

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A. I. Memo 757

February 1984

Smoothed Local Symmetries
and
Their Implementation

Michael Brady and Haruo Asada

Abstract. We introduce a novel representation of two-dimensional shape that we call *smoothed local symmetries* (SLS). Smoothed local symmetries represent both the bounding contour of a shape fragment and the region that it occupies. In this paper we develop the main features of the SLS representation and describe an implemented algorithm that computes it. The performance of the algorithm is illustrated for a set of tools. We conclude by sketching a method for determining the articulation of a shape into subshapes.

This paper will appear in the *International Journal of Robotics Research*.

Acknowledgements. This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's Artificial Intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643, the Office of Naval Research under contract number N0014-80-C-0505, and the System Development Foundation. This work was done while Haruo Asada was a visiting scientist at MIT on leave from Toshiba Corporation, Japan.

1. Introduction

We introduce a novel representation of two-dimensional shape that we call *smoothed local symmetries* (SLS). Smoothed local symmetries represent both the bounding contour of a shape fragment and the region that it subtends or encloses. In this paper we develop the main features of the SLS representation and describe an implemented algorithm that computes it. The performance of the algorithm is illustrated for a set of tools. One step in the algorithm is the construction of a representation of the significant changes in curvature along the contour of the shape. We call that representation the *curvature primal sketch* [Asada and Brady 1984], since it is analogous to the primal sketch representation of intensity changes advocated by Marr [1976]. We conclude by sketching a method for determining the articulation of a shape into subshapes. A preliminary version of this paper appeared as [Brady and Asada, 1984].

The representation of shape is a fundamental problem in computer vision that has been surprisingly neglected. It is becoming increasingly important in Robot vision as the inadequacies of currently available systems manifest themselves. So long as the position and orientation of objects is tightly constrained, so long as contrast is sufficient to allow trivial segmentation into figure and ground, and so long as objects do not overlap, simple vision systems suffice. Increasing uncertainty in any of these variables requires a substantial increase in system sophistication.

Fortunately, there has been considerable advance in our understanding of computer vision over the past decade (see [Ballard and Brown 1982, Brady 1982a, Marr 1982] for reviews). For example, edge finding, though far from being a solved problem, is increasingly reliable and able to support the accurate computation of curvature (see below). Better edge finders are a prerequisite to improved shape representations. The additional information that they provide supports richer representations of shape, such as that discussed here. Much attention has focussed on three-dimensional vision. There has been significant progress in computing depth (or local surface orientation) using stereo, photometric stereo, structured light, and bounding contour. In other work, we build on these additional capabilities of computer vision to develop better representations of three dimensional shape [Brady and Yuille 1984].

A representation of shape can be judged according to its suitability for many uses in Robotics:

- *Recognition.* Given a database of models, recognition amounts to associating an object with its model. There are two principal difficulties. First, objects may overlap. An object has to be recognised from its non-overlapped portions and these portions cannot be predicted in advance. The shape of each portion needs to be represented, and provides constraint on model identification. The richer the representation of a visible fragment, the tighter the constraint. Recognition involves the summation of the constraints from the separate visible portions. It is best suited to a highly parallel computer such as the connection machine [Hillis 1981].

Representing a shape by a set of global features such as its moments cannot deal with overlap.

Second, in general, the database of models must allow for variations in shape. For example, not all tack hammers are exactly alike. Also, tack hammers are clearly distinguishable from Warrington hammers and ball peen hammers, but they are equally all clearly identifiable as hammers. Artificial Intelligence has developed semantic networks to represent classes and subclasses of objects and their subparts. In particular, Brooks and Binford [1980] developed ACRONYM to address such problems. This issue will be taken up again in Section 6.

- *Inspection.* Since a part to be inspected may have defects, it is unreasonable to expect a perfect match between a shape description and its model. We need to be able to recognize the model corresponding to a flawed object and to determine the nature of the flaw. The richer the representation of shape that can be computed for an object, the greater the possibility of automatically generating an explanation of the flaw from the mismatch between the object description and its model. This has been a recurring theme in Artificial Intelligence [Winston 1984]. Flaws are typically restricted to a local portion of a shape. The shape representation needs to be stable with respect to local changes, yet sensitive enough to represent them [Marr and Nishihara 1978]. We say that the shape representation must have local support.
- *Grasping.* Part of our understanding of a tool such as a screwdriver is how to pick it up, and how to grasp it when we use it. For a given end-effector, one can imagine motor program links in the model of an object that contain such information. However, faced with an object that we have never seen before, we still grasp it sensibly. In order to do this, we must determine a grasp position from a representation of the shape of the object to be grasped. Boissonat [1982] and Brady [1982b] have suggested ways in which this might be done.
- *Reasoning.* Robotics is the intelligent connection of perception to action. Intelligent robots need to be able to reason about the tasks that they are to perform, for example to plan processing and assembly sequences. Currently popular models of reasoning, such as expert systems, are typically supplied with pre-digested sensory information. Shape representations should support enhanced reasoning (see for example [Brady 1984b]).

In previous papers [Brady 1982b, 1982c, 1983] we have introduced a representation of two-dimensional shape that we call *smoothed local symmetries (SLS)*. Section 2 summarizes the main ideas, and then develops a parametric analysis of the SLS of a shape. The parametric description is both region and contour based. In Section 3, we describe an implementation of the representation, illustrating the performance of our programs in Section 5 using a set of hand tools. We discuss in Section 4 the reasons for computing a representation, called the *curvature primal sketch*, of the significant curvature changes along a contour. Asada and Brady [1984] discuss the curvature primal sketch in more detail. Finally, in Section 6 we sketch a method, currently being implemented, for determining the articulation of a shape into subshapes. The emphasis of this paper is on competence: we develop an implemented representation that can support a variety of applications without

excessive concern for "real time" performance. Hardware implementation of the ideas developed here will be discussed elsewhere.

2. Smoothed local symmetries

2.1. Background on shape description

Shape is represented in most current robotics or commercially available vision systems [Agin 1980, Holland, Rossol, and Ward 1979, Perkins 1978, Rosenfeld 1982] in terms of global features such as: the center of area, number of holes, the aspect ratio of the principal axes, and the ratio of the perimeter squared to the area. Such features can be computed efficiently and they are reasonably insensitive to noise and quantization, even for low resolution binary images. Unfortunately, the computed value of a global feature for the visible portion of an occluded object bears an arbitrary relationship to the value that would be computed for the whole object. It follows that it is at best difficult, and usually impossible, to recognize occluded parts using global features.

Bolles and Cain [1982] uses "focus features" to handle limited cases of occlusion. An object model is a graph whose nodes are features such as corners and holes and whose edges are the exact distance and relative orientation of the two nodes related by the edge. Recognition is implemented by finding maximal cliques in model graphs. In essence, Bolles' representation is a compromise between global features and representations that have local support, such as smoothed local symmetries. The representation is, however, metrically exact and uses only simple features such as rectangular corners, holes, and their relative position. It is restricted to straight line contours.

Nearly every representation of shape can be classified as being either *contour based* or *region based*. Contour and region based representations of shape make different aspects of shape explicit. Contour based representations encode the bounding contour of a shape. Examples include Fourier series expansions of the contour, chain encoding, and spline approximations (see Pavlidis [1978]). Asada and Brady [1984] have developed a representation that makes explicit the significant curvature changes along a contour. Attneave [1954] showed that such changes, essentially the curvature analogue of the primal sketch (see [Asada and Brady 1984]), carry a great deal of information for the human visual system. Certainly, the human visual system is remarkably sensitive to curvature changes [Watt and Andrews 1982]. It is difficult, however, to compute or determine important regional properties, such as symmetry, or descriptions such as "elongated and curved", in a purely contour based representation of shape.

Region based approaches represent a shape by encoding the two-dimensional space occupied by the shape. Examples include quadtrees [Rosenfeld and Kak 1982], the symmetric axis transform (SAT) [Blum and Nagel 1978], and (two-dimensional) generalized cones [Nevatia and Binford 1977, Brooks 1981]. The SAT and generalized

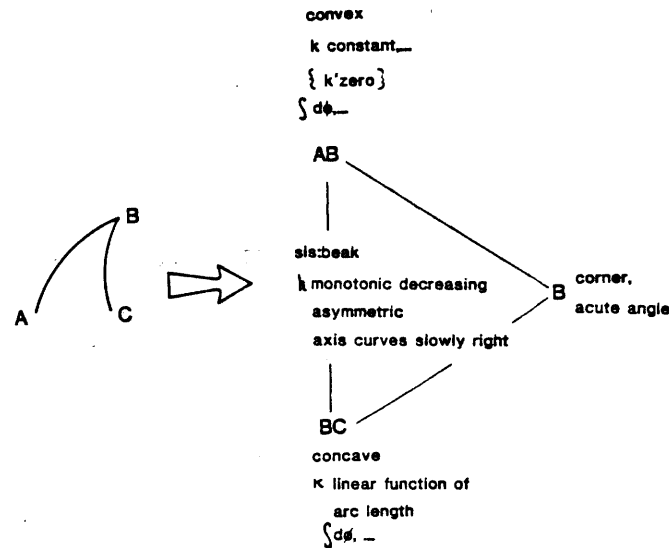


Figure 1. a. A hammer claw. b. Its geometric representation. The solid lines amount to a contour based representation, while the dotted lines encode the region occupied by the claw.

cylinders emphasize elongation, describing a shape by a one-dimensional axis and the disposition of the shape about that axis. Occlusion poses a difficult problem for purely region based representations, since recognizing partly occluded shapes requires a representation that has local support and a parametric description that is in part contour based. The smoothed local symmetries representation, described in the next section, is both contour based and region based.

