Robust Estimation of Surface Curvature from Deformation of Apparent Contours

Andrew BLAKE Roberto CIPOLLA
Department of Engineering Science
University of Oxford
Oxford OX1 3PJ, U.K.

Abstract

Surface curvature along extremal boundaries is potentially useful information for navigation, grasping and object identification tasks. Previous theories have shown that qualitative information about curvature can be obtained from a static view. Furthermore it is known that, for orthographic projection, under planar viewer-motion, quantitative curvature information is available from spatio-temporal derivatives of flow. This theory is extended here to arbitrary curvilinear viewer-motion and perspective projection.

We show that curvatures can actually be computed this way in practice, but that they are highly sensitive to errors in viewer-motion estimates. Intuitively, relative or *differential* measurements of curvature might be far more robust. Rather than measuring the absolute deformation of an apparent contour, *differential* quantities depend on the rate at which surface features are swept over an extremal boundary as the viewer moves. It is shown that, theoretically, such *differential* quantities are indeed far less sensitive to uncertainty in viewer-motion. *Ratios* of differential measurements are less sensitive still. In practice sensitivity is reduced by about two orders of magnitude. We believe this represents a significant step in the development of practical techniques for robust, qualitative 3D vision.

1 Introduction

The deformation of an apparent contour (the silhouette of a smooth surface or the image of the extremal boundary) under viewer-motion is a potentially rich source of geometric information for navigation, motion-planning and object-recognition. Barrow and Tenenbaum [*Barrow78*] pointed out that surface orientation along an extremal boundary can be computed directly from image data. Koenderink [*Koenderink84*] related the curvature of an apparent contour to the intrinsic curvature of the surface (Gaussian curvature); the sign of Gaussian curvature is equal to the sign of the curvature of the contour. Convexities, concavities and inflections of an apparent contour indicate, respectively, convex, hyperbolic and parabolic surface points. Giblin and Weiss [*Giblin87*] have extended this by adding viewer motions to obtain quantitative estimates of Gaussian and mean curvature. A surface can be reconstructed from the envelope of all its tangent planes, which in turn are computed directly from the family of silhouettes of the surface, obtained under planar motion of the viewer. By assuming that the viewer follows a *great circle* of viewer directions around the object they restricted the problem of analysing the envelope of tangent planes (a 2-parameter family) to the less general one of computing the envelope of a family of lines in a plane. Their algorithm was tested on noise-free, synthetic data (on the assumption that extremal boundaries had been distinguished from other image contours) demonstrating the reconstruction of a planar curve under orthographic projection.

This paper extends these theories further, to the general case of arbitrary non-planar camera motion under perspective projection. The Gaussian curvature of a surface, at a point on its silhouette, can be computed given some known local motion of the viewer. Curvature is computed from spatio-temporal derivatives (up to second order) of image-measurable quantities. The theory can, of course, be applied to detect extremal boundaries and distinguish them from surface markings or discontinuities. Experiments show that, with adequate viewer-motion calibration, itself computed from visual data [*Tsai87*], it is possible to obtain curvature measurements of useful accuracy.

A consequence of the theory, representing an important step towards qualitative vision, concerns the robustness of *relative* or *differential* measurements of curvature at two nearby points. Intuitively it is relatively difficult to judge, moving around a smooth, featureless object, whether its silhouette is extremal or not — whether the Gaussian curvature along the contour is bounded or not. This judgment is much easier to make for objects with feature-rich surfaces. Under small viewer-motions, features are "sucked" over the extremal boundary, at a rate which depends on surface curvature. Our theory reflects this intuition exactly. It is shown that relative measurements of curvature across two adjacent points are entirely immune to uncertainties in the viewer's rotational velocity. This is somewhat related to earlier results showing that relative measurements of this kind are important for depth measurement from optic flow [*LHiggins80, Weinshall89*] and for curvature measurements from stereoscopically viewed highlights [*Blake88*]. Furthermore, they are relatively immune to uncertainties in translational motion in that, unlike single-point measurements, they are independent of the viewer's acceleration. Only dependence on velocity remains. Experiments show that this theoretical prediction is borne out in practice. *Differential* or relative curvature measurements prove to be more than an order of magnitude less sensitive than single-point measurements to errors in viewer-motion calibration. There is some theoretical evidence that *ratios* of differential curvature measurements are less sensitive. In our experiments absolute measurements of curvature were so sensitive that they became unreliable for viewer motion errors of 0.5mm in position and 1mrad in orientation. For ratios of *differential* measurements of curvature the corresponding figures were about 50mm and 70mrad respectively.

