In Proceedings of CVPR '93

Reconstruction of HOT Curves from Image Sequences

David J. Kriegman B. Vijayakumar
Department of Electrical Engineering
Yale University
New Haven, CT 06520-1968

Jean Ponce
Beckman Institute
University of Illinois
Urbana, IL 61801

Abstract

Recently, a novel shape representation of general curved objects, which is suitable for object recognition, has been proposed; it is based on a set of surface curves, named HOT curves, defined by the locus of points where a line has high order tangency with the surface [16]. These curves determine the structure of an object's image contours and their catastrophic changes. A nat 2 ural correspondence between a point in an intensity image and some of these curves can be directly established. This correspondence can be used for pose estimation and indexing in recognition. It also permits their 3D reconstruction from feature points on the edges detected in a sequence of images under known observer motion. This paper presents an implemented reconstruction method and experimental results.

1 Introduction

In many computer vision systems, objects are represented by collections of primitives (e.g. polyhedra, quadrics, superquadrics, solids of revolution, and generalized cylinders). While implemented recognition systems have demonstrated their usefulness, the ultimate utility of a representation is limited by its scope. Recently, there has been interest in the use of algebraic surfaces since even moderate degree surfaces offer a large number of degrees of shape freedom [18], and methods are available for recognizing instances in a single image [10]. The dual motives for this representation are its wide application in computer aided geometric design and the availability of computational tools. However, any fixed set of primitives has its limitations, and so we take a different approach and focus on a representation that is directly accessible for the purposes of object recognition; since it is based on differential properties of arbitrary (generic) smooth surfaces and is encoded in a discrete fashion, wide scope is a natural consequence.

In particular, there are points on a surface where there exists a line in the tangent plane which have high order contact with the surface; the locus of such points forms curves of high order tangency (named HOT curves). These include the familiar parabolic curves as well as more exotic ones such as flecnodal curves and asymptotic bitangent curves. HOT curves were introduced in a recent paper [16], and their use in object modelling, aspect graph construction, and database indexing during recognition was discussed. In this paper, we focus on two of these curves, namely the parabolic and limiting bitangent developable curves; there exists a natural correspondence between points on HOT curves and their image. From this correspondence, a HOT curve can be reconstructed from the edges detected in a sequence of video images. While many of the underlying notions are strongly motivated by Koenderink's work [9], the reconstruction algorithm is related to that of Giblin and Weiss [5], Blake and Cipolla [2], and Vaillant and Faugeras [20] on estimating surface shape from the occluding contour.

2 HOT Curves, Bitangents, Inflections

Under pinhole perspective, the image of a point $\mathbf{p}$ is given by the intersection of an image plane with a ray (the line-of-sight) emanating from the camera center $\mathbf{c}$ in the direction of $\mathbf{p}$. The line drawing (image contours, intensity discontinuities or edges) of a smooth surface is the image of points on the surface (occluding contour, limb, rim, or contour generator) where the line-of-sight (viewing direction) grazes the surface. For a generic smooth surface and a generic viewpoint, the occluding contour is a regular surface curve while the image contours are piecewise smooth curves whose singular points are either transversal crossings (t-junctions) or cusps. The viewing direction lies in the tangent plane at all points on the occluding contour, and it is said to have second order (or higher) contact with the surface at these points.[1] At cusps, the line-of-sight has third order contact with the surface and is an asymptotic direction at the point. While any point on a surface may have second order contact with some line, only hyperbolic points may have third order contact. Contact of order four and higher only occurs along certain surface curves (i.e. parabolic and flecnodal curves), and fifth order contact only occurs at isolated points along these curves [15].

Additionally, there are other surface curves where a line grazes the surface in multiple discrete points with at least second order contact in some exceptional manner. For example, a line may contact the surface at two points and lie in their respective tangent planes. If the line-of-sight is

[1] The contact of a tangent line with a surface at a point is said to have n-point contact (or n-th order contact) iff the i-th derivative of the surface equation in the direction of the line is zero for $i < n$ and is non-zero for $i = n$. [3]

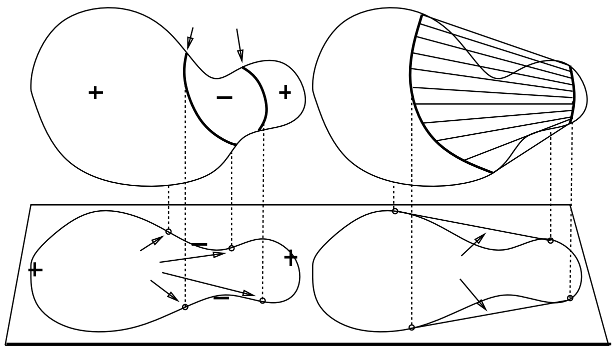

Figure 1. Parabolic curves, limiting bitangent developables, and their projections.

aligned with this line, a t-junction or crossing will generally be observed. In the special case where the surface normals at these points are aligned, the image contours will meet with a common tangent forming a tacnode. The locus of pairs of such points define two curves (or a curve in $\mathbb{R}^6$) on the surface which will be called the *limiting bitangent developable* curve [9]. Each pair of points defines a generator of a developable surface. Other lines may graze the surface at multiple points in special ways and define other HOT curves; however these will not concern us here. Interestingly, these special surface curves are the same ones used to define the visual events delineating stable views in an aspect graph [7, 9, 14]. For example, when the line-of-sight is aligned with the developable of a limiting bitangent, a tangent-crossing event is observed. A more thorough discussion of these and other HOT curves is presented in [16].

