Hierarchical Figure-Based Shape Description For Medical Imaging

November 1986

Technical Report 86-026

Stephen M. Pizer, William R. Oliver,
John M. Gauch, Sandra H. Bloomberg

The University of North Carolina at Chapel Hill
Department of Computer Science
New West Hall 035 A
Chapel Hill, N.C. 27514

To appear in Proc. NATO ASI Mathematics and Computer Science in Medical Imaging

HIERARCHICAL FIGURE-BASED SHAPE DESCRIPTION FOR MEDICAL IMAGING

Stephen M. Pizer * +, William R. Oliver *†, John M. Gauch *, Sandra H. Bloomberg *

Departments of Computer Science *, Pathology †, and Radiology +
University of North Carolina
Chapel Hill, NC, USA

ABSTRACT

Medical imaging has long needed a good method of shape description, both to quantitate shape and as a step toward object recognition. Despite this need none of the shape description methods to date have been sufficiently general, natural, and noise-insensitive to be useful. We have developed a method that is automatic and appears to have great hope in describing the shape of biological objects in both 2D and 3D.

The method produces a shape description in the form of a hierarchy by scale of simple symmetric axis segments. An axis segment that is a child of another has smaller scale and is seen as a branch of of its parent. The scale value and parent-child relationship are induced by following the symmetric axis under successive reduction of resolution. The result is a figure- rather than boundary-oriented shape description that has natural segments and is insensitive to noise in the object description.

We extend this method to the description of grey-scale images. Thus, model-directed pattern recognition will not require pre-segmentation followed by shape matching but rather will allow shape properties to be included in the segmentation itself.

The approach on which this method is based is generally applicable to producing hierarchies by scale. It involves following a relevant feature to annihilation as resolution is reduced, defining the component that is annihilating as a basic subobject, and letting the component into which annihilation takes place become its parent in the hierarchy.

1. OBJECT DEFINITION VIA HIERARCHICAL SHAPE DESCRIPTION

A common task in medical image processing and display is the definition of the pixels or voxels making up a particular anatomic object. With such a definition the object can then be displayed or analyzed, full scene analysis can begin, or parameters of image processing on the object or its image region can be chosen. The weaknesses of common methods of object definition are well known. Image noise has major effects on the object definition. Separate objects are inadvertently connected when they have similar properties. And the global variation of such features as image intensity across an object undermines definitions based on these features. These weaknesses follow from the locality of decisions on boundary or region specification, and the inability of the methods to take into account global expectations about the objects being defined. Thus, detail due to noise or normal variations interferes with the determination of an object's global properties, and global information cannot be brought to bear until after a tentative segmentation (definition of an object or subobject) has occurred.

To avoid these weaknesses, we are developing methods that

1. model expected objects hierarchically by scale so that detail is seen as a property of a subobject that does not destroy the description of the object at a larger scale,

2. use descriptors that capture global object properties that commonly go under the name *shape*, but at the same time can capture typical intensity properties such as level and profile, and

3. operate in a way that does not require finding the object before it can be described.

In our approach, shape is described using the symmetric axis transform (SAT), a descriptor that is global by depending not on the object boundary but on the figure (included pixels or voxels) and that together with a multiresolution approach induces a hierarchical subdivision of an object into meaningful objects and subobjects. Furthermore, we suggest a form of this approach that allows ''intensity shape'' to be captured as well. Thus, an object is defined by computing a description of an image or image region and matching that description against a predefined description of the object, possibly together with its environment.

We begin by facing the problem of creating an adequate shape description of an object whose boundary has already been defined and then go on to see how the method can be extended to describe images or image regions defined only by the intensity values of their pixels or voxels. In all cases the method is discussed in two dimensions but applies directly to three dimensions, not slice by slice, but by replacing two dimensional elements (pixels) by three dimensional elements (voxels) and two dimensional distance by three dimensional distance.

After reviewing other work on shape description in Section 2, we review the symmetric axis transform and its properties in Section 3. Then in Section 4 a general approach for generating hierarchies by scale using multiple resolutions is presented, and it is applied to the symmetric axis to produce a hierarchical description of shape. Section 5 covers details of the multiresolution symmetric axis transform approach such as the means of reducing resolution and of following the axis as resolution is reduced. In Section 6 the relation between the proposed shape description and other methods is discussed, and extensions are suggested. Finally, in Section 7 a method applicable to grey-scale images is presented.

2. SHAPE DESCRIPTIONS

Many methods of shape description have been previously proposed, among them many focusing on the description of boundary curvature [Koenderink, 1985; Zahn and Roskies, 1972; Richards and Hoffman, 1985], some based on a list of somewhat ad hoc features, some focusing on description of deformation from a primordial shape such as an oval [Bookstein, et al., 1985; Leyton, 1984; Leyton, 1986b], some focusing on description of the object figure, i.e., the area (or volume in 3D) of the object [Blum and Nagel, 1978; Nackman and Pizer, 1985], and some focusing on both boundary and figure [Brady and Asada, 1984]. All except those based on the feature list have difficulty with the effects of noise in the figure or boundary definition and with separating detail from more essential shape characteristics, and some, e.g., the symmetric axis transform, have especial difficulties in this regard. Many produce descriptions of questionable naturalness.

To handle detail and noise naturally, one is led to a representation of the object as a hierarchy of segments at successively smaller levels of scale. For example, a human face might be described in terms of regions and subregions as in Figure 1. An advantage of this approach is that detail that may be noise is relegated to the lower parts of the hierarchy, and if it is to be seen as noise rather than important detail, it can be ignored without disturbing the description at higher levels of scale. Another advantage is that it allows top-down (large scale subobjects first) matching of models and descriptions of data objects. This approach of producing and using a hierarchical description has been taken with attractive results in regard to grey-scale image description by multiresolution methods which focus on intensity extrema [Crowley and Sanderson, 1984; Pizer, et al., 1986].

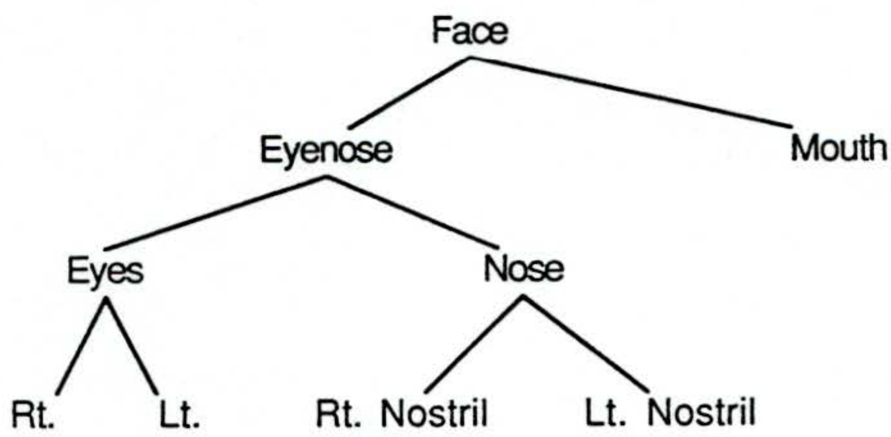

Figure 1. A Hierarchical Description of the Human Face

The symmetric (or medial) axis transform (*SAT*) [7] has elegant properties in inducing segmentation of shapes into natural components, but its major flaw has been its sensitivity to noise in the boundary or figure specification. A related problem has been that no measure of the closeness of shapes fell out from the SAT, because there was no way to discern how to group parts of the axis into major components or how to measure the importance of a component. In this paper we show how the multiresolution and symmetric axis transform approaches can be married, producing a noise-insensitive hierarchical shape description with the natural segments that had been the original promise of the symmetric axis transform.

