Multiresolution Shape Descriptions And Their Applications In Medical Imaging

June 1987

Technical Report 87-018

John M. Gauch, William R. Oliver and Stephen M. Pizer

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. Information Processing in Medical Imaging (IPMI), Netherlands, June 22-27, 1987

MULTIRESOLUTION SHAPE DESCRIPTIONS AND THEIR APPLICATIONS IN MEDICAL IMAGING

John M. Gauch*, William R. Oliver*†, Stephen M. Pizer*‡

Depts. of Computer Science*, Pathology†, and Radiology‡
University of North Carolina
Chapel Hill, NC, USA

1 INTRODUCTION

The recognition of objects in medical images and the analysis of their properties in space and time requires a representation that reflects properties of shape. These representations are independent of position, orientation and size. They can be based on ad hoc features, object boundaries, object interiors (figures), or object deformations.

Our work is based on exploiting an inherent relationship between boundary and figure based shape descriptions. In both of these categories, noise and small image details often confound shape analysis. Therefore, it is essential to generate descriptions that are hierarchical by scale using multiresolution techniques. In the following we develop techniques of this kind which are applicable to shape description of both grey-scale images and binary objects. In the sequel we describe two applications which are based on these shape descriptions.

2 GREY-SCALE IMAGE DESCRIPTION VIA LEVEL CURVES

Two dimensional grey-scale images can be viewed as a surface in three space defined by the graph (x, y, I(x,y)). One obvious way to describe the shape of such graphs is to use the tools of differential geometry to describe the surface. Another alternative is to describe the two regions of space separated by the surface.

Unfortunately, the intensity dimension is incommensurate with the spatial dimensions. There is no natural choice as to what intensity change is equivalent to what spatial distance. Shape descriptions which vary with a particular choice of equivalency must therefore be avoided. This can be accomplished by treating the intensity and spatial dimensions separately in our shape description. We do this by describing shape in terms of the level sets of the graph (see Figure 1).

The level sets for a two dimensional grey-scale image are the planar curves defined by I(x,y) = C, for all intensities in the image. These level sets act like boundaries. They partition each level into an *inside* and an *outside*. Specifically, the image surface is the union of its level sets at their respective intensity levels. The the volume *below* the surface is the union of all regions which have $I(x,y) \geq C$. Similarly, the volume *above* the image surface is the union of all regions where $I(x,y) \leq C$.

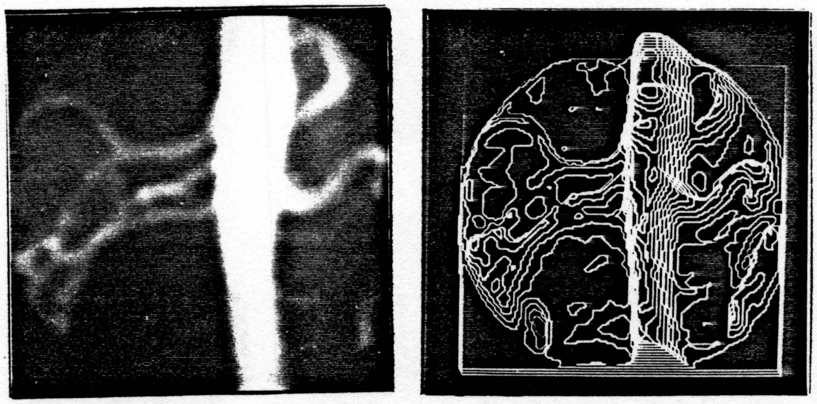

Figure 1. A digital subtraction angiogram and corresponding level sets.

Thus, it is possible to represent the image surface and the regions on either side of the surface in terms of the level curves or regions defined by level curves. To describe the shape of these level curves we use existing figure based and boundary based shape descriptions for binary images.

3 BASIC SHAPE DESCRIPTIONS

3.1 Axes of Symmetry

We begin our analysis with figure based shape descriptions for two dimensional binary images. Object symmetry is a key to understanding shape. Circles are perfectly symmetrical, so the axes of symmetry defined by circles capture this property best. When the set of circles tangent to the object's boundary is considered, we derive axes of symmetry which describe the branching and bending of the object. Several methods use this approach.

Chords of tangent circles are used to define *Smoothed Local Symmetries* (Brady and Asada, 1984). Arcs of tangent circles are used to define *Process Inferred Axes* (Leyton, 1986). Centers of tangent circles are used to define the *Symmetric Axis* (Blum, 1974). The *Internal Symmetric Axis* is defined by centers of tangent circles which are entirely within the object while the *External Symmetric Axis* is defined by centers of tangent circles which are outside the object. The *Global Symmetric Axis* is defined by the centers of all tangent circles. When the radius of each of these circles is also considered, we have the *Symmetric Axis Transform (SAT)* (Blum and Nagel, 1978).

The SAT has three attractive properties. First, the branching structure of the object is reflected by the branching of the axis. This yields a natural correspondence between components of the object and components of the shape description. Second, the bending and flaring of the object is reflected by changes in the curvature of the axis and the radius of the tangent circles. This gives us a way to compare and contrast similar shapes. Finally, this shape description is unique for an object and can be used to re-create the object.

One of the problems with the SAT is that it is very sensitive. Noise and small detail in the object can cause "unimportant" branches to appear in the axis. These confound shape analysis by introducing large numbers of axis segments and by breaking up main branches into numerous small sections. One solution is to use multiple resolution analysis to determine the *scale* of the individual components of the shape description. This approach yields the *Multiresolution Symmetric Axis* described by (Pizer et al., 1986).

To compute this shape description requires that we measure the importance of each branch in the symmetric axis. As we lower the resolution the tendency is for an object to simplify, eventually becoming an ellipse. Because the symmetric axis varies smoothly with the figure it represents, the branching structure of the axis also simplifies as we lower the resolution. Thus, we can follow axis branches to annihilation through a multiple resolution sequence of binary images. The importance of each branch is then determined by its annihilation resolution. This process also imposes a hierarchy on axis branches (see Figure 2).

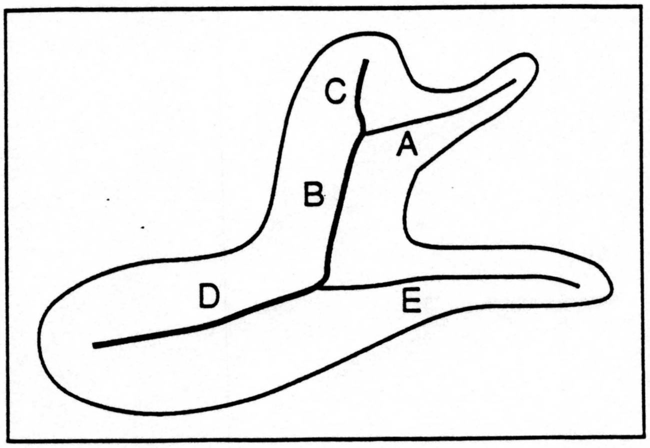

Figure 2. Branching hierarchy imposed on the symmetric axis by resolution reduction. When axis A annihilates it is labeled as a sub-branch of a new branch CB. Later, when axis E annihilates it is labeled as a sub-branch of the major axis CBD.

The order of annihilation of axis branches decomposes the SAT into limbs and twigs. When branch A annihilates two things happen. The adjacent branches C and B combine to form a single branch CB. Then branch A is labeled as a sub-object of this branch; much like a twig on a limb. When we do this for all axis branches, we obtain a description which reflects the shape of an object and also the hierarchy of sub-objects which make up the object. This multiresolution shape description can then be used to focus on image structure as a function of scale.

Now we return to the problem of obtaining a multiresolution sequence of binary images. The two alternatives we consider are boundary blurring and figure blurring. Both techniques yield acceptable results, but there can be problems. When boundary blurring is used, object topology is maintained but figural similarity is frequently not preserved. As a result, two shapes with different topologies may look similar at one scale yet look quite different at another scale (see Figure 3). This problem can be avoided by blurring the object figure.

When the characteristic function representing the object figure is blurred with a Gaussian, it gives us a grey-scale image. To obtain a binary image again requires an arbitrary selection of a level curve. To us, the most natural choice is to select the level curve which preserves the object's area, but all choices we know of can result in topological changes in the boundary. Thus, two objects may look similar at different resolutions yet have quite different topologies (see Figure 3). The problem of selecting a single level curve leads us to an important observation.

Binary images are a special case of grey-scale images; they are images which just have two values. As a consequence, binary images should be treated like grey-scale images. In particular, grey-scale shape descriptions should be applied to describe the figure. To preserve causality under resolution reduction, Gaussian blurring of the intensities should be used to impose a resolution hierarchy (Koenderink, 1984). While this further motivates our investigation of grey-scale shape descriptions, we now return to level curve shape description.

