = \rho^2. \tag{25}$$

If we assumed the source radius ρ were known, then observation of the elliptical shape of a highlight determines T^2, via equation (25). Hence T is known up to a fourfold ambiguity (the signs of the two eigenvalues of T are unknown). Then H can be computed directly from (21), so again there are four possibilities (figure 9). These four in-

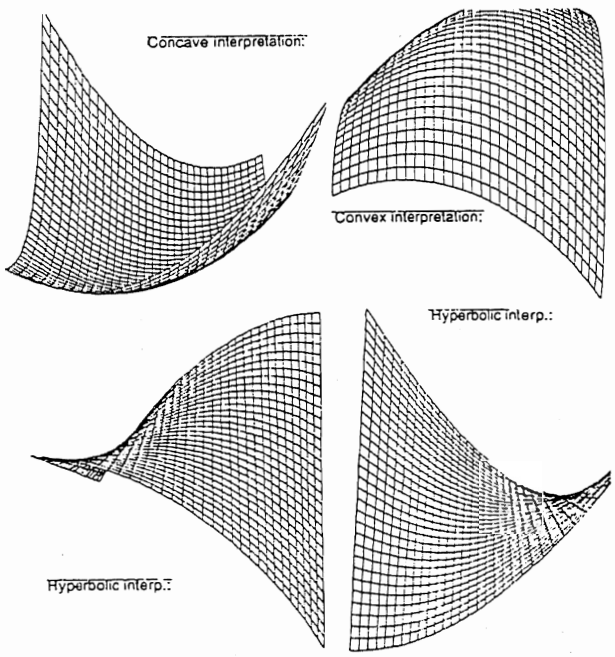

Figure 9: Observation of an elliptical specularity, with a circular source of known radius, determines local curvature up to a fourfold ambiguity, as illustrated.

terpretations generally do *not* share common lines of curvature; but the discrepancy is small if the slant is modest. Moreover, provided principal curvatures are neither small nor nearly equal in magnitude -

$$||\kappa_1| - |\kappa_2|| \gg 2\kappa_{VL} \sec\sigma,$$

the four possibilities for directions of lines of curvature collapse down to just two possibilities - this is the case in the example of figure 9.

Usually, real cameras do not form an image by polar projection, but by perspective projection. However the difference is more or less cosmetic. At a given point on the

image, perspective projection is related to polar projection by a linear transformation $\delta_m = Q\mathbf{X}$, where $\mathbf{X}$ is a position vector on the projection plane, and Q is a matrix. The elliptical specularity that appears on the image plane is given by

$$\mathbf{X}^T Q^T T^2 Q \mathbf{X} = \rho^2. \tag{26}$$

Measurements made within the image plane can provide $(Q^T T^2 Q)$, from which T^2 can be computed.

4.1 Assuming a circular source

The analysis above assumed that the angular diameter (subtended at the surface) of the source was known. Often a more reasonable assumption is that the source is circular, but of *unknown* radius. There are now two possible cases.

1. The surface is locally almost flat - the magnitudes of the principal curvatures are of the order of κ_{VL} or smaller.

2. At least one principal curvature κ_1 is large - that is, it satisfies

   $$\kappa_1 \gg \kappa_{VL} \sec\sigma.$$

   This is much the more likely state of affairs since specularities tend to cling to highly curved patches - hence the likelihood of observing a specularity on such a patch is relatively great. In this case, since the $2V\kappa_{VL}I$ term in (22) is negligible, and since T^2 is known up to an arbitrary multiplicative constant, H can be computed from (22) up to fourfold ambiguity *and* an arbitrary multiplicative constant. However, computed curvatures are accurate only to within $\kappa_{VL} \sec\sigma$.

So even when absolute source size is unknown, monocular observation of a specularity still allows some inference about surface shape. This is what might be expected intuitively; an elongated specularity, for example, seems to suggest very unequal principal curvatures - a locally cylindrical surface in the extreme case. In fact this particular example might be expected to hold good even when the source is not known to be circular, but merely to be "compact". That expectation is justified below.

