Stimulated Cores and their Applications in Medical Imaging

Daniel S. Fritsch[1,2] , David Eberly[3], and Stephen M. Pizer[2,3,4]

[1]Department of Radiation Oncology, [2]Department of Biomedical Engineering,
[3]Department of Computer Science, [4]Department of Radiology
The University of North Carolina at Chapel Hill

Dr. Daniel S. Fritsch
Dept. of Radiation Oncology
Campus Box #7512
University of North Carolina School of Medicine
Chapel Hill, NC 27599-7512 USA

Dr. David Eberly
Dept. of Computer Science
Campus Box #3175
University of North Carolina
Chapel Hill, NC 27599-3175 USA

Dr. Stephen M. Pizer
Dept. of Computer Science
Campus Box #3175
University of North Carolina
Chapel Hill, NC 27599-3175 USA

ABSTRACT:

Representing geometric properties of objects in medical images independent of object size necessarily requires a multiscale, or scale space, analysis. We describe means called *cores* for representing middle and width properties of greyscale image objects via scale space measurements and describe an efficient algorithm for their computation called *stimulated cores*. We present applications of stimulated cores in a variety of medical imaging tasks.

1. INTRODUCTION

Multiscale methods of image analysis have received increasing attention in recent years due in part to the understanding that in order to measure geometric properties of image objects independent of object size (a property we call *zoom invariance*), one must make those geometric measurements at a range of spatial scales.[†] Consequently, one of the critical issues in multiscale image analysis involves selecting from a scale space of geometric measurements those scales at which certain features are best represented; that is, choosing the *interesting*, object-relevant scales.

A multiscale method of object representation we call *cores* allows for the selection of special places in scale space where object middles are optimally manifested.[‡] Cores are computed directly from the greyscale image intensities by first deriving a measurement called *medialness* at each position in the image and at a range of scales, and then finding *scale space ridges* in the medialness function. Operationally, medialness is defined as the degree to which a given position behaves like an object middle when measured at a particular scale. Medialness is derived from the linkage of boundariness measurements across image space, where in order to provide for zoom invariance the linkage distance is proportional to the scale at which the boundariness measurements are made. Hence, medialness is always associated with a width (the linkage distance) proportional to the scale of measurement. High values of medialness indicate positions and scales that are the most consistent with the presence of an object middle of width proportional to scale. For width-varying objects, medialness values are relatively high along a track running spatially down the middle of the object and in scale proportional to local object width at each middle point. As described later, this track of especially high medialness is the core, defined as a ridge in medialness.

This paper describes an efficient means called *stimulated cores* for computing cores in 2D images when information about approximate object location and size are given (either by user-interaction or prior knowledge). This algorithm derives its efficiency by restricting the computation of medialness to a small subvolume of scale space and is described in section 2. A number of medical applications of stimulated cores are presented in section 3.

[†] In this paper, scale refers to the standard deviation of a Gaussian kernel that is convolved with the image.

[‡] In previous papers cores have been called multiscale medial axes (MMAs).

2. THEORY OF STIMULATED CORES

Because of the computational burden associated with computing medialness at every image location and at a range of scales, we have developed a more efficient method called *stimulated cores* that begins from an approximate position and scale (i.e., from knowledge that an object is "about here" and "about this large") and which generates medialness values only as needed — first to locate an initial core point and then to traverse the core in scale space. This section briefly describes how medialness is formed for a two-dimensional image and gives a detailed mathematical discussion of the stimulated core approach.

2.1 Cores and Medialness

In order to generate cores for image objects independent of object position, orientation, and size it is necessary to insure that the medialness function commutes with the operations of translation and rotation of the image (object) and that the medialness function is invariant to image or object zoom (magnification). A way to satisfy the first two invariance requirements is to let the medialness be derived via convolution of the image with scaled, rotationally-invariant operators, or alternatively to compute geometric invariants capturing medialness properties by suitable combinations of the scaled image derivatives.[†] The third requirement, zoom invariance, is satisfied if the Gaussian derivative kernels used to derive medialness are properly normalized with respect to the scale parameter σ. Let the medialness function $M(x,y,s) = K(x,y,s) * I(x,y)$ where K is the medialness kernel and I is the original image, and let $\bar{x} = mx$, $\bar{y} = my$ and $\bar{\sigma} = m\sigma$ where m is the zoom factor. Then zoom invariance is satisfied if $M(\bar{x},\bar{y},\bar{\sigma}) = M(x,y,\sigma)$. It can be shown that zoom invariance holds for any Gaussian derivative kernel K if the kernel is constructed such that its L^2 norm is independent of the scale parameter σ (i.e., the kernel K must have constant absolute integral area at any scale σ).

We have developed several means for generating medialness functions in scale space (see Pizer [1995] for a thorough mathematical discussion of a number of medialness functions). In general, we have classified medialness functions as being either *local* or *multilocal*. Local medialness functions are those which depend only on local spatial derivative information of image intensity at each point in scale space. Multilocal medialness functions are those which medialness is derived from querying boundariness at a radial distance R proportional to scale σ at each position in scale space. Furthermore, we have classified medialness functions as being either *linear*, *semilinear*, or *nonlinear*. Linear medialness functions are those derived from convolution of the image with scaled, radially-symmetric kernels. Semilinear medialness functions are those which depend on the image data for orientation (e.g., on a gauge coordinate system). Nonlinear medialness functions are completely data-dependent and involve operations such as integrating maxima of boundariness about a circle of radius R proportional to scale σ.

Figure 1 shows a cross-section of a two-dimensional local linear medialness kernel. This kernel is the normalized Laplacian of a Gaussian, $K(x,y,s) = -\sigma G_\sigma(x,y,s) = -\sigma^2 \nabla^2 G(x,y,s)$, and when convolved with an image yields a core for an object at a scale s approximately equal to the local object width R. Cores derived from such a medialness function are presented in Fritsch [1993, 1994ab].

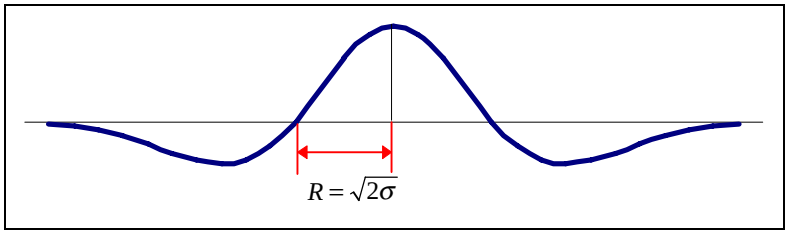

Figure 1. A local linear medialness kernel.

[†] For a thorough discussion of Gaussian scale space and differential invariants, see ter Haar Romeny [1991].

Figure 2 shows a cross section of a semilocal linear medialness kernel. This kernel sums oriented gradient measurements about a circle of radius *R* centered at a point of interest, where the scale of gradient measurement is proportional to *R*. Unlike the Laplacian of Gaussian kernel, this kernel can be "tuned", by selection of the width-to-scale ratio *k*, to limit the amount of image blur necessary to generate an object core. Moreover, because the kernel is nearly zero around its center it is less sensitive to variations in image intensity within object interiors (e.g., embedded objects, correlated noise). While this kernel can be used to better localize object boundaries and middles, a disadvantage of this kernel compared to the local linear medialness kernel is that this increased localization comes at the expense of increased sensitivity to image and boundary noise. Cores derived from such a medialness kernel are described in Morse [1993, 1994a,b].

