

AN EXTREMUM PRINCIPLE FOR SHAPE FROM CONTOUR

Michael Brady and Alan Yuille

321, MIT Artificial Intelligence Laboratory,
545 Technology Square,
Cambridge, MA 02139

Abstract

An extremum principle is developed that determines three-dimensional surface orientation from a two-dimensional contour. The principle maximizes the ratio of the area to the square of the perimeter, a measure of the compactness or symmetry of the three-dimensional surface. The principle interprets regular figures correctly, it interprets skew symmetries as oriented real symmetries, and it is approximated by the maximum likelihood method on irregular figures.

1. Introduction

An important goal of early vision is the computation of a representation of the visible surfaces in an image, in particular the determination of the orientation of those surfaces as defined by their local surface normals [Brady 1982, Marr 1982]. Many processes contribute to achieving this goal, stereopsis and structure-from-motion being the most studied in image understanding. In this paper we consider the computation of shape-from-contour. Figure 1 shows a number of shapes that are typically perceived as images of surfaces which are oriented out of the picture plane. The method we propose is based on a preference for symmetric, or at least compact, surfaces. Note that the contour does not need to be closed in order to be interpreted as oriented out of the image plane. Also, in general, contours are interpreted as curved three-dimensional surfaces.

We develop an extremum principle for determining three-dimensional surface orientation from a two-dimensional contour. Initially, we work out the extremum principle for the case that the contour is closed and that the interpreted surface is planar. Later, we discuss how to extend our approach to open contours and how to interpret contours as curved surfaces.

The extremum principle maximizes a familiar measure of the compactness or symmetry of an oriented surface, namely the ratio of the area to the square of the perimeter. It is shown that this measure is at the heart of the maximum likelihood approach to shape-from-contour developed by Witkin [1981] and Davis, Janos, and Dunn [1982]. The maximum likelihood approach has had some success interpreting irregularly shaped objects. However, the method is ineffective when the distribution of image tangents is not random, as is the case, for example, when the image is a regular shape, such as an ellipse or a parallelogram. Our extremum principle interprets regular figures correctly. We show that the maximum likelihood method approximates the extremum principle for irregular figures.

Kanade [1981, page 424] has suggested a method for determining the three-dimensional orientation of skew-symmetric figures, under the "heuristic assumption" that such figures are interpreted as oriented real symmetries. We prove that our extremum principle necessarily interprets skew symmetries as oriented real symmetries, thus dispensing with the need for any heuristic assumption to that




Figure 1. 1a. Two-dimensional contours that are often interpreted as planes that are oriented with respect to the image plane. The commonly judged slant is shown next to each shape. 1b. Some unfamiliar shapes that are also interpreted as planes that are oriented with respect to the image plane. 1c. Some shapes that are interpreted as curved three-dimensional surfaces.

effect. Kanade shows that there is a one-parameter family of possible orientations of a skew-symmetric figure, forming a hyperbola in gradient space. He suggests that the minimum slant member of the one-parameter family is perceived. In the special case of a real symmetry, Kanade's suggestion implies that symmetric shapes are perceived as lying in the image plane, that is having zero slant. It is clear from the ellipse in Figure 1 that this is not correct. Our method interprets real symmetries correctly.

First, we review the maximum likelihood method. In Section 3, we discuss several previous extremum principles and justify our choice of the compactness measure. In Section 4, we derive the mathematics necessary to extremize the compactness measure, and relate the extremum principle to the maximum likelihood method. In Section 5, we investigate Kanade's work on skew symmetry. One approach to extending the extremum principle to interpret curved surfaces, such as that shown in Figure 1c, is sketched in Section 6.

This paper is an abbreviated version of [Brady and Yuille, 1983], which should be referred to for the details of derivations. In that paper we also discuss the psychophysical literature on slant estimation and Ikeuchi's work on shape from texture.

