The Visual Motion of Curves and Surfaces

By Roberto Cipolla

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

For smooth curved surfaces the dominant image feature is the *apparent contour* or outline. This is the projection of the *contour generator* – the locus of points on the surface which separate visible and occluded parts. The contour generator is dependent of the local surface geometry and the viewpoint. Each viewpoint will generate a different contour generator. This paper addresses the problem of recovering the 3D shape and motion of curves and surfaces from image sequences of apparent contours.

For *known* viewer motion the visible surfaces can then be reconstructed by exploiting a spatio-temporal parametrization of the apparent contours and contour generators under viewer motion. A natural parametrization exploits the contour generators and the epipolar geometry between successive viewpoints. The *epipolar parametrization* (Cipolla & Blake 1992) leads to simplified expressions for the recovery of depth and surface curvatures from image velocities and accelerations and known viewer motion.

The parametrization is, however, degenerate when the apparent contour is singular since the ray is tangent to the contour generator (Koenderink & Van Doorn 1976) and at *frontier points* (Giblin & Weiss 1994) when the epipolar plane is a tangent plane to the surface. At these isolated points the epipolar parametrization can no longer be used to recover the local surface geometry. This paper reviews the epipolar parametrization and shows how the degenerate cases can be used to recover surface geometry and *unknown* viewer motion from apparent contours of curved surfaces. Practical implementations are outlined.

1. Introduction

Structure and motion from image sequences of point features has attracted considerable attention and a large number of algorithms exist to recover both the spatial configuration of the points and the motion compatible with the views. A key component of these algorithms is the recovery of the *epipolar geometry* between distinct views (Luong & Faugeras 1996). The structure and motion problem for curves and curved surfaces is more challenging. For curved surfaces the dominant image feature is the *apparent contour* which is the projection of the curve on the surface (*contour generator*) dividing visible and occluded parts. The contour generator is dependent on viewpoint and local surface geometry (via tangency and conjugacy constraints) and each viewpoint will generate a different contour generator. The image curves are therefore projections of different space curves and there is no correspondence between points on the curves in the two images.

The family of contour generators generated under continuous viewer motion

Phil. Trans. R. Soc. Lond. A (1996) **(Submitted)** 1996 Royal Society Typescript
Printed in Great Britain TEX Paper

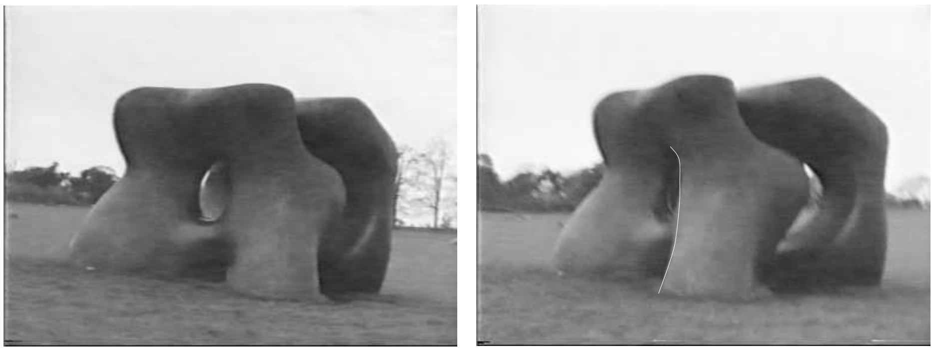

Figure 1. The apparent contours of a curved surface from two different viewpoints. The apparent contours from different viewpoints are projections of different surface curves. The apparent contour can be singular, seen here as the visible apparent contour ending abruptly.

can be used to represent the visible surface. Giblin & Weiss (1987) and Cipolla & Blake (1992) have shown how the spatio-temporal analysis of deforming image apparent contours or outlines enables computation of local surface curvature along the corresponding contour generator on the surface. To perform the analysis, however, a spatio-temporal parametrization of image-curve motion is needed, but is under-constrained . The *epipolar* parametrization is most naturally matched to the recovery of surface curvature. In this parametrization (for both the spatio–temporal image and the surface), *correspondence* between points on successive snapshots of an apparent contour and contour generator is set up by matching along epipolar lines and epipolar planes respectively. The parametrization leads to simplified expressions for the recovery of depth and surface curvature from image velocities and accelerations and known viewer motion.

There are however several cases in which this parametrization is degenerate and so can not be used to recover the local surface geometry. The first case of degeneracy occurs at a point of the surface-to-image mapping when a ray is tangent not only to the surface but also to the contour generator. This will occur when viewing a hyperbolic surface patch along an asymptotic direction. For a transparent surface this special point on the contour generator will appear as a *cusp* on the apparent contour. For opaque surfaces, however, only one branch of the cusp is visible and the contour ends abruptly (Koenderink & Van Doorn 1982).

The other case of degeneracy of the epipolar parametrization occurs when contour generators from subsequent viewpoints intersect to form an envelope. This occurs when the epipolar plane is also a tangent plane to the surface. These isolated surface points are called *frontier points* (Giblin & Weiss 1994). The surface can not be reconstructed at these points by the epipolar parametrization since the contour generator is locally stationary. However the frontier points correspond to real, fixed feature points on the surface which are visible in two views. They can be used to recover the viewer motion.

This paper addresses the problem of recovering the 3D shape and motion of curves and surfaces from image sequences of apparent contours. As with point features, the epipolar geometry plays an important role in both the recovery of the motion and in the reconstruction of the surface.

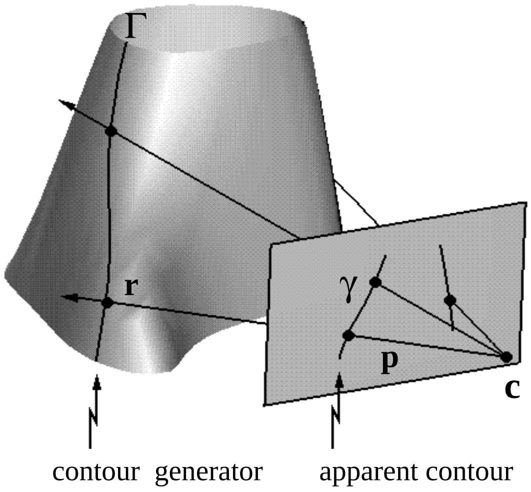

Figure 2. Viewing geometry and parametrization of the surface. For each viewpoint, **c**, the family of rays which are tangent to the surface define the contour generator, **Γ**. The image of the contour generator is called the apparent contour, **ﬂ**. A surface point, **r**, has position $\mathbf{c} + \lambda\mathbf{p}$ where λ is the distance from the viewing centre along a ray with direction given by the unit vector **p**.