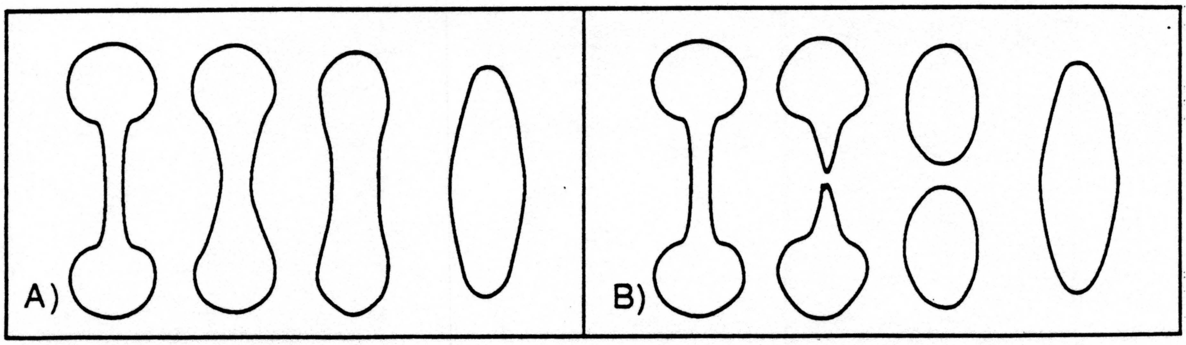

Figure 3. Three stages of a) boundary blurring and b) figure blurring. Notice the changes in topology in the object for figure blurring.

3.2 Boundary Curvature

The second major approach to shape description is based on boundary properties. Several researchers have focused on boundary curvature because it reflects the *bending* of the object, an essential aspect of shape (Brady and Asada, 1984). Extreme points of curvature (local maxima and minima) on the boundary can be used to characterize shapes.

If we decompose an object's boundary into sections bounded by two adjacent curvature minima, we obtain curve segments which (Richards and Hoffman, 1985) call *codons*. Each codon contains a single curvature maxima and can be classified into five types, depending on the signs of these curvature extrema (see Figure 4). By considering sections bounded by two adjacent curvature maxima, we obtain the *codon duals* described by (Leyton, 1986). These can be classified by simply changing the sign and type of each curvature extrema in our codon classification (see Figure 4).

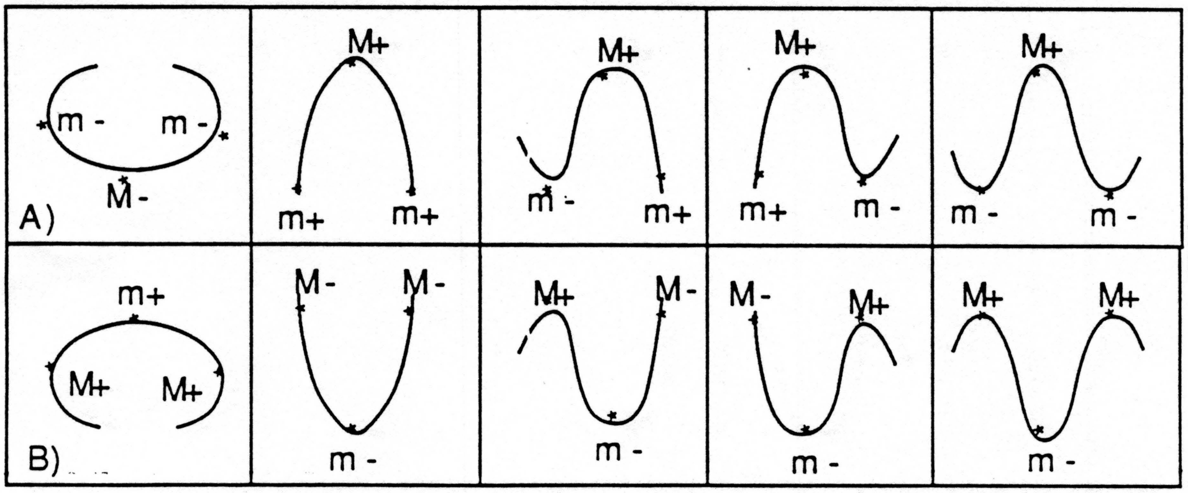

Figure 4. Boundary segments representing a) five types of codons and b) the five corresponding codon duals. The curvature extrema are labeled: M+ positive maxima, M- negative maxima, m+ positive minima, m- negative minima.

Boundary curvature is also sensitive to image detail. As a result, two objects can appear quite similar yet have different codon decompositions. Conversely, objects can have the same codon decomposition yet appear quite different. To resolve this problem, we label each codon with a measure of its importance. At the same time, we impose a hierarchy on the codons which can be used to distinguish object shapes.

Under resolution reduction, the boundary will tend to simplify. This will cause two adjacent curvature extremum (a local maximum and a local minimum) to move together and annihilate into an inflection point. When this happens, the number of curvature extremum is decreased by two. Thus two adjacent codons become a single codon. The problem is to determine which codon annihilated into which.

Recall that a codon is a boundary segment bounded by two curvature minima with a curvature maxima somewhere between these points. When we have two adjacent codons, we have a sequence of five curvature extrema of the form: (1min, 2max, 3min, 4max, 5min), where codon A consists of (1min, 2max, 3min) and codon B consists of (3min, 4max, 5min). We say that codon A annihilates into codon B if 2max and 3min are blurred into an inflection point. Here, two of the three curvature extrema which comprise codon A have disappeared (see Figure 5). Similarly, we say that codon B annihilates into codon A when 3min and 4max are blurred together.

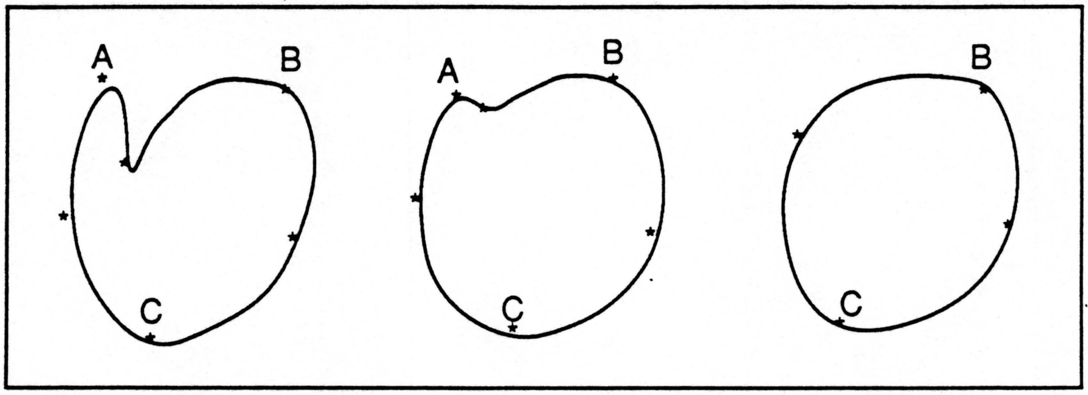

Figure 5. Codon annihilation under resolution reduction. Here codon A is determined to be a sub-object of codon B.

We record the level of resolution required to annihilate a codon as its scale. By also recording which codon blurred into which, we establish a codon hierarchy. This multiresolution shape description can then be used to focus on object curvature as a function of scale to compare and contrast objects.

3.3 Relationship Between Symmetric Axis and Boundary Curvature

Leyton has described an important relationship between these shape description methods. He has shown that *each codon has associated with it a unique axis of symmetry* and that this line terminates at or near the point of maximal curvature for the codon. Similarly, each codon dual has an symmetry line which terminates at a local curvature minima.

The type of symmetric axis associated with these boundary segments depends on the type and sign of the curvature extrema. Internal symmetric axes terminate at positive maxima, external symmetric axes terminate at negative minima, and global symmetric axes are associated with negative maxima and positive minima (see Figure 6).

This gives us a tool for studying the symmetric axis. Once we have decomposed the object boundary into codons (or at least located the curvature extrema), we know how many branch endpoints there are and also their locations. This information could be helpful for computing the SAT, but it is more important when we consider multiple resolution techniques.

We have seen that blurring imposes a scale based hierarchy on symmetric axis branches and on codons. If we recall the relationship between points of maximal curvature and axis endpoints, it is clear that these hierarchies are also related. The scale which causes the annihilation of an axis branch is equal to to the annihilation scale for the associated codon and vice versa.

