Zoom-Invariant Vision of Figural Shape: The Mathematics of Cores

Stephen M. Pizer, David Eberly, Bryan S. Morse*, Daniel S. Fritsch

Medical Image Display & Analysis Group

University of North Carolina

*Department of Computer Science

Brigham Young University

Mailing address:

Sitterson Hall, Campus Box #3175

University of North Carolina

Chapel Hill, NC 27599-3175

email: pizer@cs.unc.edu

Phone: 919/962-1785; FAX: 919/962-1799

ABSTRACT

Believing that figural zoom invariance and the cross-figural boundary linking implied by medial loci are important aspects of object shape, we present the mathematics of and algorithms for the extraction of medial loci directly from image intensities. The medial loci called cores are defined as generalized maxima in scale space of a form of medial information that is invariant to translation, rotation, and in particular, zoom. These loci are very insensitive to image disturbances, in strong contrast to previously available medial loci, as demonstrated in a companion paper. Core-related geometric properties and image object representations are laid out which, together with the aforementioned insensitivities, allow the core to be used effectively for a variety of image analysis objectives.

SYMBOLS

$\sigma, \quad \rho, \quad \theta$ Greek letter, standing for a scalar variable. Italic. Roman (italic) letters, upper case or lower case, may also be used for variables.

$\vec{x}, \quad \vec{\xi}$ Roman or Greek letter with arrow over it, standing for a vector variable. Italic.

$\xi_i, \quad \xi^i, \quad \xi^i_j, \quad \Gamma^k_{ij}, \quad R_{ijkl}$ Variable subscripted, superscripted, or both. The variable may be upper or lower case and Greek or Roman (italic), but the subscripts and superscripts will be one or more digits or lower case Roman (italic)

$P_{,k}$ or $P_{i,k}$ Variable with subscript with comma in it, with the comma possibly not preceded by any symbol, indicating a differentiated tensor

σ_{blur} variable subscripted with a word

$U_{AC}\left(\theta_i,\theta_j\right), \quad U_{S_1 S_2}\left(\theta_i,\theta_j\right)$ Upper-case letter with upper-case subscript, subscripted subscript

$\vec{p}^{\,t}$ Vector (or matrix) with superscript t, indicating transpose

$\mathfrak{I}$ Script upper case letter, indicating operator

$\mathbf{T}$ Bold upper case letter, indicating set

$f(x)$ letter followed by parenthesized expression, indicating function

$\vec{u}(\theta)$ letter with arrow over it followed by parenthesized variable(s), standing for vector function

$M(x,y,\sigma), \quad K(\vec{x};\sigma)$ letter followed by list of parenthesized variables (scalar or vector) separated by commas, with possibly one comma replaced by a semicolon

$\ln(f), \quad \cos^{-1}(\theta)$ Specific functions named by multiple letters and possibly a superscript, here the natural logarithm and the inverse cosine

$p(\Theta), \quad \det(V)$ Probability function of the set upper-case theta, determinant of the matrix upper-case V

$\forall, \quad \in, \quad \rightarrow$ Mathematical symbols for "for all", "is an element of", and "approaches"

$\bullet, \quad \circ \quad \sim, \quad \times$ Mathematical symbols for dot product (filled dot), composition (open dot), and is a neighbor of, all centered vertically on the line, and for times operation

SYMBOLS continued

$\mathfrak{R}^n$, $\mathfrak{R}^+$ Mathematical symbols for the set of reals, normally superscripted by "n", a digit, or "+". The mathematical symbol is sometimes written as an R with its vertical bar doubled.

$\int_0^{2\pi}$, $\int_{\mathfrak{R}^n}$ Mathematical symbol for integral with range of integration indicated by symbols below integral symbol and as well optionally above it

$\sum_{i=1}^{n}$, $\sum_{i \in C}$ Mathematical symbol for sumation (upper case Greek sigma) with range of summation indicated by symbols below summation symbol (of form "variable = expression" or variable is an element of expression) as well optionally above it (single variable or digit)

∇ Nabla (upside-down upper-case delta), indicating gradient

D, D^k Upper case D, possibly superscripted by a digit or variable, indicating differentiation

$\hat{D}^2 f_{ij}$ Expression beginning with a D with a circumflex over it, possibly superscripted and followed by the operand of the differentiation, indicating scale space differentiation

$|x|$ Function or variable surrounded by vertical bars, indicating magnitude

$\dfrac{x}{\sigma}$ Fraction, where vertical setting is desired

$\sqrt{\left(\dfrac{x}{\sigma}\right)^2 + \left(\dfrac{y}{\sigma}\right)^2}$ Radical symbol, covering an expression, indicating square root

$\left\langle \overset{\perp}{\theta_i} \right\rangle$ Angle brackets surrounding a vector or scalar random variable, indicating expected value

$d_{ss}^2\left[\overset{\perp}{\theta_i}\right]$ Function d (subscripted by ss and superscripted by 2) with an argument in square brackets, here indicating the squared scale space length of the operand

$\mathrm{var}_{i \in S_1, j \in S_2, i \sim j}$ Word (abbreviation of variance) with complex suscript

1   INTRODUCTION

Characterization of object shape is useful in describing objects in images or modeling objects as part of segmentation or other image analysis objectives. Two types of relationships between boundary points have been suggested as capturing aspects of object shape: the local boundary relations of slope and curvature and cross-figural relations captured by a medial locus (see Figs. 1.1 and 1.2). In this paper we focus on a means of defining and extracting medial loci directly from image intensities in a way insensitive to intensity noise, blurring, or boundary details. Beginning from the long-standing characterization of the shape of an image object as everything that is invariant to rotation, translation, and zoom and to luminance properties of the object interior, we develop in this paper a zoom invariant medial locus that exhibits the desired extractability from image intensities in a way insensitive to image disturbances. The insensitivity is verified experimentally and discussed at length in a companion paper entitled *Zoom-invariant vision of figural shape: effects on cores of image disturbances* [39]. This introduction makes this approach more precise.

Like many before us, we see an object as being a directed graph of primitive regions that we call *figures* [3, 43], such as the whole object and individual protrusions and indentations (see Fig. 1.1 for an example), and we see shape as having three aspects:

  1) Figural shape: the shape of the individual figures, derived from cross-figural linking, or equivalently, the medial locus;

  2) Interfigural shape: the size, position, and orientation of a figure relative to its parent figure in the graph; and

  3) Boundary shape: the local shape and texture of the boundary locus, relative to the information predicted by the figural description.

This paper derives mathematically a medial representation of figural shape, the *core*. The loci in scale space that form these cores have mathematical properties much different from other medial loci but still retain the powerful abilities to extract the many aspects of shape provided by an analysis in which medial and boundary extraction are coupled [43]. As a result of the sensitivity to shape at the scale of the figure and insensitivity to image disturbances, certain image analysis techniques based on cores are effective.

5

Fig. 1.1. Some of the figures making up this object are the teardrop and its two children in the graph representation shown: the protrusion on the teardrop's right and the indentation on its left. The boundary texture of being smooth in some regions and undulating in others is included in the boundary shape.


Fig. 1.2. Boundary point linking via a skeletal middle

The rest of this section clarifies the definitions and assumptions determining the core and the algorithms for its extraction from images. First we must clarify the definition of invariance that we are using. We define the invariance of a measurement M applied everywhere on a manifold $\mathbf{M}$ to a family $\mathbf{T}$ of transformations T to mean that

$$(\forall T \in \mathbf{T})(\forall \vec{x} \in \mathbf{M})\big[M(\vec{x})T = TM(\vec{x})\big]. \tag{1}$$

Second we clarify the definition of a medial locus with the required invariances. Following [1, 2], we understand the representation of figural shape to be the locus of middle and width points, $\overset{\perp}{m}, r$, linking two points $\overset{\perp}{b_i}$, i=1,2 at opposite sides of a figure such that each $\overset{\perp}{b_i}$ behaves like an object boundary with normals, $\overset{\perp}{u_i}$:, i.e., $\overset{\perp}{b_i} = \overset{\Upsilon}{m} + r\overset{\Upsilon}{u_i}$, $i = 1, 2$, where the $\overset{\perp}{u_i}$ are unit vectors. What is new is that we restrict the measurements allowed in determining the medial locus to functions of location $\overset{\perp}{m}$ and radius r that are direct functions of image intensities and are larger the more the position $\overset{\perp}{m}$ behaves like a medial point with respect to a radial width r. In addition, we insist that the measurements are translation (shift) invariant. Also, for now we will restrict the primitive image measurements to those that are linear in the intensity values (though we will relax this requirement later). Thus they must involve convolution with some set of kernels $w(\overset{\perp}{x})$ or equivalently computation of the measurement at $\overset{\perp}{x_0}$ using a set of weighting functions $w(\overset{\Upsilon}{x_0} - \overset{\Upsilon}{x})$. In addition, we will require the operations to be rotation invariant. This requirement implies that $w(\overset{\perp}{x})$ must be rotationally symmetric.

What is the implication of the further requirement, central to figural shape, that **medial locus extraction must be invariant not only to rotation and translation but also to zoom**? Any weighting function $w(\overset{\Upsilon}{x_0} - \overset{\Upsilon}{x})$ can be characterized as having an aperture size, or *scale*, defined as the root mean squared width, σ_w, of w : $\sigma_w^2 = \int_{\Re^n} |\overset{\Upsilon}{x}|^2 w(\overset{\Upsilon}{x}) / \int_{\Re^n} w(\overset{\Upsilon}{x})$. We will characterize image operations not only by their location of application $\overset{\perp}{x_0}$ but also their scale σ_w. It should be intuitively obvious that zoom invariance requires that as one zooms into the scene the aperture must be proportionately zoomed, i.e., its scale must be magnified. Proving this fact, however, involves quite some subtlety, which we will not reproduce here. The detailed analysis can be found in [23].

This requirement of scale proportionality with zoom leads to considering the kernels in families vs. σ_w. That is, the kernel w should be a function not only of offset $\overset{\perp}{x}$ but also of scale σ_w. Moreover, zoom invariance requires that the medial width r must increase proportionately with zoom, so functions of width r can equivalently be thought of as functions of σ_w.

As discussed in section 2, the theoretical studies of scale space imply that linear combinations of appropriately scale-normalized derivatives of Gaussians centered at a point satisfy the translation, rotation, and zoom requirements that we have set if the scale σ of the Gaussians increases proportionately with the zoom. Thus we will produce medial measurements that are translation, rotation, and zoom invariant if they are produced from kernels $K(\overset{\perp}{x}, \sigma)$ that are linear combinations of scale-normalized derivatives of Gaussians with scale σ with the linear combinations satisfying a rotational symmetry such that the radius of rotation is proportional to σ. Thus, $\sigma_K \propto \sigma \propto r$. Section 2 will derive some particular linear combinations of Gaussian derivatives that do the job.