2. The Sampling Approach

Witkin [1981] has treated the determination of shape-from-contour as a problem of signal detection. Recently, Davis, Janos, and Dunn [1982] have corrected some of Witkin's mathematics and proposed two efficient algorithms to compute the orientation of a planar surface from an image contour. Witkin's approach uses a *geometric model* of (orthographic) projection and a *statistical model* of (a) the distribution of surfaces in space (statistics of the universe) and (b) of the distribution of tangents to the image contour. We

utilize the geometric model, but dispense with the second part of the statistical model in favor of an extremum principle.

First, the geometric model. Suppose that a curve is drawn in the plane (σ, τ) and denote by β the angle that the tangent makes at a typical point on the curve. Let α be the tangent angle in the image plane at the point corresponding to β . Then α and β are related by $\tan(\alpha - \tau) = \tan \beta / \cos \sigma$.

We now turn to the statistical model, which consists of two assumptions called *isotropy* and *independence*. Isotropy reasonably supposes that all surface orientations are equally likely to occur in nature and that tangents to surface curves are equally likely in all directions. More succinctly, it is assumed that the quantities (σ, τ) are randomly distributed, and that their joint probability density function ("density") $D(\sigma, \tau)$ is given by [Davis, Janos, and Dunn, 1982]

$$D(\sigma, \tau) = \frac{1}{\pi} \sin \sigma$$

The independence assumption requires that the image tangents

$$\{\alpha_i: 1 \leq i \leq n\}$$

are statistically independent. That is, it is assumed that the tangent directions at different points on the image curve are independent. This is only true if the contour is highly irregularly shaped, or if the number of samples is small. In any case, the assumption of independence is an inherent weakness of the sampling approach (see for example [Witkin 1981, p 36]).

It is easy to show that the conditional density $D(\alpha|\sigma, \tau)$ of an individual image tangent angle α projected from a plane (σ, τ) is given by [Witkin 1981]:

$$D(\alpha|\sigma, \tau) = \frac{1}{\pi} \frac{\cos \sigma}{\cos^2(\alpha - \tau) + \cos^2 \sigma \sin^2(\alpha - \tau)}. \quad (2.1)$$

Denote the sample $(\alpha_1, \alpha_2, \dots, \alpha_n)$ by A (the sample is independent by assumption). It has conditional density

$$D(A|\sigma, \tau) = \prod_{i=1}^n D(\alpha_i|\sigma, \tau) \quad (2.2)$$

By Bayes' formula we obtain

$$D(\sigma, \tau|A) = \frac{D(A|\sigma, \tau)D(\sigma, \tau)}{\int \int D(A|\sigma, \tau)D(\sigma, \tau) d\sigma d\tau} \quad (2.3)$$

Observe that the numerator is independent of σ and τ . The sampling approach takes a random sample A and defines the most likely orientation of the plane (σ, τ) to be that which extremizes $D(\sigma, \tau|A)$. Witkin [1981] quantizes σ and τ , and describes an algorithm to find the maximizing (σ_j, τ_k) . Davis, Janos, and Dunn [1982] develop a more efficient algorithm that first estimates σ and τ and then uses those estimates in a Newton iterative process. They provide evidence that their method is more accurate than Witkin's. Curiously, however, they state [Davis, Janos, and Dunn 1982, p 24] that "the iterative algorithm was not used [in the experiments they report] because the initial estimates (whose computation is trivial) are very accurate and the iterative scheme often failed to converge to the solution".

3. Extremum Principles.

Brady and Horn [1983] survey the use of extremum principles in image understanding. The choice of performance index or measure to be extremized, and the class of functions over which the extremization takes place, are justified by appealing to a model of the geometry or photometry of image forming and constraints such as smoothness. For example, the use of extremum principles in surface reconstruction is based upon surface consistency theorems [Grimson 1981, Yuille 1983] and a thin plate model of visual surfaces [Brady and Horn 1983, Terzopoulos 1983]. Brady and Yuille [1983] discuss the relationship between extremum principles and Praeganz theories in perceptual psychology.