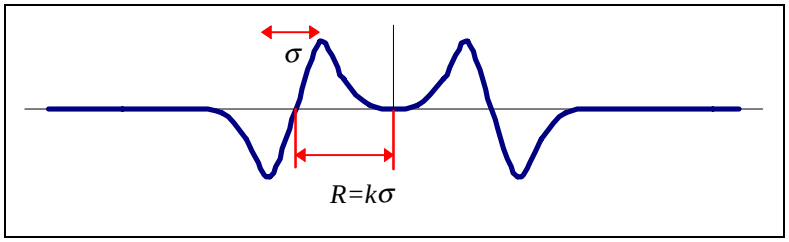

Figure 2. A semilocal linear medialness kernel.

2.2 Ridges

As discussed previously, cores are defined as *ridges* of medialness response in scale space. Ridges are a generalization of local maxima for a function. Let $f:\Re^n \to \Re$ be a C^2 function with gradient and Hessian given by *Df* and D^2f, respectively. A local maximum for *f* occurs at a point $\vec{x} \in \Re^n$ for which $Df(\vec{x}) = 0$ and $D^2 f(\vec{x}) < 0$ (the Hessian matrix is negative definite; all eigenvalues are negative). This implies that $f(\vec{y}) \le f(\vec{x})$ for all $\vec{y}$ in a sufficiently small, full *n*-dimensional neighborhood of $\vec{x}$.

The concept of local maxima can be generalized by requiring that the local maximum property only be valid in a smaller dimensional neighborhood of $\vec{x}$. The neighborhood will lie in an affine space centered at $\vec{x}$, of a specially selected dimension $n-d>0$, and spanned by specially selected linearly independent vectors $\vec{v}_k(\vec{x}), 1 \le k \le n-d$. More precisely, define the $n \times (n-d)$ matrix $V(\vec{x}) = [\vec{v}_1(\vec{x}) \cdots \vec{v}_{n-d}(\vec{x})]$ and define the function $\phi:\Re^n \times \Re^{n-d} \to \Re$ by $\phi(\vec{x},\vec{t}) = f(\vec{x} + V(\vec{x})\vec{t})$. The d-dimensional ridge points of *f* with respect to *V* are those points $\vec{x}$ for which $\phi(\vec{x},\vec{t})$ has a local maximum at $\vec{t} = \vec{0}$. Let $D\phi$ and $D^2\phi$ denote the gradient and Hessian of ϕ in the $\vec{t}$ variable. The conditions which describe *d*-dimensional ridge points are $D\phi(\vec{x},\vec{0}) = 0$ and $D^2\phi(\vec{x},0) < 0$. In terms of *f* the conditions are

$$V(\vec{x})^t Df(\vec{x}) = 0 \quad \text{and} \quad V(\vec{x})^t D^2 f(\vec{x}) V(\vec{x}) < 0.$$

2.2.1 Height ridges

A generally useful choice for a local basis of $\vec{x}$ from which the vectors $\vec{v}_k, 1 \le k \le n-d$, are chosen is the set of eigenvectors of D^2f. These are defined by $D^2 f \vec{v}_k = \lambda_k \vec{v}_k$ for $1 \le k \le n$ with $\lambda_1 \le \cdots \le \lambda_n$. The tests for a *d*-dimensional ridge point are now

$$\vec{v}_k \cdot Df(\vec{x}) = 0, \quad 1 \le k \le n-d, \quad \text{and} \quad \lambda_{n-d} < 0.$$

The second derivative tests reduce to a single test since $\vec{v}_k^{\,t} D^2 f \vec{v}_k = \lambda_k$ and the eigenvalues are assumed to be ordered. A detailed theoretical discussion of ridges and their applications can be found in Eberly [1994].

2.2.2 Optimal-scale ridges

The optimal-scale ridge definition for defining cores is designed to handle space and scale separately in their contribution to locally large medialness values. Let $M:\Re^n \times \Re^+ \to \Re$ be the medialness function in scale space, say $M(\vec{x},\sigma)$, assumed to be a C^3 function. For each position $\vec{x}$ we are interested in those scales for which $M(\vec{x},\sigma)$ is locally maximal, obtained by

solving $M_\sigma(\vec{x},\sigma)=0$ and $M_{\sigma\sigma}(\vec{x},\sigma)<0$. These conditions define a collection of *n*-dimensional surfaces, each the graph of a C^2 function $\overline{\sigma}:\Re^n \to \Re$. Each function therefore satisfies

$$M_\sigma(\vec{x},\overline{\sigma}(\vec{x}))\equiv 0,\quad M_{\sigma\sigma}(\vec{x},\overline{\sigma}(\vec{x}))<0,\ \text{and}\ -\nabla\overline{\sigma}(\vec{x})=-\nabla M_\sigma(\vec{x},\overline{\sigma}(\vec{x}))/M_{\sigma\sigma}(\vec{x},\overline{\sigma}(\vec{x}))$$

where ∇ denotes the spatial gradient operator. For the remainder of the discussion we assume that ridge construction is taking place on the graph of a given function $\overline{\sigma}(\vec{x})$. However, in the applications it is important to keep track of when the surface boundaries are encountered (these occur necessarily at places where $M_{\sigma\sigma}=0$) and when it is necessary to jump to another surface.

Optimal-scale ridges are defined as (*n*-1)-dimensional height ridges of the function $f:\Re^n \to \Re$ defined by $f(\vec{x})=M(\vec{x},\overline{\sigma}(\vec{x}))$. This function is the medialness which occurs at locally optimal scale. The idea is that the ridges of *f* contain those points for which *M* is optimal in the scale direction and in important spatial directions. This definition uses the projection of the optimal medialness values onto the spatial domain and does not take into account any curvature properties of the graph of $\overline{\sigma}(\vec{x})$. Ridge definitions that do take into account surface curvature properties are found in Eberly [1994].

To apply the height ridge definition it is necessary to compute derivatives of *f* through third order. The first derivatives are denoted f_i, the second derivatives are denoted f_{ij}, and the third order derivatives are denoted f_{ijk} where an index ℓ denotes the derivative taken with respect to component x_ℓ of $\vec{x}$. Similar notation for the derivatives of *M* is used with the addition that scale differentiation is denoted by the subscript σ. Applying the chain rule and the formula for $\nabla\overline{\sigma}$ we obtain

$$\begin{aligned}
f_i &= M_i \\
f_{ij} &= M_{ij} - M_{i\sigma}M_{j\sigma}/M_{\sigma\sigma} \\
f_{ijk} &= M_{ijk} - (M_{i\sigma}M_{jk\sigma} + M_{j\sigma}M_{ik\sigma} + M_{k\sigma}M_{ij\sigma})/M_{\sigma\sigma} \\
&\quad +(M_{i\sigma}M_{j\sigma}M_{k\sigma\sigma} + M_{i\sigma}M_{k\sigma}M_{j\sigma\sigma} + M_{j\sigma}M_{k\sigma}M_{i\sigma\sigma})/M_{\sigma\sigma}^2 \\
&\quad -(M_{i\sigma}M_{j\sigma}M_{k\sigma}M_{\sigma\sigma\sigma})/M_{\sigma\sigma}^3
\end{aligned}$$

where the evaluation of *M* and its derivatives is on the surface $(\vec{x},\overline{\sigma}(\vec{x}))$. A local (spatial) basis is chosen to consist of the eigenvectors satisfying $D^2 f\vec{v}_k=\lambda_k\vec{v}_k$ where $\lambda_1 \le \cdots \le \lambda_n$. The conditions for being an (*n*-1)-dimensional ridge point are $\vec{v}_k\cdot\nabla f=0$ and ${\vec{v}_k}^t D^2 f\vec{v}_k=\lambda_k<0$ for $1\le k\le n-1$.