4.2 Assuming a compact source

First we must define what is meant by a compact source. A source of compactness $K > 1$ is defined as one that is bounded by concentric circles, with radii in the ratio K : 1, as in figure 10. The most compact source is therefore a circle ($K = 1$). The absolute size of the source is assumed to be unknown.

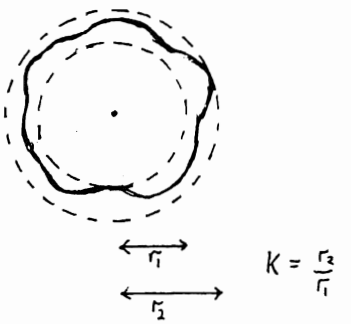

Figure 10: Definition of compact source, in terms of bounding circles.

The task now is to use the monocular specularity equations (21, 22, 23) to make some inference about surface curvature, even though the shape of the source is unknown, merely constrained. It is proposed to estimate the direction of least curvature on the surface as follows. Take the direction in which the diameter of the specular blob is greatest (figure 11), and back-project it onto the surface tangent plane. If this estimate proves reliable it provides a third constraint, in addition to the two already provided by stereoscopic analysis. Then surface curvature, represented by the three parameters of H, can be computed [2].

However, there are three problems to be addressed:

1. The direction in which the diameter of the specular blob is greatest does not correspond exactly to an eigenvector direction of T.

2. If $\mathbf{u}$ is an eigenvector of T, then in the limit that its eigenvalue $\lambda \to 0$, it can be seen from (22) that $P\mathbf{u}$ is an eigenvector of H^* (with eigenvalue 0). However, when $\lambda \neq 0$ $\mathbf{u}$ only approximates to an eigenvector of H^*.

3. Even if the eigenvectors of H^* are successfully estimated, they correspond only approximately (23) to eigenvectors of H.

What does all this amount to? Our procedure for estimating directions of lines of curvature works, but is only approximate. The approximation improves as

- the blob becomes more elongated
- the source becomes more compact
- the surface slant decreases
- the surface becomes more curved (either cylindrically or elliptically).

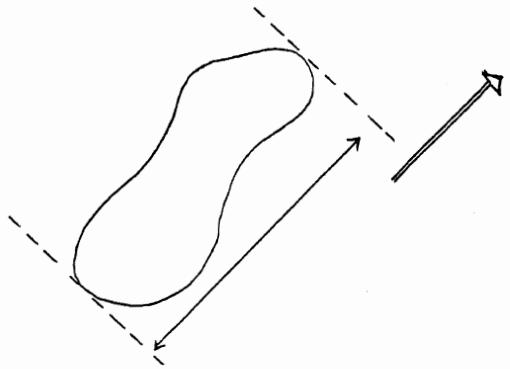

Figure 11: Measuring the direction in which the diameter of a specular blob is largest. This corresponds approximately to an eigenvector of T. The approximation improves as the blob becomes more elongated, and as the source becomes more compact.

The last of these factors deserves an additional comment. In fact our approximation will break down when the surface is close to planar (both curvatures small compared with $\kappa_{VL} \sec\sigma$). In that case nothing further can be inferred from the shape of the specularity. However this case will be relatively uncommon; specularities prefer to cling to highly curved surfaces. A specularity on a plane tends to be relatively fleeting, disappearing at the slightest viewer motion, especially when the total field of view is relatively narrow. In the remainder of the section these claims are backed up by the statement of formal results. Proofs use standard results of linear algebra, but are omitted here for lack of space.