2.2. Definition of smoothed local symmetries

Figure 1 shows a hammer claw and its geometric representation. The arc AB is described as a convex arc whose curvature is constant. In our implementations, arcs are described either by circles or by Cornu spirals, for which curvature varies linearly with arc length. The total angular change along AB is recorded. The discontinuity in curvature at B is made explicit. There are various kinds of primitive curvature discontinuities: in this case, B is an acute angled *corner*. Other primitive curvature discontinuities are defined below. They include *ends*, *cranks*, *smooth joins*, *bumps*, and *dents*. The arc BC is concave, and is also described by the best fitting circle. These descriptors form a contour based representation. The region based description is shown between the representations of AB and BC . It is described as a smoothed local symmetry of primitive type *beak*. Other primitive region types are defined below. They include *cup*, *sector*, *wedge*, *plinth*, and *flare*. The region description has several parameters associated with it, including a measure of the width of the region, and the curvature of an associated axis. In this case, we note that the width decreases monotonically in the direction approaching B , and has an axis that "slowly" curves to the right. Curvature is quantized to 5 levels, following Hollerbach [1975]. For a fuller account of parametric descriptions, see Heide [1984].

The definition of smoothed local symmetries is in three parts. First, Figure




Figure 2. a. The geometry of a local symmetry. b. The point A has local symmetries with points B_1 and B_2 .

2a shows the geometry of a *local symmetry*. A and B are points on the bounding contour of a shape. The angle α between the line BA and the outward normal n_A at A is the same as the angle between BA and the inward normal n_B at B . (A precise definition is given in the next section.) Figure 2b illustrates that in general, there may be several points B_i forming local symmetries with a given point A .

Implicit in the definition of local symmetry is the requirement that the bounding contour of the shape be extracted and the tangent angle computed sufficiently accurately. We use the edge finder developed recently by Canny [1983]. This extremizes the product of the signal-to-noise response at a step edge and a measure of the localization of the edge, while producing a single response to a single step. Earlier work used a non-directional operator based on a difference of Gaussians filter, proposed by Marr and Hildreth [1980]. The Canny operator is preferred because it gives more accurate tangent directions and is less susceptible to noise. For more details, see [Canny 1983].

Second, we consider the loci of local symmetries that are maximal with respect to forming a smooth curve. We call such loci *axes* or *spines*. Figure 3a shows the axis for a portion of a contour that contains a corner. Each axis is an alternative, locally plausible way to describe some piece of the contour and the region subtended by the axis. We call this portion of the shape, consisting of a part of the contour and a subregion of the shape, the *cover* of the axis. Notice that the covers of certain axes are wholly, that is, properly, contained in the covers of others. In such cases, we say that the axis with the contained cover is *subsumed* by the axis with the containing cover. Figure 3b shows the axes for a rectangle. The short, diagonal axes are subsumed by the horizontal and vertical axes, that are the global symmetries.




Figure 3. a. The smooth locus of local symmetries for a portion of contour containing a corner. b. The candidate axes for a rectangle. The small diagonal candidate axes are subsumed, leaving the (global symmetry) axes as smoothed local symmetries. c. The smoothed local symmetries of a square.

Notice that in case the rectangle were a square, the diagonal axes would not be subsumed (Figure 3c).

Figure 4 shows some examples of smoothed local symmetries. The smoothed local symmetries are indicated by their axes. It should be remembered that such axes are merely *graphic illustrations*, since they do not show the set of parameters and symbolic description that are associated with the smoothed local symmetry (see Figure 1 and Section 2.3).

Originally, Brady [1983, page 56] suggested that subsumed smoothed local symmetries should be deleted. Instead, we now accord them lesser importance in a hierarchical representation of the shape. The diagonal axes in Figure 3b, for example, provide information about the corners of the rectangle, which, though less important than the global symmetries of the shape, still convey useful information. For example, Figure 5a shows an acute angled corner discontinuity on a variety of contour fragments. The differences between the contour fragments is explicitly represented by region descriptors such as *beak* (see below). Figure 5b shows an occluded instance of a corner. The actual curvature discontinuity is not present, but it is implied by the monotonically decreasing width of the region and its small absolute value.

Smoothed local symmetries were developed from the ideas underlying the SAT and generalized cylinders. Brady [1983, pages 53 - 56] discusses the need to make axis smoothness explicit and to refine the definition of local symmetry implicit in the SAT. Most importantly, all of the local symmetry axes are made explicit for smoothed local symmetries.

The SAT is defined semi-constructively and topologically as the union of the




Figure 4. Examples of smoothed local symmetries.

centers of maximal disks that touch at least two points on the bounding contour of a shape. Figure 6 defines normal, branch, and end points of the SAT. The SAT is a piecewise smooth, single, forked spine. It is not obvious *a priori* that the SAT is piecewise smooth; but that has been shown by Bookstein [1979]. Branch points, where the maximal circle touches the contour in more than two points, are a major cause of the problems of the SAT, giving particularly poor responses when there is a discontinuity on the bounding contour of a shape. Figure 7 shows the SAT of a rectangle. Note that, unlike the smoothed local symmetry representation (Figure 3b), neither of the global symmetry axes are found. Blum and Nagel [1978, page 169] note that "the SAT is not the simplest description for rectilinear figures". Branch points were originally introduced into the definition of the SAT as a discriminant for isolating sub-shapes (the name comes from finding branches off the trunk of a tree). They do not perform that task reliably either. Some other differences between the SAT and smoothed local symmetries are presented in [Brady 1983].

Figure 8 compares smoothed local symmetries with the SAT. Figure 8a shows a triangle. The SAT makes explicit the center of the inscribed circle. The smoothed local symmetries are the median lines. If the triangle is isosceles, a median line is a global symmetry. The SAT of an ellipse is the major axis (Figure 8b), while the smoothed local symmetry makes both axes explicit. Figures 8c and 8d show how curvature discontinuities on the contour of a shape cause unintuitive branch points. The dotted lines are exterior portions of the SAT or smoothed local symmetries. As we shall see in Section 6, they are useful for determining subshape joins.




Figure 5. Several curvature discontinuities that subtend the same acute angle. Differences between the instances is found in the region descriptors. b. An occluded corner.

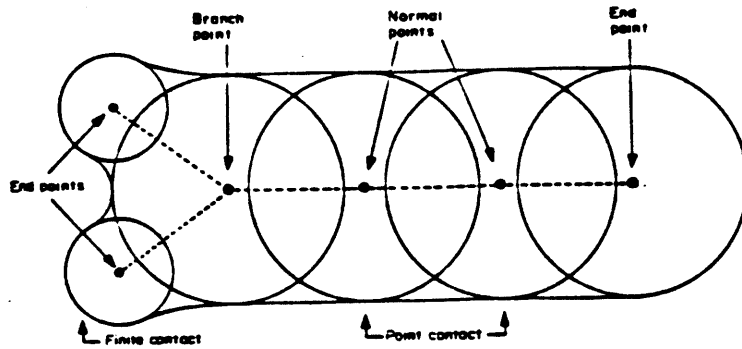


Figure 6. The symmetric axis transform. The centers of maximal circles that touch the contour at one, two, and more than two points are called end, normal, and branch points, respectively.

2.3. Analysis of smoothed local symmetries

Suppose that the point A in Figure 2a is r_A in some vector frame. Let the unit vector in the direction BA be $u(s)$ and let the length of AB be $h(s)$, where s




Figure 7. The SAT of a rectangle. Note that, unlike the smoothed local symmetry representation shown in Figure 3b, the SAT fails to make explicit either of the global symmetries of the rectangle. The SAT forks at branch points that signal the corners of the rectangle.

indicates distance along the smoothed local symmetry. The point P at which the local symmetry is recorded is $\mathbf{r}_P = \frac{1}{2}(\mathbf{r}_A + \mathbf{r}_B)$. By definition of the local symmetry, the angle α between \mathbf{u} and the outward normal at A is equal to that between \mathbf{u} and the inward normal at B . Aside from the global coordinate frame in which \mathbf{r}_A and \mathbf{r}_B are expressed, we introduce a moving frame with origin P , collinear with \mathbf{u} , and measure angles counter-clockwise from \mathbf{u} .

The tangent \mathbf{t}_A at A is given by

$$\mathbf{t}_A = \frac{d\mathbf{r}_A}{ds_A}.$$

Since ds_A is a unit speed parameter for the contour, \mathbf{t}_A is a unit vector. Similarly \mathbf{t}_B is a unit vector. Suppose that a step ds along the local symmetry axis at P corresponds to a step ds_A at A and a step $-ds_B$ at B . We use a minus sign to indicate a global direction of traversal of the curve. Bookstein [1979] shows that the steps at A and B are in opposite directions. We denote $\frac{ds_A}{ds}$ by s'_A , and $\frac{ds_B}{ds}$ by s'_B . Note that $s'_A \cdot s'_B$ is negative.