Given a set of measurements $M(\overset{\perp}{x}, \sigma)$ [or equivalently $M(\overset{\perp}{x}, r)$] produced using kernels $K(\overset{\perp}{x}, \sigma)$, we need to choose the medial locus by virtue of M being locally high there. **In analogy to extracting boundaries as a generalized maximum, or ridge, of a graded measurement of a position behaving like a boundary** (cf., for example, Canny's approach [1987] to boundary extraction)**, our approach is to extract our medial locus by a generalized maximum of the graded measurement M of a (position, width) combination behaving as medial.** We will call the graded measurement *medialness* and will write it as a function M from scene position and scale: $M: \Re^n \times \Re^+ \to \Re$. The space $\Re^n \times \Re^+$ capturing all positions and scales is called *scale space*. The generalized maximum will take the form of a directional maximum in a subset of the directions of the space in which the medial measurement is made [13; 22].

Section 2 covers means of producing a scale space of medialness that satisfies the specified invariances. In section 3 we derive scale space operators that generalize previous definitions of ridges of image functions [24] in a way so as to retain the invariances, and by using these operators on the scale space of medial information to extract cores. The geometry of this scale space and the mathematical properties of cores are then derived. After a brief summary in section 4 of the companion paper on the insensitivity of cores to image disturbances, section 5 derives core-related geometric properties and image object representations that have been shown useful for object-related image analysis operations such as registration, recognition, segmentation, and shape measurement.

2   MEDIAL THEORY

2.1 Requirement of Scale

We need to make more mathematically precise the concepts of *scale* and *scale space* , as they apply to figural shape. We make the assumptions given by [31], namely that not only do the three invariances discussed in section 1 hold but also that linearity properties hold and also that

   a) the scale space has the properties that intensity maxima must decrease with scale (maximum suppression) and that

   b) for all $\Delta\sigma > 0$ the image information at scale $\sigma + \Delta\sigma$ can be derived from that at σ by an operation of the same form as that deriving the information at scale σ from that at the inner scale (semigroup property).

Lindeberg shows that a consequence is that increasing scale corresponds to running the diffusion equation with $t = \sigma^2/2$, i.e., building our medialness functions out of Gaussians of scale σ and their derivatives.

We have seen that zoom invariance leads to a medialness operator in which the radius is proportional to the scale of the Gaussians on which it its built. Let $\rho > 0$ be the constant of proportionality between the aperture size σ and the figure's radial width r , i.e., $\sigma = \rho r$. To avoid the need to specify an arbitrary origin for the magnification, the scale-to-width relationship is more appropriately quantified in terms of how changes in radial width, dr , are related to changes in scale, σ . The relationship is concisely

$$d\sigma = \rho\, dr. \tag{2}$$

Denote the input intensities by $I(\vec{x})$. For simplicity we assume that the input is defined for all $\vec{x} \in \Re^n$. A vision system produces multiscale data $L(\vec{x},\sigma)$ for $(\vec{x},\sigma) \in \Re^n \times [\sigma_0,\sigma_1]$,where σ_0 , called the inner scale, is the minimum scale consistent with the image production process and σ_1 , called the outer scale, is the largest scale consistent with the image size. The function L can be described generally by the operator equation

$$L(\vec{x},\sigma) = \Im I(\vec{x}), (\vec{x},\sigma) \in \Re^n \times [\sigma_0,\sigma_1],$$
$$L(\vec{x},\sigma_0) = I(\vec{x}), \tag{3}$$

9

where the operator $\Im$ is determined by additional requirements imposed on the vision system, such as that it measures medialness.

As stated above, the requirements of semigroup, maximum suppression, linearity, and invariance with respect to translations, rotations, and zoom imply that scale space is generated by linear diffusion. When written in terms of σ rather than the more common t $\left(t = \sigma^2/2\right)$ and when the relation between scale and space given by the factor ρ is reflected, the diffusion equation becomes

$$\frac{\partial L}{\partial \sigma} = \frac{1}{\rho}\nabla \bullet (\sigma \nabla L), (\vec{x},\sigma) \in \Re^n \times [\sigma_0,\sigma_1],$$
$$L(\vec{x},\sigma_0) = I(\vec{x}).$$

(4)

In terms of the physical model for the heat equation, ρ is the density function and σ is the conductance function.

The diffusion process *generates* the multiscale data. The figural vision system also must interpret this data in a geometric way in order to construct the representations that later stages of processing will use. Geometric interpretation requires imposing a metric on scale space. To obtain the desired invariances, note that a measured *spatial difference* is meaningful only relative to the scale at which it is measured. Similarly, when making multiscale measurements, a measured *scale difference* is meaningful only in the context of the scales at which it is measured. These properties, which follow from zoom invariance, suggest specifying differential forms as the measurement tools. As a Euclidean space the 1-forms used for $\Re^n \times [\sigma_0,\sigma_1]$ are $dx_i, 1 \le i \le n$, and $d\sigma$. However, to retain the desired invariances, the dimensionless 1-forms to be used for scale space measurements are $\frac{dx_i}{\sigma}, 1 \le i \le n$, and $\frac{d\sigma}{\sigma}$. The geometry of scale space is determined by the metric involving these forms:

$$ds^2 = \frac{d\vec{x} \bullet d\vec{x}}{\sigma^2} + \frac{d\sigma^2}{\rho^2\sigma^2}.$$

(5)

In order to compare spatial differences $d\vec{x}$ and scale differences $d\sigma$, we need to use the proportionality constant, ρ, between width and scale. The inclusion of ρ in the metric makes the "units" of $d\vec{x}$ and $d\sigma/\rho$ the same and allows us to combine space and scale so that geometric information can be

properly interpreted. We will see in section 3 and the associated appendix that scale space with metric (5) has a non-Euclidean geometry.

The metric and the diffusion process are intimately linked. If $M(\overset{\scriptscriptstyle\perp}{x},\sigma)$ represents the real-valued measurement of interest in scale space, then changes in M are measured as

$$dM = \sum_{i=1}^{n}\frac{\partial M}{\partial x_i}dx_i + \frac{\partial M}{\partial \sigma}d\sigma = \sum_{i=1}^{n}\sigma\frac{\partial M}{\partial x_i}\frac{dx_i}{\sigma} + \rho\sigma\frac{\partial M}{\partial \sigma}\frac{d\sigma}{\rho\sigma}.$$

The natural derivatives to take in scale space are therefore $\sigma\partial M/\partial x_i, 1 \le i \le n$, and $\rho\sigma\partial M/\partial\sigma$. These quantities are dimensionless. Now the diffusion can be viewed as $\rho\sigma\dfrac{\partial}{\partial\sigma}L = (\sigma\nabla)\bullet(\sigma\nabla)L$, where the left-hand side is a single application of the scale space derivative with respect to scale and the right-hand side is a repeated application of the scale space spatial gradient. We will show in section 3 how measurements of figures, such as figure width, are based on the scale space differentiation.

2.2 Boundariness and Medialness

Our own vision system appears to recognize a figure in an image by locating and pairing opposing boundaries of the figure [3, 4, 29]. The boundaries are usually noticeable because of sharp contrast in luminance at those locations, but the contrast may be in some other property, as well. Sharp contrast is related to locally large directional derivatives in luminance. However, the pairing of opposing boundaries is a multilocal task which requires the full power of multiscale analysis and scale space. A well accepted starting point for human or computer vision is that the vision system makes a measure of *boundariness* at each position $\overset{\scriptscriptstyle\perp}{x}$, scale σ, and orientation $\overset{\scriptscriptstyle\perp}{u}$. Denote this function as $B(\overset{\scriptscriptstyle\perp}{x},\sigma,\overset{\scriptscriptstyle\perp}{u})$.

Specific choices for B might depend on the system and the task; a simple one measuring luminance change is
$$B(\overset{\scriptscriptstyle\perp}{x},\sigma,\overset{\scriptscriptstyle\perp}{u}) = \overset{\scriptscriptstyle\perp}{u}\bullet\sigma\nabla L(\overset{\scriptscriptstyle\perp}{x},\sigma) \qquad (6)$$
for bright (dark) figures on a dark (bright) background, or $\quad B(\overset{\scriptscriptstyle\perp}{x},\sigma,\overset{\scriptscriptstyle\perp}{u}) = \left|\overset{\scriptscriptstyle\perp}{u}\bullet\sigma\nabla L(\overset{\scriptscriptstyle\perp}{x},\sigma)\right| \qquad (7)$

for figures where the foreground/background intensity ratios vary through 1 as the figure boundary is traversed. More complicated and more effective boundariness measurements involve the strengthening of a simply measured boundariness according to its agreement in direction or curvature to continuations of the boundary [51].

11

The multilocal task of pairing opposing boundary points (called *involutes*) requires matching points whose boundariness in a putative boundary normal direction meet at a common distance at a putative central location of the figure (Fig. 1.2). Consistent with extracting boundaries via a graded measure of boundariness, we propose that the vision system makes a measure of *medialness* at each position $\overset{+}{x}$ and scale σ by accumulating boundariness values at locations whose common distance from $\overset{+}{x}$ is distance $r = \sigma / \rho$ in a way producing a function invariant to translation, rotation, and zoom.

There are many functions, $M(\overset{+}{x},\sigma)$, that satisfy the invariance requirements for medialness and for some class of image objects provide a local measure that increases with behaving like a medial point and width. Each function is suitable for a particular class of images or image objects. In addition to commuting with the operations of rotation and translation of an image, the medialness functions must be zoom invariant. It can be shown that these invariances hold if

  a) $M(\overset{+}{x},\sigma)$ is based on normalized Gaussian derivatives of intensity, $\sigma^k D^k L$, where

   $L(\overset{+}{x},\sigma) = G(\overset{+}{x};\sigma) * I(\overset{+}{x})$ and $G(\overset{+}{x};\sigma)$ is a Gaussian with zero mean and variance σ^2; and

  b) the set of locations at which derivatives contribute to $M(\overset{+}{x},\sigma)$ are offset from $\overset{+}{x}$ by vectors that

   scale with σ, and are positioned in a rotationally invariant fashion relative to $\overset{+}{x}$.

We classify each medialness function in two ways. First, the function is either *central* or *offset*. A central function is one for which $M(\overset{+}{x},\sigma)$ is measured using only low order, spatial derivative information of $L(\overset{+}{x},\sigma)$ at the putative spatial center $\overset{+}{x}$ at scale σ. An offset function is one for which $M(\overset{+}{x},\sigma)$ is measured by querying neighbors at some finite spatial distance from the point $\overset{+}{x}$ at scale σ. Central medialness functions have the property that they attempt to localize object boundaries by averaging spatial information about $\overset{+}{x}$ over some region whose average radius is proportional to σ. Such functions are desirable when the objects of interest contain uncorrelated, small scale noise. The functions are not as effective when the object interior contains correlated, large scale noise (in this category we include image objects interior to the object of interest). On the other hand, offset medialness functions attempt to localize object boundaries by accumulating information of a relatively small scale, σ, at neighbors which are at a distance proportional to σ from the test point $\overset{+}{x}$.