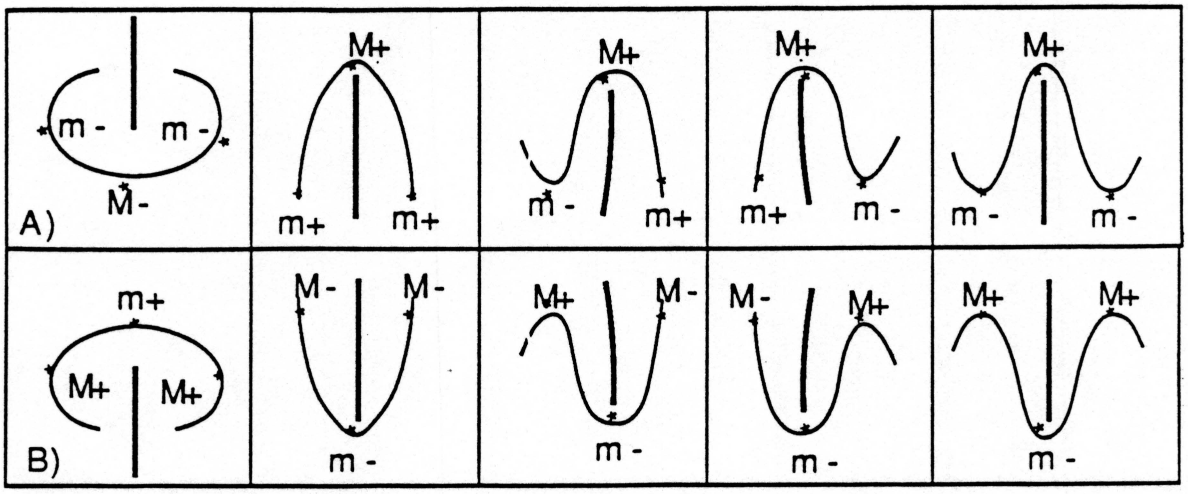

Figure 6. Relationship between the symmetric axis and a) five types of codons and b) the five types of codon duals. The axis is marked in bold.

We can use this fact to our advantage when computing these shape descriptions. Our work to date indicates that codons are easier to compute and follow through scale space than the symmetric axis. Thus, to compute the axis hierarchy we follow codons through multiple resolutions and use the correspondence between branches and codons to determine the scale of the components of the symmetric axis.

4 GREY-SCALE SHAPE DESCRIPTION

Now we return to our main task, to describe the shape of grey-scale images. In the second section, we saw the need to treat the intensity and spatial dimensions separately when describing grey-scale images. This led us to an investigation of shape description methods for level curves in the third section. In the following section we bring these ideas together. This results in two methods to describe grey-scale shape in terms of the shape of individual level curves.

4.1 Pile of Axes

First, we consider the behavior of axes of symmetry for the collection of level curves which represent a grey-scale image. To visualize this description, we embed this collection of axes in three dimensions at their respective intensity levels. This yields a figure based shape description we call the *Symmetric Axis Pile.* When scale information is also recorded, we obtain the *Multiresolution Symmetric Axis Pile.*

What does this pile of axes describe? Because the level curves of the graph (x, y, I(x,y)) partition each level into an inside and an outside, we can describe two things. The volume *below* the image surface can be described by the internal symmetric axis for each level. Similarly, the volume *above* the image surface can be described by the external symmetric axis for each level. What do these volumes represent? Since image intensity corresponds to the height of the image surface, the volume below the image surface can be used to describe the shape of *light* structures in the image. The volume above the image surface can be used to describe the shape of *dark* of regions in the image.

One of the strengths of the symmetric axis is its ability to describe the shape of individual components of an object, and combine this information to describe the whole object. Naturally, we are interested in how our new shape description behaves in this respect. This requires us to focus on four aspects of our representation:

1) the basic elements of the symmetric axis pile,
2) the bending and branching behavior of these structures,
3) the behavior of the radius function for these structures,
4) the annihilation of structures under multiple resolutions.

To simplify our analysis, we assume that the intensity function I(x,y) is smooth and continuous and that the critical points on this surface are isolated and non-degenerate.

Axis Connection

To determine the basic elements of the symmetric axis pile, we need to understand how axes are connected from one level to the next. We begin by examining the behavior of level curves. We have assumed that I(x,y) is smooth and continuous, so curves defined by $I(x,y) = C$ will vary smoothly with C, except at critical points. Because the symmetric axis varies smoothly with the region it represents, the collection of axes for these level curves will form smooth branching surfaces in three dimensions (see Figure 7). We call these surfaces ***symmetric axis sheets.***

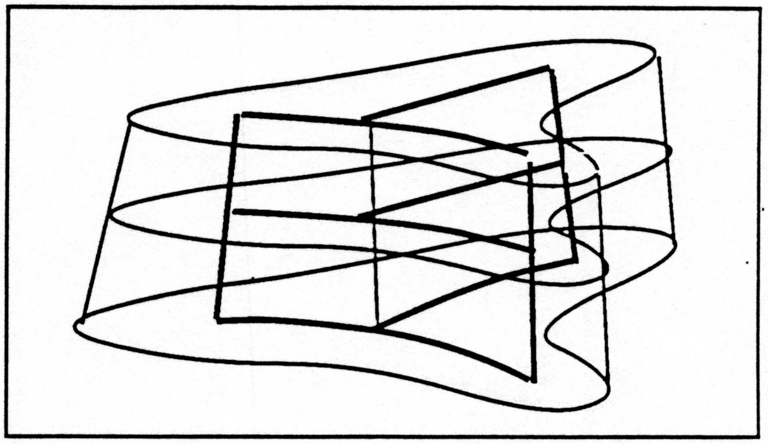

Figure 7. The level curves and corresponding symmetric axis pile for part of a synthetic grey-scale image. The shaded branching surfaces are called symmetric axis sheets.

At critical points, the topology of level curves changes abruptly. At local extrema, level curves reduce to a point and then disappear (depending on the direction from which the extremum is approached). At saddle points, level curves come together, cross and then come apart again. The symmetric axis pile near these regions also changes abruptly. The remainder of this section investigates sheet behavior near critical points and the early indications of shape this behavior provides.

At a local maximum, the symmetric axis sheet for the region under the image surface shrinks with the level curve until it disappears at the critical point. The axis sheet above the surface near a local minimum behaves similarly. These points are called ***sheet terminations*** (see Figure 8). They are an indication of locally lightest or darkest spots in the image. The behavior of sheets on the opposite side of the image surface near these extremum is more complicated.

Consider what happens to the symmetric axis for an object with a hole, as the hole gradually shrinks and then disappears. Initially, the axis for the object loops around the hole. This loop shrinks slightly as the hole shrinks, but suddenly disappears when the hole disappears. At the same time, a new piece of axis down the center of the object appears. This is exactly the situation we observe near local extremum in an image.

If we look at the symmetric axis just above a local minimum, we find that the symmetric axis sheet forms a loop around the indentation near the minimum, and that this loop disappears and another axis sheet appears as we move below the extremum. Similar axis behavior is also observed for the axis above the surface near a local maximum. These changes in symmetric axis sheets are called *loop terminations* (see Figure 8). They give us an indication of the nesting of dark regions within light regions and vice versa.

The level curves near a saddle point cross each other. If we calculate the symmetric axis pile for the region under the image surface near such crossing points, we find that the axis sheet separates into two pieces at the saddle point as we move up in intensity. These are called *axis tears* (see Figure 8). The same behavior is observed for the axis pile above the image surface except that the sheets tear apart as we go down in intensity. Thus, we call saddle points *tear points* of the symmetric axis. These points are special for two other reasons.

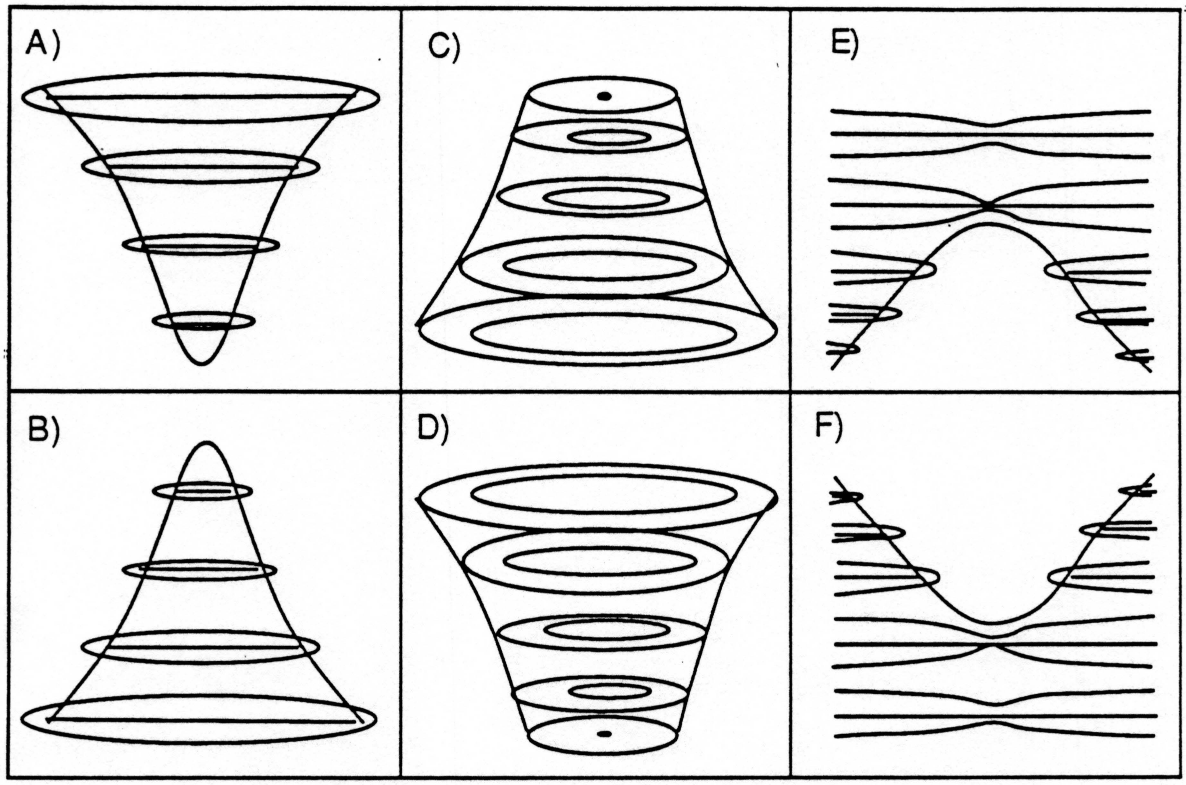

Figure 8. Symmetric axis behavior near critical points. Sheet terminations a) above and b) below the image surface. Loop terminations c) above and d) below the image surface. Axis tears e) above and f) below the image surface.

