Journal of Mathematical Imaging and Vision, 4, 303–313 (1994).

Object Shape before Boundary Shape: Scale-Space Medial Axes*

STEPHEN M. PIZER, CHRISTINA A. BURBECK, JAMES M. COGGINS, DANIEL S. FRITSCH, AND BRYAN S. MORSE
Medical Image Display Research Group, University of North Carolina, Chapel Hill, NC 27599-3175

Abstract. Representing object shape in two or three dimensions has typically involved the description of the object boundary. This paper proposes a means for characterizing object structure and shape that avoids the need to find an explicit boundary. Rather, it operates directly from the image-intensity distribution in the object and its background, using operators that do indeed respond to "boundariness." It produces a sort of medial-axis description that recognizes that both axis location and object width must be defined according to a tolerance proportional to the object width. This generalized axis is called the *multiscale medial axis*[1] because it is defined as a curve or set of curves in scale space. It has all of the advantages of the traditional medial axis: representation of protrusions and indentations in the object, decomposition of object-curvature and object-width properties, identification of visually opposite points of the object, incorporation of size constancy and orientation independence, and association of boundary-shape properties with medial locations. It also has significant new advantages: it does not require a predetermination of exactly what locations are included in the object, it provides gross descriptions that are stable against image detail, and it can be used to identify subobjects and regions of boundary detail and to characterize their shape properties.

Key words. shape description, object definition, multiscale method, medial axis, image analysis

1 Boundaries versus Medial Representations

The dominant train of thought in object shape measurement is based on boundary description. Thus for 2-D objects properties of the object edge, such as curvature, have been described, and for 3-D objects properties of the object surface, such as the loci of parabolic curves, flecnodal curves, gutterpoints, and ruffles [14], have received special attention. The difficulty of this approach is twofold. First, from the point of view of physics, for an object in an image there exists no edge locus without a tolerance since the object can exist only through imaging or visual measurements that have an associated spatial scale, and thus spatial tolerance [14], and the spatial scale that is appropriate for boundary definition is unclear. Second, shape involves certain global properties, which are not readily built into the process of describing boundaries. An important global property is that of involution, the relation between opposite points on two sides of an object (see figure 1 for examples).

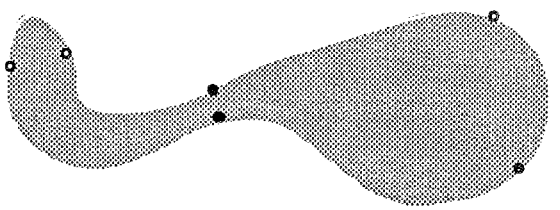

Fig. 1. Involutes: visually related opposite points on an object.

Such global shape aspects can be captured more directly by focusing on the object middle-and-width combination that arises from pairing opposite object edges [2]. Blum proposed to do this by representing the object in terms of a

*This research was partially supported by National Institute of Health grant P01 CA47982, National Aeronautics and Space Administration contract NAS5-30428, U.S. Air Force Office of Scientific Research grant 91-00-58, National Library of Medicine grant LM05508-01, National Library of Medicine/National Cancer Institute medical informatics research training grant 5 T15 LM07071, and a sabbatical grant from the Netherlands Organization for Scientific Research. It was reported orally at the NATO Advanced Research Workshop on Shape in Picture: The Mathematical Description of Shape in Greylevel Images, Driebergen, The Netherlands, in 1992.

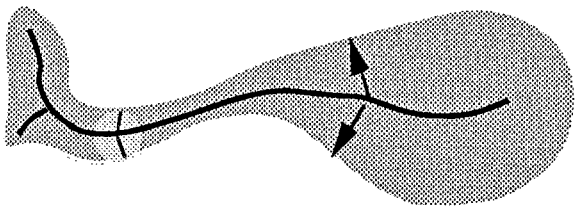

Fig. 2. Middle and width of an object and a disk defined from them.

medial axis or skeleton running down the middle of the object, together with a width value at each point on the medial axis. His axis is defined such that for each axis point a disk centered at that point and with radius equal to the width value there is tangent to the boundary at two or more boundary points and is entirely within the object (figure 2). The endpoints of these central axes correspond to corners and other object boundary locations of locally maximal curvature [15], [16], the perceptual importance of which has long been known. It has also been noted [11], [12] that subjective edge perceptions derive especially strongly from high-curvature boundary points, such as line end and corners.

The width values $w(s)$ of the middle/width representation carry straightforward access to the angle of the object boundary at each of the corresponding boundary points, relative to the axis direction at any axis point specified by arc length s: $\theta = \cos^{-1}(dw/ds)$ [3]. Moreover, the curvature of the axis and of the boundary pair relative to the axis is also straightforwardly accessible. At axis endpoints the radii perpendicular to the boundary converge to a single boundary point, which is the visually important vertex of a protrusion, i.e., a relative maximum of boundary curvature. Axis branch points correspond to indentations in the object. Thus the middle/width representation incorporates major aspects of shape.

Blum also suggested a more general "global" form of the medial axis representation in which the multiply tangent disks need not be completely inside the object. Global axis sections for which the disks overlap the object's background select boundary indentations and symmetries of larger width than the object, for example, the symmetry of the shorter sides of a rectangle.

