

Recovering Shape by Purposive Viewpoint Adjustment

Kiriakos N. Kutulakos

Charles R. Dyer

Computer Sciences Department
University of Wisconsin
Madison, Wisconsin 53706 USA

Abstract

We present an approach for recovering surface shape from the occluding contour using an active (i.e., moving) observer. It is based on a relation between the geometries of a surface in a scene and its occluding contour: If the viewing direction of the observer is along a principal direction for a surface point whose projection is on the contour, surface shape (i.e., curvature) at the surface point can be recovered from the contour. Unlike previous approaches for recovering shape from the occluding contour, we use an observer that purposefully changes viewpoint in order to achieve a well-defined geometric relationship with respect to a 3D shape prior to its recognition. We show that there is a simple and efficient viewing strategy that allows the observer to align the viewing direction with one of the two principal directions for a point on the surface. Experimental results demonstrate that our method can be easily implemented and can provide reliable shape information.

1 Introduction

There has been considerable interest in recovering information about the structure of a scene from sequences of images, assuming an observer in motion (e.g., work on optical flow and shape-from-motion). One common feature of these approaches is the use of known viewer motion in order to recover quantitative properties of the scene such as surface curvature [1]. Recently, however, there has been considerable interest in employing simple observer behaviors that either make the recovery of scene properties easier [2], or combine simple behaviors in order to perform complex tasks such as navigation and obstacle avoidance [3]. These approaches rely on maintaining specific geometric relationships between the observer and the environment. However, without knowledge of viewer motion they only recover qualitative information about the viewed scene (e.g., relative depth ordering with respect to the fixated scene point [2]).

This paper presents a new approach that combines the above two paradigms in order to recover surface curvature

from the occluding contour. We show that shape for a selected point on the viewed surface can be recovered from the occluding contour of two views. The only requirements are that (1) the surface point is projected on the occluding contour in both views, and (2) the viewing direction of the observer for one of the views has a specific relationship with the surface geometry at the selected point. The main idea of our approach is to use an active (i.e., mobile) observer that purposefully changes viewpoint in order to achieve such a well-defined geometric relationship with respect to a 3D shape prior to its recognition. We show that this relationship is characterized by specific image-computable quantities and enables an analysis similar to the one of Krotkov [4]. In addition, our approach does not require any knowledge of observer velocities or object models and assumes the use of a world-centered coordinate frame.

It is well-known that the occluding contour is a valuable source of information about surface shape [5, 6]. The occluding contour is the projection of the visible rim, the one-dimensional set of points that separates the visible from the hidden parts of a surface. There have been several approaches to deriving information about surface geometry from the occluding contour. These approaches are based on three important properties of the contour's geometry: (1) The geometry of the occluding contour is surface-dependent, (2) the occluding contour is the projection of a limited set of surface points, and (3) the geometry of the occluding contour is viewpoint-dependent.

The dependency of the occluding contour's geometry on the surface severely constrains the surface's geometry (e.g., [5, 6]). The major problem, though, is that the occluding contour is created through a projection process and in general the information it provides is ambiguous, i.e., several different surface rims can project to the same occluding contour. This has been the major motivation for using the occluding contour to derive qualitative rather than quantitative shape information [6]. The second property of the occluding contour suggests that it can provide only a limited amount of information about the complete shape of the surface. Indeed, only a one-dimensional set of surface points project to the occluding contour. The contour's dependency on viewpoint has been used to resolve both of

The support of the National Science Foundation under Grant No. IRI-9002582 is gratefully acknowledged.

the above ambiguities. A slight change in viewpoint will affect the geometry (i.e., curvature) and possibly the topology of the rim, and hence the occluding contour. Moreover, the set of rim points changes and therefore new constraining information about the surface shape becomes available. Giblin and Weiss [5] showed that if we know how the geometry of the occluding contour changes with viewpoint, we can derive a parameterization of the surface and hence determine its shape. The issue here is how to accurately measure such changes in the contour’s geometry with small viewing direction changes. For example, we must be able to measure the velocity and acceleration of surface points entering and leaving the rim [1, 7], a problem that requires first- and second-order differentiation operations of image-computable quantities, and hence is sensitive to noise.

