

Initialization, Noise, Singularities, and Scale in Height Ridge Traversal for Tubular Object Centerline Extraction

Stephen R. Aylward¹ and Elizabeth Bullitt²

¹Department of Radiology

²Department of Surgery

Computer-Aided Diagnosis and Display Lab

Medical Image Display and Analysis Group

University of North Carolina at Chapel Hill

Please direct correspondence to

Stephen R. Aylward, PhD

Department of Radiology, CB 7515

Mason Farm Road, Radiology Research Lab, MRI Building D

University of North Carolina at Chapel Hill

Chapel Hill, NC 27599-7515

(919) 966-9695

aylward@unc.edu

Support

This work was supported in part by NIH-NCI-P01A47982, NIH-NCI-R01CA67812, a Microsoft equipment and software award, an Intel equipment award, and research funds from the Department of Radiology at UNC. Portions of this software have been licensed to Surgical Navigation Technologies, Medtronic (Minn., Minn.) and to R2 Technologies, Inc. (Los Altos, CA).

INITIALIZATION, NOISE, SINGULARITIES, AND SCALE EFFECTS ON HEIGHT RIDGE TRAVERSAL FOR TUBULAR
OBJECT CENTERLINE EXTRACTION

Abstract

The extraction of the centerlines of tubular objects in two and three dimensional images is a part of many clinical image analysis tasks. One common approach to tubular-object centerline extraction is based on intensity ridge traversal. In this paper we evaluate the effects of initialization, noise, and singularities on intensity ridge traversal and present multi-scale heuristics and optimal-scale measures that minimize these effects. Monte Carlo experiments using simulated and clinical data are used to quantify how these “dynamic-scale” enhancements address clinical needs regarding speed, accuracy, and automation. In particular, we show that dynamic-scale ridge traversal is insensitive to its initial parameter settings, operates with little additional computational overhead, tracks centerlines with sub-voxel accuracy, passes branch points, and handles significant image noise. We also illustrate the capabilities of the method for medical applications involving a variety of tubular structures in clinical data from different organs, patients, and imaging modalities.

Keywords

Vessel, Tube, Skeleton, Ridge

INITIALIZATION, NOISE, SINGULARITIES, AND SCALE EFFECTS ON HEIGHT RIDGE TRAVERSAL FOR TUBULAR
OBJECT CENTERLINE EXTRACTION

I. INTRODUCTION

Many surgical and interventional radiology tasks involve visualizing and quantifying vessels, bronchi, bowels, ducts, or nerves in two and three dimensional medical data. These and other anatomic objects are tubular, i.e., they have nearly circular cross-sections and smoothly varying widths and possibly follow tortuous paths and branch. Clinical applications involving tubular anatomy in magnetic resonance angiogram (MRA), x-ray computed tomography (CT), and 3D ultrasound data include shunt length specification, virtual colonoscopy flight-path control, embolization planning, stenosis detection, radiation therapy treatment planning, and MRA / digital subtraction angiogram registration for intra-operative guidance. In this paper we present a fast and accurate method for modeling the centerlines of tubular anatomy so that computer-based methods can be more easily developed to assist clinicians with these tasks.

We begin by discussing a general method for intensity-ridge traversal for centerline extraction. We then present dynamic-scale enhancements that improve the handling of noise and local singularities (e.g., branchpoints) while decreasing the method’s sensitivity to its initial parameter values. We have found these modifications to be essential for the routine modeling of centerlines in clinical images. We support these conclusions by presenting results from Monte Carlo experiments that compare the “default” ridge traversal method with our “dynamic-scale” method using simulated and MRA data. We also provide illustrations of the application of the dynamic-scale method to clinical data from different organs, patients, and modalities. We show that the dynamic-scale method can extract the centerlines of the a variety of tubular objects occurring in head MRA, lung CT, liver CT, and liver 3D ultrasound data.

A. Specifics

It is our philosophy that centerline extraction should be an integral part of tubular object modeling. Compared to edge identification which is generally performed using small-scale measures, centerline identification can be performed by integrating over a large extent of a tubular object and therefore may be less sensitive to image noise [16]. Given a centerline, subsequent processing such as boundary detection can be stabilized by providing a spatial prior to resolve boundary ambiguities. We have also found the centerline’s geometry to provide an excellent basis for multi-modal image registration [2], [3].