There are several plausible measures of a curve that might be extremized in order to compute shape-from-contour. Contrary to what appears to be a popular belief, given an ellipse in the image plane, $\oint \kappa^2 ds$ is *not* extremized in the plane that transforms the ellipse into a circle [Brady and Yuille 1983, Appendix A]. Since ellipses are normally perceived as slanted circles, we reject the square curvature as a suitable measure.

Another possible measure is proposed by Barrow and Tenenbaum [1981, p89]. Assuming planarity (the torsion τ is zero), it reduces to

$$\oint \left(\frac{d\kappa}{ds} \right)^2 ds.$$

The first strong objection to this measure is that it involves high-order derivatives of the curve. This means it is overly dependent on small scale behaviour. Consider, for example, a curve which is circular except for a small kink. The circular part of the curve will contribute a tiny proportion to the integral even when the plane containing the curve is rotated. The kink, on the other hand, will contribute an arbitrarily large proportion and so will dominate the integral no matter how small it is compared with the rest of the curve. This is clearly undesirable. For example, it suggests that the measure will be highly sensitive to noise in the position and orientation of the points forming the contour.

A second objection to the measure proposed by Barrow and Tenenbaum is that it is minimized by, and hence has an intrinsic preference for, straight lines, for which $d\kappa/ds$ zero. This means that the measure has a bias towards planes that correspond to the (non-general) sideon viewing position. These planes are perpendicular to the image plane and have slant $\pi/2$.

We base our choice of measure on the following observations.

1. Contours that are the projection of curves in planes with large slant are most effective for eliciting a three-dimensional interpretation.
2. A curve is foreshortened by projection by the cosine of the slant angle in the tilt direction, and not at all in the orthogonal direction.

We conclude that three-dimensional interpretations are most readily elicited for shapes that are highly elongated in one direction. Another way to express this idea is that the image contour has large aspect ratio or is radially asymmetric. The measure we suggest will pick out the plane orientation for which the curve is most compact or most radially symmetric. Specifically, our measure is

$$M = \frac{(\text{Area})}{(\text{Perimeter})^2}. \quad (3.1)$$

This is a scale invariant number characterizing the curve. For all possible curves it is maximized by the most symmetric one, a circle. This gives the measure an upper bound of $1/4\pi$. Its lower bound is clearly zero and it is achieved for a straight line. It follows that our measure has a builtin prejudice against sideon views for which the slant is $\pi/2$.

In general, given a contour, our extremum principle will choose the orientation in which the deprojected contour maximizes M . For example an ellipse is interpreted as a slanted circle. The tilt angle is given by the major axis of the ellipse. It is also straightforward to show that a parallelogram corresponds to a rotated square. Brady and Yuille [1983, Appendix B] show how several simple shapes are interpreted by the measure. In particular, an ellipse is interpreted as a slanted circle, a parallelogram as a slanted square, and a triangle as a slanted equilateral triangle. In Section 5 we extend the parallelogram result to the more general case of skewed symmetry.

We note that the quantity M is commonly used in pattern recognition and industrial vision systems [Ballard and Brown 1982] as a feature that measures the compactness of an object. Furthermore, we can show that the measure M defined in Eq. (3.1) is at the heart of the geometric model in the maximum likelihood approach.

From Section 2, we see that the maximum likelihood approach maximizes the product of a number of terms of form

$$f(\alpha) = \frac{\cos \alpha}{\cos^2(\alpha - \tau) + \cos^2 \sigma \sin^2(\alpha - \tau)}. \quad (3.2)$$

Differentiating the geometric model with respect to the arc length s_I along the image curve and s_R along the rotated curve respectively we obtain

$$\frac{\kappa_I ds_I}{\kappa_R ds_R} = \frac{1}{f(\alpha)} \quad (3.3)$$

where κ_I and κ_R are the curvature at corresponding points of the image contour and its deprojection in the rotated plane respectively. In fact, $\kappa_I = d\alpha/ds_I$ and $\kappa_R = d\beta/ds_R$. There is no σ or τ dependence in the numerator of equation (3.3). We can write each term κds as $ds ds/\rho ds$ where ρ is the radius of curvature. Now observe that $\rho ds/ds$ is just a local computation of area divided by perimeter squared! Hence maximizing each $f(\alpha)$ in the maximum likelihood approach is equivalent to locally maximizing area over perimeter squared. In section (4) we will examine this connection more rigorously.