Just as points on a generic surface can be classified according to the order of contact with a tangent line, points of a generic plane curve can be similarly classified. As shown by Bruce and Giblin [3], such a curve has a discrete set of inflections with order three tangents and a discrete set of bitangents. The inflections have zero curvature and divide the curve into a discrete set of convex and concave arcs with tangents of order two.

There is a close relationship between two of the 3D HOT curves and the 2D contour inflections and bitangents. Koenderink [8] characterized the relationship between the curvature of an image contour, the two principal curvatures, and the viewing direction under orthographic projection. See [2, 20] for extensions to perspective. A consequence of this relationship is that the image of a parabolic point is generally an inflection. This defines a natural correspondence between observed inflections and parabolic points (Fig. 1.a).

Now, consider the limiting bitangent developable. The line between the two points and their common surface normal define a plane. If the line-of-sight lies in this plane, both points will be on the occluding contour, and the two corresponding image contour tangents will necessarily be aligned. Such a pair of image points defines a contour bitangent, and so once again, this yields a natural correspondence between a HOT curve and a contour feature (Fig. 1.b). The importance of bitangents and limiting bitangent developables was noted in [9], and contour bitangents have been used in invariant-based recognition of 2D objects and 3D solids of revolution [21].

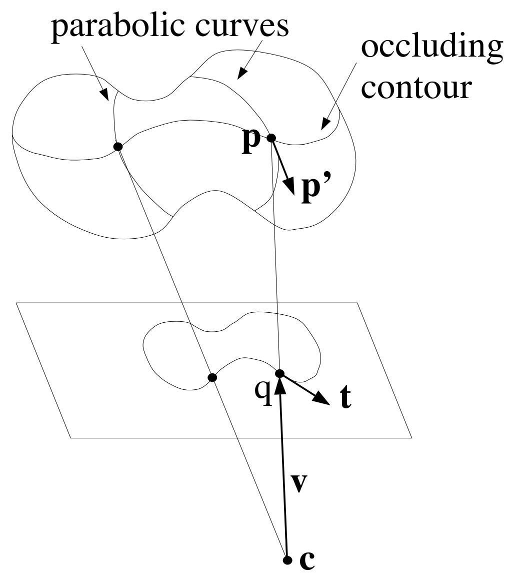

Figure 2. The geometry of curve reconstruction: Here, the parabolic line is reconstructed from inflection points.

3 Reconstruction from Video Images

We now present a method for reconstructing the 3D parabolic and limiting bitangent developable curves directly from a sequence of 2D video images. This follows the method of Giblin and Weiss [5] for reconstructing a surface from a continuous set of profiles. The 3D curve reconstruction is performed using quantities directly measurable in the image (i.e., feature points and their tangents) and derivatives directly computable from a sequence of images (the velocity of the feature). For example, the image of a parabolic point is a contour inflection, and as the camera moves, additional points on the parabolic line will be revealed as inflections. By measuring their location and motion, the 3D structure of the parabolic curve can be determined. This idea can be applied to the two endpoints of a limiting bitangent, or to any other surface curve where an image point to curve correspondence can be established.

Consider a fixed object and a moving pinhole camera with focal length f. Let ${}^w\mathbf{p} \in \mathbb{R}^3$ denote the coordinates in a global frame of the observed point $\mathbf{p}$ which lies at the intersection of the relevant surface curve and the occluding contour.[2] Also, let $\tilde{\mathbf{q}} = (u, v) \in \mathbb{R}^2$ denote the coordinates of the image of $\mathbf{p}$ under pinhole perspective projection (Fig. 2). The image coordinates are readily computed

[2] The coordinates of points and vectors are given with respect to some frame denoted by a leading superscript. Given the coordinates of a vector ${}^a\mathbf{p}$ in the a frame, a matrix ${}^b_a\mathbf{R}$ can be used to obtain the coordinates of $\mathbf{p}$ in a rotated b frame, ${}^b\mathbf{p} = {}^b_a\mathbf{R}\,{}^a\mathbf{p}$.

from pixel coordinates directly measured in an image.