2 Theoretical framework

2.1 Surface Geometry

Consider a point P on the extremal boundary of a smooth, curved surface in R^3 and parameterised locally by a vector valued function $\mathbf{r}(s,t)$. The parametric representation can be considered as covering the surface with 2 families of curves: $\mathbf{r}(s,t_0)$, and $\mathbf{r}(s_0,t)$ where s_0, t_0 are fixed for a given curve in the family. A one-parameter family of views is indexed by the time parameter t and s,t are defined so that the s-parameter curve, $\mathbf{r}(s,t_0)$, is an extremal boundary for a particular view t_0. A t-parameter curve $\mathbf{r}(s_0,t)$ can be thought of as the 3D locus of points grazed by a light-ray from the viewer, under viewer-motion. Such a locus is not uniquely defined.

Local surface geometry can be specified in terms of the basis $\{\mathbf{r}_s, \mathbf{r}_t\}$ for the tangent plane ($\mathbf{r}_s$ and $\mathbf{r}_t$ denote $\partial\mathbf{r}/\partial s$ and $\partial\mathbf{r}/\partial t$ respectively) and the surface normal - a unit vector $\mathbf{n}$.

2.2 Imaging model

The imaging model is a spherical pin-hole camera of unit radius. The image of the world point, P, with position vector $\mathbf{r}(s,t)$ is a unit vector $\mathbf{T}(s,t)$ defined by

$$\mathbf{r}(s,t) = \mathbf{v}(t) + \lambda\mathbf{T}(s,t), \tag{1}$$

where λ is the distance along the ray to the point P (figure 1).

For a given vantage position t_0 the apparent contour is a continuous family of rays $\mathbf{T}(s,t_0)$ emanating from the camera's optical centre which touch the surface so that $\mathbf{T}.\mathbf{n} = 0$. The moving observer at position $\mathbf{v}(t)$ sees a 2 parameter family of apparent contours $\mathbf{T}(s,t)$.

2.3 Properties of the extremal boundary and its projection

In [*Blake89*] we derive for perspective projection the following well-known properties of the extremal boundary and its projection [*Barrow78*, *Koenderink82*, *Brady85*, *Giblin87*].

1. The orientation of the surface normal, $\mathbf{n}$, can be recovered by measuring the direction of the ray $\mathbf{T}$ of a point on an extremal boundary and the tangent to the apparent (image) contour, $\mathbf{T}_s$.
$$\mathbf{n} = \frac{\mathbf{T} \wedge \mathbf{T}_s}{|\mathbf{T} \wedge \mathbf{T}_s|}. \tag{2}$$

2. The ray direction, $\mathbf{T}$, and the tangent to the extremal boundary, $\mathbf{r}_s$, are in *conjugate* directions with respect to the second fundamental form
$$\mathbf{T}.\mathbf{n}_s = 0 \tag{3}$$
The ray direction and the extremal boundary will only be perpendicular if the ray $\mathbf{T}$ is along a *principal* direction.

