A Hough-Like Medial Axis Response Function

Bryan S. Morse[a], Stephen M. Pizer[a], and Christina A. Burbeck[b]

[a]Department of Computer Science
[b]Department of Psychology
University of North Carolina at Chapel Hill
Chapel Hill, North Carolina 27514

Abstract

Previous papers from this reseach group have suggested that we group object boundaries not by tracking around the boundary but by pairing boundary points across the object. Such pairings can be used to compute the medial axis as described by Blum. In this paper, the results of a fuzzy (non-binary) operator sensitive to object boundaries are combined to create a response function of three variables: two spatial and one of scale. The value of this response function at any spatial position and scale is the likelihood of that spatial position being on the medial axis where the scale corresponds to the width of the object at that spatial position. These boundary-sensitive operators are applied at every spatial position and at a number of scales, and the results of these operators are combined using a technique similar to the Hough transform to produce the desired response function.

Introduction

Pizer and colleagues have suggested that we observe object boundaries by pairing corresponding points on opposite sides of the boundary [1]. These pairings can then be used to compute the medial axis of the object [2].

One simple approach to this might involve applying edge-sensitive operators (each producing a characteristic function describing the “edginess” of the point), making some type of decision to turn this fuzzy characteristic function into a binary one, and then computing the medial axis of this binary structure in a straightforward manner. Unfortunately, this approach throws out information at early stages of processing, thus making it unavailable to later steps.

Another approach is to retain the fuzzy nature of the computation until the highest level possible. This could be done by using the fuzzy characteristic function for edges to compute a similar characteristic function for points on the medial axis. This approach produces an axis characteristic function that can be used by still higher order processes or by a decision process that extracts the optimal binary form of the axis. We have termed this axis characteristic function a “medialness response function” in the sense that it responds to medial properties.

This paper presents one such method for combining the outputs of the individual edge measurements into a medialness response function that captures both the axis position and the associated radius (scale) of the object at that position.

Why Use Edges at All?

A response function can also be created by applying a single operator centered at the potential axis point that reports its result at that position and scale. Such an operator must be sensitive to changes at the boundaries of the object. The response of the operator is, in a sense, the fit of the operator to a particular position in the object. An operator of the scale of "best fit" should produce a greater response than a slightly larger or slightly smaller operator. Similarly, an operator positioned exactly on the axis should produce a greater response than a neighboring off-axis operator. Such a best fit operator is termed an "axis-centered operator". An example of this is shown in Figure 1a. Compare this for a moment to the alternative of applying individual "edge-centered" operators and combining their results as shown in Figure 1b.