Fig. 2.1 Binary sawtooth object

Second, medialness functions are classified as *linear* or *adaptive*. In each case $M(\overset{+}{x},\sigma)$ is computed by some weighting of information from the original image $I(\overset{+}{x})$ in a neighborhood of $\overset{+}{x}$ whose size is proportional to the scale σ. In the linear case the neighborhood weights are necessarily radially symmetric to satisfy the rotational invariance and are data-independent. Thus, a linear medialness function can be computed as a linear convolution of some kernel $K(\overset{+}{x};\sigma)$ with $I(\overset{+}{x})$. In the adaptive case the neighborhood weights are data-dependent, with properties such as the orientation or extent of the weights depending on the image data in the neighborhood. The next subsections illustrate the ideas with specific medialness functions for two spatial dimensions. Fig. 2.1 is a binary sawtooth object whose medialness is computed using the specific functions.

2.2.1. Central Linear Medialness

An example of a central linear medialness function is the linear convolution

$M(x,y,\sigma) = K(x,y;\sigma) * I(x,y)$, where the kernel is the normalized Laplacian of a Gaussian[*] $G(x,y;\sigma)$:

$K(x,y,\sigma) = -\sigma^2 (G_{xx} + G_{yy}) = -\dfrac{2-R^2}{\sigma^2} e^{-\frac{R^2}{2}}$, where $R = \sqrt{(\dfrac{x}{\sigma})^2 + (\dfrac{y}{\sigma})^2}$. Fig. 2.2 shows a rendered

graph of the kernel and medialness of a binary object at small, medium, and large scale. This medialness is extensively used in [16, 17] for extracting anatomic objects with nonparallel sides, approximately uniform interiors, edges of fixed contrast polarity, and possibly low signal to noise ratio in portal radiographs, CT scans, MR images, and other medical images.

---

[*] Throughout this paper a subscript on a function indicates a derivative with respect to the subscripted variable or in the subscripted direction. Double subscripts indicate second derivatives, etc.

Fig. 2.2. Upper left: linear, central medialness kernel. Upper right: medialness of sawtooth object at r=4 pixels, in a 128 x 128 image. Lower left: medialness of sawtooth object at r=8. Lower right: medialness of sawtooth object at r=16.

2.2.2 Offset Linear Medialness

An example of an offset linear medialness function is the linear convolution

$M(x, y, \sigma) = K(x, y; \sigma) * I(x, y)$, where the kernel acts as an integrator of directional Gaussian derivatives around a circle with center (x,y) and whose radius r is proportional to scale. Define the oriented boundariness measurement $B(x, y, \sigma, \theta) = -\sigma \vec{u}^{\perp}(\theta) \bullet \nabla L(x, y, \sigma)$, where $\vec{u}(\theta) = \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix}$. The medialness is defined by $M(x, y, \sigma) = \int_{0}^{2\pi} B((x, y) + r\vec{u}(\theta), \sigma, \theta) d\theta$, where $r = \sigma / \rho$ for some constant ρ. Therefore, the kernel is defined by $K(x, y; \sigma) = \sigma \int_{0}^{2\pi} \left[-\vec{u}(\theta) \bullet \nabla G((x, y) + r\vec{u}(\theta), \sigma) \right] d\theta$.

Fig. 2.3 shows a rendered graph of the kernel and medialness of a binary object at small, medium, and large scale. This medialness has been analyzed in detail in [35, 36,37,38]. Because the kernel is near zero in the central region but is sensitive to boundary pairs at all angles with the same contrast polarity at

14

both edges, this kernel has been found effective for analyzing general medical image objects that have a variety of shapes, significant internal intensity structure, and a fixed polarity of contrast around the object. The cortex of the brain in CT head images is an example.


Fig. 2.3. Upper left: linear, offset medialness kernel. Upper right: medialness of sawtooth object at r=4. Lower left: medialness of sawtooth object at r=8. Lower right: medialness of sawtooth object at r=16.

2.2.3 Central Adaptive Medialness

We now relax the linearity requirement. This allows the choice of medialness operators that respond only to involute pairs and are selective to the image properties at the involutes, while retaining the required invariances. For example, the orientation of the operator can be adjusted to respond most sensitively to the orientation of parallel boundaries, or to respond most sensitively to the the inter-boundary angulation (i.e., the what Blum called the object angle) at the involutes. Or its elongation can be adjusted to respond to the elongation of the gap between the figure of interest and the a neighboring figure. Or the operator can be set to respond optimally to the polarity of intensity change at the boundaries.

An example of a local adaptive medialness function is $M(x,y,\sigma) = -\sigma^2 L_{pp}^{rr} = -\sigma^2 \hat{p}^t D^2 L \hat{p} = -\sigma^2 \lambda$,

15

where D^2L is the 2x2 matrix of second-order spatial derivatives of L and where $D^2 L\vec{p} = \lambda\vec{p}$, and λ is the largest magnitude eigenvalue of D^2L. At a single location the neighborhood weights, thought of as a kernel, are $K(x,y;\sigma) = -\sigma^2 G_{xx}(x,y;\sigma)$, but the kernel is oriented in the direction of vector $\vec{p}$. Fig. 2.4 shows a rendered graph of the kernel and medialness of a binary object at small, medium, and large scale.


Fig. 2.4. Upper left: adaptive, central medialness kernel; the maximum response over all orientations of this kernel is chosen. Upper right: medialness of sawtooth object at r=4. Lower left: medialness of sawtooth object at r=8. Lower right: medialness of sawtooth object at r=16.

This medialness is analyzed in [16, 20]. It is particularly effective for objects with parallel sides and uniform interior intensity, such as a blood vessel (see Fig. 3.1.2). It is less sensitive to intensity variations along the object axis than the linear, central medialness shown in Fig. 2.2.

2.2.4 Offset Adaptive Medialness

An example of an offset adapative medialness function is

$$M(x,y,\sigma) = -\sigma\vec{u}(\theta) \bullet \nabla G((x,y) + r\vec{u}(\theta);\sigma) + \sigma\vec{u}(\theta) \bullet \nabla G((x,y) - r\vec{u}(\theta);\sigma),$$

where the angle θ is selected to maximize the right-hand side over all possible angles. At a single location the neighborhood weights, thought of as a kernel, are

$K(x, y; \sigma) = -\sigma G_x(x + r, y; \sigma) + \sigma G_x(x - r, y; \sigma),$ but the kernel is oriented in the direction of the vector $\vec{u}(\theta)$. Fig. 2.5 shows a rendered graph of the kernel and the medialness of a binary object at small, medium, and large scale. This kernel is the sum of two boundariness operators, each with a scale proportional to the medial radius.


Fig. 2.5. Upper left: adaptive, offset medialness kernel; the maximum response over all orientations of this kernel is chosen. Upper right: medialness of sawtooth object at r=4. Lower left: medialness of sawtooth object at r=8. Lower right: medialness of sawtooth object at r=16.

This medialness is especially effective for objects with parallel edges but with internal structures of scale significant compared to the object of interest. It is less sensitive to intensity variations along the object axis than the linear, offset medialness shown in Fig. 2.3. It was studied in [44] to help measure stenosis in arteries which are imaged under various noisy or blurry conditions.

Another example of offset adaptive medialness made by summing two boundariness results adapts not only the medial orientation but also the object angle. Such a medialness operator, discussed in [43],

particularly closely matches the Blum definition of the medial axis given earlier in the paper. It is not discussed further here because the results with this operator have not been fully demonstrated. A variant on this operator, due to McAuliffe [1996], optimizes the elongation of the medialness weighting function at each involute to avoid too much effect from neighboring objects.

Another example of an offset adaptive medialness function is obtained by integrating the absolute value of boundariness around the circle of radius r [35, 38]:

$$M(x,y,\sigma) = \int_0^{2\pi} | B((x,y) + r\vec{u}(\theta),\sigma,\theta) | \, d\theta.$$

In the linear case the kernel is designed to detect bright objects on a dark background. This nonlinear function is useful if an object's intensity near the boundary is in some locations brighter than the background and in some places darker than the background.

3   CORES

In this section we define the core as a generalized maximum, i.e., ridge, of medialness,. We argue that for the core based on such an invariant medialness to itself be invariant, the generalized maximum operator must itself satisfy these invariances. We then give the mathematics for producing such ridges.

Medialness collects bilocal information about linkable boundariness into a local measure of the degree to which a point in scale space behaves like a figural middle at a width that is a fixed multiple of the scale. We desire a medial locus, which we call a *core*, made from centrally located positions and a scale and based on medialness. Such a locus is a track at which medialness is locally large compared to nearby locations and scales. At such a point the medialness should decrease as you move towards the involutes and as scale increases or decreases. That is, it is a generalized local maximum of $M(\vec{x},\sigma)$, generalized so that the relative maximum need only be taken in a subset of the directions forming the space (here, a scale space). In other directions at such a point, e.g., in a direction tangent to the medial locus, the medialness need not be locally maximal. Local maxima of a function with respect to a restricted number of directions have been called *ridges* [24, 9]. The mentioned works involve the study of ridge definitions in Euclidean space. However, to attain the desired invariances of cores, the definition of ridge used here

18

must be invariant to rotation, translation, and zoom. Also, we need extensions of the definitions to scale space with its imposed metric (eq. 5). That is, the derivatives involved in finding the generalized maxima must be those that reflect the desired invariances. This will first require a description of the geometry of scale space that is very mathematical and thus is given in the appendix.

We believe it is necessary to lay out this mathematics in detail because the results are generally applicable to problems of locus definition in scale space and because the success of using medialness ridges to obtain medial loci in scale space depends on doing the mathematics correctly. Crowley et al. [1984] inspired our effort with their idea of using ridges in scale space for this purpose, but they achieved limited success because the definitions of medialness and ridge they used failed to have the necessary invariances and also their ridge definition had other mathematical weaknesses.

3.1 Ridges of Medialness

To produce a d-dimensional core manifold from a medialness function in a scale space of dimension $n+1$, we must choose the subset of $n+1-d$ directions in which medialness is locally maximal (has a zero first directional derivative and a negative second directional derivative). The choice of these directions, which determines our definition of ridge, must be invariant to rotation, translation, and zoom. Two options that we have proposed is that the two directions are

1) the two orthogonal directions of greatest convexity of medialness, leading to the *maximum convexity ridge* [24, 9];

2) the direction of pure scale change and the direction of greatest convexity of the projection of medialness at the optimal scale onto position, leading to the so-called *optimal scale ridge* [15, 1993; Pizer, 1994a].