3. THE SYMMETRIC AXIS TRANSFORM

The symmetric or medial axis (*SA*) of a 2D object is intuitively the set of points within the object figure that are medial between the boundaries. More precisely, the SA is the locus of the centers of all maximal disks in the object, where a maximal disk is a disk entirely contained within the object figure but which is not contained by any other such disk. Figure 2 shows an example.

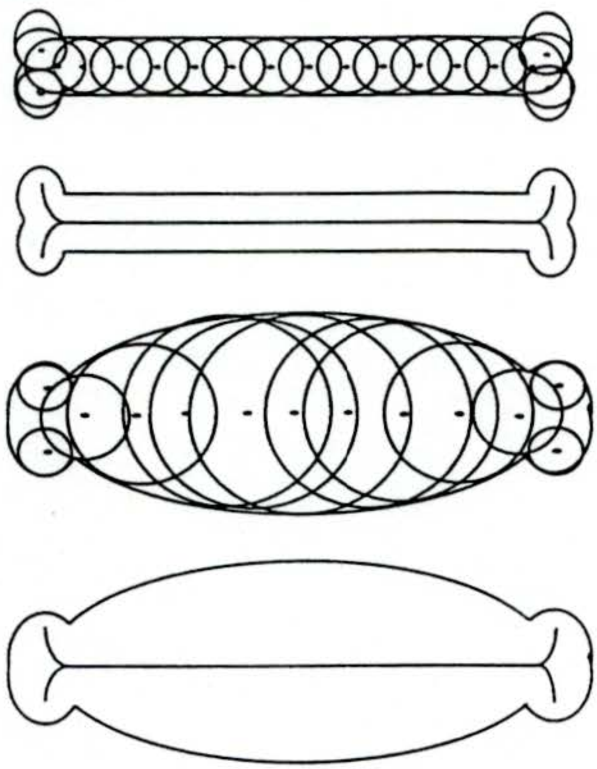

Figure 2. Maximal Disks in Simple Objects and the Corresponding Symmetric Axis

The SA forms a graph (a tree if the object has no holes). Segmenting the SA at branch points produces so-called simplified segments. Associated with each simplified segment is a part of the object figure made up of the union of the maximal disks of the points on the simplified segment. The object segmentation thus induced is frequently very natural. For example, in Figure 2 the bone shape is segmented into a rod and four knobs.

Associated with each point on a simplified segment is the radius of the maximal disk at that point. The SA together with this radius function of position on the axis is called the *symmetric axis transform (SAT)*. The radius function can be analyzed in terms of curvature properties [Blum and Nagel, 1978] which characterize the behavior of the width of the object at that point, e.g., as flaring or cupping. One of the attractive properties of the SAT is that it separates these width properties from the curvature properties of the axis. These two sets of curvature properties can be used to further segment the SA and thus the object.

Nackman [1985] has shown how the ideas of the SAT and the associated width and axis curvatures generalize to three dimensions. The ''axis'' becomes a locus of the centers of maximal balls, which in general is a branching surface. This surface can be subdivided into simplified segments at the branch curves, and again a natural subdivision is frequently produced.

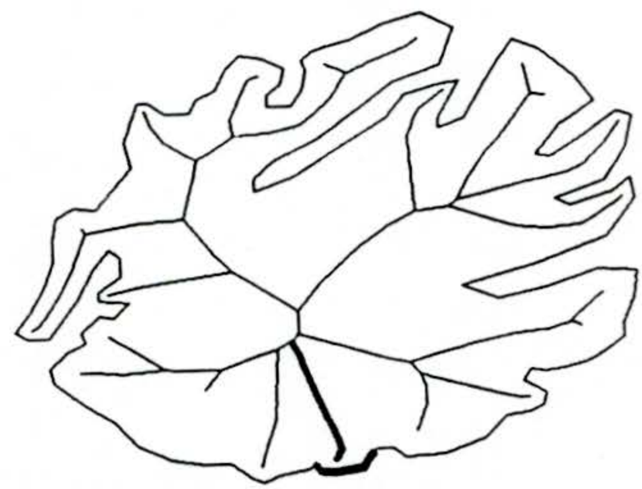

Figure 3. Sensitivity of the SA to Figure Noise in an image of a glomerulus. Note, for example, the large ratio of boundary to axis arc length in the region marked in bold.

The major weakness of the SAT is its sensitivity to properties of the detail of the object boundary. That is, changes in the figure or its boundary that are small in terms of distance can produce major changes in the SA (see Figure 3). A boundary feature that has a short arc length may result in a long symmetric axis branch, which moreover distorts the branch to which it is connected. The result is not only that branches that describe only detail are difficult to discern as such but also that major branches are split in such a way that a portion of axis that should naturally be viewed as a unit (a *limb* of the SA tree) is broken into unassociable portions. This weakness of the SAT has been so great as to destroy interest in it despite its otherwise elegant properties. Attempts to ameliorate it by pre-smoothing the boundary or by analysis, after SAT calculation, of properties such as axis arc length to boundary arc length ratios have foundered on the arbitrary thresholds that had to be imposed (''One man's noise is another man's detail.''). However, we suggest that the imposition of a scale-based hierarchy on the symmetric axis segments solves these problems and thus allows one to take full advantage of the attractive properties of the SAT of inducing segmentation strongly related to our sense of shape and of separating width curvature from axis curvature.

4. THE SYMMETRIC AXIS BRANCH HIERARCHY

A useful paradigm for creating scale-based hierarchies for describing a complex distribution of components, such as an image or image object, is to find an important component feature that smoothly changes as the underlying distribution is blurred and that annihilates after an appropriate amount of blurring, then becoming part of another component. This approach has been fruitful with grey-level images, where, in a generalization of the pyramid approach, Koenderink [1984] and Pizer [1986] have suggested following intensity extrema under blurring until they annihilate. There, the amount of blurring necessary for a particular extremum to annihilate is taken as the scale of the extremum; in the process a region surrounding the extremum is associated with the extremum, producing its *extremal region*; in addition, when upon annihilation an extremum melts into another, the former is associated with the latter as its child in the hierarchy. The result is that the image is described by a tree of extremal regions, each labeled with a scale, where larger scale regions have tree descendants that are smaller scale regions contained by it.

We can apply this paradigm to the problem of describing object shape by focusing on the branches of the symmetric axis. We have found empirically that reasonable methods of smoothing of the object boundary cause the branches of the symmetric axis to change smoothly, such that at certain levels of smoothing a branch will disappear (see Figure 4). According to the general approach laid out above, we associate with the branch, as a measure of its scale, the amount of resolution reduction necessary to achieve annihilation, and we say that the annihilating branch is a subobject of the branch into which it disappears. This process is continued until only a branch-free SA remains.

Every annihilating branch can be traced back to the part of the SA of the original, unblurred object from which it was smoothly generated. In the case of all but the branches at the frontier of the tree, these SA parts consist of a limb of the original SA, i.e., a sequence of simplified segments from which twigs have been removed at the one or more branch points where smaller scale branches that annihilated earlier were attached. The result is that the method has the property of defining naturally associated axis components from the associated segments in the sequence.