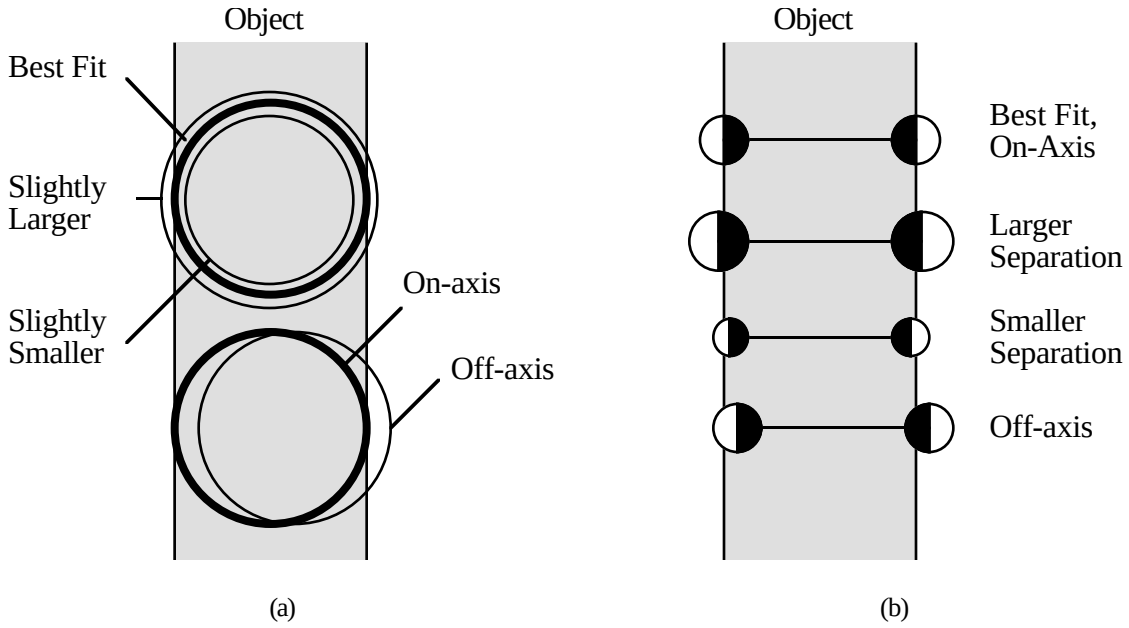

Figure 1. Axis-centered (a) vs. edge-centered (b) operators. In both cases, the best response is for operators of the correct scale that are centered on the axis. Response decreases for larger scales, smaller scales, and points off the axis.

Using an axis-centered operator is attractive in that it is the best fit to a region rather than to edges. It has several advantages and has been used for specific applications [1] [3] [4]. However, combinations of individual edge-centered operators have the following advantages:

- Axis-centered operators require integration over the entire width of the object. This makes it sensitive to internal variations within the object. It is also physiologically

improbable as a human visual model. Edge-centered operators can integrate information over smaller areas.

- Axis-centered operators are equally sensitive to changes on either side of the operator. Small differences in the strength or nature of the boundary on the two sides of the object can pull the response to one side or the other. Combinations of edge-centered operators are not necessarily sensitive to small differences between the individual operators—these differences can change the overall strength of the response, but do not change the relative spatial pattern of the response.
- Edge-centered operators can rectify the response of individual operators, producing pairings between boundary transitions of opposite polarity. For example, it can pair boundaries that are lighter than the background on one side and darker than the background on the other. This property cannot occur with simple linear operators.
- Edge-centered operators, since they operate on each edge individually, can combine edges of different natures to produce a medial response. That is, it can detect objects that are bounded by different types of boundaries at different parts of the object. These might include luminence boundaries, texture boundaries, line boundaries, or any other detectable type of edge.

There is also psychophysical evidence that indicates that the human visual system (for sufficiently large separations) individually localizes each target when measuring distances between scene targets. We will revisit this point in later discussion.

Edges cannot then be rejected in a segmentation model, but rather are a necessary part of scene perception. The key is to use edges in a proper way—as a fuzzy characteristic function rather than explicit edge detection.

Relating Edge Scale to Object Width

A key characteristic of the human visual system is that it uses scale information appropriate to the task performed [5]. The significance of small-scale fluctuations must be interpreted in the context of the overall scale of the task. This behavior is incorporated into a medial axis response function by using edge operators of a scale proportional to the radius associated with the medial axis.

It must be emphasized that this refers to the scale over which the edge or boundary is localized. The underlying information may be of smaller scale (higher frequency) [6].

Computing The Response Function

Let R(x,y;r) be a measure of how “medial” a particular spatial position (x,y) is with respect to a certain object half-width (r). By Blum’s definition of the medial axis, the circle of radius r

centered at (x,y) must be at least doubly tangent to the object boundary. So, R(x,y;r) is related to the amount of edge response for points on this circle. Furthermore, since this circle is tangent to the object boundary, R(x,y;r) is related to the *directional* response for points on this circle in the direction from the point to the center of the circle. This is illustrated in Figure 2.

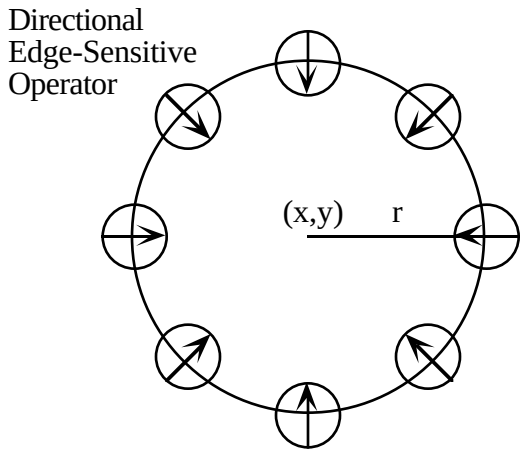

Figure 2. Individual edge-sensitive directional operators combine to produce a medial response.

