

# Extracting the Geometry of Branching Slabs via Cores

Yonatan Fridman, Stephen M. Pizer

Medical Image Display & Analysis Group  
University of North Carolina, Chapel Hill, NC  
fridman@cs.unc.edu

**Abstract.** Our objective is to extract branching objects such as the kidney and its ureter from 3D medical images. Our approach is based on cores, height ridges of a graded measure of an image location's medial strength. Cores' strength lies in their providing a powerful representation of shape while being relatively insensitive to image noise and boundary details. The novelty in this paper is the extension of cores from existing methods that handle branching tubular objects to a broader ability to handle non-tubular, or slab-like, objects. The slab-like core extraction method presented utilizes a coarse-to-fine grid-based marching scheme with branch handling and automated end detection. Performance is analyzed using sets of synthetic images and demonstrated on CT images of kidneys and adjoining vasculature. Results are good on relatively low noise medical images and show promise for use on more difficult images.

## 1 Introduction

The structure and geometry of anatomic objects in three-dimensional medical images are often desired, but are not explicitly available in the image data. Creating a geometric representation of the objects of interest from image intensity data can provide such structural information and aid in analysis and patient treatment. This paper considers a method for computing such a representation using 2D cores of branching objects in 3D images.

The 2D cores method presented here extracts object information directly from image data by marching along an object's medial locus (axis or sheet), using object boundaries' high image gradient as a means of directly following the medial locus. As the marching occurs, potential branches are flagged and indications that the object is ending are sought. Then potential branches are investigated further, and core marching resumes along valid branches.

This method is an extension of the successful tubular cores method presented in (Fridman et al., 2004), which provided a means to robustly extract the branching geometry of tubes from 3D images without the need for a prior segmentation. Both the tubular method and the current extension to slabs rely on cores, which are medial axes, at scale, computed using techniques developed in the Medical Image Display and Analysis group at the University of North Carolina at Chapel Hill (Fritsch et al., 1995; Furst, 1999).

Cores of objects in 3D images can be either 1D or 2D (Fig. 1). A core of a generic 3D object is a surface in three-space. If the object of interest is assumed to be tubular, i.e., to have circular cross-sections, the desired core is a curve in three-space. Such 1D cores of branching tubes have been studied in (Aylward and Bullitt, 2002; Fridman et al., 2004). This paper builds on those methods to compute 2D cores of non-tubular, or slab-like, objects.



**Fig. 1.** (Left) A 1D core of a tubular object. (Right) A 2D core of a slab-like object.

Section 2 of this paper provides motivation and an overview. Section 3 introduces medial atoms and describes a marching method for extracting 2D cores of slab-like objects. Sections 4 and 5 describe methods for branch handling and automatic termination of core following, respectively. Sections 6 and 7 provide analyses, results and conclusions.

## 2 Method Motivation and Overview

Methods for image segmentation can roughly be separated into two categories: deformable methods and marching methods. Deformable methods require prior knowledge of the object being analyzed, as they are based on a model which is deformed according to image data to fit the specific object instance. For example, m-reps (Pizer et al., 2003) deform models that bear similarities to cores in their structure and in the way they interrogate the image.

While deformable methods can be quite robust, the prior knowledge on which they rely is not always available. Such methods cannot handle structures like blood vessel trees, which are highly variable from person to person, or even simple structures for which a model has not been created. Rather, marching methods are appropriate for these situations.

Marching methods can be further divided into boundary methods and medial methods. Boundary marching methods are susceptible to the effects of image noise due to the local manner in which they interrogate the image. To counter this weakness, a series of methods stemming from Blum's Medial Axis Transform (Blum, 1967) were developed that derive an object's medial locus from its boundary, and that interrogate the image multi-locally. This multi-locality makes such methods more resistant to image noise, but they tend to be overly sensitive to object boundary perturbations.