3. The curvature of the apparent contour, κ^p, (more precisely the geodesic curvature of the curve, $\mathbf{T}(s,t_0)$ which can be computed from spatial derivatives in image measurements), can be written in terms of the *normal curvature* of the extremal boundary, κ^s:
$$\kappa^p = \frac{\mathbf{T}_{ss}.\mathbf{n}}{|\mathbf{T}_s|^2} \tag{4}$$
$$= \lambda\frac{\kappa^s}{1-(\mathbf{T}.\frac{\mathbf{r}_s}{|\mathbf{r}_s|})^2} \tag{5}$$
Equations (4) and (5) shows that an apparent contour is smooth except for a special viewing geometry when the ray direction runs along the extremal boundary. At such points the apparent contour may have a cusp. For opaque surfaces only one branch of the cusp is visible, however, corresponding to a contour-ending [*Koenderink82*, *Koenderink84*].

2.4 Choice of Parameterisation

There is no unique spatio-temporal parameterisation of the surface. The mapping between extremal boundaries at successive instants is undetermined. The problem of choosing a parameterisation is an "aperture problem" for contours on the spherical perspective image ($\mathbf{T}(s,t)$) or on the Gauss sphere ($\mathbf{n}(s,t)$), or between space curves on the surface $\mathbf{r}(s,t)$). A natural parameterisation is the "epipolar parameterisation" defined by

$$\mathbf{r}_t \wedge \mathbf{T} = 0. \tag{6}$$

For this parameterisation the tangent-plane basis vectors $\mathbf{r}_s$ and $\mathbf{r}_t$ are in *conjugate* directions (from (3)).

Differentiating (1) with respect to time and enforcing (6) leads to the "matching" condition[1]

$$\mathbf{T}_t = \frac{(\mathbf{v}_t \wedge \mathbf{T}) \wedge \mathbf{T}}{\lambda}. \tag{7}$$

Points on different contours are "matched" by moving along great-circles on the image sphere with poles defined by the direction of the viewer's instantaneous translational velocity $\mathbf{v}_t$. This induces a natural *correspondence* on the surface between extremal boundaries from different viewpoints. If the motion is linear corresponding points on the image sphere will lie on an epipolar great-circle. This is equivalent to Epipolar Plane matching in stereo. For a general motion, however, the epipolar structure rotates continuously as the direction of $\mathbf{v}_t$ changes and the space curve, $\mathbf{r}(s_0, t)$, generated by the movement of a contact point will be non-planar.

The parameterisation will be degenerate when $\{\mathbf{r}_s, \mathbf{r}_t\}$ fails to form a basis for the tangent plane. This occurs if the contour is not an extremal boundary but a 3D rigid space curve (when $\mathbf{r}_t = 0$) or at a cusp/contour-ending in the projection (when $\mathbf{r}_s \wedge \mathbf{r}_t = 0$, see earlier)[*Blake89*]. The parameterisation degrades gracefully and hence these conditions pose no special problems.

2.5 Information available from the deformation of the apparent contour

We show in [*Blake89*] that local surface geometry can be recovered from spatio-temporal derivatives (up to 2nd order) of image measurable quantities and known viewer motion. We summarise the most important results below.

1. **Recovery of depth**
   Depth (distance along the ray, λ) can be computed from the rate of deformation ($\mathbf{T}_t$) of the apparent contour under known viewer motion (translational velocity $\mathbf{v}_t$)[*Bolles87*]. The depth is given by
$$\lambda = -\frac{\mathbf{v}_t.\mathbf{n}}{\mathbf{T}_t.\mathbf{n}} \tag{8}$$

[1] If we choose the reference frame to be the instantaneous camera co-ordinate system we can express $\mathbf{T}$ and $\mathbf{T}_t$ in terms of an spherical image position vector, $\mathbf{Q}$, and image velocities (optic flow) $\mathbf{Q}_t$. Namely $\mathbf{T} = \mathbf{Q}$ and $\mathbf{T}_t = \mathbf{Q}_t + \omega \wedge \mathbf{Q}$. Equation (7) reduces to the equation of motion and structure from motion from optic flow [*Maybank85*].

