

Motion from the frontier of curved surfaces

Kalle Åström Roberto Cipolla Peter J. Giblin

November 1994

Abstract

The *frontier* of a curved surface is the envelope of contour generators showing the boundary, at least locally, of the visible region swept out under viewer motion. In general, the outlines of curved surfaces (apparent contours) from different viewpoints are generated by different contour generators on the surface and hence do not provide a constraint on viewer motion. We show that *frontier* points, however, have projections which correspond to a real point on the surface and can be used to constrain viewer motion by the epipolar constraint.

We show how to recover viewer motion from frontier points for both continuous and discrete motion, calibrated and uncalibrated cameras. We present preliminary results of an iterative scheme to recover the epipolar line structure from real image sequences using only the outlines of curved surfaces. A statistical evaluation is also performed to estimate the stability of the solution.

1 Introduction

Structure and motion from the images 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. Structure and motion from the outlines of curved surfaces, on the other hand, has been thought to be more difficult because of the *aperture problem*, i.e. it is not possible to get the correspondence of points between two images of the same curve.

For a smooth arbitrarily curved surface an important image feature is the outline or apparent contour. This is the projection of the locus of points on the surface which separates the visible from the occluded parts (figure 1). Under perspective projection this locus – the critical set or contour generator, Σ – can be constructed as the set of points on the surface where rays through the projection centre \mathbf{c} are tangent to the surface. Each viewpoint will generate a different contour generator with the contour generators ‘slipping’ over the visible surface under viewer motion (figure 2).

Under *known* viewer motion, the deformation of apparent contours can be used to recover the surface geometry (structure) [9, 4, 16]. This requires a spatio-temporal parameterization of image-curve motion. The latter is of course underconstrained since the mapping between contour generators, and hence between apparent contours, at successive instants can not be uniquely determined since the image contours are projections of different 3D space curves. 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 apparent contours is set up by matching along epipolar lines. 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. The parametrization leads to simplified expressions for the recovery of depth and surface curvature from image velocities and accelerations and known viewer motion.

In this paper we address the problem of recovering the viewer motion from the deformation of apparent contours. A solution can be found by considering the cases in which the epipolar parameterization is degenerate and so can not be used to recover the local surface geometry. These are:

1. Cusps or singular apparent contours

This occurs when viewing a hyperbolic patch along an *asymptotic* direction. The ray is not only tangent to the surface but also to the contour generator and the effect is to generate a cusp in the apparent contour. For opaque surfaces, only one branch of the cusp is visible and the contour ends abruptly [11, 10]. Although cusps can be detected and tracked under viewer motions Cipolla and Giblin [5]

have shown that they do not provide any constraints on viewer motion. They can however be used to recover the surface geometry by the image motion of the cusp to induce an alternative parameterisation of the surface in the vicinity of the *cusp generator* locus on the surface.

2. Rigid space curves

The second case of degeneracy is when the contour generator does not *slip* over the surface with viewer motion but is fixed. The contour generator is not an extremal boundary but is fixed to the surface or is 3D rigid space curve (surface marking or discontinuity in depth or orientation). Despite the *aperture* problem – only the normal component of image velocity can be measured from local measurements of a curve – Faugeras et al [7] and Cipolla [3] have shown how in principle that the spatio-temporal image of a space curve under viewer can be used to derive a constraint on the viewer motion from second order spatio-temporal derivatives. This has not yet been proved to be of practical use due to the difficulty in accurately extracting second order spatio-temporal derivatives. These methods can be used to derive a constraint at all points on a curve. An alternative is to use the normal velocities at special, isolated points (see below).

3. Frontier points

The remaining case of degeneracy of the epipolar parameterisation occurs for epipolar planes (spanned by the direction of translation and the ray) which coincide with tangent planes to the surface. The points of contact on the surface are called *frontier points* – so called because they are locally on the boundary of the visible region swept out by the contour generators. The frontier is in fact the locus of intersections of consecutive contour generators in an infinitesimal sense (figure 2).

The surface can not be reconstructed by the epipolar parameterisation at these points since the contour generator is locally stationary. However instantaneously frontier points correspond to real, fixed feature points on the surface and hence can provide a constraint on viewer motion.

