To appear, MedIA Journal, 2004
Special MICCAI issue

Extracting Branching Tubular Object Geometry via Cores

Yonatan Fridman, Stephen M. Pizer, Stephen Aylward, and Elizabeth Bullitt

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

Abstract. Blood vessels and other anatomic objects in the human body can be described as trees of branching tubes. The focus of this paper is the extraction of the branching geometry in 3D, as well as the extraction of the tubes themselves, via skeletons computed as cores. Cores are height ridges of a graded measure of medial strength called medialness, which measures how much a given location resembles the middle of an object as indicated by image intensities. Object bifurcations are detected using an affine-invariant corner detector and computations on the core's medialness values. The methods presented in this paper are evaluated on synthetic images of branching tubular objects as well as on blood vessels in head MR angiogram data. Results show impressive resistance to noise and the ability to detect branches spanning a variety of widths and branching angles. An extension that allows cores to extract general branching structures, not only branching tubes, is introduced.

Keywords: Segmentation; Branching; Tube; Core; Vessel.

1 Introduction

Three-dimensional medical images are often difficult for a physician to view both due to the noise introduced by the imaging process and because of the problems inherent in visualizing any volume image. Tasks such as surgical planning and guiding a catheter through a blood vessel tree can be facilitated with an accurate segmentation (Bullitt et al., 2001). This paper considers a method for computing such a segmentation using 1D cores of branching tubular objects, such as blood vessels, in 3D images of the quality found in MR angiograms. The paper also introduces an extension from 1D cores to 2D cores for the purpose of segmenting general (slab-like) branching objects. Cores are medial axes, at scale, computed using methods developed in the Medical Image Display and Analysis group at the University of North Carolina at Chapel Hill (Fritsch et al., 1995; Furst, 1999). Not only are cores insensitive to image noise and small object boundary perturbations (Morse et al., 1998), but they are computed directly from image grayscale information, so they can be used as an automated segmentation tool. Cores carry with them radius and orientation information of the object(s) they represent, providing additional geometric information (Pizer et al., 1998).

The method presented can be compared with other methods for the segmentation of tubular objects. Frangi et al. (1999) use a model-based method where an object's approximate medial axis is coupled to its boundary and then refined. Lorigo et al. (1999) use a second order level set method to rapidly segment a whole image. Vasilevskiy and Siddiqi (2002) use a geometric flow method in which a surface evolves under image-based constraints so as to cling to object boundaries. Aylward and Bullitt (2002) use a multi-scale image intensity ridge traversal method with branch handling. While the method of Aylward and Bullitt and the method presented in this paper are both based on cores, the former separately searches for position and width information and defines orientation implicitly, whereas the latter simultaneously determines position, width, and orientation information as a ridge of medialness.

The work presented in this paper is built upon the work of Furst (1999), Aylward and Bullitt (2002), and Aylward et al. (1996). It extends the work of Furst by improving the robustness of his core following and termination and by augmenting his method with the ability to handle branching objects. It differs from the work of Aylward et al. by relying on only one seed point per tree structure and no post-processing to segment a tubular branching object.

Section 2 of this paper describes a predictor-corrector method for marching along 1D cores of tubular objects in 3D images. Section 3 describes a method for determining when a core of a tubular object reaches a branch in that object. The cores defined in this paper mathematically do not branch (Miller, 1998), so once such a branch in an object is encountered, the method locates and follows the cores of the two new branches. Section 4 introduces a method for terminating core following when an object ends while still managing to traverse small breaks in the object. Section 5 provides results and evaluates methods applied both to synthetic images and to clinical images. Section 6 briefly introduces the extension of cores from 1D to 2D for extracting slab-like branching objects from 3D images. Section 7 gives a discussion and conclusions on the methods introduced.

2 Core Following

Before beginning core following, the image to be analyzed is pre-processed using the 3D oriented adaptive filtering technique of Westin et al. (2001) in order to suppress image noise while enhancing objects of interest. Once the image is filtered in this way, it is segmented using a marching algorithm based on medial atoms and medialness.