This formula is an infinitesimal analogue of triangulation with stereo cameras. The numerator is analogous to baseline and the denominator to disparity. The result also holds for a rigid space curve or an occluding boundary. It is independent of choice of parameterisation [*Blake89*].

2. **Local surface curvature**
The normal curvature at P in the direction of the ray $\mathbf{T}$, κ^t (which is the same as the normal curvature of the space curve, $\mathbf{r}(s_0, t)$) can be computed from the rate of deformation ($\mathbf{T}_t$) of the apparent contour under viewer motion, and its temporal derivative. This requires knowledge of viewer motion (translational and rotational velocity and acceleration)

$$\kappa^t = \frac{(\mathbf{T}_t.\mathbf{n})^3}{(\mathbf{T}_{tt}.\mathbf{n})(\mathbf{v}_t.\mathbf{n}) + 2(\mathbf{T}.\mathbf{v}_t)(\mathbf{T}_t.\mathbf{n})^2 - (\mathbf{v}_{tt}.\mathbf{n})(\mathbf{T}_t.\mathbf{n})} \tag{9}$$

The sign and magnitude of Gaussian curvature can then be computed from the product of the normal curvature κ^t, and the curvature of the apparent contour, κ^p,(measured by (4)) scaled by inverse-depth [*Koenderink84*]

$$\kappa_{gauss} = \frac{\kappa^p \kappa^t}{\lambda} \tag{10}$$

3 Experimental Results: Determining curvatures from absolute measurements

Figure 2 shows 3 views from a sequence of a scene taken from a camera mounted on a moving robot-arm whose motion has been accurately calibrated from visual data [*Tsai87*]. Using a numerical method for estimating surface curvatures from 3 discrete views (see [*Blake89*]) we can estimate the radius of curvature of the *normal section* R (where $\kappa^t = 1/R$) for a point on an extremal boundary of a cup, B. The method is repeated for a point which is not on an extremal boundary but is a surface marking, A. This is a degenerate case of the parameterisation. A surface marking can be considered as the limiting case of a point with infinite curvature and hence ideally will have zero "radius of curvature". If the measurements are error-ridden and the motion is not known accurately, however, surface markings will appear as extremal boundaries on surfaces with high curvature.

	measured (mm)	actual (mm)	error (mm)
surface marking A	1.95	0.0	1.95
extremal boundary B	45.7	44.4	1.3

Table 1. Radius of curvature (of *normal section* defined by the ray direction) estimated from 3 distinct views of a point on a surface marking and a point on an extremal boundary.

The radius of the cup was measured using calipers as 44.4 ± 2mm. The estimated curvatures agree with the actual curvatures. However, the results are very sensitive to perturbations in the motion parameters (figure 3a and 3b).

4 Differential measurement of curvature

We have seen that although it is perfectly feasible to compute curvature from the observed deformation of an apparent contour, the result is highly sensitive to motion calibration errors. This may be acceptable for a moving camera mounted on a precision robot-arm or when a grid is in view so that accurate visual calibration can be performed. In such cases it is feasible to determine motion to the accuracy of around 1 part in 1000 that is required. However, when only crude estimates of motion are available another strategy is called for. It is sometimes possible in such a case to use the crude estimate to bootstrap a more precise visual egomotion computation [*Harris87*]. However this requires an adequate number of identifiable corner features, which may not be available in an unstructured environment. Moreover, if the estimate is too crude the egomotion computation may fail; it is notoriously ill-conditioned [*Tsai84*].

The alternative approach is to seek qualitative measurements of geometry that are much less sensitive to error in the motion estimate. In this section we show that relative or *differential* measurements of curvature have just this property. Differences of measurements at two points are insensitive to errors in rotation and in translational acceleration. Typically, the two features might be one point on an extremal boundary and one fixed surface point. The surface point has infinite curvature and therefore acts simply as a stable reference point for the measurement of curvature at the extremal boundary. Intuitively the reason for the insensitivity of differential curvature is that global additive errors in motion measurement are cancelled out.