4. Extremizing the Measure

We now write down the measure for a curve with arbitrary orientation and then extremize with respect to the orientation. Let the unit normals to the image plane and the rotated plane be $\mathbf{k}$ and $\mathbf{n}$ respectively. The slant σ of the rotated plane is given by the scalar product $\cos \sigma = \mathbf{k} \cdot \mathbf{n}$.

Let Γ_R and Γ_I be the contour in the rotated and image planes. A vector $\mathbf{r}$ in the image plane satisfies $\mathbf{r} \cdot \mathbf{k} = 0$, and is the projection of a vector $\mathbf{v}$ in the rotated plane that satisfies $\mathbf{v} \cdot \mathbf{n} = 0$. Now Γ_R and Γ_I have (vector) areas A_R and A_I are related by

$$\|A_R\| = \frac{\|A_I\|}{\cos \sigma}$$

We find

$$P_R = \oint_{\Gamma_R} \left\{ (d\mathbf{r})^2 + \frac{(\mathbf{n} \cdot d\mathbf{r})^2}{(\mathbf{n} \cdot \mathbf{k})^2} \right\}^{\frac{1}{2}}$$

In general there is no simple relationship between the perimeters analogous to that holding between the areas. Nevertheless, we have $A_R/P_R^2 = A_I/P_I^2 \cos \sigma$.

and so our extremum principle is equivalent to extremizing $\cos^{\frac{1}{2}} \sigma P_R$, which we write as

$$I = \oint_{\Gamma_R} \left\{ (\mathbf{n} \cdot \mathbf{k}) d\mathbf{r}^2 + \frac{(\mathbf{n} \cdot d\mathbf{r})^2}{(\mathbf{n} \cdot \mathbf{k})^2} \right\}^{\frac{1}{2}}$$

We extremize this with respect to the orientation $\mathbf{n}$ of the rotated plane, maintaining the constraint that $\mathbf{n}$ is a unit vector by a Lagrange multiplier Λ . After algebraic manipulation [Brady and Yuille 1983], this reduces to

$$2 \oint \{ \cos^2 \sigma + (\mathbf{n} \cdot \mathbf{t})^2 \}^{-\frac{1}{2}} (\mathbf{t} \cdot \mathbf{n}) d\mathbf{r} = \mathbf{k} \times (\mathbf{k} \times \mathbf{n}) \oint \{ \cos^2 \sigma + (\mathbf{n} \cdot \mathbf{t})^2 \}^{\frac{1}{2}} d\mathbf{r}$$

where $\mathbf{t} = d\mathbf{r}/|d\mathbf{r}|$ is the unit tangent to the image contour. Let the unit vectors in the x and y directions in the image plane be $\mathbf{i}$ and $\mathbf{j}$ and let $\mathbf{k}$ be the normal to the image plane. The tangent vector $\mathbf{t}$ and the normal $\mathbf{n}$ can be written:

$$\mathbf{t} = \cos \alpha \mathbf{i} + \sin \alpha \mathbf{j}$$

$$\mathbf{n} = \sin \sigma \cos \tau \mathbf{i} + \sin \sigma \sin \tau \mathbf{j} + \cos \sigma \mathbf{k}$$

where α is the tangent angle in the image. We now form the scalar products of equation above with $\mathbf{i}$ and $\mathbf{j}$ to obtain eventually