A frame, whose coordinates are written with respect to some global or world frame, can be attached to the moving camera with the focal point or camera center at the origin ${}^w\mathbf{c}$ and with the first two rows of the rotation matrix ${}^w_c\mathbf{R}$ spanning the image plane. The coordinates in the camera frame of an image point $\tilde{\mathbf{q}} = (u, v)$ are ${}^c\mathbf{q} = (u, v, f)$, and its world coordinates are ${}^w\mathbf{q} = {}^w_c\mathbf{R}\,{}^c\mathbf{q} + {}^w\mathbf{c}$. Since the camera is moving, ${}^w_c\mathbf{R}$ and ${}^w\mathbf{c}$ are functions of time t. Because $\mathbf{p}$ lies on the ray joining $\mathbf{c}$ and $\mathbf{q}$, we have:

$$ {}^w\mathbf{p} = {}^w\mathbf{c} + \lambda\,{}^w\mathbf{v}, \tag{1} $$

where ${}^w\mathbf{v} = {}^w_c\mathbf{R}\,{}^c\mathbf{q}$ is the direction of the ray (the *line-of-sight*), and $\lambda \in \mathbb{R}$ is an unknown scalar. Note that the line-of-sight is different at each image point.

Let us parameterize the observed surface curve by t; its 3D tangent is $\mathbf{p}'(t) = d\mathbf{p}(t)/dt$. In addition, the image velocity $\mathbf{q}'$ of $\mathbf{q}$ can be estimated from a sequence of images. Let ${}^c\mathbf{t} = (t_x, t_y, 0)$ denote the measured tangent to the contour, and ${}^w\mathbf{t} = {}^w_c\mathbf{R}\,{}^c\mathbf{t}$ denote its coordinates in the global coordinate system. Since $\mathbf{p}$ lies on the occluding contour, the tangent $\mathbf{p}'$ lies in the plane spanned by the image contour tangent $\mathbf{t}$ and the line-of-sight $\mathbf{v}$. This can be written as:

$$ (\mathbf{t} \times \mathbf{v}) \cdot \mathbf{p}' = 0. \tag{2} $$

where $\mathbf{t}, \mathbf{v}$, and $\mathbf{p}'$ are written in the global frame. Differentiating (1) to get ${}^w\mathbf{p}'$ and substituting into (2) yields

$$ (\mathbf{t} \times \mathbf{v}) \cdot [\mathbf{c}' + \lambda'\mathbf{v} + \lambda\mathbf{v}'] = (\mathbf{t} \times \mathbf{v}) \cdot [\mathbf{c}' + \lambda\mathbf{v}'] = 0, \tag{3} $$

where $\mathbf{t}, \mathbf{v}, \mathbf{v}'$, and $\mathbf{c}'$ are again in world coordinates, and

$$ {}^w\mathbf{v}' = ({}^w_c\mathbf{R}\,{}^c\mathbf{q})' = {}^w_c\mathbf{\Omega}\,{}^w\mathbf{v} + {}^w_c\mathbf{R}\,{}^c\mathbf{q}'. \tag{4} $$

${}^w_c\mathbf{\Omega}$ is the skew symmetric angular velocity matrix (i.e., $\mathbf{R}' = \mathbf{\Omega}\mathbf{R}$). Note that ${}^w\mathbf{v}'$ is directly computable from the known camera orientation, camera rotation, feature location and feature velocity. Solving (3) for λ yields:

$$ \lambda = -\frac{(\mathbf{t} \times \mathbf{v}) \cdot \mathbf{c}'}{(\mathbf{t} \times \mathbf{v}) \cdot \mathbf{v}'}, \tag{5} $$

and once λ is computed, ${}^w\mathbf{p}$ is easily determined from (1). Note also that the surface normal at $\mathbf{p}$ is directly available in world coordinates as ${}^w\mathbf{n} = {}^w\mathbf{t} \times {}^w\mathbf{v}$.

Thus, a camera (or equivalently the object) can be systematically moved to reveal new points on the surface curve. Since the features are stable, almost any camera motion will do. A few remarks related to potential failings of the reconstruction procedure are in order.

When reconstructing parabolic lines, there may be isolated surface points $\mathbf{p}$ where the line-of-sight becomes aligned with the asymptotic direction at $\mathbf{p}$, and a visual event (lip or beak) occurs [7]. This should be detectable, and almost any motion of the camera center in the tangent plane at $\mathbf{p}$ (defined by the measured line-of-sight and contour tangent) will both keep $\mathbf{p}$ on the occluding contour and lead to a generic viewpoint. A similar problem occurs for the limiting bitangents when the viewpoint becomes aligned with the bitangent developable, and a visual event called a tangent crossing occurs. Again, the solution is to move the camera center within the tangent plane.

Another question requiring consideration is when does the denominator of (5) vanish? Since reconstruction is only attempted at regular image contour points, the contour tangent $\mathbf{t}$ is always non-zero, and the line-of-sight $\mathbf{v}$ is well defined for all image measurements. Since $\mathbf{t}$ and $\mathbf{v}$ are never collinear, the tangent plane at $\mathbf{p}$ can always be determined from a single image. Thus, the denominator of (5) vanishes when either $|\mathbf{v}'| = 0$ or when $\mathbf{v}'$ lies in the tangent plane. There are two types of camera motions that can lead to $|\mathbf{v}'| = 0$: first, when the camera center $\mathbf{c}$ moves along the line-of-sight; second, when camera motion is a pure rotation about an axis through $\mathbf{c}$, and the line-of-sight remains unchanged. Now, when $\mathbf{v}'$ lies in the tangent plane and $\mathbf{v}' \neq \mathbf{0}$, a given surface point $\mathbf{p}$ remains on the occluding contour and $\mathbf{p}'(t) = \mathbf{0}$. Though (5) cannot be used, the standard stereo equations can be applied; the tangent plane defines the epipolar plane. More globally, the camera motion may cause the surface point $\mathbf{p}$ to become occluded even though there may exist a viewpoint where it could be visible. Again, camera motion in the tangent plane will reveal the occluded point.