3.1.1 The Maximum Convexity Ridge in Scale Space

The maximum convexity ridge definition is based on finding local maxima of functions $f\left(\vec{\xi}\right)$ where the domain is restricted to the span of the maximum convexity subset of eigenvectors of the second-derivative (Hessian) matrix. This definition applies also for f in the scale space $\vec{\xi} = \vec{x}, \sigma$, where we use

19

covariant derivatives rather than partial derivatives. The formal tensor products in the definition are manipulated so that scale space derivatives occur. This is essential for the numerical implementation of core construction. The mathematical details are given in section A.6 of the appendix.

For 2D images the scale space is three-dimensional, and the core is a curve (1-manifold [8]) in scale space. Thus there are two maximum convexity directions in which the medialness must be a maximum. For 3D images the scale space is four-dimensional, and the core is most commonly a 2-manifold in scale space. For objects with a nearly circular cross-section, a core which is a 1-manifold forming a skeleton of the tube also makes sense [22]. For the core that is a 2-manifold there are two maximum convexity directions in which the medialness must be a maximum, and for the core that is a 1-manifold there are three such directions.

3.1.2 Optimal scale ridge

Instead of choosing the vectors v_1^i and v_2^i in which f must be maximum as the directions of greatest convexity of f, in the case where f is medialness we may choose v_1^i as $(0,0,...,0,1)$, the unit vector in the σ direction. At each spatial location $\overset{\perp}{x}$ call those σ for which $f_\sigma(\overset{\perp}{x},\sigma) = 0$ and $f_{\sigma\sigma}(\overset{\perp}{x},\sigma) < 0$ *optimal scales* at $\overset{\perp}{x}$. Generally these optimal scales are a continuous function $\sigma(\overset{\perp}{x})$. For any of these optimal scale manifolds, we wish to choose v_2^i as the spatial direction of greatest convexity of the projection of f at optimal scale. That is, $(\overset{\perp}{x},\sigma)$ is on an optimal scale ridge of f if $f_\sigma(\overset{\perp}{x},\sigma) = 0$, $f_{\sigma\sigma}(\overset{\perp}{x},\sigma) < 0$, $f_{v_2^{\overset{r}{x}_i}}(\overset{\perp}{x},\sigma) = 0$, and $f_{v_2^{\overset{r}{x}_i}v_2^{\overset{r}{x}_i}}(\overset{\perp}{x},\sigma) < 0$, where v_2^i is the maximum convexity eigenvector of the scale space Hessian of M restricted to the image space with σ fixed at its optimum scale value (in the terms of the appendix $v_2^i h_{,i} = 0$ and $v_2^i h_{,ij} v_2^j < 0$, where $v_2^i = (w_1^i, 0)$, $h(\overset{r}{x}) = f(\overset{r}{x}, \sigma(\overset{r}{x}))$, and w_1^i is the eigenvector with the most negative eigenvalue of $D^2 h$). The optimal scale ridge can be shown to be equivalent to the maximum convexity ridge with v_1^i and v_2^i the solutions of the generalized eigenvector equation $A \circ (\overset{\perp}{x},\sigma) = \lambda B \circ (\overset{\perp}{x},\sigma)$ with A = the scale-space Hessian and $B = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}$ with the identity matrix of size equal to the number of spatial dimensions.

Our experience is that the maximum convexity ridge and the optimal scale ridge produce qualitatively similar results.

3.2 Core Extraction from 2D Images

Let $f_{,ij}u^j = \alpha u_i$, $f_{,ij}v^j = \beta v_i$, and $f_{,ij}w^j = \gamma w_i$, where $\alpha \leq \beta < \gamma$. Consistent with maximum convexity ridges, we will assume that the eigenvectors form an orthonormal system whose derivatives satisfy

$$\begin{bmatrix} u^i_{,j} \\ v^i_{,j} \\ w^i_{,j} \end{bmatrix} = \begin{bmatrix} 0 & 0 & a_j \\ 0 & 0 & b_j \\ -a_j & -b_j & 0 \end{bmatrix} \begin{bmatrix} u^i \\ v^i \\ w^i \end{bmatrix}$$ for some choice of continuous tensors a_j and b_j. Define $P = u^i f_{,i}$,

$Q = v^i f_{,i}$, and $R = w^i f_{,i}$. Core points satisfy $P = 0$, $Q = 0$, and $\alpha \leq \beta < 0$.

All of the techniques for core extraction begin from an approximate core point in scale space and flow to (see Appendix A.7.1) or search for a nearby point on the ridge. There are four techniques that we have used to follow the core locus from such a starting point:

1) Integrating the differential equation $P = 0$, $Q = 0$ (see Appendix A.7.2).

2) Extending the ridge in the w direction and finding the root of $P = 0$, $Q = 0$ in the plane of the maximum convexity eigenvectors at the extended point (see Appendix A.7.3 and [13]).

2a) Extending the ridge in the w direction and minimizing $DM^t(D^2M - \gamma I)DM$ in the plane of the maximum convexity eigenvectors at the extended point (see Appendix A.7.4 and [13]).

3) Sampling the local scale space in a Cartesian fashion and applying zero-trapping techniques combined with convexity constraints, resulting a marching ridges approach [22] that is comparable to the well-known marching cubes approach [32].

The relative strengths of these methods are as follows:

1) Following the ridge direction explicitly increases the chance that the ridge will not be lost.

2) Avoiding the solution of the differential equation lowers the required order of differentiation by one and allows much larger steps, considerably speeding the process. Methods 1 and 2 are in most common use in our laboratory.

2a) This method avoids having to decide whether a Hessian eigenvector or its negative is the

    continuous extension of the that eigenvector used at the previous step.

3) This method is speedier than the others because it allows larger steps and requires second

    derivatives of P and Q only at points that satisfy the zero test. It also generalizes nicely to

    following higher dimensional ridges in higher dimensional ridges.

3.3 Core Extraction Examples

Figs. 3.1 shows a number of examples of cores extracted from objects in 2D test images or medical

images. The medial axes (the projections of the cores onto image space) are shown as tracks on the

images, and the core widths are indicated by darkening or brightening the region formed by the union of

disks at each core point with each disk having the radius given by the core point at which it is centered.

These cores were extracted on various workstations in a period of just under a second to 3 seconds.

The results shown in Figs. 3.1 all use method 1 described in section 3.2 to find optimal scale ridges. The

medialness function used for the test object and the MRI brain slice was obtained via the linear, central

medialness kernel described in section 2.3.1. The medialness function used for the blood vessels and the

portal image was obtained via the adaptive, central medialness kernel described in section 2.3.3.

3.4 2D-from-3D Core Extraction

The aforementioned theory generalizes to intensity functions of $\Re^3$. This produces a 4-dimensional

scale space, $\Re^3 \times \Re^+$, in which normally cores are 2D manifolds. Each point on this manifold gives a

medial position between two opposing boundary faces as well as a radius (scale) of the fuzzy tangent

sphere.

Fig. 3.1: Original images and selected cores (light grey or black) and width envelopes (dark grey or white): Gun-shaped test object, MRI head, artery tree in angiogram, bones in portal radiograph.

We call these *2D-from-3D cores*, because they are 2D manifolds extracted from 3D images (via a 4D scale space). The underlying mathematics is given in section A.8 of the appendix. An algorithm for 2D-from 3D core extraction with an approach similar to the marching cubes algorithm [32] is reported in a separate paper [22].

4   INSENSITIVITIES TO IMAGE DISTURBANCES

Images and their content typically suffer from various forms of disturbance, the results of variations in

    the imaging process, e.g., blur, which produce different images of the same scene when imaged by

       different modalities, devices, or settings;

    random variables in the imaging process, e.g., intensity noise, which produce different images of the

       same scene even with the same device and settings; or

    the image content, which produce different images of scenes that are identical except for small object

       deformations or changes in boundary texture, or possibly the addition of protrusions or

       indentations.

Each of these disturbances may be thought of as having an associated spatial scale, i.e., spatial magnitude of the disturbance. Image blur clearly has a spatial extent. Greater levels of noise have a larger spatial influence when low-pass filtered to a fixed mean amplitude. And minor variations between objects may often be described as indentations or protrusions of varying size.

The companion paper entitled *Zoom-invariant vision of figural shape: effects on cores of image disturbances* [39] presents mathematics and Monte Carlo experiments that indicate that the effect of a disturbance on cores is determined by the scale of the disturbance relative to the scale of the core. Disturbances that are much smaller than the scale(s) used to make measurements in the image have little effect, while those comparable to the core scale have significant effect. The reason for this behavior is that the core is extracted from medialness that uses an aperture whose size is proportional to the figural width there. Measurements using this aperture strongly damp variations in intensity at contrasts comparable to or less than that of the figure and at scales distinctly smaller than that of the aperture.

It is the stability of cores against small scale disturbances, together with the way in which medial loci capture figural shape, that leads to its usefulness in image analysis.

Spatial integration of information using an aperture is what allows cores to capture certain relationships much more easily than in other approaches that begin with boundaries measured with a single fixed (usually small) aperture. Consider the example shown in Fig. 4.1. Hierarchical approaches to shape representation that begin with an initial fixed-scale contour [27, 40, 45] perform an initial separation of the figure into interior and exterior, and separately represent each—they can't realize that the notch is a piece missing from the interior instead of a part of the exterior. They represent this shape as two touching rectangles (admittedly, a valid interpretation in some settings) instead of a rectangle with a notch cut out from it. As shown, cores are able to represent both the notch and the overall properties of the object without the notch (the rectangle).

Fig 4.1. The spatial projections of cores of a notched rectangle.

5   CORE-BASED IMAGE ANALYSIS

The invariances of the core and its insensitivities to image disturbances make it a good basis for or aid in various image analysis tasks, including figural shape description and shape change measurement, segmentation, recognition, and registration. Each of these are described briefly in the following paragraphs.

In segmentation the core leads to an approximate object boundary that aids in methods, such as deformable contours, which require a good initial boundary estimate [33, 43]. It also allows the identification of protrusions and subfigures that necessarily begin near this approximate boundary [20].

Also, simultaneous extraction of an object's core and boundary using a model that reflects all of figural, interfigural, and boundary shape can stabilize a deformable contour method [48, 49, 50, 43].

Registration methods involving fiducials are based on either matching corresponding fiducial loci in a target image and a reference image or optimizing an integrated image property on the target image on a fiducial extracted from a reference image [18, 19]. Cores provide a fiducial locus that is insensitive to image disturbances. When the core is chosen as a fiducial, the locus matching requires the measurements of distance between loci in scale space, and the property optimization approach involves the measurement of medialness in the target image on the geometrically transformed locus from the reference image.

Recognition involves measuring how the loci and spatial relations representing shape accord with those of a model. Cores are useful in in characterizing the figural shape in a model and in extracting a model from training images. Shape change measurement and the measurements of shape variability that form a useful part of a model benefit from involving the figural shape measures locally extractable from a sampled core.