The special case of frontier points under orthographic projection and object rotation about a single axis was considered by Rieger [15] and Giblin et al [8]. Porrill and Pollard [13], although primarily concerned with stereo calibration from 3D space curves, noted that the intersection of the two contour generators from two discrete viewpoints generated a real point visible in both images which could also be used to generate an epipolar constraint.

In this paper we show how frontier points can be detected in image sequences and used to recover viewer egomotion. We analyse both the continuous (infinitesimal) and discrete viewer motion cases as well as considering calibrated and uncalibrated cameras. We present preliminary experimental results obtained from real image sequences of curved surfaces from unknown viewpoints. An iterative technique is implemented which recovers the epipolar line structure (*essential matrix*) from the image motion of frontier points. A statistical evaluation is also performed to estimate the stability of the solution.

2 Frontier Points

Consider a surface $\mathbf{r}(u, v)$ parametrized locally by two parameters u and v , and a camera motion with projection centre $\mathbf{c}(t)$ parametrized by time t . The condition on the three variables u, v, t that the point $\mathbf{r}(t)$ lies on the contour generator at time t is simply

$$(\mathbf{r}(u, v) - \mathbf{c}(t)) \cdot \mathbf{n}(u, v) = 0, \quad (1)$$

where $\mathbf{n}(u, v)$ is the normal to the surface at $\mathbf{r}(u, v)$. The equation (1) can be thought of as defining a family of curves in the u, v parameter plane. As such the *envelope* of the family of curves, which can be thought of as the locus of intersections of ‘consecutive’ curves of the family, is given by differentiating (1) with respect to t (compare [1, p.102]). This gives

$$\mathbf{c}_t \cdot \mathbf{n} = 0, \quad (2)$$

the suffix t denoting differentiation. What this says is that points $\mathbf{r}(u, v)$ obtained by eliminating t between (1) and (2) are precisely the points of the *envelope of contour generators* on the surface. We call this envelope the *frontier* of the surface relative

to the given motion. Over a short period of time, the part of the surface covered by the contour generators is *on one side* of this frontier. Note that at frontier points the epipolar parametrization breaks down since the contour generators cannot form part of a coordinate system on the surface near the frontier as they do not cross one another transversally. This is illustrated in Figure 3, where the frontier is a curve with one point corresponding to each value of t .

Suppose we exclude the case where the direction of translation, \mathbf{c}_t , is parallel to the view vector $\mathbf{r} - \mathbf{c}$. (In this case the camera is heading straight for the contour generator point.) Then the epipolar plane is spanned by the vectors \mathbf{c}_t and $\mathbf{r} - \mathbf{c}$. The second of these is automatically perpendicular to the surface normal \mathbf{n} (by (1)) so the epipolar plane coincides with the tangent plane to the surface precisely when the first vector, \mathbf{c}_t , is also perpendicular to the normal. This is the condition (2). Thus frontier points can also be described as *epipolar tangency points*: the epipolar plane is the tangent plane.

Note that if the motion is *linear* then the frontier degenerates: for \mathbf{c}_t is then a *constant* vector, and the condition (2) becomes independent of t . If a point $\mathbf{r}(u, v)$ lies on the frontier at some time t then this same point continues to satisfy the frontier condition at subsequent times: the motion lies *in* the tangent plane at this isolated frontier point. This is illustrated in Figure 4. The same would hold for any motion which lay entirely in the tangent plane to a surface at a particular point of the surface.

3 Motion constraints at frontier points

3.1 Infinitesimal motion

As shown above frontier points are points on a smooth surface where the tangent plane coincides with the epipolar plane. From their definition as points which satisfy $\mathbf{c}_t \cdot \mathbf{n} = 0$, it is obvious that if they can be detected they provide a simple constraint on the direction of translation.

Consider the projection of a frontier point onto the image sphere. Since the surface normal projects to the normal to the apparent contour and the epipolar plane projects to a great-circle with poles on the unit sphere, $e = \pm \mathbf{c}_t / |\mathbf{c}_t| \in S^2$, the following property

holds.