With respect to the moving coordinate frame centered on P , the tangent \mathbf{t}_B at B is at angle $\pi/2 - \alpha$ and \mathbf{t}_A is at $\alpha - \pi/2$. It follows that

$$\mathbf{u} \cdot (\mathbf{t}_A - \mathbf{t}_B) = 0.$$

Hence \mathbf{u} is perpendicular to the unit vector

$$\hat{\mathbf{t}} = \frac{(\mathbf{t}_A - \mathbf{t}_B)}{2 \cos \alpha}.$$




Figure 8. Comparison of smoothed local symmetries (left) and the SAT (right). a. A triangle. b. An ellipse. c. and d. Illustrating the deterioration of the SAT when curvature discontinuities cause spurious branch points.


We assume that $|\alpha| < \pi/2$ for a local symmetry. In general, \mathbf{t} is not the tangent \mathbf{t}_P to the smoothed local symmetry at P since

$$\begin{aligned} \mathbf{t}_P &= \frac{d}{ds} \mathbf{r}_P \\ &= \frac{1}{2} (\mathbf{t}_A s'_A - \mathbf{t}_B s'_B) \end{aligned} \quad (2.1)$$

Suppose that \mathbf{t}_P is at angle $\phi(s)$ counter-clockwise from \mathbf{t} , that is, at angle $\phi - \frac{\pi}{2}$ from \mathbf{u} . Then, forming the dot product of Equation (2.1) with \mathbf{t} gives

$$\|\mathbf{t}_P\| \cos \phi = (s'_A + s'_B) \frac{\cos \alpha}{2}.$$

By forming the cross product of Eq. (2.1) with \mathbf{t}_A and \mathbf{t}_B , we find



$$\frac{ds_A}{ds_B} = \frac{R}{r}$$

Figure 9. The smoothed local symmetry of two concentric circles of radii R and r . The ratio ds_A/ds_B is equal to R/r , and can be arbitrarily large.

$$\begin{aligned} s'_A &= 2 \frac{\|t_P \times t_B\|}{\|t_A \times t_B\|} \\ &= 2 \|t_P\| \frac{\sin \alpha + \phi}{\sin 2\alpha} \end{aligned} \quad (2.2a)$$

and

$$\begin{aligned} s'_B &= 2 \frac{\|t_P \times t_A\|}{\|t_B \times t_A\|} \\ &= 2 \|t_P\| \frac{\sin \alpha - \phi}{\sin 2\alpha}. \end{aligned} \quad (2.2b)$$

The ratio $s'_A/s'_B = ds_A/ds_B$ can be arbitrarily large, as can be seen in Figure 9, which shows the smoothed local symmetry of two concentric circles, the ratio of whose radii is large.

It follows from Eqs. (2.2) that

$$\frac{s'_A}{s'_B} = \frac{\sin \alpha + \phi}{\sin \alpha - \phi},$$

apart from the special case (like concentric circles) when $\alpha = \phi$. It follows that

$$\tan \phi = \tan \alpha \frac{s'_A - s'_B}{s'_A + s'_B}. \quad (2.3)$$

From $\mathbf{r}_A - \mathbf{r}_B = h(s)\mathbf{u}$, we deduce

$$\begin{aligned} h'(s) &= (\mathbf{t}_A \cdot \mathbf{u})s'_A + (\mathbf{t}_B \cdot \mathbf{u})s'_B \\ &= \sin \alpha (s'_A + s'_B) \\ &= 2\|\mathbf{t}_P\| \tan \alpha \cos \phi \end{aligned} \quad (2.4)$$

For example, if $\phi = 0$, as it is for a global or flexed symmetry, then $h' = 2 \tan \alpha$, as expected. Suppose that \mathbf{t}_A is in direction ϕ_A with respect to some base direction, and \mathbf{t}_B is in direction ϕ_B . Since $\phi_A - \phi_B = 2\alpha - \pi$, we find

$$\begin{aligned} 2\alpha' &= \phi'_A - \phi'_B \\ &= \kappa_A s'_A - \kappa_B s'_B, \end{aligned} \quad (2.5)$$

where, κ_A , respectively κ_B , is the curvature of the contour at A , respectively B . Similarly, the curvature of the smoothed local symmetry axis, κ_P is given by

$$\kappa_P = \phi' + \frac{1}{2}(\kappa_A s'_A + \kappa_B s'_B) \quad (2.6)$$

Of the parameters introduced so far, $\alpha, \alpha', h, h', \kappa_A$, and κ_B are directly available. In general, ds_A and ds_B are small, so their computation is poorly conditioned because of quantization. It seems that ϕ is inevitably computed indirectly from points P and the tangents at local symmetries. Finally, κ_P and ϕ' involve second differential quantities. We solve the linear equations (2.4) and (2.5) to find

$$\frac{s'_A}{s'_B} = \frac{h'\kappa_B + 2\alpha' \sin \alpha}{h'\kappa_A - 2\alpha' \sin \alpha} \quad (2.7)$$

so by Eq. (2.3)

$$\tan \phi = \tan \alpha \frac{h'(\kappa_B - \kappa_A) + 4\alpha' \sin \alpha}{h'(\kappa_A + \kappa_B)} \quad (2.8)$$

Consider, as a special case, a (non-flexed) global symmetry. The axis curvature κ_P is everywhere zero, as are ϕ and ϕ' . Also, $-\kappa_B = \kappa_A = \kappa$ say. From Equations (2.5) and (2.6)

$$\alpha' = \kappa_A s'_A = -\kappa_B s'_B.$$

We find that $s'_A = 1/\cos \alpha$, so that $\|\mathbf{t}_P\| = 1$, hence $h' = 2 \tan \alpha$, and $\alpha' = \kappa/\cos \alpha$. It follows that extrema of h occur when α is zero and they provide natural places to segment a symmetric object into its components. The points where $\alpha = 0$ do not coincide with extrema of κ_A in general. Figure 10 illustrates this. It shows a fragment of a symmetric shape whose curvature consists of a Cornu spiral smoothly

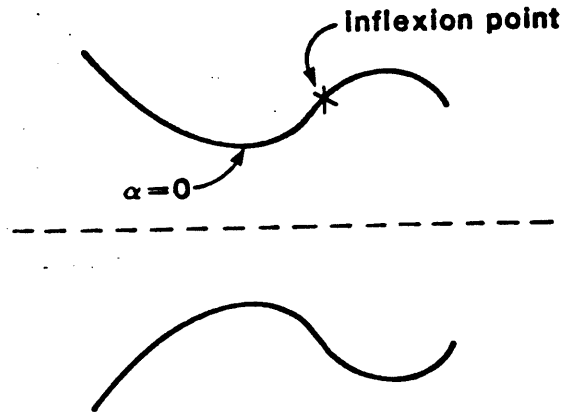


Figure 10. Segmenting a symmetric shape. The contour is the smooth join of a Cornu spiral and a circle. The inflexion point coincides with the curvature maximum of the upper part of the contour. The minimum width occurs where α is zero.

joined at an inflexion to a circle. The inflexion points are marked. The curvature of the upper part of the contour attains a maximum at the inflexion point. However, the shape fragment is more naturally segmented where h is smallest (compare [Hollerbach 1975]).

As a second special case, consider a *worm* [Blum and Nagel, 1978], for which h is constant (Figure 11). It follows that $h' = h'' = 0$, and that $\alpha = \phi = \phi' = 0$. We find that

$$\begin{aligned} s'_A &= \frac{\kappa_P}{\kappa_A} \\ s'_B &= \frac{\kappa_P}{\kappa_A} \\ \rho_P &= \frac{\rho_A + \rho_B}{2} \\ \rho_A &= \rho_B + h \end{aligned}$$

where $\rho = \frac{1}{\kappa}$ denotes radius of curvature.

Not surprisingly, the description of a worm amounts to the description of a plane curve that is the axis, and a statement that the cross-section function is constant. More generally, consider a flexed symmetry [Blum and Nagel, 1978]. It is easy to show that a shape is a flexed symmetry if and only if $\phi = 0$. The shape description is neatly factored into (a) a description of the shape of the flexed symmetry axis, (b) a description of the shape orthogonal to the flexed symmetry. That is, in the special case of a flexed symmetry, smoothed local symmetries coincide with (two-dimensional) generalized cones. In general, however, the smoothed local symmetry representation accords more closely with human perception than do generalized cones (see for example [Brady 1983, Figure 10]).




Figure 11. A flexed symmetry.

Indeed, it is clear that ϕ is a local measure of the *lack* of symmetry of a shape. This is because the angle between t_A and t_P is $\alpha - \phi$, while the angle between t_B and t_P is $\alpha + \phi$. The difference between these two angles is 2ϕ , and is zero for a flexed symmetry. Generalized cone parameterizations are not well suited to describing asymmetric shapes.