The following three theorems answer, in turn, the three problems raised above, by providing error bounds. First, some definitions and assumptions. The source has compactness K. The specularity has aspect ratio $A > K^2$ (A is the ratio of maximum diameter to minimum diameter). The surface has curvatures $\kappa_{max}, \kappa_{min}$ with

$$|\kappa_{max}| > |\kappa_{min}|.$$

The surface is non-planar:

$$|\kappa_{max}| > 5\kappa_{VL} \sec\sigma.$$

Theorem 1 *The angular error in taking the direction in which the diameter of the specularity is greatest to be an eigenvector of T is bounded by*

$$\tan^{-1}\left\{K\left(\frac{K^2-1}{A^2-K^4}\right)^{\frac{1}{2}}\right\}. \qquad (27)$$

Theorem 2 *If $\mathbf{u}$ is an eigenvector of T then the angular difference between $P\mathbf{u}$ and an eigenvector of H^* is bounded by*

$$\min\Big[\ \sin^{-1}\left\{\sin\sigma\left(1-\sec^2\sigma\tfrac{K}{A}\right)^{-\frac{1}{2}}\right\},\ \sin^{-1}\left\{\tan\sigma\left(\cos^2\sigma\tfrac{A}{K}-1\right)^{-\frac{1}{2}}\right\}\ \Big]. \qquad (28)$$

Theorem 3 *The angle between eigenvectors of H, H^* is bounded by*

$$\sin^{-1}\left\{\sin\sigma\left(3\frac{|\kappa_{min}|}{|\kappa_{max}|}+8\frac{K}{A}\sec^2\sigma\right)^{\frac{1}{2}}\right\}. \qquad (29)$$

5 Results from the implemented system

Some of the principles described above have been incorporated in a software system. The main components are:

Feature detection Features for conventional stereo matching are obtained from an edge detector employing directional derivative-of-gaussian operators [6]. A combination of global and local analysis of the spatial distribution of intensity is then applied to identify specular reflections [4]. Those edge features which subsequently proved to be specular are then pruned from the list of "surface" features. Remaining surface features can then be matched using standard methods, and computed depths can safely be regarded as representing points that lie on a real surface. The danger of treating matched specularities as surface points is thus removed.

Feature description Specular features form "blobs" which may be irregular, but are often elliptical. The feature description consists of the direction of elongation of the blob, its length and its aspect ratio. In the case of an elliptical blob this is more or less equivalent to measuring the directions and lengths of the ellipse axes. For each specular feature, a nearby surface feature must also be identified, to serve as a "reference" for depth or disparity. The system simply chooses the surface feature that is closest to the specularity (in the left image); the closer it is the more accurate it is as a reference. However, features consisting of almost horizontal line segments are avoided as they are well known to yield inaccurate depth measurements.

Stereoscopic correspondence Our system employs the PMF software for stereoscopic correspondence and depth computation. This delivers depths of surface points, including those to be used as disparity references. Correspondence must also be established between specularities in right and left images. A conventional matcher is unsuitable here because epipolar constraints do not apply. On the other hand,