Consider two visual features whose projections on the image sphere are $\mathbf{T}(s_i, t)$, $i = 1, 2$ which we will abbreviate to $\mathbf{T}^i$, $i = 1, 2$. Think of them as two points on extremal boundaries, which trace out curves with (normal) curvatures κ^{t1} and κ^{t2} as the viewer moves. The first temporal derivatives of $\mathbf{T}^i$ are dependent only on image position, viewer velocity and rotation and depth:

$$\mathbf{T}^i_t = \frac{(\mathbf{v}_t \wedge \mathbf{T}^i) \wedge \mathbf{T}^i}{\lambda} \tag{11}$$

Second order temporal derivatives are,

$$\mathbf{T}^i_{tt}.\mathbf{n} = \frac{1}{\lambda}\left[-\frac{(\mathbf{T}^i_t.\mathbf{n})^2}{\kappa^{ti}} + 2(\mathbf{T}^i.\mathbf{v}_t)(\mathbf{T}^i_t.\mathbf{n}) - \mathbf{v}_{tt}.\mathbf{n}\right] \quad i = 1, 2. \tag{12}$$

Let us define two relative quantities. The *differential curvature* $\Delta\kappa^t$ of the feature pair is defined by

$$\frac{1}{\Delta\kappa^t} = \frac{1}{\kappa^{t1}} - \frac{1}{\kappa^{t2}}. \tag{13}$$

Note that it is not an infinitesimal quantity but a difference of inverse curvature. The *relative view vector* is defined to be

$$\delta(t) = \mathbf{T}(s_2, t) - \mathbf{T}(s_1, t) \tag{14}$$

Consider the two features to be instantaneously spatially coincident, that is, initially, $\mathbf{T}(s_1, t) = \mathbf{T}(s_2, t)$. Moreover assume they lie at a common depth λ, and hence, instantaneously, $\mathbf{T}^1_t = \mathbf{T}^2_t$. In practice, of course, the feature pair will only coincide exactly if one of the points is a surface marking which is instantaneously on the extremal boundary. Now, taking the difference of equation (12) for $i = 1,2$ leads to a relation between these two *differential* quantities:

$$\delta_{tt}.\mathbf{n} = \frac{(\mathbf{v}_t.\mathbf{n})^2}{\lambda^3}\frac{1}{\Delta\kappa^t}. \tag{15}$$

From this equation we can obtain *differential curvature* $\Delta\kappa^t$ as a function of depth λ, viewer velocity $\mathbf{v}_t$, and the 2nd temporal derivative of δ. Absolute measurement of curvature (12) depended also on the viewer's translational acceleration $\mathbf{v}_{tt}$. Uncertainty from practical measurements (based on finite differences, for example) of the lower derivative will, of course, be much reduced. Hence the relative measurement should be much more robust. Moreover, because δ is a relative measurement on the projection sphere, unlike the image vectors $\mathbf{T}^i$ which occur in the absolute measurement of curvature, it is unaffected by errors in viewer rotation.

In the case that $\mathbf{T}^1$ is known to be a fixed surface reference point, with $1/\kappa^{t1} = 0$, then $\Delta\kappa^t = \kappa^{t2}$ so that the differential curvature $\Delta\kappa^t$ constitutes an estimate, now much more robust, of the normal curvature κ^{t2} at the extremal boundary point $\mathbf{T}^2$. Of course this can now be used in equation(10) to obtain a robust estimate of gaussian curvature.

Our experiments confirm this. Figures 4a and 4b show that the sensitivity of the differential curvature to error in position and rotation computed between points A and B (2 nearby points at similar depths) is reduced by an order of magnitude. This is a striking decrease in sensitivity even though the features do not coincide exactly as the theory required.