In addition, the edge response must be measured at a scale appropriate to the size of the object. Letting the scale of the edge operator (σ) be proportional to the radius (r) gives

$$\sigma = kr \tag{1}$$

Where k is the proportionality constant, typically 2. In practice, this is approximated by

$$\sigma = kr + c \tag{2}$$

with a small value of c. This allows a slightly larger scale to be used at extremely small object widths where the inner scale of the image itself limits the size of meaningful operators. The effect of this constant becomes negligible for larger values of r so that σ/r approaches a constant value.

Adding the directional response along the points in the circle in Figure 2 gives the medial response for the point (x,y) and radius r. This response is

$$R(x,y;r) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \left| D_{\theta}(x',y',\sigma) \right| W((x,y),(x',y'))\, dx'\, dy \tag{3}$$

where

$\sigma = kr + c$ (see Eq. 2),

θ is $\tan^{-1}\left(\frac{y - y'}{x - x'}\right)$, the direction between (x,y) and (x',y').

$D_\theta(x',y',\sigma)$ is the directional response at position (x',y'), scale σ, and direction θ, and

$W((x,y),(x',y'))$ is a weighting function of the distance between (x,y) and (x',y') that has value one at distance r from (x',y') and positive value less than one elsewhere, decreasing as the distance between (x'y') and (x,y) increases.

Equation 3 is a continuous form of the response equation. For discrete spatial sampling, this is approximated by

$$R(x,y;r) = \frac{1}{r} \sum_{x'} \sum_{y'} |D_\theta(x',y',\sigma)| W((x,y),(x',y')) \qquad (4)$$

Equation 4 may be interpreted as the summation of the directional edge response of the appropriate scale at all points approximately distance r from (x,y). Since the number of such points is approximately proportional to r, the summation is normalized by a factor of 1/r.

If the directional edge response is linearly separable (e.g., luminance gradient), then it can be calculated by:

$$D_\theta(x',y',\sigma) = G(x',y',\sigma) \cos(\theta - T(x',y',\sigma)), \qquad (5)$$

where

D_x and D_y are the directional response in the x and y directions respectively (for luminance, these are the components of the gradient at scale σ),

$G(x',y',\sigma) = \sqrt{D_x^2 + D_y^2}$, the magnitude of the (gradient) vector whose components are D_x and D_y, and

$T(x',y',\sigma) = \tan^{-1}(D_y/D_x)$, the direction of the (gradient) vector whose components are D_x and D_y.

This can be used to reduce the computation required for such linearly separable directional edge response, since the edge must only be sampled in two directions rather than all possible orientations.

For edge-sensitive operators that are not linearly separable, three approaches are possible. The applicability of these approaches depends on the type of operator. One approach is to compute a single optimal edge direction and choose an appropriate tuning curve for the angular difference from this direction. The case of linear separability is a subset of this approach using a cosine tuning curve. Another approach would be to sample the directional response in each of a number of directions and interpolate for intermediate angles. The final approach would be to actually compute the directional response in exactly the direction required.

Contribution of an Individual Spatial Point

Equation 5 can be used to describe the influence of edge operators at an individual spatial point. Consider a spatial position (x,y). At this position directional edge-sensitive operators are applied in the x and y directions at some scale σ. From this, the edge magnitude $G(x,y,\sigma)$ and orientation $T(x,y,\sigma)$ are computed. Using Equation 2, an appropriate radius r is determined. That is, this edge response only affects the axis response for an axis of radius r. The weighting function W also limits the spatial influence to points at or near r. According to Equation 5, the response is also weighted by the strength of the edge magnitude. Equation 5 also says that the effect is weighted by the cosine of the angle between the affected point and the optimal direction $T(x,y,\sigma)$. That is, the affected area is circular, with strongest effect in the direction across the edge and zero effect along the edge direction itself. All of these factors combine to limit the extent of the influence of a single edge response as shown in Figure 3.