The difficulty with Blum's definition is that whereas it tackles the problem of global shape, it still requires an object boundary that is defined with zero tolerance. No method that requires such a boundary can be expected to be adequately insensitive to small-scale image properties, and, indeed, Blum's method has been heavily criticized for this sensitivity.

2 Multiscale Geometry Detectors

Many investigators have suggested that notions of shape must be based on measurements in scale space, i.e., by sets of operators that sense a regional rather than curvilinear (e.g., edge or medial axis) property, with each operator sensing the same property but at different spatial scales. Among the operator kernels suggested have been derivatives of Gaussians [13], [19], differences of Gaussians [5], [26], Gabor functions [6], [23], Wigner operators [24], and wavelets [17], [18]. A persuasive argument for the form of operators, by ter Haar Romeny et al. [10], is that the system must be invariant to translation, rotation, and size change and that this implies multiscale operators h with kernels that are solutions to the diffusion equation $\nabla \bullet [c(\mathbf{x}; t)\nabla h(\mathbf{x}; t)] = (\partial/\partial t)h(\mathbf{x}; t)$, where t is half the square of the spatial scale σ, $\mathbf{x}$ is a spatial location in $\Re^2$, and c is a conductance function that can vary in space and scale. Linear combinations of derivatives of a Gaussian with standard deviation σ satisfy this equation for $c = 1$.

These operators, or combinations of them, can be thought of as giving the degree to which a point in scale space $\mathbf{x}$, σ has the properties expected of a particular geometric feature. For example, we will say that *boundariness* is the degree to which the point behaves like a boundary and *cornerness* is the degree to which the point behaves like a corner. Similarly, we will say that *medialness* is the degree to which the point behaves like the middle of an object.

Boundariness at a particular location $\mathbf{x}$ and scale σ has typically been associated with variations in luminance about that location, i.e., with combinations of first or second partial

derivatives in some direction **u** of the intensity function after convolution with a Gaussian with standard deviation σ (see [4], Sobel's method in [7], and [25]). However, there are many other possible cues to boundariness. Among them are measures of *endness* such as the corner detector of Blom [1], measures responding to an outline surrounding an object, measures of texture change, surface slant (giving depth change), and measures of velocity change. Each of these boundariness measures $B(\mathbf{x}, \sigma, \mathbf{u})$ is a function of position **x**, scale σ, and direction given by a unit vector **u**; each gives the degree to which this point in scale space behaves like a boundary with normal direction **u**.

An edge with tolerance proportional to σ may be taken to be a ridge in $\max_u B(\mathbf{x}, \sigma, \mathbf{u})$. We define a ridge of a function $f(\mathbf{x})$ to be the locus of positions with the following property. Let $\mathbf{w} = \nabla f(\mathbf{x})/|\nabla f(\mathbf{x})|$ be the orientation of the gradient of f at **x**. Let **v** be a unit vector orthogonal to **w**, i.e., tangent to the level curve of f through **x**. Then **x** is a ridge point of f if the rate of change of the gradient orientation in the **v** direction $D_{\mathbf{v}}\mathbf{w}$ has a relative maximum for a step in the **v** direction. This is a place where a level curve of f has maximal curvature. Another possible, but nonequivalent, definition is the locus of maxima of $f(\mathbf{x})$ in the eigendirection with the largest-magnitude negative eigenvalues of the Hessian. Unlike alternative definitions, these two definitions have all of the following properties: they are local, they are invariant to rotation, translation, and affine transformations of intensity, they do not, in fact, depend on the global shape of level curves, and they do not require intensity to be commensurate with spatial distance. The definitions generalize to 3-D.

3 Medialness

Collectively, the preceding ideas have led us to the development of a new model for visual region formation and description of object shape that accords with results from visual psychophysics and neurophysiology, as discussed in [22]. It is based on the idea that just as explicit boundaries (if they are ever needed) must be derived from boundariness in scale space, so middles and widths must be derived from a scale-space measure that we call "'medialness." Medialness $M(\mathbf{x}_A, \sigma_A)$ is the degree to which a point in scale space $\mathbf{x}_A$, σ_A has the property of being an object middle at a specified width.

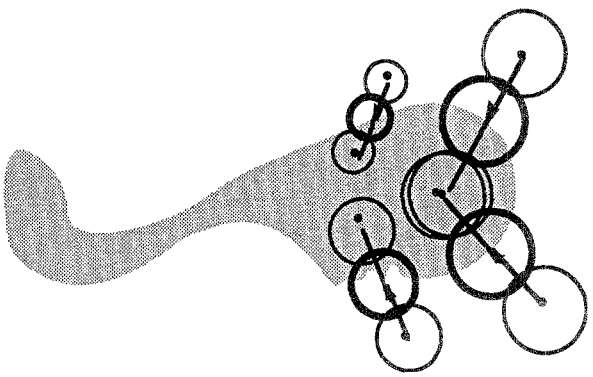

Fig. 3. Boundariness responses at the positions of the arrowheads, in an orientation indicated by the arrows, and at scales indicated by the sorrounding heavier circles contribute to medialness around the points indicated by the bullets and at scales indicated by the lighter circles around the bullets. Note that boundariness kernels at any point in space exist for all orientations, including the ones that are nonorthogonal to the edge, and at all scales.

