

# Duals, Invariants, and the Recognition of Smooth Objects from their Occluding Contour

David Renaudie<sup>1</sup>, David Kriegman<sup>2</sup>, and Jean Ponce<sup>2</sup>

<sup>1</sup> Ecole Nationale Supérieure d'Informatique et Mathématiques Appliquées de  
Grenoble,  
681 Rue de la Passerelle, St Martin d'Hères, France  
David.Renaudie@ensimag.imag.fr

<sup>2</sup> Department of Computer Science  
University Of Illinois  
Urbana, IL 61820 USA  
{kriegman,ponce}@cs.uiuc.edu

**Abstract.** This paper presents a new geometric relation between a solid bounded by a smooth surface and its silhouette in images formed under weak perspective projection. The relation has the potential to be used for recognizing complex 3-D objects from a single image. Objects are modeled by showing them to a camera without any knowledge of their motion. The main idea is to consider the dual of the 3-D surface and the family of dual curves of the silhouettes over all viewing directions. Occluding contours correspond to planar slices of the dual surface. We introduce an affine-invariant representation of this surface that can be constructed from a sequence of images and allows an object to be recognized from arbitrary viewing directions. We illustrate the proposed object representation scheme through synthetic examples and image contours detected in real images.


## 1 Introduction

Most approaches to model-based object recognition are based on establishing correspondences between viewpoint-independent image features and geometric features of object models [9, 15]. For objects with smooth surfaces, few surface markings and little texture, the most reliable image feature is the object's silhouette, i.e., the projection into the image of the curve, called the *occluding contour*, where the cone formed by the optical rays grazes the surface [11]. The dependence of the occluding contour on viewpoint makes the construction of appropriate feature correspondences difficult. Appearance-based methods do not rely on such correspondences, and they are suitable for recognizing objects bounded by smooth surfaces, but they generally require a dense sampling of the pose/illumination space to be effective [17]. Methods for relating image features

to the 3-D geometric models of curved surfaces have been developed for surfaces of revolution [8, 13], classes of generalized cylinders [14, 19, 21, 25], and algebraic surfaces [12, 18, 20, 22]. Two limitations of these approaches are that there must be some means to obtain the 3-D model, and more critically, that only a limited number of objects are well approximated by a single primitive.

An alternative is to replace the explicit description of the entire surface of the object of interest by the representation of some relatively sparse set of features directly useful for recognition. It is possible to represent either the 3-D geometry of these features [10] or to derive invariants from the image coordinates of detected features [24]. In these two methods, objects are modeled from a sequence of images obtained as a camera moves over a trajectory, yet they can be recognized from novel viewpoints; in [10] the camera motion must be known whereas the motion is not needed in [24]. The setting considered in this paper generalizes [24] which only considered a sparse set of features of the silhouette (inflections, bitangents, parallel tangents), and therefore offered limited discriminatory power. In contrast, the proposed approach uses nearly the entire silhouette for recognition. This study builds on geometric insights about the occluding contour and silhouettes of smooth surfaces [7, 11], and their use in determining structure from sequences of images [1, 2, 4, 5, 23].


The basic processing steps for each image include detecting the silhouette curve, computing its dual, and then computing an *HD-curve* (high-dimensional curve) which is invariant to rigid transformations. When modeling an object from a camera moving over a trajectory of viewpoints, these HD-curves sweep out a surface (an HD-surface) which could have been computed directly from the object's 3-D geometry if that geometry were available. Each object is then represented by an HD-surface. During recognition in a single image from a novel viewpoint, an HD-curve can be computed, and this HD-curve should lie on the HD-surface of the corresponding object. The relation of these geometric entities are shown in Figure 1, and we now define these entities and discuss their properties in the subsequent sections.



**Fig. 1.** This diagram summarizes the relation of the original curves, surfaces, pedal curves, pedal surfaces, HD curves and the HD surfaces.

## 2 Duals of curves and surfaces

We start the development by defining some standard geometric concepts about curves, surfaces and duals, which can be found in [3, 6].



**Fig. 2.** Pedal curve and  $\delta(s)$  construction.

Consider a planar curve  $\alpha : I \subset \mathbb{R}^2 \rightarrow \mathbb{R}^2$  parameterized by its arc length  $s$ . At each point of the curve,  $T(s) = \alpha'(s)$  and  $N(s)$  will denote the (unit) tangent and normal vectors of the curve.