Property 1 *Frontier points project to points on the image sphere where an epipolar great-circle (with poles defined by the direction of translation) is tangent to the apparent contour.*

Thus the projection of a frontier point in the image provides a great-circle constraint on the direction of translation, \mathbf{c}_t . If several such frontier points exist, the corresponding great circles all intersect at the same two points $\pm e \in S^2$. It is thus possible to determine the direction of translation, \mathbf{c}_t .

How can the projection of the frontier points be detected when viewer motion is unknown? A solution follows by considering the image motion.

The normal image velocity for a point on an apparent contour on the image sphere, \mathbf{q} , is given by [4]:

$$\mathbf{q}_t \cdot \mathbf{n} = \frac{-\mathbf{c}_t \cdot \mathbf{n}}{\lambda} - (\boldsymbol{\Omega} \times \mathbf{q}) \cdot \mathbf{n} \quad (3)$$

where λ is the distance from the projection centre to the contour generator point on the surface and $\boldsymbol{\Omega}$ is the angular rotational velocity. (Note that this is the same as for a point on a static space curve at depth λ – contour generator points and fixed points differ in their image accelerations.)

Consider the simpler case of no viewer rotation about the projection centre. (If the rotation is known the rotationally component of image motion can be subtracted since it is independent of scene structure). From the definition of frontier points and (3) we see that the normal velocity at the image of a frontier point is zero. The *envelope* of the apparent contours is then the projection of the frontier. This can be used to detect frontier points and hence recover the direction of translation. For a discrete motion the bitangents at two consecutive apparent contours are epipolar great circles. The projection of the frontier point has the same epipolar tangent but a different image position (figure 5).

What happens when viewer rotation is unknown? From (3) we see that the image velocity at frontier points has a special structure with zero translational component.

Property 2 *At the projection of frontier points, (points where the tangents of the apparent contours go through the epipole e) the normal velocity is determined by the rotational motion only and is independent of viewer translation and scene structure.*

$$\mathbf{q}_t \cdot \mathbf{n} = -(\boldsymbol{\Omega} \times \mathbf{q}) \cdot \mathbf{n} \quad (4)$$

This property can be used in an iterative scheme to search for frontier points. After hypothesising a direction of translation as a point e on the image sphere, (4) evaluated at 4 or more epipolar tangency points can be used to check for a consistent solution to the rotation velocity (three unknowns). In matrix notation (4) can be expressed by:

$$[Q] \begin{bmatrix} 1 \\ \Omega_1 \\ \Omega_2 \\ \Omega_3 \end{bmatrix} = 0 \quad (5)$$

where Q is a $N \times 4$ matrix where for each row the first element is the normal velocity, $(\mathbf{q}_t \cdot \mathbf{n})$, and the next three elements are the components of $(\mathbf{n} \times \mathbf{q})$. For a consistent solution the determinant of $Q^T Q$ should be zero. The value of this determinant can be used to iteratively solve for the position of the epipole, e . This technique was exploited for 3D rigid space curves features [2]. We have shown that it is also applicable to the apparent contours of curved surfaces. This is extended and explored below for the case of discrete viewer motion.

3.2 Discrete motion with calibrated cameras

Discrete motion can be considered as a special case of linear motion (figure 4). The frontier degenerates to a frontier point which is the intersection of the two contour generators from the two viewpoints. It is the point where the epipolar plane is tangent to the surface. Its projection in both image spheres will be where the epipolar great circles are tangent to the apparent contours. The projections of the frontier point will correspond to the same real point on the surface. In general, of course, apparent contours of the same surface in two viewpoints will not be projections of the same surface contour.

The epipole and a set of epipolar tangency points in each image determine a pencil of epipolar plane which are tangent to the surface at frontier points. Both images will in fact generate the same pencil of epipolar planes and this can be exploited

to determine the direction of translation (2 parameters), \mathbf{T} , and the rotation R (3 parameters) between the viewpoints. This information is usually represented by the *essential matrix* [12], $E = (\mathbf{T} \times)R$ where $(\mathbf{T} \times)$ is the matrix that represents cross product with \mathbf{T} . For the optimisation analysis in the next section we introduce a more convenient representation of the motion between viewpoints. We use the position of the epipoles in the left and right image, e_1 and e_2 (4 parameters) and the connection between oriented lines (planes) through the epipoles, i.e. basically a planar rotation, (1 parameter). In all representations 5 independent parameters are needed to calibrate the two viewpoints if the camera intrinsic parameters are known. Iterative techniques using the coefficients of the essential matrix [13] estimate 8 parameters.