In this paper, we explore the roles of initialization, noise, singularities, and scale in ridge-based methods for tubular-object centerline extraction. Ridge methods operate by considering an N-dimensional image to be a surface in an (N+1)-dimensional space by mapping intensity to height; tubular object centerlines will exist as 1-dimensional height ridges on that surface. The effectiveness of height-ridge traversal for centerline modeling is application dependent, but to quantify the general capabilities of these methods we identified three common criteria which are described next.

Speed: Many clinical imaging protocols (especially those involving multi-detector CT and high-gradient MR scanners) acquire pre and post-contrast images at sub-millimeter spatial resolutions. As a result, multiple mega-voxel datasets may be generated per patient per hospital visit. Our goal is to automatically process these datasets within thirty minutes; thereby making our approach applicable to all but the most urgent-care clinical situations. The presented methods are

able to extract a centerline passing thru 20 voxels in ~ 0.3 seconds on a 733Mhz Pentium III PC running Windows 2000. A patient's intracranial vasculature captured by MRA can be modeled in under 15 minutes.

Accuracy: Critical to any tubular object centerline extraction technique is the robust handling of image noise, branch points, and widths approaching the inner-scale of the data. For example, interventional radiologists require the definition of 1 mm feeding vessels for the embolization of tumors and arteriovenous malformations, and yet many MR angiograms are acquired using 1x1x1 mm voxels. Even in noisy images, the methods presented in this paper can accurately extract the centerlines of one-voxel-wide tubes: average error is less than 0.5 voxel, branchpoints are consistently passed, and the full length of the centerline is consistently extracted.

Automation: The performance of many extraction techniques is highly dependent on the initial specification of multiple parameter values. Automation is greatly facilitated if a method has few parameters and if it is insensitive to the initial value of those parameters. The minimal information required to designate a centerline-of-interest is a localization of the tube in image space $x' \in \mathbb{R}^N$ and scale space σ' . Using the dynamic-scale ridge traversal method, centerline extraction accuracy is not statistically significantly dependent on the values of x' and σ' . Automatic processing of lung CT, head MRA, and liver CT data is demonstrated.

For this paper, the default and dynamic-scale methods were evaluated via Monte Carlo experiments using simulated data and clinical MRA data. For the simulated data, we generated images at three different noise levels containing a tortuous tubular object that branches and whose radius varies from 0.5 to 4.0 voxels. For 200 random starting locations x' using three different starting scale values σ' we measured the time to extract a centerline passing through 20 voxels in order to quantify speed. We measured the average error, maximum error, percent of points within 0.5 voxel of ideal, and percent of points within one voxel of ideal as measures of accuracy. We measured ridge length and frequency of branchpoint crossing as measures of automation. Using these measures we tested the Null Hypotheses that (a) there is no difference between the mean performance of each method and (b) there is no difference between the mean performance of a method given different starting scale values. Levels of significance were measured. In summary, we concluded that the dynamic-scale enhancements significantly improve performance, and the enhanced implementation's performance is not significantly effected by its initial scale value. For the MRA data, given 300 random starting points near a vessel, any two extractions are within 1/10th of a voxel of one another.

At the end of this paper we provide illustrations of the application of the dynamic-scale method to clinical applications involving different tubular structures, organs, patients, and imaging modalities. For detailed discussions on specific clinical uses of our centerline modeling method and related technology developments in visualization, 3D/2D registration, 3D/3D registration, vascular connectivity description, and surgical cutting path analysis, the reader should consult our previous publications [7], [8], [9] and our web pages at <http://caddlab.rad.unc.edu>. The data used in the experiments presented in this paper are also available on those web pages.