Medial Atoms. The method described in this section extracts a sampled core composed of discrete medial atoms. A medial atom $\underline{m}$ in 3D is a structure defined by four parameters ($\underline{x}$, r, $\underline{\mathbf{F}}$, θ) – the coordinates $\underline{x}$ of the atom in 3-space, the radius r of the object of interest at location $\underline{x}$, a frame $\underline{\mathbf{F}}$ that describes the orientation of the medial atom, and an object angle θ that describes the rate of widening or narrowing of the object (Fig. 1a). In this paper the object angle is constrained to $\pi/2$, assuming objects with nearly parallel sides. For a tube these four parameters imply a set of

concentric vectors, known as spokes, that extend from the medial location $\underline{x}$ to the implied object boundary. Ideally each medial atom would have an infinite number of spokes, but for computational efficiency there are a finite number of spokes greater than or equal to two. Results shown here are achieved using eight spokes, since tests have shown that fewer spokes produce inconsistent results, while more spokes increase running time without providing a significant improvement in results. Furthermore, the spokes are all constrained to be of the same length, imposing the limiting assumption that the object of interest is perfectly tubular. The method can also segment objects that are not quite circular in cross-section, but the more elongated the cross-section of the object of interest, the less likely it is to be correctly segmented.

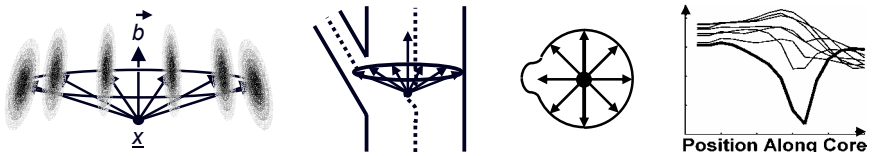

Fig. 1. (a) The set of concentric spokes of length r that make up a medial atom in 3D. The vector **b** is the core tangent direction. The atom shown has the angle θ between **b** and the spokes less than but not too far from π/2. Weighting functions are placed at the tip of each spoke. (b) A medial atom at a branch of a tubular object with cores shown as dotted lines. (c) A cross-sectional view of the tube and medial atom as a branch starts forming on the left. (d) The spoke at 9 o'clock (curve shown in bold) gives the weakest response near a branch among the eight spokes. This plot shows responses along a section of vessel from the image in Figure 6.

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. In 3D, this is given by the equation

$$\int_{\vec{v}\in V} D_{\hat{\vec{v}}} I(\underline{x}+\vec{v},\sigma) \qquad \textbf{(1)}$$

where **V** is the set of concentric spokes. σ 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).

In the method described here, to select a medial atom as being on the core, a position $\underline{x}_0$ is predicted by extending the path from the previous atom. Then medialness is optimized over the parameter space (r, $\underline{\mathbf{F}}$). The location $\underline{x}$ of the medial atom is then found by further optimizing medialness over the spatial plane that is normal to the

core tangent 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 predictor-corrector marching algorithm (Eberly, 1996) to extract a core, beginning with a manual initialization and then taking a step from a given medial atom i in the tangent direction of atom i to predict the location of atom i+1. The parameters (r, $\underline{\mathbf{F}}$) of atom i+1 are predicted by simply setting them equal to the parameters of atom i. Assuming the step size is small enough this is a reasonable approximation. Finally, the location, radius, and orientation of atom i+1 are corrected using the described optimization. For a more detailed discussion of cores see (Pizer et al., 1998).

3 Branch Handling

The cores defined in this paper do not have branch points in theory. As a core approaches a branch in a tubular object the core does not bifurcate to follow both new branches, nor does the core follow either of the two branches individually. Rather, the core bisects the branch. Initializing core following at each of the three branches results in three separate cores, each of which approximately bisects the other two branches (Fig. 2). These three cores all pass through the object's branching region but they generically do not have a mutual intersection point (Damon, 1998). To handle this, the core following method attempts first to detect when the core being followed reaches the object's branching area and subsequently to locate the two new branches and to resume following each of the cores of these new branches.

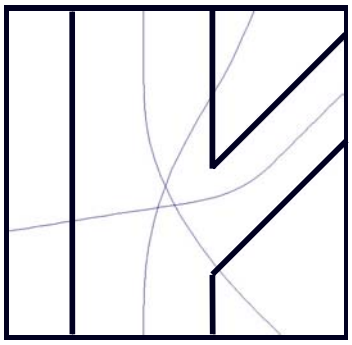

Fig. 2. A synthetic image showing the outline of an ideal branching object in 2D. The three blue curves are the three cores of this object as computed starting from the outer end of each of the three branches.