The complete multiresolution process defines a tree (hierarchy) of limbs and twigs (axis portions) in the original SA. The root is the portion of the axis to which the final branch-free SA traces back, and descendants of the axis portion at any node are the portions of axis which annihilated into that axis portion. The axis portion at each node is either a single SA simplified segment or a limb made up of a sequence (without branching) of simplified segments.

With every axis portion in the SA tree there is the original radius function on that axis portion. The union of the maximal disks centered at each point on the axis portion and with radii given by the respective radius function value is a subobject associated with the axis portion. The description tree can then be thought of as a tree of subobjects, of decreas-

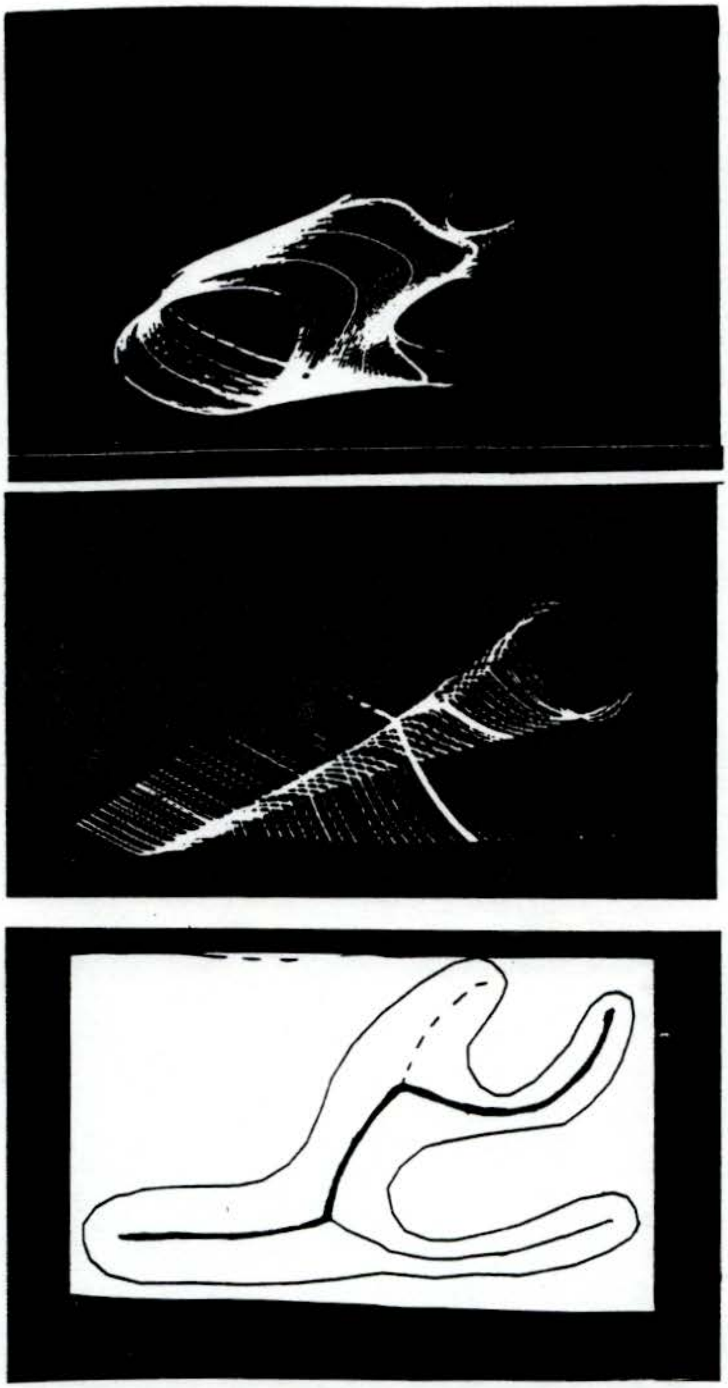

Figure 4. The SA Hierarchy based on B-spline Boundary Smoothing a) stack of boundaries, b) stack of SA's, c) induced description.

ing scale (but not necessarily area or volume) as you move down the tree. Each node in the tree (subobject) can be labeled by its scale together with properties describing the width (radius function) curvature and the axis curvature.

These ideas generalize straightforwardly to three dimensions. The axis a.

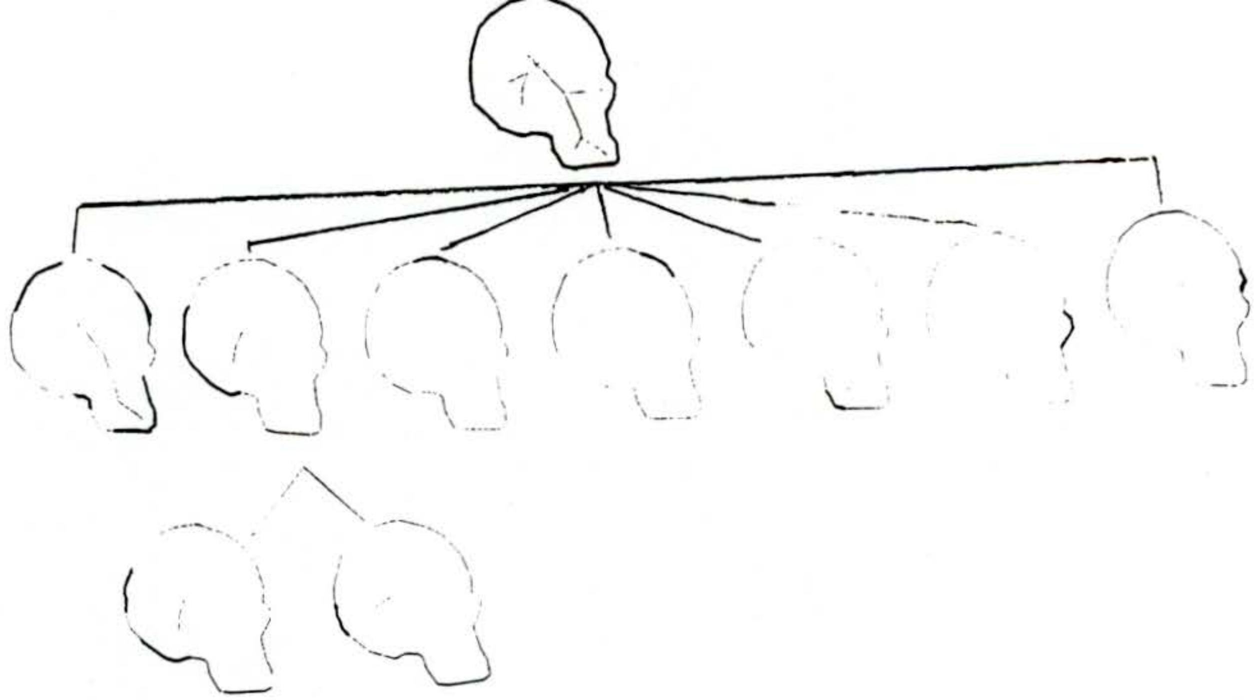

b.