## 2.1 Dual and Pedal Curves

The dual of a point on a curve is the tangent line to the curve at that point, and over the entire curve  $\alpha$ , the dual is also a curve.

**Definition 1 (Dual curve).**

$$DCurve(\alpha) = \{(N(s), -N(s) \cdot \alpha(s)), s \in I\} \quad (1)$$

Since  $N(s)$  is a unit vector, a point on a dual curve lies on a cylinder. Because visualizing dual curves (and later dual surfaces) is sometimes difficult, we consider the *pedal curve* which is embedded in  $\mathbb{R}^2$  (respectively  $\mathbb{R}^3$ ) rather than on a cylinder.


**Definition 2 (Pedal curve – P-Curve).**

$$PedalCurve(\alpha) = \{\delta(s), s \in I\} \quad (2)$$

where  $\delta : I \rightarrow \mathbb{R}^2$  is defined by

$$\delta(s) = (N(s) \cdot \alpha(s))N(s). \quad (3)$$

In other words, the pedal curve is the set of points swept by the tip of the unit normal scaled by the (signed) distance between the origin and a curve point as this point varies across the curve. Figure 2 shows how  $\delta(s)$  is defined, and Figure 3 shows an example of a closed planar curve and its pedal curve. A disadvantage of considering pedal curves over the duals themselves is that pedal



**Fig. 3.** Initial curve and its corresponding Pedal curve.

curves are “ramified at the origin”, that is a circle of tangent lines to  $\alpha(s)$  passing through the origin map to the origin of the pedal curve [3].

The following useful properties of pedal curves are readily derived from the definitions.

*Property 1.* There is a one-to-one mapping between the dual and the pedal curve.

*Property 2.* An inflection of the initial curve  $\alpha$  maps onto a cusp of the pedal curve.

*Property 3.* A set of  $p$  points of the initial curve  $\alpha$  with a common tangent line maps onto a point of multiplicity  $p$  of the pedal curve.

For example, bitangent lines on a curve map onto crossings of the pedal curve.

*Property 4.* A set of  $p$  points of the initial curve with the same tangent direction maps onto collinear points of the pedal curve, and these points are collinear with the origin. Reciprocally, the intersection points between a line passing through the origin and the pedal curve correspond to parallel tangent lines of the initial curve.

Consider for example the curve and its pedal in Figure 3; the cusps in the pedal correspond to the two inflections while the crossing corresponds to the common horizontal tangent line at the top of the heart.

It is important to note that the pedal curve depends on the choice of the coordinate system in which the initial curve  $\alpha$  is described. This is significant because the curves we will consider are image contours, and a rigid image-plane transformation of the object will lead to a geometric change to the pedal curve. Nonetheless, Properties 1 to 4 are independent of rigid transformations, and we will subsequently define a mapping from the pedal curve to another curve which is invariant to rigid transformations of the original curve  $\alpha$ .

## 2.2 The Pedal Surface

The dual of a point on a surface is a representation of the tangent plane, and the concept of pedal curve can be extended to non-singular surfaces in  $\mathbb{R}^3$ . Let  $\sigma : U \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$  be a parameterization of a surface  $S$  defined by  $\sigma(s, t)$ . We define the 3D *P-surface* of  $S$  as follows.

**Definition 3 (Pedal Surface – P-Surface).** *The pedal surface associated with the parameterized surface  $\sigma : U \rightarrow \mathbb{R}^3$  is the parameterized surface  $\delta : U \rightarrow \mathbb{R}^3$  defined by  $\delta(s, t) = (N(s, t) \cdot \sigma(s, t))N(s, t)$ .*

## 3 Occluding Contours and Silhouettes

We consider an object bounded by a non-singular surface  $S$  with parameterization  $\sigma$ , and assume that images are formed under scaled orthographic projection (weak perspective).

Given a point  $P \in \mathbb{R}^3$  and a camera whose origin is in  $O$  and whose image plane is spanned by the orthogonal vectors  $\mathbf{i}$  and  $\mathbf{j}$ , the coordinates  $(X, Y)$  of the projected point are given by:

$$\begin{cases} X = \chi \overrightarrow{OP} \cdot \mathbf{i} \\ Y = \chi \overrightarrow{OP} \cdot \mathbf{j} \end{cases} \quad (4)$$

where  $(O, \mathbf{i}, \mathbf{j})$  is the camera's coordinate frame, and  $\chi$  is the inverse of the depth of some reference point. The vector  $\mathbf{v} = \mathbf{i} \times \mathbf{j}$  is the *viewing direction*. For *pure orthographic projection*,  $\chi$  is taken to be constant and without loss of generality we shall choose  $\chi = 1$ .