3.1 Branch Detection

Branch detection is accomplished in two steps. The first step is to apply an affine-invariant corner detector, $L_{\mathbf{uu}}L_{\mathbf{v}}$ (Blom, 1991; Lindeberg, 1994; ter Haar Romeny, 2002), to the three-dimensional image, where **v** is the image gradient direction and **u**

is the eigenvector corresponding to the largest eigenvalue of the Hessian of L in the plane normal to **v**. This operator is applied at the same scale at which medialness is computed (see section 2). This consistently gives strong responses at corners on the boundary of branches, but it also gives strong responses in undesirable places such as the inside edge of a sharply bending object or places that have voxel jaggedness due to noise or aliasing (Fig. 3). Medial atoms whose spoke tips are at maxima of cornerness are potential branch points. Once these potential branch points are determined the second step is to identify and discard false positives as follows.

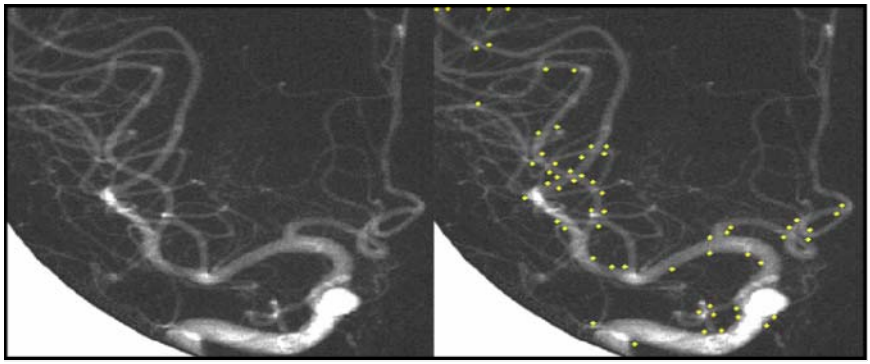

Fig. 3. (a) A 2D DSA projection image of the head. (b) Local maxima of the affine-invariant corner detector $L_{uu}L_v$ (*yellow dots*) displayed on the original image, where only maxima above a given (arbitrary) threshold of cornerness are shown. This example is shown in 2D to facilitate visualization; similar results are achieved in 3D.

In the discrete computation as a core nears a branch, it begins bisecting the branching angle, but then typically it snaps back to one of the two children (Fig. 1b). In snapping back to one of the children, the core jumps from a local optimum in medialness at the original core to a stronger optimum at the new core. Just before this snapping occurs, the responses of one or two spokes decrease dramatically with respect to the rest of the spokes (Aylward and Bullitt, 2002) (Fig. 1c, d). If this decrease does not occur, the potential branch is rejected as a false positive. If this decrease does occur, core following must resume at the two child branches.

3.2 Continuation of Core Following

Since the core automatically snaps to one of the two branches, the algorithm only needs to locate the other branch. As described, the medial atom spoke that lies on the intersection of the two branches gives a significantly weaker response in medialness (see Fig. 1). The location of this spoke tip provides an approximate starting location for the core of the new branch and the direction in which it points provides an approximate tangent direction for the new core. Based on the assumption that the area of the tube cross-section splits between its two branches, the radius of the new core is estimated as ½ times the radius of the original core. Once the new medial atom is

estimated it is refined by the corrector portion of the predictor-corrector method. These estimates are sufficiently accurate for the corrector portion of the method to locate the new core, as shown in section 5.1.

4 Core Termination

When marching along an object's core there are two possible reasons to terminate core following: 1) The object being segmented comes to an explicit end, where the object is clearly capped; or 2) the object being segmented comes to an implicit end due to a decreased signal-to-noise ratio, as in the case of a narrowing blood vessel. As the signal-to-noise ratio becomes small or the width of the blood vessel approaches the inner resolution of the image, the ability to follow the core decreases. This case of implicit ends is the more difficult one to handle. Thus, although the method presented has been tested on both types of object ends, only implicit ends are discussed here.