2. Reconstruction under known viewer motion

(*a*) *Viewing geometry*

Consider a smooth surface M. For each vantage point, **c**, the sets of points, **r**, on the surface for which the visual ray is tangent to M can be defined. This is called the *contour generator*, Γ, and is the set of points **r** for which

$$(\mathbf{r} - \mathbf{c}).\mathbf{n} = 0 \tag{2.1}$$

where **n** is the unit normal to the surface at **r** (figure 2). The contour generator is usually (but not always, see section 3) a smooth curve on the surface separating the visible from the occluded parts and can be parametrized using say s as a parameter.

The image, **ﬂ**, of the contour generator, **Γ**, is called the apparent contour or outline and is the intersection of the set of rays which are tangent to the surface and the imaging surface. Without loss of generality and for mathematical simplicity, we consider perspective projection onto the unit sphere. An apparent contour point, **p**, (a unit vector specifying the direction of the visual ray) satisfies:

$$\mathbf{r} = \mathbf{c} + \lambda\mathbf{p} \tag{2.2}$$

$$\mathbf{p}.\mathbf{n} = 0 \tag{2.3}$$

where λ is the distance along the ray to the surface point, **r**, from the position of the centre of projection **c**.

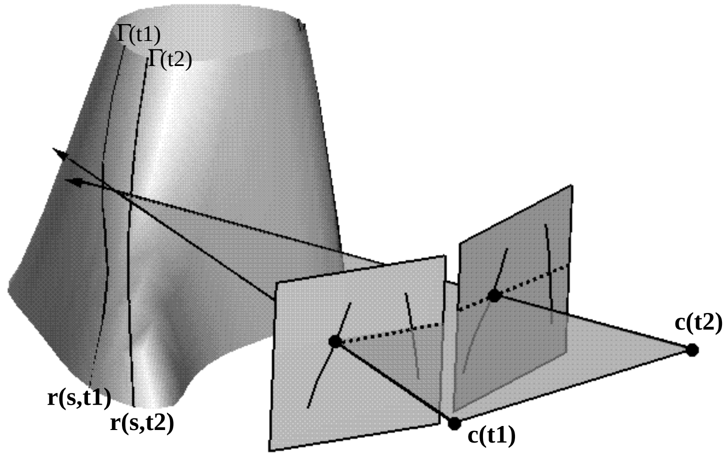

Figure 3. Epipolar parametrization. The surface is parametrized locally by the contour generators from successive viewpoints and *epipolar* curves defined by the intersection of the pencil of epipolar planes and the surface. The parametrization is only possible if the contour generators are not singular and are transverse to the epipolar curves. The parametrization is degenerate at cusp and frontier points.

(*b*) *Spatio-temporal parametrization*

Each viewpoint will generate a different contour generators on the surface M. A moving monocular observer with position at time t given by $\mathbf{c}(t)$, will generate a one parameter family of contour generators, indexed by time, $\Gamma(t)$. It is natural to attempt a parametrization of M which is 'compatible' with the motion of the camera centre, in the sense that contour generators are parameter curves. We want there to exist a regular (local) parametrization of M of the form $(s, t) \rightarrow \mathbf{r}(s, t)$, where the set of points $\mathbf{r}(s, t_0)$, for fixed t_0, (i.e the s-parameter curve) is the contour generator from viewpoint $\mathbf{c}(t_0)$. The set of points $\mathbf{p}(s, t_0)$ is the corresponding apparent contour on the unit sphere at the viewpoint; the actual apparent contour points in space are $\mathbf{c}(t_0) + \mathbf{p}(s, t_0)$. Note that (2.2) and (2.3) become

$$\mathbf{r}(s, t) = \mathbf{c}(t) + \lambda(s, t)\mathbf{p}(s, t) \tag{2.4}$$

$$\mathbf{p}(s, t).\mathbf{n}(s, t) = 0. \tag{2.5}$$

However the spatio-temporal parametrization of the apparent contours, $\mathbf{p}(s, t)$, and the surface, $\mathbf{r}(s, t)$, is not unique. The choice of the t-parameter curves, $\mathbf{p}(s_0, t)$ and $\mathbf{r}(s_0, t)$, for fixed s_0, is under-constrained.

(*c*) *Epipolar parametrization*

A natural choice of parametrization is the *epipolar* parametrization. In this parametrization the *correspondence* between points on successive snapshots of an apparent contour and contour generator are set up by matching along epipolar planes. Namely the corresponding ray in the next viewpoint (in an infinitesimal sense), is chosen so that it lies in the epipolar plane defined by the viewer's translational motion and the ray in the first viewpoint (figure 3).

The epipolar parametrization is defined by (Cipolla & Blake 1992):

$$\mathbf{r}_t \wedge \mathbf{p} = 0 \tag{2.6}$$

and leads to the following *epipolar matching* condition

$$[\mathbf{p}_t, \mathbf{c}_t, \mathbf{p}] = 0 \tag{2.7}$$

such that the t-parameter curves are defined to lie instantaneously in the epipolar plane defined by the ray and direction of translation. The parametrization leads to simplified expressions for the recovery of depth and surface curvature. By simple manipulation of equations (2.4) to (2.7) and their spatial and temporal derivatives (denoted below by subscripts s and t) it is easy to show that the local surface geometry can be recovered from spatio-temporal derivatives (up to second order) of the apparent contours and the *known* viewer motion as follows:

(i) The orientation of the surface normal, $\mathbf{n}$, can be recovered from a single view from the vector product of the ray direction, $\mathbf{p}$ and the tangent to the apparent contour, $\mathbf{p}_s$:

$$\mathbf{n} = \frac{\mathbf{p} \wedge \mathbf{p}_s}{|\mathbf{p} \wedge \mathbf{p}_s|}. \tag{2.8}$$

(ii) If the contour generator is smooth at $\mathbf{r}$ then its tangent is in a *conjugate* direction (with respect to the second fundamental form) to the visual ray $\mathbf{p}$ and:

$$\mathbf{p}.\mathbf{n}_s = 0 \tag{2.9}$$

(iii) Under viewer motion and the epipolar parametrization, a given point on a contour generator, $\mathbf{r}$, will *slip* over the surface with velocity given by $\mathbf{r}_t$ and which depends on the distance, λ, and surface curvature (normal curvature in direction of the ray):

$$\mathbf{r}_t = -\left(\frac{\mathbf{c}_t.\mathbf{n}}{\lambda\kappa^t}\right)\mathbf{p} \tag{2.10}$$

Note that the velocity is inversely proportional to the surface curvature and is zero in the limiting case of viewing a space curve or crease. The latter can be simply treated as apparent contours with infinite curvature along the ray direction.