2.3 Tracking Cores on Optimal-Scale Manifolds

In this paper we restrict our attention to finding cores in 2-dimensional greyscale images, so $\vec{x}\in\Re^2$. The optimal-scale ridge definition requires finding 1-dimensional height ridges of $f(\vec{x})=M(\vec{x},\overline{\sigma}(\vec{x}))$. Let $D^2 f\vec{u}=\alpha\vec{u}$ and $D^2 f\vec{v}=\beta\vec{v}$ where the eigenvalues are ordered such that $\alpha\le\beta$ and the corresponding eigenvectors $\vec{\mu}$ and $\vec{\upsilon}$ form a right-handed orthonormal system. Define $P=\vec{u}\cdot Df$ and $Q=\vec{\upsilon}\cdot Df$. According to the height ridge definition, a point $\vec{x}\in\Re^2$ is a 1-dimensional ridge point if $P(\vec{x})=0$ and $\alpha(\vec{x})<0$.

The eigenvectors of $D^2 f$ are smooth whenever $\alpha<\beta$. At an umbilic point $(\alpha=\beta)$ the eigenvectors can become discontinuous. This can cause problems in ridge tracking. If a ridge contains an umbilic point, a branching of the ridge into two or more curves may occur. The technical details of how to detect umbilics and what to do when they are encountered are not discussed in this paper. A detailed discussion can be found in Eberly [SPIE Vision Geometry paper]. In the present case, we simply track ridges in regions for which $\alpha<\beta$ and $\alpha<0$.

2.3.1 Ridge flow

Computing medialness $M(\vec{x},\sigma)$ can be very expensive since medialness requires integrating information over a spatial area proportional to scale. Methods which attempt to calculate *M* everywhere in space for a given scale are inefficient since generally cores are *sparse* in scale space. It is more desirable to localize calculations to regions which are small and contain

cores, thereby saving computation time. In a region for which $\alpha < \beta$ and $\alpha < 0$, ridges can be located by a step we call *ridge flow*. The ridge curve is a solution to $P = 0$ and $M_\sigma = 0$. These can be combined into a single function $(P^2 + M_\sigma{}^2)/2$ whose zeros are the ridge points. The path from a point $(\vec{x}_A, \sigma_A)$ to a ridge is via gradient descent applied to this function, $d(\vec{x},\sigma)/dt = -\nabla(P^2 + M_\sigma{}^2)/2$. shown below separately in space and scale:

$$\frac{d\vec{x}}{dt} = -P\nabla P - M_\sigma \nabla M_\sigma, \quad \vec{x}(0) = \vec{x}_A,$$
$$\frac{d\sigma}{dt} = -M_\sigma M_{\sigma\sigma}, \quad \sigma(0) = \overline{\sigma}_A.$$

Ridge flow terminates at a point $\vec{R}$ if $P(\vec{R}) = 0$ (or is within some specified tolerance of zero), in which case a ridge point has been found, or if $\nabla P(\vec{R}) = \vec{0}$. The latter case may occur at a point for which $P^2 > 0$, in which case the gradient descent algorithm needs modification. Such occurrences of positive local minima appear to be very infrequent in the applications.

Ridge flow requires computing ∇P. In Eberly [1994] it is shown that

$$\nabla P = \alpha\vec{u} + \frac{Q}{\alpha - \beta}(D^3 f\vec{v})\vec{\mu}$$

where $D^3 f$ is the tensor of third derivatives of f and $D^3 f\vec{v}$ is a tensor contraction which yields a 2×2 matrix. This formula avoids having to formally calculate derivatives of eigenvectors (and replaces such calculations by the third derivative terms). Note that the evaluations of P and ∇P involves computing derivatives of f through third order, which in turn requires evaluating higher-order derivatives of medialness on the surface $\overline{\sigma}(\vec{x})$.

2.3.2 Ridge (core) traversal

Let $(\vec{x}_R, \sigma_R)$ be a ridge point found by ridge flow. The ridge itself can be traversed if its tangent vectors are known. The ridges are defined as the curves of intersection of the cylinder $P(\vec{x}) = 0$ and the surface $M_\sigma(\vec{x},\sigma) = 0$, so the ridge tangents are given by the cross product

$$DP \times DM_\sigma = (\nabla P, 0) \times (\nabla M_\sigma, M_{\sigma\sigma}) = (M_{\sigma\sigma}\nabla P^\perp, \nabla P \times \nabla M_\sigma)$$

where D indicates the gradient in both space and scale where $(a,b)^\perp = (b,-a)$. Ridge traversal is determined by solving a system of differential equations

$$\frac{d\vec{x}}{dt} = \pm\nabla P^\perp, \quad \vec{x}(0) = \vec{x}_R$$
$$\frac{d\sigma}{dt} = \pm\frac{1}{M_{\sigma\sigma}}\nabla P \times \nabla M_\sigma, \quad \sigma(0) = \sigma_R.$$

The ± indicates that the ridge can be traversed in either of two directions. These equations can be solved via application of numerical ODE solvers such as the Runge-Kutta method. Ridges terminate if the eigenvalue α becomes positive, in which case the function appears locally convex, or if the direction vector changes too much from one iteration to the next. The latter provides a simple means for detecting umbilics, as these are places where the eigenspaces can swap.

2.4 Discrete Implementation

In order to implement the above core tracking algorithm, one needs to compute derivatives of the discrete image at subpixel locations (e.g., to allow for detection of zeros that indicate initial core points and to provide stability in the ODE solvers). Given simple medialness functions, such as those derived from the linear, radial-symmetric Laplacian of Gaussian medialness kernel, one can derive closed-form expressions for medialness and its derivatives in terms of convolutions of the image with various Gaussian derivative kernels. Moreover, one can use these Gaussian derivative kernels to perform the subpixel

interpolation assuming the original image is comprised of a sum of shifted impulses. For example, we can define the continuous medialness function $M(x,y,\sigma) = -\sigma L_\sigma(x,y,s)$ where

$$L(x,y,\sigma) = \sum_i \sum_j G(x-i,y-j,\sigma)I(i,j) \quad \text{and} \quad L_\sigma(x,y,\sigma) = \sum_i \sum_j G_\sigma(x-i,y-j,\sigma)I(i,j).$$

In the above expressions G is a Gaussian kernel of standard deviation (scale) σ, and $I(i,j)$ is the intensity of the original sampled image at pixel indices i and j. Similar expressions can be used to derive values for the derivatives of the medialness function needed for ridge construction at arbitrary spatial locations. Note that these derivatives need not be computed globally (e.g., via Fourier transforms), but only at a single location and scale as needed during ridge flow and traversal. Alternative means for generating subpixel approximations involve the application of 2D spline fits to optimal-scale manifolds where analytic expressions are no longer needed to compute higher order derivatives. Instead, the derivatives are obtained by differentiation of the spline equations. Spline interpolation is convenient when using semilinear and nonlinear medialness functions where the expressions for derivatives of medialness become intractable.