Using cores for these tasks requires core-based measurements, loci, and models. In particular, it is useful for the distance between cores and core-based image coordinates to satisfy the conceptual basis of cores, that all distances are relative to object width (core scale). A core-based approximation of the figural boundary, called the *boundary at the scale of the core* (BASOC), is useful to link medial information with boundary and subfigure information. And segmentation and recognition benefits from a means of representing a) structured collections of figures to form object models and b) all aspects of the shape variations of an object from a model or within a family of scenes or images. The computer vision methods based on these ideas and the performance of these methods are reported in separate papers. The mathematics of the necessary core-based measurements, loci, and models are covered in the following sections.

5.1 Distance in Scale Space

The principle of zoom invariance of objects implies that both spatial distances and scale changes must be measured relative to object width. This is the relation stated by equation 5. From this relation geodesic distances in scale space can be determined. The result is that the distance between $(\vec{x}_1, \sigma_1)$ and $(\vec{x}_2, \sigma_2)$

is given by $\dfrac{1}{\rho}\ln\left(\dfrac{\sigma_2}{\sigma_1}\dfrac{1+\sqrt{1-(a\sigma_1)^2}}{-ad+1+\sqrt{1-(a\sigma_1)^2}}\right)$, where $d = \rho|x^2 - x^1|$ (Euclidean norm) and

$$a = \dfrac{2d}{\sqrt{\left(\sigma_1{}^2 - \sigma_2{}^2\right)^2 + d^2\left[d^2 + 2\left(\sigma_1{}^2 + \sigma_2{}^2\right)\right]}}.$$

One use for this distance is in core-based registration methods, where the distance between two cores that are supposed to match can be measured as the integral of interpoint distances between the cores. However, registration in real images requires that these distances be normalized by their variability under the imaging and scene disturbances.

5.2 The Boundary at the Scale of the Core, and Multifigure Objects

The core specifies a medial locus in spatial coordinates, a locus that we call the *core middle*. The core also specifies the radius of a fuzzy n-sphere at each point of the medial locus. If the sphere is made hard and the envelope of these spheres is taken, the result is a locus we call the *boundary at the scale of the core* and abbreviate as the *BASOC*. The core is the Blum medial manifold of the BASOC, so the core is isomorphic to the BASOC. The BASOC gives, at the scale of the core, an approximation to the boundary of the figure represented by the core.

As illustrated in Fig. 5.1, the normal to the BASOC (*a BASOC ray*) meets the core middle at an angle given by $\cos^{-1}\left(\dfrac{dR}{dy}\right) = \cos^{-1}\left(\dfrac{1}{\rho}\dfrac{d\sigma}{dy}\right)$ for a 2D image. For a 3D image the normal to the BASOC meets the core middle surface at angle $\cos^{-1}\left(\dfrac{1}{\rho}|\nabla\sigma|\right)$ in the direction $\nabla\sigma/|\nabla\sigma|$, where the gradient is taken with respect to motion in the core middle surface.
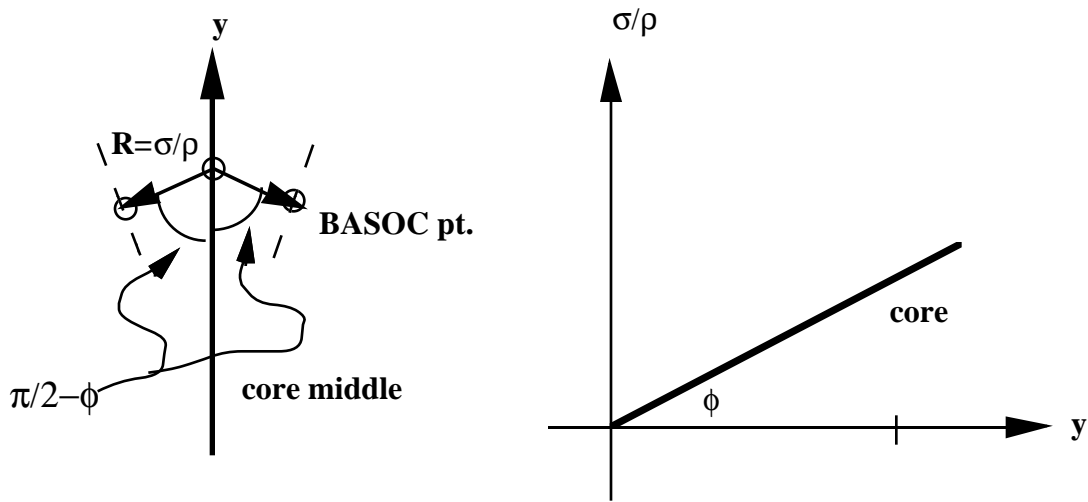
Fig. 5.1. Relation between core middle and BASOC. The coordinate y measures arc length along the core middle.

By this means each point on the core is associated with one (for core endpoints) or two (the normal case) continuous regions of boundary. It has been proven that cores cannot generically branch [8], so generically no core point corresponds to more than two boundary regions.

A full object is typically represented by a collection of cores, one for the main figure and others for subfigures, sub-subfigures, etc. These subfigures can be either protrusions or indentations. The BASOC can be generalized for a collection of figure and subfigure cores by taking the boundary of the envelope of union of the interiors of the BASOCs of the main core and protrusion figures and taking the set subtraction of the union of the interiors of the BASOCs of the indentation figures.

Fritsch [1995b] has shown how subfigure cores can be extracted by using as new stimulation points maxima of medialness as a function of scale and BASOC arc length. Each subfigure core is taken as a child of the core whose BASOC stimulates it. In general, the resulting structure is a directed acyclic graph (see Fig. 5.2 for an example).

McAuliffe [1996] has shown that the BASOC forms an effective starting point for an active contours boundary refinement. Thus core-based segmentation involves first core extraction, then BASOC computation, and finally active contour refinement (see Fig. 5.3).

```
                      global ventricle core

         lower        upper        lower        upper
         left         left         right        right
         horn         horn         horn         horn
         core         core         core         core
```

Fig. 5.2. DAG of brain ventricle computed automatically from single stimulation at center of ventricle in the image in Fig. 5.3.


Fig. 5.3. Cores, BASOC, and evolved boundary of brain ventricle in MR image

Each BASOC point is associated with a single ray from a single core point, defined in Fig 5.1. Actual object boundary points are usefully thought of as being associated with a nearby BASOC point [42]. In practice, continuing the ray from the core middle of the narrowest figure defining the BASOC locally, to its intersection with the boundary is used to associate boundary points with BASOC and core points.

5.3 Core-Based Gibbs Object Models

The recognition of image objects involves starting from a set of training images and using the shape information of an object of interest in these training images to extract the corresponding object in a new image. The core, in combination with a related boundary representation, provides a powerful means of

representing the statistics of object shape from the training images [48, 49, 50, 43] and for extracting the object in this training phase. Medialness is a useful measure in the recognition and segmentation phase.

In this approach, the shape of an object is taken to be characterized by the three aspects of shape listed early in this paper:

1) figural shape, each given by a connected, unbranching medial locus;

2) interfigural shape, given for a protrusion or indentation subfigure by the positional, angular, and size relation of its medial locus endpoint near the parent BASOC and a medial locus point in the parent core;

3) boundary shape, given by each point on the object boundary relative both to adjacent boundary points and to the medial point with which it is associated.

Wilson has shown how all of these can be captured by a Markov random field, with each site being a point on a medial locus or a boundary point and the random variables at each site being the location in scale space of the corresponding point (see Fig. 5.4). Corresponding to the three aspects of shape, five types of neighbor relations, or links, are established:

1) The relation between a medial point and an adjacent medial point on the same figure's medial locus. The scale space vector between the locations of these neighbors characterizes the local figural shape.

2) The relation between a medial point of one figure and the proximal endpoint of its subfigure's medial locus. The scale space vector between the locations of these neighbors, and the angle of the first link from the subfigure endpoint relative to the links from the parent figure medial point characterize the position, width, and orientation of the proximal end of the protrusion or indentation relative to the parent figure there.

3) The relation between a medial point and either of the two corresponding boundary points along the associated BASOC rays. The difference between the width in the medial point and the length of the image space vector between the locations of these neighbors characterizes the the level of detail or texture of the boundary there.

3) The relation between a medial endpoint and the corresponding boundary vertex. The scale space vector between the locations of these neighbors characterizes the local shape of the boundary at

the vertex relative to the core. The importance of vertices in shape characterization in human vision has been emphasized by Hoffman [1984] and by Leyton [1992].

4) The relation between adjacent boundary points. This is the relation commonly used in deformable contours methods in computer vision.


Fig 5.4. Sites and scale space vectors whose statistics define a brain ventricle. [Courtesy of Liyun Yu]

The approach to recognition described by Wilson [1995a,b,c] and by Pizer [1996]is Bayesian, with the prior containing the statistics on the relations just listed as derived from the family of training images. Wilson and Johnson have shown how a Gibbs distribution can be used for this prior, with potentials measuring the variance-normalized scale space squared distance between the vector between neighbors and the vector between the corresponding means for the families :

$$log\, p(\Theta) = log\, Z - \sum_{i \sim j} U\!\left(\theta_i - \theta_j, \theta_j, \hat{\theta}_i - \hat{\theta}_j, \hat{\theta}_j\right),$$

where Θ is the set of sites, $\hat{\Theta}$ is the corresponding set of sites in the model, Z is a normalizing constant, each θ_i represents a site, $\sim$ represents the neighbor relationship, - represents a difference in image space or scale space, as appropriate, and the potentials U measure the scale space difference between the interneighbor link $\theta_i - \theta_j$ and the corresponding link in the model relative to the standard deviation of this link in the training set. In the case of interfigural links, the scale space difference is relative to properties at the parent site. Cores are used to extract the medial loci used in the training stage.

[43] discusses how recognizing and segmenting an object in a new image, as well as object -based registration of a pair of images, can be done by optimizing the posterior of linked medial and boundary loci. The log likelihood function used in this deformable loci expresses the likelihood of the image as a sum of the medialness at the medial points C, boundariness at the boundary points B, and vertexness at the vertex points V:

$$\log p(I \mid \Theta) = k_C \sum_{i \in C} M(\theta_i) + k_B \sum_{i \in B} B(\theta_i) + k_V \sum_{i \in V} V(\theta_i).$$

Such posterior optimization has produced a stable, automatic recognition and segmentation in early trials on ventricles in MR images [43]. The power of this variant of a deformable contours technique [e.g., 26; 47, 6, 46, 34] is that medial loci's ability to capture figural shape and the stability of medialness against image disturbances, causes a stabilization of the localization of the boundary, relative to methods that are purely boundary based, while retaining the advantages of previous methods in the boundary being measured at small scale when the image information supports that. It also provides a means of identifying where and how much the resulting boundary deviates from the model shape.