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. At each step along the core 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 flagged. If at least four out of five consecutive medial atoms are flagged as having weak medialness, core following is terminated. Four weak medial atoms are required to terminate core following rather than one to allow for traversal of small breaks that can occur either because of an actual break in the object or because of image noise. Similarly, the method allows for the possibility that one out of five medial atoms gives a strong response due to background noise even when the object is no longer being followed.

5 Results

The method described was tested in two ways. First, several sets of synthetic images were created to simulate blood vessels. Each set of synthetic images was created to test a different part of the method (core following, branch handling, and termination) and each of these parts of the method was tested separately. Second, the method was tested as a whole on clinical MRA data of the head.

5.1 Results in Synthetic Images

The basic core following, without accounting for branching and ending, is quite robust in the presence of image noise. Figure 4 shows one slice of a three-dimensional image of a curved binary tube that has a radius of three voxels. Gaussian noise is added to the image with a magnitude per voxel of approximately five times what would be expected in an MRA (Aylward and Bullitt, 2002). Even in such a noisy image the core of the entire tube was extracted consistently over several test runs. Similar results were achieved for tubes of different sizes, curvatures, and torsions. In images with a level of noise commonly seen in an MRA cores were extracted consistently from tubes with a radius of slightly less than one voxel.

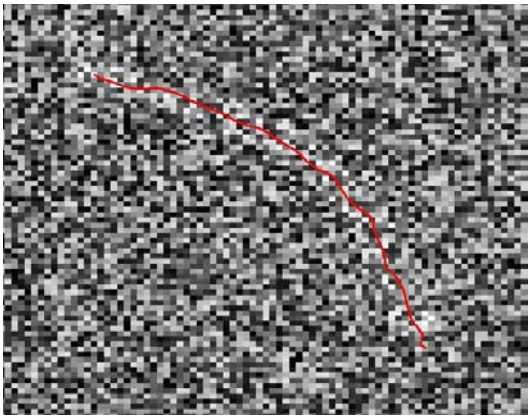

Fig. 4. One slice of a noisy synthetic 3D image of a tube of radius three pixels whose axis lies entirely in the plane displayed. The red curve is the core of the tube.

Branch handling was tested on a set of 36 simulated images of branching tubular structures spanning a range of branching angles, branch sizes, and background noise levels. The code was run on each image ten times. A branching angle of 0 degrees in Figure 5 indicates that the two child branches have axes that share a tangent direction at the point where they branch and which then diverge. The branch size labels indicate the cross-sectional area of the smaller of the two child branches relative to that of the parent, where the cross-sectional areas of the two child branches sum to the cross-sectional area of the parent. Three different levels of Gaussian noise were added to the image – low (approximately half of what would be expected in an MRA), normal (approximately what would be expected in an MRA), and high (approximately twice what would be expected in an MRA). The code detected the branch correctly in all 120 test runs with low background noise. Of the 120 runs with normal background noise the branch was missed on three runs, all three of which had a small branch size (0.1 times the cross-sectional area of the parent). Of the 120 test runs with high background noise the branch was missed on 24 test runs, 17 of which had a small branch size. The results for the 120 high noise test runs are summarized in Figure 5.

In the cases where the branch point was detected correctly, continuation of core following in the two new branches was accomplished correctly 98% of the time.

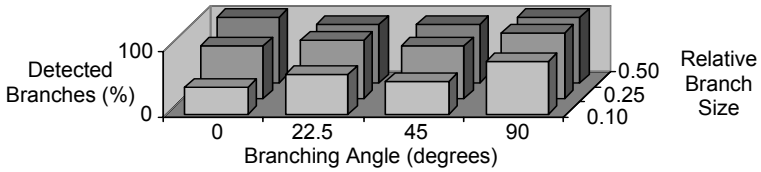

Fig. 5. The success of the branch detection code on synthetic images with a noise level of approximately twice what would be expected in an MRA. Each bar shows the percentage of test runs on which the branch was correctly detected for a given branch angle and a given branch size.

As described in section 4, core following should terminate when there is a significant decrease in the image signal-to-noise ratio. The difficulty is in following the object as long as possible but stopping when there is no basis in the image information for following the object further. The method was tested on a set of synthetic images of tubes of varying curvature and torsion that gradually narrow to a radius of zero. Varying amounts of image noise were added. In none of the images did the core continue following past where the tube came to a point. In 87% of the cases the core stopped within one voxel of the actual end of the tube. In the remaining 13% of the cases the core stopped between two and five voxels early.