Consider two images ω_1 and ω_2 of rigid objects. The relation between image features in the two images fulfill the following two equivalent properties (figure 5).

Property 3 *The features of $R * \omega_1$ and ω_2 have the same tangents through \mathbf{T} .*

Property 4 *The tangents of ω_1 through e_1 and the tangents through of ω_2 through e_2 are related to each other with a one-dimensional rotation.*

Remark. The two images ω_1 and ω_2 are features on the viewing sphere. The rotation of such an image is well defined. The tangent through a point is also well defined.

The two properties are equivalent. The first one is asymmetric with respect to the images and the formulation is similar to the continuous velocity case. The second property is symmetric with respect to the two images and the formulation is similar to the uncalibrated case.

Notice also that all features can be used, i.e. points, rigid planar curves, rigid space curves and the contour generators of curved surfaces of surfaces. \square

3.3 The discrete motion case with uncalibrated cameras

Again consider two images ω_1 and ω_2 of rigid objects. In the uncalibrated case we can hope to determine the fundamental matrix F . The fundamental matrix F [6] is a rank 2 matrix defined up to scale, (7 parameters). We can represent this by the epipoles in the left and right image, e_1 and e_2 (4 parameters) and the connection

between oriented lines through the epipoles, i.e. basically an one-dimensional oriented projective transformation, (3 parameters).

The relation between the two images fulfills the following property.

Property 5 *The tangents of ω_1 through e_1 and the tangents of ω_2 through e_2 are related to each other with one-dimensional oriented projective transformation.*

We now show to exploit these properties to solve for the direction of translation and rotation. An iterative scheme is now described.

4 Optimization

In the experiments planar images are taken and B-splines are fitted to apparent contours. With a calibrated camera these image curves determine a family of rays. We project the B-splines onto the viewing sphere. Given an image and an epipole, i.e. a point on the viewing sphere $e \in S^2$, it is straightforward to find the points on the contours where the tangents go through the epipole. Notice that on the viewing sphere the tangents are great circles. If the epipole is the north pole on the earth then tangents through the epipole are longitudes. After choosing a particular longitude or zero meridian, (the one through Greenwich), all tangents can be identified with a planar angle α . See Figure 5.

Given two images and an essential matrix E , the epipoles e_1 and e_2 in the left and right images can be computed. Furthermore the essential matrix relates the lines through e_1 with lines through e_2 . In particular a choice of zero meridian in one image determines the choice of zero meridian in the other. The difference between a tangent angle α_i in image 1 and corresponding tangent angle α'_i in image 2

$$\alpha_{err,i}(E) = \alpha_i - \alpha'_i$$

can thus be calculated. This difference is independent of the choice of zero meridian in the first image.

Ideally $\alpha_{err,i}(E)$ should be zero for all i at the correct essential matrix E , but due to measurements error it is not. The errors are roughly uncorrelated and usually of

the same magnitude. As a preliminary experiment we have chosen to estimate E with a non-linear least squares estimator with equal weights, i.e. we estimate E with the matrix \tilde{E} that minimizes

$$f = \sum_i (\alpha_{err,i}(E))^2.$$

Thus we have assumed that the error variances of $\alpha_{err,i}(E)$ are approximately equal. A proper investigation of these variances can and will be included in the future.

We are minimizing over a manifold of essential matrices

$$M = \{E = (\mathbf{T} \times) R \mid \text{with } \mathbf{T} \in S^2 \text{ and } R \in SO(3)\}.$$