The edge response at (x,y) is not sampled only at a single scale s, but at a series of scales $\sigma_1, \sigma_2, \ldots, \sigma_n$, each with their corresponding $r_1, r_2, \ldots, r_n$. Each edge response sampled at (x,y) and scale σ_i produces an effect in $R(x,y;r_i)$ that is a circle of approximate radius r_i centered at potential axis position (x,y). The total influence of all scale samples at a single spatial position is a cone in R(x,y;r) as shown in Figure 4. The influences from multiple edge points interact with each other in an additive fashion to produce the axis response function R(x,y;r).

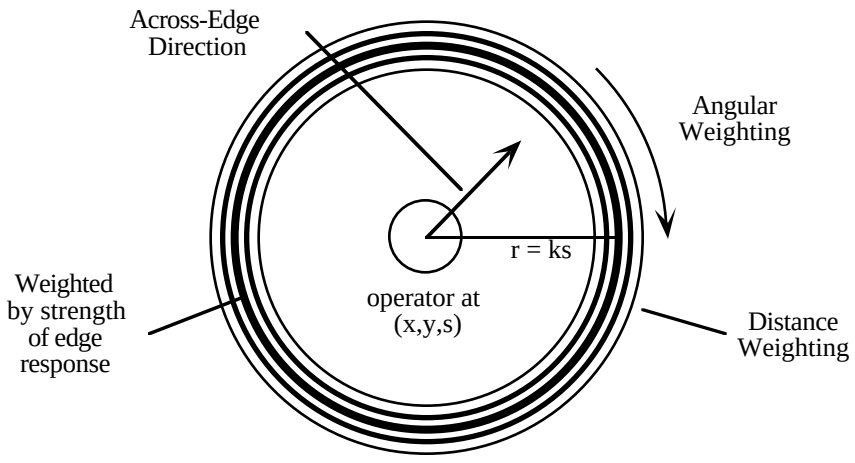

Figure 3. Influence of a single edge response at a single scale. Combination of response strength weighting, angular weighting, and distance weighing.

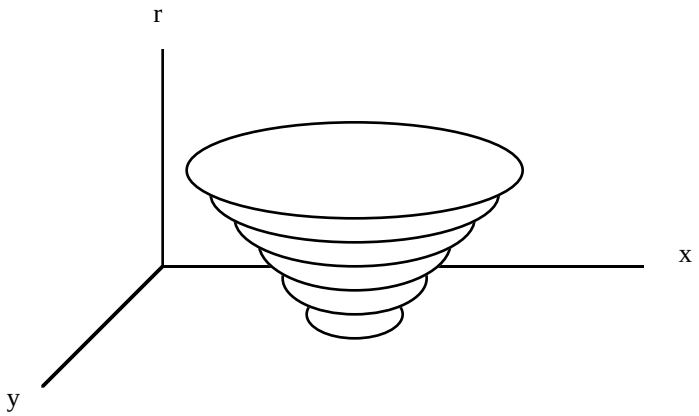

Figure 4. Influence in R(x,y;r) space of the responses to a single edge position at multiple scales.

Algorithm

This pattern of influence and the cumulative effect of individual spatial points is the basis for an algorithm similar to the Hough transform. In the case of linearly separable directional operators,

1. Apply edge-sensitive operators at every position in the image at a number of scales;
2. For each position and scale, compute the magnitude and direction of the edge; and
3. For each position and scale, compute the contribution of the response at that position and scale to the axis response R(x,y;r) using Eq. 4 and Eq. 5 and add that response to an accumulator for R(x,y;r).

The result of this algorithm is a transformation from a $(x,y;\sigma)$ edge-response scale-space to a (x,y;r) axis-response scale-space. Note that this response gives information not only on axis position (x,y), but also object width (r).

The multiscale medial axis can be determined as a "scale-space ridge" of the medialness response according to the methods described in Pizer, et. al. [1]

Results

The following results were computed using first derivative Gaussian gradient operators. The constant of proportionality between the scale of the Gaussian (σ) and the radius of influence (r) was two. Scale-space was sampled in a self-similar fashion [7] such that

$$s_i = s_0 k^i \qquad (6)$$

where, for these experiments $s_0 = 1.0$ and $k = 1.1$. Enough scales were sampled to provide for a corresponding radius as large as the size of the image.

