

Lecture 10: Shape Description (Regions, cont'd)

©Bryan S. Morse, Brigham Young University, 1998–2000
Last modified on February 2, 2000 at 10:00 AM

Contents

10.1 Distance Transforms (or Distance Maps)	1
10.1.1 Using Distance Maps	1
10.2 Why Medial?	2
10.3 Blum's Medial Axis	2
10.3.1 Definition	2
10.3.2 Medial Hierarchies	3
10.3.3 Radius Values	3
10.3.4 Types of Points	3
10.3.5 Sensitivity to Noise	3
10.3.6 Building and Pruning Hierarchies	3
10.4 Other Medial Representations	4
10.5 Computational Methods	4
10.5.1 Thinning	4
10.5.2 Distance Maps	4
10.5.3 Voronoi Methods	5
10.6 Higher Dimensions	5

Reading

SH&B, 6.3.4–6.4
Castleman, 19.3.4.3

10.1 Distance Transforms (or Distance Maps)

Recall from our earlier discussions that a *distance transform* or *distance map* assigns to each point \bar{x} in the image a minimal distance $D(\bar{x})$ from some locus of points (usually object boundaries). If we build a distance transform from the boundary of the object, we end up with zero at the boundary, increasing values just inside it, yet higher values farther inside, etc. Similarly, we get distances increasing as we go outward from the boundary. (A large number of applications use a positive distance map outside the object and a negative one inside, or vice versa.)

10.1.1 Using Distance Maps

Distance maps have several applications:

- Interpolation
Powerful techniques have been developed for interpolating between two shapes by interpolating between their (positive outside, negative inside) distance maps. By using the 0-set of the map as the boundary at each step in the iteration, one can find intermediate boundaries between the two original ones. This is called interpolation by *implicit functions*.
- Shape description
Distance functions can also be used to help create representations for shape. Since shape and spatial relationships (distance) are so integrally related, this would seem to be an obvious application. One specific such application is the *medial axis*, a shape representation that we'll spend the rest of this lecture discussing.

10.2 Why Medial?

All of the methods we've discussed so far can be used to differentiate between shapes, but they don't describe the shape well in any way that people readily understand.

The boundary- and region-based representations we've seen so far give ready access to certain parameters of the shape but not to others. Suppose we wanted to know how wide the object is, or how many branches it has, etc. These are things not readily available from the boundary description. In a sense these are *structural properties* of the shape that go beyond simple boundaries.

To infer larger structure from a shape requires that we be able to group together parts of the boundaries that aren't necessarily adjacent. For example, there is an immediacy to the way we see the two sides of an object. If we look at telephone pole, for example, we can make immediate judgments of how thick or thin it is. We do so by pairing up the two opposite sides and comparing them, even without having to trace completely around the object to ensure that they are indeed part of the same contour.

Once we pair sides in this fashion, it's reasonable to start talking about object middles. These middles give structure to the shape, allowing us to identify trunks, arms and branches, narrowing or widening sections, etc. Such middle-related structures are called *medial* representations.

10.3 Blum's Medial Axis

10.3.1 Definition

The most common form of medial representation is Blum's *medial axis*. (In later years, Harry Blum wrote that he preferred to call it the *symmetric axis*, but others continue to call it the medial axis as part of the larger class of medial representations.)

Blum's medial axis can be defined in a few different intuitive ways, all equivalent:

Grassfire Analogy. Consider a fire started at the boundary that spreads with uniform speed and burns everything in its path. As the fire spreads into the middle, each part eventually meets other parts of the fire started at other parts of the boundary. When these two wavefronts meet, they quench each other because there is nothing left for either to burn. Since the fire spreads at uniform speed, these "quench points" must be equidistant from two different parts of the boundary. The locus of these quench points is the medial axis.

Maximal Discs. Suppose that at each point of the boundary you placed a circle on the inside of the boundary such that it was tangent to it. If you start with a very small circle and expand it outwards, it will eventually "bump into" other parts of the boundary. When you've grown it as large as possible, you'll have a circle that has the following properties:

- It is entirely within the object, and
- It is tangent to the object at more than one point

The locus of centers of all such circles is the medial axis.

Ridges in a Distance Transform. As we progress from away from a particular point on the boundary, we find increasing values in the distance map until we reach a point at which we are now closer to some other part of the boundary than the place that we left. So, the distance increases as you progress inwards until it forms a *ridge*, a locus of points with values higher than those just off the ridge on either side. The locus of these ridges in the distance transform is the medial axis.

There are, of course, various mathematical definitions in the literature as well. They come in different forms, but they typically correspond to one of these three intuitive definitions.

All three of these definitions can be used to construct morphological approaches to finding medial axis.

- If we strip away boundary pixels, we eventually get down to single-pixel thin curves that cannot be safely removed—these approximate the medial axis (grassfire or distance transform definitions).