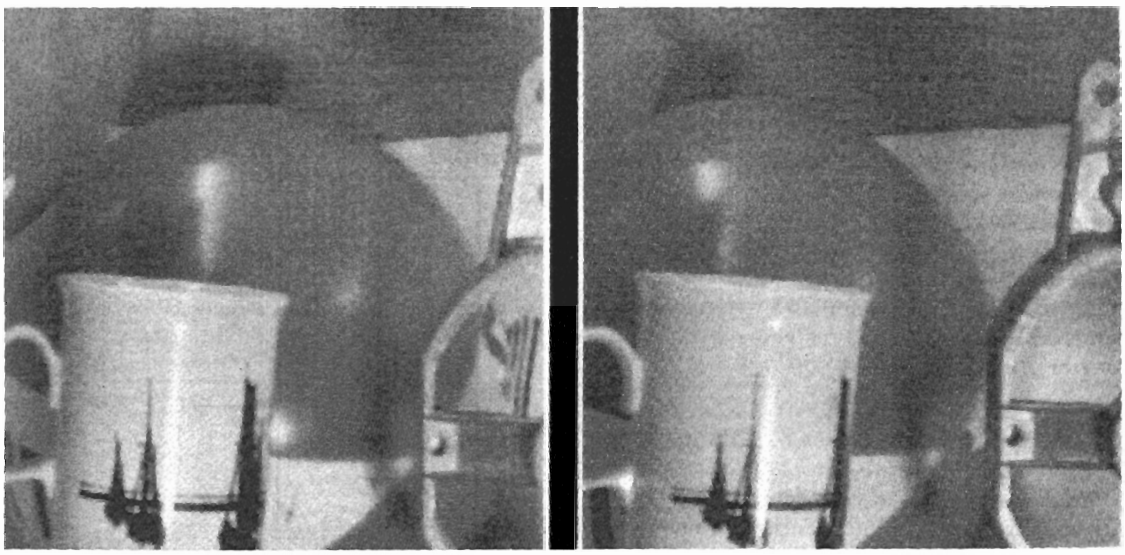

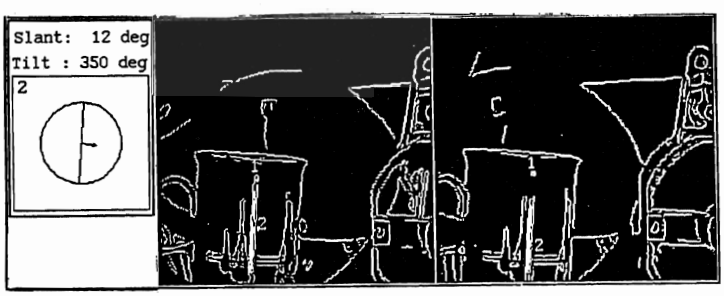

Figure 12: An interactive implementation of the analysis of specularities. (a) Stereo image pair. (b) Results of processing: edge detection output is labelled with specularities (marked 1 and 2); slant, tilt and line of (least) curvature for specularity number 2 are displayed at the bottom left.

the problem of false matches negligible for specular features since they are usually sparse. A simple matcher, employing rough comparison of blob features, has proved sufficient in our experiments.

Geometric inference Stereoscopic and monocular inference of geometry proceed separately, as detailed in previous sections, and finally inferences are pooled. In some cases local curvature is completely determined, in others merely constrained. Results are displayed by means of appropriate graphics, together with error bounds. They are also accessible at program-level for use in model-matching, geometric reasoning or other applications.

Error treatment Each measured quantity in the stereo analysis has uncertainty associated with it; this can be represented crudely by propagation of error bounds. By combining the errors it is possible to quantify the uncertainty of the quantities in the constraint equation (14). This can be done by summing square errors [22] at each step in the analysis. This method of combination strictly only applies to independent sources of uncertainty. As the uncertainties involved here are unlikely to be completely independent, room exists for refinement.

An example of the system in operation is shown in figure 12. Subwindows allow user intervention at various levels, from selection of images to tracing the inference steps. Line drawings at the bottom show detected, matched specularities, labelled 1 and 2. Results of geometric inference are summarised in the display at the lower left. The ellipse and needle indicate surface orientation. Numerical slant and tilt values are also shown. The line indicates the direction of the line of least curvature - note that this appears to coincide with the axis of the cylindrical cup on which the specularity (number 2) lies. In fact the cup is not quite cylindrical, and the system has inferred (see COMBINED EVIDENCE window) that the surface is hyperbolic. Further detail can be obtained by selecting "graph" to illustrate stereoscopic and monocular inferences, and their combination. The graph for stereoscopic inference is shown in figure 13. In an ideal, error-free world it would be simply a hyperbola as in figure 6. The effect of allowing for error in the components of the Hessian H is that the hyperbola is "thickened". So on the basis of stereoscopic information, the principal curvatures may take any values in the shaded set. Note that the *directions* of the lines of curvature are not fixed, but differ for different points in the set.