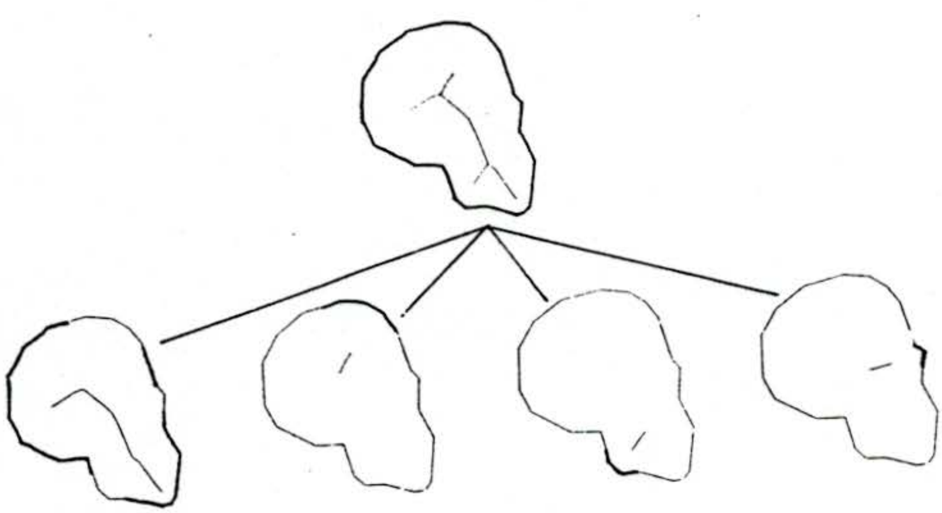

Figure 5. Shape Descriptions Produced by the SA Hierarchy. The descriptions above are lateral views of two different human skulls. Each node of a description tree shows a component of the SA and shows in bold the part of the boundary corresponding to that component. The leftmost child of a node is its principal component axis. The other children of that node represent axis complexes branching from the principal component, arranged in order of decreasing scale.

components at nodes in the tree are simple surfaces, and the subobjects associated with a node are corresponding unions of maximal balls. Axis curvature and width curvature properties, as described by Nackman [1985], as well as scale, label each node.

Examples given in Figure 5 suggest that this method produces natural descriptions. Furthermore, our experience is that objects that we see as similar, such as outlined human skulls viewed laterally, produce similar descriptions. Problems with this description arise from four facts. First, with some types of resolution reduction the topology of the figure and of the axis is not maintained: for example, a simplified segment of SA can split into two, or two can join into one. Second, shape features related to boundary concavities are not directly represented by this approach. In Section 4, where types of resolution reduction are discussed, we show how including other axes of symmetry in the representation seems to handle both of these problems.

The third problem is that the sensitivity of the symmetric axis to small changes in the boundary can cause implementation difficulties in axis segment following. In particular, axis segments that do not exist at one resolution level can be artifactually created, or those that decrease smoothly in size for an underlying smooth boundary are artifactually removed at some resolution levels, only to reappear at later levels. Therefore, following the axis segments across steps is made difficult. This problem is discussed in Section 5.

Finally, there is one case in which similar objects have dissimilar descriptions. When the two branches emanating from a branch point are similar in scale in that when one annihilates the other is also almost gone, a small change in the scale (length or width) of one of the branches can change which of these two branches annihilates first and thus change which is considered part of the limb and which the attached twig. The tree changes that result from such small object changes are predictable and are discussed further in Section 6.

5. RESOLUTION REDUCTION AND SYMMETRIC AXIS FOLLOWING

What method should be used to continuously reduce the object resolution to produce the multiresolution stack of symmetric axes that induce the SAT hierarchy? The natural first thought is to focus on the boundary of the object to be described by applying some smoothing operator to its curvature. The result of such an approach, in which boundary points at one level of resolution were used as control points for a B-spline which forms the boundary at the next level, is shown in Figure 4. Koenderink [1986] discusses why it is preferable to focus on first blurring the figure and then computing a consequently smoother boundary from the result, rather than directly to smooth the boundary. In essence the argument is that figure properties better capture the global relationships which we call shape than do boundary properties, which are too local. This very argument is the basis of the appeal of the symmetric axis method of shape description over methods based on describing boundary curvature.

Koenderink's suggestion [1986] of the means for figure-oriented resolution reduction starts by treating the figure as a characteristic function, i.e., an image which is 0 outside the figure and 1 inside. He then would convolve the result with an appropriate Gaussian and compute a new figure boundary as a level curve in the result. Koenderink suggests that the level curve be taken at some fixed intensity, but this requires the choice of some arbitrary intensity, and the choice affects the shape description. Moreover, the approach of choosing a fixed intensity level causes the figure to shrink as resolution is reduced, so that after some amount of resolution reduction the figure disappears entirely. While there are many means of level choice that avoid this disappearance, we take our cue from the accepted definition of shape that it is what is left after normalization for size and orientation. Thus, we normalize to constant size at each amount of resolution reduction, choosing the figure-defining level such that the area (in 2D, volume in 3D) of the figure remains constant.

Resolution reduction based on figure blurring does indeed behave more intuitively correctly than direct boundary smoothing. Furthermore, it does have the additional advantage that it is in principle directly applicable to grey--scale object representations and not just characteristic function representations. However, there are two difficulties with this type of resolution reduction. First, topology is not maintained: connected components can split under blurring, disconnected components can join, indentations in the figure can turn into holes and vice-versa, and holes can disappear (see Figure 6). Second, with an implementation using a piecewise linear boundary the smoothness of the SA branch disappearance is more affected than with direct boundary smoothing. Let us discuss each of these difficulties in turn.