Medialness at $\mathbf{x}_A$, σ_A must be derived from boundariness at various $\mathbf{x}_B$, σ_B, so the tolerance of loci derived from medialness will be proportional to the tolerance (scale) of the boundariness values that contribute to it. All else follows from this property of human vision: *the tolerance for the width of an object and for its middle location must be proportional to the object width there.* In fact, this property that the scale for object-middle measurement is proportional to object width $|\mathbf{x}_A - \mathbf{x}_B|$ allows the medialness to separate information about object features at different scales and to be invariant to magnification change. Stated mathematically,

(a) $M(\mathbf{x}_A, \sigma_A)$ must be derived from $B(\mathbf{x}_B, \sigma_B, \mathbf{u}_B)$ at various $\mathbf{x}_B$ but with the scale σ_B satisfying $\sigma_A = c\sigma_B$ for some constant of proportionality c, and

(b) the separation $|\mathbf{x}_A - \mathbf{x}_B|$ between the boundariness position and the medialness position must satisfy $|\mathbf{x}_A - \mathbf{x}_B| = k\sigma_B$ for some constant of proportionality k (see figure 3).

In addition, for $B(\mathbf{x}_B, \sigma, \mathbf{u}_B)$ to contribute to $M(\mathbf{x}_A, \sigma_A)$, $\mathbf{u}_B$ must be approximately in the direction $\mathbf{x}_A - \mathbf{x}_B$. Thus

$$
\begin{aligned}
&M(\mathbf{x}_A, \sigma_A) \\
&\quad = \int\int\int B(\mathbf{x}_B, \sigma_B, \mathbf{u}_B) \\
&\qquad \times \Bigg[W\Bigg(\frac{\mathbf{x}_A - (\mathbf{x}_B + k\sigma_B \mathbf{u}_B)}{\sigma_B}, \frac{\sigma_A - c\sigma_B}{\sigma_B}, \\
&\qquad\qquad \frac{\mathbf{x}_A - \mathbf{x}_B}{|\mathbf{x}_A - \mathbf{x}_B|} - \mathbf{u}_B \Bigg) \\
&\qquad\quad + W\Bigg(\frac{\mathbf{x}_A - (\mathbf{x}_B - k\sigma_B \mathbf{u}_B)}{\sigma_B}, \frac{\sigma_A - c\sigma_B}{\sigma_B}, \\
&\qquad\qquad \frac{\mathbf{x}_A - \mathbf{x}_B}{|\mathbf{x}_A - \mathbf{x}_B|} + \mathbf{u}_B \Bigg) \Bigg] d\mathbf{x}_B d\sigma_B d\mathbf{u}_B .
\end{aligned}
$$

The integration over $\mathbf{x}_B$ and σ_B is over all of scale space, and the integration over $\mathbf{u}_B$ is over the semicircle of orientations. W is an effect-smearing function in position, scale, and boundary orientation, such as a zero-mean Gaussian in its three variables. It allows a given boundariness to affect medialness at points in scale space near but not exactly equal to the target position and scale, $\mathbf{x}_A, \sigma_A$; see figure 4 for an example. In figure 4 $B(\mathbf{x}, \sigma, \mathbf{u}) = |\nabla(G(\mathbf{x}; \sigma) * L(\mathbf{x})) \bullet \mathbf{u}|$, where $G(\mathbf{x}; \sigma)$ is a zero-mean Gaussian with standard deviation σ and $L(\mathbf{x})$ is the image; because $W(\Delta\mathbf{x}, \Delta\sigma, \Delta\mathbf{u}) = \max(1 - |\Delta\mathbf{x}|, 0)$, linear interpolation between adjacent pixels is done.

The effect is that for a point $\mathbf{x}_A$ inside the object and near a boundary, the medialness $M(\mathbf{x}_A, \sigma_A)$ as a function of σ_B with $\sigma_A = c\sigma_B$ will have the two-humped shape shown for point E in figure 5(a). At small scales σ_B the medialness will be low because there is no boundariness to be found at small scales at positions at distance $k\sigma_B$ from $\mathbf{x}_A$. As $k\sigma_B$ approaches the distance to the nearer edge, the boundariness originating from the edge (and oriented orthogonal to the edge) will have increased effect on the medialness. For a somewhat larger distance $k\sigma_B$, the correspondingly oriented boundariness will be smaller; moreover, where the edge is crossed at distance $k\sigma_B$ from $\mathbf{x}_A$, the boundariness oriented towards $\mathbf{x}_A$ will be low because that orientation will be far from orthogonal to the edge. The boundariness will remain small until $k\sigma_B$ approaches the distance to the far object edge, when the boundariness, and thus the medialness, will increase and then decrease as σ_B increases.

On the other hand, for positions x_A nearly equidistant from the two edges, there will be a single relative maximum of $M(x_A, c\sigma_B)$ with respect to σ_B, because there the two equidistant edges will both be contributing their boundariness at the same scale. Moreover, the medialness maximum will be higher at the middle than nearer the edge because of the combination of the boundariness contributions from the two edges. Figure 5(a) shows this behavior of the medialness-versus-scale curves as one moves from near the middle to near the boundary. Figure 5(b) shows how the scale at which the maximum occurs at a middle point increases linearly with the width of the object.

Medialness can also be computed by using kernels that respond to two equidistant boundaries simultaneously rather than from each boundary separately. An example of such a kernel is the normalized Laplacian of a Gaussian (Crowley and Parker [5] use a similar normalization on a difference of Gaussians). Details can be found in [8], [9].