These original images are all 128 x 128 8-bit images. Each is a simple foreground shape with a non-branching axis. Each object is constant intensity with anti-aliased edges. The shapes have been given descriptive names for simplicity of reference.

Tube

The object with the simplest medial axis (other than the trivial case of a circle) is one with a straight-line axis and constant radius. Such a shape is seen in Figure 5 and is called a “tube”. First-derivative Gaussian operators were applied in the horizontal and vertical directions for 32 scales. These edge responses were combined using the previously described algorithm to produce the R(x,y;r) medial axis response shown in Figure 6. At small scales, the only effect is near the boundaries. As the scale and corresponding radius increases, the influence of the edge moves farther from the boundary. The influence of the edge can be thought of as a wave propagating outward from the edge and upward in scale space. Even at scales less than the radius of the object, the wave starts to compress on itself at the rounded top and bottom of the tube. At successively larger scales, this wave compresses more and more until it ultimately forms the single high response at the end of the axis. At the same time, the two sides of the boundary are moving in towards the center until they meet to form a high response signifying to the medial axis.

Each of the subimages in Figure 6 shows the response at a fixed r. Similarly, we can take cross-sections of the response for a fixed value of x or y. Figure 7a shows the response at all scales along the horizontal midline of the image. Observe the influence of each edge spreading outward and upward through scale space. Figure 7b shows the response at all scales along the vertical midline of the image. Since the medial axis of the object itself lies along this vertical midline, Figure 7b can also be interpreted as a plot of radius (ordinate) vs. arc length (abscissa) along the axis. Since the radius increases exponentially but is displayed linearly, the ordinate should actually be interpreted as log r. Compare this to the graph of log r vs. arc length shown in Figure 8.

Teardrop

Different shapes can have the same axis but a different associated set of radius values. It is critical that R(x,y;r) capture not only the position of the axis, but the associated radius as well. The shape in Figure 9 has the same straight-line axis as the tube but a linearly increasing radius function. Such a shape is called a “teardrop”.

The calculated response function R(x,y;r) for this teardrop is shown in Figure 10. As with the tube, the influence of the edges propagates outward from the edges as the radius parameter increases. Since the teardrop has a zero radius (a sharp corner or point) at the tip, the two edges add to produce a point of high response there immediately—this is shown by the bright spot at the

tip of the teardrop at the smallest scales. As the radius parameter increases, this high response moves down the axis until it eventually reaches the point where the axis terminates. As with the tube, the rounded end of the teardrop has been compressing and increasing in response until it meets with the point moving down the axis to form the high response at the end of the axis.

Again, as with the tube, cross-sections of the response function can be taken along the horizontal and vertical midlines of the image (i.e., across and along the axis) as shown in Figure 11. This time, the vertical cross-section shows the increasing radius of the teardrop. Compare this to the actual radius as shown in Figure 12.

Dumbbell

The shape in Figure 13 also has the same straight-line axis as the tube and the teardrop, but has a parabolic radius function. Such a shape is called a “dumbbell”.

The response function for the dumbbell is shown in Figure 14. Like the response for the teardrop starting at the tip and moving to the other end, observe the high response in the dumbbell starting at the pinched middle and moving outward towards the ends as the radius increases. Also observe the now-familiar compression of the rounded ends ultimately forming the very high response at the end points of the axis.

Again taking cross-sections as shown in Figure 15, the cross-section along the axis shows the parabolic shape of the radius. Compare this to the actual radius shown in Figure 16.

Paisley

Different shapes can also have the same radius function but a different axis. Such shapes can be thought of as produced by bending another shape. The shape in Figure 17 has the same linearly increasing radius function as the teardrop, but a curved axis. Such a shape is called “paisley”.

The response function for the paisley pattern is similar to that for the teardrop and is shown in Figure 18. Again, the high response starts at the tip and moves along the axis to the other end as the radius increases. Notice that this time the point of high response moves not in straight line down the axis, but curves as the axis itself does. Thus, the response function R(x,y;r) accurately captures changes in position as well as radius.