(iv) Depth (distance along the ray, λ) can be computed from the deformation of the apparent contour, ($\mathbf{p}_t$), under known viewer motion (translational velocity $\mathbf{c}_t$):

$$\lambda = -\frac{\mathbf{c}_t.\mathbf{n}}{\mathbf{p}_t.\mathbf{n}} \tag{2.11}$$

(v) The Gaussian curvature at a point on the apparent contour, K, can be recovered from the depth, λ, the normal curvature κ^t along the line of sight and the *geodesic* curvature of the apparent contour, κ^p:

$$K = \frac{\kappa^p\kappa^t}{\lambda} \tag{2.12}$$

Since the normal section in the direction of the ray must always be convex at a point on the apparent contour, the sign of the Gaussian curvature is determined by the sign of the curvature of the apparent contour. Convexities, concavities and

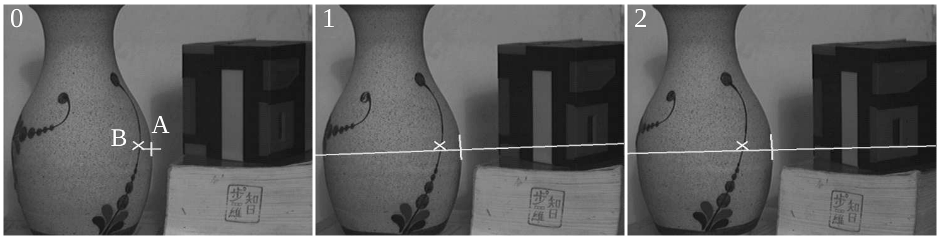

Figure 4. Estimating surface curvatures from three discrete views. Points are selected on image contours in the first view (t_0), indicated by crosses A and B for points on an extremal boundary and surface marking respectively. For the *epipolar parametrization* of the surface corresponding features lie on epipolar lines in the second and third view (t_1 and t_2). Measurement of the three rays lying in an epipolar plane can be used to estimate surface curvatures. Point B can be considered as a degenerate surface point with infinite curvature.

inflections of an apparent contour corresond to elliptic, hyperbolic and parabolic surface points respectively (Koenderink 1984).

Figures 4 – 6 illustrate the epipolar parametrization and the reconstruction of a strip of surface at the contour generator. Details of the camera calibration and the detection and tracking of the apparent contours with B-spline snakes can be found in Cipolla & Blake (1992) and Boyer & Berger (1997).

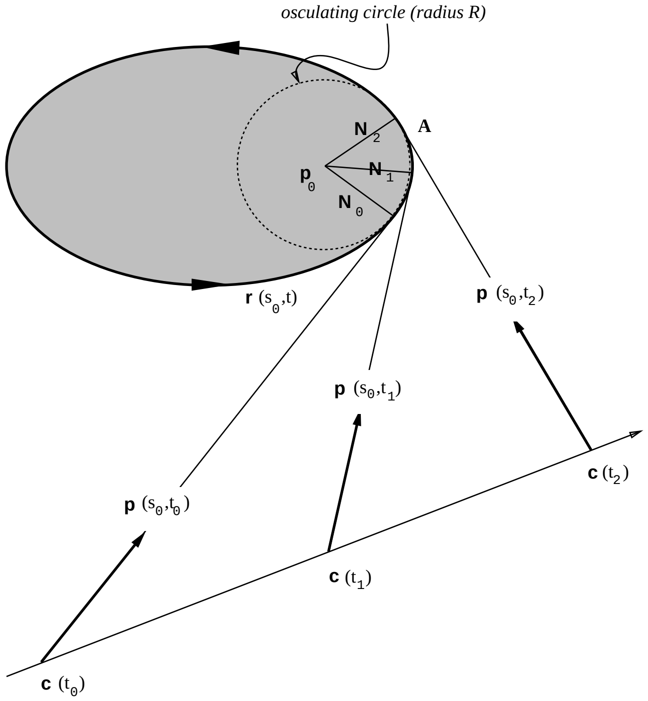

Figure 5. The epipolar plane. Each view defines a tangent to $\mathbf{r}(s_0, t)$. For linear camera motion and epipolar parametrization the rays and $\mathbf{r}(s_0, t)$ lie in a plane. If $\mathbf{r}(s_0, t)$ can be approximated locally by its osculating circle, it can be uniquely determined from measurements in three views. For curvilinear motion the epipolar geometry is continuously changing and the epipolar curve is no longer planar.

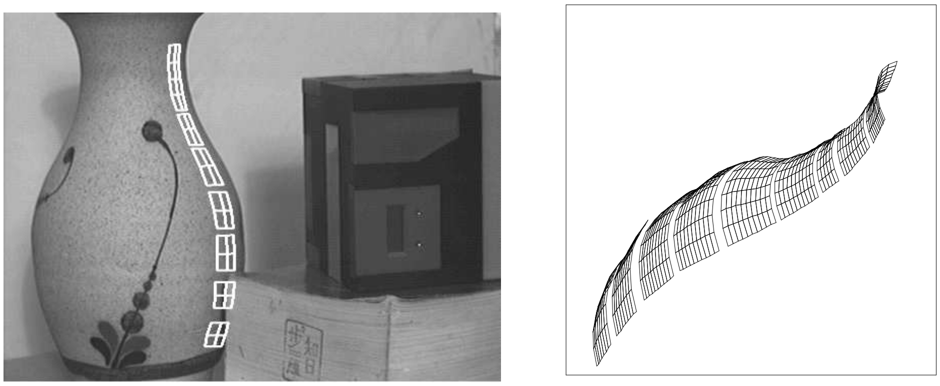

Figure 6. Recovery of surface strip in vicinity of apparent contour. The surface is recovered as a family of s-parameter curves – the contour generators – and t-parameter curves – portions of the *osculating* circles measured in each epipolar plane. The strip is shown projected into the image of the scene from a different viewpoint and after extrapolation (Cipolla & Blake 1992).