Let  $v$  be a viewing direction in  $\mathbb{R}^3$ , the occluding contour is the set of points in  $\mathbb{R}^3$  that lie on the surface  $S$  and whose normal vectors are orthogonal to the viewing direction. The occluding contour is generally a nonplanar curve on  $S$ , and obviously depends upon the viewpoint [11]. The silhouette is the projection of the occluding contour into the image. Note that this definition treats the objects as being translucent and does not account for self-occlusion of sections of the occluding contour by other portions of the object.

## 4 Links Between P-curves and P-surfaces

We now consider the relation between pedal curves of the silhouettes for some viewing direction and the pedal surface. We denote by  $v$  the viewing direction, and by *OccCont* the corresponding occluding contour.

*Property 5.* The occluding contour maps onto a curve given by the intersection of the pedal surface with the plane  $v^\perp$  passing through the origin and perpendicular to the viewing direction  $v$ .

This property leads us to consider the links between the P-curve of a silhouette and some planar slice of the P-surface. Let us consider our surface  $S$  and its corresponding P-surface  $S'$ . Now choose some viewing direction  $v$ ; it defines an occluding contour  $OccCont$  and a silhouette  $Sil$  under pure orthographic projection. Let  $A'$  denote the projection of the origin of the object's coordinate system used to define  $S'$ . We have the following theorem:

**Theorem 1.** *The P-curve of the silhouette  $Sil$  with respect to the point  $A'$  has the same shape as the intersection (planar slice) of the P-surface  $S'$  with the plane  $v^\perp$ .*

*Proof.* Property 5 says that the P-surface of the occluding contour is a plane curve, entirely embedded in the affine plane  $v^\perp$ . Since the silhouette is obtained by pure orthographic projection of the occluding contour onto the image plane, the tangent planes to  $S$  at the points of the contour are projected into the tangent lines of the silhouette. Thus computing the P-curve of the silhouette with respect to  $A'$  is exactly equivalent to taking the planar intersection of the P-surface with the plane  $v^\perp$ .  $\square$

This property leads to the following critical insight. Consider a camera moving over a trajectory during modeling such that the viewing direction  $v(t)$  is known. Then from the previous property, one could reconstruct a subset of the object's P-surface from the P-curves if  $A'(t)$  can be determined. If for example,  $v(t)$  covers a great circle, then the entire P-surface would be reconstructed. Furthermore, for some other viewing direction  $\tilde{v}$  not in  $v(t)$ , the corresponding P-curve would simply be a planar slice of the reconstructed P-surface.

#### 4.1 A Representation of the Pedal Curve and the Pedal Surface

The previous properties could form the basis for a recognition system in which smooth surfaces are modeled from a sequence of images with known camera motion, and then objects could be recognized from an arbitrary viewing directions  $\tilde{v}$ . However, this requires establishing correspondence of  $A'(t)$  during modeling and detecting the projection of the 3-D origin  $\tilde{A}$  during recognition; we know of no properties for doing this. Instead, we now define a representation (HD-curves and HD-surfaces) that is invariant to the choice of the origin (and in fact affine image transformations) and furthermore eliminates the necessity for knowing the camera motion  $v(t)$  during modeling. In particular, we take advantage of the previously mentioned invariance of properties 1 to 4 to rigid plane transformations.

**Definition 4 (Signature).** *Given a planar curve  $\mathcal{PC}$  parameterized by  $\delta : I \rightarrow \mathbb{R}^2$  and a point  $C$  (called the scanning center), we call the signature of the curve  $\mathcal{PC}$  with respect to a point  $C$  the couple  $((\theta_1, \dots, \theta_p), (d_1, \dots, d_p))$  such that:*

$$(1) \sum_{k=1}^{p-1} (\theta_{k+1} - \theta_k) = 2\pi$$

$$(2) \forall k \in 1 \dots p - 1 \quad \forall \theta \in [\theta_k, \theta_{k+1}] \quad \#(\mathcal{D}(C, \theta) \cap \mathcal{PC}) = d_k + 1$$

where  $\mathcal{D}(C, \theta)$  denotes the straight line passing through  $C$  with orientation  $\theta$ , and  $\#E$  denotes the number of elements of the set  $E$ .