- If we find explicitly maximal disks by using morphological openings, we can also find the skeleton. In this case, though, we may find that because of discrete pixel locations we end up with a disconnected set of points. This is sometimes called the *morphological skeleton*. Like the medial axis it can be used to reconstruct an object, but it loses important shape information.

10.3.2 Medial Hierarchies

The medial axis is a continuous structure that may branch in places. Each branch corresponds to a different piece of the object. It these branches that give the structural decomposition of the medial axis by forming a hierarchy—each bump on an arm of the shape is a branch of a branch.

10.3.3 Radius Values

At each point of the medial axis, we also have a corresponding r value: the radius of the maximal disc, the distance from the grassfire, or the distance from the distance map.

If we let s be arc length along any given segment of the axis, we have a one-dimensional function $r(s)$ along the segment.

If $\frac{dr}{ds}$ is positive, this part of the object widens as we progress along the axis segment; if it's negative the part narrows.

10.3.4 Types of Points

The easiest definition for which to consider the various properties of the medial axis is the maximal-disk one. Consider one such disk. As mentioned earlier, it touches the boundary in more than one place. These points can come in different flavors:

Two disjoint sets of points. Most of the axis points fall into this category.

Three disjoint sets of points. These are branch points.

Exactly one contiguous set of points. These are end points.

Notice that its possible for the medial disc to touch the boundary in a set of contiguous points rather than a single one. In other words, the boundary is locally circular.

These sets of points don't really change the properties of the resulting axis—it's the number of disjoint sets of points, not the actual number of points themselves, that are important.

10.3.5 Sensitivity to Noise

Sounds great so far, right? Well, the medial axis has one problem—it is extremely sensitive to noise in the boundary. Since the axis has a branch for each structural piece of the object, it has to spawn a new branch for each little bump and wiggle.

The length of the new branch isn't necessarily reflective of the significance of the variation. It must extend all the way out to the base of the bump itself.

10.3.6 Building and Pruning Hierarchies

Because of this sensitivity to noise, many researchers have studied ways to prune the noisy tree-like structure produced by the medial axis. Ways to do this include

- Smoothing the boundary before computing the axis,
- Pairing the boundary points to the corresponding axis segments and pruning based on length of the boundary (Ogniewicz, 1992).

By ordering the segments according to the order in which they would be pruned, you can organize them into a top-down hierarchy.

Pruning away segments or arranging them in order of significance can (to a limited degree) let you compare similar but slightly different objects.

10.4 Other Medial Representations

Let's consider the maximal disc in Blum's definition. As already discussed, it touches the boundary in at least two places.

- The center of the circle lies on the Blum symmetric axis.
- The midpoint of the chord between the two boundary points lies on Brady's *Smooth Locus of Symmetries* or SLS.
- The midpoint of the chord between the two boundary points lies on Leyton's *Process-Induced Symmetric Axis* of PISA.

Brady's SLS definition seems to be more intuitive than Blum's because the axis really does lie directly between the related points. It is more complex to compute, however.

Leyton's PISA definition isn't quite as intuitive, but the resulting representation seems to describe the way you have to push in or push out the boundary to produce the shape (hence the "process-induced" in the name).

Another method, which approaches the Blum definition when taken to the continuous limit, is the *Voronoi Skeleton*. I'll come back to this later.

10.5 Computational Methods

One can use the definitions of the axis to precisely determine the medial axis of a continuous curve. There are, however, two problems:

1. Computational expense.
2. The discrete pixel grid.

For this reason, various algorithms have been produced that produce medial axes that are approximations to Blum's.

10.5.1 Thinning

One way to find the middle of something is through the *thinning* process we discussed earlier. Suppose that we successively peel away thin layers from the object, but only if they don't get rid of single-pixel thick parts of the remaining object. This thinning method is analogous to the grassfire definition of the medial axis. What is left from this thinning is a structure called a *skeleton*, similar in principle to the medial axis.

10.5.2 Distance Maps

Another way to compute medial axes is to use the distance-map definition. We can explicitly compute a distance map and then analyze it to find ridges. This requires two steps:

1. Generate the distance map, and
2. Find its ridges.

Finding Ridges

We won't talk right now about finding ridges in functions, but we'll come back to it later when we discuss geometric methods for local analysis of functions (images).

10.5.3 Voronoi Methods

If we distribute discrete points along the boundary and compute the Voronoi diagram for the points (the locus of all points that are no nearer to one point than another), contained within the Voronoi diagram is a skeletal structure for the shape. If we take this to the limit by adding more and more points along a continuous contour, this Voronoi skeleton approaches the Blum medial axis. [Make sure to convince yourself why.]

We don't have to go to this limit, though, to get a usable skeletal structure.

10.6 Higher Dimensions

The medial structure of a three-dimensional shape isn't composed of one-dimensional curve segments but two-dimensional manifolds within the three-dimensional space. These *medial manifolds* also have useful applications for three-dimensional vision.

Vocabulary

- Distance
- Distance Transform
- Skeleton
- Medial Axis
- Voronoi Diagram