Now when monocular information is taken into account, only a small portion of the shaded set remains feasible - shown in black on figure 14. The combination of information works very simply: the specularity is observed to be very elongated, and this more or less fixes the directions of the lines of curvature (see earlier discussion). The black region consists of those solutions from the shaded set which correspond approximately to those computed directions.

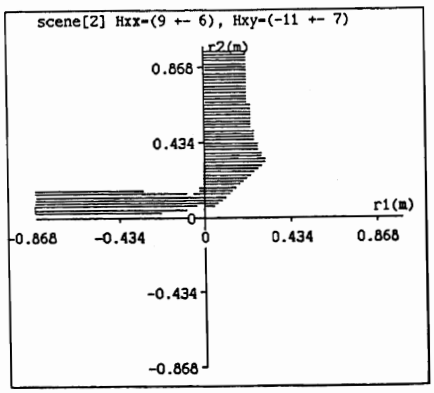

Figure 13: Constraints from stereoscopic analysis of specularity number 2, in figure 12.

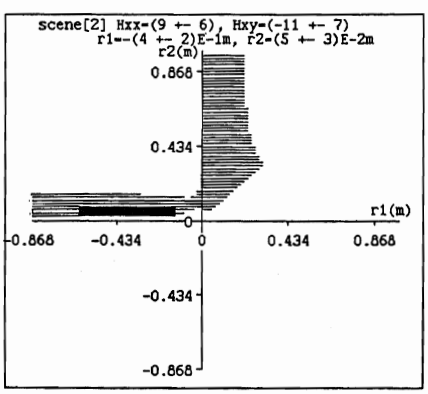

Figure 14: Combined stereoscopic and monocular constraints restrict the set of possible solutions for surface curvature to the black region.

Stereoscopic analysis for the other specularity (number 1 in figure 12) is shown in figure 15. In this case the

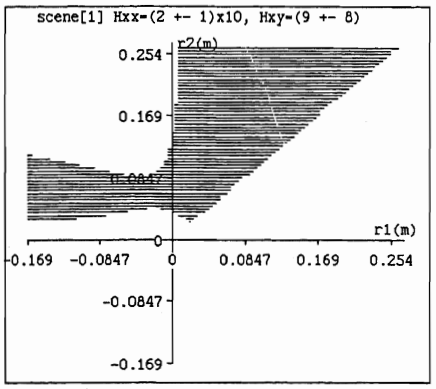

Figure 15: Stereoscopic analysis of specularity number 1 in figure 12 restricts local surface curvatures to values within the shaded set.

specularity appears as a small blob, too small reliably to determine its shape. Hence (in the absence of information about the absolute size of the source) monocular analysis does not constrain local shape any further.

The pair of stereo images in figure 16 is computer generated, using a narrow field of view to exaggerate disparities. This makes it easy to see relative displacement of specularities, and how they relate to the shape of the underlying surface. The specularity is displaced both horizontally *and* vertically relative to surface markings. That is, there are both horizontal and vertical relative disparities - a vivid illustration of the earlier assertion that specularities may break epipolar constraints. (The specularity tends to cling to the line of greatest curvature, hence the relative motion is oblique.) Indeed, the non-zero relative vertical disparity is *evidence*, in principle, that the bright blob is indeed specular. The horizontal disparity is positive, leading correctly to a prediction that the surface is convex [27]. The true values of principal curvatures (denoted by the white cross) lies within the black region that indicates predicted curvatures with error bounds. Of course the data is computer generated and free of noise, but this does at least indicate that the error bound computations, which determine the extent of the feasible set (in black) in terms of image measurement errors, are correct.