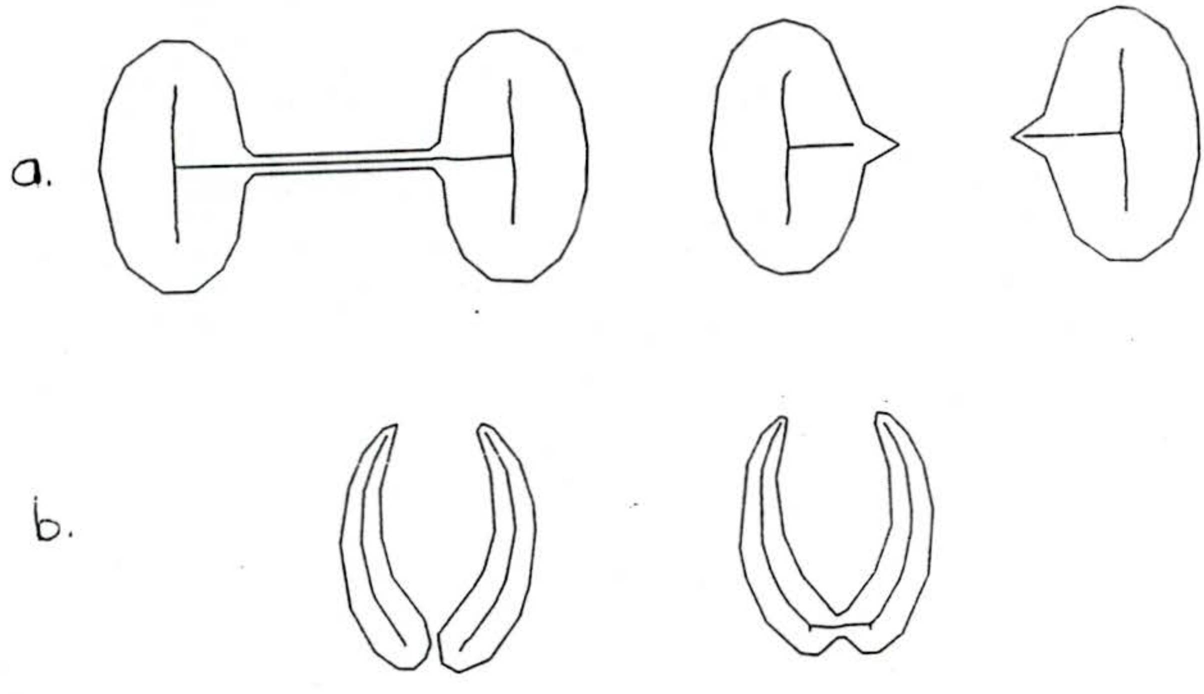

Figure 6. a) Splitting, and b) Joining of the SA under Resolution Reduction

From the point of view that the figure is essential and the boundary is derivative, the non-maintenance of topology is no problem at all. It is easy to argue that the object figures shown in Figure 6a are indeed close and that it is natural that a decrease in resolution should cause the isthmus between the two disks to be broken (eventually to be rejoined at yet lower resolution). Similarly, the object figures in Figure 6b should naturally combine as the resolution is lowered; other natural transitions are the closing of two nearby points of land around a bay to form a lake, the melting of a narrow strip between a lake and the sea to form a bay, and the drying up of a lake. We see it as unnatural to insist that topology be maintained under resolution reduction; instead, we must arrange our shape description not to be too sensitive to topology.

On the other hand, we are more disturbed by nonsmooth appearances or disappearances of large pieces of the SA. These appear to happen under some of the changes in topology listed above. For example, when two ellipses, each with a horizontal major axis and one just above the other, are blurred, a vertical segment of SA will appear nonsmoothly as the two ellipses join (see Figure 7a). We can avert many of these difficulties by including the external symmetric axis or the global SA as part of the SA.

The external SA of a figure is the SA of its complement. If we take the overall SA as the union of the internal and external SA (see Figure 7b), we find first that with resolution reduction as a piece of internal SA breaks, a corresponding pair of external axis pieces come together, and second that the overall SA now reflects concavities in the figure boundary directly. We therefore suggest that an improved shape description can be obtained by following this overall SA under resolution reduction. Note that as the resolution is reduced the object eventually becomes ovoidal and the external part of the SA becomes null.

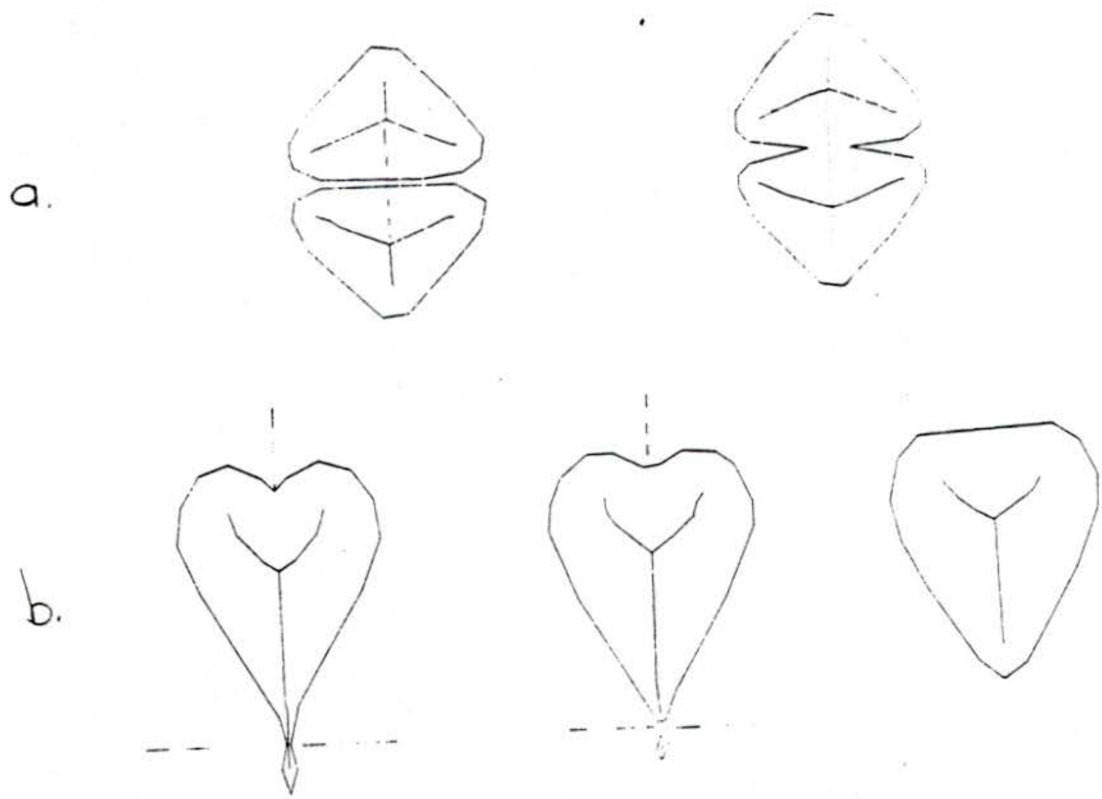

Figure 7. The Change Under Resolution Reduction of a) The Ordinary (——) and One Other (- - -) Component of the Global Symmetric Axis of Two Nearby Joining Figures, and b) The Internal (——) and External (- - -) SA of Two Splitting Figures

The global SA [Blum, 1979] is formed by the locus of the centers of all disks tangent to two or more disconnected regions on the figure boundary. The ordinary (first order) SA and the external SA are subsets of the global SA. We have observed that when a new segment of ordinary SA appears nonsmoothly as two pieces of figure join or a hole is eliminated, it actually forms smoothly from a piece of global SA that is transformed into ordinary SA. On the other hand, when a hole fills in, the ordinary SA segment loop around the hole does disappear nonsmoothly. Further research is needed to catalogue these transitions and to determine the usefulness of computing the global SA.

The other disadvantage of resolution reduction based on figure blurring, as compared to direct boundary smoothing, is that following SA branches to annihilation is more difficult to implement using a piecewise linear boundary approximation. In our implementation the histogram of the image at each level of resolution is used to find the intensity such that the number of pixels with greater or equal intensity is equal to the original figure area (or volume). Points on the isointensity contour at that intensity are connected by linear segments, with the points selected by the recursive splitting method described in Ballard [1982, Algorithm 8.1]. Our SAT algorithm essentially computes the Voronoi diagram of this collection of line segments and takes the SA to be the curves separating the regions in the Voronoi diagram, less those separating curve segments that touch the object boundary. The latter are removed because they are an artifact caused by using a piecewise linear approximation to a smooth object boundary. The result is an SA that is piecewise made up of linear and parabolic components.

The computed axis depends on the boundary approximation. In particular, small pieces of SA are either artifactually added or omitted depending on the approximation, and the added pieces are not necessarily small compared to SA pieces reflecting the ''true'' boundary at the given level of resolution, especially since the ''true'' branches get small as resolution is reduced. The artifactual addition or omission of pieces that are unrelated across stages of resolution reduction complicates following the SA branches to annihilation.