The signature partitions the plane into angular sectors centered at  $C$  and bounded by  $\theta_k, \theta_{k+1}$  such that the number of intersections between any line passing through  $C$  and the curve  $\mathcal{PC}$  is constant within each sector. The critical points where the number of intersections changes is given by the singular points of the pedal curve (See Properties 2 and 3) and points where the tangent to pedal curve pass through the  $C$ .

We now define the *HD-curve of a planar curve  $\mathcal{PC}$  with regards to a point  $C$*  as the multi-dimensional set of curves ( $HDComp_1, \dots, HDComp_p$ ) given by:

**Definition 5.**  $\forall k \in 1 \dots p$ ,  $HDComp_k$  is a curve embedded in  $\mathbb{R}^{d_k}$  with coordinates  $(y_1 = x_2 - x_1, \dots, y_{d_k} = x_{d_k+1} - x_1)$  with  $(x_1, \dots, x_{d_k+1})$  being the abscissas of the intersection points between the oriented line  $\mathcal{D}(C, \theta)$  and  $\mathcal{PC}$ .


In other words, in each angular sector  $k$  defined previously, we take the  $d_k + 1$  intersection points between  $\mathcal{PC}$  and lines passing by  $C$ , order these intersection points by relative abscissas on the line, and compute the  $d_k$  differences by the first (thus the smallest) abscissa.

We now consider an important property of the *HD-curve* that makes them useful for modeling and recognition. Consider an initial silhouette curve  $\mathcal{C}$ , parameterized by  $\alpha : I \rightarrow \mathbb{R}^2$ . Now compute its corresponding P-curve  $\mathcal{PC}$ , parameterized by  $\delta : I \rightarrow \mathbb{R}^2$ , constructed relative to the origin of the image plane.

*Property 6.* The HD-curve of  $\mathcal{PC}$  of  $\mathcal{C}$  is invariant with respect to the choice of the origin of  $\mathcal{C}$  and any rigid transformation of  $\mathcal{C}$ .

The important thing here is the fact that the P-curve gathers all parallel tangency and inflection information under an easily exploitable form but its shape depends upon the choice of the origin. By choosing the HD-curve scanning center as being the same point as the P-curve origin, this dependence on the choice of the origin is suppressed. Here is another perspective: consider the silhouette curve  $\mathcal{C}$ ; its tangency features (inflection points, parallel tangents...) are intrinsic properties of the curve. For example, consider a set of four points on a curve with parallel tangents; they will be mapped to four collinear points of the P-curve. Now, consider the relative distances between these four aligned points. They correspond to the relative distances between tangent lines to the initial curve, thus it is not surprising that we succeed in eliminating the dependency with respect to the P-curve center.

In a similar manner, one can define an *HD-surface* derived from the P-surface. Unlike the P-surface whose shape depends upon the choice of the origin, the HD-surface is invariant to rigid transformations of the original surface. Based on Property 1 and its implications as well as the invariance property of HD-curves to rigid transformation, the entire HD-surface can be determined from the sequence of silhouettes formed when the viewing direction covers a great circle. Hence, the HD-surface can serve as a representation for recognition, and it can be constructed without knowledge of the camera motion.



**Fig. 4.** Contour detected by tresholding and the computed pedal curve.


Finally, under weak perspective, the scale  $\chi(t)$  is a function of time. In constructing the HD-curve components, we can treat the  $d_k$  coordinates as homogeneous coordinates or normalize the coordinates as  $(z_1 = y_1/y_{d_k}, \dots, z_{d_k-1} = y_{d_k-1}/y_{d_k})$  with  $(y_1, \dots, y_{d_k})$  being the coordinates of points on  $HDComp_k$  of the HD-curve. This creates a curve embedded in  $\mathbb{R}^{d_k-1}$  rather than in  $\mathbb{R}^{d_k}$ , and we call this the *Quotient HD-curve*. The Quotient HD-curve of  $\mathcal{PC}$  is invariant with respect to any homothetic transformation of the initial silhouette curve  $\mathcal{C}$  in the image plane. Over a sequence of images, the family of Quotient HD-curves sweeps out a *Quotient HD-surface*.

## 5 Implementation

Here, we present some details of our Matlab implementation for computing the geometric entities described previously.