In general the denominator of (5) does not actually vanish, but the equation can become ill-conditioned. This accentuates image noise in reconstruction. More precisely, if measurement noise is small and normally distributed, the variance in λ can be determined from $\sigma_\lambda = \nabla\lambda^t \Sigma_m \nabla\lambda$ where $\nabla\lambda$ is taken with respect the measurements $({}^c\mathbf{q}, {}^c\mathbf{q}', {}^c\mathbf{t})$ and Σ_m is a covariance matrix describing measurement noise. By evaluating σ_λ for each point, the quality of the reconstruction can be determined; this has been used to automatically prune poorly constructed results and can be used when planning observer motion.

4 Implementation and Results

To fully reconstruct the HOT curves of complicated objects, it must be possible to move the object (or camera) with three degrees of freedom in order to place the line-of-sight in the tangent plane and then orient it within the tangent plane. However, to demonstrate the feasibility of reconstruction, the object is rotated on a turntable with a fixed camera; most points on the surface appear on the occluding contour. The equations presented in sec. 3 are easily rewritten in the fixed camera frame. Before reconstruction can commence, three issues must be considered: calibration, segmentation, and feature tracking.

To apply (1)–(5), it is necessary to compute the viewing direction for a particular feature point in world coordinates and the relative camera motion. Thus, the intrinsic camera parameters and the extrinsic relationship of the camera to the axis of rotation must be determined. Tsai's method is used to compute the intrinsic parameters [19]. The four additional parameters characterizing the camera-to-axis relationship are obtained from a sequence of images

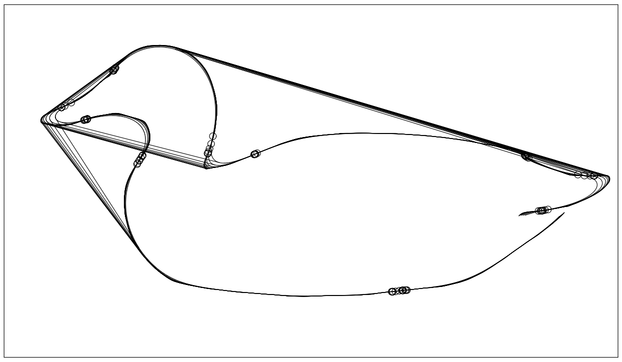

Figure 3. The edges of a duck image smoothed with a series of Gaussian filters and inflections and bitangents tracked through scale space.

of a calibration fixture rotating by a known amount. The origin of the fixture's frame sweeps out a circle. The 3D coordinates of the origin are easily determined for each image; a circle is fit, and the axis is then readily determined.

In the experiments, images were acquired using a CCD camera with a $25mm$ lens at a resolution of 480 by 640 pixels, and the edges were found using an implementation of Canny's edge detector [4]. For reliability, only prominent features that can be stably extracted from an image are tracked through the sequence and used in reconstruction. Following Asada and Brady's curvature primal sketch, linked edges are smoothed with a sequence of Gaussian filters with increasing variance [1]. However, like Mackworth and Mokhtarian, the curve is parameterized by arc length $(x(s), y(s))$ at each scale, and the filter is separately applied to each of these functions [13]. To estimate curvature at each scale, a window of constant arc length is moved over the smoothed curve, and a cubic polynomial is fit. The tangent direction and curvature at a point are then determined from the polynomial coefficients. At each scale, all zero-crossings of curvature are marked as inflections. These are tracked though scale space using a greedy algorithm. Only those inflections that are preserved through scale space are retained for reconstruction.

Similarly, the bitangent end-points are tracked through scale space, and only the stable ones are retained. Note that tracking bitangents through scale space is inherently more reliable than tracking inflections since the coordinates of *both* end-points are available. Also, the bitangents at each scale can be efficiently found in time linear in the total number of edge points n in the image. Suppose that the inflection points found above partition the contour into m convex and concave branches such that connecting the two end-points of a branch defines a convex polygon. For each pair of branches, there can be at most two external bitangents and two internal ones. Since each branch defines a convex polygon, these bitangents can be found in linear time using the methods described in Preparata and Hong's convex hull algorithm [17]; thus, the overall time complexity for computing bitangents is $O(nm^2)$. Rather than estimating the tangent direction from the curve parameters at each point, the segment connecting the two end-points yields a more accurate estimate.