Scale space medial methods have been developed to reduce these sensitivities to image and object noise. For example, Näf et al. (1996) produce a Voronoi diagram from an object's boundary, prune diagram edges corresponding to small scale details, and then impose a topology to produce a coherent medial locus, at

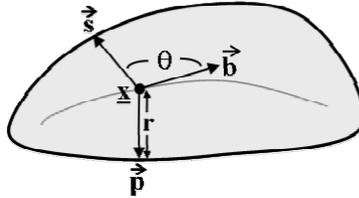
scale. Siddiqi et al. (1999) use a curve evolution method in which an object's boundary propagates inward according to a Hamilton-Jacobi equation. Medial locations are given as those places where shocks of the propagating curve occur. These methods produce excellent results but rely on a segmentation.

Cores are resistant to both image and object noise and produce a medial representation directly from image data. The branching tubular cores method presented in (Fridman et al., 2004) shows impressive resistance to noise due in part to the constrained medial atom structure it uses to simultaneously sample eight locations on the object boundary. The method presented here extends that work in an attempt to extract general objects from 3D images. The method follows the core of an object by marching along a medial grid, where a bi-local medial atom is fit to the image data at each grid vertex. Object branches are handled using an affine-invariant corner detector that identifies potential branches by giving a strong response at boundaries of branches. Potential branches are examined and then either discarded or saved and extracted using the core following method. Core following is automatically terminated when local statistics on the fit of medial atoms to the image data indicate that the atoms no longer lie on the core.

### 3 Core Following

#### 3.1 Medial Atoms

The core following method extracts a sampled core composed of discrete medial atoms. A medial atom  $\underline{m}$  is a structure defined by four parameters  $(\underline{x}, r, \underline{F}, \theta)$ —the coordinates  $\underline{x}$  of the atom in 3-space, the radius  $r$  of the object of interest at  $\underline{x}$ , a frame  $\underline{F} = (\vec{b}, \vec{n}, \vec{b}_{perp})$  that describes the orientation of the medial atom, and an object angle  $\theta$  that describes the rate of widening or narrowing of the object (Fig. 2). For slab-like objects these four parameters imply a pair of vectors, known as spokes, which extend from the medial location  $\underline{x}$  to the implied object boundary. In this paper  $\theta$  is constrained to  $\pi/2$ , resulting in collinear spokes. The reasons for and consequences of this are discussed in section 3.3.



**Fig. 2.** The geometry of a medial atom.  $\underline{x}$  is the medial atom's location;  $r$  is its radius;  $\theta$  is its object angle;  $\vec{b}$  is the core's tangent direction at  $\underline{x}$ ;  $\vec{p}$  and  $\vec{s}$  are the atom's spokes.

### 3.2 Medialness

The medialness  $M(\underline{m})$  of a medial atom is a scalar function that measures the fit of the medial atom to image data. For the purposes of this paper it is measured by placing a derivative of a Gaussian at the tip of each spoke, where the derivative is taken in the direction of the spoke. These functions provide weights on the image and the results of the weighting are integrated to give the medialness value. This is given by the equation

$$D_{\frac{\underline{x}}{\bar{p}}} I(\underline{x} + \bar{p}, \sigma) + D_{\frac{\underline{x}}{\bar{s}}} I(\underline{x} + \bar{s}, \sigma) \quad (1)$$

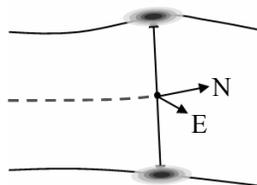
where  $\underline{x}$  is the medial atom location,  $\bar{p}$  and  $\bar{s}$  are the spokes, and  $\sigma$  is the scale of interrogation, i.e., the standard deviation of the Gaussians. The methods presented in this paper use a scale of interrogation that is constant across object widths. This choice is meant as a compromise between the detection rate and the localization accuracy of medial atoms, since detection is improved by increasing the scale of interrogation while localization is improved by decreasing the scale of interrogation (Witkin, 1983).

### 3.3 Marching

In the method described here, to select a medial atom as being on the core, a position  $\underline{x}_0$  is predicted from previously computed atoms. Then, fixing  $\underline{x}_0$ , medialness is optimized over the parameter space  $(r, \underline{F})$ . The location  $\underline{x}$  of the medial atom is then found by further optimizing medialness over the spatial direction that is parallel to the atom's spokes and that passes through the predicted position  $\underline{x}_0$ . This choice of optimization space defines what is known as an optimum parameter core (Fritsch et al., 1995; Furst, 1999).