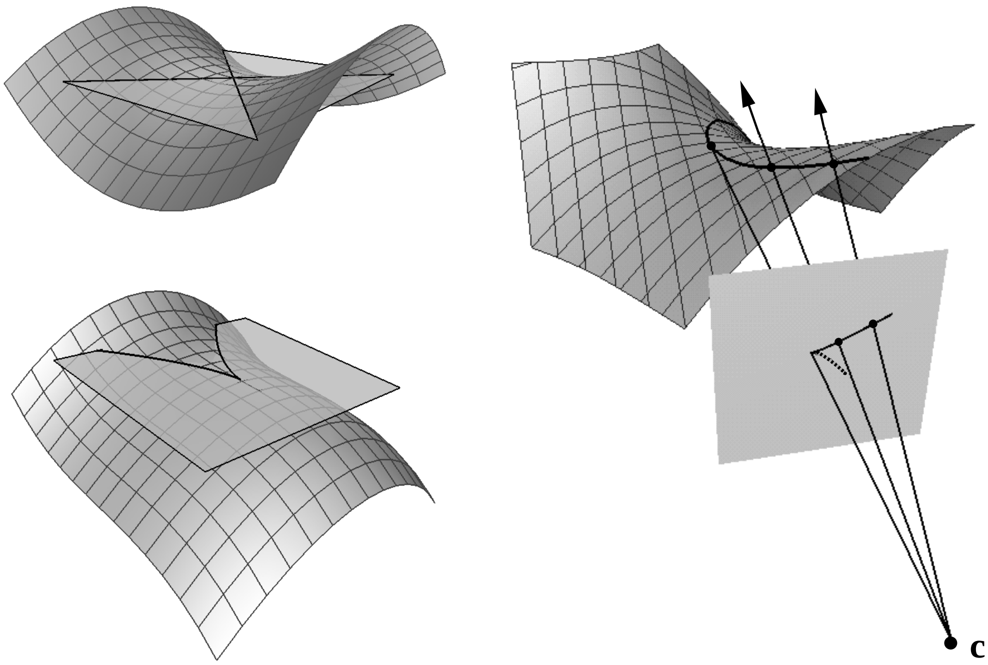

Figure 7. The asymptotic directions of a hyperbolic and parabolic surface point (left). The visual ray has 3-point contact with the surface when viewing a hyperbolic surface point along one of its asymptotic directions. The contour generator and ray are parallel and a cusp is generated in the apparent contour. Only one branch of the cusp is visible for an opaque surface and the apparent contour ends abruptly.

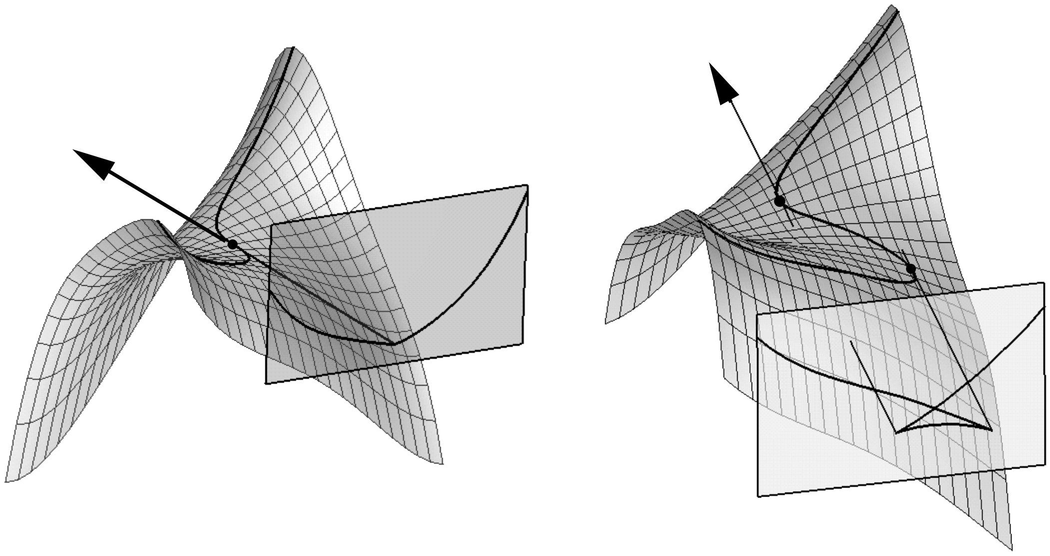

Figure 8. The swallowtail transition. A smooth apparent contour develops a pair of cusps and a T-junction under viewer motion. For opaque surfaces only one branch of one of the cusps is visible and the smooth apparent contour comes to a sudden end. The transition occurs at an asymptotic direction of a *flecnodal surface point*. The ray has 4-point contact with the surface.

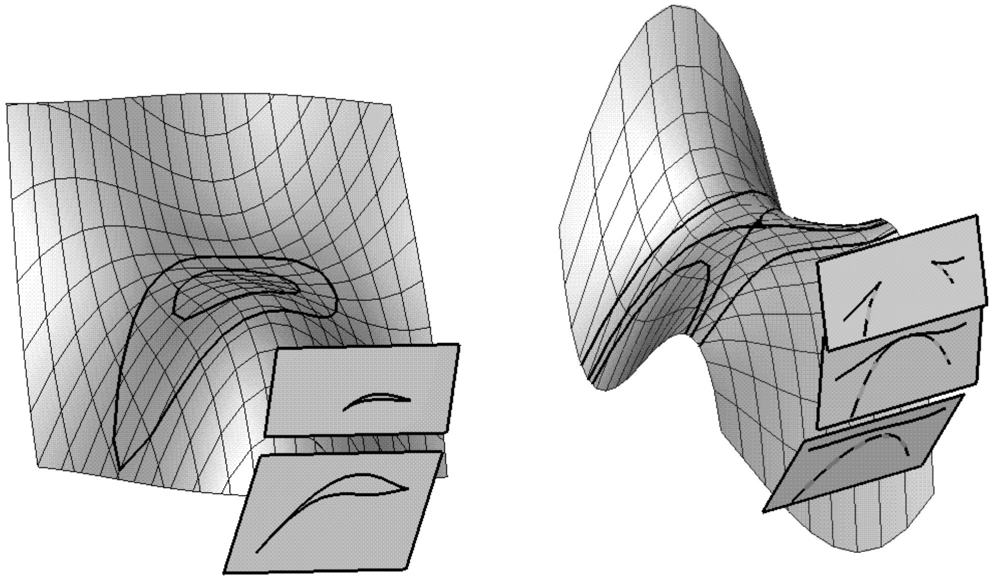

Figure 9. The "lips" and "beaks" visual events. The contour generators are singular (an isolated point in the lips transition; a crossing in the case of the beaks transition) when viewing a parabolic surface point along its asymptotic direction. A pairs of cusps are generated under viewer motion.