It is not possible to choose a good global parametrization so at each iteration with $E = E_k$ a new parametrization of a neighbourhood of the manifold M around the point E_k was chosen, i.e. a function E_k

$$R^5 \ni x \mapsto E_k(x) \in M, \quad E_k(0) = E_k.$$

We chose a parametrization $x = (x_1, x_2, x_3, x_4, x_5)^T$, in which the first two parameters describe changes in the first epipole e_1 , the next two parameters describe changes in the second epipole e_2 and the last one roughly describe a rotation θ around one epipole. Using this parametrization we determine a Taylor expansion of first order of $\alpha_{err}(E_k(x))$

$$\alpha_{err}(E_k(x)) \approx \alpha_0 + Bx$$

and a Taylor expansion of second order of $f(E_k(x))$

$$f(E_k(x)) \approx f_0 + (\partial f)^T x + x^T (\partial^2 f) x / 2.$$

Standard optimization techniques were then used to find the a better choice of essential matrix E_{k+1} . Since we are only searching through a five parameter space and since quadratic convergence is achieved through the use of Newton-Raphson like techniques a local optimum can be found quite fast.

When a local minima has been found standard techniques can also be used to estimate the magnitude of the image errors and also to estimate the covariance matrix of the estimated essential matrix.

For the linear problem it can be shown [14], that the minimal sum of squared errors can be used to estimate the variance $\sigma^2[\alpha_{err}]$ of the individual errors, i.e.

$$\tilde{\sigma}^2 = \frac{f_{\min}}{(c - d)}$$

where c is the number of measurements and d is the number of estimated variables. In the linear case this is an unbiased estimate, but in the non linear case this estimate although biased will give us some clues on the magnitude of the image errors. For small errors the linearization is adequate and the bias should not be a problem. A plot over the residuals α_{err} can also be helpful. Sometimes just by looking at such a plot irregularities and large outliers can be found.

Using the linearization of the parametrization around the minima

$$\alpha_{err}(E(x)) \approx \alpha_0 + Bx$$

an estimate of the covariance matrix of the estimated parameters x , can be found,

$$V[x] \approx \sigma^2(B^T B)^{-1}.$$

The eigenvectors and eigenvalues of V gives useful information about the directions in which the epipoles are badly estimated. The two largest eigenvalues are illustrated by uncertainty ellipses in Figure 8.

The only technical part is the choice of parametrization $x \mapsto E(x)$ and the calculation of ∂f , $\partial^2 f$ and B . It turns out that the gradient ∂f and the matrix B is exactly the same as for point features. The curvature at the tangent point only affects the hessian $\partial^2 f$. The calculation are simpler in the planar case. A general feeling can be obtained by studying the following example.

Example. Assume that the epipole is at the origin $e = (0, 0)$ a curve patch $\mathbf{p}(s)$ has a tangent through the origin for $s = 0$, along the positive x-axis at distance X , i.e. $\mathbf{p}(0) = (X, 0)$ and $\mathbf{p}'(0) = (1, 0)$. Assume also that the curve has curvature k . Then the gradient of the angle α with respect to changes in the epipole is

$$\partial\alpha = \begin{bmatrix} 0 \\ -1/X \end{bmatrix}$$

and the second derivatives are

$$\partial^2 \alpha = \begin{bmatrix} 0 & -1/X^2 \\ -1/X^2 & -1/(kX^3) \end{bmatrix}.$$

The general planar case can be obtained by transforming the gradient and the second derivatives by a rotation matrix, and similarly the spherical case is obtained by a slightly more complicated transformation, but qualitatively the expressions are the same. \square

5 Experimental results

A few experiments on real images are presented in this section. In every experiment an approximate internal calibration matrix has been used. B-spline curves are fitted to apparent contours. A couple of different starting points were chosen for the essential matrix and then the iterative scheme described above was used to find a local minima. The residuals at the minima are used to evaluate the solution. A statistical evaluation was then performed to estimate the stability of the solution.

In the first experiment, see Figure 6, five apparent contours of corresponding fruits were used to estimate the essential matrix. The solution is illustrated in Figure 7 by using the estimated essential matrix to predict the epipolar lines for all tangencies in the left image. The epipolar line structure is qualitatively excellent. Epipolar tangencies in the left image remain epipolar tangencies in the right view indicating the correct detection of frontier points and displacement and rotation between the two viewpoints. A rough estimate of the standard deviation $\sigma[\alpha_{err,i}]$ for the errors gave

$$\sigma = 0.94 \cdot 10^{-3}.$$

A study of the stability of the estimated essential matrix showed that it was particularly poorly estimated in two directions. Using the local parametrisation, the five estimated standard deviations for x around the minima, i.e. square root of eigenvalues of $V[X]$, were

$$\lambda = (0.125, 0.033, 0.003, 0.001, 0.0002).$$

The worst directions (the one corresponding to eigenvalue 0.125) correspond to a combined change in epipoles toward or away from the group of image curves. The second direction correspond to a combined change of epipoles sideways with respect to the group of image curves. This is illustrated in Figure 8.