$$\frac{\partial}{\partial \sigma} \left(\frac{1}{\cos^{\frac{1}{2}} \sigma} \oint (\cos^2 \sigma + \sin^2 \sigma \cos^2(\alpha - \tau))^{\frac{1}{2}} d\mathbf{r} \right) = 0 \quad (4.1)$$

$$\frac{\partial}{\partial \tau} \left(\frac{1}{\cos^{\frac{1}{2}} \sigma} \oint (\cos^2 \sigma + \sin^2 \sigma \cos^2(\alpha - \tau))^{\frac{1}{2}} d\mathbf{r} \right) = 0 \quad (4.2)$$

to emphasize that they correspond to extremizing with respect to σ and τ .

To conclude this Section, we show that these equations are similar, though not identical, to those obtained by the maximum likelihood method in the limit as the number of sampled tangents tends to infinity. To see this we recall that this method involves extremizing $D(A|\sigma, \tau)$ with respect to σ and τ . Since the denominator is independent of σ and τ , this amounts to extremizing $D(A|\sigma, \tau)D(\sigma, \tau)$. This is the same as extremizing $\log D(A|\sigma, \tau)D(\sigma, \tau)$. We find

$$E = n \log \cos \sigma + \log \sin \sigma - \sum_{i=1}^n \log(\cos^2(\alpha_i - \tau) + \cos^2 \sigma \sin^2(\alpha - \tau))$$

where we have ignored factors of π which will vanish on differentiation. Dividing E by n and taking the limit as n tends to infinity gives:

$$F = \log \cos \sigma \oint d\mathbf{r} - \oint \log(\cos^2(\alpha - \tau) + \cos^2 \sigma \sin^2(\alpha - \tau)) d\mathbf{r}.$$

Using the identity:

$$\cos^2(\alpha - \tau) + \cos^2 \sigma \sin^2(\alpha - \tau) = \cos^2 \sigma + \sin^2 \sigma \cos^2(\alpha - \tau)$$

gives

$$F = \log \cos \sigma \oint d\mathbf{r} - \oint \log(\cos^2 \sigma + \sin^2 \sigma \cos^2(\alpha - \tau)) d\mathbf{r}$$

This formula is similar to Eqs 4.1 and 4.2. Thus we expect the Extremum Method to give similar results to the Sampling Method when the contour is sufficiently irregular. We are currently carrying out experiments to verify this.

5. Skew Symmetry

We now consider a more general class of shapes for which the maximum likelihood approach is not effective. Kanade [1981, sec. 6.2] has introduced *skewed symmetries*, which are two-dimensional linear (affine) transformations of real symmetries. There is a bijective correspondence between skew symmetries and images of symmetric shapes that lie in planes oriented to the image plane. Kanade proposes the heuristic assumption that a skew symmetry is interpreted as an oriented real symmetry, and he considers the problem of computing the slant and tilt of the oriented plane.

Denote the angles between the x -axis of the image and the images of the symmetry axis and an axis orthogonal to it (the skewed transverse axis) by α and β respectively. The orthogonality of the symmetry and transverse axes enable one constraint on the orientation of the plane to be derived. Kanade uses gradient space (p, q) (see Brady [1982] for references) to represent surface orientations. He shows [Kanade 1981, p. 425] that the heuristic assumption is equivalent to requiring the gradient (p, q) of the oriented plane to lie on the hyperbola

$$p_1^2 \cos^2 \frac{(\alpha - \beta)}{2} - q_1^2 \sin^2 \frac{(\alpha - \beta)}{2} = -\cos(\alpha - \beta) \quad (5.1)$$

where

$$\begin{aligned} p_1 &= p \cos\left(\frac{\alpha + \beta}{2}\right) + q \sin\left(\frac{\alpha + \beta}{2}\right), \\ q_1 &= -p \sin\left(\frac{\alpha + \beta}{2}\right) + q \cos\left(\frac{\alpha + \beta}{2}\right). \end{aligned} \quad (5.2)$$