The basic assumption used by all of the above approaches was that the viewing direction is *arbitrary*. This means that the viewing direction is not related in any way to the geometry of the rim or the occluding contour. (For example, there are surfaces for which their rim is planar when viewed from a particular set of directions.) This is a reasonable approach, however, only if the observer cannot control the direction of gaze. When the observer has the ability to control the viewing direction, the choice of viewpoint(s) does not have to be arbitrary. Our approach uses an active observer to obtain a view based on the observed object’s geometry in order to recover exact shape information from the occluding contour.

1.1 Active shape recovery

Our goal is to actively derive a quantitative shape description for surface points in the vicinity of the rim. We accomplish this goal by using properties of the occluding contour. The basic step of our approach involves selecting a point on the rim and recovering the surface shape (i.e., principal curvatures and principal directions) at that point. In addition, we present a strategy for applying this shape recovery step to neighboring surface points. The surface description is therefore incrementally extended by successively including new points on the rim and recovering the surface geometry for those points.

The main step of our approach is based on a relation between the geometries of a surface in a scene and its occluding contour: If the viewing direction of the observer is along a principal direction for a selected surface point whose projection is on the contour, the corresponding principal curvature at the point can be recovered. Hence, even though in general surface curvature calculation from the occluding contour of a single view is an underconstrained problem, for any given point there do exist viewing directions that make this recovery problem well-defined. If the observer can move to one of those special viewing directions, the ambiguities caused by the projection process can be resolved. We show that the observer can in fact deterministically find these special viewing directions by simply maximizing or minimizing a geometric quantity of the oc-

cluding contour (curvature at a point) while changing viewing direction in a constrained way. Furthermore, we show that we can recover the shape of the surface at the selected point (i.e., both principal curvatures) from the occluding contour of one additional view for which the selected point is projected onto the contour. Thus an active observer selects a point on the surface rim and purposefully moves to one of the special viewpoints in order to make shape recovery a well-defined problem.

The significance of our method lies in the use of purposive observer motion to achieve and maintain purely geometrical relations between a surface and its occluding contour in order to recover surface shape. Hence, there is no need to perform any velocity or acceleration measurements for image points in the vicinity of the rim, a process requiring point-to-point correspondences and precise knowledge of viewer motion. Furthermore, there is a well-defined procedure to reach the desired viewpoint, so the observer does not need to perform a complicated search to find it.

Even though our approach is limited to the recovery of surface shape in the vicinity of a single point on the rim, we show that there is an important special case for surfaces of revolution, for which we can derive shape information for the complete set of rim points. In this case the observer is actively “aligned” with the viewed surface in order to find a viewpoint giving complete surface information (i.e., one perpendicular to the axis of rotation). Furthermore, we present an extension to the above approach that recovers the shape of points in the vicinity of the rim. After the shape of a selected rim point is recovered, the observer changes viewpoint in order to bring a new surface point onto the rim and to recover its shape. This is a major difference from approaches using “passive” motion, where the points selected for reconstruction cannot be controlled.

The rest of this paper is organized as follows. The next section reviews basic terminology. Section 3 discusses the relation between the geometries of the occluding contour and the surface, and presents the major result enabling us to actively recover surface geometry from the occluding contour. Section 4 uses this result to describe the main shape recovery step of our approach. Our results are then extended in the following two sections. Section 5 discusses shape recovery for the case of surfaces of revolution, and Section 6 outlines the viewing strategy used to select a new point for shape recovery. Finally, Section 7 presents experimental results on synthetic images to demonstrate the applicability of our theoretical results.