First, level curves through saddle points describe the nesting of hills and valleys in the image (Blicher, 1985). Thus, the axis sheets can be partitioned at these levels to obtain descriptions of local light and dark regions of the image. Second, saddle points are the only points in common to both axis piles, so they act as connection points between the axis pile below the surface and the axis pile above the surface. This adds coherence to our shape description which can be exploited to describe the relationship between local light and dark regions in the image.

Axis Bending and Branching

While critical point behavior yields a basic understanding of the relationships between light and dark regions in the image, additional information about the shape of these regions is conveyed by the branching and bending of symmetric axis sheets. How individual sheets bend gives us an indication of the shape of individual light and dark regions of the image. How these sheets combine to form branching surfaces captures the global branching structure of the image being described. By combining these shape properties, we can describe the general shape of the grey-scale image.

The bending of axis sheets reflect two different image properties. When the bending is in the spatial dimensions, it captures how ridges (or valleys) in the image are bending (see Figure 7). This is similar to the bending of binary images described by the 2D symmetric axis, so a similar classification scheme can be used. When the bending is in the intensity dimension, it reflects the *asymmetry* of ridge (or valley) profiles (see Figure 9). Once we have a description of the shape of individual sheets, we need to consider how these sheets are connected.

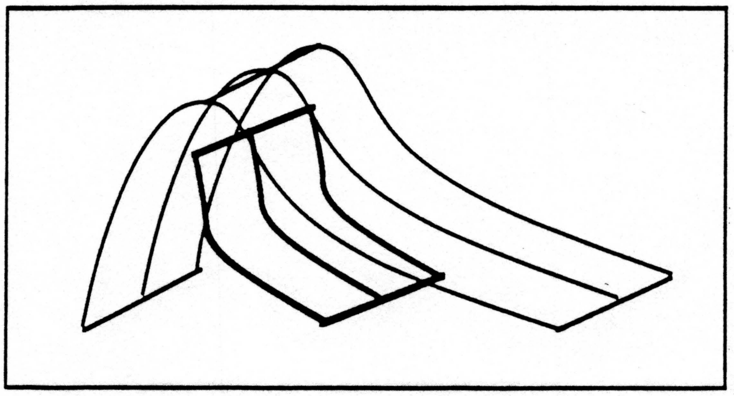

Figure 9. Asymmetry of ridge profiles reflected by axis bending.

The branching structure of sheets above and below the image surface corresponds to the structure of dark and light regions in the image respectively. Each light ridge-like structure in the grey-scale image is described by an axis sheet *below* the image surface. The connections between the axis sheets reflect the connections between these ridges. Similarly, the branching of each dark valley-like structure in the image is described by the symmetric axis pile *above* the surface. While axis piles capture the bending and branching of light and dark regions of the image in a natural way, we still need to consider the width of the regions.

Axis Radius Function

Recall that each point on the 2D symmetric axis has associated with it the radius of the maximal circle at that point. This function is used to describe the widening and narrowing of individual axis branches. When this notion is extended to our symmetric axis pile, we have a radius function defined for each point on our axis pile which describes the *width* of the image at that point. Changes in the radius function along axis sheets gives us two indications of the shape of the grey-scale image.

Width changes in the spatial dimensions reflect the *widening* and *narrowing* of ridges and valleys. This happens when the first derivative of the radius function is positive and negative respectively. The second derivative behavior describes the *flaring* and *cupping* of ridges and valleys. The seven combinations of these derivative properties correspond to those described by Blum and Nagel for the radius function of the 2D symmetric axis.

The *sharpness* and *roundness* of ridges or valleys corresponds to width changes in the intensity dimension (see Figure 10). Consider the radius function as we go down the sheet for a ridge. Because the image surface is described by a function, the radius increases monotonically as we go down. Hence, the first derivative is always positive. When the second derivative is positive, the ridge appears sharp at the top, because it is flaring as we do down. Conversely, ridges which appear round have a negative second derivative; they are cupping as we go down. This analysis extends to valleys by considering the radius function as we go up the sheet corresponding to the valley.

When the shape properties provided by the radius function are combined with the other properties we have described, we obtain an overall description of the shape of a two dimensional grey-scale image at a single resolution. Next, we discuss the means of imposing a hierarchy on the symmetric axis sheets by multiresolution analysis.

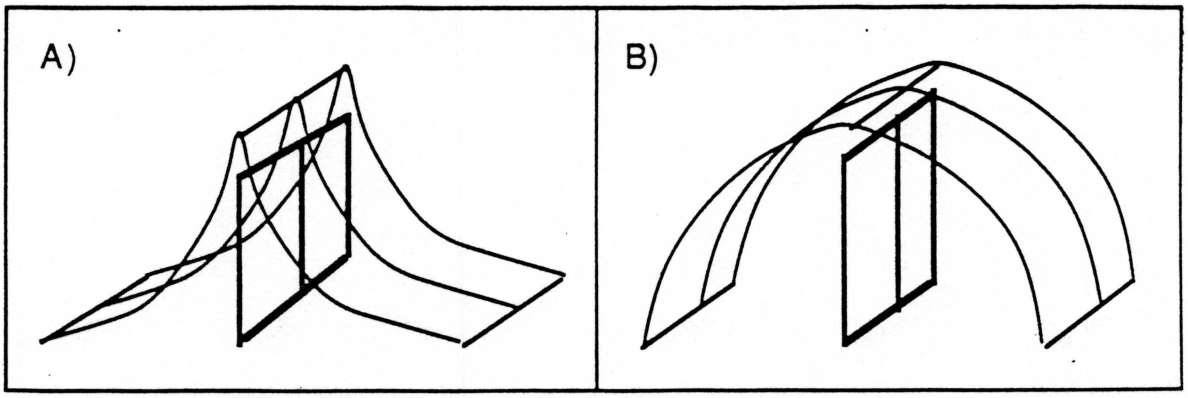

Figure 10. Roundness properties determined by second derivative behavior in the intensity direction. Second derivative is a) positive and b) negative for the ridge shown.

Axis Hierarchy

The symmetric axis pile we have described captures many aspects of an image's shape. Unfortunately, as with the ordinary symmetric axis, it is too sensitive. Noise and small image details often produce "unimportant" axis sheets in our shape description. Our solution is to label each sheet with a measure of its importance and define a hierarchy on axis sheets. Then application programs can focus on image structure as a function of scale.