4 The Multiscale Medial Axis

For a position in scale space $(\mathbf{x}, \sigma)$ to correspond to a middle point and width of an object, it must first be at an optimal scale—a scale maximizing medialness at that $\mathbf{x}$. That is, a variation in width (scale) must result in a decrease in medialness. Secondly, the medialness at optimal scale must spatially have the ridge property. We call the loci of points in scale space $(\mathbf{x}, \sigma)$ that have the preceding two properties the *multiscale medial axis* (MMA).[2] The $\mathbf{x}$ component of such a point specifies a point located on the middle of the object, and the σ component of such a point simultaneously gives (with appropriate constants of proportionality) the object-width property at $\mathbf{x}$ and the tolerances of both the location and width of the medial point.

Mathematically stated, $(\mathbf{x}, \sigma)$ is on the multi-

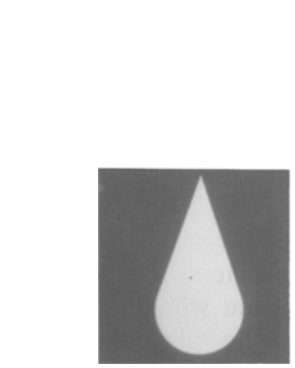

(a)

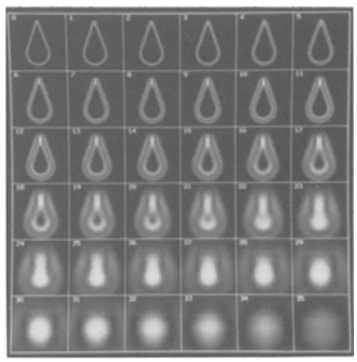

(b)

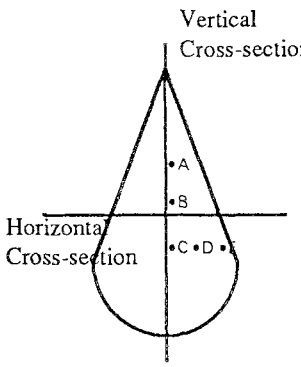

(c)

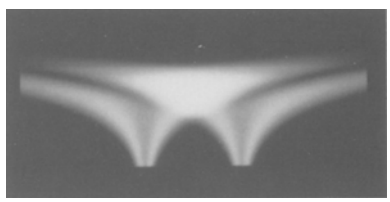

(d)

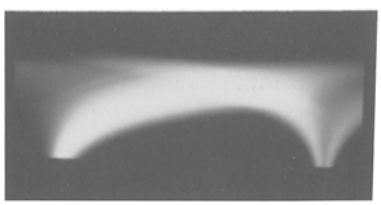

(e)

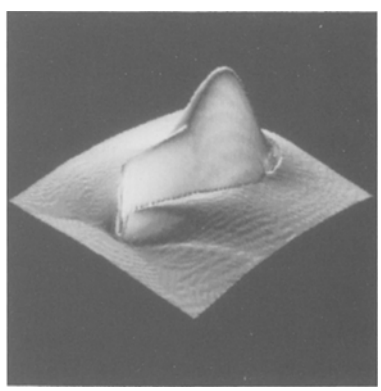

(f)

Fig. 4. (a) Image to be analyzed; (b) medialness versus scale σ_A (image number) and position $\mathbf{x}_A$; (c) cross sections and points relevant to (d)–(e) and figure 5; (d) and (e) medialness versus position along central cross sections through $\mathbf{x}_A$ (on the abscissa) and versus σ_A (on the ordinate) along (d) horizontal and (e) vertical cross sections of the image; (f) optimal scale medialness versus image space seen as a height. The results shown are for an object with a sharp boundary, but similar results are obtained for an object with a blurred boundary.

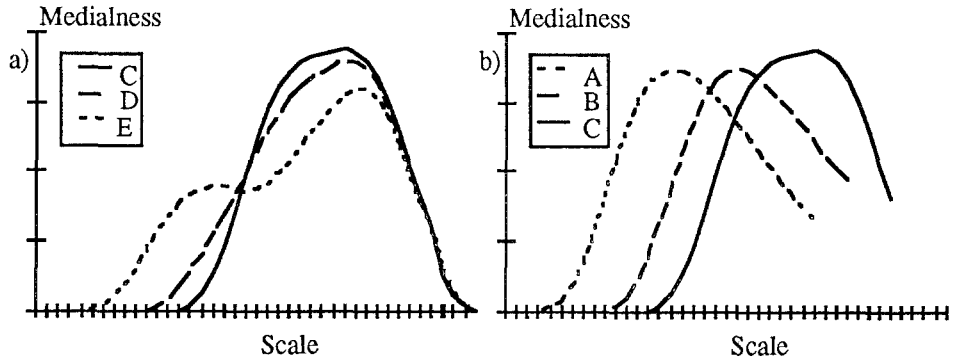

Fig. 5. Medialness at a point versus scale for points (a) across and (b) along the object middle [see figure 4(c)].

scale medial axis if the following hold:

(a) $M(\mathbf{x}, \sigma)$ has a relative maximum with respect to σ at $\mathbf{x}$ (σ is an optimal scale at $\mathbf{x}$). Let S be the set of $(\mathbf{x}, \sigma)$ such that $M(\mathbf{x}, \sigma)|_{\mathbf{x}}$ is such a relative maximum with respect to σ. Partition S into its connected subsets, S_i, $i = 1, 2, \ldots$. In each S_i there exists a connected region of image points $\mathbf{x}$ not necessarily covering the whole image space and there exists at most one scale σ associated with any such position $\mathbf{x}$. Figure 6 shows the loci S_i for a cross section across the narrow dimension of a 2-D object [cf. figure 4(d)] or a 1-D image of a bar.