Interestingly, just the *signs* of h' , κ_A , κ_B , and α' allow us to define more than the basic shapes proposed by Blum and Nagel [1978]. Figure 12 assigns symbolic descriptors to primitive shapes, defined by the signs of the flanking curvatures. The *beak* primitive occurred in Figure 1. Brady [1983, page 62] compares the descriptors with those proposed by Blum and Nagel [1978]. The primitive shapes have associated descriptors, such as α' and h' , which define, for example, the rate of flaring of a flare shape. Heide [1984] has investigated the computation of these symbolic descriptors.

2.4. An $O(n^2)$ algorithm for computing smoothed local symmetries.

We have constructed one algorithm to investigate the competence of the smoothed local symmetry representation without worrying unduly about performance. The next section describes a more practical algorithm we are developing in a parallel effort. The time complexity of the algorithm described in the remainder of this section is $O(n^2)$, and consists simply of testing every contour point against every other to find local symmetries.

More precisely, consider *fixing* contour point A in Figure 13, and testing whether the point B forms a local symmetry with A . As in the previous section, we set up a local coordinate frame BA and measure angles counter clockwise. Denote the unit vector in the direction BA by $u(t)$, where t denotes arclength measured at B , and let the angle between u and the outward normal n_B be β . B forms a local










| | | K_A | | | |
|-------|-------|---|---|---|-------|
| | | $h' > 0$ | > 0 | $= 0$ | < 0 |
| K_B | > 0 |  cup |  sector |  beak | |
| | $= 0$ |  sector |  wedge |  plinth | |
| | < 0 |  beak |  plinth |  flare | |

Figure 12. Symbolic descriptors for primitive shapes defined by the signs of the curvature at the points A and B defining a local symmetry. The sign of h' defines the direction of the shape primitive. The values of α' and h' estimate the rate of increase of the cross section of the shape.

symmetry with A if and only if $\pi - \beta = \alpha$, which is if and only $\alpha + \beta - \pi = 0$. Define the function $f(t) = \alpha(t) + \beta(t) - \pi$. More properly, f should be denoted f_A to show that A is fixed. Finding local symmetries of A corresponds to finding zero crossings of $f(t)$.

The simplest version of the algorithm considers every point A , and tests every point B as a local symmetry with A . Figure 14 illustrates two problems that cause us to include in the implementation of local symmetry two cut-off tests at A and B . The low curvature of the contour between A and B in Figure 14a gives rise to the tiny candidate axis shown. In practice, there are many such tiny axes along the contour, and so we suppress them by requiring the angle between the normals at A and B to be above a threshold. In Figure 14b the curvature at point B is large, and so a small portion of the contour at B gives rise to local symmetries along a relatively long portion of the contour at A . We currently solve this problem by requiring the length $|BA|$ to be no greater than the radius of curvature at points A and B . Figure 15 shows some examples of candidate axes produced by the algorithm.




Figure 13. Geometry for determining a local symmetry in the $O(n^2)$ algorithm.




Figure 14. a. Local symmetries in low curvature portions of the contour. b. Local symmetries from a high curvature portion of the contour.

One way to make the $O(n^2)$ algorithm more useful in practice is to sample the contour (Figure 16). This is particularly useful for long portions of the contour where the curvature is small. It follows from the analysis in the previous section that the curvature of the axis is also small. By the mean value theorem, if $f(a) \cdot f(b) < 0$ there is a zero-crossing between $t = a$ and $t = b$, which can be found by binary subdivision. If $f(c) \cdot f(d) > 0$, we cannot conclude that there is no zero-crossing between c and d . We can approximate f between c and d , say by a function f_{approx} , and sample more closely to find a zero crossing of f only if there is a zero crossing of f_{approx} . We can use a cubic approximation to f if we can estimate $f'(c)$ and $f'(d)$. Now $f'(t) = \alpha' + \beta'$. It can be shown that $\alpha' - \beta' = -\kappa_B$. Also, $\beta' = \kappa_B + (h'/h) \cot \beta$, from which we find

$$f' = \kappa_B + 2 \frac{h'}{h} \cot \beta.$$




Figure 15. Examples of candidate axes found by the algorithm.




Figure 16. Sampling the contour to increase the speed of the $O(n^2)$ algorithm. The dotted line indicates an approximation to the axis computed from the sampled points.

In the next section we discuss a more ambitious algorithm that uses a curvilinear approximation to the contour.

An axis describes a portion of the region subtended by the shape, and does so by interpreting the relationship between two portions of the contour of the shape.




Figure 17. Example of the smoothed local symmetries found by Heide's program. a. The axes found for the carving knife. b. The cover of the blade axis. c. The cover of the handle axis.

Let A be a candidate axis, and denote by Λ_A and Φ_A the left and right portions of contour that give rise to the contour. Just as an individual point can have several local symmetries, so a portion of the contour of a shape can be the left or right portion of several candidate axes. Let \mathcal{R}_A be the region subtended by A . Now suppose that B is another candidate axis, and suppose that Λ_B overlaps Λ_A . We say that A *subsumes* B if Λ_B is contained in, but not equal to, Λ_A and if \mathcal{R}_B is contained in \mathcal{R}_A . The mapping from axis to contour is in fact a diffeomorphism [Bookstein 1979], and so subsumption can be determined from the contour. Heide [1984, forthcoming] has implemented the subsumption operation, and demonstrated its use on a variety of shapes. Figure 17 illustrates the performance of Heide's algorithm on an image of a carving knife. Figure 17a shows the axes of the three smoothed local symmetries found by the program; Figure 17b shows the cover of the blade, and Figure 17c the cover of the handle.

3. Computing smoothed local symmetries by approximating the contour

We are developing a program that computes the SLS of a shape by approximating the contour. One reason for doing this is that the computation of smoothed local symmetries can be noisy; computing smoothed local symmetries from an approximation to the contour can significantly reduce the noise. The basic approach is easily stated: choose a set of points on the contour, and then use them to construct a piecewise smooth approximation to the contour using the chosen points as "knots". We then compute the smoothed local symmetries of the approximation to the contour. We currently approximate the contour using straight lines and circles. Figure 18 shows some example smoothed local symmetries found using the algorithm.




Figure 18. Candidate smoothed local symmetries found using the approximation algorithm. a. hammer. b. a pair of pliers. c. a carving knife.

The knot points should not be chosen completely at random. Consider for example a corner (Figure 19). The knot should be placed at the corner rather than being placed on one side of it, in which case a smooth approximating curve has to include the discontinuity at the corner. More generally, the set of knots should mark significant changes in the curvature along the contour. Finding significant curvature changes is a difficult problem exactly analogous to computing a rich representation of intensity changes in one direction across an image ("the primal sketch" of [Marr 1976]). In the next section we sketch a method for computing such a representation of curvature changes. Asada and Brady [1984] give more details.

Once we have a set of knot points, we can compute a piecewise smooth approximation to the contour. Asada and Brady [1984] observe that, at least for the class of tool shapes that we have considered so far, locating the knot points accurately is more important than the choice of smoothing functions to approximate the contour between the knots. To date, we have found that fitting circles between the knots is efficient and gives good results (see [Asada and Brady 1984] for examples).

We try to fit circles to the points $\{(x_i, y_i) | 1 \leq i \leq n\}$ between adjacent knot points. Suppose the circle is $(x - x_0)^2 + (y - y_0)^2 - r^2 = 0$. We form the square error term

$$E(x_0, y_0, r) = \sum_{i=1}^n \left\{ (x_i - x_0)^2 + (y_i - y_0)^2 - r^2 \right\}^2,$$

and minimize E with respect to its parameters. We have

$$r^2 = \frac{1}{n} \left\{ \sum (x_i - x_0)^2 + \sum (y_i - y_0)^2 \right\},$$




Figure 19. Knot points should mark significant changes of curvature along the contour. a. A corner fragment on a contour. b. Approximation when knot points (marked by crosses) are placed to either side of the perceived corner. c. Approximation when the corner is made explicit.

and, after algebraic manipulation, find that two linear equations must be solved for x_0 and y_0 :

$$\begin{aligned} x_0 \sum (x_i - \bar{x})^2 + y_0 \sum (x_i - \bar{x})(y_i - \bar{y}) &= \frac{1}{2} \sum (x_i^2 + y_i^2)(x_i - \bar{x}) \\ x_0 \sum (x_i - \bar{x})(y_i - \bar{y}) + y_0 \sum (y_i - \bar{y})^2 &= \frac{1}{2} \sum (x_i^2 + y_i^2)(y_i - \bar{y}) \end{aligned}$$

where \bar{x} is the average of the x_i , and \bar{y} the average y_i . The circular approximation can become poorly conditioned when the radius of the circle is large (equivalently, the curvature is small), in which case the contour is better approximated by a straight line. Denote the matrix of coefficients of x_0, y_0 by \mathbf{A} :

$$\mathbf{A} = \begin{bmatrix} \sum (x_i - \bar{x})^2 & \sum (x_i - \bar{x})(y_i - \bar{y}) \\ \sum (x_i - \bar{x})(y_i - \bar{y}) & \sum (y_i - \bar{y})^2 \end{bmatrix}$$

The condition number of the matrix $\mathbf{A} = (a_{ij})$ is defined (see for example [Strang 1980, p. 282]) to be the ratio $\lambda_{max}/\lambda_{min}$ of its eigenvalues. Since the eigenvalues of \mathbf{A} are the roots of a quadratic in the coefficients, it is more convenient to work with the ratio:

$$\begin{aligned} L &= \frac{|\lambda_1 - \lambda_2|}{\lambda_1 + \lambda_2} \\ &= \frac{\sqrt{(a_{11} - a_{22})^2 + 4a_{12}^2}}{\text{Tr}(\mathbf{A})} \end{aligned}$$




Figure 20. Geometry for local symmetry of a point P and a circular arc. The point (x_c, y_c) is the center of the circular arc, and R is the local symmetry of the point P and the circular arc.

where the λ_i are the eigenvalues of A . If L is less than a threshold (currently 0.96), corresponding to the condition number of A being less than 5, we accept the circular approximation, else we fit a straight line to the points (x_i, y_i) .