Kanade [1981, p. 426] further proposes that the vertices of the hyperbola, which correspond to the least slanted orientation, are chosen within this one-parameter family. This proposal is in accordance with a heuristic observation of Stevens [1980]. In the special case that the skew symmetry is a real symmetry, that is in the case that $\alpha - \beta = \pm \pi/2$, the hyperbola reduces to a pair of orthogonal lines [Kanade 1981, page 426] passing through the origin. In such cases the slant is zero. In other words, Kanade's proposal predicts that real symmetries are inevitably interpreted as lying in the image plane, and hence having zero slant. Inspection of Figure 1 shows that this is not the case. A (symmetric) ellipse is typically perceived as a slanted circle, particularly if the major and minor axes do not line up with the horizontal and vertical.

Although Kanade's minimum slant proposal does not seem to be correct, there is evidence (for example [Stevens 1980]) for Kanade's assumption that skew symmetries are interpreted as real symmetries. We can show that the assumption can in fact be *deduced* from our Extremum Principle [Brady and Yuille 1983, Section 5]. As a corollary we can determine the slant and tilt of any given skewed-symmetric figure [Brady and Yuille 1983, Appendix C]; only in special cases does it correspond to the minimum slant member of Kanade's one-parameter family.

6. Interpreting Image Contours as Curved Surfaces

Figure 1c shows a number of contours that are interpreted as curved surfaces. In this section we discuss one method for extending our extremum principle to this general case. The key observation, as it was for Witkin [1981], is that our method can be applied locally. To do this, we assume that the surface is locally planar. At the surface boundary, corresponding to the deprojection of the image contour, the binormal coincides with the surface normal. The idea is to compute a local estimate of the surface normal by the extremum principle described in the previous sections and then to use an algorithm, such as that developed by Terzopoulos [1982] to interpolate the surface orientation in the interior of the surface. Details of one implementation can be found in [Brady and Yuille 1983].

Acknowledgements

The authors thank Ruzena Bajcsy, Chris Brown, John Canny, Eric Grimson, Ellen Hildreth, Tommy Poggio, Demetri Terzopoulos, and Andy Witkin for their comments.

References

- Ballard, D. H., and Brown, C. M., [1982], *Computer vision*, Prentice-Hall, New Jersey.
- Barrow H. G. and Tenenbaum J. M., [1981], "Interpreting line drawings as three dimensional surfaces," *Artificial Intelligence*, 17, 75-117.
- Brady, Michael, [1982], "Computational approaches to image understanding," *Computing Surveys*, 14, 3-71.
- Brady, Michael, and Horn, B. K. P., [1983], "Rotationally symmetric operators for surface interpolation," *Computer Graphics and Image Processing*.
- Brady, Michael, and Yuille, Alan, [1983], An extremum principle for shape from contour, MIT, AI Lab., MIT-AIM 711.
- Davis, L. S., Janos, Ludwig, and Dunn Stanley, [1982], Efficient recovery of shape from texture, Computer Vision Laboratory, University of Maryland, TR-1133.
- Grimson, W. E. L., [1981], *From images to surfaces: a computational study of the human early visual system*, MIT Press, Cambridge.
- Horn, B. K. P., [1981], The curve of least energy, MIT, AIM-610.
- Kanade, T., [1981], "Recovery of three-dimensional shape of an object from a single view," *Artificial Intelligence*, 17, 409-460.
- Marr D., [1982], *Vision*, Freeman, San Francisco.
- Stevens, K. A., [1980], Surface perception from local analysis of texture (also, AITR512), MIT.
- Terzopoulos, D., [1983], "Multi-level reconstruction of visual surfaces," *Computer Graphics and Image Processing*.
- Witkin, Andrew P., [1981], "Recovering surface shape and orientation from texture," *Artificial Intelligence*, 17, 17-47.
- Yuille, Alan, [1983], Zero crossings on lines of curvature, MIT, forthcoming.