This method of determining medial atoms is used in a coarse-to-fine grid-based marching algorithm to extract a core. The algorithm marches in four directions – call these directions North, South, East, and West – that are orthogonal to the spokes, spanning the tangent plane to the core (Fig. 3). Given a medial atom, which is manually initialized for the first iteration, the algorithm takes four separate steps from the given atom's location, one in each compass direction, to predict the locations of four potential new atoms. These potential atoms are assigned the same radius and orientation as the initial atom and are placed on a queue. On each subsequent iteration a potential atom is taken from the head of the queue, optimized over position, radius, and orientation as described earlier, and then tested to determine if it meets two conditions: 1) it is still on the core, i.e., has not fallen off the end of the core. This test is discussed further in section 5; and 2) it is not too close in position to an existing atom. If the atom fails ei-

ther of these two conditions, it is discarded. Otherwise, the atom is saved and its own North, South, East, and West atoms are computed and placed on the queue.



**Fig. 3.** Marching along a slab-like core in 3D. The North and East marching directions span the tangent plane to the core.

This marching scheme works well in noise-free images, but the introduction of noise leads to the introduction of error. After a large number of steps, minor inaccuracies compound and become unacceptable. This effect is reduced by first computing the core at a coarse sampling, not allowing enough steps for the errors to compound significantly. This coarse rectangular sampling is then sub-sampled by placing a new medial atom in the center of each rectangle and optimizing it. In sub-sampling the core, new medial atoms are also added outside of the border of the original, coarsely sampled, core. This differentiates this marching scheme from the coarse-to-fine paradigm used in m-reps (Pizer et al., 2003). The two are further differentiated in that the marching scheme presented here is coarse-to-fine only in sampling and not in the scale of interrogation used to locate the medial atoms. It uses the same scale at each level of sampling.

This coarse-to-fine approach decreases potential errors in core extraction, but the medial atom structure is nonetheless much more susceptible to image noise than the tubular medial atom described in (Fridman et al., 2004). For this reason  $\theta$  is constrained to  $\pi/2$ , improving the method's stability. This constraint also assumes that opposing object boundaries are parallel; as the boundaries deviate from parallel, the fit of the Gaussian derivatives to the image data deteriorates and the resulting core following becomes less robust. Fortunately, this effect only becomes significant near an object angle of  $\pi/4$ , allowing for effective extraction of most anatomic objects.

To further stabilize core extraction, two penalty terms are added to Equation 1 that penalize changes in geometry between adjacent atoms. In particular, medialness is penalized if a medial atom's radius or the orientation of its spokes differ from those of its neighbor. For each of these two parameters, a small difference in values has no effect, but beyond a chosen threshold further differences are subtracted linearly from computed medialness values. During the first (coarse) level of marching, these penalty terms are applied to differences in parameter values between the new potential atom and the atom from which it was predicted. During subsequent levels of sub-sampling, the penalty terms are applied to differences between the parameter values of the new potential atom and the average parameter values of the four surrounding atoms. Using the coarse-

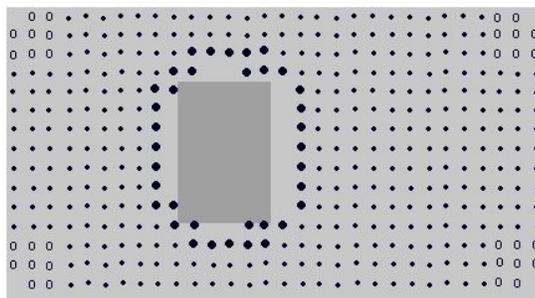
to-fine marching scheme, constraints, and penalties, core extraction is fairly robust, as shown in section 5.

## 4 Branch Handling

Complex objects that have branches, protrusions, or sub-figures cannot be represented by a single core. Rather, each branch is represented by an individual core, and connections are made between the cores in order to represent the entire object. The core following method presented in section 3 does not handle sub-figures and must be augmented with branch handling techniques.