### 5.1 Image Processing

For each image of the sequence, we detect the silhouette of the object using thresholding to obtain a connected curve, i.e. a discrete list of all successive points on the contour.



**Fig. 5.** A contour after recursively applying the Gaussian filter with  $\sigma = 2.0$  five times and the computed pedal curve.



The discrete contour extracted by thresholding introduces artifacts due to aliasing, leading to poor results when used in the subsequent steps of the algorithm. Consequently, we recursively smooth the contour using Gaussian filters parameterized by arc length [16]. First, the arc length to each point on the contour is computed. For each point  $\alpha(t_0)$  on the original contour, we obtain new  $x$  and  $y$  coordinates according to:

$$\forall t_0 \in 0 \dots t_{max} \quad Smoothed.\alpha(t_0) = \frac{\int_{t_0-4\sigma}^{t_0+4\sigma} G(t_0-t)\alpha(t)dt}{\int_{t_0-4\sigma}^{t_0+4\sigma} G(t_0-t)dt} \quad (5)$$

where

$$G(t) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{t^2}{2\sigma^2}\right) \quad (6)$$

Since the curve is discrete  $\alpha(t)_{t \in 0 \dots t_{max}}$ , we approximate the above integral using a trapezoid method with irregularly spaced interpolation points.

*Example 1.* Figure 4 shows the initial contour and its corresponding P-curve, and Fig. 5 shows the results after five consecutive smoothings.

## 5.2 Pedal Curve Computation

Given the discrete form of the possibly smoothed curve, we compute the normal vector for each point on the contour. For this, we use linear least squares to estimate of the direction of the tangent. Then Equation 3 can be directly applied to compute the pedal curve; this phase is very quick once the normals are available. The only difference with Eq. 3 is that the contour is parameterized by a discrete parameter instead of a continuous one.

## 5.3 HD-Curve Computation

The method for computing the HD-Curve follows:

1. An oriented line constitutes the reference line of the angles. For reasons described in Section 4.1, we choose for scanning center the same point that was used for computing the P-Curve, and in practice, the origin of the angles is the horizontal, right-oriented-line passing through the scanning center. Changing the origin of the angles only cycles the storage of the points of the HD-Curve, but it does not affect it otherwise.
2. We regularly sample  $[0, 2\pi]$  and obtain the angles  $\{\theta_i\}_{i \in 1 \dots N}$
3. For each of the sampled angles  $\theta_i$ :
  - Consider the line passing through the scanning center and having an angle  $\theta_i$  with the reference line, counted anti-clockwise.
  - Calculate the signed abscissas (regarding this line) of the intersection points of this line and the P-Curve, and sort them.
  - Calculate the distances between these intersection points, i.e. the differences between the previously calculated abscissas and the first abscissa (the smallest).
  - Store these distances.




Fig. 6. A heart, a squash and a banana.





Fig. 7. A single image of the squash and its corresponding HD-Curve and Quotient HD-Curve.

## 6 Results

We demonstrate these algorithms with three series of images: a heart, a butternut squash and a banana (Fig. 6), rotating about a vertical axis. While the heart images are synthetic, please note that the squash and banana images have been gathered with a camera and *real* fruits and vegetables.

### 6.1 An Example of HD-Curve and Quotient HD-Curve

We consider a single image of the squash, and the HD-Curve and Quotient HD-Curve that have been extracted from this image. More precisely, Fig. 7 shows only the 3D part of the HD-Curve and the 2D part of the Quotient HD-Curve,



**Fig. 8.** Images of the heart.

since they are more easily representable than higher dimensional curves. Note that in this particular case, the HD-Curve had only a 1D and a 3D components, but the banana, which is a more “complex” object, has a 5D component that we cannot readily draw.

Note that for any translation and/or rotation of the initial image, the HD-Curve (thus also the Quotient HD-Curve) remain unchanged, and any zoom or definition change leaves the Quotient HD-Curve unchanged.

## 6.2 HD-Surfaces

Below we show the result of computing the HD-surfaces for the heart, squash and banana. Please note that we have drawn sample points on the HD-surfaces rather than a rendering of an interpolated surface.