3. Degenerate cases of the epipolar parametrization

There are two possible cases where degeneracy of the parametrization arises. These occur when the contour generators and epipolar curves are singular or not transverse, i.e. $\{\mathbf{r}_s, \mathbf{r}_t\}$ fails to form a basis for the tangent plane of the surface.

$$\mathbf{r}_s \wedge \mathbf{r}_t = 0 \tag{3.1}$$

The first case occurs when $\mathbf{r}$ is a hyperbolic surface point viewed along an asymptotic direction. The apparent contour is singular (a cusp is generated), seen as a contour-ending for opaque surfaces. The second case occurs when $\mathbf{r}$ is a *frontier point*. The epipolar plane (spanned by the velocity vector $\mathbf{c}_t(t)$ of the viewpoint and the ray $\mathbf{p}$) coincides with the tangent plane to M and the contour generators form an envelope on M.

(a) *Singular apparent contours*

The most common case of degeneracy occurs when viewing a hyperbolic surface point along an *asymptotic* direction (figure 7). The ray grazes the surface with 3-point contact and is tangent to the contour generator, $\mathbf{r}_s \wedge \mathbf{p} = 0$. For a transparent surface this special point on the contour generator (the cusp generator point) will appear as a *cusp* on the apparent contour with $\mathbf{p}_s = 0$. For opaque surfaces, however, only one branch of the cusp is visible and the apparent contour ends abruptly (Koenderink & Van Doorn 1982).

Cusps (or contour-endings as they appear for opaque surfaces) are stable phenomena and they persist under viewer motion. In any generic view of a curved surface we expect to see smooth apparent contours with a finite number of singularities (cusp or contour-ending) and T-junctions. Under viewer motion the generic events that are visible consist of cusps being created or annihilated in pairs. These are described in (Koenderink & Van Doorn 1976; Koenderink 1990) and illustrated in figures 8 and 9.

The epipolar parametrization can no longer be used to recover the depth and surface curvature at cusp points. In fact the cusp points have a component of motion out of the epipolar plane and so can not be localised on corresponding epipolar lines in the different views. The cusp point in the image can, in principle, be localised and tracked under viewer motion, and if its trajectory is smooth and can be parametrized by t, $\mathbf{p}(t)$, it can be used to induce an alternative parametrization of the surface in the vicinity of the cusp generator (Cipolla, Fletcher & Giblin 1997).

At a cusp, the cuspidal tangent is given by $\mathbf{p}_{ss}$ (since $\mathbf{p}_s = 0$) and the surface normal, depth and Gaussian curvature can be recovered as follows:

$$\mathbf{n} = \frac{\mathbf{p} \wedge \mathbf{p}_{ss}}{|\mathbf{p} \wedge \mathbf{p}_{ss}|} \tag{3.2}$$

$$\lambda = -\frac{\mathbf{c}_t.\mathbf{n}}{\mathbf{p}_t.\mathbf{n}} \tag{3.3}$$

$$K = \frac{-(\mathbf{p}_t.\mathbf{n})^4}{[\mathbf{p}, \mathbf{c}_t, \mathbf{p}_t]^2} \tag{3.4}$$

Note that the Gaussian curvature can be recovered from first-order derivatives

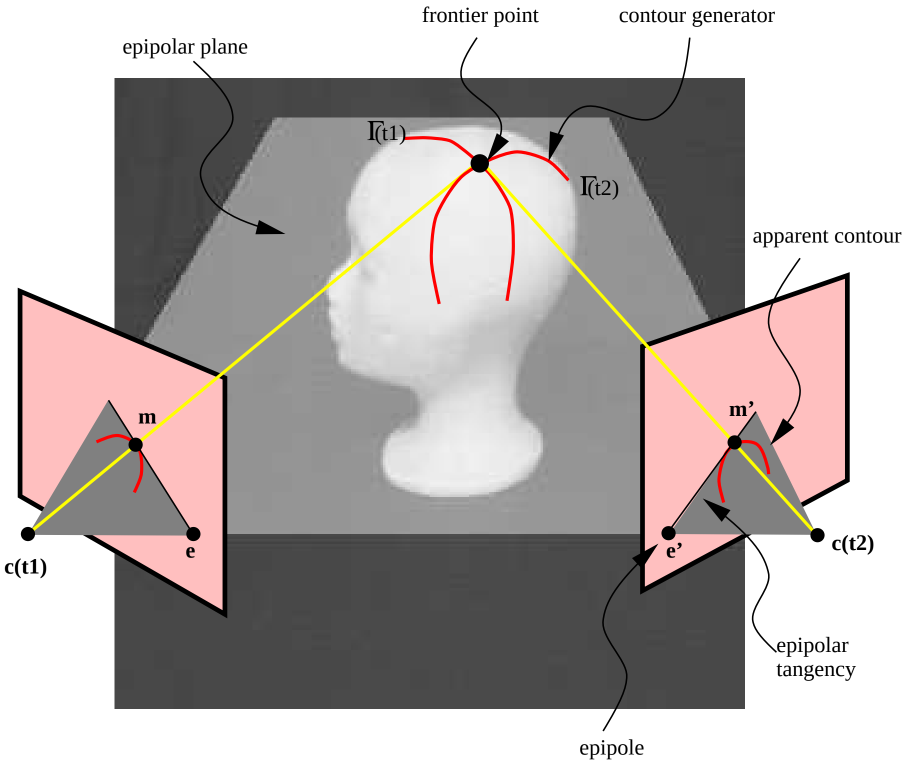

Figure 10. A frontier point appears the intersection of two consecutive contour generators and is visible in both views. The frontier point projects to a point on the apparent contour which is an epipolar tangency point.

only. Compare this to a normal apparent contour point that requires second-order spatio-temporal derivatives of viewpoint and apparent contour.

(*b*) *Frontiers and epipolar tangencies*

The remaining case of degeneracy of the epipolar parametrization occurs for epipolar planes (spanned by the direction of translation and the ray) which coincide with tangent planes to the surface. This will occur at a finite set of points on the surface where the surface normal $\mathbf{n}$ is perpendicular to the direction of translation:

$$\mathbf{c}_t.\mathbf{n} = 0. \tag{3.5}$$

From (2.10) we see that the contour generator is locally stationary ($\mathbf{r}_t = 0$). In fact (see figure 10) consecutive contour generators will intersect at points where the epipolar plane is tangent to the surface.

The points of contact on the surface are called *frontier points* because for continuous motion the locus of intersections of consecutive contour generators in an infinitesimal sense define a curve on the surface which represents the boundary of the visible region swept out by the contour generators under viewer motion.