Finally, a real image (figure 17) is shown below, together with the results of geometric inference from specularity analysis. Surface curvatures have measured approximately, and fall within error bounds computed by the system. There are stereoscopic constraints as shown, but no monocular constraints; this is because the specularity is so nearly circular that ellipse axes cannot be reliably computed. Nonetheless, the shaded set does contain (just!) the measured values of surface curvature. Moreover, the black window at the bottom of the graph indicates that the system has found that the surface could be locally spherical (umbilic), according to the test described earlier (13).

6 Conclusions

Analysis of specularity is of potential assistance to geometric inference in machine vision. Stereoscopic and monocular analysis are complementary, and together can entirely determine local shape. A further role for vertical disparity has been demonstrated, in addition to its role in calibration of viewing geometry. Analysis of specularity has been incorporated into a stereoscopic vision system, and shown to yield usably accurate results. Some questions remain. How can such local shape measurements be integrated? Quantitative methods involving stereoscopic reconstruction are certainly available [11, 23] - can qualitative methods, bypassing depth maps, be found? Another question concerns motion: how much surface information can be extracted from specularities under extended displacement of the viewer?

A number of questions are raised too for human vision.

- Is vertical disparity - violation of epipolar constraints - used to identify bright features as specular?
- Can perceived surface curvature be manipulated by adjusting the disparities of a specularity?
- Can predicted fourfold ambiguity of curvature be realised in monocular views of specularities?
- Do monocular and stereoscopic analysis of specularities combine to fix perceived curvature, as predicted theoretically?
- Can the direction of displacement of a specularity in right and left stereoscopic views, resolve reversal ambiguities, as is theoretically predicted?

Acknowledgements

Discussion with A.Zisserman and with J.Mayhew, J.Frisby and S.Pollard is gratefully acknowledged. The PMF stereo system from AIVRU, Sheffield University proved invaluable and for the basis for experimentation. The support of IBM UK Scientific Centre, the Royal Society of London, the SERC and the University of Edinburgh are gratefully acknowledged.