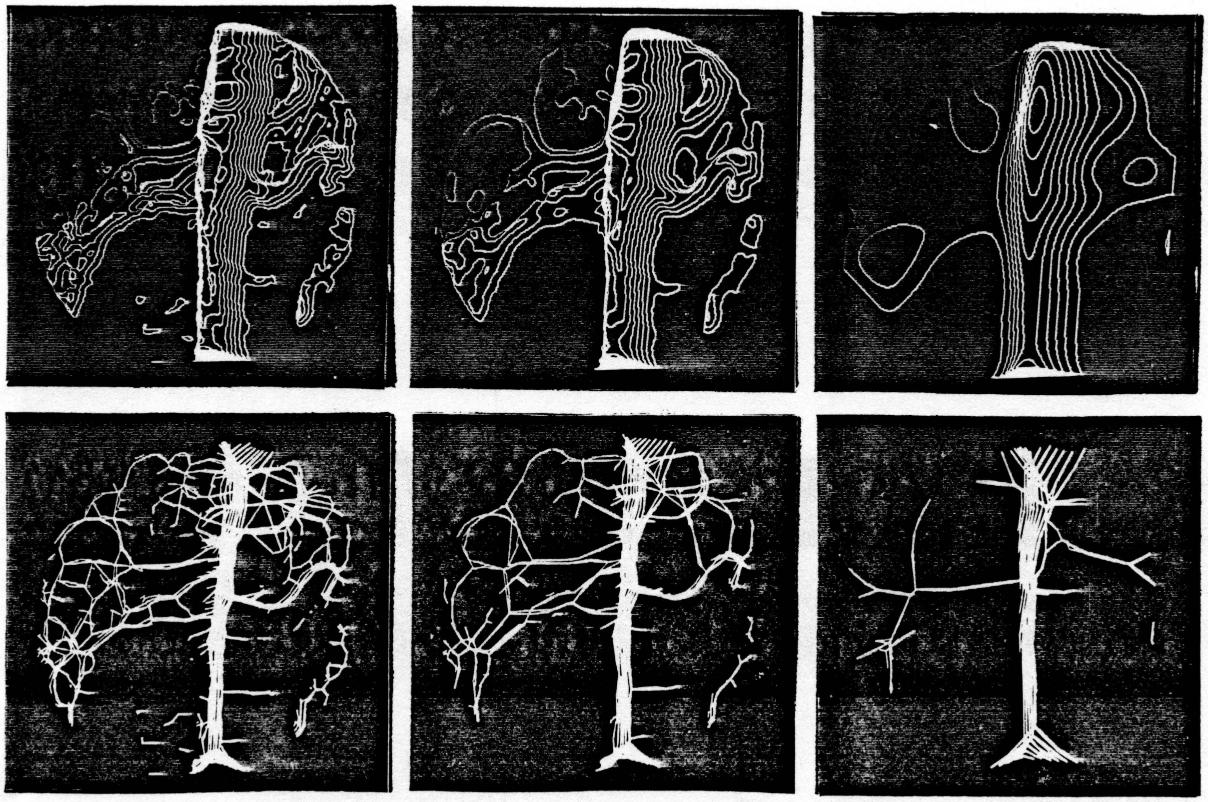

Figure 11. Level curves and corresponding symmetric axis piles for a digital subtraction angiogram reflecting three levels of blurring.

We determine the *scale* of axis sheets by continuously blurring the grey-scale image with a Gaussian and detecting the level of blurring required to cause each sheet to annihilate. After a small amount of blurring, the image structures corresponding to noise and small details will be annihilated. Because the symmetric axis pile varies smoothly with the region it represents, the axis sheets for these unimportant details will also disappear. Larger image structures persist and so do their corresponding axis sheets (see Figure 11). We continue this blurring until only one axis sheet remains (this occurs when the image becomes an elliptical blob). This process is also used to define a hierarchy on axis sheets.

When one sheet blurs into another, the former sheet is defined to be a sub-object of the latter. When this is applied to every sheet in the pile, we obtain a hierarchical representation of the branching of axis sheets. While this hierarchical shape description yields a useful tool for studying an object's shape as a function of scale, it is difficult to compute because of the high dimensionality of the symmetric axis pile followed through scale space. To simplify this computational problem, we can exploit the relationship between axes of symmetry and boundary curvature.

4.2 Vertex Curves

We need to extend our analysis of boundary curvature to handle grey-scale images. To study curvature properties of the image surface, we extend our analysis of boundary curvature to all the level curves which represent the image. Again, we use the extreme points of boundary curvature (vertices) to characterize the image surface. By following the movement of vertices from one level curve to the next, we obtain curves on the image surface we call *Vertex Curves*. Since curvature extrema are the endpoints of codons, these curves define the boundaries of an image decomposition we call *Codon Districts*. Multiresolution analysis then yields an image representation which describes the spatial curvature properties of the image as a function of scale.

To understand the structure of vertex curves and the shape information they convey, we must examine the following:

1) the connection of curvature extrema from one intensity level to the next,
2) the branching of vertex curves,
3) the surface partition defined by vertex curves,
4) the behavior of vertex curves under multiple resolutions.

Again, we simplify our analysis by assuming that the intensity function is smooth and that critical points are isolated and non-degenerate.

Connection of Curvature Extrema

First, we investigate the behavior of curvature extrema. When the surface (x, y, I(x,y)) is smooth and continuous, the level curves defined by I(x,y) = C also vary smoothly, except at critical points. It follows that the curvature for these level curves will vary smoothly and that the points of local maximum and minimum curvature are connected and form curves on the image surface.

Since vertex curves are based on level curve behavior, they reflect the *spatial* curvature properties of the image. They tell us where the bending of level curves is most extreme. Vertex curves consisting of points which have positive maximal curvature correspond to tops of ridges in the image. Similarly, the bottoms of valleys in the image are marked by vertex curves consisting of negative curvature minima. The relationships among ridges and valleys are reflected in the branching structure of the vertex curves.

Branching of Vertex Curves

We begin our analysis of branching by considering vertex curve behavior near critical points (see Figure 12). At local intensity extrema, the topology of level curves changes. Slightly below a local intensity maxima, the level curve is generally elliptical. Thus, it contains four curvature extrema, two maxima and two minima. As we move up in intensity, the level curves shrink and these four vertices approach each other. At the local maximum, the level curve becomes a point and the four vertex curves defined by these points meet. Similarly, four vertex curves meet at each intensity minima in the image.

At saddle points, the topology of level curves also changes abruptly. The level curves at the saddle point cross while the level curves slightly above and below each saddle point are generally hyperbolic. Thus, we have two points of locally maximal curvature on the level curves above the critical point and two points of locally minimal curvature on the level curves below. As a result, four vertex curves meet at each saddle point in the image. Two curves of local curvature maxima go uphill and two curves of local curvature minima go downhill.

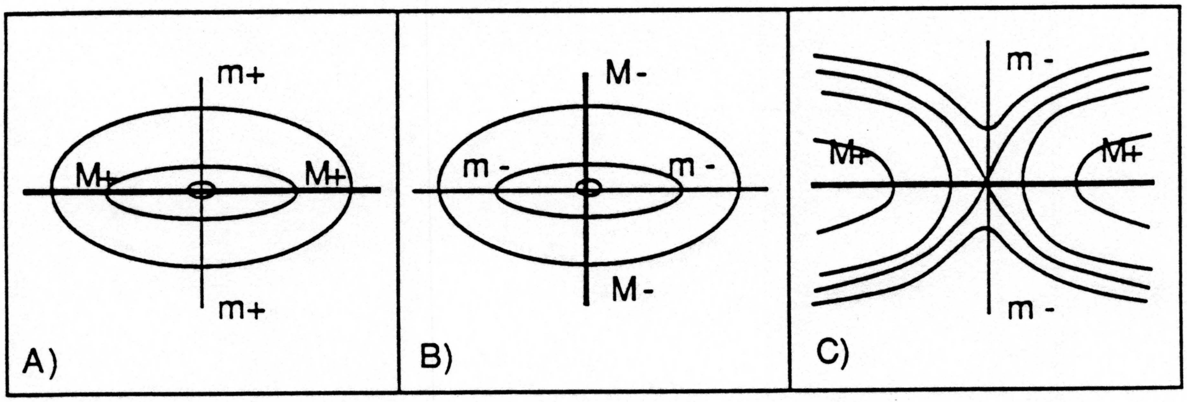

Figure 12. Vertex curve behavior near a) intensity maxima, b) intensity minima, and c) saddle points in an image.

Finally, we consider curvature inflection points on level curves. If we consider the level curves slightly above and below a level curve which has an inflection point, we find that one curve has two curvature extrema (one maxima and one minima) near the inflection point while the other has none. Thus, two vertex curves originate at each inflection point on the image surface (see Figure 13). Since one vertex curve consists of local curvature maxima and the other local minima, these curves mark the tops of ridges and the bottoms of valleys respectively.

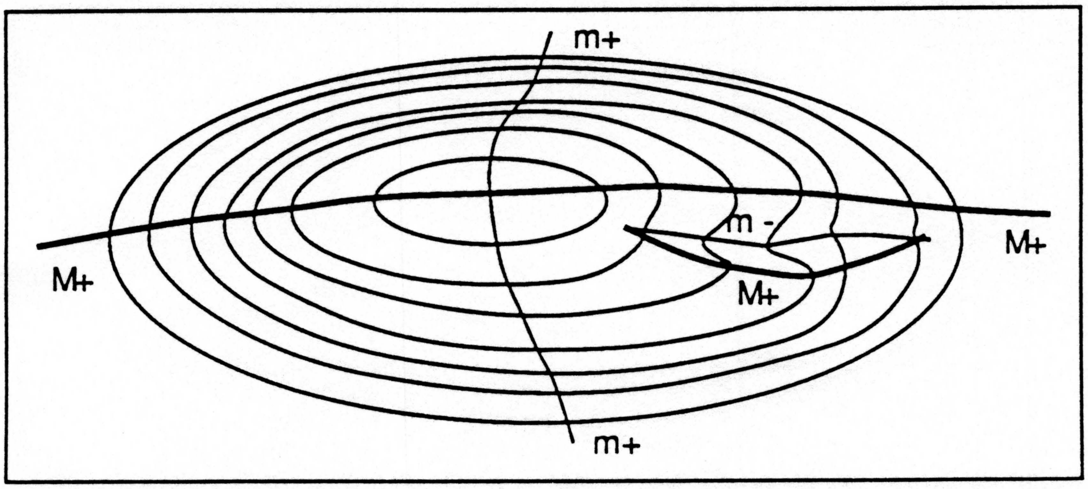

Figure 13. Vertex curve behavior near curvature inflection points. Here we have a small valley on the side of a hill whose extent is marked by two vertex curves.