To estimate the velocity $\mathbf{v}'$, features are tracked through the image sequence yielding a discrete curve $\tilde{\mathbf{q}}_i = (u_i, v_i)$ where $\tilde{\mathbf{q}}_i$ in image i is observed at turntable angle θ_i. Because images are densely sampled (every $1°$ of turntable rotation), image motion is typically small (< 5 pixels). Additionally, the images are sparsely populated with features. The next location of the feature is predicted by first smoothing the tracked feature curve with an infinite impulse response filter and then extrapolating. A greedy algorithm is used to match the features in the next image. Before estimating image velocity, the curve is smoothed by applying a non-causal filter.

4.1 Results

A preliminary Common Lisp implementation has been applied to a few image sequences, and here we consider an unpainted duck decoy shown in fig. 6 with the results overlaid. The decoy is a rather complicated surface that would be very difficult if not impossible to accurately model with a computer aided design system.

After calibration, 220 images were acquired as the turntable was rotated with $1°$ increments. The scale space method described previously was used to reliably locate features as shown in fig. 3. The detected features were then tracked through the image sequence. Fig. 4 shows the edges and features found in twelve images at twenty degree increments. The trajectory of the tracked bitangent and inflection points are respectively shown in figs. 5.a and 5.b. 84 inflections and 29 bitangents were tracked through the sequence for at least 10 images.

These tracked features are the input to the reconstruction algorithm. Figs. 5.c and 5.e show orthogonal views of the reconstructed bitangent developables from overhead and side views while figs. 5.d and 5.f show the reconstructed parabolic lines from the same viewpoints. Six of the bitangent developables have been dropped from figs. 5.c and 5.e to reduce clutter. From calibration, the transformation between the turntable and camera frames is known for each image; this is used to reproject the reconstructed points onto the images shown in fig. 6.

A few points concerning the reconstruction results are in order. First, without a good deal more experimental work, it is difficult to precisely determine the reconstruction accuracy. It is clear that bitangent developables are more accurately estimated than parabolic lines. This is probably due to the fact that the direction of bitangents are more accurately measured. Additionally, the reconstruction error appears to be on the order of $\pm 1cm$ which for the experimental camera object distance of $170cm$ yields an error of less than 1%. While this may seem small, the size of the duck's bill is about $2cm$, and so this 1% error may still be significant. This accounts for the noisy appearance of the reconstructed curves on the bill.