(b) For each S_i, project $M(\mathbf{x}, \sigma)$ for $(\mathbf{x}, \sigma) \in S_i$ onto $\mathbf{x}$ to form the image or subimage $M^{\max_i}(\mathbf{x}) = M(\mathbf{x}, \sigma)$ for $(\mathbf{x}, \sigma) \in S_i$. The intensity for each of these images is an optimal-scale medialness at the corresponding image point. Then $(\mathbf{x}, \sigma)$ is in the multiscale medial axis if $\mathbf{x}$ is a ridge point in any such portion of $M^{\max_i}(\mathbf{x})$ for any i [see figure 4(f)].

The max-over-scale surfaces S_i are separated in scale space. As illustrated in figures 4(d) and 6, for points on the object from its right edge to some point near its middle there are two maximal scales, one of smaller scale (below in the graph) corresponding roughly to the distance to the right (near) edge and the other of larger scale (above in the graph) corresponding roughly to the distance to the left (far) edge. Similarly,

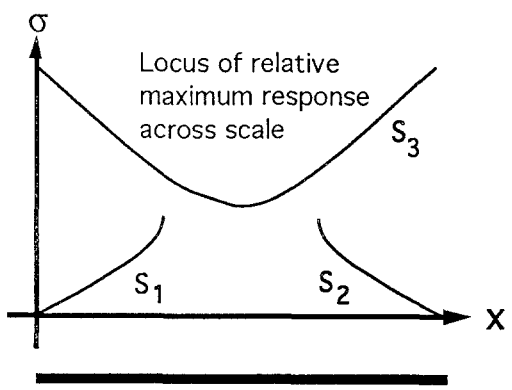

Fig. 6. Max-over-scale surfaces in 1-D scale space for the input bar shown below the graph. Note the continuity for second-nearest edge response.

for points from the object's left edge to some point near its middle there are two maximal scales, one of smaller scale (below in the graph) corresponding roughly to the distance to the left (near) edge and the other of larger scale (above in the graph) corresponding roughly to the distance to the right (far) edge. For object points between these two intervals there is a region of only a single medialness maximum with respect to scale; we have found experimentally that it is continuous with the far edge responses and that the ridge of optimal-scale medialness occurs on this branch (S_3 in figure 6). Morse theory guarantees that in generic situations the loss of one of the maxima of $M(\mathbf{x}, \sigma)|_{\mathbf{x}}$ will occur at an $(\mathbf{x}, \sigma)$ position that is separated from the other trace. Thus the result is three separated loci S_i.

The ridge of optimal scale medialness [see figure 4(f)] is a (normally unbranching but possibly branching) trace in scale space—the MMA. The image space (**x**) positions of the MMA form a medial axis for an object, and their scales specify its width and tolerances at each axis point.

As is shown in figure 7, the long component of the axis at a large scale describes the gross orientation and width properties of the object. It establishes the boundary of the object only to a tolerance proportional to the width of the object. Figure 7 also shows another component that is an axis of object symmetry at yet larger scale—comparable to Blum's global medial axis. The components at smaller scales correspond to smaller boundary detail or objects within the main object, either with boundary tolerances proportional to their widths. Even tighter tolerances on the boundaries can be obtained from smaller-scale operators responding to single boundaries within the boundary regions associated with the medial ridge.

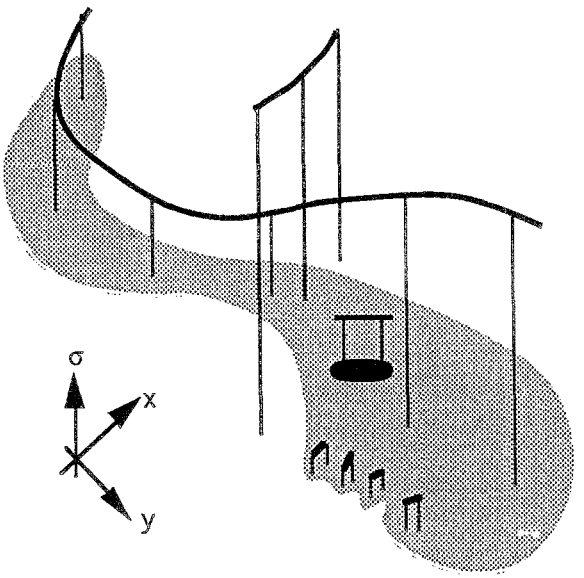

Fig. 7. Multiscale medial axes in scale space for an object, for detail on that object, for an object within that object, for a larger-scale symmetry of that object.

5 Boundariness–Medialness Interactions

The smoothness or wiggliness of boundaries has little effect at the scales proportional to the width of an object that determine its main MMA. Thus these shape properties cannot be reflected in the MMA itself but rather are reflected directly in boundariness properties. To describe the object shape fully, i.e., to show both the medial and boundary behavior and their relation, it is necessary to identify the medial location at which (and thus the object to which) a particular boundariness is bound and the scale at which the boundariness is relevant. That is, the boundariness properties must be put into correspondence with MMA locations.