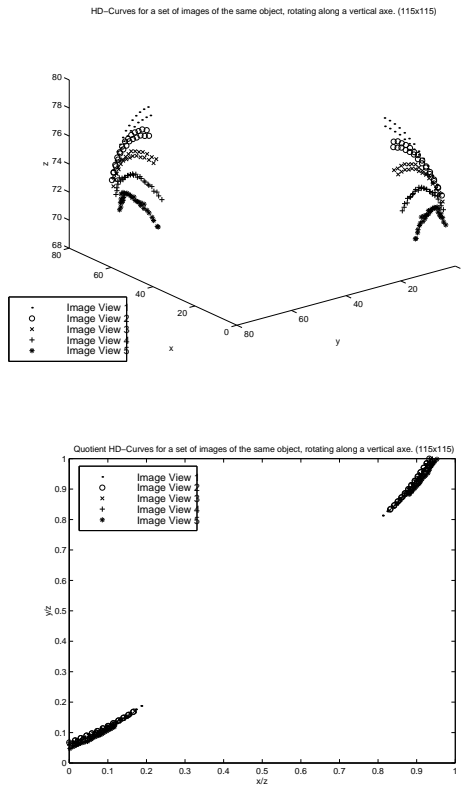
**The Heart Sequence** Figure 8 shows a series of five artificially generated images (115x115) of the heart, rotating along a vertical axis, and Figure 9 shows the corresponding computed HD-Surface (3D part only) and the Quotient HD-Surface (2D part only).

**The Squash Sequence.** Figure 10 shows a series of five *real* images (300x300) of a real squash, and Figure 11 shows the corresponding HD-surfaces.

**The Banana Sequence** Finally, Figure 12 shows a series of four images of a banana, and Figure 13 shows the corresponding HD-surfaces. Please note that the angle of rotation over the four images is significant (over 120 degrees), and so four images yields a very sparse sampling of the HD-surfaces.

## 7 Conclusion

In this paper, we have introduced a new relation between the dual of a smooth surface and the dual of the silhouettes formed under orthographic projection. While the dual (or pedal) curve/surface depends upon the choice of the origin, the HD-curves (HD-surfaces) are invariant to rigid transformations. We have




**Fig. 9.** Sample points on HD-Surface and Quotient HD-surface computed from five images of the heart.


illustrated how these geometric entities can be computed from synthetic and real images, and have discussed how they could be used for object recognition. Unlike nearly all methods for recognizing smooth curved 3-D objects, this technique does not assume that objects come from a limited class such as surfaces of revolution, generalized cylinders or algebraic surfaces.

Note that we have not yet incorporated this representation within an object recognition system. An important issue to be addressed will be the combinatorics of determining whether a measured HD-curve lies on a model HD-surface when there is occlusion during either the modeling or recognition phase.

The basis for this method is that the set of points on an object's surface with parallel tangent planes project under orthographic projection to image curve points with parallel tangent lines. Between a test image and each model image, the HD-curves and HD-surfaces are essentially being used to identify candidate stereo frontier points [7]. This paper generalizes our earlier results [24] which defined invariants from silhouette tangent lines that were parallel to the tangent lines at inflections or bitangents. It turns out that these features are the



**Fig. 10.** Five images of the squash.



**Fig. 11.** Squash's HD-Surface (3D part) and Quotient HD-Surface (2D part).


singularities (cusps and crossings) of the dual of the silhouette. In [24], this lead to representing an object by a set of “invariant curves” while in this paper, an object is represented by a set of “invariant surfaces,” namely the HD-surfaces. Consequently, the presented representation retains much more information about the object’s shape and should provide greater discriminatory power than the curves used in [24].

### Acknowledgments

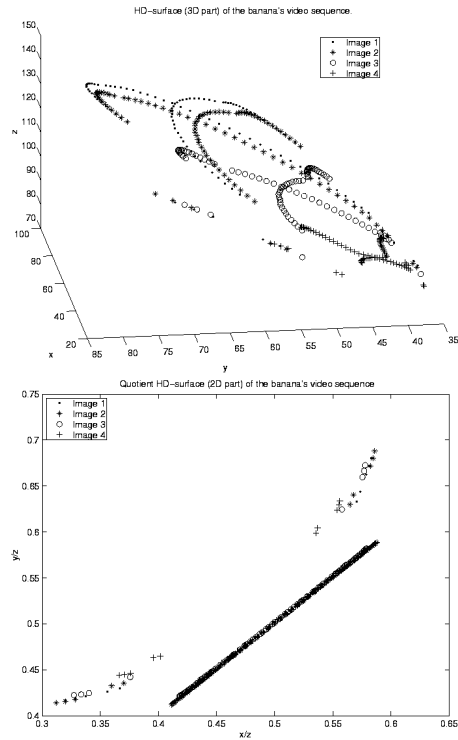
This work was supported by ARO DAAG55-98-1-0168.