This difficulty does not occur as greatly with direct boundary smoothing, as the boundary pieces at one step are related to those at the next. Nevertheless, the fundamental attractiveness of resolution reduction by figure blurring leads us to cope with its difficulty rather than resort to direct boundary smoothing. We believe that we need either to develop a method of figure-blurring-based resolution reduction in which the piecewise linear approximation at one step has related pieces to those in the previous step, or to replace the piecewise linear boundary approximation by a smooth piecewise approximation and develop a method of calculating the SAT for such approximations.

In our results to date, including those shown in Figure 5, the matching of SA segments from level to level has been done by hand, though the large majority of the matchings can be correctly made by a straightforward algorithm.

6. RELATIONS AND EXTENSIONS OF THE MULTIRESOLUTION SAT

Brady [1984] has defined another form of an axis of symmetry that he calls *Smoothed Local Symmetries (SLS)*. He combines boundary curvature and figural properties in his definition. We find the boundary curvature aspects unattractive, as they involve an arbitrary degree of boundary smoothing, but the SLS, defined as the set of smooth loci of centers of chords that have angular symmetry with the boundaries they touch, is a set of axes that have some advantages (and some disadvantages) over the SA. We note here only that the multiresolution approach to inducing a hierarchy on the SA could equally well be applied to the SLS (in our modified definition).

The hierarchical SA or SLS descriptions have a relation to shape descriptions based on boundary curvature, such as the codons of Richards [1985] or the boundary deformations of Leyton [1986b]. Codons describe simple convexities or concavities of the boundary. Leyton has shown [1986a] that these simple boundary segments have simple SA segments, i.e., correspond to the outermost branches of the overall SA. That work begins to relate the boundary and figure points of view, and more mathematics in this direction would be valuable.

Leyton's and Richards' methods, as well as others, describe the boundary as a sequence of curvature features or deformations to obtain them, but they frequently have difficulty determining an order of features or deformations to be applied. The multiresolution SAT method described above could be used to induce the order by scale for these methods, if the descriptions produced by latter are deemed superior to those produced by the multiresolution SAT.

Whether the multiresolution SAT is used directly to produce a shape description or as an order-inducing auxiliary to another method of producing a shape description, its weakness of having small changes in the image produce discontinuous change in the description must be dealt with. Recall that this behavior results from a close decision in deciding which of two branches or limbs emanating from a branch point forms the branch and which forms part of the branching limb. This behavior can happen at any branch point at which the two SA branches have almost equal scale. For example, in Figure 5a, the axis piece corresponding to the base of the skull is more prominent, while in 5b, the axis piece corresponding to the back of the skull is more prominent. It seems to be an application-dependent question whether trees differing in one or more of these close decisions should be considered instances of the same shape or of different shapes. In fact, it might be decided that the ''close call'' case is a single shape, independent of which way the call goes, while each of the cases, in either direction, of a ''distinct difference'' in scale between the respective branches might be called two yet different shapes. Of course, this decision would require an arbitrary threshold to be established between each of the shapes, and it also would require some means to be developed to define the scale of the non-annihilating branch that becomes part of the limb.

An alternative seems to be to let an annihilation cause both the annihilating branch and the ''other'' branch at the fork both to be declared as branches. The problem here is which of the (normally) three branches incident to a branch point to declare as the limb and which the two branches. Without a means to make such a declaration, the strength of the method in associating axis segments into a limb is lost. With such a means, one could detect when two branches are nearly equally strong and then produce alternative description trees corresponding to each of these branches annihilating first. Further research is needed on this question.

7. THE MULTIRESOLUTION SAT FOR GREY-SCALE IMAGES

Finally, we move on to the problem of creating a shape-based description of intensity-varying images or objects. With a method for creating such a description we can avoid the preliminary step of defining the pixels in the object before the shape is described, so the shape description can contribute to the object definition. Furthermore, the description can reflect the levels and ''shape'' of changes in the intensity dimension as well as spatial shape.

A common way to visualize a grey scale image is as the surface in three dimensions defined by the image intensity function. This has two benefits. First, the ''shape'' of the grey-scale image can be described by a three-dimensional shape description method. And second, the familiar tools of differential geometry can be used to study properties of grey scale images and the shape description.

Our first idea was to describe the image surface using the 3D SAT. This approach has the problem that the intensity dimension is being treated as commensurate with the spatial dimensions; some choice must be made as to what intensity change is equivalent to what spatial distance. Unfortunately, two different choices may result in very different shape descriptions. Thus, we must design a shape description which treats the intensity and spatial dimensions separately.

To meet this requirement, we look at the image intensity function as a collection of isointensity contours. Since the level curves for the intensity function provide a complete representation of the image, the collection of SAT's of all of these level curves will also provide a complete representation. We describe each isointensity contour with the 2D SA of the region with intensity greater than or equal to that of the contour. Thus the whole image is described by a collection of 2D SA's obtained by varying the intensity defining the contour (see Figure 8). We then continuously blur the image and follow the annihilation of major components of this pile of SA's to induce a hierarchical description of the image based on a natural subdivision. Let us examine some of the properties of our description.

First, let us examine the connection of SA's from one intensity to the next, at a fixed degree of blurring. Consider functional surfaces which have a finite number of discrete critical points. Any selected intensity level i determines a collection of isointensity contours at that intensity. Except at levels at which critical points occur, the set of contours varies smoothly with i. Since an SA varies smoothly with the region it represents, the SA's of the regions corresponding to the two sets of contours are also very similar. For this reason, the set of SA's for all intensities in the (2D) image will form a branching surface in three dimensions. We call this structure the *SA-pile* for a grey scale image, and we call its

Figure 8. A Digital Subtraction Angiogram and Corresponding Contour Pile and SA-Pile

branches *SA-sheets*. For a 3D image the SA-pile appears in four-dimensions and its sheets form a three-dimensional hyper-surface in this space.

To visualize the behaviors of SA-piles, view the image as terrain and the regions whose SA's we are computing as horizontal cross-sections through this solid mass. Thus near a maximum a cross-section is a closed region, while near a minimum a cross-section is a region with a hole in it.

For an image with only one smooth oval bright spot (maximum), the SA-pile consists of only one SA-sheet. For images with bright spots of a more complex shape, there may be a branching structure of SA-sheets (very much like the branching structure in the 2D SA for non-oval 2D objects). What about more general images? We know that the critical points of the surface cause catastrophic changes in the isointensity contours. Contours appear and disappear at local maximum points and local minimum points, and contours connect at saddle points. The SA-sheets also change drastically at these points.

Local extrema are the first two types of critical points. Consider the terrain near a local maximum. As illustrated in Figure 9a, the SA sheet at this position ends at the intensity of the hilltop. Now consider the terrain near a local minimum. As illustrated in Figure 9b, the SA forms a loop around the hole formed by the minimum, which instantaneously transforms into an SA without a loop as we move below the intensity of the minimum. That is, the SA-pile consists of a cylindrical sheet which instantaneously turns into a simpler sheet at the local minimum. Following our discussion in section 5, this abrupt appearance of the simpler sheet can be avoided if the global SA rather than just the first-order SA is used.