Experiments were also performed on a sequence of images of a Henry Moore statue, see Figure 9. Two plausible local minima were found, one with eight tangencies, (Figure 9) and one with six tangencies, (Figure 10). Both epipolar line structures are in agreement with the image data. However they yield different solutions. The statistical evaluation indicate in both cases very unstable solutions. This is probably caused by the low number of tangencies, the small baseline and perhaps a poorly calibrated camera. These are the same problems that plague all structure from motion algorithms.

The problem of local minima is also illustrated in the still life sequence, Figure 11 and 12. Choosing the centre of the images as starting points for iteratively estimating the positions of the epipoles gives a bad local minima, whereas choosing the starting point for the pair of epipoles roughly in the right direction gives a better solution to the minimization problem. A good initial estimate of the direction of translation is therefore important. In the reconstruction of surfaces from the deformation of apparent contours, egomotion is usually known, albeit not to the accuracy required to recover accurate curvatures [4]. The technique presented here can be used to refine the estimate of viewer motion.

6 Conclusions and future work

We have shown how to recover viewer motion from frontier points in both the continuous and discrete motion cases and proposed an iterative method for both calibrated and uncalibrated cameras. We have presented preliminary results of an iterative scheme to recover the epipolar line structure from real image sequences of families deforming apparent contours of curved surfaces.

The apparent contour and its deformation under viewer motion is known to be a rich source of surface geometric information which can be used in visual navigation and object manipulation. Here we have shown how the frontier points of apparent contours can also be used to recover the viewer motion.




Figure 1: Perspective projection: the contour generator Σ with a typical point \mathbf{r} , the image sphere with centre \mathbf{c} and the corresponding apparent contour point $\mathbf{c} + \mathbf{p}$. Thus \mathbf{p} is the unit vector joining the centre \mathbf{c} to the apparent contour point. Also \mathbf{n} is normal to the surface at \mathbf{r} .




Figure 2: Degenerate case of epipolar parameterisation. The epipolar plane is a tangent plane of the surface at a frontier point. Movement of the viewpoint will cause the contour generators to sweep over the surface. At a frontier point the contour generators from consecutive viewpoints intersect.




Figure 3: The frontier. A surface and five contour generators for nearby (equally spaced) intervals of time. Each of these corresponds to a definite value of t in equation (1). Contour generators for nearby values of t intersect approximately on the frontier. The frontier is actually the limit of intersection points as the time interval tends to 0.




Figure 4: For translational motion in a straight line the frontier degenerates to a point through which all the contour generators pass. The motion is in the tangent plane at the isolated frontier point.

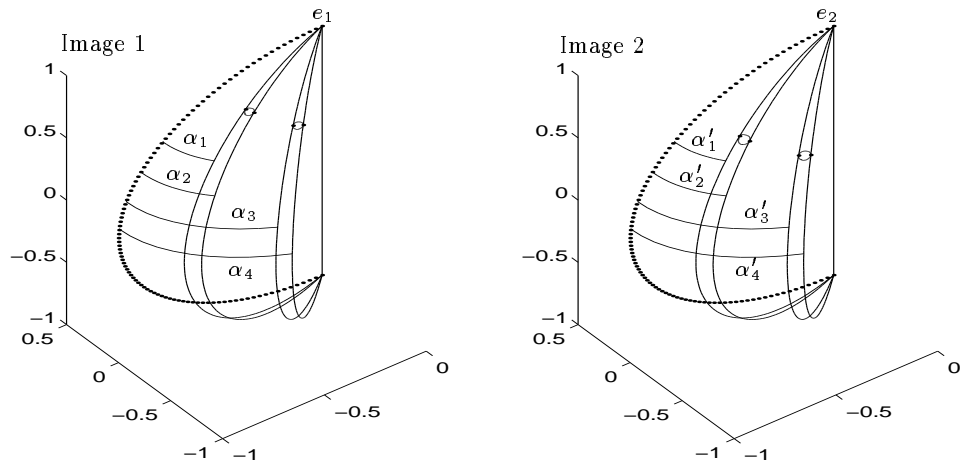




Figure 5: Once a base direction is chosen, see the dotted line in the images, an angle α is associated with every epipolar tangency plane.