For larger discrete motions the contour generators defined by the discrete viewpoints also intersect at points on the surface where the epipolar plane is tangent to the surface. This is easily seen if we consider the motion to be linear. $\mathbf{c}_t$ is then a *constant* vector, and the frontier point on the surface at time t satisfies

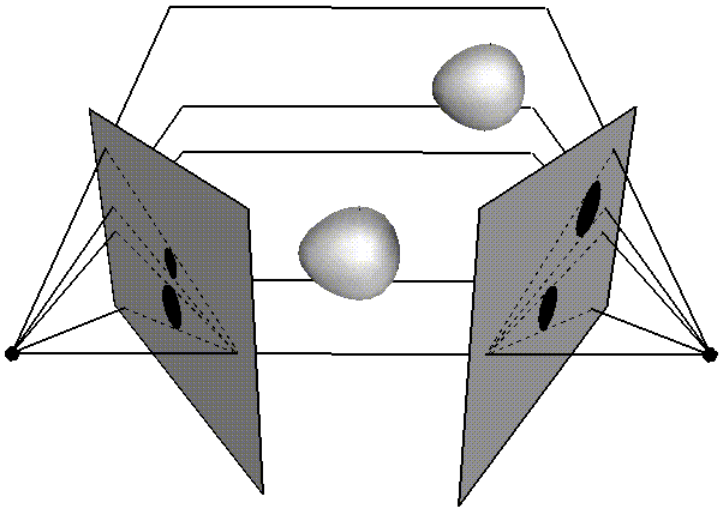

Figure 11. Epipolar geometry and epipolar tangencies under arbitrary motion

the frontier condition at subsequent times. The frontier degenerates to a point on the surface. In the discrete case the frontier points are defined by the condition

$$\Delta \mathbf{c}.\mathbf{n} = 0 \tag{3.6}$$

where $\Delta \mathbf{c} = \mathbf{c}(t_2) - \mathbf{c}(t_1)$ and $\mathbf{n}$ is the surface normal at the point in which in the two contour generators for each viewpoint intersect.

The frontier point projects to a point on the apparent contour in both views such that its tangent passes through the epipole. It is an epipolar tangency point since the tangent plane is also the epipolar plane.

The surface curvature can not be recovered by the epipolar parametrization at these points since the contour generator is locally stationary. However frontier points correspond to real, fixed feature points on the surface which are visible in both views, once detected they can be used to provide a constraint on viewer motion. In fact they can be used in the same way as points in the recovery of the epipolar geometry via the epipolar constraint.

4. Recovery of viewer motion

(*a*) *Parametrization of the fundamental matrix*

The epipolar geometry between two uncalibrated views is completely determined by 7 independent parameters: the position of the epipoles in the two views $(u_e, v_e, 1)^T$ and $(u'_e, v'_e, 1)^T$ and the 3 parameters of the homography relating the pencil of epipolar lines in view 1 to those in view 2

$$\tau' = -\frac{h_2\tau + h_1}{h_4\tau + h3} \tag{4.1}$$

where τ and τ' represent the directions of a pair of corresponding epipolar lines in the first and second images respectively. The transformation is fixed by three pairs of epipolar line correspondences (Luong & Faugeras 1996).

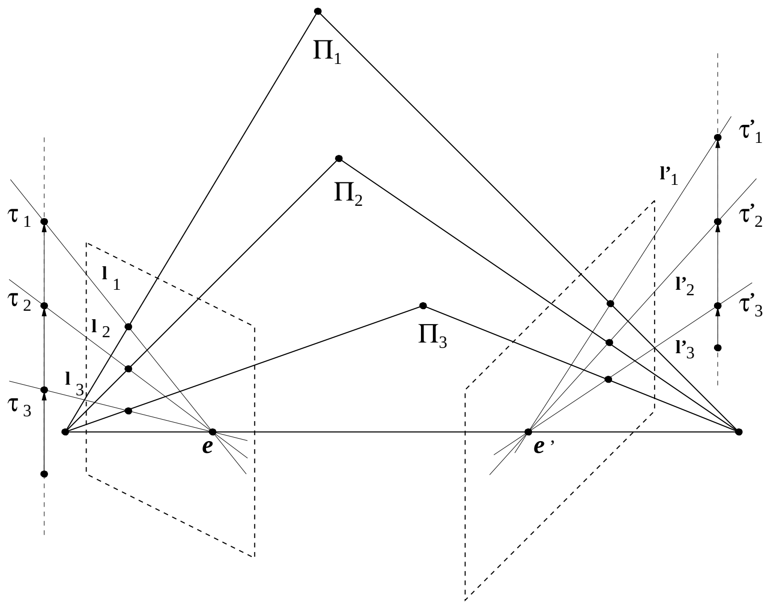

Figure 12. The epipolar geometry of an uncalibrated stereo pair of images is completely specified by the image positions of the epipoles and 3 pairs of corresponding epipolar lines. The projective parameters τ and τ' represent the intersection of the epipolar line and the line at infinity. The directions in the two views are related by a homography.

The epipolar geometry can be conveniently specified by the Fundamental Matrix, F, (3×3 matrix defined up to an arbitrary scale and of rank two) such that the image co-ordinates (projective representation) of a pair of corresponding points, $\mathbf{m}$ and $\mathbf{m}'$, must satisfy the epipolar constraint:

$$\mathbf{m}'^T \mathbf{F} \mathbf{m} = 0 \tag{4.2}$$

and where the left and right epipoles ($\mathbf{e}$ and $\mathbf{e}'$) are given by the null space of F and F^T respectively.

This gives the following parametrization of the fundamental matrix:

$$\mathbf{F} = \begin{pmatrix} h_1 & h_2 & -u_e h_1 - v_e h_2 \\ h_3 & h_4 & -u_e h_3 - v_e h_4 \\ -u'_e h_1 - v' e h_3 & -u'_e h_2 - v' e h_4 & u_e u'_e h_1 + v_e u'_e h_2 + u_e v' e h_3 + v_e v' e h_4 \end{pmatrix}$$

(*b*) *Finding the epipoles*

Under pure translation the epipolar geometry is completely determined by the position of the epipole in a single view. The position of the epipole is the same in both views if the intrinsic parameters do not change and the epipolar lines have the same directions and are *auto-epipolar*. The bitangents at two consecutive apparent contours are epipolar tangencies and hence the projection of frontier points. The intersection of at least two distinct tangencies (epipolar lines) determines the position of the epipole. See figure 13 (Sato & Cipolla 1998).

The solution is no longer trivial in the case of arbitrary motion with rotation. There is in fact no closed form solution since the epipoles are needed to define the epipolar tangency points (and frontier points) and these are needed to determine the epipoles.