In the previous section we described a $O(n^2)$ algorithm for computing the smoothed local symmetries of a shape. For each point along the contour we compute the set of local symmetries. We can use the contour approximation to speed the computation by a large factor. Let A be an arbitrary point on the contour. Instead of comparing A against all other points along the contour, we compare A against each of the circular arcs that approximate the contour. Figure 20 shows the geometry for determining the local symmetry of a point and a circular arc. The Figure treats the case of positive curvature κ of the circle. We assume that the point A is at the origin and that the x -axis is aligned with the tangent at A . Let the center of the circular arc be (x_0, y_0) , and let the radius be r . Let the local symmetry between A and the point B on the circular arc occur at angle θ measured counter clockwise from the radius that is perpendicular to the y axis. It can be shown that

$$\theta = 2 \tan^{-1} \frac{x_0 + r + y_0}{x_0 - (r + y_0)}$$

$$|AP| = \frac{r^2 - (x_0^2 + y_0^2)}{2(r + y_0)}$$

In these expressions it is assumed that $0 < |AP| < r$, which is the curvature cut-off condition referred to above. Similar analyses are possible in the case $\kappa < 0$ and for a straight line segment. Figure 21 shows the result of this on the image of a pair of pliers. The global line of symmetry is quite visible, but is noisy. The problem is that a single point can form local symmetries with several other points because of discretization and round-off.




Figure 21. Approximating the smoothed local symmetries of a pair of pliers. The local symmetries are produced by comparing each point on the contour against the circular arcs forming the approximation. The result should be compared to Figure 22b, which is the result of computing smooth local symmetries from pairs of approximating arcs.

We can further reduce the noise in the computation of the smoothed local symmetries, as well as speeding their computation, by directly comparing two circular arcs and computing an analytic expression for the smoothed local symmetry. This allows successive axes to be chained together. Figure 22a shows the geometry for computing the analytic form of the smoothed local symmetry of two circular arcs. There are various cases to consider, depending on the signs of the curvature of the arcs. In region I (see Figure 22a) we can show that the locus of the intersection of the normals at local symmetries is an ellipse. In region II the locus is a hyperbola. Figure 22b shows the smoothed local symmetries of the pliers (Figure 21) computed by comparing arcs directly.

4. Representing significant changes in curvature

There are two reasons for discovering and representing the significant changes in curvature along a contour. First, they provide a set of knot points for constructing a perceptually close approximation to the contour. Second, as we pointed out in Section 2.1, smoothed local symmetries are both a region based and contour based representation of shape. A contour based representation requires that the curvature changes be detected, localized, and described correctly. This section sketches a representation of the significant curvature changes along a contour, a representation that we call the *curvature primal sketch*. Further detail can be found in [Asada and Brady 1984].

Our approach follows that of Marr [1974]. We take ideal parameterized versions of a particular set of curvature changes as our models $f_i(s)$, where f_i is described




Figure 22. a. Geometry for computing the smoothed local symmetry from two circular arcs. b. The smoothed local symmetries of the pliers computed by directly comparing pairs of circular arcs.

as a function of the orientation of the tangent to the contour. We then determine analytic forms for $G''_{\sigma} * f_i(s)$, where

$$G_{\sigma}(t) = \frac{1}{\sqrt{2\pi}\sigma} \exp -\frac{t^2}{2\sigma^2}$$

is a Gaussian of standard deviation σ . We use Gaussian second derivatives G''_{σ} (approximately equal to the difference of Gaussians [Marr and Hildreth 1980]) as filters because of the analogy between curvature and intensity changes. The algorithm filters the contour with a set of Gaussians of increasing standard deviation (scale) to construct a *scale space* representation of the contour [Witkin 1983, Yuille and Poggio 1983]. Using the analytic form of $G''_{\sigma} * f_i$ as a guide, we construct a program to find instances of the curvature change model along actual contours. We have identified a set of models, developed the analytic form of their convolution with the second differential of a Gaussian at several scales, and implemented an algorithm that matches models against the differentiated curvature along an actual contour.

Figures 23, 24, and 25 show the set of ideal curvature changes that form our current set of models. There is an exact analogy between the problem of representing the significant curvature changes along a contour and the "primal sketch" representation advocated by [Marr 1982] for significant intensity changes




Figure 23. a. The *corner* model consisting of two circular fragments of curvatures κ_1 and κ_2 enclosing an angle ϕ . b. The corner model in orientation space, relating the orientation of the tangent to the curve to distance along the curve. c. The result of convolving the corner model with the second derivative of a Gaussian.




Figure 25. The remaining models in orientation space. a. The *end* b. The *dent* or *bump*. c. The *smooth join*, which is often difficult to detect or localize.

in a single direction across an image. For example, the *corner* shown in Figure 23 is analogous to a step change in intensity, while the *crank* shown in Figure 24 is analogous to a thin bar.

Consider the corner fragment, defined (Figure 23b) by

$$f_{corner}(s) = \begin{cases} \kappa_1 s + c & \text{if } s < 0; \\ \kappa_2 s + c + \phi & \text{if } s > 0. \end{cases}$$




Figure 24. a. The *crank* model consisting of three circular fragments of curvatures κ_1 , κ_2 and κ_3 enclosing angles ϕ_1 and ϕ_2 , the radius of the crank being $2a$. b. The crank model in orientation space c. The result of convolving the crank model with the second derivative of a Gaussian where $\sigma/a \geq 1$.

It is tedious, though straightforward, to show that

$$\sqrt{2\pi}\sigma(G_\sigma * f_{corner})''(s) = \frac{-\phi}{\sigma^2} s \exp\left(-\frac{s^2}{2\sigma^2}\right) + (\kappa_2 - \kappa_1) \exp\left(-\frac{s^2}{2\sigma^2}\right).$$

As Figure 23c shows, this function has a zero crossing near the corner separating two peaks of opposite sign. The distance along the s -axis between the zero crossings is

$$d_{corner} = \frac{\sigma^2}{\phi} \sqrt{(\kappa_1 - \kappa_2)^2 + \frac{4\phi^2}{\sigma^2}}$$

The formula for the height of the side lobes is complex. In the special case $\kappa_1 = \kappa_2$, it reduces to

$$h_{corner} = \frac{1}{\sqrt{2\epsilon\pi}} \frac{|\phi|}{\sigma^2}$$

Figure 26 shows some typical variations of the corner model.



Figure 26. Variations of the corner model.

As a second example, consider the crank model, whose analytic form (see Figure 24b) is

$$f_{crank}(s) = \begin{cases} \kappa_1(s+a) + c & \text{if } s < -a; \\ \kappa_2(s+a) + c + \phi_1 & \text{if } -a < s < a; \\ \kappa_3(s-a) + (2a\kappa_2 + c + \phi_1 + \phi_2) & \text{if } s > a. \end{cases}$$

The crank is analogous to a "thin bar" intensity change in the intensity primal sketch. It is even more tedious to show that

$$\begin{aligned} \sqrt{2\pi}\sigma(G_\sigma * f_{crank})''(s) = & \frac{-1}{\sigma^2} \left\{ \phi_1(s+a) \exp -\frac{(s+a)^2}{2\sigma^2} \right. \\ & \left. + \phi_2(s-a) \exp -\frac{(s-a)^2}{2\sigma^2} \right\} \\ & + (\kappa_3 - \kappa_2) \exp -\frac{(s-a)^2}{2\sigma^2} \\ & + (\kappa_2 - \kappa_1) \exp -\frac{(s+a)^2}{2\sigma^2} \end{aligned}$$

The qualitative form of the function $(G_\sigma * f_{crank})''(s)$ depends on the relative values of its parameters, in particular the degree of smoothing σ and the contour length a of the step of the crank. Figure 27 shows typical variation when $\kappa_1 = \kappa_2 = \kappa_3 = 0$. When the ratio σ/a is less than 1/2, the crank is essentially signalled by two independent corner responses. However, when σ/a is greater than one, it produces a strong central peak with two side peaks of opposite sign that are at most half the height of the central peak.