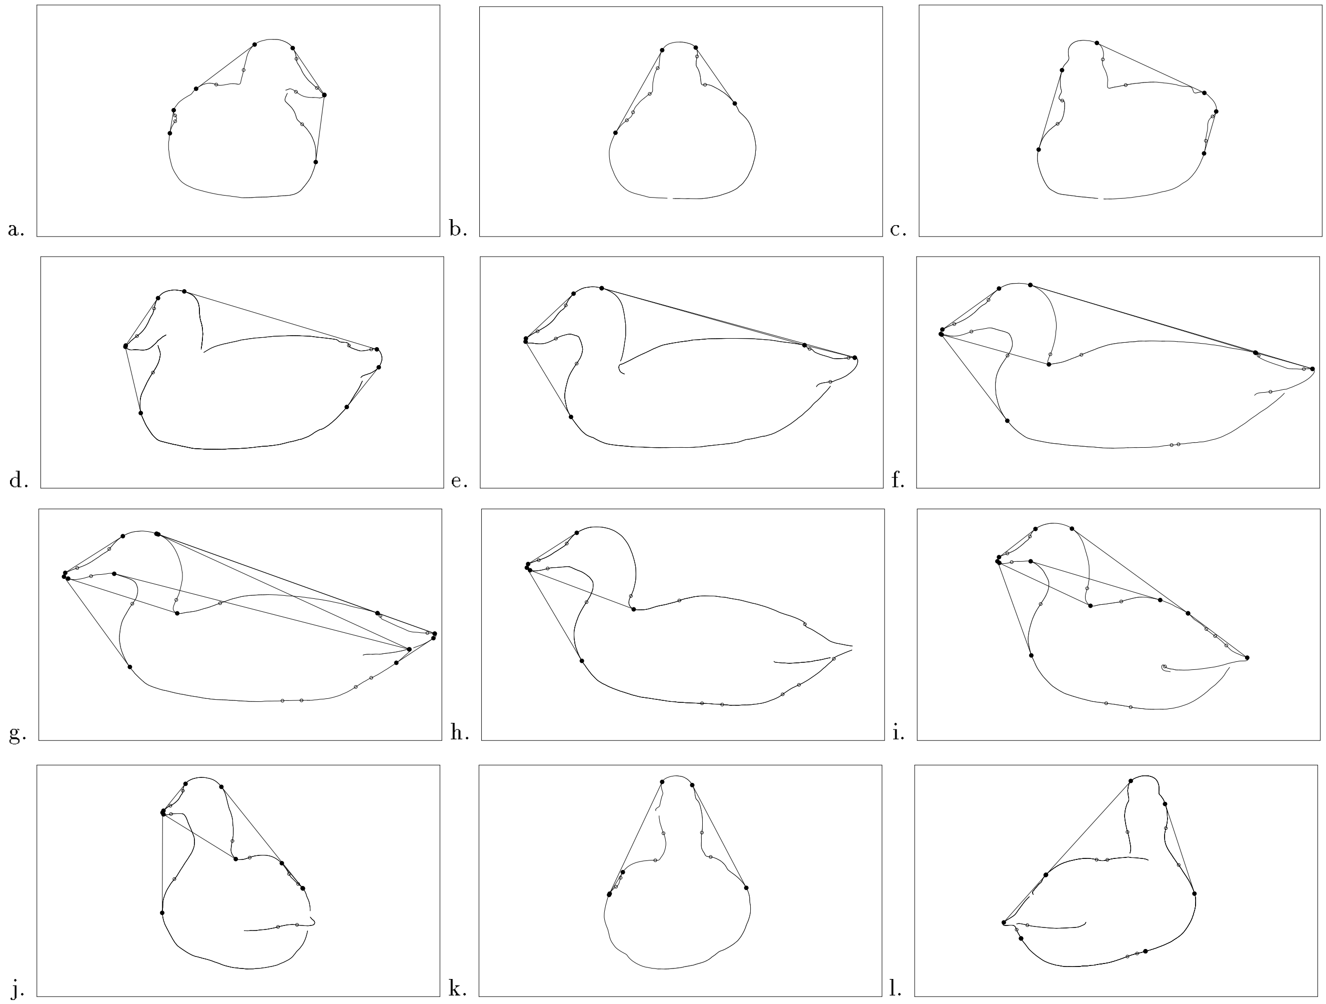

Figure 4. The edges, inflections, and bitangents detected every 20° in a 220° image sequence.

5 Towards recognition

The reconstructed parabolic and limiting bitangent curves are directly useful for pose estimation and consequently object recognition. Many successful approaches for recognition of polyhedral objects establish a correspondence of image features (e.g. corner) to 3D model features (e.g. vertices) which are verified using the so-called rigidity or viewpoint consistency constraints [6, 12]. A direct point-to-point correspondence based solely on feature type cannot be easily made for curved objects since image features are viewpoint dependent. However, a point-to-curve correspondence can often be established. In [11], an approach for pose estimation from a set of viewpoint dependent image features was presented and can be applied here.

The essential observation is that given two points on the surface, the intersection of their tangent planes uniquely determines a viewing direction for which both points will lie on the occluding contour under orthographic projection. For an image formed with this viewing direction, the contour tangent at each point will be in the direction of the intersection of the tangent plane with the image plane. A pair of image points and their tangents define a triangle; the angles between the legs are invariant under rotation, translation, and scaling in the image plane. Thus, a correspondence between two measured features and two surface points can be verified by comparing these angles.

Now, for every pair of points on the discrete reconstructed curves, the pair of invariant angles and the corresponding viewing are computed and stored in a table. Online, a pair of features is extracted from another image, and this table is indexed by the measured angles to find the corresponding pair of 3D curve points and the viewing direction; the rest of the pose parameters are then readily calculated. While this approach has been implemented for algebraic surfaces [11], the reconstructed curves are already in discrete form. Thus, the reconstructed HOT curves can be used directly in this table-based pose estimation scheme. By considering a third feature, the table can be enhanced for recognition.