2 Viewing geometry

Let S be a smooth, oriented surface in \mathfrak{R}^3 , viewed under orthographic projection along a viewing direction ξ . Given a parametrization x of S and a point $p = x(u, v)$, the partial derivatives $x_u(p), x_v(p)$ of x with respect to u and v define $T_p(S)$, the plane tangent to S at p . Local surface shape (i.e., curvature) is completely expressed by the first and second fundamental forms of S with respect

to x [8]. Specifically, let $N(p)$ be the surface normal at p . The normal section of S along a direction ξ in $T_p(S)$ is the plane curve produced by intersecting S with the plane of ξ and $N(p)$. The second fundamental form, $II(p)$, gives an expression for the curvature of this curve at p . $II(p)$ has a single maximum and minimum, k_{n_1} and k_{n_2} , along the principal directions e_1 and e_2 , respectively. We can use these two quantities, called the principal curvatures of S at p , to compute the curvature of the normal sections along any other direction using Euler’s formula:

$$k_n(\phi) = k_{n_1} \cos^2 \phi + k_{n_2} \sin^2 \phi \quad (1)$$

where ϕ is the angle between the new direction and e_1 . Hence, we can recover the local shape of S completely from the principal curvatures of S . In the following we assume that p is not an umbilic point (i.e. a point where all directions are principal directions), and consider the parametrization $x(u, v)$ for which the curves $x(u, v_0)$ and $x(u_0, v)$ (u_0, v_0 constant) are *lines of curvature*, i.e., have tangents along the principal directions.

3 Local shape from occluding contour

Under orthographic projection, the *rim* of S is the set of those points p for which $T_p(S)$ contains a line parallel to the viewing direction ξ . The *occluding contour* of S is the projection of the visible rim on the image plane. The problem of recovering surface geometry from the occluding contour has been mainly studied under the assumption that the viewing direction is *arbitrary*. This means that the viewing direction is not related in any way to the geometry of the rim or the occluding contour. Under this assumption there are three main results describing what can be recovered from the shape of the occluding contour when orthographic projection is assumed: (1) We can recover the surface orientation at p (i.e., $T_p(S)$ or $N(p)$) from ξ and the tangent to the occluding contour at q , the projection of p [6], (2) if k_o is the curvature of the occluding contour at q , then k_o and the Gaussian curvature $K = k_{n_1} k_{n_2}$ of S at p have the same sign [6], and (3) if k_n is the normal curvature of S at p along ξ , then $K = k_n k_o$ [5, 6]. Because K is defined as the product of two curvatures on the surface (i.e., k_{n_1}, k_{n_2}), these results suggest that if we know k_o then we only need to measure one curvature on the surface instead of two. In fact k_n and k_o determine the second fundamental form at p . This was the main idea behind the surface reconstruction approach of Cipolla and Blake [1].

The above results are important but they also imply that if we have no additional information about the shape of the viewed surface, the information provided by the occluding contour is primarily qualitative. However, when the observer can actively control the viewing direction, we can exploit the existence of directions that allow the derivation of complete information about the surface. We show this by presenting three simple corollaries to a result of Blaschke (proofs can be found in [9]). Blaschke’s result is

analogous to Euler’s formula and relates the curvature of the occluding contour with the principal curvatures of S at the rim:

Theorem 1 (Blaschke [8]) *Let ϕ be the angle between ξ and e_1 . Then,*

$$k_o^{-1}(\phi) = k_{n_1}^{-1} \sin^2 \phi + k_{n_2}^{-1} \cos^2 \phi \quad (2)$$

Corollary 1 *If ξ is along the principal direction e_1 at p , then $k_o = k_{n_2}$.*

Corollary 2 *Let ξ, ξ' be two distinct viewing directions in $T_p(S)$ from which p is visible, and let k_o, k'_o be the curvatures of the occluding contour at the projections of p . If (1) $K \neq 0$ at p , (2) $\xi = e_1$, and (3) the angle between ξ and ξ' is known, then we can compute k_{n_1} , e_2 , and K at p .*

Corollary 3 *Let p be a non-umbilic point on the visible rim of S with $K \neq 0$. Let $\phi \in [0, 2\pi)$ be the angle between $\xi \in T_p(S)$ and e_1 . (1) If p is elliptic (i.e., $K > 0$), the function $k_o(\xi)$ takes its minimum and maximum values only when ξ coincides with one of the principal directions. (2) If p is hyperbolic (i.e., $K < 0$), $k_o(\xi)$ is well-defined only when $|\phi| < \arctan \sqrt{k_{n_1}/(-k_{n_2})}$ for $\phi < \pi$, or $|\phi| - \pi < \arctan \sqrt{k_{n_1}/(-k_{n_2})}$ for $\phi \geq \pi$. For these directions, $k_o(\xi)$ takes its maximum value when ξ coincides with e_1 and it has no minimum value.*