The third type of critical point, the saddle point, comes in three generic varieties: passes between two hilltops and between two pit-bottoms (see Figure 9c) a pass between a hilltop and a hole in the hillside (see Figure 9d), and a pass between a pit-bottom and a peak on the pitside (see figure 9e). These figures illustrate that as we move from below near these critical points, the components of the SA-piles behave as follows. In the first and third cases an SA-sheet tears at the saddle point. In the second case a tear is formed in an SA-cylinder, turning the cylinder into an

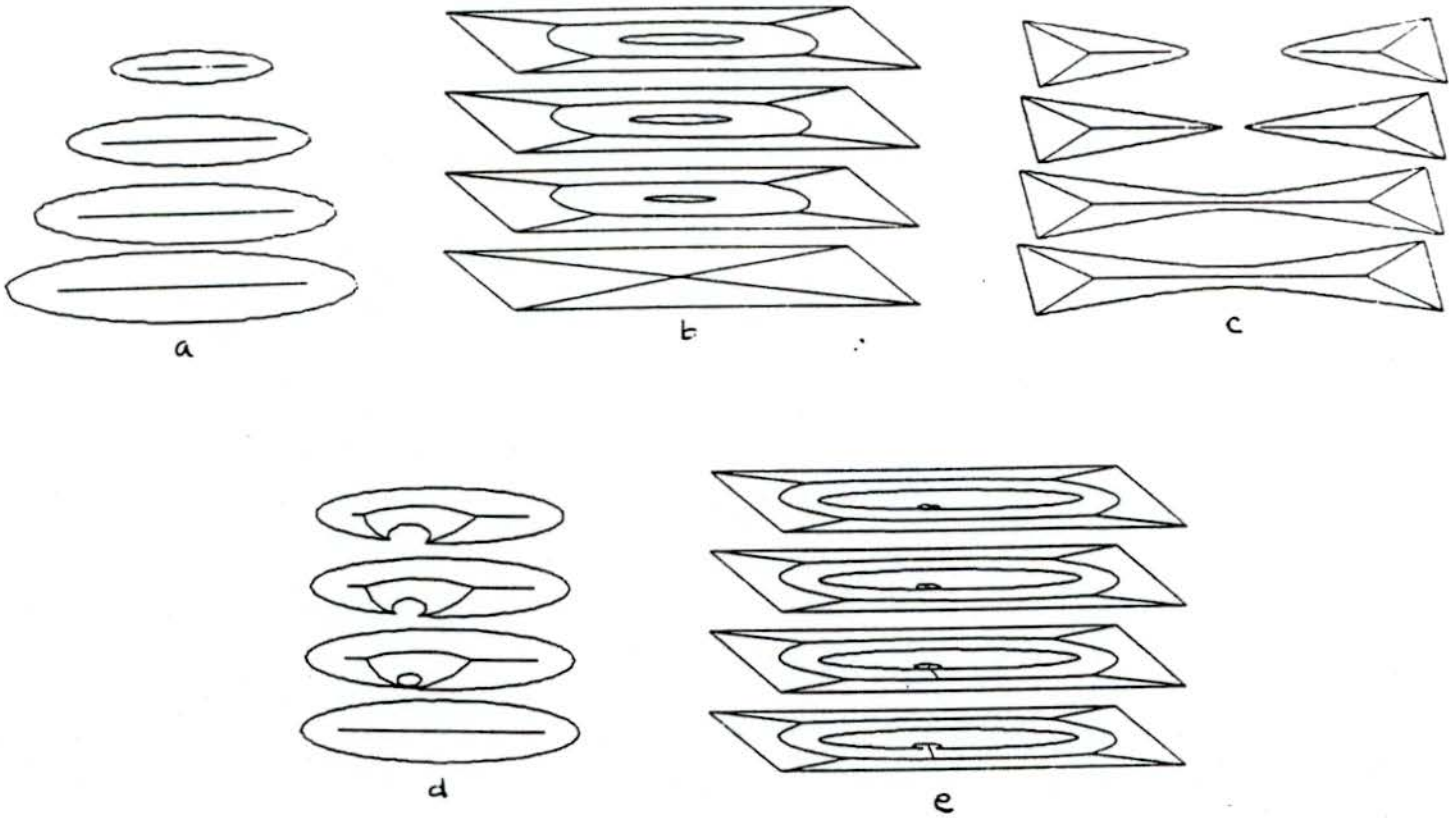

Figure 9. Components of SA-Piles near Intensity Critical Points: a) Maximum; b) Minimum; c) Saddle Point between Two Hills and between Two Pits; d) Saddle Point between Hill and Hole in Its Side; and e) Saddle Point between Pit and Peak on Its Side.

ordinary sheet. In summary, the SA-pile consists of a branching structure of possibly torn and SA-sheets and cylinders.

Now we must address the problem of imposing a natural hierarchy on the branching sheets in our description. We know that Gaussian blurring of an image annihilates extrema and saddle points in pairs [Koenderink, 1984]. Because the contours associated with the image will simplify with blurring, the corresponding SA -pile will also simplify. Our experience to date confirms the inference that the process of blurring will cause successive disappearance of SA-sheets and SA-cylinders. We then impose a hierarchy based on the order of disappearance of SA-sheets and SA-cylinders, in a way analagous to the multiresolution SA of defined objects described in Section 4.