We have implemented the above stimulated core tracking algorithm in an interactive tool, shown in Figure 3. This interface allows the user to select a location in image space with a variable width cursor that indicates a corresponding scale, thereby indicating a point in scale space at which to begin ridge flow and subsequent core traversal. In the top left of Figure 3, the user has selected a set of cores corresponding to the scalp, brain stem, cerebellum and jaw in an image of an MRI head. The corresponding stimulated set of cores in a rotated and zoomed image demonstrate the invariance of cores to such image (object) transformations. Figure 3 also illustrates the insensitivity of cores to relatively small-scale noise and blurring and to changes in relative object contrast.

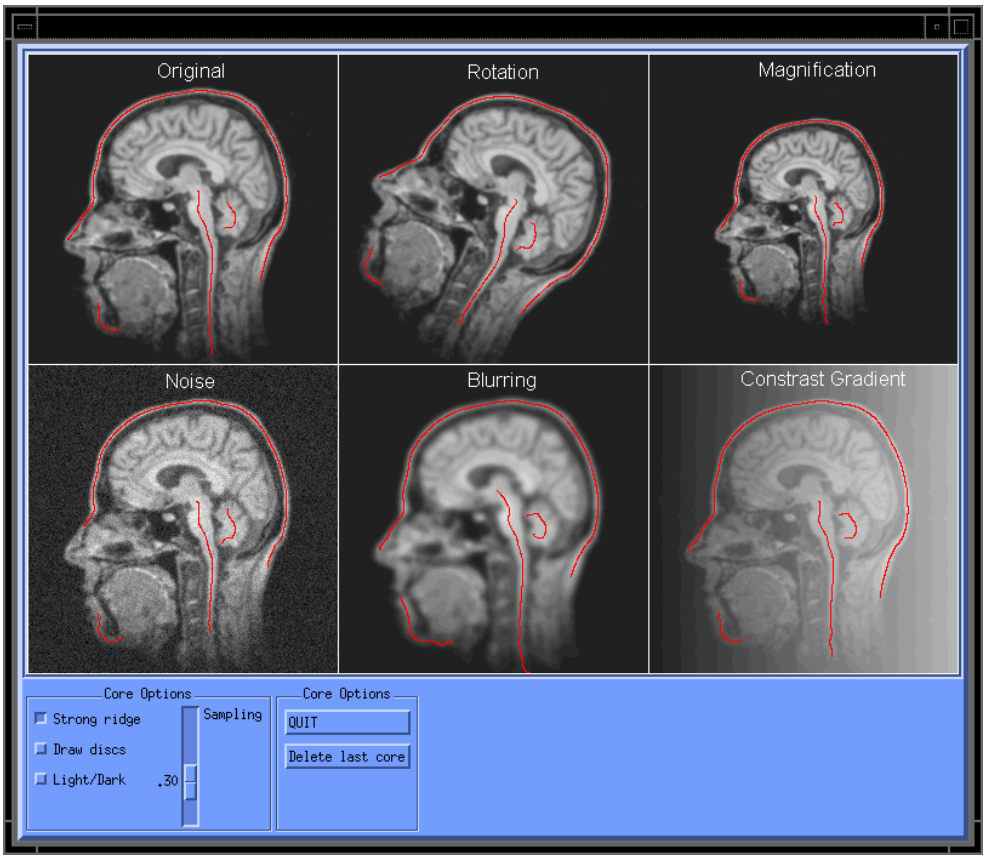

Figure 3. An user interface for stimulated core selection. Stimulated cores for an MRI image of the head illustrate the invariance of cores to rotation and zoom and the insensitivity of cores to noise, blurring and changes in contrast.

3. APPLICATIONS OF STIMULATED CORES

The ability to represent individual image objects via the stimulated core approach coupled with the properties of cores (sparseness in scale space, invariance under translation, rotation and zoom and insensitivity to noise and blurring) provide a general means for accomplishing a variety of medical imaging tasks. Several of these tasks currently under development in our laboratory are described below.

3.1 Core-Based Image Registration

A variety of image registration techniques have been described in the literature, including methods based on fiducial (e.g., landmark or curve) matching, intensity correlation, and statistical (e.g., inertial moment) matching.[†] In terms of general applicability, methods based on the correspondence of fiducials are arguably the most useful since they can take explicit advantage of user-selected features, and they can be made robust in the presence of noise or interfering objects.

To perform *automatic* image registration based on fiducial matching one needs reliable, automatic means for identifying common fiducials and means for insuring that undesirable structures (i.e., structures that may deform non-rigidly from image to image or structures that appear in only one of the images) are not considered in the registration process. Because cores capture object shape properties independent of object position, orientation and size, they are ideally suited to registering images differing by translation, rotation and magnification. Our approach to automatic image registration involves allowing the user to select, using the stimulated core interface, a collection of fiducial cores representing "stable" objects in a reference image. Assuming the misregistration between the two images is not too great, or if rough estimates of the transformation relating the images can be known, the set of reference cores can be used to find the approximate location in scale space of the corresponding set of cores in the misaligned image. This is accomplished by an optimization method where the integrated medialness in the misaligned image is maximized as a function of the parameters of the transformation (e.g., rotation, translation and magnification) applied to the reference cores. In general, a maximum is found when the transformed cores of the reference image are best aligned with the medialness ridges in the misaligned image.

Once the approximate transformation relating the cores is found by the above approach, a representative point along each transformed reference core is used as a starting place to stimulate the formation of its corresponding core in the misaligned image. Registration of the collection of corresponding cores then proceeds using a variation of the chamfer matching algorithm [Barrow, 1977; Borgefors, 1988], an iterative registration method which attempts to minimize the sum of the distances between corresponding cores in either image space or in scale space.

We have applied the core-based registration method to a variety of image pairs in both 2D (see Fritsch [1993, 1994a,b, 1995] for examples) and 3D [Liu, 1994]. Figure 3 shows the result of an automatic core-based registration applied to a pair of radiotherapy portal images, where the task is to determine the degree of patient misalignment during cancer treatment with high-energy photon beams. A portal image is acquired during treatment by recording radiation exiting the patient and shows the collimated beam along with patient anatomy. The goal then is to insure that the position of anatomical features (e.g., bones) stable with respect to the tumor are in the proper position with respect to the beam edges. Results of a study using portal images simulated from CT data [Fritsch, 1995], where exact patient misalignments were known, showed that automatic core-based registration has accuracies on the order of a single pixel in translation and on the order of one degree in rotation (pixel sizes were 0.5 mm at the patient midplane). Moreover, the registration can be accomplished automatically in under 10 seconds, given a set of precomputed reference cores, making it suitable for on-line, clinical verification of treatment setup geometry in radiotherapy.

[†] For a thorough review and classification of registration methods, see van den Elsen [1993].

