

THREE-DIMENSIONAL SHAPE FROM LINE DRAWINGS

Stephen T. Barnard and Alex P. Pentland

SRI International, 333 Ravenswood Ave., Menlo Park, California 94025

ABSTRACT

The problem of interpreting the shape of a three-dimensional space curve from its two-dimensional perspective image contour is considered. Observation of human perception indicates that a good strategy is to segment the image contour in such a way as to obtain approximately planar segments. The orientation of the osculating plane (the plane in which the space curve lies) can then be estimated for these segments, and the three-dimensional shape recovered. The assumption of spatial isotropy is used to derive the theoretical results needed to formulate such an estimation strategy. The resulting estimation strategy allows a single three-dimensional structure (up to a single Necker reversal) to be assigned to any smooth image contour. An implementation is described and shown to produce an interpretation that is quite similar to the analytically correct one in the case of a helix, even though a helix has substantial torsion. The general applicability of the algorithm is discussed.

I Introduction

Much recent vision research has emphasized the importance of image contour for shape interpretation [1,2,3,4,5,6,7]. Tenenbaum and Barrow [1] argue that image contour, for example, is dominant over shape from shading. Pentland [8] has presented examples in which the addition of a contour substantially improved the interpretation of a shaded surface. It seems that contour is one of the strongest sources of information for shape perception.

One source of evidence of the strength of contour information is line drawings. When we examine a line drawing, our perception of the three-dimensional shape implied by such a drawing is nearly always clear and unambiguous. How can we account for this, given that purely geometrical constraints admit of an infinite number of valid interpretations?

A. An Observation About Human Perception

When we observe line drawings such as those in Figure 1 (a), we have a clear perception of a non-planar three-dimensional structure. Notice that if we were to segment each of these drawings at the circled points, each of the resulting segments would have the same shape as they did when they were still hooked together and would be approximately planar, as is shown in Figure 1 (b). Thus, for these line drawings the problem of recovering the three-dimensional structure can be reduced to the problems of (1) segmenting the curve into perceptually planar segments, and (2) finding the plane that contains each of the curve segments (the *osculating plane*) [9]. Once we know the orientation of the plane which contains a curve segment we can then easily determine its three-dimensional shape.




Figure 1. (a) Some Line Drawings, (b) Their Planar Subregions.

If we "by hand" try to segment image contours into planar regions, we find that the strategy can be successfully applied to a surprisingly large number of naturally-occurring image contours. For some contours, however, it is not obvious how well this strategy will work, primarily because there are no points which segment the space curve into planar regions. An example of such a curve is the helix shown in Figure 2 (a). Nonetheless, it may still be possible to obtain a good approximation of the three-dimensional structure of such a curve using this strategy.

B. A Strategy For Recovering Three-Dimensional Shape

This observation about human perception leads to the following processing strategy:

- (1) Segment the image contour in such a way that each segment is likely to comprise a projection of a planar segment of the space curve.
- (2) Calculate the planes implied by the segments from (1).
- (3) Assemble the results of (2) into an estimate of the shape of the entire space curve.

The specific criteria for the initial segmentation are not dealt with here. It is clear, however, that the image contour should be segmented at singular points of curvature (maxima, minima, and inflection points). Hoffman and Richards [10] have presented a theory of curve segmentation that addresses this issue. Our approach will be to temporarily ignore the segmentation problem and to simply estimate the orientation of parts of the space curve from many local parts of the image contour. If valid results are forthcoming with this approach the method can only be improved with more elaborate segmentation.

C. Modeling the Space Curve

We shall model a space curve in the conventional way, as a three-dimensional vector function $\mathbf{x}(s)$ of one parameter s which is assumed to be a natural parameter, i.e., $|\mathbf{dx}(s)/ds| = 1$. The shape of such a curve is completely determined by two properties that are scalar functions of s : curvature, $\kappa(s)$, and torsion, $\tau(s)$ [9]. Curvature is always nonnegative; only straight lines and inflection points have zero curvature. Torsion may be intuitively defined as the amount of "twist" in the curve at a point s . Another way to visualize torsion is as the degree to which the osculating plane (the plane which contains the curve) is changing. Only planar

* The research reported herein was supported by the Defense Advanced Research Projects Agency under Contract No. MDA 903-83-C-0027; this contract is monitored by the U. S. Army Engineer Topographic Laboratory. Approved for public release, distribution unlimited.

curves have zero torsion everywhere. Unlike curvature, torsion may be either negative or positive.