Figure 27. Second derivative of a crank after Gaussian smoothing. The edges of the crank are assumed to be straight. The responses vary according to the ratio σ/a , the exact values being shown to the side of the figures.




Figure 28. Response of a straight line end as it varies with σ/a . We assume that each κ_i in Figure 23 is zero, and that the $\phi_j = \pi/2$. The responses are for $\sigma/a = 0.5, 0.75, 1.0, 1.25$.

We have similar analyses for the three models shown in Figure 25. A smooth join generates a peak of height $(\kappa_2 - \kappa_1)/(\sigma\sqrt{2\pi})$. Only when the curvatures flanking the smooth join are very different can this be reliably found. This is consistent with human perception of such smooth joins. Figure 28 shows the second differential of a Gaussian convolved with an end whose edges are straight. As in Figure 27, the response varies with σ/a . Finally, Figure 29 shows the response to a bump. If the bump is long enough, the response is equivalent to two separate edges of opposite sign.




Figure 29. Response of a bump or dent. The bump in the case shown here is circular, with initial and final angles $\pi/4$.

Figure 30a shows the second derivative of the smoothed orientation at a variety of scales σ along the contour of a hammer. The model instances described above are marked on the graphs shown in Figure 30b. In Figure 30c we show the knot points that mark the significant curvature changes along the contour of the hammer at the different scales. Finally, in Figure 30d we show the smoothed local symmetries resulting from applying the approximation algorithm described in the previous section given the knot points shown in Figure 30c.

5. Examples

In this section we show some examples of the smoothed local symmetries computed for a set of hand tools. Figure 31 shows a screwdriver. The global axis of symmetry is found by the algorithm. Note that at the coarsest scale, the two corners forming the end of the handle, and the (primitive) end of the blade, are discovered, as are the two cranks defining the join of the handle and blade. Figure 32 shows a tack hammer. The head of the hammer, consisting of the two striking surfaces, is easily found. So is the handle. The join of the handle and the head is a side-end join, and can be found using the method sketched in the next section. Figure 33 shows a carving knife (compare Figure 17) and a cheese knife.

6. Finding sub-shapes

The shapes that populate our world have identifiable sub-shapes. A hammer, for example, has a head and a handle. The head consists of a striking part, that has a flat surface, and a counterweight part, whose shape varies considerably between

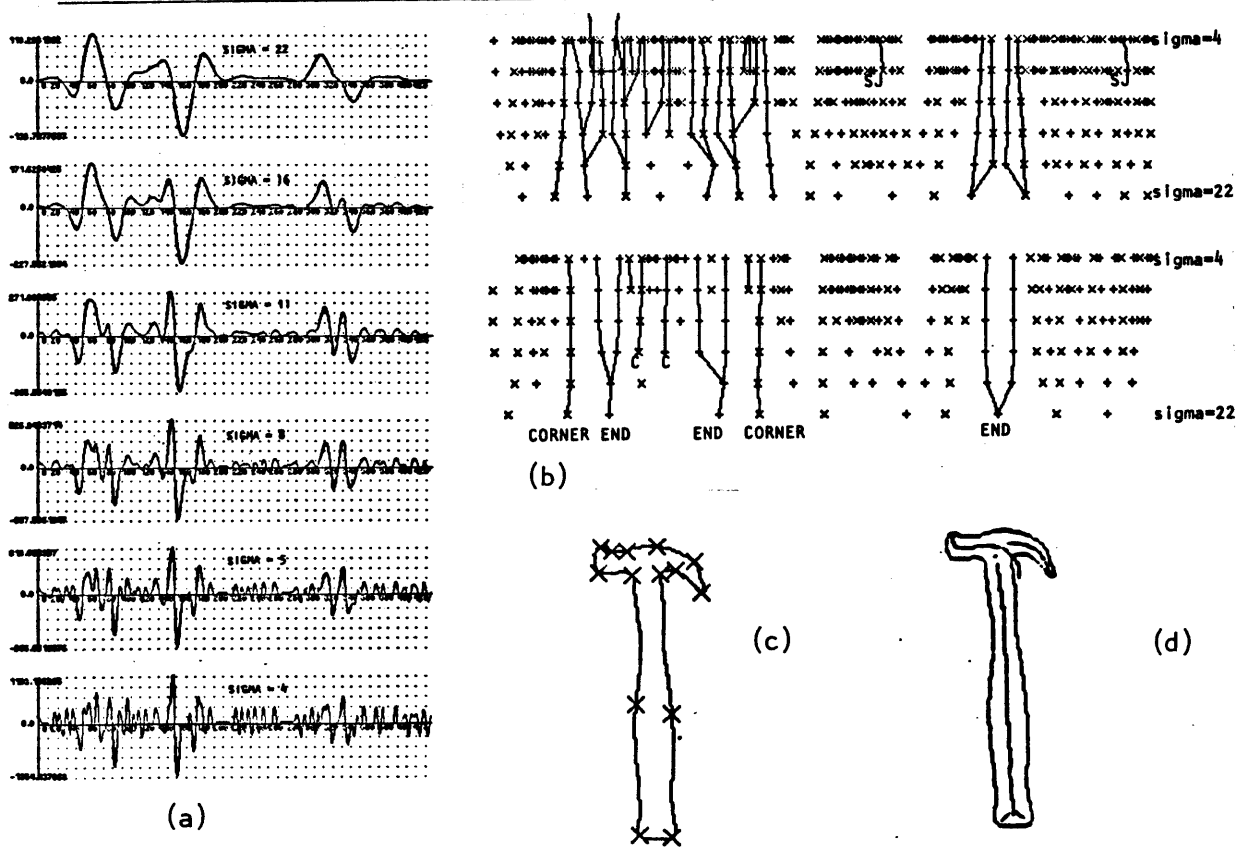


Figure 30. a. The second derivative of the orientation along the contour of a hammer, smoothed at a variety of scales, corresponding to $\sigma = 5, 8, 11, 22$. b. The model instances are shown in a tree diagram as they are tracked over the different scales. c. The knot points that correspond to the model instances in (b). d. The smoothed local symmetries resulting from applying the approximation algorithm described in the previous section given the knot points shown in (c).

hammers. Our current work aims to make sub-shapes explicit. In this section, we report our progress. Figures 34 through 36 show the representation that we are aiming at. In Figure 34 we show how sub-shape information at a single scale, indicated by the *part-of* links, is integrated with the smoothed local symmetry representation discussed in the previous sections. Figure 35 shows how a shape can be represented at several scales.

Figure 36 suppresses geometric detail to show other kinds of relationships, particularly *part-of* relationships, *a-kind-of* specialization relationships, and *constraints* between elements of the representation. *A-kind-of* links indicate class constraints. *Instance* links enable us to index model parts from smoothed local symmetry representations of shape fragments. Other attributes of a shape are under investigation. One such is a set of *motor programs* that suggests for a model such as a hammer how it might be used and for what purpose. *Part-of* and *a-kind-of*





Figure 31. The smoothed local symmetries found for a screwdriver.




Figure 32. The smoothed local symmetries found for a tack hammer and a Warrington hammer.

relationships have been explored extensively in knowledge representation studies in Artificial Intelligence. In particular, ACRONYM [Brooks and Binford 1980] is a major contribution to the development of representations of object models, and it has heavily influenced our work at the level indicated in Figure 36. All parts in ACRONYM models are (simple instances of) generalized cones, whereas in our work they are smoothed local symmetries, as shown in Figures 34 and 35. Following ACRONYM, we incorporate constraints, in the form of bounds, between links. For example, there are bounds on the relative sizes of the head and handle of a hammer. It would be surprising to find a sledge hammer with a two-inch long handle, or to find a half-inch wrench with a six foot handle. An important part of ACRONYM




Figure 33. The smoothed local symmetries found for a carving knife and a cheese knife.

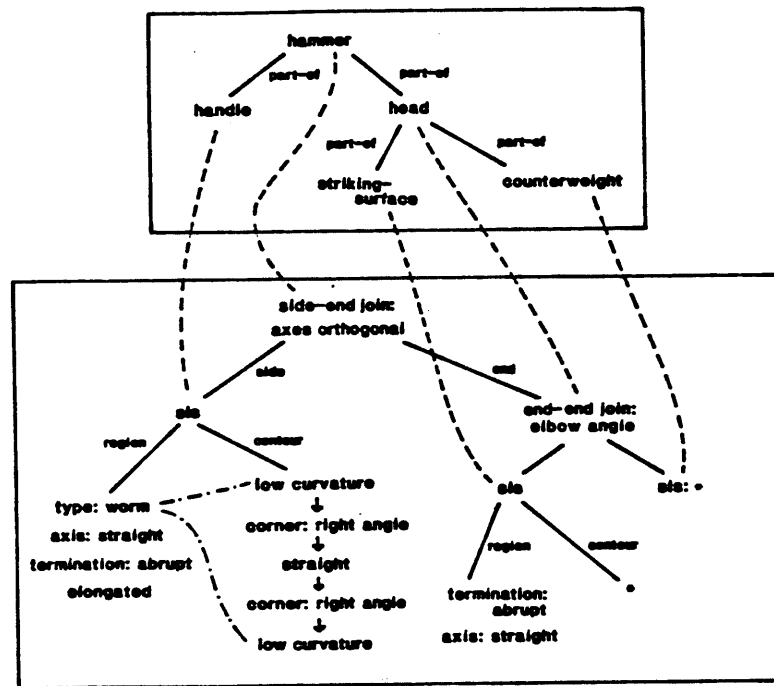


Figure 34. The representation of the shape of a hammer. The integration of *part-of* links defining sub-shapes with the smoothed local symmetry representation of the contours of and regions subtended by the sub-shapes.

was the program of Brooks [1982] for determining bounds on expressions given bounds on their variables.