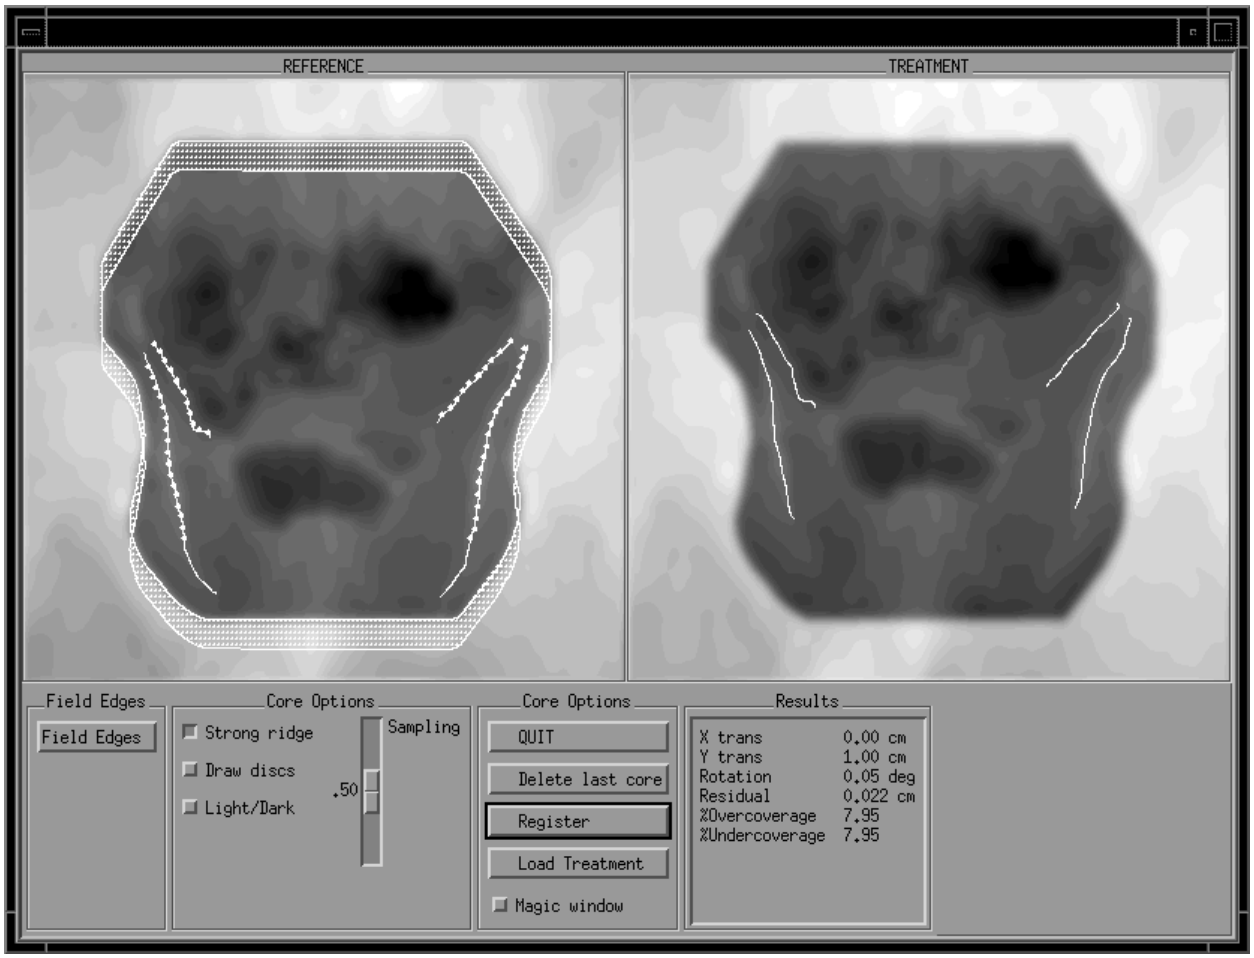

Figure 4. Core-based automatic registration of a pair of radiotherapy portal images for a patient undergoing treatment for prostate cancer. Fiducial cores representing the projection of pelvic bones have been interactively selected in the reference portal image in a one-time preprocessing step. The corresponding set of cores has been automatically computed in the acquired treatment image using a stimulated core approach. Registration of the corresponding cores yields a transformation (lower right) indicating patient misalignment of 1 cm. Shown on the reference image are transformed points from the misaligned treatment image overlaid on the reference cores. Areas of overcoverage and undercoverage are indicated by the patterned dots on the reference image.

3.2 Hierarchical Object Description and Model-Based Object Recognition

A typical object in a medical image is comprised of a collection of unbranched cores, each corresponding to what we call object *figures*. The figures for a typical object can be described in a hierarchy ordered by scale, with the largest-scale core (the *global* figure core) being a parent to its child subfigure cores which typically corresponding to smaller-scale object protrusions and indentations. As illustrated in Figure 5, this type of analysis leads to a directed acyclic graph (DAG) description of an object, the nodes of which contain statistical information related to the individual cores and the arcs of which contain information related to the subfigure type (e.g., protrusion or indentation) and the relationship of the subfigure core to its parent (e.g., relative position, size and orientation).

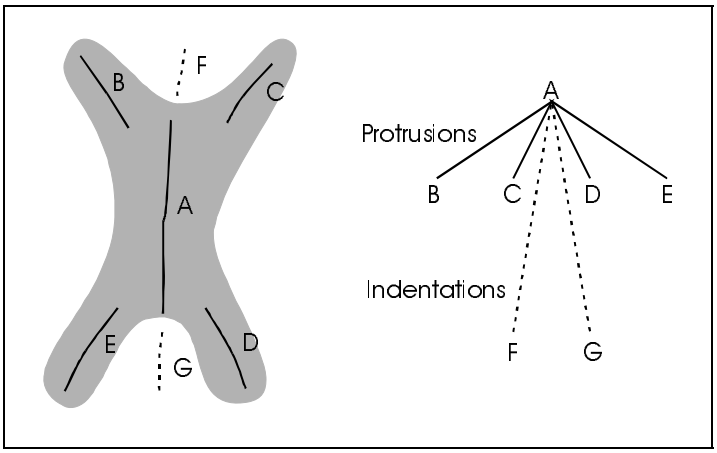

Figure 5. Core-based hierarchical object description. Shown at left is an object and its collection of figural cores. Shown at right is the directed acyclic graph (DAG) representing the core hierarchy. Solid arcs indicate object protrusions and dotted arcs indicate object indentations.

If one computes a DAG for a given object from a population of different images, one can generate a model of the object that can in turn be used to automatically recognize the same object in another image. The model DAG contains both means and variances in its nodes and arcs which can be used to facilitate the recognition stage. For example, given that one can locate the global figure core for an object, the model allows subfigure cores to be found by searching in a location (given by the model) relative to the parent core for a subfigure core and comparing the arc and node descriptors of that subfigure to those represented by the model.

As a pilot study, we have applied this procedure to a set of 2D MRI slices containing a cross section of the ventricle (Figure 6). The model was constructed by from a set of 6 similar images from different patients by representing each of the cores by an average position, scale and orientation and normalizing position and scale with respect to the central global core (i.e., the global core establishes the object coordinate system). The arcs of the DAG contain average figure position, orientation, and size (scale) information relative to the global parent core along with variances of each.

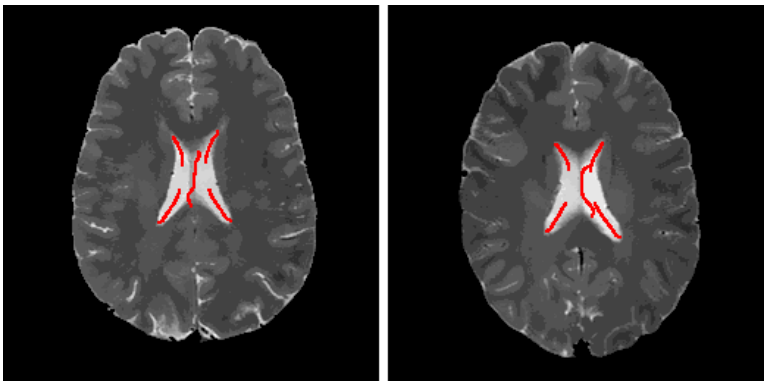

Figure 6. Core representing the ventricle in MRI scans of two different patients.