To handle branches, a cornerness operator is applied to the image, and strong responses in cornerness are noted during core following. After the core of the current figure is extracted, locations of high cornerness are revisited. From these locations, new branches are predicted using the geometry of the computed core and then the cores of these new branches are extracted in turn. This is analogous to the method presented in (Fridman et al., 2004) that effectively handled branches of tubular objects.

In the method presented here for handling branches of slabs, each new medial atom produced during core following is examined to determine if it is situated in the vicinity of a branch. This is accomplished with the application of an affine-invariant corner detector,  $L_{\bar{u}\bar{u}}L_{\bar{v}}$  (Blom, 1991; Lindeberg, 1994; ter Haar Romeny, 2002), to the three-dimensional image, where  $\bar{v}$  is the image gradient direction and  $\bar{u}$  is the eigenvector corresponding to the largest eigenvalue of the Hessian in the plane normal to  $\bar{v}$ . This operator is applied at the same scale at which medialness is computed (see section 3.2). Medial atoms whose spoke tips are at locations of high cornerness are flagged as branch points (Fig. 4).



**Fig. 4.** One slice of a synthetic image of a slab (*light gray*). The dark gray area represents a rectangular sub-figure that protrudes from the main figure toward the viewer. The dots show the computed medial atoms of the main figure and their associated cornerness values. Open dots represent negative cornerness values, small dots represent both positive and negative cornerness values near zero, and large dots represent high cornerness values (above a chosen threshold). In this example, all medial atoms indicated by large dots were flagged as branch points.

Once extraction of the main figure's core is complete, one of the flagged atoms is arbitrarily chosen, and a sub-figure atom is predicted from the flagged atom's geometry. The core of the sub-figure is then followed in the same manner as was the core of main figure, starting from the predicted atom. If core following terminates (see section 5) before finding five medial atoms, the sub-figure is discarded as a false positive. Once core following of the sub-figure is complete, another flagged atom from the main figure is chosen, and a new sub-figure atom is predicted. If this predicted atom falls within the volume defined by an existing sub-figure, or if core following from this predicted atom leads into an existing sub-figure, it is discarded. This process continues until all flagged atoms have been considered.

## 5 Core Termination

Core following should be terminated as soon as it is determined that the core reaches the end of the object. The method used for deciding when to terminate core following is based on the confidence with which the desired core is detected, determined by the strength of the medialness value of the core relative to local statistics on the expected medialness value of the core. For each new potential atom a large number of randomly positioned and oriented medial atoms are sampled in the spatial vicinity of the core to get a sense for the range of expected medialness values. It is assumed that the large majority of these medial atoms will not be aligned with an object. Due to the way in which medialness is defined in this paper (see Eq. 1) medialness can be either positive or negative and the mean medialness value of random medial atoms should be approximately zero. If the mean is not near zero, the image is resampled. If the medial atom along the detected core is more than three standard deviations above the mean then it is accepted as a valid core point, otherwise it is discarded.

## 6 Results and Analysis

### *6.1 Analysis of Methods Using Synthetic Images*

The properties of the method described were analyzed on sets of synthetic images of slab-like objects. Core following and core termination were tested on images of both narrowing objects and tortuous objects. Branch handling was then tested on images having a main figure and one attached sub-figure.

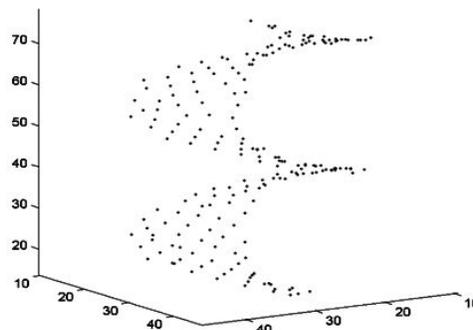
#### 6.1.1 Core Following of Narrowing Objects

