

Robust Shape Recovery from Occluding Contours Using a Linear Smoother

Richard Szeliski

Digital Equipment Corporation
Cambridge Research Lab
One Kendall Square, Bldg. 700
Cambridge, MA 02139

Richard Weiss

Computer and Information Science Department
University of Massachusetts at Amherst
Amherst, MA 01003

Abstract

The occluding contour of a curved surface is an important source of information about its shape. However, recovering the shape of an object from triangulation fails at occluding contours of smooth objects because the contour generators are view dependent. For three or more views, shape recovery is possible, and several algorithms have recently been developed for this purpose. Our approach uses a linear smoother to optimally combine all of the measurements available at the contours (and other edges) in all of the images. This allows us to extract a robust and dense estimate of surface shape, and to integrate shape information from both surface markings and occluding contours.

The problem of reconstructing a smooth surface from its *profiles* (also known as extremal boundaries or occluding contours) has been explored for known planar motion by Giblin and Weiss [3] and subsequently for more general known motion by Vaillant [6] and Cipolla and Blake [1]. These approaches require second derivatives of edge point locations. Unfortunately, determining differential quantities reliably in real images is difficult. Cipolla and Blake use relative measurements in order to cancel some of the error due to inadvertent camera rotation, and B-snakes to smooth the contours in the image, which requires the initialization of each contour to be tracked.

To overcome these limitations, we apply estimation theory (Kalman filtering and smoothing) to make optimal use of each measurement without computing differential quantities. In [5], we derive a *linear* set of equations between the unknown shape (surface point positions and radii of curvature) and the measurements. We then develop a robust linear smoother [2] to compute statistically optimal current and past estimates from the set of contours. Smoothing allows us to combine measurements on both sides of each surface point.

Our technique produces a complete surface description, i.e., a network of linked 3D surface points, which provides us with a much richer description than just a set of 3D curves. Due to self-occlusion and occlusion by other surfaces, some parts of the surface may never

appear on the profile. Since the method presented here also works for arbitrary surface markings and creases, a larger part of the surface can be reconstructed than from occluding contours of the smooth pieces alone.

The surface being reconstructed from a moving camera can be parametrized in a natural way by two families of curves [3, 1]: one family consists of the *critical sets* (also known as contour generators or limbs) on the surface; the other is tangent to the family of rays from the camera focal points. The latter curves are called *epipolar curves*. The problem is that any smooth surface reconstruction algorithm which is more than a first order approximation requires at least three images and, that in general, the three corresponding tangent rays will not be coplanar. However, there are many cases when this will be a good approximation. One such case is when the camera trajectory is almost linear.

Given three or more edgels tracked with our technique, we compute the location of the surface and its curvature by fitting a circular arc to the lines defined by the view directions at those edgels. In general, a space curve will have a unique circle which is closest to the curve at any given point. This is called the *osculating circle*, and the plane of this circle is called the *osculating plane*. It is easy to see that the epipolar plane is an estimate of the osculating plane [1], and the lines defined by the view directions can be projected onto this plane.

The overall sequence of processing steps is the following. Initially, we perform a batch fit to the first three frames, using the last frame as the reference frame. Next, we convert the local estimate into a global 3D position and save it as part of our final surface model. Then, we predict the 3D surface point and its radius onto the next frame, i.e., into the frame defined by the next 2D edgel found by the tracker. We repeat the above process so long as a reliable track is maintained (i.e., the residuals are within an acceptable range). If the track disappears or a robust fit is not possible, we terminate the recursive processing and wait until enough new measurements are available to start a new batch fit.

The generalization of the Kalman filter to update previous estimates is called the *linear smoother* [2]. The three commonly used types of smoothing are *fixed-interval smoothing*, *fixed-point smoothing*, and *fixed-lag smoothing* [2]. For our contour-based shape recovery algorithm, we have developed a new fixed-lag smoother,

*The second author acknowledges the support of grants TACOM DAAE07-91-C-RO35, NSF IRI-9208920, and NSF IRI-9116297.




Figure 1: Image from synthetic truncated ellipsoid sequence, reconstructed profile curves (oblique and top view) and portion of meshed surface.

which fits in naturally with the batch and Kalman filter approaches. Our fixed-lag smoother begins by computing a *centered* batch fit to $n \geq 3$ frames. The surface point is then predicted from frame $i-1$ to frame i as with the Kalman filter, and a new measurement from frame $i+L$, $L = \lfloor n/2 \rfloor$ is added to the predicted estimate. The addition of measurements ahead of the current estimate is straightforward using the projection equations for the least-squares (batch) fitting algorithm.

To determine the performance of our shape reconstruction algorithm, we generated a synthetic motion sequence of a truncated ellipsoid rotating about its z axis (Figure 1). The camera is oblique (rather than perpendicular) to the rotation axis, so that the trajectories of the pixels are not linear, and the reconstruction plane is continuously varying over time.

When we run the edge images through our least-squares fitter or Kalman filter/smoothing, we obtain a series of 3D curves. The curves corresponding to the surface markings and ridges (where the ellipsoid is truncated) should be stationary and have 0 radius, while the curves corresponding to the occluding contour should continuously sweep over the surface.

Figure 1 (middle two images) shows all of the 3D curves overlaid in a single image. As we can see, the 3D surface is reconstructed quite well, except for parts which do not appear in profile. These results were obtained using the linear smoother with $n = 7$ window size. The rightmost image show a portion of the surface patch created by linking successive matched edges along their epipolar curves. Most points on the smooth portion of the surface will be covered by two such meshes, since their corresponding profiles will have been seen from two different viewpoints (once as the point appears and once as it disappears). We are currently developing an algorithm to merge these meshes, along with the recovered surface marking curves, into a single surface description.

We have also applied our algorithm to real image sequences [5] which were obtained by placing an object on a rotating mechanized turntable whose edge has a

Gray code strip used for reading back the rotation angle [4]. The camera motion parameters for these sequences were obtained by first calibrating the camera intrinsic parameters and extrinsic parameters to the turntable top center, and then using the computed turntable rotation. Our results indicate that the overall shape of the objects is reconstructed quite well.

To summarize, this paper extends previous work on both the reconstruction of smooth surfaces from profiles and on the epipolar analysis on spatiotemporal surfaces. The ultimate goal of our work is the construction a complete, detailed geometric and topological model of a surface from a sequence of views. A more complete version of this paper will appear in [5].

References

- [1] R. Cipolla and A. Blake. Surface shape from the deformation of apparent contours. *Int. J. Computer Vision*, 9(2):83-112, November 1992.
- [2] A. Gelb, editor. *Applied Optimal Estimation*. MIT Press, Cambridge, Massachusetts, 1974.
- [3] P. Giblin and R. Weiss. Reconstruction of surfaces from profiles. In *ICCV'87*, pp. 136-144, London, England, June 1987.
- [4] R. Szeliski. Shape from rotation. In *CVPR'91*, pp. 625-630, Maui, Hawaii, June 1991.
- [5] R. Szeliski and R. Weiss. Robust shape recovery from occluding contours using a linear smoother. In *Image Understanding Workshop*, Washington, D. C., April 1993. Morgan Kaufmann Publishers. An extended version is available as a technical report from either author.
- [6] R. Vaillant and O. D. Faugeras. Using extremal boundaries for 3-D object modeling. *IEEE Trans. PAMI*, 14(2):157-173, February 1992.