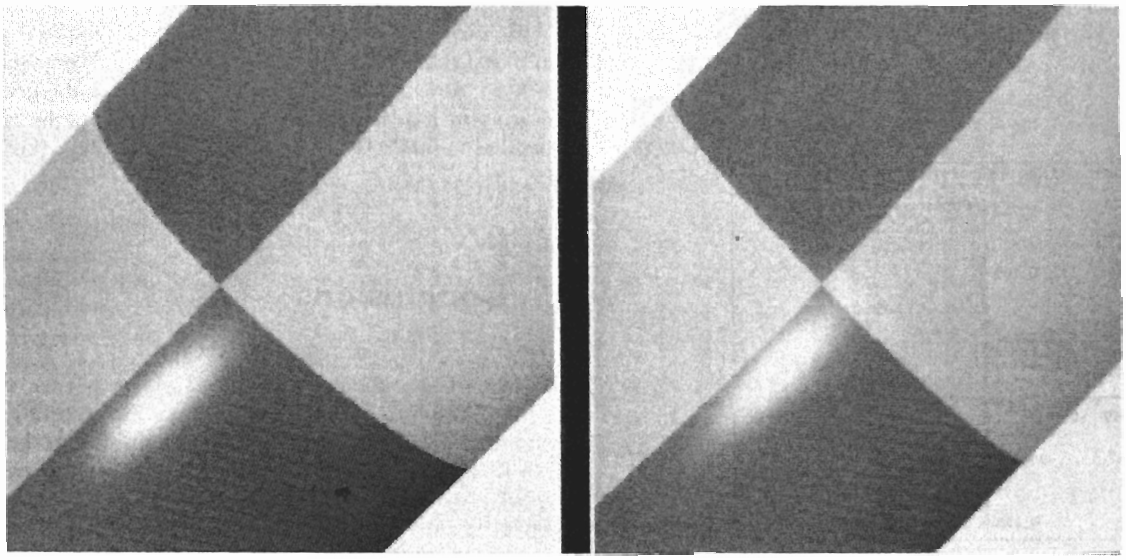

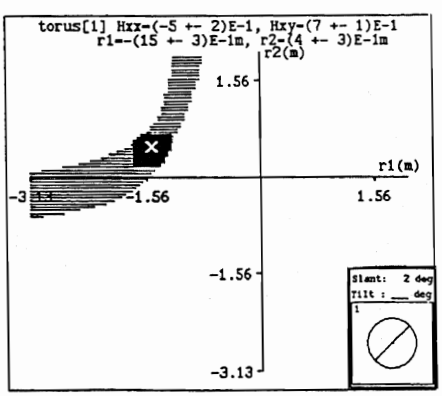

Figure 16: Inner surface of of a torus ring with surface markings (artificial images). Stereoscopic constraints, shown shaded, are restricted to the black region by monocular analysis. The correct values of curvature (used in the geometric modeller) are indicated by a white cross.

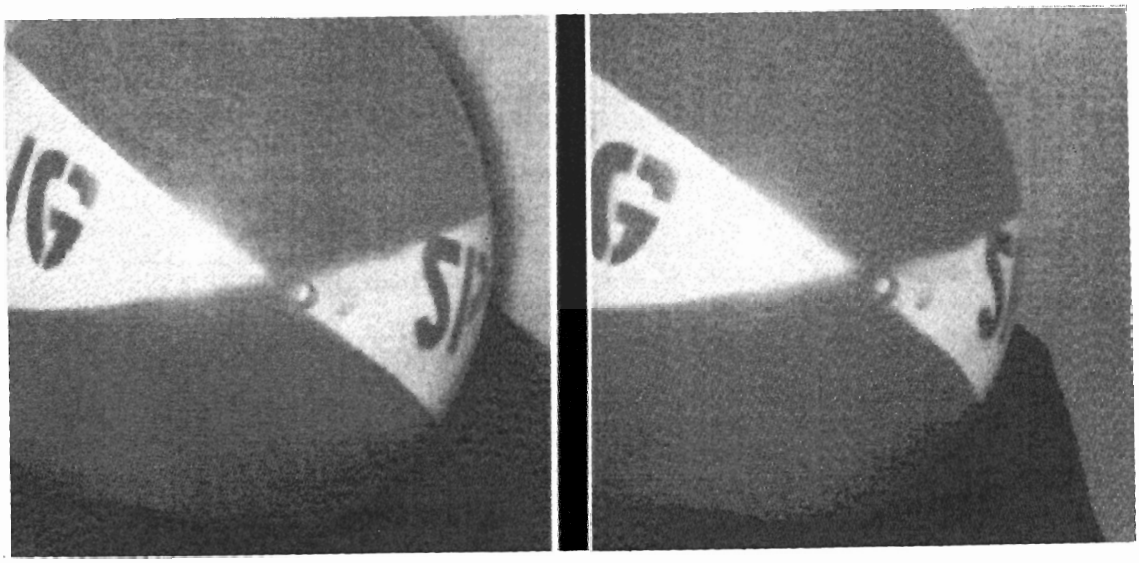

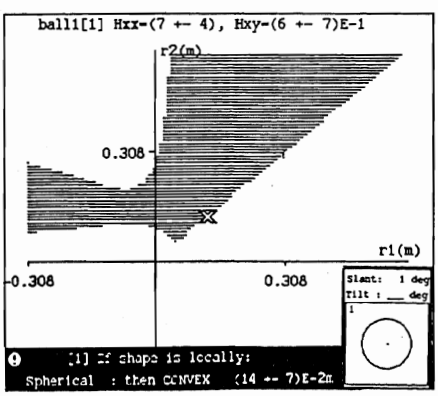

Figure 17: Beach ball of 12cm radius.

References

1. Babu, M.D.R., Lee, C-H. and Rosenfeld, A. (1985). Determining Plane Orientation From Specular Reflectance, *Pattern Recognition*, 18, 1, 53-62.