Further robustness can be obtained by considering the ratio of differential curvatures. Ratios of two-point differential curvature measurements are, in theory, completely insensitive to viewer motion [*Blake89*]. This is because in a *ratio* of $\Delta\kappa^t$ measurements for two *different* point-pairs, terms depending on absolute depth λ and velocity $\mathbf{v}_t$ are cancelled out in equation (15). This result corresponds to the following intuitive idea. The rate at which surface features rush towards or away from an extremal boundary is proportional to the (normal) curvature there. The constant of proportionality is some function of viewer-motion and depth; it can be eliminated by considering only ratios of curvatures. Results (figures 5a and 5b) show another striking decrease in sensitivity – of another order of magnitude.

5 Conclusion

We conclude that in theory and practice given just one surface reference point, highly robust relative curvature measurements can be made at points on apparent contours. Moreover, the technique can achieve by motion analysis something which has so far eluded photometric analysis: namely discrimination between fixed surface features and points on extremal boundaries.

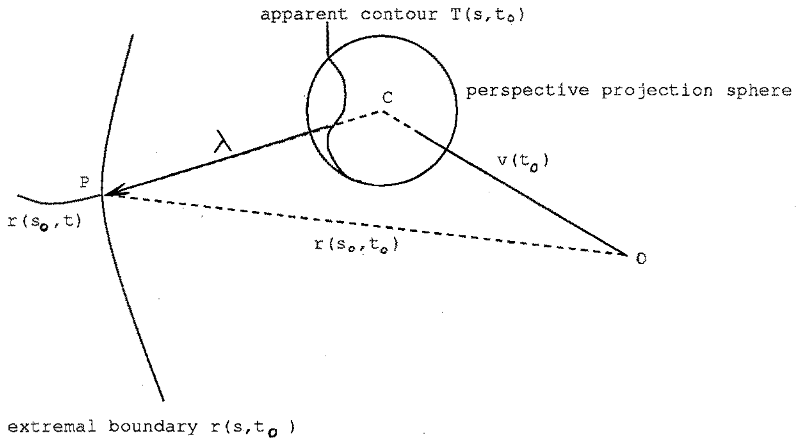

Figure 1. Surface and Viewing Geometry.

P lies on a smooth surface which is parameterised by $\mathbf{r}(s,t)$. For a given vantage point, $\mathbf{v}(t_0)$, the family of rays emanating from the viewer's optical centre (C) that touch the surface defines an s-parameter curve $\mathbf{r}(s,t_0)$ – the extremal boundary from vantage point t_0. The spherical perspective projection of this extremal boundary – the apparent contour, $\mathbf{T}(s,t_0)$ – determines the direction of rays which grazes the surface. The distance along each ray, CP, is λ. A moving observer at position $\mathbf{v}(t)$ sees a 2 parameter family of extremal boundaries $\mathbf{r}(s,t)$ whose spherical perspective projections are represented by a 2 parameter family of apparent contours $\mathbf{T}(s,t)$. t-parameter curves ($\mathbf{r}(s_0,t)$ and $\mathbf{T}(s_0,t)$) are not uniquely defined.