The presence of torsion is not directly evident in the image. It simply results in more or less foreshortening as the osculating plane of the contour varies. The effects of torsion, therefore, can be exactly mimicked by changes in curvature, and vice versa.

II Theory of Contour Interpretation

Not all three-dimensional interpretations of an image contour are equally likely. If we assume that spatial isotropy holds, then we know that viewer position is independent of the shape of the curve — which allows us to make a reasonable guess about the latter's three-dimensional shape [8]. The first step towards a guess at the space curve's shape is the following proposition:

Proposition (Zero Torsion). *The maximum-likelihood estimate of the torsion of the space curve is zero (i.e., no "twisting" of the curve).*

This proposition follows because the assumption of spatial isotropy implies that the viewer's position and the shape of the space curve are mutually independent. Thus, not only is it unlikely that significant features of the curve will be hidden from view by coincidental alignment of the viewer and the curve, but, conversely, it is likely that the viewed scene will not change much with small changes in viewing position.* The appearance of a curve with substantial torsion† will change considerably with small changes in viewer position; if we assume spatial isotropy, therefore, we must expect that the torsion of the curve will be small.

Furthermore, given that spatial isotropy implies that the viewer position and the shape of the curve are mutually independent, the torsion of the curve must then also be independent of viewer position. Consequently, the torsion of the curve is as likely to be positive as negative, and thus the mean value (and maximum-likelihood estimate) for the magnitude of the torsion is zero†. The probability that the torsion is small implies this estimate will generally be a good one.

A. Estimation With The Assumption Of Zero Torsion

Even if we assume that torsion is zero (i.e., the space curve is planar), there is still a two-parameter set of space curves that could have generated that imaged contour. The two parameters correspond to the two degrees of freedom of the osculating plane.

Assume that we are given a small portion of an imaged contour, and asked to estimate the three-dimensional shape of the space curve which generated that image. If we measure the position and curvature at three points on the imaged contour, then we can uniquely define an elliptical arc that fits the image data. By the previous proposition, this elliptical arc is most likely caused by a space curve that is either an arc of a circle or of an ellipse, as those are the two planar (zero torsion) shapes which can project to an ellipse††.

Previous research ([2], [12]) has shown that the maximum-

likelihood estimate of the space curve's shape is given by the following proposition (see also [2]):

Proposition (Planar Interpretation). *Given an elliptical segment of an image contour and that the space curve is planar, the maximum likelihood estimate of the space curve's three-dimensional shape is a segment of a circle.*

Barnard [12] has constructed a maximum entropy estimator that implements this proposition for perspective images and that is tolerant of digitization noise. Operating under the assumption that the space curve has zero torsion, it chooses the orientation that maximizes the entropy of backprojected image contour curvature measurements. That is, curvature is first measured at several points in the image contour, then the curvatures of hypothetical planar space curves of essentially all orientations are computed by backprojection, and, finally, the orientation that leads to the space curve of most uniform curvature (in the sense of maximum entropy) is selected. In general, three image contour curvature measurements are sufficient for an unambiguous maximum-entropy interpretation (up to a Necker reversal).

III Three-Dimensional Estimation

Now let us return to the general problem of estimating the shape of the space curve, given a smooth imaged contour. Let us first take three curvature measurements along the imaged contour. These three measurements define an ellipse. As just described, this leads to a circular interpretation of the space curve. Now suppose that we have additional image contour curvature measurements. There are, then, two cases to consider:

First case: the new points fit on the same ellipse. In the first case we have quite strong evidence of the space curve's shape. For, if the osculating plane were changing, the curvature would have to be changing also — and in just such a manner as to exactly cancel (in the image) the effect of the changing osculating plane. Similarly, if the curvature of the space curve were changing, the osculating plane would have to change just exactly enough to cancel the effect of the changing curvature. As such a "conspiracy" to cancel the visible effects of change is unlikely (a direct violation of general position), we must conclude that there was neither torsion nor change in curvature, and, thus, there is a great (in fact, maximum) likelihood that the new image curvature measurements result from the same circular space curve defined by the first three measurements.

Second case: the new points don't fit on the same ellipse. What if the additional measurements lie off the ellipse defined by the first three measurements? Then we can be certain that either the curvature or the osculating plane (or both) of the space curve has changed. This new point is, therefore, a possible place to segment the curve. What we must do when we encounter such a point is advance along the image contour until we are completely past the point, and obtain a new estimate of the space curve's osculating plane. If the new osculating plane has the same orientation as the previous osculating plane, then we have evidence that the space curve continues to be planar, and we should not segment the curve. If, however, we obtain a different orientation for the osculating plane, then we should segment the space curve and begin a new planar segment of the curve.