As described in section 3.3, core following becomes less robust as the boundaries of the core being followed deviate from parallel. To better quantify this, a

set of 120 synthetic images were created of objects having opposing walls that narrow or widen at differing rates. The objects spanned a range of sizes, shapes, and curvatures. Three different levels of Gaussian noise were added to the images – low (approximately half of what would be expected in an MRA (Aylward and Bullitt, 2002)), normal (approximately what would be expected in an MRA), and high (approximately twice what would be expected in an MRA). Core following was tested by running the code ten times on each of the 120 images. Core following performed similarly whether the object was narrowing or widening in the marching direction. As expected, the success rate was inversely proportional both to image noise level and to the rate of narrowing/widening. For example, in images with normal noise, 79 of the 80 test runs with parallel object boundaries correctly extracted the desired core, while only 58 of the 80 test runs with rapidly narrowing objects found the desired core. Figure 6a details these results, where each data point represents 80 test runs. The narrowing rates represent the deviation of each of the two opposing object boundaries from parallel, e.g., for an object with a narrowing rate of 30 degrees, the angle between the two boundary tangent directions is 60 degrees.

### 6.1.2 Core Following of Tortuous Objects

To test the stability of cores on tortuous objects, a set of 24 images were created of helical ribbons having varying amounts of curvature and torsion, and ten test runs were performed on each image. 97% of the test runs on low noise images correctly extracted the core, 89% of the normal noise cases were accurate, and 71% of the high noise cases were accurate. Figure 5 shows an example core.



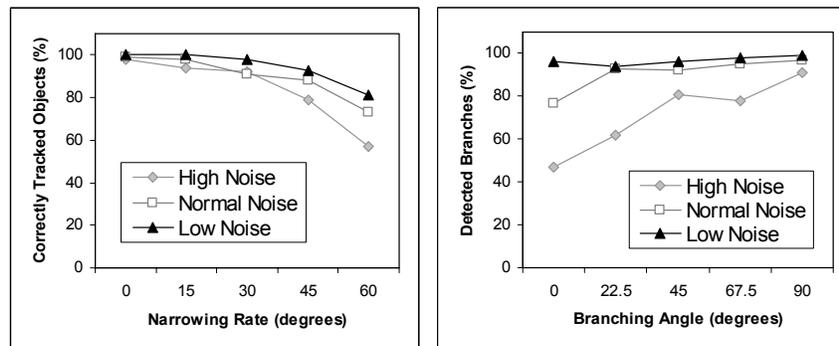
**Fig. 5.** 3D view of the medial atoms computed from a synthetic image of a helical ribbon. The axis units are image voxels.

### 6.1.3 Branch Handling

To analyze the performance of branch handling, branches are first divided into three categories: crest branches, face branches, and abutting objects. A crest branch occurs when a sub-figure is attached to the crest of the surface of the main figure, where the crest of a surface is a curve along which the first princi-

ple curvature of the surface has a local maximum in the first principal direction. A finger is a good example of such a crest branch, as it is attached to the crest of the hand's surface. Crest branches have the property that the core of the main figure and the core of the sub-figure connect end to end. Consequently, branch handling may not be required for this type of branch, as core following extracts both figures at once.

A face branch occurs when the core of a sub-figure is connected to the core of the main figure anywhere but at an end atom. Such a branch might appear as a finger protruding from the palm of the hand, and the branch handling method described in section 4 is required to identify this type of branch. To test the performance of the branch handling method in this situation, a set of 120 images were created of objects with a main figure and one sub-figure, where the branching angle between the two figures varies. A branching angle of 0 degrees in Figure 6b indicates that the two figures have medial axes that share a tangent plane at the point where they branch and which then diverge. In computer graphics terminology, curves exist spanning the boundaries of both figures that are  $G^1$  continuous. This differs from the situation for a non-zero branching angle, in which case there is a discontinuity where the two figure boundaries intersect. The code was run ten times on each of the 120 synthetic images. As shown in Figure 6b, branch handling performed consistently well over all branching angles in low image noise cases, but in the presence of high image noise performance deteriorated significantly for smaller branching angles. Not shown in the graph is the tendency for narrow sub-figures to be missed more often than wide ones.



**Fig. 6.** (Left) The success of the core following method on synthetic images of narrowing slab-like objects. Each data point shows the percentage of test runs in which the object was correctly extracted for a given narrowing rate and image noise level. (Right) The success of the branch handling code on synthetic images of branching slab-like objects. Each data point shows the percentage of test runs in which the branch was correctly detected for a given branching angle and image noise level.