Once the model is constructed from the training set, it can be used to automatically generate the set of cores for the object in a different image. Figure 7 (left) shows a graphical representation of the model of the ventricle superimposed on an image not used in model generation. The position and width of the large central circle indicate the average point in scale space of the global central core. The small circles indicate the position and width of the subfigure cores relative to the central core, in units of distance proportional to the central core width. Double-sided arrows indicate orientation of the subfigure cores relative to the orientation of the central core. To use the model in automatic object recognition, we first locate the global core

via user interaction or knowledge of the position of the ventricle relative to the larger scale brain. For this core, we compute its average position, scale and orientation and place the model such that the central cores are aligned. This establishes a common coordinate system that allows subfigure cores to be found. Subfigure cores are then stimulated using the positional and scale information from the model. Figure 7 (right) shows the result of automatic object recognition using this approach.

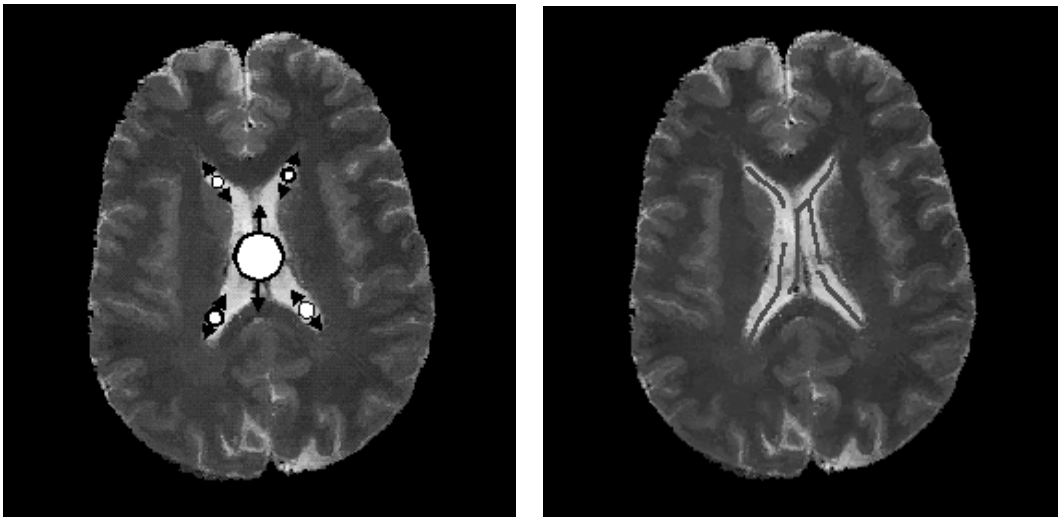

Figure 7. Model based object recognition based on cores. The figure on the left shows a graphical representation of the model of a ventricle generated from 6 cases overlaid on a test image. The figure on the right shows the result of automatic ventricle recognition by model-based core stimulation.

3.3 Vessel Tracking

A number of imaging tasks involve describing vessels in 2D x-ray projection images or in photographic images. Typically, these tasks require defining vessel centers and the associated vessel widths along the center locations — properties that are directly captured via the core. We are currently applying core-based analysis methods to the problem of identifying vessels in angiography and in photographic images of retinas. These applications are described below.

3.3.1 Angiography vessel tracking

An area of current interest in the medical imaging community is the reconstruction of arterial vasculature from two or more angiographic projection images. This application typically requires defining and corresponding vessels in a set of angiograms acquired at known angles. We have developed an angiographic vessel tracking method that uses the stimulated core approach with a local semilinear medialness function given as

$$M(x, y, \sigma) = -\sigma^2 L_{\vec{p}\vec{p}} = -\sigma^2 \vec{p}^{\,t} D^2 L \vec{p} = -\sigma^2 \lambda$$

where D^2L is the 2x2 matrix of second-order spatial derivatives of the image L at position (x,y) and scale σ and where $D^2 L\vec{p} = \lambda\vec{p}$, and λ is the largest magnitude eigenvalue of D^2L. The kernel has the shape of a symmetric second-order Gaussian derivative oriented in the direction $\vec{p}$ (see Figure 1) and falls off like a Gaussian in the orthogonal direction $\vec{p}^{\perp}$. It can be described as a scaled, oriented bar detector and performs particularly well for objects of relatively constant width and approximately parallel sides.

Figure 8 shows a portion of an angiogram of the cerebral vasculature and a set of cores computed using the above medialness function and the interactive stimulated core algorithm. This algorithm adequately handles branching and crossing so long as the width of the vessel being tracked is large compared to the branching or crossing vessel. Work is continuing in our laboratory to use extracted vessels from multiple-view 2D angiograms to reconstruct 3D vessels with the aid of a 3D magnetic resonance angiogram (MRA).

Figure 8 shows a portion of an angiogram of the cerebral vasculature and a set of cores computed using the L_{pp} medialness function and the interactive stimulated core algorithm. This algorithm adequately handles branching and crossing so long as the width of the vessel being tracked is large compared to the branching or crossing vessel. Work is continuing in our laboratory to use extracted vessels from multiple-view 2D angiograms to reconstruct 3D vessels with the aid of a 3D magnetic resonance angiogram (MRA).

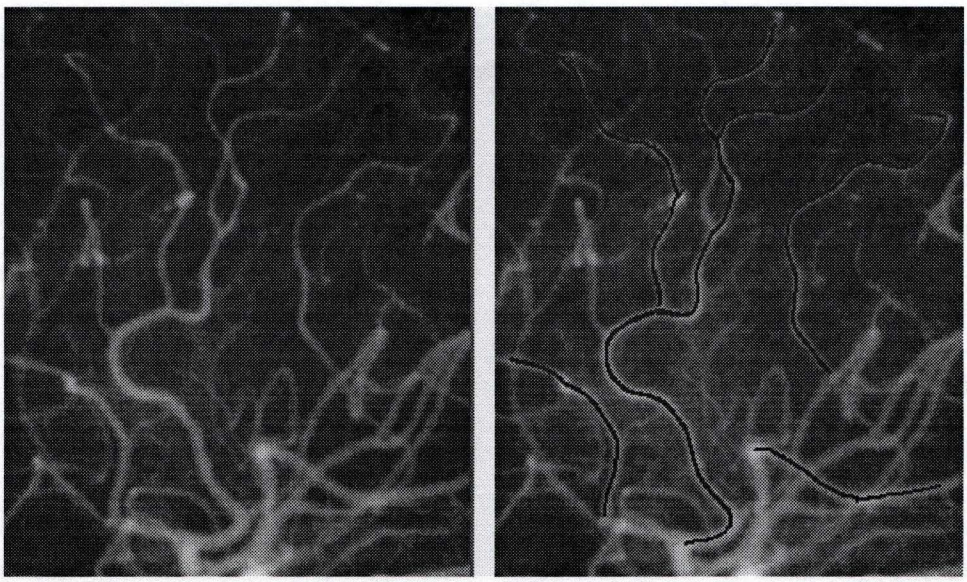

Figure 8. Core-based vessel tracking in a cerebral angiogram. Shown at left is a portion of the original angiogram. Shown at right is the projection of the spatial core coordinates (computed via interactive stimulation using L_{pp} medialness) superimposed on the angiogram.