As any smooth image contour may be closely approximated by portions of ellipses and straight lines*, this interpretation strategy will yield a single interpretation for the three-

*This is often referred to as the assumption of general position. Thus, spatial isotropy implies general viewing position.

**As a function of position on the image contour rather than as a function of s

†Note that at places where the curvature is zero — straight segments and inflection points — the torsion is not defined and may arbitrarily be taken to be zero. That is, the osculating plane may be changed freely at these points without affecting the shape of the space curve.

††This is true of both perspective and orthographic projection, however, we will deal exclusively with the more general case of perspective foreshortening.

*Only the third and higher derivatives of the imaged contour that will fail to be exactly matched. People, it should be noted, are very poor observers of changes in the third derivatives of an image contour.

dimensional shape of the space curve (up to Necker reversals). Further, this interpretation will be the most likely interpretation on a point-by-point basis. It should be noted that the first two steps of this estimation strategy are similar to the strategy proposed in [1].

IV An Example

The interpretation strategy has been implemented and applied to a synthetic image of a helical space curve. The helix example is a good test because a helix has significant torsion everywhere, thus, distinguished segmentation points do not exist and it is not clear what the estimation strategy will do. If we can recover the helical shape of the space curve with some accuracy, we shall have demonstrated that the estimation strategy can perform even when no good segmentation is available.

Figure 2 (a) shows a perspective image of a helix. Figure 2 (b) shows a plot of the spherical indicatrix of the helix. The spherical indicatrix is a plot of the orientation of the osculating plane of the space curve. The axes in this plot corresponds to the azimuth and elevation of the osculating plane. As mentioned previously, knowledge of the orientation (azimuth and elevation) of the osculating plane at each point, together with the imaged contour, uniquely determines the shape of the space curve. Thus, the spherical indicatrix is a method of displaying the three-dimensional shape of the space curve. Figure 2 (c) shows the spherical indicatrix estimated for the contour in (a). When this is compared with the actual indicatrix shown in (b), it is evident that the three-dimensional shape of the space curve has been fairly accurately recovered.

Summary. We have developed a theory for assigning a three-dimensional interpretation to any smooth image contour. The theory has been implemented and is undergoing evaluation, which may lead to further development. The results reported above indicate that the estimation strategy performs reasonably well even for cases such as a helix, where the presence of substantial torsion might have led one to expect poor performance.

REFERENCES

- [1] H. G. Barrow and J. M. Tenenbaum, Interpreting Line Drawings as Three-Dimensional Surfaces, *Artificial Intelligence* 17 (1981), 1-47.
- [2] A. Witkin, Recovering Surface Shape And Orientation From Texture, *Artificial Intelligence* 17 (1981), 17-47.
- [3] J. Kender, Shape From Texture, *Ph. D. Thesis, Computer Science Department, Carnegie-Mellon University* (1980).
- [4] T. Kanade, Recovery of the 3-D Shape of an Object from a Single View, *Artificial Intelligence* 17 (1981), 409-460.
- [5] K. A. Stevens, The Visual Interpretation of Surface Contours, *Artificial Intelligence* 17 (1981), 47-73.
- [6] M. Brady and A. Yuille, An extremum principle for shape from contour, *In this proceedings*.
- [7] D. G. Lowe and T. O. Binford, The Interpretation of Three-Dimensional Structure from Image Curves, *Proceedings of the 7th IJCAI, Aug. 24-28, University of British Columbia, Vancouver, B.C., Canada.* 2, 613-618.
- [8] A. Pentland, Local Inference of Shape: Computation From Local Features, *Ph.D. Thesis, Psychology Department, Massachusetts Institute of Technology* (1982).
- [9] M. M. Lipshutz, *Differential Geometry*, McGraw-Hill, New York, New York, 1969.
- [10] D. D. Hoffman and W. A. Richards, Representing Smooth Plane Curves for Recognition: Implications for Figure-Ground Reversal, *Proceedings of the AAAI, Aug. 18-20, 1982, Carnegie-Mellon University, Pittsburgh, Pennsylvania*, 5-8.
- [11] A. Witkin, Shape From Contour, *Ph.D. Thesis, Psychology Department, Massachusetts Institute of Technology* (1982).
- [12] S. Barnard, Interpreting Perspective Images, *SRI Artificial Intelligence Center Technical Note 271 to appear in Artificial Intelligence* (1982).




Figure 2. (a) An image of a helix, (b) the actual spherical indicatrix, (c) the recovered spherical indicatrix.