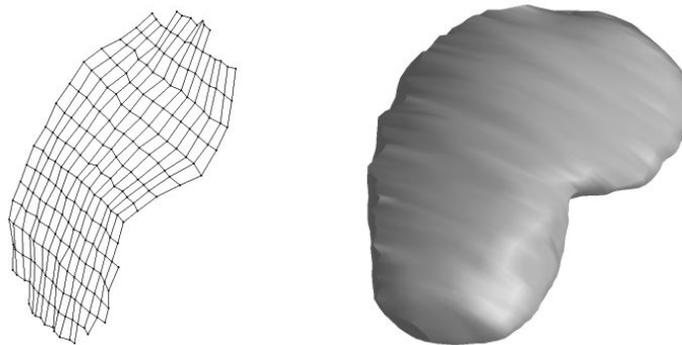
The third possible branch type occurs in the case of two separate but abutting objects. For example, in CT images the kidney and liver are often displayed at indistinguishable intensities, causing them to appear as a single object with no

defined border when pressed side by side. With this type of branch the sub-figure is usually not of interest, since it is a separate object. 24 synthetic images were created to simulate this situation, and ten test runs were performed on each image, with core following initialized in the kidney. In 83% of the test runs the liver was automatically identified as a sub-figure and in 17% it was ignored. On only one test run out of 240 did core following incorrectly extract the kidney and liver together as one figure. In the case where one abutting object is recognized as the main figure and the other is chosen as a sub-figure, the user must manually indicate that the sub-figure should not be considered, i.e., that it is a separate object. There is currently no method for identifying abutting objects as such automatically.

## 6.2 Demonstration of Methods on Clinical Images

The methods for core following with branch handling and core termination were run on five clinical CT images of the abdomen. The images each consist of 256 x 256 x 128 isotropic voxels with resolutions ranging from 0.9 to 1.25 mm/voxel.

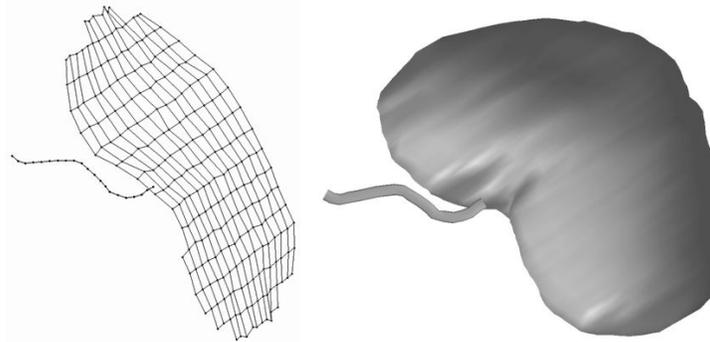
For each kidney in each of the five images, core following with branch handling and end detection was performed ten times for a total of 100 test runs. Each test run was initialized at a different location within the given kidney. Of these 100 test runs, the abutting liver was detected as a branch 12 times. In another 7 test runs the core clearly lost track of the kidney. The remaining 81 test runs gave results similar to those shown in Figure 7.



**Fig. 7.** (Left) 3D view of the 2D core of a kidney, computed from CT. Medial atoms are shown as dots connected by line segments. (Right) The 3D surfaces implied by the computed core.

Furthermore, branch handling was extended to integrate the 2D slab-like cores described here with the 1D tubular cores described in (Fridman et al., 2004). This was accomplished by running the 2D core following method on the main figure, but forcing sub-figures to be tubular. Figure 8 shows an example in

which the kidney was extracted in such a manner and the renal artery was automatically located as a branch.



**Fig. 8.** (*Left*) 3D view of the 2D core of a kidney, treated as a slab-like object, along with the 1D core of the attached renal artery, treated as a tubular object. (*Right*) The 3D surfaces implied by the computed cores, viewed from a slightly different angle.

## 7 Conclusions and Discussion

This paper presents the use of 2D cores for extracting the geometry of branching slab-like structures in three-dimensional grayscale images. The methods described are capable of handling a wide variety of object shapes and structures, but lack consistency in noisy images.