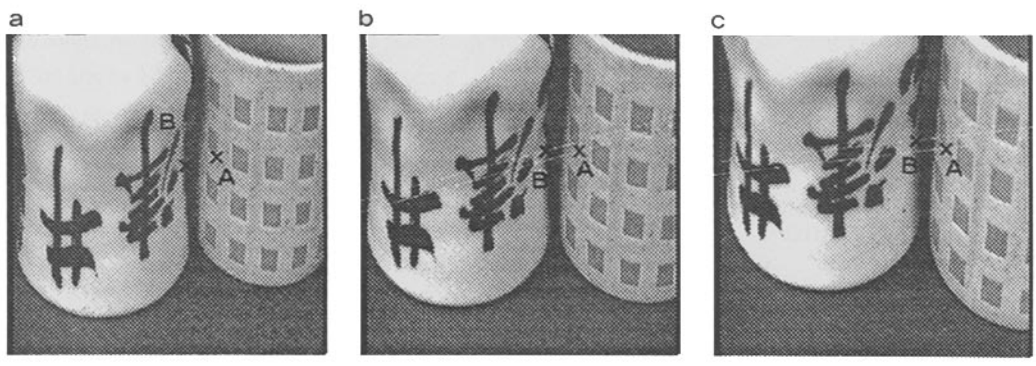

Figure 2. Estimating surface curvatures from 3 discrete views

Points are selected on image contours in view 1 (a), indicated by crosses A and B for points on a surface marking and extremal boundary respectively. For *Epipolar parameterisation* of the surface corresponding features lie on epipolar lines in views 2 and 3 (figures 2b and 2c). Measurement of 3 view vectors lying in an epipolar plane can be used to estimate surface curvatures.

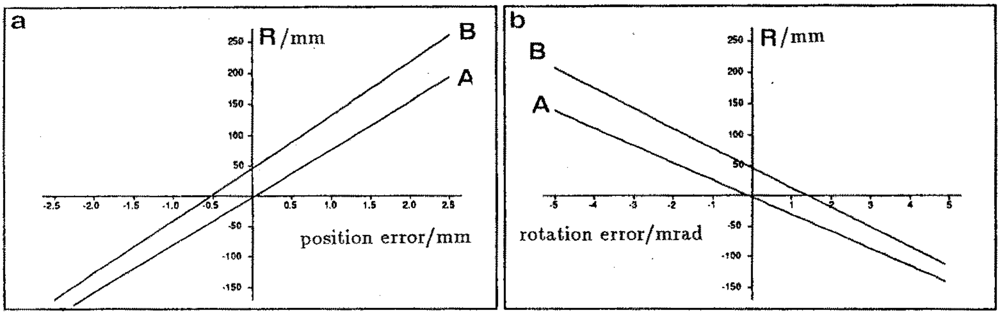

Figure 3. Sensitivity of curvature estimated from absolute measurements to errors in motion.
The radius of curvature (mm) for both a point on a surface marking (A) and a point on an extremal boundary (B) is plotted against error in the estimate of position (a) and orientation (b) of the camera for view 2. The estimation is very sensitive and a perturbation of 1mm in position produces an error of 190% in the estimated radius of curvature for the point on the extremal boundary. A perturbation of 1mrad in rotation about an axis defined by the epipolar plane produces an error of 70%.

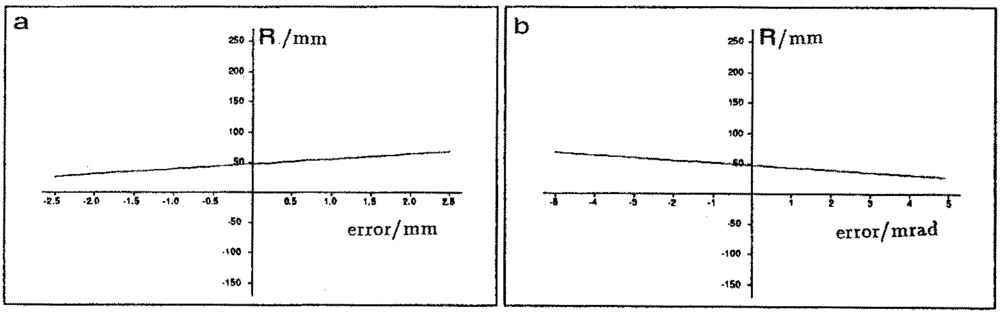

Figure 4. Sensitivity of *differential curvature*
The difference in radii of curvature between a point on the extremal boundary and the nearby surface marking is plotted against error in the position (a) and orientation (b) of the camera for view 2. The sensitivity is reduced by an order of magnitude to 17% per mm error and 8% per mrad error respectively.