3.3.2 Vessel tracking in retinal images

Certain eye disorders can be diagnosed by observing properties of vessels in photographic images of the retina. Figure 9 shows such an image along with a set of stimulated cores generated via user interaction using the L_{pp} medialness. Work is continuing on the development of an automatic system for the diagnosis of retinal disorders based on core analysis of vessels.

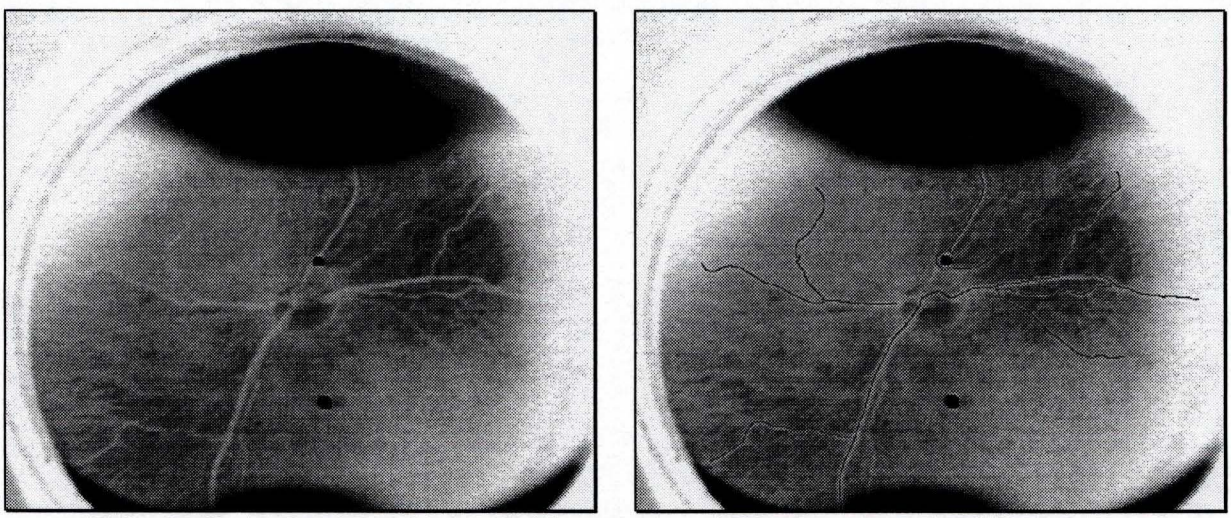

Figure 9. Vessel definition in a photographic image of an infant retina using stimulated cores from L_{pp} medialness. Cores are used to derive measures of vessel tortuosity ("wiggliness") and vessel width that are in turn used to aid in the diagnosis of an infant eye disease called retinopathy of prematurity (ROP).

3.4 Boundary Claiming

While cores provide robust representations of image objects via middle and width properties, they do so at a scale proportional to object width and so give boundary locations only to a tolerance proportional to this scale. Many image analysis applications require more accurate boundary definition than provided by the core — for example, applications that require measurement of object area or volume, or tasks which require object rendering (segmentation).

Given a core representation of a figure, one can generate a type of boundary description called the *boundary at the scale of the core*, or BASOC for short. The BASOC is computed as the boundary of the binary object whose Blum [1978] medial axis is the core. More precisely, parameterize the core by its spatial arclength s such that the core is the locus $(\vec{m}(s),\sigma_C(s))$, where $\vec{m}(s)$ is the medialness track in image space and $\sigma_C(s)$ is the corresponding scale of the core. The BASOC is the locus of points whose envelope is the union of the set of disks centered at $\vec{m}(s)$ with respective radius $k\sigma_C(s)$. For medialness functions generated using the kernel in Figure 1, the width-to-scale ratio k can be taken to be unity. For the multilocal medialness kernel in Figure 2, the k used here is the same as the value used to form the medialness function.

For normal BASOC points (those corresponding neither to branchpoints nor endpoints), the perpendiculars from the BASOC points make an angle of $\theta(s)=\sin^{-1}\left[1/\left(1+(k\sigma'_C(s))^2\right)\right]$ with the core. Thus, each normal core point is associated with two opposing BASOC points (called *involutes*) and these points can be assigned a scale value $\sigma_B(s)$ proportional by a factor of k_B to the scale of the core point $\sigma_C(s)$ (see Figure 10).

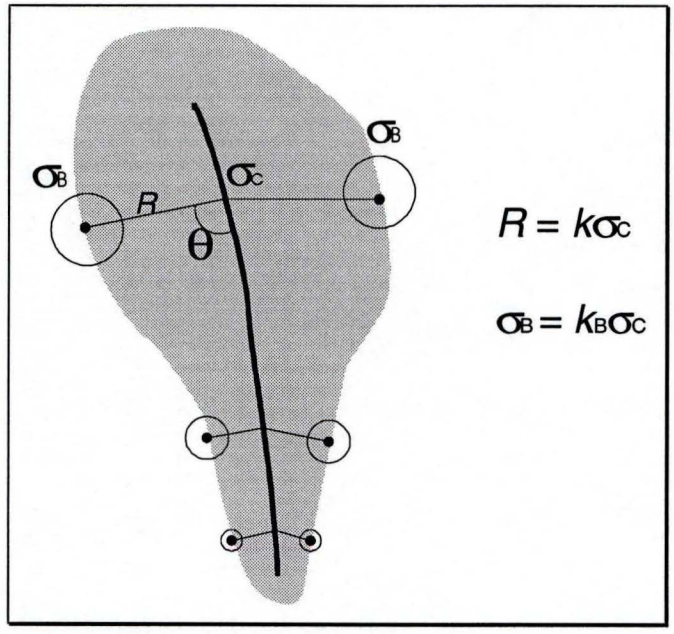

Figure 10. Associating core points with BASOC points. The scale σ_B of the opposing BASOC points (involutes) is chosen proportional to the scale σ_C of the corresponding core point.

The basic idea behind core-based boundary definition is to use the scale information of the core to give the proportional scale for boundary measurement at the associated BASOC points (hence insuring that the zoom invariance is satisfied) and to use the BASOC location as a starting place for boundary refinement. At this early stage of research, we use a simple but effective means based on the active contours (snakes) approach [Kass, 1987]; at each iteration, a given BASOC point is allowed to move along the path normal to the BASOC in the direction that maximizes gradient magnitude measured at the scale of the BASOC point. We call this process *core-based active contours*. We currently do not include internal energy terms in the active contours algorithm as the use of multiscale information for boundary measurement tends to insure smoothness in the final boundary contour.

Figure 11 shows the final boundary configuration resulting from application of core-based active contours to the same BASOC (derived from a single global core) using several values of the BASOC-scale to core-scale ratio k_B. For k_B=1 the resulting boundary is essentially the same as the initial BASOC. A k_B of ½ produces a boundary that nicely reflects the global shape of the tube and which only slightly reflects the smaller scale variations in the object's boundary. Reducing k_B further allows the boundary to conform to the smaller scale indentations and protrusions.