Surface Partition

Marking ridge tops and valley bottoms vertex curves partition the image into regions which are similar to the *Slope Districts* described by (Nackman, 1984). Maxima vertex curves mark ridge tops much like the *ridge lines* which define watersheds. Similarly, minima vertex curves behave like *course lines.* Thus, vertex curves through the critical points in the image partition the image into hills and valleys. Vertex curves through inflection points reflect a different type of nesting of ridges and valleys.

Consider a small valley on a ridge side. In this case, a new maxima vertex curve marks the sub-ridge while the sub-valley is marked by a new minima vertex curve. The vertex curves defining this sub-region emerge at one curvature inflection point and disappear at another on the image surface (see Figure 13). Because these vertex curves do not pass through intensity extrema, this type of image structure would not be captured by slope districts. This is one of the advantages of our shape description.

More natural subdivision in terms of ridge areas and valley areas rather than sides of them is given in terms of the codon structure centered on vertex curves. These *codon districts* are each bounded by two vertex curves of the same type (either maxima or minima) and contain one vertex curve of the opposite type. This decomposition captures the extent of ridges and valleys in the image. Fundamental changes in the codon structure for level curves occur as we pass through saddle points. In particular, four minima vertex curves above the saddle point terminate at the level curve through the saddle point. Two new minima vertex curves start at the saddle point (see Figure 12). Hence, the topology change in the level curves at this point is reflected by the reduction of the number of level curve codons by two.

Multiresolution Analysis of Vertex Curves

The complexity of this new shape description and its sensitivity to small structures or noise can again be avoided by imposing a scale based hierarchy on the vertex curves and codon districts. Applications can then focus on the curvature properties of an image as a function of scale.

Figure 14. Images and corresponding vertex curves under blurring.

When we use Gaussian blurring to obtain a multiresolution sequence of grey-scale images, two things happen to simplify the image structure. First, the critical points of the image move together and annihilate in pairs (one saddle and one extremum). Second, the "curvature" of the surface is smoothed. Hence, the branching structure of vertex curves simplifies with blurring (see Figure 14).

While much work remains to categorize vertex curve behavior under blurring, it appears that the annihilation of vertex curves corresponds to the annihilation of branching structures in the image. This observation is explained by the relationship between our two grey-scale shape descriptions.

4.3 Relationship Between Axis Piles and Vertex Curves

Both the symmetric axis pile and codon districts reflect the shape of grey-scale images on a level by level basis. For this reason, these methods retain the fundamental relationships between codons and axes of symmetry. In particular, *each codon dis-*

trict has associated with it an axis sheet which terminates at or near the interior vertex curve. Thus, positive maxima vertex curves mark the ends of the branches of the symmetric axis pile under the image surface. Similarly, negative minima vertex curves mark the ends of branches above the image surface. The other vertex curves mark the extent of image surface associated with each axis branch.

As a result, this correspondence gives us a way to impose a scale based hierarchy on one description given the other. We have found that following the symmetric axis pile through scale space is a difficult task compared to computing surface curvature properties. It is our hope to calculate the annihilation scale of each axis branch by following the annihilation of the corresponding vertex curves. Then, the one to one correspondence between the original vertex curve segments and axis sheets can be used to derive the scale of axis sheets. Thus, we can compute the axis hierarchy.

Finally, this correspondence gives us two ways to look at the shape of an object. We can focus on the symmetric axis to get an understanding of the branching and bending of image components, or we can focus on the codon districts to partition the image surface into nested ridges and valleys. Both methods allow us to study the branching structure of ridge tops and valley bottoms.

5 APPLICATIONS OF SHAPE DESCRIPTION METHODS

The two shape descriptions described in section four capture many aspects of grey-scale shape which can be exploited by computer vision applications. In this section we describe how these shape descriptions can be used to segment images and study the deformation of binary objects.

5.1 Segmentation

Segmentation is the process of partitioning an image into "natural" regions. Lifshitz (Lifshitz, 1987) and others have shown that multiresolution analysis can be applied to obtain more successful segmentations than conventional techniques based on local pixel properties or measures of edge strength (Ballard and Brown, 1982).

Following Crowley (Crowley and Parker, 1984), we suggest that yet better segmentations can be obtained by taking image shape into account. Therefore, we propose to use our multiresolution shape descriptions as a tool for segmenting two dimensional images. These segmentation techniques hinge on our association of pixels to components of the shape description for these images. We begin with techniques to segment binary images and then extend these techniques to grey-scale images.

Binary Image Segmentation

To segment a binary image, we use the multiresolution symmetric axis. The association between pixels and segments is of prime importance. Consider a point on the symmetric axis. If we draw the maximal circle centered at that point, we define the region of the figure associated with that point of the axis. Extending this notion to a branch of the two dimensional symmetric axis, we find that the union of all maximal circles centered on the axis defines the region of the object associated with that axis segment.

A problem with this technique is that axis regions can overlap. Pixels near branch points in the axis can be covered by several maximal circles. One solution is to use the scale of each symmetric axis component to determine which axis the point is associated with. There are several options here. We can associate pixels with the largest scale axis or the smallest scale axis of which it is a part. These result in quite different image segmentations (see Figure 15). When pixels are included as part of the largest scale axis to which they belong, the object is naturally decomposed into branches and sub-branches. When the pixels are associated with the smallest scale axis, we obtain a less natural decomposition because sub-branches subsume part of larger scale branches.

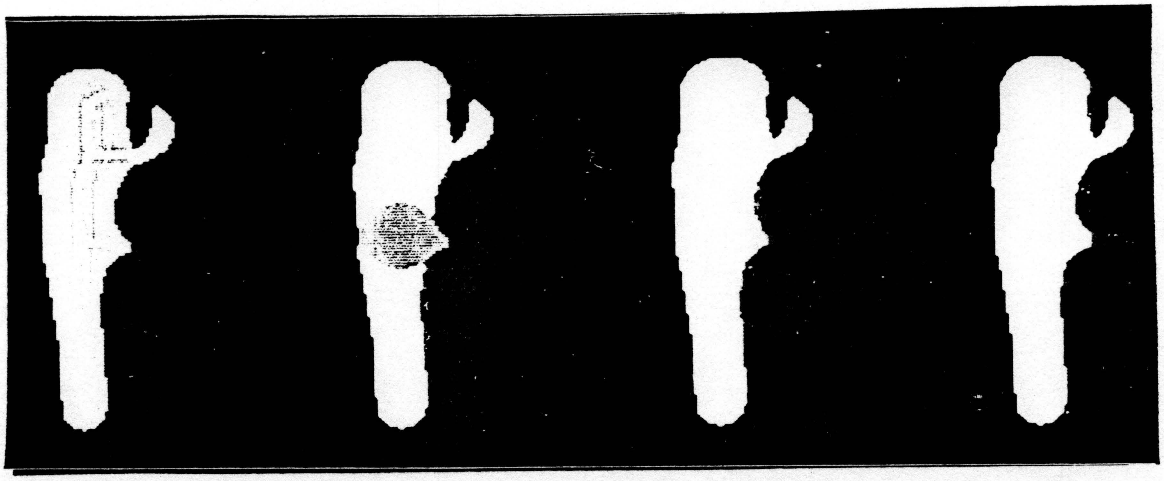

Figure 15. Symmetric axis for a binary object and three segmentations based on pixel associations to axis segments: a) pixels displayed as part of smallest scale axis, b) pixels displayed as largest scale axis, and c) small scale regions displayed as part of larger scale regions.

The axis hierarchy our shape description imposes can be used to segment the image into regions according to scale. Once a scale is specified, objects with smaller scales can be ignored or associated with the segment they "belong to". We use the axis hierarchy to determine this relationship. In this way, pixels which reflect small image detail can be omitted or displayed as part of larger scale segments in a natural way (see Figure 15). This technique can be generalized to handle grey-scale images.

Grey-Scale Image Segmentation

The figure of a grey-scale object (the region above or below the image surface) can be segmented using the multiresolution symmetric axis pile. The key to this process is associating voxels in these regions with sheets of the symmetric axis pile. Extending the ideas from binary images, we define the region associated with an axis sheet to be the union of all maximal circles for all points on the sheet. To associate this volume with the original grey-scale image requires that we project this region in the intensity dimension onto the x-y plane and integrate its volume. This technique will give us a grey-scale image R(x,y) which represents the region.