The direction in scale space of the MMA at spatial location $\mathbf{x}_A$ and scale σ_A provides the information to associate boundary regions with that MMA point. The situation is as shown in figure 8. The direction of the projection onto image space of the MMA bisects the angle made by connecting the two corresponding boundariness regions to $\mathbf{x}_A$. If σ_A is scaled to object width, the angle ϕ_A between the direction perpendicular to the MMA direction and the directions of the vectors $\mathbf{u}_{A+}$ and $\mathbf{u}_{A-}$ connecting $\mathbf{x}_A$ to the two clusters of contributing boundariness is equal to the angle of the MMA with the image space plane, $\cos^{-1}[(k/c)(d\sigma_A/ds)]$, where s is spatial arc length along the MMA. The angle ϕ_A can be interpreted as the image-space angle between the boundary at the scale of the MMA and the axis. The distance of the boundary region from $\mathbf{x}_A$ along the directions $\mathbf{u}_{A+}$ and $\mathbf{u}_{A-}$ is proportional to σ_A.

The MMA thus induces a corresponding boundaries by placing smears (e.g., Gaussian) of variance proportional to its scale σ_A centered at the boundary positions determined as just described: $\int_{(\mathbf{x}_A,\sigma_A)\in\mathrm{MMA}} G[\mathbf{x}_A + k_1\sigma_A\mathbf{u}_{A\pm}; k_2\sigma_A]$, where $G[\mathbf{x}; \sigma]$ indicates an isotropic Gaussian with mean $\mathbf{x}$ and variance σ. This MMA-induced boundariness can be used to permit directly measured boundarinesses at the positions in question and at scales smaller than that of the corresponding MMA point.

Initially, boundariness in many regions contributes to the medialness that underlies an MMA, but, ultimately, only boundariness in boundary regions associated with the object should contribute to its medialness. Roughly,

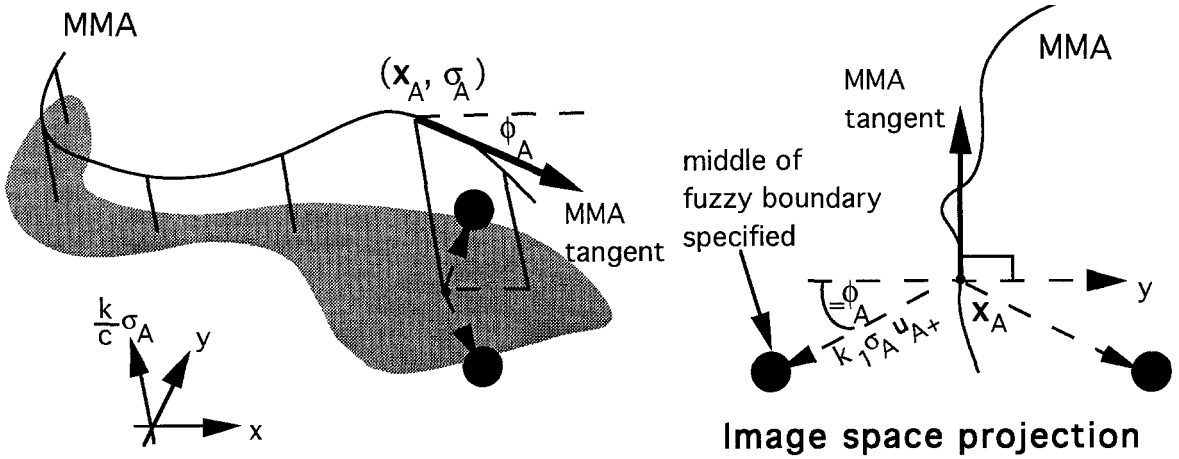

Fig. 8. Association of boundary regions with MMA points.

the boundariness at a position and scale should contribute only in that direction for which the medialness at the corresponding scale is greatest (for more detail see [20]). The result is that only a few points contribute to a winning medialness and thus determine its direction and, consequently, their position. This feedback was part of the computation leading to figure 9.

Neighbor interference poses an additional difficulty with the approach as specified so far. Large-scale-boundariness kernels appropriate for characterizing objects of large width overlap objects adjacent to or within the object being analyzed. But boundariness derived from medialness can be used to restrict the boundariness receptive fields to the region of the object. This is accomplished by letting the scale for boundariness be defined not according to the time of a uniform diffusion equation (the variance of a Gaussian envelope) but according to the time of a variable-conductance diffusion equation. [25], where the conductance is monotonic decreasing with the medialness-based boundariness. This behaves as though the space near object boundaries is stretched before the medialness measures are made.