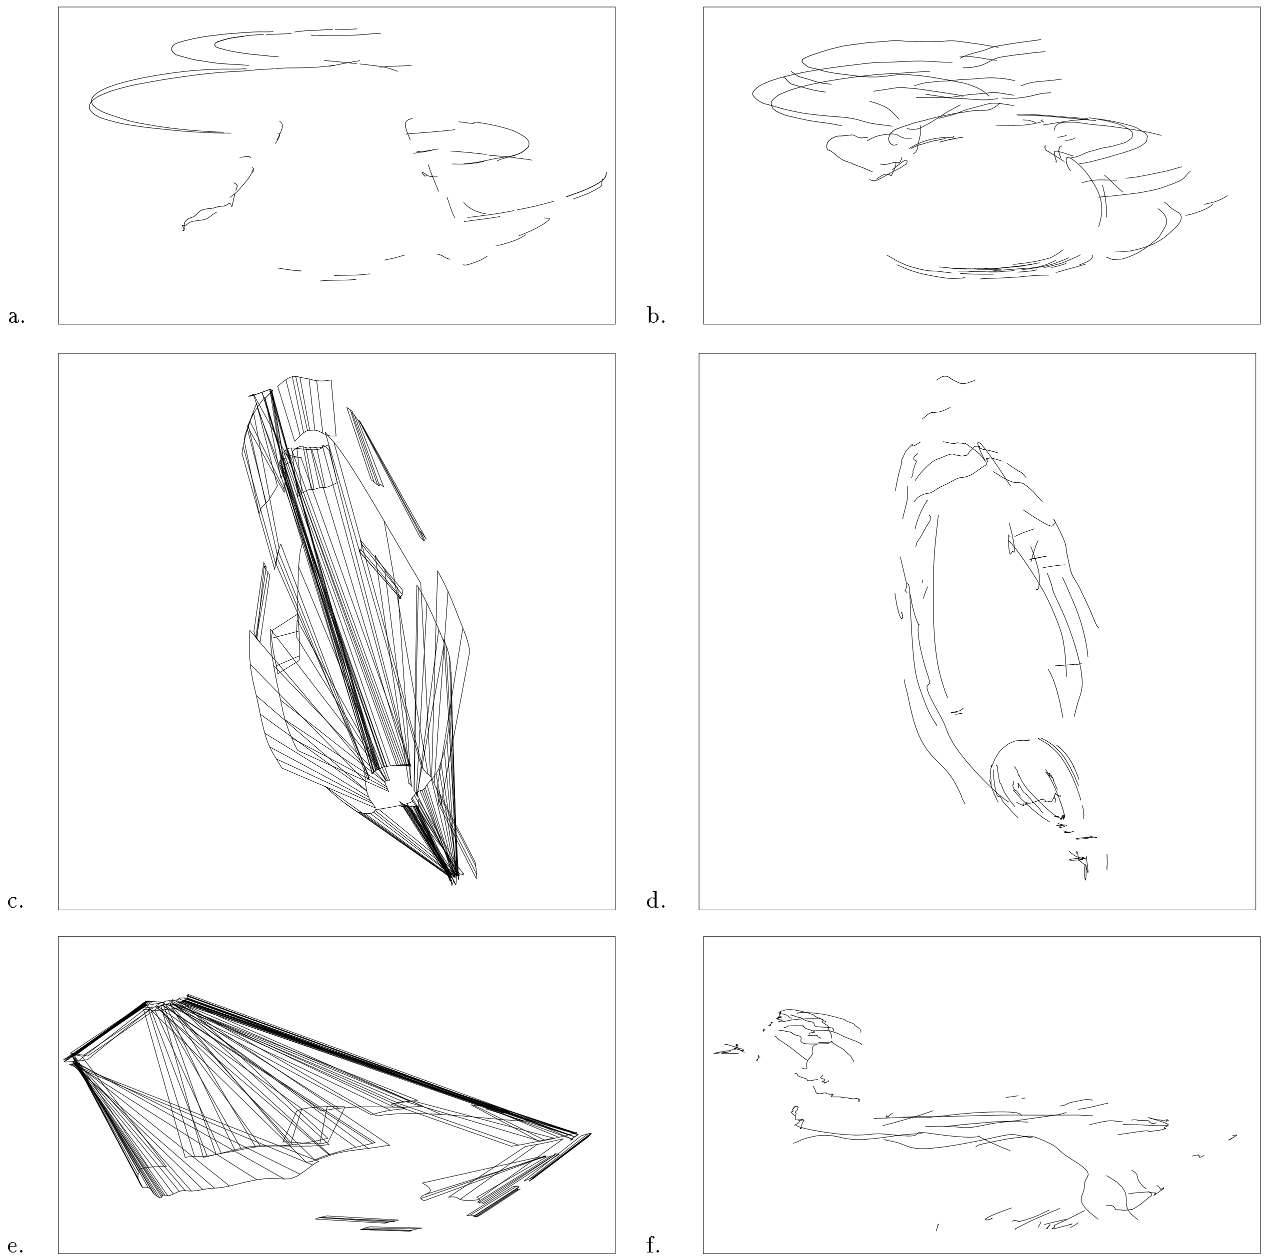

Figure 5. Reconstruction of parabolic and limiting bitangent curves on a mallard decoy: a,b: Image trajectories of bitangents (a) and inflections (b). c,e: Overhead and side views of the reconstructed bitangent developables; the line segments indicate every fifth developable generator. d,f: Overhead and side views of the reconstructed parabolic curves.

6 Conclusions and Discussion

In this paper, an implemented approach was presented for reconstructing two types of 3D HOT curves from a sequence of images, and these curves are useful for object recognition. The reconstruction results for bitangents are particularly encouraging in comparison to those of inflections. They are probably more accurately reconstructed because they are readily located in images and their common tangent is very accurately estimated from the point locations. This also makes bitangents a good feature choice for recognition.

More intriguing is the effect of the choice of camera motions on reconstruction accuracy. From the denominator of (5) the reconstruction error for a fixed feature localization uncertainty is related to the error of the inner product of the normal to the measured tangent plane and the measured feature velocity. It will be minimized when $\mathbf{v}'$ is aligned with the curve normal for a given camera motion. So not all camera motions yield the same accuracy, and this observation can form the basis for planning camera