Corollary 1 suggests that the principal directions at p form a special set of directions providing explicit information about surface geometry in the vicinity of p . Now assume that we are viewing p from a particular viewing direction and can measure the curvature of the occluding contour at p ’s projection. If somehow we can adjust our viewing direction to coincide with a principal direction at p and know what this adjustment is, Corollary 2 shows that we can derive the second fundamental form of S at p . This solves the shape recovery problem for p . The most important result is given by Corollary 3. It shows that the problem of finding the principal directions at a point can be treated as a simple maximization (or minimization) problem. We describe the implications of this result in the next section and show how it can be used by an active observer to find the principal directions at p .

4 Recovering local surface geometry

The basic step of our surface reconstruction approach is to select a point on the occluding contour and recover the local surface geometry for its corresponding rim point. We do not address the point selection problem directly. The reason for this is that we cannot decide *a priori* which point on the occluding contour will prove the most useful. This will depend on the context in which this approach is used. However, there are specific types of points for which our reconstruction method may not work. Therefore, our task will be to select a point on the rim for which we can ensure that our approach is effective. Below we first outline the main ideas and, in Section 6, present more details.

4.1 The active reconstruction approach

Suppose we have selected a point p on the rim. For simplicity we will assume that p is at the origin. Corollary 3 says that if p is a non-umbilic elliptic or a non-umbilic hyperbolic point, there are only two viewing directions in $T_p(S)$ for which k_o obtains a local maximum value (i.e., those for which the angle with e_2 is 0 or π) and two directions for which k_o obtains a local minimum. Our goal is to find one of these directions since they correspond to e_1 and e_2 . We discuss the problem of finding e_2 ; e_1 is treated similarly.

Viewing directions in the plane $T_p(S)$ can be thought of as points on a unit circle C , defined by the intersection of a unit sphere centered at p with $T_p(S)$ (see Figure 1). As the observer changes viewing direction on $T_p(S)$, the corresponding point moves on C . Our goal is to smoothly move this point on C until the viewpoint with maximum k_o is found. To do this we must answer two questions: (1) Which direction should the observer move on the unit circle, and (2) how can the observer detect when the viewing direction is equal to e_2 ?

We only have two possibilities for moving on the unit circle, either clockwise or counterclockwise. Obviously, we prefer the minimal motion solution in which the desired extremum is attained with the smallest possible change in viewing direction. From Corollary 3 we know that if we move in the direction of increasing k_o , the first extremum we reach is a maximum. It follows from the local geometry of elliptic and hyperbolic points that this strategy will produce the smallest viewing direction change (Figure 2). However, parabolic points do not have this property.

The second question, detecting when the viewing direction is equal to e_2 , is partly answered by Corollary 3. It says that we can detect this event by detecting a local maximum of k_o . However, in order to detect this local maximum, p must be visible; k_o cannot be measured otherwise. The visibility of p is affected by the local surface geometry at p as well as by the global geometry of S . Ignoring for a moment the case where p is occluded by some distant point on S , we arrive at the following two conclusions: (1) If p is




Figure 1: Viewing directions on $T_p(S)$.




Figure 2: Finding the principal directions. Top views of the tangent plane are shown. The curve D is Dupin's indicatrix for p . (a) p is an elliptic point. (b) p is a hyperbolic point. Shaded areas, delimited by the asymptotes of the point, represent the directions where p is occluded. (Note that the axis labels have been reversed.)

elliptic, we can align the viewing direction with either e_1 or e_2 . Furthermore, the maximum possible direction change before the alignment takes place is $\pi/2$ (Figure 2(a)). (2) If p is hyperbolic, we can align the viewing direction only with e_1 . The maximum possible direction change in this case is determined by the point's asymptotes (Figure 2(b)).

These results suggest a simple algorithm to align the observer's viewing direction with e_2 :

Step 1: Perform a small change of viewing direction on $T_p(S)$ and measure the difference between the previous and current value of k_o . If it increases, continue to change the viewing direction in the same way so that e_2 will be reached first. If it decreases, move the viewing direction in the opposite way.