6 Geometry from the MMA and Boundariness

While working from image intensities, the multiscale medial axis and the associated boundari-

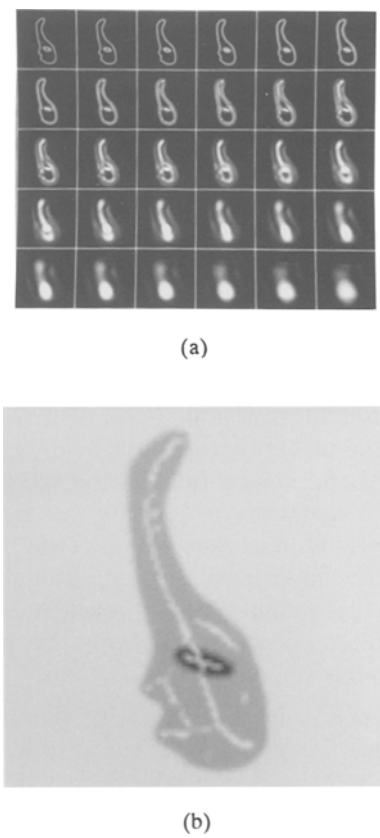

Fig. 9. (a) Medialness values and (b) multiscale medial axis superimposed on an image. Medialness and boundariness feedback by boundary/MMA correspondence.

nesses communicate much about the shape of the object:

(a) The MMA direction in scale space determines both the direction of the object in space and the angle of the boundary relative to the MMA at the scale of its local width. The tolerance of both of these values is also determined. Derivatives of these values with respect to distance along the MMA in image space determine the curvature of the axis and the curvature of the boundaries relative to the axis direction.
(b) Boundary detail is given by the curvature of a ridge in boundariness in scale space, and this detail is associated with the corresponding MMA points. If the boundary is wiggly, smaller-scale MMAs will be found corresponding to the protrusions at those scales. The means of determining the lower limit of the scale at which the image data support such boundariness measurements is under study but is beyond the scope of this paper.
(c) Subobjects are defined by MMAs spatially at a smaller scale than and inside the region define by a larger-scale MMA.
(d) Certain symmetries, namely associations between involutes, at all scales are defined by the pairs of boundary points associated with the MMA at its scale. This includes not only the principal symmetry of the object and symmetries of its detail and subobjects but also external symmetries (object indentations) and symmetries larger than the principal symmetry, as with the global medial axis of Blum.

Like Blum's medial axis, the MMA separates object curvature from width properties, thus preserving shape measures across small changes in local orientation produced by warping or bending; allows the identification of the visually important ends of protrusions and indentations, i.e., points of extremal boundary curvature; naturally incorporates size constancy and orientation independence; and generalizes to 3-D. However, unlike Blum's medial axis, it provides this information at a scale appropriate to the object width, and so it is a more stable property of the object—there is low sensitivity to noise in the boundary because this appears at a smaller scale than the axis. There is also stability relative to edge detectors, deriving from the fact that the MMA is tied to the center of the object and so cannot get lost, as an edge boundary can. Moreover, the MMA induces a natural hierarchy within objects by level of geometric detail and between objects and subobjects.

As for boundary properties, we have seen that the understanding of them must follow (and interact with) the characterization of object shape by multiscale medial properties. Only with medial information can we determine the boundary region that may belong to a particular object and the object locations and scales that can affect the boundariness in that region. This is very different from the standard view, in which the boundary is determined first.

Acknowledgments

The authors thank David Eberly, Robert Gardner, Christine Scharlach, and James Damon for collaboration on mathematical aspects of the multiscale medial axis. We are grateful to Andrea van Doorn, Guido Gerig, Wim van de Grind, Jan Koenderink, Andre Noest, and the editors of [21] for helpful suggestions. We thank Carolyn Din for help in manuscript preparation, Bo Strain for photography, and Graham Gash for computer-systems support.

Notes

1. In our subsequent work we have come to call the multiscale medial axis, the "core."
2. In [20] we give an alternative locus as defining the MMA. Instead of finding a maximum in scale of medialness and a spatial ridge of the result, we find the 1-D ridge of medialness directly in 3-D scale space by using scale-normalized derivatives. Both loci have been found to be effective.