We have not yet investigated the incorporation of yet other sorts of information into our models. Such information might make explicit surface markings or textures,




Figure 35. The hammer is represented at several scales, here two. At the coarsest scale, the handle is simply represented as a symmetric shape with constant cross section. At the finer scale, the shapes of the handle and head are refined.

for example. As we noted in the Introduction, in practical applications, the model database could be very large.


Our representation proposes four major modifications to the ACRONYM representation. First, we use smoothed local symmetries instead of generalized cones. The reasons for this are developed in [Brady 1983]. Second, we represent a shape at multiple scales, for the reasons discussed in Section 5 and in [Asada and Brady 1984]. Third, we incorporate additional links, particularly *instance* links and *motor programs*, into the representation. Finally, unlike ACRONYM, we explicitly represent sub-shape connections (Figure 34). As we discuss more fully below, sub-shapes can be joined together in a variety of ways, in particular in ways we call *side-end* and *end-end* [Marr 1977].

In general, deciding that a portion of a shape is a sub-shape requires the mobilization of application-specific, semantic information. Here, we are concerned with lower level processes, embodying weaker semantics, that nevertheless isolate sub-shapes in a usefully wide class of instances. Several methods have been proposed to determine the sub-shapes of a shape.


For example, the original motivation for the concept of branch point in the symmetric axis transform was to determine sub-shapes. The idea does not work well in practice. First, as we noted in Section 2.2, branch points often occur when there is an orientation discontinuity on the contour, in which case they do not signal



Figure 37. A side-end join, one of whose concavities is not signalled by a significant negative minimum of curvature.



(a)



(b)

Figure 38. a. Many shapes have alternative descriptions. The shapes shown on the left and right are typically interpreted as shown; the center shape is ambiguous. b. A related problem concerns ambiguity about what is joined to what.

join. We postpone investigation of several important questions, such as: how to choose among different possible "parsings" of a shape into sub-shapes (Figure 38a); and how to determine which is the sub-shape and which the main shape (Figure 38b). Figure 39 shows our model of sub-shape formation: two shapes may overlap with optional smoothing of their contours. Alternatively, the shapes may be close and their contours smoothly joined. Notice that it is the smoothing of a contour near a real subpart join that invalidates naive expressions of the matched concavity heuristic.




Figure 39. Model of (side-end) join formation. The contours of two overlapping or nearby shapes may be smoothly joined.

Terminations

In order to assert subparts, we need to define terminations of a shape, so that we can distinguish side and end. In our approach, a shape is represented by a (finite) smoothed local symmetry. By definition, a smoothed local symmetry axis is a maximal locus of local symmetries. It follows that a smoothed local symmetry axis partitions a shape into two *sides*, which are those portions of the contour that generate the axis, and two *terminations*, which are the remaining parts of the contour. Switching attention between different smoothed local symmetries correspondingly changes the parts of a contour that are perceived as side and perceived as end. Sides are typically elongated, which is an alternative way of saying that, all other things being equal, we prefer descriptions based on a shape's longest axes. Notice that the description of a join is *not* a symmetric relationship between sub-shapes. In Figure 40, the "handle" *A* is on the side of *B* but *B* is on the end of *A*. Such alternative descriptions can be extremely useful in adequately characterizing a join.

Currently, we distinguish four types of termination (Figure 41). A termination is considered *open* if it is the distal end of a negative smoothed local symmetry corresponding to a concavity. Typically the contour "turns away" at an open termination. All other terminations are *closed*, and are of three types. First, in a *tapered termination* (Figure 41a) the width function $h(s)$ slowly tends to zero. Typically, the symmetry angle α is nearly constant, and a tapered termination gives rise to a (single) curvature discontinuity of type *end*. Second, for a *blunt termination*, the width function tends to zero but "quite abruptly" (compare Hollerbach's [1975] and Heide's [1984] use of symbolic qualifiers corresponding to ranges of values). Typically, there is an abrupt change in the symmetry angle α . A blunt termination usually generates a curvature discontinuity that is an *end* at small scales but a




Figure 40. *A* is on the side of *B* but *B* is on the end of *A*.




Figure 41. Three types of closed termination of a smoothed local symmetry. a. tapered. b. blunt. c. abrupt.

pair of *corners* at larger scales. Note that both tapered and blunt terminations are signalled by curvature discontinuities of type *end*. Finally, an *abrupt termination* (Figure 41b) gives rise to a pair of curvature discontinuities of type *corner*. In general, nothing can be asserted about $h(s)$ or α . Of course, there is no absolute distinction between a blunt and tapered end, or between a blunt and tapered end. Positive smoothed local symmetries ("figure") have two closed terminations. Negative ones ("ground"), corresponding to concavities in a shape, have an open distal termination and a closed proximal one.

Subpart Joins

We propose two rules that assert the existence of a subpart join, prior to determining whether the join is of type side-end or end-end. As we have noted, and as we shall see further below, this is because the notions of side and end are intimately associated with (the global property of) elongation. Deciding that part of a shape is side or end derives from choosing an sls to organise one's description of a shape.

The first rule generalizes the familiar matched concavity heuristic:

Rule 1

if there are two corresponding concavities,
then the cover of the sls associated with the concavities is a subpart.

The crucial problem is *how to assert the existence of concavities*. There are various possibilities, and they rely on the fact that our representation is both contour and region based. First, there may be a significant curvature change of type corner. Second, there may be a primitive curvature change of the crank type that indicates a (small) concavity. This suffices, for example, to isolate the handle and shaft of a screwdriver.

Finally, a concavity can be asserted even if there is no significant curvature change, if the contour is corrupted by noise, or if there is a significant curvature change that is overlapped. In such cases we rely on the region component of our representation. A concavity is a "negative" shape that has a smoothed local symmetry axis. Figure 42a shows a matched pair of concavities and the graph of contour fragments and smoothed local symmetry axes to which it gives rise. We assert the existence of a subpart, whose sls is μ if (i) we find an instance of the sub-graph shown in Figure 42a; and (ii) the contour fragments *a* and *d* are smoothly extensible without a significant change in curvature.

There is an analogous rule for matched convexities asserting the existence of a negative subpart such as an inlet (see Figure 42b). For example, the axes λ and ν in Figure 42a, corresponding to the concavities, have been labelled with a minus sign. Notice that reversing the sense of figure and ground corresponds to interchanging minus and plus. Figure 42b is interpreted as a subpart that is a negative shape, that is, an inlet.

Suppose that Rule 1 is applicable. The subpart it defines is called a *limb* if it is elongated and has a closed termination. If it is not elongated, it is described as an *attached subpart*. Finally, if it does not have a closed termination it is called a *subpart join*. Figure 43 shows examples.

Rule 1 is sufficient but not necessary for asserting the existence of subparts. Figure 44 shows three example shapes in each of which there is only one concavity. Figures 44a and 44b are end-end joins, but Figure 44c is the side-end join shown in Figure 37. Figure 45 shows once more the crucial role that elongation plays in asserting the existence of a subpart. The fragment shown in Figure 45a is perceived as a join in Figure 45b but not in Figure 45c.

Rule 2




Figure 42. a. A concavity is asserted if there is an instance of the sub-graph of the smoothed local symmetry representation shown on the right. b. Interchanging plus and minus enables us to assert negative subparts such as "inlets".




Figure 43. a. A limb. b. an attached subpart. c. a subpart join. See text for details.

if there is a concavity,
 and there is an sls, one of whose generating contours flanks the concavity,
 such that the sls is either elongated or has a closed termination,
 then the cover of the sls is a subpart.

If the subpart has a closed termination, it is described as a *limb* if it is elongated or an *attached subpart* if it is not. Otherwise the (elongated) subpart is called a *join*. Figure 44 shows examples of object joins. As in the case of Rule 1, there is an analogous rule for convexities that signal inlet subparts. Other rules are under




Figure 44. Matched concavities provide a sufficient, but not necessary, for asserting the existence of subparts. a. an end-end join. b. an end-end join generated by a crank. c. the side-end join shown in Figure 37.