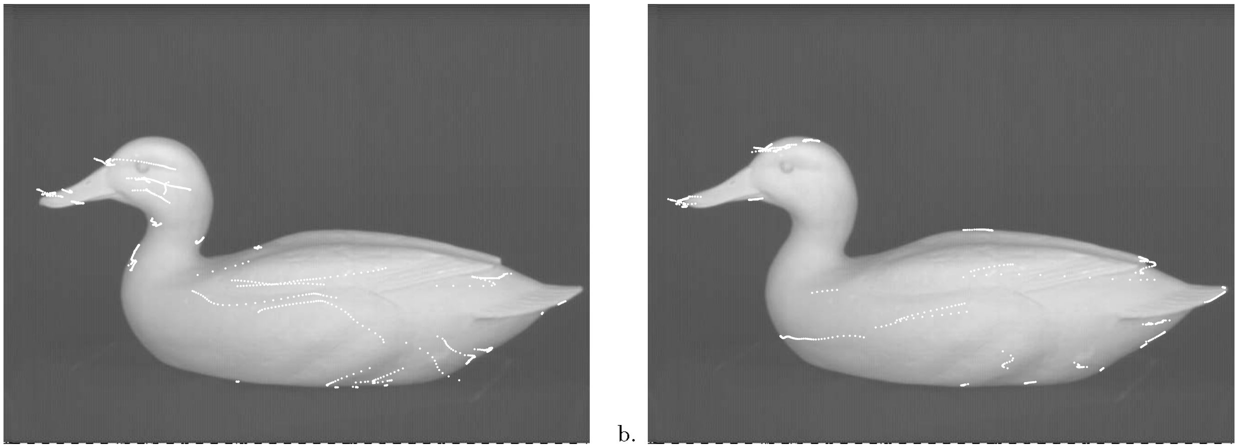

a. b.

Figure 6. Reconstructed parabolic (a) and bitangent (b) curves are reprojected onto the duck image taken at 120°.

motions. Another related topic is determining a sequence of camera motions that will completely map out all visible parabolic and bitangent developable curves. We are also investigating the reconstruction of other HOT curves. Note, for example, that a visual event called a swallowtail transition occurs between the images shown in fig. 4.e and 4.f. The surface points projecting onto these singular points lie on flecnodal curves, one of the other HOT curves.

Acknowledgments Thanks to C.J. Taylor, Steve Sullivan Greg Hager, and Sean Engelson for their help with calibration and pan/tilt heads. Many thanks to Al Harlow, Jr. of Artomation, Inc. for providing the decoy. This work was supported by the National Science Foundation under IRI-9015749, DDM-9112458, and by an NSF Young Investigator Award to D. Kriegman, IRI-9257990.

References

[1] H. Asada and J. Brady. The curvature primal sketch. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 8(1):2–14, 1986.

[2] A. Blake and R. Cipolla. Robust estimation of surface curvature from deformation of apparent contours. In *European Conference on Computer Vision*, April 1990.

[3] J. Bruce and P. Giblin. Generic geometry. *American Mathematical Monthly*, pages 529–545, October 1983.

[4] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intelligence*, pages 679–98, Nov. 1986.

[5] P. Giblin and R. Weiss. Reconstruction of surfaces from profiles. In *International Conference on Computer Vision*, pages 136–144, London, U.K., 1987.

[6] W. E. L. Grimson. *Object Recognition by Computer: The Role of Geometric Constraints*. MIT Press, 1990.

[7] Y. L. Kergosien. La famille des projections orthogonales d'une surface et ses singularités. *C.R. Acad. Sc. Paris*, 292:929–932, 1981.

[8] J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13, 1984.

[9] J. Koenderink. *Solid Shape*. MIT Press, Cambridge, MA, 1990.

[10] D. Kriegman and J. Ponce. On recognizing and positioning curved 3D objects from image contours. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 12(12):1127–1137, 1990.

[11] D. Kriegman, B. Vijayakumar, and J. Ponce. Constraints for recognizing and locating curved 3D objects from monocular image features. In *European Conference on Computer Vision*, 1992.

[12] D. G. Lowe. The viewpoint consistency constraint. *Int. J. Computer Vision*, 1(1), 1987.

[13] A. Mackworth and F. Mokhtarian. The renormalized curvature scale space and the evolution properties of planar curves. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 318–326, 1988.

[14] S. Petitjean, J. Ponce, and D. Kriegman. Computing exact aspect graphs of curved objects: Algebraic surfaces. *Int. J. Computer Vision*, 1992. Submitted. Also Beckman Institute Tech. Report UIUC-BI-AI-RCV-92-02, April 1992.

[15] O. Platonova. Singularities of the mutual disposition of a surface and a line. *Russian Mathematical Surveys*, 36(1):248–249, 1981.

[16] J. Ponce, T. Joshi, D. Kriegman, and B. Vijayakumar. HOT curves for modelling and recognition of aribtrary smooth 3D shapes. Technical report, University of Illinois, Nov. 1992.

[17] F. Preparata and S. Hong. Convex hulls of finite sets of points in two and three dimensions. *Communications of the ACM*, 2(20):87–93, 1977.

[18] G. Taubin, F. Cukierman, S. Sullivan, J. Ponce, and D. Kriegman. Parameterizing and fitting bounded algebraic curves and surfaces. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, 1992.

[19] 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*, 3(4):323–344, Aug. 1987.

[20] R. Vaillant and O. Faugeras. Using extremal boundaries for 3D object modeling. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 14(2):157–173, February 1992.

[21] A. Zisserman, D. Forsyth, J. Mundy, and C. Rothwell. Recognizing general curved objects efficiently. In J. Mundy and A. Zisserman, editors, *Geometric Invariance in Computer Vision*, pages 228–251. MIT Press, Cambridge, Mass., 1992.