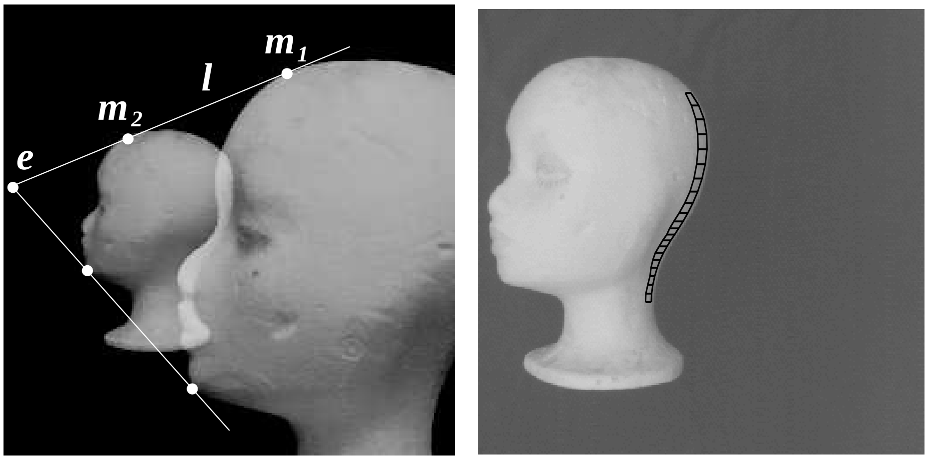

Figure 13. Under pure translation the frontier/epipolar tangency point moves along the epipolar line since the position of the epipole and the direction of the epipolar lines do not change. From a minimum of two bitangents of the apparent contour in two views (a) it is possible to recover the epipole, $\mathbf{e}$, and hence to reconstruct the strip of visible surface at the apparent contour (b).

(*c*) *Optimisation*

The solution proceeds as a search and optimisation problem to find the position of the epipoles in both views such that the epipolar tangencies in the first view are related to the set of epipolar tangencies in the second view by a one-dimensional homography (Cipolla, Åström & Giblin 1995).

A suitable cost function is needed. A geometric criterion (distance) is used in the estimation of the fundamental matrix from point correspondenes and can also be used in the case of curves. The geometric distance is computed as the sum over all tangency points of the square of the distance between the image point and the corresponding epipolar line from the tangency point in the other view.

The key to a successful implementation is to ensure that the search space is reduced and that the optimisation begins from a good starting point using approximate knowledge of the camera motion or point correspondences. The solution proceeds as follows:

(i) Start with an initial guess or estimate of the epipoles in both views.

(ii) Compute the epipolar tangencies, $\mathbf{m}(\mathbf{e})$ and $\mathbf{m}'(\mathbf{e}')$, in both views respectively. These are points on the apparent contours with tangents passing through the epipole.

(iii) Estimate the elements of the homography between the pencil of tangencies in both views. This can be done linearly by minimising

$$\sum_i (h_4 \tau_i \tau_i' + h_3 \tau_i' - h_2 \tau_i - h_1)^2 \tag{4.3}$$

by least squares over all pairs of correspondences (τ and τ').

(iv) The fundamental matrix is now given by the parametrization above and the distance criterion (i.e sum of squared distances between tangency point and corresponding epipolar line) can be computed.

$$C = \sum_i \left(\frac{1}{(\mathbf{Fm}_i)_1^2 + (\mathbf{Fm}_i)_2^2} + \frac{1}{(\mathbf{F}^T\mathbf{m}'_i)_1^2 + (\mathbf{F}^T\mathbf{m}'_i)_2^2} \right) (\mathbf{m}'^T_i \mathbf{Fm}_i)^2$$

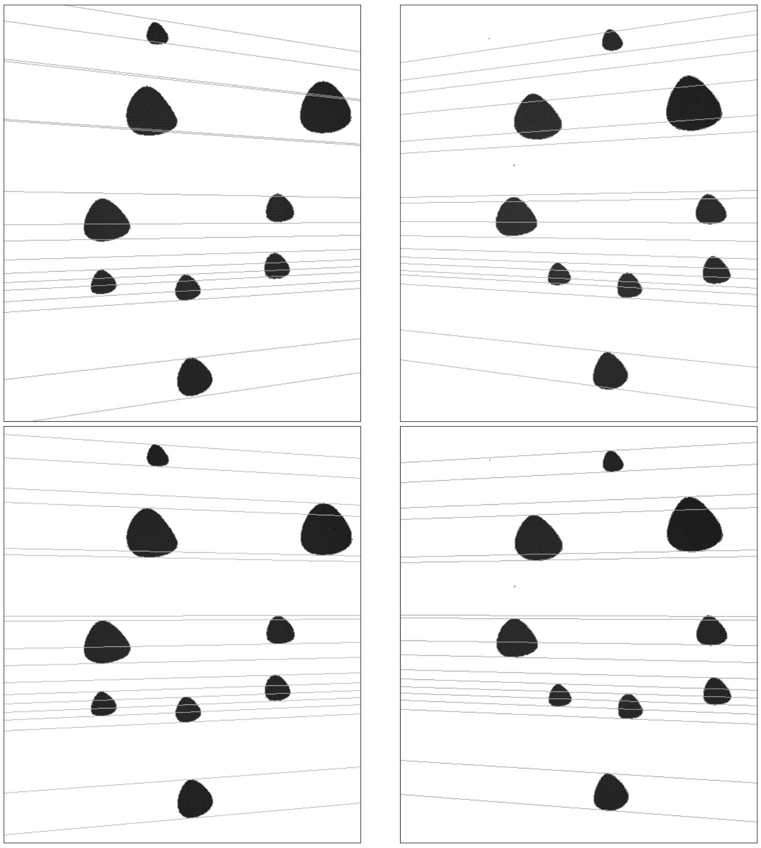

Figure 14. Starting point for optimisation (above). An initial guess of the position of the epipoles is used to determine a finite set of epipolar tangencies in both views and the homography relating the two set of epipolar lines. For each tangency point the corresponding epipolar line is drawn in the other view. The distances between epipolar lines and tangecy points is used to search for the correct positions of the epipoles. Convergence to local minimum after 5 iterations (below). The epipolar lines are tangent to apparent contours in both views.