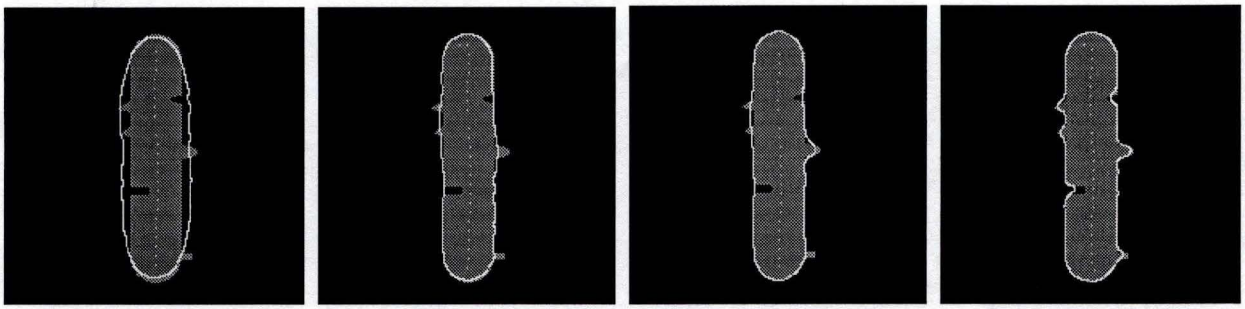

k_B = 1 k_B=1/2 k_B=1/4 k_B=1/8

Figure 11. Core-based active contours applied to a tube object with various indentations and protrusions for several values of the BASOC-scale to core-scale ratio k_B.

Figure 12 shows BASOCs for several objects in an MRI image of the head and the boundaries that result from application of the core-based active contours method. The final boundaries are arguably improved over those provided by the BASOC. Work is ongoing to investigate ways of incorporating subfigure cores into the segmentation process.

Figure 12. Core-based active contours applied to BASOCs corresponding to the brain stem, cerebellum and ventricle in an MRI image of the head. Shown at left are the initial BASOCs generated from a Laplacian-of-Gaussian medialness function. Shown at right are the boundaries resulting from application of core-based active contours using a BASOC-scale to core-scale ratio of ¼.

Because core-based active contours begins from a stable, object-based representation, and because we use object-relevant scale information in the measurement of boundaries, we have the advantage of constraining the number of possible boundary configurations that can result from application of the active contour approach. We envision using core-based boundary-finding methods as part of a (possibly interactive) tool for the fast and accurate extraction of anatomical objects for purposes such as radiotherapy or surgical treatment planning. Coupled with the ideas of building models of image objects based on cores, we believe that automatic core-based segmentation of anatomy is an achievable goal.

4. Discussion and Conclusion

We have developed a theory of object representation in greyscale images called cores that is derived directly from the image intensities and which captures object middle and width information independent of object position, orientation and size. Because the core is defined as a track of medialness in scale space, it is stable against noise, blurring, and boundary texture occurring at scales small relative to object widths.

We have described an efficient means for computing cores for image objects that takes advantage of sparseness properties of cores and which uses prior knowledge about object location and size to restrict the region over which measurements are made. This method, called stimulated cores, tracks the ridges in scale space that are the core by computing measurements of medialness only as needed and is capable of extracting a core in a very few seconds.

The invariances of cores to translation, rotation, and zoom coupled with their stability and sparseness properties give rise to core-based solutions to a number of medical applications that require robust means of object representation. We have demonstrated the power of core-based analysis in image registration, model-based object recognition, vessel tracking in angiograms and retinal images, and in segmentation (boundary-finding) tasks. While various extensions of these ideas to 3D images have been shown on a pilot basis and programs for core extraction from 3D images have been completed, 3D applications are at an early stage of development. Generalization of the stimulated core approach (described here for 2D images) to 3D images will allow cores to be interactively extracted using current workstation technology and will facilitate development of 3D core-based applications.

5. References

Barrow, H.G.; Tenenbaum, J.M.; Bolles, R.C.; Wolf, H.C. Parametric correspondence and chamfer matching: Two techniques for image matching. Proc. 5th Int. Joint Conf. Artificial Intell. 659-663; 1977.

Blum, H.; Nagel, R.N. Shape description using weighted symmetric axis features. Pattern Recognition 10:167-180; 1978.

Borgefors, G. Hierarchical chamfer matching: A parametric edge matching algorithm. IEEE Trans. PAMI 10(6):849-865; 1988.

Eberly, D.H.; Gardner, R.B.; Morse, B.S.; Pizer, S.M.; Scharlach, C. Ridges for image analysis. J. Math. Imaging and Vision 4:351-371; 1994a.

Eberly, D. Fast algorithms for ridge construction. To appear in Proc. SPIE Photonics East 1994: Vision Geometry III 2356; 1994b.

Fritsch, D.S. Registration of radiotherapy images using multiscale medial descriptions of image structure. Doctoral dissertation, Department of Biomedical Engineering, University of North Carolina at Chapel Hill; 1993.

Fritsch, D.S.; Pizer, S.M.; Chaney, E.L.; Liu, A. Cores for image registration. Proc. SPIE Medical Imaging '94: Image Processing 2167:128-142; 1994a.

Fritsch, D.S.; Pizer, S.M.; Morse, B.S.; Eberly, D.H.; Liu, A. The multiscale medial axis and its applications in image registration. Pattern Recognition Letters 15:445-452; 1994b.

Fritsch, D.S.; Chaney, E.L.; Boxwala, A.; McAuliffe, M.J.; Raghavan, S.; Thall, A.; Earnhart, J.R.D. Core-based portal image registration for automatic radiotherapy treatment verification. Submitted to International Journal of Radiation Oncology, Biology, Physics, 1995.

Kass, M.; Witkin, A.; Terzopolous, D. Snakes: Active contour models. Proc. First International Conf. on Computer Vision:259-268, 1987.

Liu, A; Pizer, S.M.; Eberly, D.; Morse, B.S.; Rosenman, J.; Chaney, E.L.; Bullitt, E.; Carrasco, V. Volume registration using the 3D core. Visualization in Biomedical Computing, Proceedings of the SPIE 2359:217-226; 1994.

Morse, B.S.; Pizer, S.M.; Liu, A. Multiscale medial analysis of medical images. In: Barrett, H.H., Gmitro, A.F., eds. Information Processing in Medical Imaging, Lecture Notes in Computer Science 687. Berlin: Springer-Verlag; 1993:112-131.

Morse, B.S.; Pizer, S.M.; Fritsch, D.S. Robust object representation through object-relevant use of scale. To appear in Proc. SPIE Medical Imaging '94: Image Processing, 1994.

Pizer, S.M.; Eberly, D.; Fritsch, D.S. Objects, boundaries, and medical image analysis. Submitted to ICCV, 1995.

ter Haar Romeny, B.M.; Florack, L.M.J; Koenderink, J.J.; Viergever, M.A. Scale-space: Its natural operators and differential invariants. In: Colchester, A.C.F. and Hawkes, D.J., eds. Information Processing in Medical Imaging, Lecture Notes in Computer Science 511. Berlin: Springer-Verlag; 1991:239-255.

van den Elsen, P.A.; Pol, E.D.; Viergever, M.A. Medical image matching — A review with classification. IEEE Eng. Med. Biol., 1992.

6. Acknowledgments

This research was supported in part by the National Library of Medicine under the HPCC grant LM05508-1, by a National Library of Medicine/National Cancer Institute Medical Informatics Research Training grant 5 T15 LM07071, and by an NIH program project grant P01 CA47982.

The authors gratefully acknowledge the contributions of Bryan Morse, Edward Chaney, Andrew Thall, Alan Liu, and Aziz Boxwala to this work.