Noise

To test the behavior of the algorithm in the presence of noise, zero-mean Gaussian noise was added to the image of the tube as shown in Figure 19. In the first test, the standard deviation of the noise in each pixel was half the intensity difference between the foreground and background. Observe the small deviations in the boundary caused by the noise. The computed response function

for this image is shown in Figure 20 with cross-sections in Figure 21. Notice how at smaller scales, the gradients caused by the noise have significant influence, but at larger scales more appropriate for the width of the object, the response function is little changed from that of the noise-free version of the image.

In the second test, the standard deviation of the noise in each pixel was equal to the intensity difference between the foreground and background. The results are identical to those with the lesser noise, but the effect of the noise extends to higher scales. In this case, the noise is just strong enough to have an effect at the scale of the object.

Figure 5. A tube. Simple medial axis shape with straight, non-branching axis and constant radius.

Figure 6. Response function R(x,y;r) for a tube. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 7. Cross-sections of the response function R(x,y;r) shown in Figure 6 along the horizontal midline (top) and vertical midline (bottom).

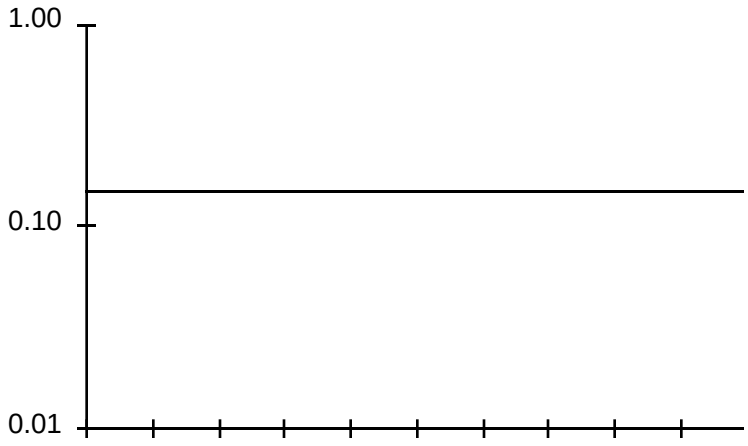

Figure 8. Log radius vs. arc length for the tube.

Figure 9. A teardrop. A medial axis shape with straight-line, non-branching axis and linearly increasing radius (from top to bottom).

Figure 10. Response function R(x,y;r) for a teardrop. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 11. Cross-sections of the response function R(x,y;r) shown in Figure 10 along the horizontal midline (top) and vertical midline (bottom).

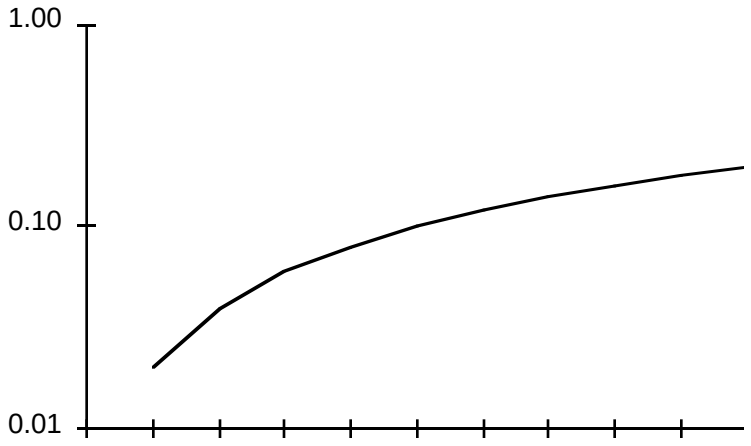

Figure 12. Log radius vs. arc length for the teardrop.

Figure 13. A dumbbell. A medial axis shape with straight-line, non-branching axis and parabolic radius.

Figure 14. Response function R(x,y;r) for a dumbbell. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 15. Cross-sections of the response function R(x,y;r) shown in Figure 14 along the horizontal midline (top) and vertical midline (bottom).