Step 2: Continue moving in the same direction (or perform a binary search along the lines of [4]) until k_o reaches a maximum. This viewpoint corresponds to e_2 and therefore the observer can stop moving and use the current value of k_o for k_{n_1} .

Step 3: Measure the total change of direction between the initial and final viewing direction. Corollary 2 says that this angle along with k_{n_1} and the initial value of k_o can be used to determine k_{n_2} .

4.2 Selecting points for reconstruction

Any observer motion minimizing or maximizing k_o must take into account the effects of global surface geometry: Irrespective of its local structure, p may become occluded by distant points on S . The following proposition shows that (a) there are at least some points on the visible rim of S that cannot be occluded by S if the observer changes viewing direction as described above, and (b) these points are easily detected on the occluding contour (see [9] for the proof):




Figure 3: Selecting points for surface recovery. Solid lines on the occluding contour of a candlestick show the points that cannot become occluded while changing viewing direction in their tangent planes.

Proposition 1 (1) *Let p be a visible, elliptic point on the rim of a smooth surface S when viewed from direction ξ under orthographic projection. Let q be the projection of p in the image plane and let l be the tangent to the occluding contour at q . Then, p is visible from every direction on $T_p(S)$ iff l does not intersect the occluding contour and is not tangent to it at any point other than q .*

(2) *Let C be the occluding contour of S when viewed from direction ξ . Then there is at least one point on S projected in C that is visible from every direction in $T_p(S)$.*

Figure 3 shows the results of applying Proposition 1 to the occluding contour of a candlestick. The proposition implies that the only points ensuring the correctness of the algorithm are elliptic. However, this is a *necessary* requirement for the presence of occlusion but not a *sufficient* one. This means that there are cases where the geometry of hyperbolic points can be recovered with our approach. In fact, shape recovery for hyperbolic points requires less observer motion on average since the extent of the visibility of these points is limited by their asymptotic directions.

5 Surfaces of revolution

In the last section we presented an algorithm for recovering the shape of a single point on the surface rim. Surfaces of revolution, however, present an ideal example of surfaces for which the local shape of a single rim point reveals global properties of the surface. Such surfaces are completely described by their axis and their generating curve. Approaches for recovering the axis of a surface of revolution have been mainly geared towards detecting symmetries in their occluding contour or outline (e.g., [10]), or utilizing their viewpoint-invariant properties (e.g., [11]). The problem with detecting symmetries in the occluding contour is that their existence depends on viewpoint. On the other hand, the identification, detection and utilization of viewpoint-invariant properties is a non-trivial task [11]. Furthermore, the axis is severely foreshortened for near-“top” views.

Our active approach neatly lends itself to these problems in order to make them easier to handle. The idea is that if the viewer can align the viewing direction with a principal direction of a rim point, then shape and symmetry analysis of the occluding contour becomes especially simple. This is because the parallels of a surface of revolution (i.e., the lines generated by rotating a point of the generating curve around the axis) are lines of curvature. Hence, one of the principal directions corresponds to a “side” view of the surface (i.e., a view for which the viewing direction is perpendicular to the axis of rotation). If the generating curve of the surface can be written in the form $y = f(x)$, then the recovery of the axis of rotation allows us to recover the generating curve directly from a side view. The occluding contour from a side view is symmetric. Therefore, the axis of rotation (as well as the generating curve) can be recovered by simply using existing symmetry-seeking approaches which are well-defined for such a viewpoint. In fact, the direction and position of the axis can also be determined by recovering the principal curvatures for two points on the rim of a side view [9].

It is also possible to extend the above discussion to the case of regular, straight, homogeneous generalized cylinders, by aligning the viewing direction with the principal direction of points belonging to geodesic parallels [9]. In either case, the observer must choose between moving towards the principal direction of minimum curvature or moving towards the one of maximum curvature. Surfaces of revolution whose generating curve is of the form $y = f(x)$ have two properties that can help in this decision: (1) Hyperbolic rim points are visible only from the principal direction corresponding to a side view, and (2) if the viewing direction smoothly changes on the tangent plane of a selected rim point, this point will not become occluded if the viewing direction is approaching the direction of a side view.