Just as the regions associated with branches in the two dimensional symmetric axis can overlap, we find that the volumes associated with axis sheets can also overlap. Again, we solve this problem by using the scale of axis sheets to associate voxels in the figure with only one axis sheet. Now, the sum of all volumes in our segmentation defines the figure of the grey-scale object. As a consequence, the sum of all region images R(x,y) equals our original image I(x,y). This is an important result. It gives us two novel ways to segment grey-scale images.

When voxels are associated with the largest scale axis sheet to which they belong, we obtain image segments which factor out small image details. When voxels are associated with the smallest scale axis sheet, the image decomposition looks less natural because small scale objects subsume larger scale objects. When the hierarchy on symmetric axis sheets is exploited, we can derive two segmentations based on scale. Objects below a specified scale can either be ignored or displayed as part of their parent objects. This approach is compelling because noise and small detail can be filtered ot or displayed as part of the larger scale objects.

5.2 Deformation Analysis of Binary Images

A singularly interesting avenue of investigation in medical imaging is the question of shape deformation. While the diagnostic applicability of the extraction of shape data as vectors for structural or statistical pattern recognition and shape comparison are obvious, in many instances it would be illuminating to have tools to describe how one shape is transformed into another — the change from carcinoma in situ of the skin to invasive cancer, for instance, or the examination of developmental processes and abnormalities.

Examination of shape change as deformation was proposed in the biologic literature as early as the beginning of this century by Sir D'Arcy Thompson. Bookstein has described this deformation in terms of tensors, the "biorthogonal grid" (Bookstein, 1984, 1985, 1986). The applications of the statistical evaluation of shape change described in this manner have been elegantly demonstrated in such diverse systems as craniofacial growth and post-ischemic cardiac kinetics. Such applications are limited, however, in that the objects being evaluated must have recognizable landmarks, either inherent in the object itself (such as the anatomic landmarks of the skull) or imposed (as by surgical implantation of radiodense pellets onto the surface of the heart).

This is a severe limitation in areas of medical diagnosis which examine morphologic changes in soft tissues, such as anatomic pathology, because true or at least recognizable landmarks do not exist — there are no fins on a neoplastic gland. The application of the multiresolution symmetric axis transform may provide a method of allowing the investigator to use tools such as biorthogonal grid deformation analysis, however, by allowing the imposition of pseudo-landmarks based on the SAT.

As described previously, progressive blurring of any image allows the establishment of a hierarchy of SA limbs and twigs and in turn a hierarchy of vertex curves. In binary images, each vertex curve collapses into a point, the vertex, in the original unblurred image. Similarly each pile of SA branch points collapses into a single branch point. As a result, the hierarchy imposed can be taken to apply to the vertices and branch points. These landmarks may then provide shape-based, if not truly biologic landmarks that can serve as the requisite data for biorthogonal grid evaluation.

Inherent in the concept of a hierarchical object description is that the higher level descriptions contain both more important and more global information about the shape than the lower limbs and twigs of the hierarchy. Similarly, changes in the higher level landmarks provide relatively more information for deformation analysis. Of course, using only the highest levels of the SA provides an incomplete description of the shape deformation.

Transforming one symmetric axis into the other does not provide the true deformation, because the SA of the figure after deformation is not equal to the deformed SA. An iterative method of approximating the true shape deformation was recently suggested to us by Bookstein. This approach has the following steps:

1) Establish pseudo-landmarks in the original and deformed shapes by utilization of the hierarchical SAT.

2) Utilize the Bookstein method of geometric construction for the establishment of shape deformation tensors.

3) Perform the inverse of that deformation on the deformed shape to achieve a first approximation of the original shape. A deformation measure for boundary points which are not SAT derived landmarks is interpolated from the existing landmarks.

4) Calculate another SAT for the approximation of the original shape calculated in step 3. The difference between this SAT and the original SAT is used to measure object similarity.

5) Take as a new deformed shape the result of applying the approximate inverse deformation to the previous deformed shape. Repeat steps 1 – 4, and continue the process until the approximation converges.

The overall deformation is obtained by concatenating the series of SA-based deformations from the successive repetitions. An example of the application of this method is given in Figure 16.