5.2 Results in Clinical Images

Core following with branch handling and core termination was tested on three clinical MRA images of the head acquired using a Siemens 3T system. Figure 6 shows such an MRA and the corresponding 3D cores model automatically extracted from the MRA. This particular image is 512 x 512 x 91 voxels with a resolution of 0.42 x 0.42 x 1.25 mm.

The core following algorithm was manually initialized in three places – the left and right carotids and the basilar artery. These three arteries are clearly visible at the bottom of the coronal view of the model. From these three initializations branch handling and core termination were employed without further user interaction to achieve the results shown. Running time was four hours on a 2 GHz PC in the MATLAB environment. Initial tests show that the same computation would take approximately 30 minutes when implemented in C++. Figure 7 displays additional detail as the computed core traverses a branch of a narrow vessel in the same MRA.

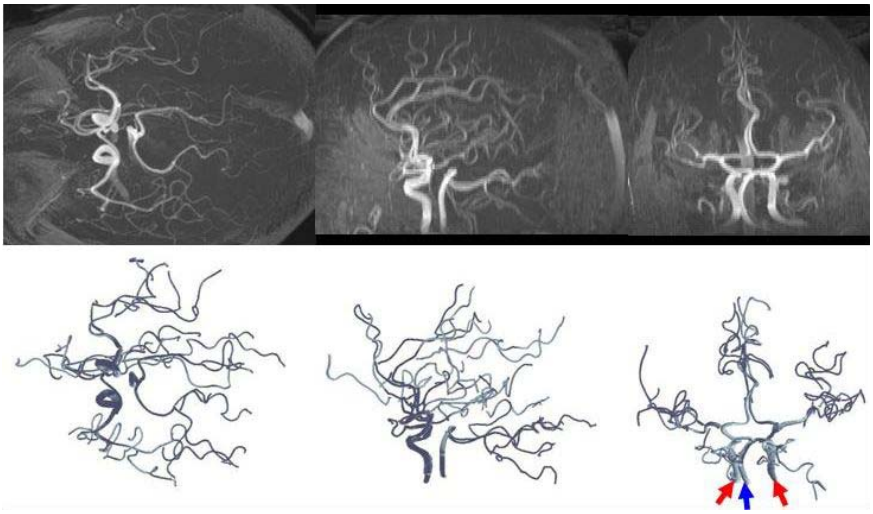

Fig. 6. Maximum intensity projection images of a 389 x 300 x 91 voxel portion of an MRA of the head (*top*) and corresponding 3D cores model (*bottom*) shown from slightly rotated axial (*left*), sagittal (*middle*), and coronal (*right*) viewpoints. Arrows indicate right and left carotid arteries (*red*) and basilar artery (*blue*).

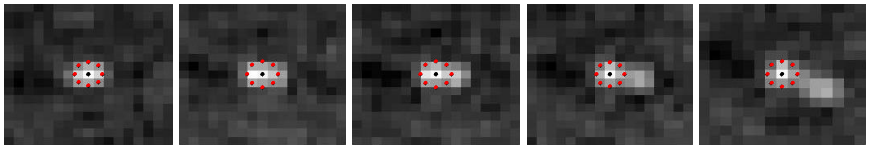

Fig. 7. Five slices through a 3D image of a branching vessel, where each slice is taken normal to the tangent direction of the computed core. In each slice the central black dot shows the location $\underline{x}$ of the medial atom and the eight red dots show the locations of the eight spoke tips. These images are taken on five consecutive iterations of the core following algorithm as the branching region is traversed. In the second image the blood vessel begins to widen as the branching region is entered, and by the fifth image the two new branches have nearly separated. In this example the core snaps to the brighter of the two branches, as in Figure 1b.

6 2D Cores for Slab-Like Objects

Cores of branching tubular structures as presented in this paper show impressive resistance to noise, benefiting in part from the constraints imposed on the tube extraction process. To be of general applicability, though, these constraints must be relaxed to allow for the extraction of non-tubular, or slab-like, branching structures. Such cores of slabs show promise for the analysis of a variety of anatomic structures including the pelvis, the liver, and the kidneys with adjoining vasculature. This section introduces the use of cores for extracting such slab-like structures.