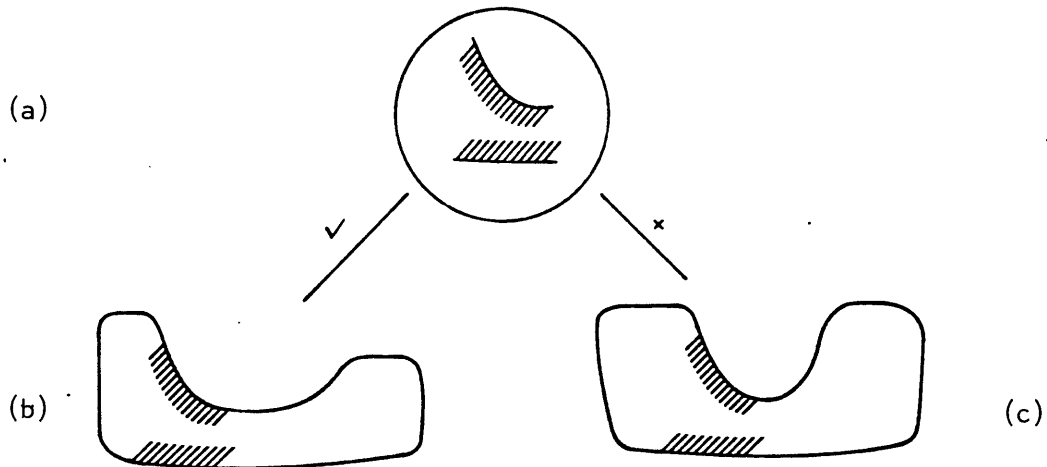


Figure 45. a. a shape fragment. b. the shape fragment in (a) corresponds to a subpart join. c. the shape fragment in (a) does not correspond to a subpart join.

investigation.

Recursive application of the rules

When a subpart is asserted, the smoothed local symmetries of the two parts are then (re-)computed. (Of course, there are efficient ways to use the existing information.) The crucial point is that we *inhibit* local symmetries between points that are from different sub-shapes. This removes many apparently spurious axes from the initial investigation of a compound shape. The process is applied recursively. This generates a further hierarchical shape organisation that is similar to that




Figure 46. a. A notched rectangle. b. The smoothed local symmetry axes for the notched rectangle. c. The result of inhibiting spines from the rectangle and notch shapes found by rule 1.

advocated by Marr and Nishihara [1978].

Consider, for example, Figure 46. It shows the notched rectangle, discussed by many authors in connection with the symmetric axis transform. Figure 46b shows the smoothed local symmetry axes computed for the notched rectangle. There are several small axes that seem to be spurious. There is also a negative inlet (triangle) shape that can be found using Rule 1. Crucially, every spurious axis derives from a portion of the contour of the inlet subpart and a contour of the rectangle. Inhibiting them leaves precisely the description: an inlet into a rectangle.

7. Conclusion

We have introduced smoothed local symmetries as a representation of two-dimensional shape that is both contour and region based and which has local support. We described two implementations, one of which efficiently computes an approximation to the smoothed local symmetry representation, and relies upon computing the curvature primal sketch representation of the significant curvature changes along a contour. The previous section reported some initial findings on the isolation of subobjects.

Currently we are developing a database of parametric models of objects that enable them to be recognized even when they are occluded. We are investigating the relationship between function and form to support reasoning about objects.

Acknowledgements

The authors thank Steve Bagley, Ruzena Bajcsy, Dana Ballard, Bob Bolles, Dave Braunegg, Chris Brown, John Canny, Olivier Faugeras, Marek Ghallab, Georges Giralt, Laurent Gouzenes, Eric Grimson, Scott Heide, Tommy Poggio, James Trevelyan, Patrick Winston, and Alan Yuille for their comments. Eric Grimson, as ever, carefully read a draft of the paper, and made many comments that significantly improved it.

References

- Agin, G. J, [1980], "Computer vision systems for industrial inspection and assembly," *Computer*, **13**, 11-20.
- Asada, Haruo, and Brady, Michael, [February, 1984], The curvature primal sketch, MIT Artificial Intelligence Laboratory, AIM 758.
- Attneave, Fred, [1954], "Some informational aspects of visual perception," *Psych. Review*, **61**, 183 - 193.
- Ballard, D. H., and Brown, C. M, [1982], *Computer Vision*, Prentice-Hall, New Jersey.
- Blum, Harry, and Nagel, Roger N, [1978], "Shape description using weighted symmetric axis features," *Pattern Recognition*, **10**, 167-180.
- Bolles, R. C., and Cain, R. A, [1982a], "Recognizing and locating partially visible objects: the local focus feature method," *The International Journal of Robotics Research*, **1**, (3), 57 - 82.
- Bookstein, Fred L, [1979], "The line-skeleton," *Comp. Graph. and Image Proc.*, **11**, 123-137.
- Brady, Michael, [1982a], "Computational approaches to image understanding," *Computing Surveys*, **14**, 3-71.
- Brady, Michael, [1982b], Parts description and acquisition using vision, *Robot Vision*, Rosenfeld A. (ed.), SPIE, 20 - 28.
- Brady, Michael, [1982c], "Smoothed local symmetries and local frame propagation," *Proc. Patt. Rec. and Im. Proc., Las Vegas, June*, , 629-633.
- Brady, Michael, [1983], Criteria for representations of shape, *Human and Machine Vision*, Rosenfeld, A., and Beck, J., (eds.), Academic Press, New York.
- Brady, Michael, [1984a], Representing shape, Proc. IEEE Int. Conf. Robotics, Atlanta, Ga., 256 -265.
- Brady, Michael, [1984b], "Artificial Intelligence and Robotics," *Artificial Intelligence*, (to appear).

- Brady, Michael, and Asada, Haruo, [1984], Smoothed local symmetries and their implementation, *Proceedings of the First International Symposium on Robotics Research*, Michael Brady and Richard Paul (eds.), MIT Press, Cambridge.
- Brady, Michael, and Yuille, Alan, [June 1984], Representing three-dimensional shape, Proc. IFToMM Conf., Udine, Italy.
- Brooks, R. A., and Binford, T. O., [1980], Representing and reasoning about partially specified scenes, *Proc. DARPA Image Und. Workshop*, Baumann, Lee S. (ed.), College Park, Md., 95 — 103.
- Brooks, R.A, [1981], "Symbolic Reasoning Among 3-D Models and 2D Images," *Artificial Intelligence*, 17, 285-348.
- Canny, J. F, [1983], Finding lines and edges in images, MIT, Artificial Intelligence Laboratory, Cambridge, Mass., AI-TR-720 .
- Heide, S, [1984], A hierarchical representation of shape, MIT, (in preparation).
- Hillis, W. Daniel, [1981], The Connection Machine (Computer Architecture for the New Wave), MIT Artificial Intelligence Laboratory, AIM-646.
- Holland, Steven W., Rossol, L., and Ward, Mitchell R, [1979], CONSIGHT 1: a vision controlled robot system for transferring parts from belt conveyors, *Computer Vision and Sensor Based Robots* eds. Dodd, G. and Rossol, L. Plenum Press.
- Hollerbach, J. M, [1975], Hierarchical shape description of objects by selection and modification of prototypes, MIT, also AI-TR-346.
- Marr, D, [1974], The low level symbolic representation of intensity changes in an image, MIT Artificial Intelligence AIM-325.
- Marr, D, [1976], "Early processing of visual information," *Phil. Trans. Roy. Soc. Lond. B*, 275, 483 -524.
- Marr, D, [1977], "Analysis of occluding contour," *Proc. Roy. Soc. Lond. B*, 197, 441 - 475.
- Marr, D, [1982], *Vision*, Freeman, San Francisco.
- Marr, D., and Hildreth, E, [1980], "Theory of edge detection," *Proc. Roy. Soc. Lond. B*, 207, 187 - 217.
- Marr, D., and Nishihara, H. K, [1978], "Representation and recognition of the spatial organisation of three-dimensional structure," *Proc. Roy. Soc. B*, 200, 269 — 294.
- Nevatia, R., and Binford, T.O, [1977], "Description and Recognition of Curved Objects," *Artificial Intelligence*, 8, 77-98.
- Nishihara, H. K., and Larson, N. G, [1981], Towards a real time implementation of the Marr-Poggio stereo matcher, *Techniques and applications of image understanding*, Pearson, J. J. (ed.), SPIE 281.
- Pavlidis, T, [1977], *Structural Pattern Recognition*, Springer, New York.

- Perkins, W. A, [1978], "A model-based vision system for industrial parts," *IEEE Trans. Computers*, C-27 (2), 126 - 142.
- Rosenfeld, A, [1982], *Robot Vision*, SPIE vol. 336.
- Rosenfeld, A., and Kak, A. C., [1982], *Digital Picture Processing*, Academic Press, New York.
- Strang, Gilbert, [1980], *Linear Algebra and its Applications*, Academic Press (2nd ed.), New York.
- Watt, R. J., and Andrews, D. P, [1982], "Contour curvature analysis: hyperacuties in the discrimination of detailed shape," *Vision Research*, 22, 449 - 460.
- Winston, Patrick H, [1984], *Artificial Intelligence (second edition)*, Addison Wesley.
- Witkin, A., [1983], "Scale-Space Filtering," *Proc. 7th Int. Jt. Conf. Artif. Intell., Karlsruhe*, 1019 - 1021.
- Yuille, A.L. and Poggio, T, [1983], "Fingerprints Theorems for Zero-Crossings," *MIT Artificial Intelligence Laboratory AIM-790*.