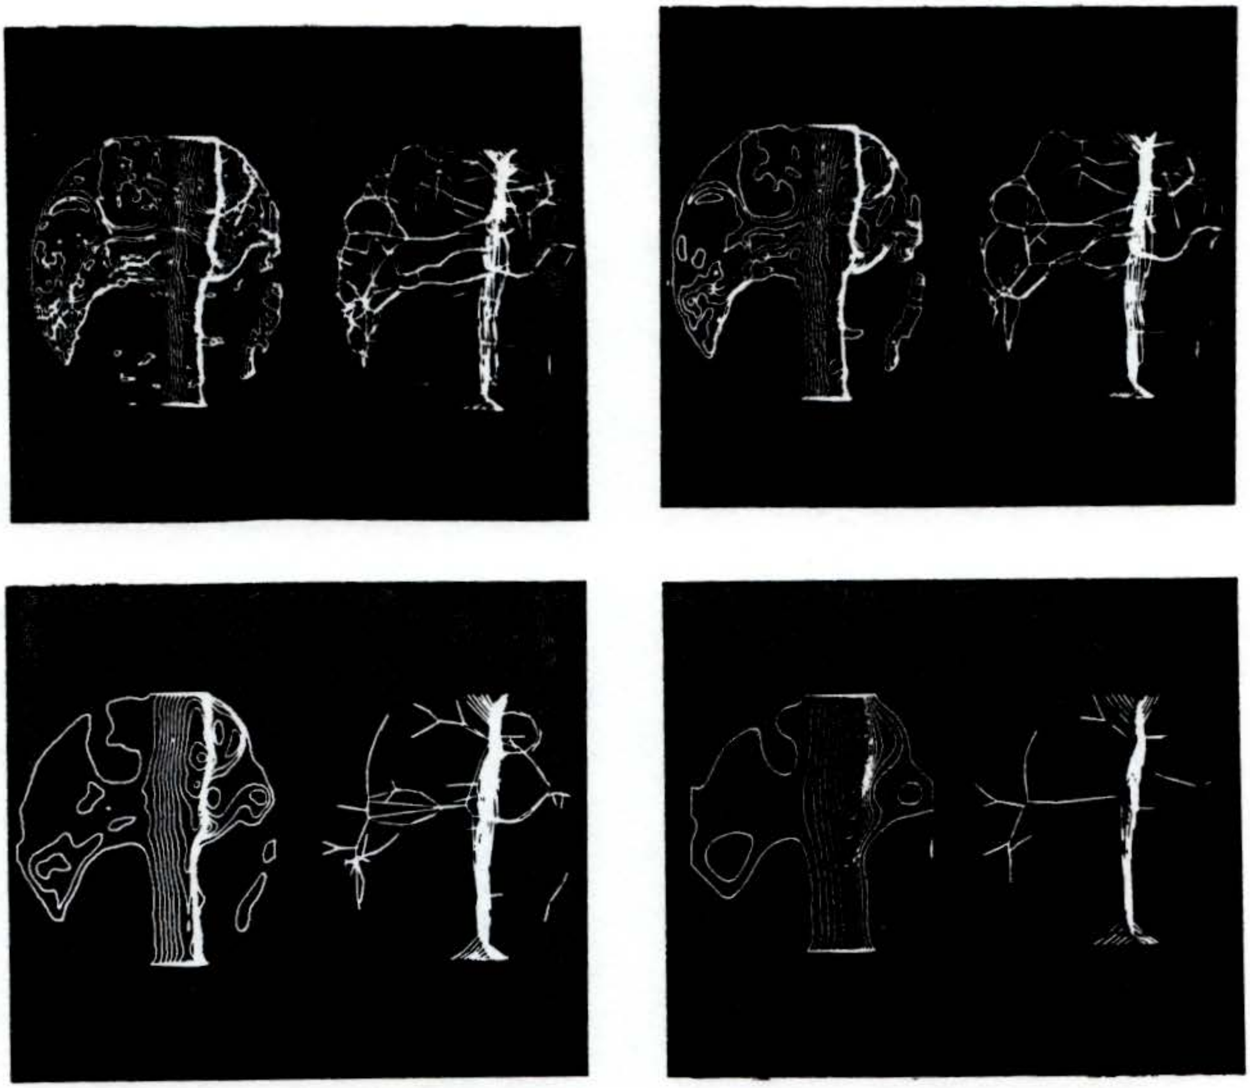

Figure 10. A Sequence of Blurred Digital Subtraction Angiograms, Their Contour Piles, and Their SA-Piles.

We have implemented this process using subroutines that convolve an image with a Gaussian, threshold the image at a selected intensity, and compute an SA by connecting centers of maximal disks in the thresholded region as one follows the boundary of the region. The results of successively larger threshold intensities are displayed as a pile of SA's on a vector graphics screen. Similarly, the successive isointensity contours can be displayed as a pile. Figure 10 shows the contour piles and associated SA-piles of a sequence of blurred versions of a digital subtraction angiogram. The predicted behavior can be seen.

Further mathematical investigation and algorithmic development of the multiresolution grey-scale SA-pile is required. The usefulness of the external and global SA deserves attention. In addition, the relationship of this method to natural vision might be enhanced if the original image was first transformed into ''perceived intensity" by reflecting the effect of edges and intensity diffusion [17].

8. SUMMARY AND CONCLUSIONS

In closing, we suggest that the multiresolution symmetric axis transform is a main challenger in the possibilities for shape description. However, more work is needed in a number of areas. Mathematical study is needed of how symmetric axes behave under figure- or boundary-based resolution reduction and how SA-piles behave under Gaussian blurring. Algorithms must be developed to connect SA's into an SA-pile and to follow branches of SA's and SA-piles through resolution reduction to annihilation. This will require methods of computing the SA that limit the creation of artifactual branches and the premature omission of real branches.

In addition, we hope to have focused attention onto the multiresolution approach of following features to annihilation and thus inducing scale-based hierarchies. This appears to be a powerful method of deriving descriptions that are suitable for model-directed pattern recognition.

ACKNOWLEDGEMENTS

We are indebted to Fred L. Bookstein, Jan J. Koenderink, and Lee R. Nackman for useful discussions. We are grateful to Sharon Laney for manuscript preparation and to Bo Strain and Karen Curran for photography. The research reported in this paper was carried out under the partial support of NIH grant number R01-CA39060 and NATO Collaborative Research Grant number RG. 85/0121.

REFERENCES

Ballard, D.H. and C. M. Brown *Computer Vision*, Prentice-Hall, Englewood Cliffs, N.J., 1982, p. 234.

Blum, H. ''3-D Symmetric Axis Co-ordinates: An Overview and Prospectus'', draft of talk at Conf. on Representation of 3D Objects, U. of Penn., 1979.

Blum, H. and R. N. Nagel ''Shape Description Using Weighted Symmetric Axis Features'', *Pattern Recognition* 10: 167-180, Great Britain: Pergamon Press Ltd., 1978.

Bookstein, F., B. Chernoff, R. Elder, J. Humphries, G. Smial and R. Strauss *Morphometrics in Evolutionary Biology*, Special Publication 15, The Academy of Natural Sciences of Philadelphia, 1985.

Brady, M. and H. Asada ''Smoothed Local Symmetries and Their Implementation'', Tech. Rept., Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1984.

Crowley, J.L. and A. C. Sanderson ''Multiple Resolution Representation and Probabilistic Matching of 2-D Gray-Scale Shape'', Tech. Rept. No. CMU-RI-TR-85-2, Carnegie-Mellon Univ., 1984. Also see Crowley, J. L. and A. C. Parker, ''A Representation for Shape Based on Peaks and Ridges in the Difference of Low-Pass Transform'', *IEEE Trans. PAMI*, 1984.

Grossberg, S. and E. Mingolla ''Neural Dynamics of Perceptual Grouping: Textures, Boundaries, and Emergent Segmentations'', *Perception and Psychophysics, 38* (2): 141-171, 1985.

Koenderink, J.J. ''The Structure of Images'', *Biological Cybernetics* 50: 363-370, 1984.

Koenderink, J.J. ''The Internal Representation of Solid Shape Based on the Topological Properties of the Apparent Contour'', *Image Understanding*, W. Richards, ed., MIT Press, 1985.

Koenderink, J.J. and A. J. van Doorn ''Dynamic Shape'', Tech. Rept., Department of Medical and Physiological Physics, Physics Laboratory, State University of Utrecht, The Netherlands, 1986.

Leyton, M. ''Perceptual Organization as Nested Control'', *Biological Cybernetics* 51: 141-153, 1984.

Leyton, M. ''A Theorem Relating Symmetry Structure and Curvature Extrema'', Tech. Rept., Department of Psychology and Social Relations, Harvard Univ., 1986a, submitted for publication.

Leyton, M. ''Smooth Processes on Shape'', Tech. Rept., Department of Psychology and Social Relations, Harvard Univ., 1986b.

Nackman, L.R. and S. M. Pizer ''Three-Dimensional Shape Description Using the Symmetric Transform, I: Theory'', *IEEE Trans. PAMI* 7(2): 187-202, March 1985.

Pizer, S. M., J. J. Koenderink, L. M. Lifshitz, L. Helmink, and A. D. J. Kaasjager *Inf. Proc. in Med. Imaging*, S. L. Bacharach, ed., Martinus Nijhoff, Dordrecht, The Netherlands, 1986, pp. 24-37.

Richards, W. and D. D. Hoffman, ''Codon Constraints on Closed 2D Shapes'' *Computer Vision, Graphics, and Image Processing* 31: 265-281, 1985.

Zahn, C.T. and R. Z. Roskies ''Fourier Descriptors for Plane Closed Curves'', *IEEE Trans. Comp. 21*(3): 269-281, 1972.