The result of all the image analysis methods discussed in this section, as applied to medical images, is covered in [17, 18, 19, 20, 21] and [43].

SUMMARY

We have argued that the invariance requirements commonly associated with the concept of shape, the use of a medial approach to figural shape description, and a recognition that measurement aperture size is an essential aspect of image analysis lead to the core as the representation of figural shape. After giving the mathematics associated with the definition and extraction of cores, we sketched why and how the core is stable against smaller-scale image disturbances. We then sketched ways in which cores are proving to be useful for image analysis tasks, and we gave mathematics relevant to core-based image analysis.

ACKNOWLEDGMENTS

APPENDIX: THE GEOMETRY OF SCALE SPACE AND CORES

Let scale space be denoted by $\Re^n \times \Re^+$ with typical element denoted by $\overset{\perp}{\xi} = (\overset{\Upsilon}{x}, \sigma)$. For indexing purposes, we have $\xi_i = x_i$ for $1 \le i \le n$ and $\xi_{n+1} = \sigma$.

A.1 Tensor Notation

Tensor calculus is used throughout. Indices are used both for denoting the components of a tensor (scalar, vector, matrix, etc.) and for denoting derivatives with respect to independent variables ξ_i. Subscripted indices are used for covariant components and superscripted indices are used for contravariant components. If a tensor quantity contains covariant derivatives of indexed quantities, a comma is used to separate the subscripted indices from derivatives. For example, if T is a zero-order tensor (depending on ξ), then $T_{,i}$ is the first-order covariant tensor whose components are the derivatives of T with respect to ξ_i. The covariant derivative of the first-order covariant tensor U_i is the second-order covariant tensor $U_{i,j}$. The covariant derivative of the first-order contravariant tensor V^i is denoted $V^i_{,j}$.

The Einstein summation convention will be used. If an index occurs twice in a term, once as a subscript and once as a superscript, it indicates a summation over the repeated index. Occurrence of an index more than two times in a term is not allowed, and contraction of two subscripted indices or of two superscripted indices is not allowed. For example, if a_{ij} is a second-order covariant tensor and b^j is a first-order contravariant tensor, then $c_i = a_{ij}b^j$ is a first-order covariant tensor obtained by summing over the j index.

A.2 Metric Tensor

The metric (eq. 5) can be written in tensor form as $ds^2 = g_{ij}d\xi^i d\xi^j$, where g_{ij} is called the *metric tensor*.

It is a diagonal matrix whose first n diagonal entries are $1/\sigma^2$ and whose last diagonal entry is $1/(\rho\sigma)^2$.

In contrast, the metric tensor for standard Euclidean space is $g_{ij} = \delta_{ij}$, the Kronecker delta which is 1 if

both indices are equal, but 0 otherwise. The inverse metric tensor is denoted by g^{ij} and has the property

$g_{ik}g^{kj} = \delta_i^j$ The metric can be used to raise or lower indices through contraction. If u_i is a covariant

tensor and u^i is its contravariant counterpart, they are related by $u^i = g^{ij}u_j$ (raising an index) and

$u_i = g_{ij}u^j$ (lowering an index).

A.3 Christoffel Symbols

The analysis involves tensor quantities which need to be differentiated. In general, the partial derivative

of a tensor is not necessarily a tensor. Tensor differentiation requires the use of some nontensorial

objects called Christoffel symbols, $\Gamma_{ij}^k = \frac{1}{2} g^{kl}\left(\frac{\partial g_{jl}}{\partial x_i} + \frac{\partial g_{il}}{\partial x_j} - \frac{\partial g_{ij}}{\partial x_l}\right)$. Let $\overset{r}{e}_k \in \Re^{n+1}$ denote the $(n+1)\times 1$

unit length vector whose components are all 0 except for the k^{th} component which is 1. For the scale

space metric, define the matrix $\Gamma^k = [\Gamma_{ij}^k]$; then $\Gamma^k = -\frac{1}{\sigma}\begin{Bmatrix} \overset{r}{e}_k \overset{r}{e}_{n+1}^t + \overset{r}{e}_{n+1}\overset{r}{e}_k^t, & 1 \le k \le n \\ \overset{r}{e}_{n+1}\overset{r}{e}_{n+1}^t - \rho^2 \sum_{l=1}^n \overset{r}{e}_l\overset{r}{e}_l^t, & k = n+1 \end{Bmatrix}$. For

example, if $n = 1$, $\quad \Gamma^1 = -\frac{1}{\sigma}\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ and $\Gamma^2 = -\frac{1}{\sigma}\begin{bmatrix} -\rho^2 & 0 \\ 0 & 1 \end{bmatrix}$, and if $n = 2$,

$\Gamma^1 = -\frac{1}{\sigma}\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \Gamma^2 = -\frac{1}{\sigma}\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \Gamma^3 = -\frac{1}{\sigma}\begin{bmatrix} -\rho^2 & 0 & 0 \\ 0 & -\rho^2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. Each Γ^k is a symmetric matrix.

A.4 Riemann Tensor and Curvature

The implication of using the scale space metric is that the geometry is Riemannian. The Riemann tensor

of the second kind is defined by $R_{jkl}^i = \frac{\partial \Gamma_{jl}^i}{\partial x_k} - \frac{\partial \Gamma_{jk}^i}{\partial x_l} + \Gamma_{jl}^m\Gamma_{mk}^i - \Gamma_{jk}^m\Gamma_{ml}^i$. The Riemann curvature relative

to the metric g_{ij} is defined for each pair of contravariant vectors u^i and v^i as $K\left(\overset{r}{\xi};u,v\right) = \frac{R_{ijkl}u^i v^j u^k v^l}{G_{ijkl}u^i v^j u^k v^l}$,

where $R_{ijkl} = g_{im}R_{jkl}^m$ and $G_{ijkl} = g_{ik}g_{jl} - g_{il}g_{jk}$. It can be shown that the only nonzero independent

components are $R_{ijij} = -\rho^2/\sigma^4$ and $G_{ijij} = 1/\sigma^4$, for $1 \le i < j \le n+1$. The other components are either

0 or are determined by the values of the terms mentioned above. Consequently, the Riemann curvature

of the space is identically a constant, $K = -\rho^2$. When $\rho = 1$, this particular Riemannian geometry is called *hyperbolic geometry* and has been studied extensively in the differential geometry literature.

A.5 Covariant and Scale Space Derivatives

Covariant differentiation provides the generalization of partial differentiation to scale space. The covariant derivative of a tensor is itself a tensor. The scale space derivatives we define are dimensionless quantities but turn out not to be tensors. Define the non-tensorial quantity a_{ij} to be the square-root of the metric tensor g_{ij}. The first n diagonal entries of a_{ij} are $1/\sigma$, and the last diagonal entry is $1/(\rho\sigma)$. Similarly, define the non-tensorial quantity a^{ij} to be the inverse of a_{ij}. Thus, $g_{ij} = a_{ki}a_{kj}$, $g^{ij} = a^{ki}a^{kj}$, and $a_{ki}a^{kj} = a^{ki}a_{kj} = \delta_{ij}$, where the summation convention is used despite the terms not being tensors. The idea is that the metric can be written symbolically as $ds^2 = g_{ij}d\xi^i d\xi^j = \left(a_{ki}d\xi^i\right)\left(a_{kj}d\xi^j\right)$, where the summation convention is still used. The terms $a_{ki}d\xi^i$ are the natural dimensionless 1-forms used for scale space measurements.

Let $f\!\left(\overset{\perp}{\xi}\right)$ be a smooth real-valued function. The calculation of cores will require third-order derivatives. The first covariant derivative of f is given by $f_{,i}$. The total derivative of f is $df = f_{,i}d\xi_i = f_{,i}a^{ki}a_{kj}d\xi_j = \left(a^{ki}f_{,i}\right)\left(a_{kj}d\xi_j\right)$, so the dimensionless *scale space first derivatives* are defined by $\hat{D}f_i = a^{ik}f_{,i}$, which is not a tensor. The covariant derivative of $f_{,i}$ is given by $f_{,ij} = \dfrac{\partial f_{,i}}{\partial \xi^j} - \Gamma_{ij}^k f_{,k}$, where the Christoffel symbols were defined earlier. The dimensionless *scale space second derivatives* are defined by $\hat{D}^2 f_{ij} = a^{ik}a^{jl}f_{,ij}$, which is not a tensor. The covariant derivative of $f_{,ijk} = \dfrac{\partial f_{,ij}}{\partial \xi^k} - \Gamma_{ik}^l f_{,lj} - \Gamma_{jk}^l f_{,il}$, and the dimensionless *scale space third derivatives* are defined by $\hat{D}^3 f_{ijk} = a^{li}a^{mj}a^{nk}f_{,lmn}$, which is not a tensor.

Define λ_{ij} to be the matrix whose diagonal values are the ordered eigenvalues of $f_{,ij}$, say $\lambda_1 \le \lambda_2 \le ... \le \lambda_n$, and whose off-diagonal values are zero. Let v_{ij} be the matrix whose columns are the covariant eigenvectors of $f_{,ij}$. The contravariant eigenvectors are obtained by raising an index, $v_j^i = g^{ik}v_{kj}$. Additionally assume that the selected eigenvectors form an orthonormal system with respect to the metric. That is, $v_{ki}v_j^k = \delta_{ij}$. The eigensystems are collectively represented by $f_{,ij}v_k^j = v_{il}\Lambda_k^l$.

35

A.6 Maximum Convexity Ridges of Medialness

We define the maximum convexity ridge of dimension $n-1$ since that provides the $n-1$-dimensional core manifold that is appropriate for general objects in $\Re^n$. The definition is easily restated for a ridge of arbitrary dimension d by modifying the number of dimensions in which derivative constraints must hold from 2 to $n+1-d$. Define the directional derivatives $P_j = v_j^i f_{,i}$ for $1 \le j \le n+1$. The n-dimensional core points are those ridge points of medialness, $(\overset{\perp}{x}, \sigma)$, for which f has a local maximum when restricted to the 2-dimensional geodesic surface whose tangents at $(\overset{\perp}{x}, \sigma)$ are the first two eigenvectors v_1^i and v_2^i [12]. This is equivalent to finding those points $(\overset{\perp}{x}, \sigma)$ such that $P_1(\overset{\perp}{x}, \sigma) = 0$; $P_2(\overset{\perp}{x}, \sigma) = 0$, $\lambda_1(\overset{\perp}{x}, \sigma) \le \lambda_2(\overset{\perp}{x}, \sigma) < 0$. We are solving 2 equations in $n+1$ variables, so we generally expect the solutions to lie on $n-1$-dimensional manifolds.