6 Extending surface recovery

Our main objective is to recover the complete shape description for a single rim point. In this section we consider an extension to this approach—selecting a new point and applying the shape recovery process to that point. We must consider two important issues in order to demonstrate the effectiveness of such an extension:

- (*Efficiency*) The extent of the viewing direction adjustments needed to align the viewing direction with one of the principal directions at the newly selected point.
- (*Reliability*) The extent of the viewing direction adjustments *required* by our basic shape recovery algorithm in order to produce reliable shape information for the newly selected point. This is because if the viewing direction adjustment is close to zero, numerical problems are introduced in the calculations of the principal curvatures from Corollary 2.

Let p be the previously selected point. After applying the shape recovery step, the viewing direction ξ of the observer is aligned with one of the principal directions at p , say e_2 . We have seen that if we change directions in $T_p(S)$, p will not leave the rim. Therefore, we must change viewing direction in some other plane containing e_2 . The important issues here are (1) which plane should be selected for changing the viewing direction, and (2) how much should the viewing direction change in that plane? We show in [9] that if the selected plane P is the normal plane (i.e., the plane defined by the viewing direction and the surface normal at p) and if the viewing direction change on this plane is small, then the viewing direction adjustments during the shape recovery step will in fact be smooth and depend entirely on intrinsic properties of the surface. Specifically, we show that these adjustments are (to a first approximation) proportional to the geodesic curvature of the lines of curvature at p and inversely proportional to the normal curvature of the lines of curvature at p . This is an important result because it allows us to predict the performance of our active viewing strategy based on knowledge of the intrinsic properties of the surface.

7 Experimental results

We have implemented a prototype system that (1) automatically selects points on the rim of an object, (2) tracks these points while changing viewing direction on their tangent plane, and (3) computes the curvature of the occluding contour at the selected points in order to detect the viewpoints where it obtains an extremum value. Our experiments were performed by synthetically generating images from 3D object models (Figure 4). The viewing direction changes were constrained to lie on a horizontal plane perpendicular to the plane of the page. Our system automatically detected points on the occluding contour having tangent planes parallel to this plane and subsequently labeled and tracked them while the viewing direction changed smoothly (Figure 4). The points were initially detected and labeled for the viewing direction $\xi = 0$. Note that after a rotation of 3.93 radians the only unoccluded points for the candlestick model are the points 0 and 6, exactly as predicted by Proposition 1 (i.e., the tangents to the occluding contour at these points do not intersect the contour).

Because we used polyhedral models, curvature computations were performed by first approximating the occluding contour in the neighborhood of the selected points using cubic B-splines. Figure 5 shows that the major peaks and valleys of the curvature estimates are clearly visible even in the presence of the discontinuities caused by the polyhedral approximation. Both our models were surfaces of revolution with their “side” view corresponding to a principal direction of maximum curvature for all selected rim points. View 2 shows the occluding contour from the viewpoint of maximum curvature for point 0.

8 Concluding remarks

Current limitations of the approach are (1) the use of orthographic projection, (2) the requirement that viewing directions change on arbitrary planes, and (3) its applicability to only elliptic or hyperbolic surface points. We believe, however, that our active approach of moving towards viewpoints that are closely related to the geometry of the viewed surfaces is a very important and general one. Consider, for example, the problem of obtaining a “face-on” view of a planar curve or a texture element. This problem has been studied extensively in the past and several approaches exist that *hypothesize* face-on views, based on information from a single viewpoint (e.g., [10]). We are currently investigating an approach that enables the observer to change viewpoint in order to obtain a face-on view of a planar curve. We are also trying to extend our results to perspective projection and are investigating possible uses of this approach in the context of active surface-exploration.