When generalizing cores in this manner two major changes are made from the tubular cores presented. First, the structure of a medial atom must be altered since the set of concentric spokes described in section 2 assumes the extraction of tubular objects. A medial atom used to extract a slab-like structure has only two spokes, which are collinear when the object angle is constrained to $\pi/2$ (Fig. 8). These spokes are vectors having their tails positioned at the medial atom location $\underline{x}$ and their heads positioned at opposing object boundary locations. Such a configuration assumes that the 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 tracking becomes less robust. This effect only becomes significant near an object angle of $\pi/4$, allowing for effective tracking of most anatomic objects.

The other major change that is made for slab-like objects regards the dimensionality of the computed cores. Whereas cores of tubes in 3D images are necessarily 1D, i.e., space curves, cores of slabs are 2D. This change has the implication that the marching algorithm presented in section 2 gains a degree of freedom and increases in complexity. Instead of simple 1D marching, the method must now march in two mutually orthogonal directions that are also orthogonal to the spokes and that define the tangent plane to the core (Fig. 8).

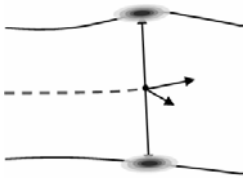

Fig. 8. Marching along a slab-like core in 3D, in which case the core is 2D and each medial atom has only two collinear spokes.

To handle this increased complexity, a coarse-to-fine grid-based marching algorithm has been developed. Other changes are made to the tubular core extraction method, including the addition of a penalty term to Equation 1 that penalizes changes in geometry between adjacent atoms. Branch handling is augmented to deal with slabs, while the method for core termination remains virtually unchanged.

Results appear promising but are not as robust as in the tubular case. Figure 9 shows an example of a 2D core of a kidney computed from a 3D CT image, integrated with the 1D core of the adjoining renal artery, which was automatically detected and extracted using branch handling techniques. Further algorithm details, results, and analyses of the properties of slab-like cores will be presented in a future paper.

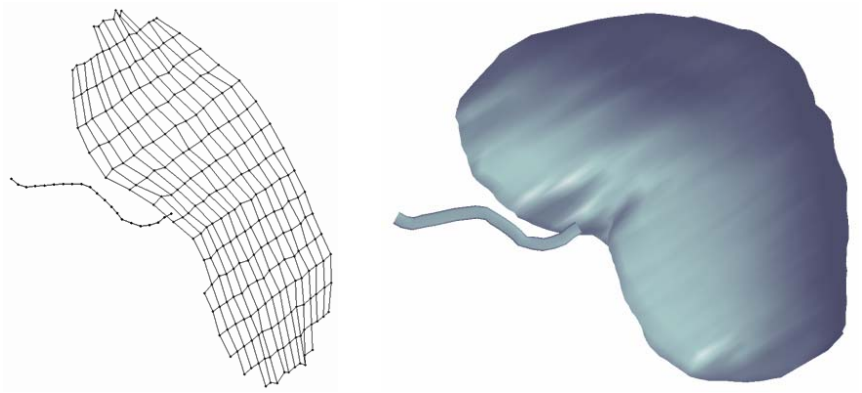

(a) (b)

Fig. 9. (a) 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, computed from CT. Medial atoms are shown as dots connected by line segments. (b) The 3D surfaces implied by the computed cores, viewed from a slightly different angle.

7 Conclusions and Discussion

This paper describes the use of cores for extracting the geometry of branching tubular structures from three-dimensional grayscale images, with a focus on branches. The method described is robust and capable of handling a wide variety of tubular structures in noisy images. For the case of cerebral vasculature, results for vessel segment extraction are highly resistant to noise, consistently tracking vessel segments in images with a level of noise several times what would be expected in an MRA. This robustness is primarily due to the use of Gaussian derivatives as edge detectors at scale and the combination of constraints imposed on the structure of each medial atom.

The branch handling method developed in this paper is fairly robust, allowing for the automated extraction of tree-like structures. A limitation of the method is its inconsistency in correctly detecting narrow tubes that bifurcate from wider tubes. Tubes with a diameter of one to two voxels are occasionally missed, mainly in cases with a small branching angle and a low signal-to-noise ratio in the vicinity of the bifurcation. This occurs because the branch detection method described in section 3.1 looks both for a corner in the image information and for a significant break in the tube wall caused by a bifurcation. In the case of a narrow tube and a small branching angle these two indicators can be mistaken for image noise. Nonetheless, such omissions are rare, and false positive branches are nearly nonexistent.