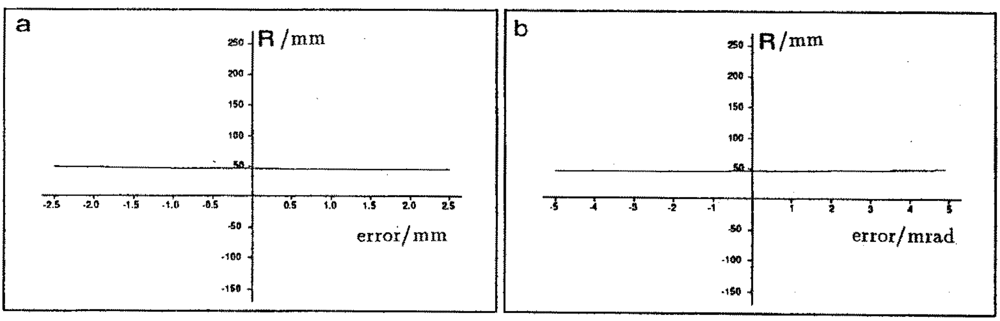

Figure 5. Sensitivity of ratio of *differential curvatures*
The ratio of *differential curvatures* measurements made between 2 points on an extremal boundary and the same nearby surface marking is plotted against error in the position (a) and orientation (b) of the camera for view 2. The sensitivity is further reduced by an order of magnitude to 1.5% error for a 1mm error in position and 1.1% error for 1mrad error in rotation. The vertical axes are scaled by the actual curvature for comparision with figures 3 and 4.

References

[*Barrow78*] H.G. Barrow and J.M. Tenenbaum. *Recovering Intrinsic Scene Characteristics from Images.* A.I Center Technical Report 157, SRI International, 1978.

[*Blake88*] A. Blake and G. Brelstaff. Geometry from specularities. In *Proc. 2nd Int. Conf. on Computer Vision*, pages 394–403, 1988.

[*Blake89*] A. Blake and R. Cipolla. *Robust Estimation of Surface Curvature from Deformation of Apparent Contours.* Technical Report OUEL 1787/89, University of Oxford, 1989.

[*Bolles87*] R.C. Bolles, H.H. Baker, and D.H. Marimont. Epipolar-plane image analysis: an approach to determining structure. *International Journal of Computer Vision*, vol.1:7–55, 1987.

[*Brady85*] M. Brady, J. Ponce, A. Yuille, and H Asada. Describing surfaces. *Computer Graphics Image Processing*, 32:1–28, 1985.

[*Giblin87*] P. Giblin and R. Weiss. Reconstruction of surfaces from profiles. In *Proc. 1st Int. Conf. on Computer Vision*, pages 136–144, London, 1987.

[*Harris87*] C.G. Harris. Determination of ego - motion from matched points. In *3rd Alvey Vision Conference*, pages 189–192, 1987.

[*Koenderink82*] J.J. Koenderink and A.J. Van Doorn. The shape of smooth objects and the way contours end. *Perception*, 11:129–137, 1982.

[*Koenderink84*] J.J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.

[*LHiggins80*] H.C. Longuet-Higgins and K. Pradzny. The interpretation of a moving retinal image. *Proc.R.Soc.Lond.*, B208:385–397, 1980.

[*Maybank85*] S.J. Maybank. The angular velocity associated with the optical flow field arising from motion through a rigid environment. *Proc. Royal Society, London*, A401:317–326, 1985.

[*Tsai84*] R.Y. Tsai and T.S. Huang. Uniqueness and estimation of three-dimensional motion parameters of a rigid objects with curved surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6(1):13–26, 1984.

[*Tsai87*] R.Y. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, 1987.

[*Weinshall89*] D. Weinshall. *Direct computation of 3D shape and motion invariants.* AI Memo 1131, MIT, 1989.

Acknowledgments

The authors acknowledge discussions with Professor Mike Brady, Dr Andrew Zisserman, Dr Peter Giblin, Dr David Murray and Dr David Forsyth. Roberto Cipolla acknowledges the support of the IBM UK Scientific Centre.