References

- [1] R. Cipolla and A. Blake, “The dynamic analysis of apparent contours,” in *Proc. 3rd Int. Conf. on Computer Vision*, pp. 616–623, 1990.
- [2] D. H. Ballard, “Animate vision,” *Artificial Intelligence*, vol. 48, pp. 57–86, 1991.
- [3] J. Aloimonos, “Purposive and qualitative active vision,” in *Proc. Int. Conf. on Pattern Recognition*, pp. 346–360, 1990.
- [4] E. Krotkov, “Focusing,” *Int. J. Computer Vision*, vol. 1, pp. 223–237, 1987.
- [5] P. Giblin and R. Weiss, “Reconstruction of surfaces from profiles,” in *Proc. 1st Int. Conf. on Computer Vision*, pp. 136–144, 1987.
- [6] J. J. Koenderink, “What does the occluding contour tell us about solid shape?,” *Perception*, vol. 13, pp. 321–330, 1984.
- [7] R. Vaillant and O. D. Faugeras, “Using extremal boundaries for 3-d object modeling,” *IEEE Trans. PAMI*, vol. 14, no. 2, pp. 157–173, 1992.
- [8] J. J. Koenderink, *Solid Shape*. MIT Press, 1990.
- [9] K. N. Kutulakos and C. R. Dyer, “Recovering shape by purposive viewpoint adjustment,” Tech. Rep. 1035, Computer Sciences Department, University of Wisconsin-Madison, August 1991.
- [10] R. Horaud and M. Brady, “On the geometric interpretation of image contours,” *Artificial Intelligence*, vol. 37, pp. 333–353, 1988.
- [11] J. Ponce, D. Chelberg, and W. B. Mann, “Invariant properties of straight homogeneous generalized cylinders and their contours,” *IEEE Trans. PAMI*, vol. 11, no. 9, pp. 951–966, 1990.




Figure 4: Snapshots of the occluding contour for a candlestick and two tori as viewing direction changes. The numbered points are the points automatically selected and tracked.

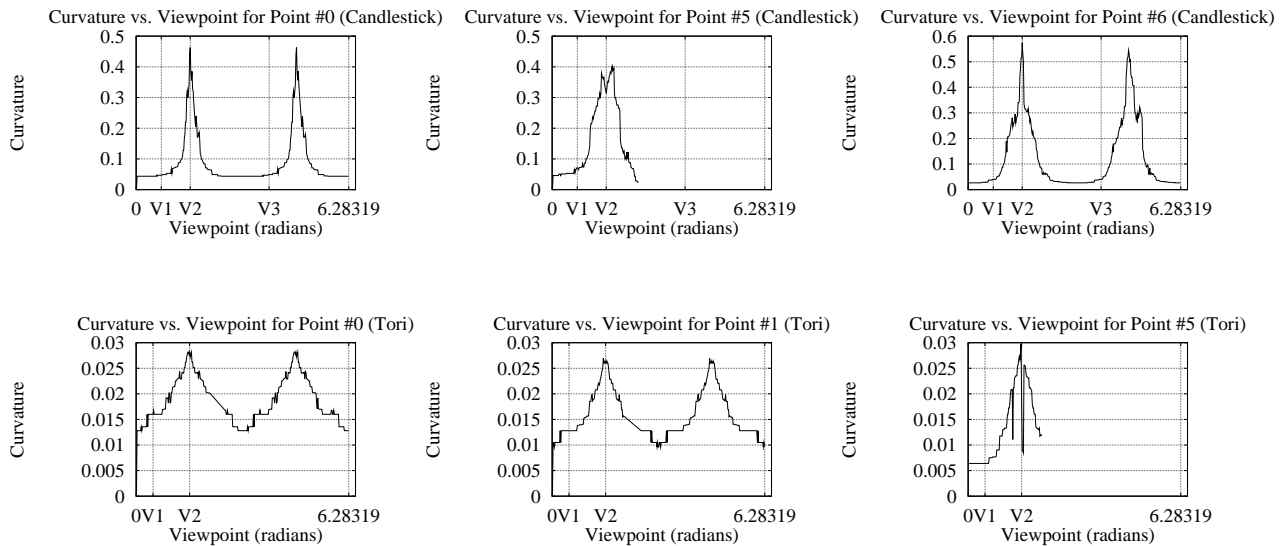


Figure 5: Variation of the absolute curvature with respect to viewpoint at the selected points on the occluding contour. The models were rotated a total of 2π radians. The curves for point 5 on the candlestick and the tori end at the viewpoint where their occlusion is detected.