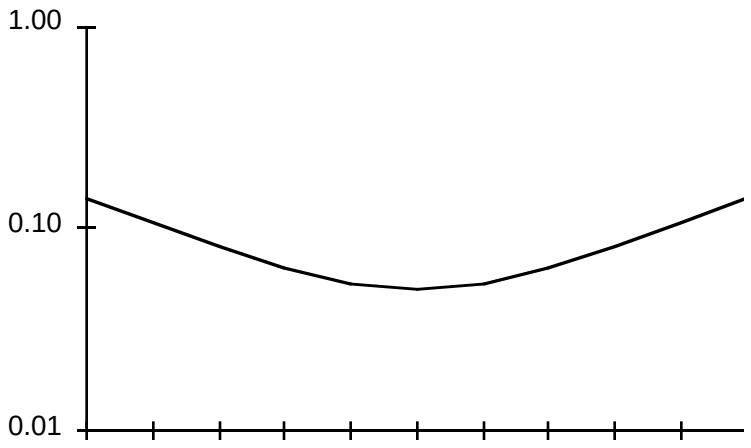

Figure 16. Log radius vs. arc length for the dumbbell.

Figure 17. Paisley. A medial axis shape with curved, non-branching axis and linearly increasing radius. This can be considered a curved or bent teardrop.

Figure 18. Response function R(x,y;r) for paisley. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 19. A tube with noise added. The standard deviation of the noise is half the difference between the foreground and the background.

Figure 20. Response function R(x,y;r) for a tube in the presence of noise. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 21. Cross-sections of the response function R(x,y;r) shown in Figure 20 along the horizontal midline (top) and vertical midline (bottom).

Figure 22. Another tube with noise added. The standard deviation of the noise is equal to the difference between the foreground and the background.

Figure 23. Response function R(x,y;r) for a tube in the presence of noise. Each subimage is the response at each spatial position for increasing radius values. The upper right subimage is for the smallest radius with radius increasing from left to right and top to bottom.

Figure 24. Cross-sections of the response function R(x,y;r) shown in Figure 23 along the horizontal midline (top) and vertical midline (bottom).

Discussion

It is apparent from the results in Figures 6–18 that this axis response function does capture the positional as well as the scale parameters of simple shapes. This representation should allow determination of the actual axis and radius functions themselves.

Figures 19-24 show that by using gradient operators proportional to the width of the object, small fluctuations and deviations in the boundary are smoothed away. This has been one of the principal criticisms of the medial axis since under normal conditions even the tiniest fluctuation in the boundary can produce long branching arms extending from the principal axis. However, by adding an appropriate sense of scale based on the width of the principal axis, such minor detail is ignored when judging the overall shape of the object. Such detail should, however, be captured at some lower scale. This is the subject for future investigations.

Another observation made during the presentation of the results involved the collapsing of rounded object ends to form extremely high response axis endpoints. The reason for this comes from the definition of the medial axis. Since the medial axis is the locus of all maximal circles, endpoints of the medial axis naturally correspond to locally circular parts of the boundary. In other words, the maximal circles at normal points on the axis are bitangent to the object boundary and therefore receive input from more than one edge point. However, maximal circles at the end point of a medial axis can be tangent to the boundary not only at two points, but along an entire arc. Each of the points along this arc contributes to the axis response at the end point. The same should in theory be true of points where the medial axis branches—there the maximal circles are tritangent or more depending on the degree of the branching and thus also produce an axis response that is stronger than the rest of the axis. This property is a useful one in that the points where the response is highest are the point that most define the axis: end points and branch points.

Future Work

As mentioned, detection of details at appropriately small scales is a subject for future investigation. It is also suspected that this detail axis may not be connected in scale space with the principal axis. This relation between large and small scale information is still unknown.

The principal item for future research is the extraction of the axis itself from the response function presented here. Research is proceeding using geometric as well as connectionist means. A connectionist algorithm would make use of excitatory connections to allow points on the axis to stimulate other axis points, thus filling in points that may have been lost due to noise, occlusion, or other interference with the boundary. This filling in could be the basis for producing illusory contours based on figural, rather than edge, interpolation.

One problem with this and any other Hough transform method is that edge points are allowed to cast votes in the accumulator indiscriminately. Often, the interpretation of one feature of the image is lost in the noise produced by other, stronger features. It is possible to make multiple passes through the transformation process, limiting the voting of edge points at subsequent stages based on the results of previous transformations [8]. This allows voting in the accumulator to be more selective based on the previous voting pattern. Such a mechanism could be used here, where at each subsequent iteration each input edge point downgraded its votes for axis points where there was little agreement with other input points. Such downgrading should produce a sharpening of the response function making interpretation more accurate and in some cases, making features detectable where they otherwise might not have been. Such sharpening should also allow the system to settle into one of several possible interpretations in the case of bistable images.