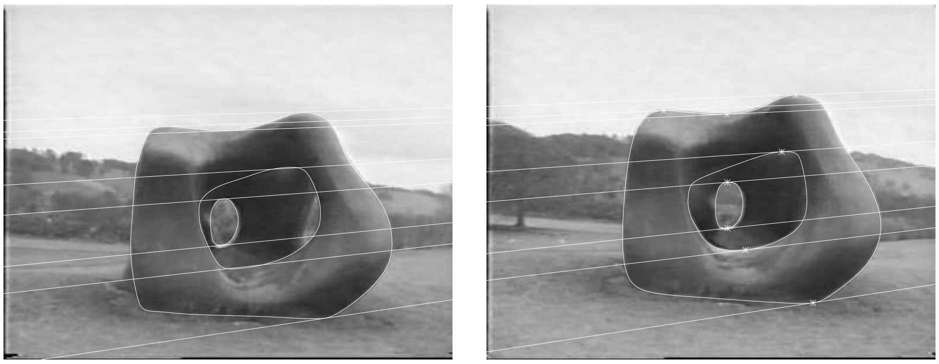

Figure 15. Local minimum obtained by iterative scheme to estimate the epipolar geometry from 8 epipolar tangencies.

(v) Minimise the distance by the conjugate gradient method. The search space is restricted to the four co-ordinates of the epipoles only. This requires the first-order partial derivatives of the cost function with respect to the co-ordinates of the epipoles which can be computed analytically but are more conveniently estimated by numerical techniques.

At each iteration of the algorithm, steps 1 to 4 are repeated, and the positions of the epipoles are refined. The search is stopped when the root-mean-square distance converges to a minimum (usually less than 0.1 pixels). It is of course not guaranteed to find a unique solution.

A number of the experiments were carried out with simulated data (with noise) and known motion (figure 14). The apparent contours were automatically extracted from the sequence by fitting B-splines (Cham & Cipolla 1996). 5-10 iterations each for 4 different initial guesses for the position of the epipole were sufficient to find the correct solution to within an root-mean-square error of 0.1 pixel per tangency point.

Figure 15 shows an example with real data whose apparent contours are detected and automatically tracked using B-splines snakes. A solution is found very quickly which minimises the geometric distances but as with all structure from motion algorithms, a limited field of view and small variation in depths result in a solution which is sensitive to image localisation errors.

5. Conclusions and future work

The recovery of the epipolar geometry between views is a key part of any algorithm to recover the 3D structure and motion compatible with the views. The structure and motion problem for curves and surfaces is more challenging since the apparent contours are viewpoint dependent and the correspondence of points between the two viewpoints in not given.

We have shown how the viewer motion can be recovered from the outlines (apparent contours) of curved surfaces by searching for epipolar tangency points. The results of initial experiments using these algorithms have been promising but the performance of the algorithm remains to be fully evaluated and it is still unclear whether the extraction of the motion leads to a unique solution. After computing the motion the epipolar geometry can be exploited to parameterise the apparent contours and to recover the visible surface.

We have only used apparent contours in the motion estimate. In practice one would use a combination of image features to estimate motion and could then use the apparent contours to reconstruct the surface. An important test of the usefulness of the proposed theories will be the accuracy of the reconstruction of an arbitrarily curved surface from uncalibrated viewer motion.

Roberto Cipolla acknowledges the support of the EPSRC. The research on following cusps and frontiers was carried out at the Isaac Newton Institute for Mathematical Sciences and the Department of Engineering, Cambridge in collaboration with Karl Åström, Andrew Blake, Gordon Fletcher, Peter Giblin, Stefan Rahmann and Jun Sato. The figures of visual events were produced with the Liverpool Surfaces Modelling Package.

References

Åström, K., Cipolla, R. & Giblin, P.J. 1996 Generalised epipolar constraint. In *Proc. 4th . Conf.*

on Computer Vision, Cambridge (England), volume 2, pages 97–108, 1996.

Boyer, E. & Berger, M.O. 1997 3D Surface reconstruction using occluding contours. *Int. Journal of Computer Vision*, 22(3):219–233.

Cham, T.J. & Cipolla, R. 1996 MDL-based curve representation using B-spline active contours. In *Proc. British Machine Vision Conference*, Edinburgh, pages 363–372, (September) 1996.

Cipolla, R. & Blake, A. 1992 Surface shape from the deformation of apparent contours. *Int. Journal of Computer Vision*, 9(2):83–112, 1992.

Cipolla, R., Åström, K. & Giblin, P.J. 1995 Motion from the frontier of curved surfaces. In *Proc. IEEE 5th Int. Conf. on Computer Vision*, pages 269–275, Boston, (June) 1995.

Cipolla, R., Fletcher, G.J. & Giblin, P.J. 1997 Following Cusps. *Int. Journal of Computer Vision*, 23(2):115–129.

Faugeras, O.D., Luong, Q.-T, & Maybank, S.J. 1992 Camera self-calibration: theory and experiments. In *Proc. 2nd European Conference on Computer Vision*, pages 321–334, Santa Margherita Ligure, Italy May 1992, Lecture Notes in Computer Science 588, Springer–Verlag.

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

Giblin, P.J. & Weiss, R. 1994. Epipolar fields on surfaces. In *Proc. 3rd European Conference on Computer Vision*, volume 1, pages 14–23, Stockholm, May 1994, LNCS 800, Springer–Verlag.

Giblin, P.J., Pollick, F.E. & Rycroft, J.E. 1994. Recovery of an unknown axis of rotation from the profiles of a rotating surface. J. Opt. Soc. America, 11A:1976–1984.

Koenderink, J.J. & Van Doorn, A.J. 1976. The singularities of the visual mapping. *Biological Cybernetics*, 24:51–59.

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

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

Koenderink, J.J. 1990 *Solid Shape.* MIT Press.

Longuet-Higgins, H.C. 1981 A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135.

Luong, Q.-T, & Faugeras, O.D. 1996 The fundamental matrix: theory, algorithms, and stability analysis. *Int. Journal of Computer Vision*, 17(1):43–76.

Porrill, J. & Pollard, S. 1991 Curve matching and stereo calibration. *Image and Vision Computing*, 9(1):45–50.

Rieger, J.H. 1986 Three dimensional motion from fixed points of a deforming profile curve. *Optics Letters*, 11:123–125.

Sato, J. & Cipolla, R. 1998 Affine Reconstruction of Curved Surfaces from Uncalibrated Views of Apparent Contours. In *Proc. IEEE 6th Int. Conf. on Computer Vision*, Bombay, pages 715–720, (January) 1998.

Vaillant, R. & Faugeras, O.D. 1992 Using extremal boundaries for 3D object modelling. *IEEE Trans. Pattern Recognition and Machine Intelligence*, 14(2):157–173.