2. Blake, A. (1985). Specular Stereo. *Proc. IJCAI-85*, 2, 973-976.

3. Blake,A. (1985). Boundary conditions for lightness computation in Mondrian world. *CVGIP, 32*, 314-327.

4. Brelstaff,G.J. (1988). PhD thesis, University of Edinburgh.

5. Buchanan, C.S.B. (1987). Determining Surface Orientation from Specular Highlights, RCBV Tech. Report, RCBV-TR-87-19, University of Toronto.

6. Canny,J.F. (1983). Finding edges and lines in images. S.M. thesis, MIT, Cambridge, USA.

7. Coleman, E.N.Jr., and Jain, R. (1982) Obtaining 3-Dimensional Shape of Textured and Specular Surfaces Using Four-Source Photometry, Computer graphics and image processing, 18, 309-328.

8. Forbus, K. (1977). Light Source Effects A.I. Memo 422, A.I. Lab., M.I.T.

9. Gershon, R., Jepson, A.D. and Tsotsos, J.K. (1987). Highlight Identification Using Chromatic Information, *Proc. IJCAI87 Milan*, 752-754.

10. Grimson, W.E.L. (1981). *From Images to Surfaces* M.I.T. Press, Cambridge Mass. U.S.A..

11. Grimson, W.E.L. (1982). Binocular Shading and Visual Surface Reconstruction. M.I.T. A.I.Lab. Memo No 697.

12. Healey, G., and Binford, T.O. (1987). Local shape from specularity. Proc. First Int. Conf. on Computer Vision, London, 151-160

13. Horn, B.K.P. (1974). Determining Lightness from an Image. Computer Graphics and Image Processing, 3, 277-299.

14. Ikeuchi, K. (1981). Determining Surface Orientation of Specular Surfaces by Using the Photometric Stereo Method. IEEE Trans. PAMI, 3, 6, 661-669.

15. Koenderink, J.J. and van Doorn, A,J. (1980). Photometric invariants related to solid shape. Optica Acta, 27, 7, 981-996.

16. Land, E.H. (1983). Recent advances in retinex theory and some implications for cortical computation: Color vision and natural images. *Proc. Natl. Acad. Sci. U.S.A.* 80, (1983), 5163-5169.

17. Klinker, G.J., Shafer, S,A., and Kanade, T. (1987). Using a color reflection model to separate highlights from object color. Proc. First Int. Conf. on Computer Vision, London, 145-150.

18. Mayhew,J.E.W and Frisby,J.P. (1981). Towards a computational and psychophysical theory of stereopsis. *AI Journal*, 17, 349-385.

19. Mayhew,J.E.W. and Longuet-Higgins,H.C. (1982). A computational model of binocular depth perception. *Nature*, 297, 376-378.

20. Sanderson, A.C., Weiss, L.E. and Nayar, S.K. (1988). Structured Highlight Inspection of Specular Surfaces. *IEEE Trans. PAMI* 10, 1, 44-55.

21. Stevens,K.A. (1979). Ph.D. Thesis, MIT, Cambridge, USA.

22. Squires, G.L. (1968). *Practical Physics*, European Physics Series, McGraw-Hill, London.

23. Terzopoulos, D. (1985). Multilevel computational processes for visual surface reconstruction. *Computer Vision Graphics and Image Processing*, 24, (1985), 52-96.

24. Thrift, P. and Lee, C-H. (1983). Using Highlights to Constrain Object Size and Location. *IEEE Trans. Sys, Man and Cybernetics*, 13, 3, 426-431.

25. Ullman, S. (1976). On Visual Detection of Light Sources. *Biol. Cybernetics*, 21, (1976), 205-212.

26. Woodham, R.J. (1980). Photometric method for determining surface orientation from multiple images. Optical Engineering, 19, 1, 139-144.

27. Zisserman,A., Giblin,P.J. and Blake,A. (1988). The information available to a moving observer from specularities. *Proc. Alvey Conf. 1988.*