Core following and core termination are fairly robust in most situations, including highly tortuous objects. Core following becomes poor in situations where an object narrows or widens rapidly, particularly when opposing object boundaries deviate more than 30 to 45 degrees from parallel, depending on the level of image noise. Although this does not significantly affect the extraction of most anatomic objects, it is a worthwhile topic for future research, as cores are particularly susceptible to this problem near object ends.

Branch handling shows promise for identifying sub-figures of slab-like objects but occasionally fails in certain situations. Narrow sub-figures are often missed, mainly when the branching angle is small and the signal-to-noise ratio is low in the vicinity of the branch. This occurs because the branch detection method described in section 4 looks for a corner in the image information caused by an attached sub-figure. In the case of a narrow sub-figure and a small branching angle this indicator can be mistaken for image noise.

The combined extraction of slab-like figures and adjoining tubular figures shows good initial results and is potentially valuable to the field of medical imaging. In addition to the demonstration in section 6.2 of the extraction of a kidney and renal artery, possible applications of this method include the heart and the liver along with adjoining vasculature.

## Acknowledgements

We are grateful to Stephen Aylward and Elizabeth Bullitt for their extensive work on tubular methods, which is the foundation for the methods presented here, as well as for insights on the extension to slabs. We also acknowledge Guido Gerig and James Damon for helpful discussions regarding our problem. We thank Edward Chaney for the CT datasets on which the methods were tested. This work was done under the support of NIH grants R01 GM68966, P41 EB002025 20 and P01 EB002779 NIBIB.

## References

- [1] Aylward, SR, E Bullitt (2002). Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction. *IEEE Transactions on Medical Imaging*, **21**: 61-75.
- [2] Blom, J (1991). *Affine invariant corner detection*. Ph.D. Thesis, Utrecht University.
- [3] Blum, H (1967). A transformation for extracting new descriptors of shape. *Models for the Perception of Speech and Visual Form*, W Whaten-Dunn, ed., MIT Press: 362-380.
- [4] Fridman, Y, SM Pizer, S Aylward, E Bullitt (2004, 2003). Extracting Branching Tubular Object Geometry via Cores. To appear in *Medical Image Analysis Special MICCAI issue*. This is an extended version of Segmenting 3D branching tubular structures using cores. *Medical Image Computing and Computer Assisted Intervention*, RE Ellis, TM Peters, eds., Lecture Notes in Computer Science, **2879**: 570-577.
- [5] Fritsch, DS, D Eberly, SM Pizer, MJ McAuliffe (1995). Stimulated cores and their applications in medical imaging. *Information Processing in Medical Imaging*, Y Bizais, C Barillot, R DiPaola, eds., Kluwer Series in Computational Imaging and Vision: 365-368.
- [6] Furst, JD (1999). *Height Ridges of Oriented Medialness*. Ph.D. Dissertation, Department of Computer Science, University of North Carolina at Chapel Hill.
- [7] Lindeberg, T (1994). *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Dordrecht, Netherlands.
- [8] Näf, M, O Kubler, R Kikinis, ME Shenton, G Székely (1996). Characterization and recognition of 3D organ shape in medical image analysis using skeletonization. *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*: 139-150.
- [9] Pizer, SM, PT Fletcher, S Joshi, A Thall, Z Chen, Y Fridman, D Fritsch, G Gash, J Glotzer, M Jiroutek, K Muller, G Tracton, P Yushkevich, EL Chaney (2003). Deformable m-reps for 3D medical image segmentation. *International Journal of Computer Vision Special UNC-MIDAG issue*, **55(2/3)**: 85-106.
- [10] Siddiqi, K, S Bouix, AR Tannenbaum, SW Zucker (1999). The hamilton-jacobi skeleton. *International Conference on Computer Vision*: 828-834.
- [11] ter Haar Romeny, BM (2002). *Front-end vision and multi-scale image analysis*. Kluwer Academic Publishers.
- [12] Witkin, AP (1983). Scale-space filtering. Proceedings of the Eight International Joint Conference on Artificial Intelligence: 1019-1022.