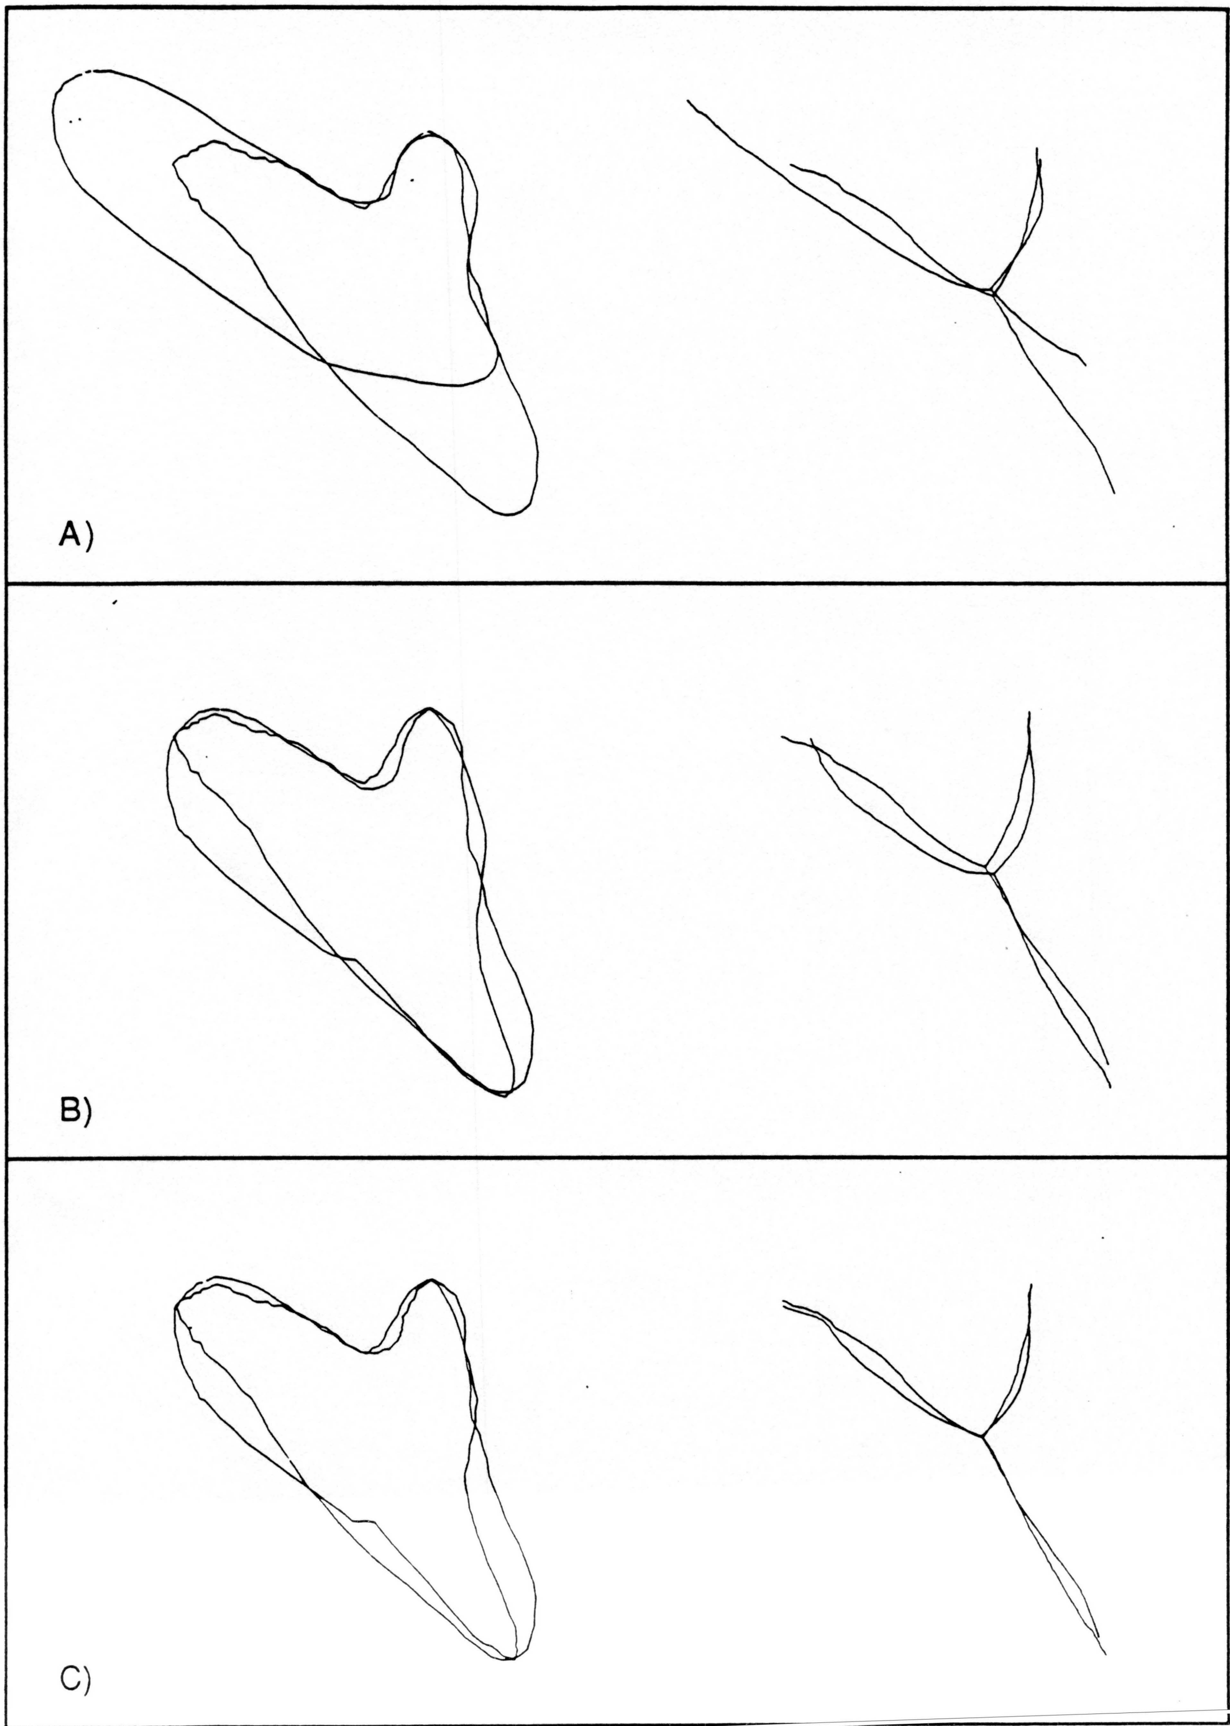

Figure 16. Deformation analysis based on pseudo-landmarks, a) original and deformed objects with corresponding symmetric axes, b) result after one iteration, result after convergence.

Our continued work in this area will involve the following considerations:

1) Utilization of the first one or two levels of the hierarchy ignores the contribution of the lesser limbs of the SA to the deformation. While the lowest level twigs of the hierarchy may indeed be only noise, a way must be developed to integrate secondary information into the deformation analysis.

2) The current method of interpolation of deformation ignores that part of shape information carried by the width properties of the SA. For instance, shapes 1 and 2 in Figure 16 contain the same landmark data. A way must be found to integrate SA width properties into the iterative deformation algorithm.

3) Another possible complication with the hierarchical shape description method of imposing landmarks in a time series is a discontinuous change of hierarchy induced by a large shape change, e.g. when a previously dominant axis segment which would survive as part of a surviving limb decreased in importance, while a previously less important limb becomes dominant. Two answers immediately present themselves. The first is that in dealing with some processes, such as time series, the changes which occur between frames may be small enough to make the point moot. The second possibility is that in a given biologic system such as cell shape change, the change in the shape hierarchy may in fact reflect a concomitant change in the biologic importance of that portion of the cell, and that the change in hierarchical dominance is biologically appropriate.

We have currently investigated the iterative reverse deformation on only a few selected shapes, and need to pursue further experience to investigate whether the progressive approximation of shape features does, in fact, converge for an arbitrary figure.

In addition to providing a method of evaluating shape change, the SA also provides a convenient coordinate system for shape space. In a manner similar to that described for SA hierarchy-based segmentation, one may describe a location within the shape in terms of distance along the most dominant limb and distance along that radius of the tangent circle associated with the limb position which intersects the boundary. By use of the iterative inverse transformations, it is possible to describe the location of a subobject within a deformed shape in terms of the SA of the original shape. The description of subobject location in terms of the original SA provides a description of subobject movement which is invariant over shape change.

An application of this deformation analysis to a biologic process is the tracking of subobjects that move within the primary object while the primary object itself is undergoing a deformation. For instance, consider the subclass of white blood cells, the polymorphonuclear leukocyte, or neutrophil. The neutrophil is that cell which is concerned with finding and destroying various types of contaminants (bacteria, foreign bodies, senescent and dead somatic cells) within the body. As such, it is one of a small number of classes of cells which is capable of independent amoeboid movement within the body, and the ability of this cell to move in an appropriate manner is essential to its function in fighting disease. An organelle called the centriole is the origin of much of the cytoskeleton of the cell, and is thought to interact intimately, if not control, the dynamics of both cell movement and movement of organelles within the cell. As the cell moves or is otherwise activated, the centrioles themselves move within the cell. An understanding of the movement of the centrioles within the cells while the cell in turn is undergoing shape deformation may increase our understanding of cell motility. We plan to attempt to use the methods described above to examine the movement of the centrioles within the cell.

6 CLOSING REMARKS

In summary, we have extended two shape description methods for binary images to describe grey-scale images hierarchically by scale. Since binary images are a special case of grey-scale images, these descriptions also describe binary images hierarchically by scale. One shape description, the symmetric axis pile, allows us to focus on the branching and bending of image structures. The other description, vertex curves, gives us insight into the curvature of the image surface. We have also demonstrated that the relationship between these methods can be exploited to obtain a powerful shape description tool.

We still need to analyze the behavior of vertex curves under multiple resolutions. We also need to complete our implementation of the symmetric axis pile, incorporating the hierarchy information provided by multiresolution vertex curves. Following the completion of this work for 2D images, we could extend these shape descriptions to 3D grey-scale images by studying the behavior of the higher dimension level surfaces which represent the 3D image.

We have prototyped two applications of the symmetric axis pile. Image segmentation is accomplished by associating pixels in the image with components of the shape description. Then the hierarchical nature of the representation is exploited to produce image segments based on shape and scale. We have also suggested how landmarks based on axis endpoints can be used to determine deformation maps. These can then be used to define a coordinate system for studying aspects of shape change.

Our future work will also address new applications of our shape descriptions. For example, *object analysis* applications could make use of the descriptive features of the symmetric axis pile to compare and contrast the structures of grey-scale objects. *Object recognition* could be facilitated by defining a metric on shape space (to measure the difference between two shapes) and using the scale based hierarchy defined by our shape descriptions to implement top down matching of object models.

ACKNOWLEDGEMENTS

We are indebted to Fred L. Bookstein for useful discussions on deformation analysis. Many thanks to Susan Gauch for document preparation — a job which is deceptively difficult. We also thank Bo Strain and Karen Curran for photography. This research reported in this paper was carried out under partial support of NIH grant number R01-CA39060.

REFERENCES

Ballard, D. H. and Brown, C. M. (1982). Computer Vision, Prentice Hall, Englewood Cliffs, NJ.

Blicher, A. P. (1985). Edge Detection and Geometric Methods in Computer Vision, Stanford Report STAN–CS–85–1041, 189–205.

Blum, H. (1974). A Geometry for Biology, Annals of the New York Academy of Sciences, 231, 19–30.

Blum, H. and Nagel, R. N. (1978). Shape Description using Weighted Symmetric Axis Features, Pattern Recognition, 10, 167–180.

Bookstein, F. (1984). A Statistical Method for Biological Shape Comparisons, J. Theor. Biol. 107,475-520.

Bookstein, F., Chernoff, B. Elder, R. Humphries, J., Smith, G., Stauss, R. (1985). Morphometrics in Evolutionary Biology. Special Publication 15. The Academy of Natural Sciences of Philadelphia, Philadelphia, PA.

Bookstein, F. L. (1986). From Medical Images to the Biometrics of Form, Center for Human Growth and Development, University of Michigan.

Brady, M. and Asada, H. (1984). Smoothed Local Symmetries and their Implementation, MIT AI memo 757.

Crowley, J. L. and Parker, A. C. (1984). A Representation for Shape Based on Peaks and Ridges in the Difference of Low-Pass Transform, IEEE PAMI, 6, No. 2, 156–170.

Koenderink, J. J. (1984). The Structure of Images, Biological Cybernetics, 50, 363–370.

Leyton, M. (1986). Smooth Processes on Shape, Draft Report, Harvard University.

Lifshitz, L. M. (1987). A Multiresolution Hierarchical Approach to Image Segmentation based on Intensity Extrema, Proceedings of the Xth IPMI International Conference, Plenum Publishing, New York.

Nackman, L. R. (1984). Two-Dimensional Critical Point Configuration Graphs, IEEE PAMI, 6, No. 4, 442–450.

Pizer, S. M., Oliver, W. R., Gauch, J. M., and Bloomberg, S. H. (1986). Hierarchical Figure-Based Shape Description For Medical Imaging, Chapel Hill Technical Report 86-026.

Richards, W. and Hoffman, D. D. (1985). Codon Constraints on Closed 2D Shapes, Computer Vision, Graphics, and Image Processing, 31, 265–281.