The tests $P_1 = 0$ and $P_2 = 0$ implicitly assume the continuity of eigenvectors v_1^i and v_2^i, which is true when λ_1 and λ_2 are distinct eigenvalues, but may not hold otherwise. This problem can be circumvented, but the details are very technical and tedious, so are not presented here [11]. The basic idea is to restrict our attention to regions for which $\lambda_2 < \lambda_3$ to avoid eigenvector swaps between the first 2 eigenspaces and the last $n-1$ eigenspaces. The eigenvectors v_j^i can be replaced by smoothly varying orthonormal eigenvectors u_j^i for which (1) the span of u_1^i and u_2^i is equal to the span of v_1^i and v_2^i and (2) the span of u_3^i through u_{n+1}^i is equal to the span of v_3^i through v_{n+1}^i. Moreover, the vectors can be chosen so that (1) the derivatives $u_{j,k}^i$ for $j = 1,2$ are linear combinations of the u_j^i for $3 \le j \le n+1$ and (2) the derivatives of $u_{j,k}^i$ for $3 \le j \le n+1$ are linear combinations of the u_j^i for $j = 1,2$. Defining $Q_j = u_j^i f_{,i}$, the tests $P_1 = 0$ and $P_2 = 0$ can be replaced by $Q_1 = 0$ and $Q_2 = 0$. Core calculation requires computing derivatives of the Q_j. These derivatives are used for initially locating a core point and traversing the core by following tangents to the core. In [11] it is shown that the direction of continuous flow to the ridge can be calculated using only the eigenvalues and eigenvectors, despite the discontinuities that occur in the eigenvectors. It is also shown that continuous tangent vectors can be calculated using only the eigenvalues and eigenvectors. The discontinuity of eigenvectors is captured by a single coefficient in the tangent vector equations. This coefficient is the determinant of the matrix $\left[v_j^i \right]$ and is either 1 or -1. In the numerical calculations it is necessary to keep track of the determinant

36

and adjust the computed tangent vectors in the event the determinant is -1. In light of this result, in the next sections we will treat the eigenvectors v_j^i as if they have the properties assigned to the smooth u_j^i, to simplify the development.

A.7 Flow to and Along a Core

Differentiating P and Q yields $\qquad P_{,k} = u^i f_{,ik} + u^i_{,k} f_{,i} = \alpha u_k + R a_k, \quad Q_{,k} = v^i f_{,ik} + v^i_{,k} f_{,i} = \beta u_k + R b_k.$

The tensors a_k and b_k are determined by the eigensystem for w. Differentiate this eigensystem to obtain

$f_{,ij} w^i_{,k} + f_{,ijk} w^j = \gamma w_{i,k} + \gamma_{,k} w_i$. Substitute for $w^i_{,j} = -a_j u^i - b_j v^i$, $w_{i,j} = -a_j u_i - b_j v_i$ to obtain

$(\alpha - \gamma) u_i a_k + (\beta - \gamma) v_i b_k = f_{,ijk} w^j - \gamma_{,k} w_i$. Contracting with u^i and v^i, then solving for a_k and b_k yields

$a_k = \dfrac{1}{\alpha - \gamma} f_{,ijk} u^i w^j, \quad b_k = \dfrac{1}{\beta - \gamma} f_{,ijk} v^i w^j.$ Substituting back into the gradients for P and Q yields

$P_{,k} = \alpha u_k + \dfrac{R}{\alpha - \gamma} f_{,ijk} u^i w^j, \quad Q_{,k} = \beta v_k + \dfrac{R}{\beta - \gamma} f_{,ijk} v^i w^j.$

Our core extraction algorithms all consist of finding an approximation in scale space to a core point, moving to the ridge of medialness that is the core, and then traversing the core. These formulas will be used both in finding an initial core point and in determining the core direction.

A.7.1 Finding a Core Point from an Approximation

Given an approximation $\vec{u}_0 = (\vec{x}_u, \sigma_u)$ to a core point, finding the core point involves solving the equation $P(\vec{u}) = 0, Q(\vec{u}) = 0$ subject to the constraints $\lambda_1(\vec{u}) \le \lambda_2(\vec{u}) < 0$. Recall that P and Q are defined in terms of the vectors $v^1(\vec{u})$ and $v^2(\vec{u})$. The numerical approach is then to take a step to a root of $P(\vec{u}) = 0, Q(\vec{u}) = 0$ in the space spanned by $v^1(\vec{u})$ and $v^2(\vec{u})$ for the $\vec{u}$ that is the present approximation to the core point, and then to iterate this process. We have found that the stability of a root-trapping method is necessary, so a method based on bisection is used for this root trapping. However, a variety of other root finding methods have been tried, and some are also in use.

A.7.2 Traversing the Core by Integration of the Differential Equation

Let $\overset{\perp}{u}$ be a core point. To find successive points on the core, we can use the fact that the core is a solution to a differential equation, or we can take a step in a direction approximately along the core and then return to the core using the method described in section A.7.1.

Let us first derive the differential equation for which the core is the solution. If $\overset{\mathrm{r}}{T}\left(\overset{\perp}{\xi}\right)$ is a tangent vector to the core, the core can be traversed by solving a system of ordinary differential equations, $\dfrac{d\xi_i(t)}{dt} = \pm T_i(\xi(t)), \quad \xi_i(0) = u_i$, where two traversals are required, one in the direction $+\overset{\mathrm{r}}{T}\left(\overset{\perp}{\xi}\right)$ and one in the direction $-\overset{\mathrm{r}}{T}\left(\overset{\perp}{\xi}\right)$. The problem is now to construct a closed formula for the ridge direction $\overset{\perp}{T}$.

The core curve is the intersection of surfaces implicitly defined by $P(\xi) = 0$ and $Q(\xi) = 0$. The curve direction is therefore a vector which is orthogonal to both surface normals. The contravariant vectors $N^i = g^{ij} P_{,j}$ and $M^i = g^{ij} Q_{,j}$ are surface normals, so a core direction is the cross product of the two vectors, $T_i = e_{ijk} N^j M^k$, where e_{ijk} is the permutation tensor defined by $e_{123} = e_{231} = e_{312} = 1, e_{132} = e_{213} = e_{321} = -1, 0$ otherwise. If the sign of the tangent vectors is handled carefully, the core tangent is a continuous function despite discontinuities in the eigenvectors.

A.7.3 Traversing the Core by Approximation of Ridge Tangent and Root Finding

We have found that it frequently is more efficient and more stable to use the fact that typically the third eigenvector of the Hessian of medialness approximately points along the core so a step proportional to the extracted radius at a core point (proportional to the extracted scale) produces an approximation to a core point. From this approximation, the method of section A.7.1 can be applied to return to the core.

A.7.4 Traversing the Core by Approximation of Ridge Tangent and Use of Gradient Orthogonality

Expressing the gradient of M in terms of the eigenvectors of $D^2 M$ leads to $DM = Pu + Qv + Rw$. It follows that $\left(D^2 M - \gamma I\right) DM = (\alpha - \gamma) Pu + (\beta - \gamma) Qv$. Thus the ridge point can be found as a zero of $\left(D^2 M - \gamma I\right) DM$, which can be computed by minimizing $DM^t \left(D^2 M - \gamma I\right) DM$ and assuring that the minimum value in magnitude is below some threshold.

All of the methods that we have programmed for core extraction calculate medialness derivatives using $n+1$-dimensional splines fit to local medialness values. Only medialness values needed for this purpose, i.e., in the neighborhood of the core, are evaluated.

A.8 2D Maximum Convexity Ridges of Medialness from 3D Images

Let $f_{,ij}u^j = \alpha u_i$, $f_{,ij}v^j = \beta v_i$, $f_{,ij}w^j = \gamma w_i$, and $f_{,ij}\phi^j = \delta\phi_i$, where $\alpha \le \beta < \gamma \le \delta$.

$$
\begin{bmatrix} u^i_{,j} \\ v^i_{,j} \\ w^i_{,j} \\ \phi^i_{,j} \end{bmatrix} = \begin{bmatrix} 0 & 0 & a_j & c_j \\ 0 & 0 & b_j & d_j \\ -a_j & -b_j & 0 & 0 \\ -c_j & -d_j & 0 & 0 \end{bmatrix} \begin{bmatrix} u^i \\ v^i \\ w^i \\ \phi^i \end{bmatrix}
$$
for some choice of continuous tensors a_j, b_j, c_j, and d_j. Define

$P = u^i f_{,i}$, $Q = v^i f_{,i}$, $R = w^i f_{,i}$., and $S = \phi^i f_{,i}$. Core points satisfy $P = 0$, $Q = 0$, and $\alpha \le \beta < 0$.

Differentiating P and Q yields

$$
P_{,k} = u^i f_{,ik} + u^i_{,k} f_{,i} = \alpha u_k + Ra_k + Sc_k, \quad Q_{,k} = v^i f_{,ik} + v^i_{,k} f_{,i} = \beta u_k + Rb_k + Sd_k.
$$

The tensors a_k, b_k, c_k, and d_k are determined by the eigensystem for w and ϕ. Differentiate the eigensystems to obtain $f_{,ij}w^i_{,k} + f_{,ijk}w^j = \gamma w_{i,k} + \gamma_{,k}w_i$, $f_{,ij}\phi^i_{,k} + f_{,ijk}\phi^j = \delta\phi_{i,k} + \delta_{,k}\phi_i$. Substitute for $w^i_{,j} = -a_j u^i - b_j v^i$, $w_{i,j} = -a_j u_i - b_j v_i$, $\phi^i_{,j} = -c_j u^i - d_j v^i$, and $\phi_{i,j} = -c_j u_i - d_j v_i$ to obtain

$(\alpha - \gamma)u_i a_k + (\beta - \gamma)v_i b_k = f_{,ijk}w^j - \gamma_{,k}w_i$, $(\alpha - \delta)u_i a_k + (\beta - \delta)v_i b_k = f_{,ijk}\phi^j - \delta_{,k}\phi_i$. Contracting with

u^i and v^i, then solving yields

$$
a_k = \frac{1}{\alpha - \gamma} f_{,ijk}u^i w^j, \quad b_k = \frac{1}{\beta - \gamma} f_{,ijk}v^i w^j, \quad c_k = \frac{1}{\alpha - \delta} f_{,ijk}u^i \phi^j, \quad d_k = \frac{1}{\beta - \delta} f_{,ijk}v^i \phi^j.
$$

Substituting back into the gradients for P and Q yields

$$
P_{,k} = \alpha u_k + \frac{R}{\alpha - \gamma} f_{,ijk}u^i w^j + \frac{S}{\alpha - \delta} f_{,ijk}u^i \phi^j, \quad Q_{,k} = \beta v_k + \frac{R}{\beta - \gamma} f_{,ijk}v^i w^j + \frac{S}{\beta - \delta} f_{,ijk}v^i \phi^j.
$$

These formulas will be used both in finding an initial core point and in determining the core tangent spaces.