It might even be possible to allow interaction between this sharpening process and the previously described axis-excitation process. Such interaction could be thought of as simultaneous neural processes. One fills in gaps in the axis, while the other allows those filled-in gaps to cause their inputs to be more selective, thus perhaps allowing a cleaner response in another part of the image, leading to further refining of another axis, and so on.

Another area for future work involves the inputs used for the edge-sensitivity measurements. These not only could be luminance edges as implemented here, but could be bar detectors (for line-drawn boundaries), texture difference operators (for textural boundaries), or any other operator capable of producing a magnitude and orientation for the boundary. Some, such as the texture difference operators, could be composite operators. That is, operators that consist of one operator applied to the outputs of other, perhaps smaller scale, operators. Such composite operators are consistent with psychophysical evidence that seems to indicate that we have both detection and localization mechanisms, operating on different scales. For example, high-frequency textures with zero low-frequency content could be detected by high-frequency mechanisms and used as input into a larger scale localization operator. It might even be possible to use the axis response itself as input into calculating the axis of larger, composite objects.

One difficulty with the edge-centered operators is the relatively large separation of edges required. If a radius to scale constant of two is used, then the smallest measurable object width is four times the scale of the smallest edge operator. However, at such a small separation, it becomes reasonable to use axis-centered operators. There is evidence that the human visual system might to a similar switch between mechanisms at large and small separations [5]. A hybrid large-scale edge-centered, small-scale axis-centered response function might be interesting.

Conclusion

The edge-centered response function R(x,y;r) is able to capture both the positional and scale information of the medial axis. A significant characteristic of this axis response function is that it is computed entirely without explicit (binary) edge determination using scales appropriate for the width of the object. Such a system should be both more accurate and more robust than calculations based on explicit edge determination, since it allows more information to be retained at higher levels of processing. There is much more work to be done, but these initial results are encouraging.

References

1. Pizer, S. M., Coggins, J. M, Burbeck, C. A, Morse, B. S., and Fritsch, D., "Image Object Description Without Explicit Edge-finding", Technical Report TR91-049, Department of Computer Science, University of North Carolina at Chapel Hill. Submitted to the *Second European Conference on Computer Vision* (ECCV), 1992.
2. Blum, H. and Nagel, R. N., "Shape Description Using Weighted Symmetric Axis Features", *Pattern Recognition*, Vol. 10 (1978), pp.167-180.
3. Crowley, J. L. and Parker, A. C., "A Representation fr Shape Based on Peaks and Ridges in the Difference of Low-Pass Transform." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 6, No., 2, March 1984, pp.156-170.
4. Fritsch, D. S., Coggins, J. M., and Pizer, S. M. "A Multiscale Medial Description of Greyscale Image Structure", Advances in Intelligent Robotic Systems, Intelligent Robots and Computer Vision X: Algorithms and Techniques, SPIE, 1991.
5. Levi, D. M. and Westheimer, G., "Spatial-interval Discrimination in the Human Fovea: What Delimits the Interval?", *J. Opt. Soc. Am. A*, Vol. 4, No. 7, July 1987, pp.1304-1313.
6, Burbeck, C. A., "Position and Spatial Frequency in Large-scale Localization Judgements", *Vision Research*, Vol. 27, No. 3 (1987) pp. 417-429.
7. ter Haar Romeny, B. M. and Florack, L., "A Multiscale Geometric Model of Human Vision", in Perception of Visual Information (Hendee, B. and Wells, P. N. T. Editors.), Spinger Verlag—Berlin, 1991.
8. Gerig, G., "Linking Image-Space and Accumulator-Space: A New Approach for Object Recognition", in Proceedings First International Conference on Computer Vision (ICCV'87), pp.112-122.

Acknowledgements

The authors would like to acknowledge the support and contributions of James Coggins, Dan Fritsch, and Jonathan Marshall.

This research was supported by the UNC Graduate School, by NIH grant #P01 CA47982, and in part by Air Force grant #AFOSR-91-0058.