References

1. J. Blom, "Affine invariant corner detection," in *Topological and Geometrical Aspects of Image Structure*, chap. 5,

Ph.D. dissertation, University of Utrecht, Utrecht, The Netherlands, 1992.
2. H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form*, W. Wathen-Dunn, ed., MIT Press: Cambridge, MA, 1967, pp. 363–380.
3. H. Blum and R.N. Nagel, "Shape description using weighted symmetric axis features," *Patt. Recog.*, vol. 10, pp. 167–180, 1978.
4. J. Canny, "A computational approach to edge detection," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. PAMI-8, pp. 679–698, 1986.
5. J.L. Crowley and A.C. Parker, "A representation for shape based on peaks and ridges in the difference of low-pass transforms," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. PAMI-9, pp. 158–170, 1984.
6. J.G. Daugman, "Two-dimensional spectral analysis of cortical receptive field profiles," *Vis. Res.*, vol. 20, pp. 847–856, 1980.
7. L.S. Davis, "A survey of edge detection techniques," *Comput. Vis., Graph., Image Process.*, vol. 4, pp. 248–270, 1975.
8. D.S. Fritsch, "A medial description of greyscale image structure by gradient-limited diffusion," *Proc.*, vol. 1808, pp. 105–117, 1992.
9. D.S. Fritsch, "Registration of radiotherapy images using multiscale medial descriptions of image structure," Ph.D. dissertation, Department of Biomedical Engineering, University of North Carolina, Chapel Hill, NC, 1993.
10. B.M. ter Haar Romeny, L.M.J. Florack, J.J. Koenderink, and M.A. Viergever, "Scale-space: its natural operators and differential invariants," in *Information Processing in Medical Imaging*, A.C.F. Colchester and D.J. Hawkes, eds., Lecture Notes in Computer Science, vol. 511, Springer-Verlag: Berlin, 1991, pp. 239–255.
11. F. Heitger, L. Rosenthaler, R. von der Heydt, E. Peterhans, and O. Kübler, "Simulation of neural contour mechanisms: from simple to end-stopped cells," Image Science Division, Communications Technical Laboratory, ETH Zurich, Tech. Rep. BIWI-TR-126, 1991.
12. R. von der Heydt, E. Peterhans, and G. Baumgartner, "Illusory contours and cortical neuron responses," *Science*, vol. 224, pp. 1260–1262, 1984.
13. J.J. Koenderink and A.J. van Doorn, "Receptive field families," *Biol. Cybernet.*, vol. 63, pp. 291–297, 1990.
14. J.J. Koenderink, *Solid Shape*, MIT Press: Cambridge, MA, 1990.
15. M. Leyton, "Symmetry–curvature duality," *Comput. Vis., Graph., Image Process.*, vol. 38, pp. 327–341, 1987.
16. M. Leyton, *Symmetry, Causality, Mind*, MIT Press: Cambridge, MA, 1992.
17. S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. PAMI-11, pp. 674–693, 1989.
18. S. Mallat, "Zero-crossings of a wavelet transform," *IEEE Trans. Informat. Theory*, vol. IT-37, pp. 1019–1033, 1991.
19. D. Marr, *Vision*, Freeman: San Francisco, CA, 1982.
20. B.S. Morse, S.M. Pizer, and A. Liu, "Multiscale medial analysis of medical images," in *Information Processing in Medical Imaging (IPMI) '93)* H.H. Barrett and A.F. Gmitro, eds., Lecture Notes in Computer Science, vol. 687, Springer-Verlag: Berlin, 1993, pp. 112–131. To appear in revised form in *Image & Vis. Comp.*
21. Y.-L. O, A. Toet, H. Heijmans, D.H. Foster, and P. Meer (eds.), *Shape in Picture: The Mathematical Description of Shape in Greylevel Images*, Springer-Verlag: Berlin, 1994.
22. C.A. Burbeck and S. M. Pizer, "Object representation by cores," in preparation. University of North Carolina Dept. of Comp. Sci. Tech. Report TR94–160.
23. A.B. Watson, "The cortex transform: rapid computation of simulated neural images," *Comput. Vis., Graph., Image Process.*, vol. 39, pp. 311–327, 1987.
24. H. Wechsler, *Computational Vision*, Academic Press: San Diego, CA, 1990.
25. R.T. Whitaker, "Geometry-limited diffusion in the characterization of geometric patches in images," *Comput. Vis., Graph., Image Process.*, vol. 57, pp. 111–120, 1993.
26. H.R. Wilson and J.R. Bergen, "A four-mechanism model for threshold spatial vision," *Vis. Res.*, vol. 19, pp. 19–32, 1979.

Stephen M. Pizer received the Ph.D. degree in computer science from Harvard University in 1967. A Kenan Professor of Computer Science, Radiation Oncology, Radiology, and Biomedical Engineering at the University of North Carolina at Chapel Hill, he heads the multidepartmental Medical Image Display Research Group and co-leads the Department of Computer Science's Graphics and Image Laboratory. His research, focused since 1962 on medical image processing and display, covers human vision, image analysis, interactive 3-D graphics, and contrast enhancement. He has active collaborations with laboratories in The Netherlands, Switzerland, and the United States and is associate editor for display of *IEEE Transactions on Medical Imaging.*

Christina A. Burbeck is research associate professor of computer science and psychology at the University of North Carolina at Chapel Hill. She received her B.A. and M.A. degrees in mathematics from the University of California at San Diego and her Ph.D. in psychology with an emphasis in cognitive science from the University of California at Irvine. Her research focuses on spatial vision in humans, with particular emphasis on spatial relations in 2-D images. Before coming to Chapel Hill in 1990, Dr. Burbeck was director of the visual sciences program at SRI International in Menlo Park, California.

James M. Coggins is an associate professor and associate chairman in the Department of Computer Science at the University of North Carolina at Chapel Hill. He obtained graduate degrees in computer science from Michigan State University, the M.S. in 1977 and the Ph.D. in 1983. His main interest is in developing multiscale, geometric statistical methods for image pattern recognition that present a unified approach to image analysis.

Daniel S. Fritsch obtained his Ph.D. in biomedical engineering from the University of North Carolina at Chapel Hill. He has worked there with members of the Departments of Computer Science and Radiation Oncology to develop methods for the automatic registration of radiotherapy portal verification images. He has contributed to the development of a method for describing objects in images, called the multiscale medial axis (MMA), and he has worked on the development of image-registration methods based on the MMA. Dr. Fritsch is currently a postdoctoral fellow in the Department of Biomedical Engineering at the University of North Carolina and is involved with the development of real-time, fully automated methods of radiation-treatment verification.

Bryan S. Morse is presently completing his Ph.D. in computer science at the University of North Carolina at Chapel Hill under the direction of Stephen M. Pizer. He received his B.S. (magna cum laude) and M.S. degrees in computer science from Brigham Young University Provo, Utah in 1986 and 1990, respectively. He has also worked for International Business Machines Corporation in the design and implementation of image-based document storage and retrieval systems and in the development of multimedia-based educational software. His present research interests involve techniques for image analysis and interactive display.