A secondary contribution of this paper is automatic core termination in objects that do not have explicit ends. The method described effectively traverses small breaks in the object or small areas of weak image information without losing track of the object. The core termination method has shown good results on a variety of images, although

it cannot handle the extreme level of noise that the basic tubular core following method can handle.

Finally, methods for extending cores to 2D for the purposes of extracting geometric information from general branching slab-like objects show promise.

Acknowledgements

We are grateful to Russell M. Taylor II for driving problems and helpful discussions and to James Damon for mathematical insights into our problem. We also acknowledge Carl-Fredrik Westin and Raul Estepar, Brigham and Women's Hospital, for their gracious help in image filtering. This work was done under the support of ONR MURI grant N00014-98-1-0597 as well as the partial support of NIH grants R01 EB000219 NIBIB, R01 HL69808 NIHLB, P01 CA47982 NCI, and P01 EB002779 NIBIB.

References

[1] Aylward, SR, SM Pizer, E Bullitt, D Eberly (1996). Intensity ridge and widths for tubular object segmentation and description. *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, **56**: 131-138.

[2] 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.

[3] Blom, J (1991). *Affine invariant corner detection*. Ph.D. Thesis, Utrecht University.

[4] Bullitt, E, SR Aylward, K Smith, S Mukherji, M Jiroutek, K Muller (2001). Symbolic description of intracerebral vessels segmented from magnetic resonance angiograms and evaluation by comparison with X-ray angiograms. *Medical Image Analysis*, **5**: 157-169.

[5] Damon, JN (1998). Generic structure of two-dimensional images under Gaussian blurring. *SIAM Journal on Applied Mathematics*, **59(1)**: 97-138.

[6] Eberly, D (1996). Ridges in image and data analysis. *Computational Imaging and Vision Series*. Kluwer Academic Publishers, Dordrecht, Netherlands, 1996.

[7] Frangi, AF, WJ Niessen, RM Hoogeveen, T van Walsum, MA Viergever (1999). Model-based quantitation of 3D magnetic resonance angiographic images. *IEEE Transactions on Medical Imaging*, **18**: 946-956.

[8] 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.

[9] Furst, JD (1999). *Height Ridges of Oriented Medialness*. Ph.D. Dissertation, Department of Computer Science, University of North Carolina at Chapel Hill.

[10] Lindeberg, T (1994). *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Dordrecht, Netherlands.

[11] Lorigo, LM, O Faugeras, WEL Grimson, R Keriven, R Kikinis, CF Westin (1999). Co-dimension 2 geodesic active contours for MRA segmentation. *Information Processing in Medical Imaging*, A Kuba, M Sámal, A Todd-Pokropek, eds., Lecture Notes in Computer Science, **1613**: 126-139.

[12] Miller, JE (1998). *Relative Critical Sets and their Application to Image Analysis*. Ph.D. Dissertation, Department of Mathematics, University of North Carolina at Chapel Hill.

[13] Morse, BS, SM Pizer, DT Puff, C Gu (1998). Zoom-invariant vision of figural shape: effects on cores of image disturbances. *Computer Vision and Image Understanding*, **69**: 72-86.
[14] Pizer, SM, D Eberly, BS Morse, DS Fritsch (1998). Zoom-invariant vision of figural shape: The mathematics of cores. *Computer Vision and Image Understanding*, **69**: 55-71.
[15] ter Haar Romeny, BM (2002). *Front-end vision and multi-scale image analysis*. Kluwer Academic Publishers.
[16] Vasilevskiy, A, K Siddiqi (2002). Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **24**: 1565-1578.
[17] Westin, CF, L Wigstrom, T Loock, L Sjoqvist, R Kikinis, H Knutsson (2001). Three-dimensional adaptive filtering in magnetic resonance angiography. *Journal of Magnetic Resonance Imaging*, **14**: 63-71.
[18] Witkin, AP (1983). Scale-space filtering. *Proceedings of the Eight International Joint Conference on Artificial Intelligence*: 1019-1022.