### References

- [1] E. Arbogast and R. Mohr. 3D structure inference from image sequences. *Journal of Pattern Recognition and Artificial Intelligence*, 5(5), 1991.



**Fig. 12.** Four images of the banana.



**Fig. 13.** The computed HD-surface and Quotient HD-surface from the four images of the Banana.

- [2] E. Boyer and M. Berger. 3d surface reconstruction using occluding contours. *Int. J. Computer Vision*, 22(3):219–233, 1997.
- [3] J. Bruce and P. Giblin. *Curves and Singularities*. Cambridge University Press, 1992.
- [4] R. Cipolla, K. Aström, and P. Giblin. Motion from the frontier of curved surfaces. In *Int. Conf. on Computer Vision*, pages 269–275, 1995.
- [5] R. Cipolla and A. Blake. Surface shape from the deformation of the apparent contour. *Int. J. Computer Vision*, 9(2):83–112, November 1992.
- [6] M. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, N.J., 1976.

- [7] P. Giblin and R. Weiss. Epipolar curves on surfaces. *Image and Vision Computing*, 13(1):33–44, 1995.
- [8] R. Glachet, M. Dhome, and J. Lapersté. Finding the perspective projection of an axis of revolution. *Pattern Recognition Letters*, 12:693–700, 1991.
- [9] D. Huttenlocher and S. Ullman. Recognizing solid objects by alignment with an image. *IJCV*, 5(2):195–212, November 1990.
- [10] T. Joshi, B. Vijayakumar, D. Kriegman, and J. Ponce. HOT curves for modelling and recognition of smooth curved 3D shapes. *Image and Vision Computing*, pages 479–498, July 1997.
- [11] J. J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13, 1984.
- [12] 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.
- [13] D. J. Kriegman and J. Ponce. Computing exact aspect graphs of curved objects: Solids of revolution. *Int. J. Computer Vision*, 5(2):119–135, 1990.
- [14] J. Liu, J. Mundy, D. Forsyth, A. Zisserman, and C. Rothwell. Efficient recognition of rotationally symmetric surfaces and straight homogeneous generalized cylinders. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 123–128, 1993.
- [15] D. G. Lowe. The viewpoint consistency constraint. *Int. J. Computer Vision*, 1(1):57–72, 1987.
- [16] 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.
- [17] H. Murase and S. Nayar. Visual learning and recognition of 3-D objects from appearance. *Int. J. Computer Vision*, 14(1):5–24, 1995.
- [18] S. Petitjean, J. Ponce, and D. Kriegman. Computing exact aspect graphs of curved objects: Algebraic surfaces. *Int. J. Computer Vision*, 9(3):231–255, 1992.
- [19] J. Ponce and D. Chelberg. Finding the limbs and cusps of generalized cylinders. *Int. J. Computer Vision*, 1(3), October 1987.
- [20] J. Ponce, A. Hoogs, and D. Kriegman. On using CAD models to compute the pose of curved 3D objects. *CVGIP: Image Understanding*, 55(2):184–197, Mar. 1992.
- [21] M. Richetin, M. Dhome, J. Lapresté, and G. Rives. Inverse perspective transform from zero-curvature curve points: Application to the localization of some generalized cylinders from a single view. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 13(2):185–191, February 1991.
- [22] J. Rieger. Global bifurcations sets and stable projections of non-singular algebraic surfaces. *Int. J. Computer Vision*, 7(3):171–194, 1992.
- [23] R. Vaillant and O. Faugeras. Using extremal boundaries for 3D object modeling. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 14(2):157–173, February 1992.
- [24] B. Vijayakumar, D. Kriegman, and J. Ponce. Invariant-based recognition of complex curved 3-D objects from image contours. *Computer Vision and Image Understanding*, pages 287–303, Dec. 1998.
- [25] M. Zeroug and G. Medioni. The challenge of generic object recognition. In M. Hebert, J. Ponce, T. Boult, and A. Gross, editors, *Object Representation for Computer Vision*, pages 271–232. Springer-Verlag, 1995.