(a)

(b)

Figure 6: Left (a) and right (b) images showing 5 apparent contours and 10 frontier points used to estimate the epipolar line structures. Epipolar tangents are shown in the left image with their corresponding epipolar lines in the right image. If the estimate of motion between the views is correct the set of epipolar tangencies define the frontier points on the curved surfaces which are visible in both views.



(a)



(b)

Figure 7: Estimated epipolar line structures shown for all frontier points.





Figure 8: Estimated epipole and uncertainty shown on image sphere viewed from above. Apparent contours are also shown.




Figure 9: Henry Moore sculpture. 5 epipolar tangencies are used to estimate the motion between frames. Although they yield a consistent estimate of the motion between frames (judged by quality of epipolar line structure) the solution is ill-conditioned.




Figure 10: Henry Moore sculpture. 8 epipolar tangencies lead to convergence to different local minima. Due to the small field of view and because the direction of translation is outside the image frame the solution is and very sensitive to image contour localisation errors.




Figure 11: The optical axis is used as the initial guess for the epipole in the left and right images. Top row shows left and right images with epipolar tangencies in the left image and their corresponding epipolar lines in the right image. The error in the initial guess is seen by the fact that the epipolar lines are not tangent to the apparent contour. The bottom row shows the left and right images after convergence. Convergence to a local minima follows due to a poor starting position







Figure 12: Convergence to global minima.

References

- [1] J.W. Bruce and P.J. Giblin. *Curves and Singularities*. Cambridge University Press, 1984.
- [2] S. Carlsson. Sufficient image structure for 3D motion and shape estimation. In J-O. Eklundh, editor, *Proc. 3rd European Conf. on Computer Vision*, volume I, pages 83–91. Springer–Verlag, 1994.
- [3] R. Cipolla. *Active Visual Inference of Surface Shape*. PhD thesis, University of Oxford, 1991.
- [4] R. Cipolla and A. Blake. Surface shape from the deformation of apparent contours. *Int. Journal of Computer Vision*, 9(2):83–112, 1992.
- [5] R. Cipolla and P.J. Giblin. Following cusps. Technical Report CUED/F-INFENG/TR 194, University of Cambridge, October 1994.
- [6] O.D. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig. In G. Sandini, editor, *Proc. 2nd European Conference on Computer Vision*, pages 563–578. Springer–Verlag, 1992.
- [7] O.D. Faugeras and T. Papadopoulo. A theory of the motion fields of curves. *Int. Journal of Computer Vision*, 10(2):125–156, 1993.
- [8] P.J. Giblin, F.E. Pollick, and J.E. Rycroft. Recovery of an unknown axis or rotation from the profiles of a rotating surface. *J. Opt. Soc. America*, A11:1976–1984, 1994.
- [9] P.J. Giblin and R. Weiss. Reconstruction of surfaces from profiles. In *Proc. 1st Int. Conf. on Computer Vision*, pages 136–144, London, 1987.
- [10] J.J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.
- [11] J.J. Koenderink and A.J. Van Doorn. The shape of smooth objects and the way contours end. *Perception*, 11:129–137, 1982.

- [12] H.C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, 1981.
- [13] J. Porrill and S.B. Pollard. Curve matching and stereo calibration. *Image and Vision Computing*, 9(1):45–50, 1991.
- [14] C. R. Rao. *Linear Statistical Inference and Its Application*. John Wiley & Sons, Inc. 1965.
- [15] J.H. Rieger. Three dimensional motion from fixed points of a deforming profile curve. *Optics Letters*, 11:123–125, 1986.
- [16] R. Vaillant and O.D. Faugeras. Using extremal boundaries for 3D object modelling. *IEEE Trans. Pattern Analysis and Machine Intell.*, 14(2):157–173, 1992.