The 2-dimensional tangent space of a 2D-from-3D core is determined by the system of differential equations defining steps in the tangent space to the core. We restrict the expressions to the maximum convexity ridge, though the generalization is similar for the optimal scale ridge. The core surface is the intersection of two 3-dimensional surfaces implicitly defined by $P(\xi) = 0$ and $Q(\xi) = 0$. The core

tangent space is therefore the orthogonal complement (in $\Re^4$) of the space spanned by the normals $N^i = g^{ij} P_{,j}$ and $M^i = g^{ij} Q_{,j}$. Define $T_{ij} = \det(V) e_{ijkl} N^k M^l$, where e_{ijkl} is the permutation tensor on four symbols and V is the matrix of eigenvectors of $D^2 M$ so $\det(V) = \pm 1$.

The four eigenvectors of the matrix $T = T_{ij}$ give two normals and two tangents to the core. The two tangents correspond to the nonzero eigenvalue of multiplicity 2.

REFERENCES

Blum, H (1967). A transformation for extracting new descriptors of shape, in W. Wathen-Dunn, ed., *Models for the Perception of Speech and Visual Form*, MIT Press, Cambridge, MA: 363-380.

Blum, H, RN Nagel (1978). Shape description using weighted symmetric axis features, *Patt. Recog.*, **10**: 167-180.

Burbeck, CA & Pizer, SM (1995a). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research* **35**(13):1917-1930.

Burbeck, CA, SM Pizer, BS Morse, D Ariely, GS Zauberman, JP Rolland (1996). Linking object boundaries at scale: a common mechanism for size and shape judgments.*Vision Research*, **36**(3):361-372.

Canny, J (1987). A computational approach to edge detection. *IEEE Trans. PAMI*, **PAMI-6**: 679-698.

Chaney, EL, V Johnson, D Fritsch, L Yu, SM Pizer (1966). Core-based Bayesian segmentation of objects in medical images. In preparation, UNC Dept. of Radiation Oncology.

Crowley, JL & AC Parker (1984). A representation of shape based on peaks and ridges in the difference of low-pass transform. *IEE Trans. PAMI* **6**(2): 156-170.

Damon, JN (1995). *Properties of Ridges and Cores for Two Dimensional Images*. Tech. Report, Dept. of Mathematics, University of North Carolina at Chapel Hill.

Eberly, D, RB Gardner, BS Morse, SM Pizer, C Scharlach (1994a). Ridges for image analysis, *Journal of Mathematical Imaging and Vision,* **4**: 351-371.

Eberly, D (1994b). A differential geometric approach to anisotropic diffusion. Chapter 14 in *Geometry-Driven Diffusion in Computer Vision*, Bart ter Haar Romeny ed., Kluwer Publishers.

Eberly, D (1994c). Fast algorithms for ridge construction, to appear in *Proceedings SPIE Photonics East 1994: Vision Geometry III* , **2356**.

Eberly, D (1994d). *Geometric Methods for Analysis of Ridges in N-Dimensional Images*. Ph.D. Dissertation, Dept. of Computer Science, University of North Carolina at Chapel Hill.

Florack, L (1994). *The Syntactic Structure of Scalar Images*,. Ph.D. dissertation, University of Utrecht, Netherlands, 3D Computer Vision Group.

Fritsch, DS (1992). Medial description of gray-scale image structure by gradient-limited diffusion. *Visualization in Biomedical Computing 1992*, R Robb, ed. SPIE **1808**: 105-117.

Fritsch, DS (1993). Registration of Radiotherapy Images using Multiscale Medial Descriptions of Image Structure. Ph.D. dissertation, Department of Biomedical Engineering, University of North Carolina.

Fritsch, DS, SM Pizer, BS Morse, DH Eberly, A Liu (1994a). The multiscale medial axis and its applications in image registration. *Pattern Recognition Letters* , **15**:445-452.

Fritsch, DS, SM Pizer, EL Chaney, A Liu, S Raghavan, T Shah (1994b). Cores for image registration. *Medical Imaging '94: Image Processing*, SPIE **2167**:128-142.

Fritsch, DS, EL Chaney, A Boxwala, MJ McAuliffe, S Raghavan, A Thall, JRD Earnhart (1995a). Core-based portal image registration for automatic radiotherapy treatment verification. To appear, *International Journal of Radiation, Oncology, Biology, Physics.*

Fritsch, D, D Eberly, SM Pizer (1995b). Stimulated cores and medical applications, poster and short paper, IPMI '95: *Information Processing in Medical Imaging*, Y Bizais, C Barillot, R DiPaola, eds. Kluwer Series in Computational Imaging and Vision, Dordrecht: 365-368.

Fritsch, DS, SM Pizer, D Eberly, M McAuliffe, A Liu, L. Yu, J Furst, C Gu, DT Chen (1996). Medical applications of core-based image analysis, in preparation.

Furst, JD, SM Pizer (1996). Marching ridges: a method for extracting mD ridges in nD spaces. Tech. Report. UNC Department of Computer Science. Submitted revision of Furst, JD, SM Pizer, DH Eberly, Marching cores: a method for extracting cores from 3D medical imges, Proc. Workshop on Math. Methods in Biomed. Image. Anal., IEEE Cat# 96TB100056: 124-130.

Furst, JD (1996). Zoom-invariant kernels and Gaussian derivatives. Tech. Report TRxxx UNC Department of Computer Science.

Haralick, R (1983). Ridges and valleys in digital images. *Computer Vision, Graphics, and Image Processing* **22**: 28-38.

Hoffman, DD, WA Richards (1984). Parts of recognition. *Cognition*, **18:** 65-96.

Kass, M, A Witkin, D Terzopoulos (1987). Snakes: Active contour models. *Proc. 1st ICCV*. IEEE Catalog #84CH2465-3: 259-268.

Kimia, BJ, AR Tannenbaum, SW Zucker (1992). Shapes, shocks, and deformations I: the components of shape and the reaction-diffusion space. Technical Report LEMS-105, Division of Engineering, Brown University.

Koenderink, JJ (1990). The brain a geometry engine. *Psychol. Res.*

Lee, TS (1995). Neurophysiological evidence for image segmentation and medial axis computation in primate V1. Dept of Brain and Cognitive Sciences, MIT; submitted to 4th Comp'l Neurosci. Mtg.

Leyton, M. (1992). *Symmetry, Causality, Mind*, The MIT Press, Cambridge, Massachusetts.

Lindeberg, T (1994a). *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Dordrecht, International Series in Engineering and Computer Science.

Lindeberg, L & BM ter Haar Romeny (1994b). Linear scale space, Chapters 1, 2 in ter Haar Romeny, B *Geometry-Limited Diffusion*, Kluwer Series on Computational Imaging and Vision and , Dordrecht.

Lorensen, WE, HE Cline (1987). Marching Cubes: A High Resolution 3D Surface Construction Algorithm, *Computer Graphics*, **21**(4),: 163-169.

Morse, BS, SM Pizer, AL Liu (1993). Multiscale medial analysis of medical images, in *Information Processing in Medical Imaging (IPMI'93)*, HH Barrett and AF Gmitro, eds. Lecture Notes in Computer Science **687**:112—131, Springer-Verlag. In revised form in *Image & Vision Computing*, **12:6**:327-338.

Morse, BS, SM Pizer, DS Fritsch (1994a). Robust object representation through object-relevant use of scale, in *Proceedings SPIE Medical Imaging 94: Image Processing* **2167**:104—115, SPIE.

Morse, BS, SM Pizer CA Burbeck (1994b). General shape and specific detail: context-dependent use of scale in determining visual form, in *Aspects of Visual Form Processing*, C Arcelli, LP Cordella, G Sanniti di Baja, eds. 374-383. World Scientific.

Morse, BS (1994c). *Computation of Object Cores from Greyscale Images*. Ph.D. Dissertation, Department of Computer Science, University of North Carolina at Chapel Hill.

Morse, BS, SM Pizer, DT Puff, C Gu (1995). Zoom-invariant vision of figural shape: effects on cores of image disturbances. Submitted to *Computer Vision and Image Understanding* as a companion paper to this one.

Ogniewicz, RL (1992). *Discrete Voronoi Skeletons*, Ph.D. Dissertation, ETH Zurich.

Pizer, SM, AE Todd-Pokropek (1976). Noise Character in Processed Scintigrams. *Information Processing in Scintigraphy* , edited by C. Raynaud and A. Todd-Pokropek. Commissariat a l'Energie Atomique, Dept. de Biologie, Service Hospitalier Frederic Joliot, Orsay, France, March 1976 :1-16. Also in *Biomedical Computing*, W. J. Perkins, ed., Pitman Medical Publishing Co., Ltd., London, 1977: 198-213.

Pizer, SM, CA Burbeck, JM Coggins, DS Fritsch, BS Morse (1994a). Object shape before boundary shape: scale-space medial axes. Presented at *Shape in Picture,* (NATO Advanced Research Workshop), *Journal of Mathematical Imaging and Vision*, **4**:303-313.

Pizer, SM, S Murthy, DT Chen (1994b). Core-based boundary claiming. *Medical Imaging '94: Image Processing*, SPIE **2167**:151-159.

Puff, D, (1995) Human vs. Model Performance for Medical Image Estimation Tasks. Ph.D. dissertation, Dept. of Biomedical Engineering, University of North Carolina.

Rom, H, G Medioni (1993). Hierarchical decomposition and axial shape description, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15**(10):973—981.

Caselles, V, R Kimmel, G Sapiro (1995). Geodesic active contours. *Proc. 5th ICCV*. IEEE Catalog #95C35744: 694-699.

Tek, H, BB Kimia (1995). Image segmentation by reaction-diffusion bubbles. *Proc. 5th ICCV*. IEEE Catalog #95C35744: 156-162.

Whitaker, RT (1993). *Geometry-Limited Diffusion*. Ph.D. Dissertation, Dept. of Computer Science, University of North Carolina at Chapel Hill.

Whitaker, RT, DT Chen (1994). Embedded active surfaces for volume visualization. *Proc. Medcial Imaging 1994: Image Processing*, SPIE **2167**: 340-352.

Wilson, A (1995a). Statistical Models for Shapes and Deformations. Ph.D. dissertation, Institute of Statistics and Decision Sciences, Duke University

Wilson, A and VE Johnson (1995b). Models for shape deformation, to appear in *Bayesian Statistics* **5**, edited by J. Berger, J. Bernardo, P. Dawid, and A. Smith.

Wilson, AG, VE Johnson, L Yu, SM Pizer. (1995c). Scale-Space Image Models for Shape Deformation, invited paper to appear in *Computing Science and Statistics: Proceedings of the 27th Symposium on the Interface*, Springer-Verlag, New York.

Zucker, SW, A Dobbins, L Iverson (1989) Two stages of curve detection suggest two styles of visual computation. *Neural Computation***1**: 68 - 81.