II. BACKGROUND

Centerline models can be generated (a) explicitly, (b) implicitly, or (c) via post-processing by vessel modeling methods. An analysis of several previously published methods from each of these categories follows and is summarized in Table I. It is important to realize that this analysis is inherently biased against methods for which centerlines are implicit or

defined via post-processing. Those methods were not intended to be directly used for centerline modeling. They are included in this analysis, however, because of their potential for centerline modeling. Additionally, for this review, we only considered methods that do not assume a global structure to the tubular network being extracted. The topologies of brain, lung, and liver vascular networks have too much inter-patient variability to be fit by global models.

(a) Explicit centerline extraction has existed for decades as a basis of tubular object modeling [5], [16], [22]. Explicit centerline extraction is the basis of our approach to tubular-object modeling [1]. Our method starts from an initial point on or near a vessel and subsequently performs a multi-scale extraction of the vessel’s centerline via ridge traversal and width estimation. The dynamic-scale enhancements in this paper grew in part from the “core” work of Stephen Pizer [11], [13], [19], [21]. Core methods simultaneously solve for middles and widths during traversal and have excellent theoretic insensitivity to noise by using measurement apertures (i.e., scales) proportional to the width of the object, but core finding can be hampered by having to seek extrema in high-dimensional, correlated parameter spaces. For tubular object centerline modeling, we have found that by uncoupling the middle and width estimation processes and yet continuing to use apertures proportional to the object’s width (implimented in the dynamic-scale ridge traversal method), most of the insensitivity properties remain and finding extrema in the separate optimization spaces is tractable and fast. Other groups have also developed explicit centerline modeling methods. For example, Gerig [14] developed steerable filters to track the centers and widths of vessels. More recently, Frangi and Niessen [12] have presented a method that given a pair of endpoints use the connected vessel voxels (defined by a binary object mask) to coarsely define a path that is iteratively refined to determine a centerline (Bitter [4] has published a similar method that does not require endpoint specification.) For some applications, endpoint and object mask specification is appropriate – an excellent example is abdominal aortic-aneurysm stent planning using contrast CT. For other applications, however, these initialization requirements might be cumbersome and hinder automation. For example, given complex, branching tubes, such endpoint identification or object segmentation may actually be the goal subsequent to centerline extraction.

(b) Tubular object segmentation methods in which a centerline extraction is implied include anisotropic diffusion [10], [20] and region growing / level-set evolution [17], [18], [27] methods. These techniques have produced excellent segmentations of intracranial vasculature and bronchi. One of the most interesting of these implicit methods is the level-set method developed by Lorigo [17]. It uses 2nd-order level-set information to rapidly extract entire vascular trees. However, it can be very difficult to control (e.g., specify the stopping criteria for) certain diffusion and region growing methods when attempting to model vessels that exist at the inner scale of the data.

(c) Post-processing is necessary to generate a centerline representation when tubular objects are identified on a voxel-by-voxel basis. One well-developed method for voxel labeling uses spatially-adaptive histogram analysis to automatically determine local thresholds to identify vessel voxels in MRA data [25]. That method and most voxel-labeling methods are particularly useful when the vascular structure of interest is non-tubular. For example, that method has demonstrated clinical utility for aneurysm visualization and quantification for treatment planning [26]. The absence of shape constraints, however, makes handling noise and small vessels more difficult. Additionally, if centerlines are desired, skeletonization must be applied to the labeled voxels, and some skeletonization methods such as binary erosion should be avoided since they will only produce voxel-level centerline representations.

Method	Centerline	Speed	Centerline Accuracy	Automation	Application, Organ & Modality
Aylward [this paper]	Explicit: Centerline traversal	2	3	3 Application-specific seed-point identification	Surgical Planning and Guidance - Arbitrary tubes - Arbitrary modality
Frangi & Niessen [12]	Explicit: Centerline refinement	2	2 No small vessels	1 Both endpoints & object segmentation required	Abdominal Aortic Aneurysms - Arbitrary tubes probable - Modality limited by obj. seg.
Lorigo [17]	Implicit: Level-set evolution	3	2 May be difficult to balance noise with small vessels	3 Application-specific seed-point identification	Neurosurgical Planning - Arbitrary tubes probable - Arbitrary modality
Wilson & Nobel [25]	Post-Process: Adaptive thresholding	3	2 Small vessels difficult & centerline requires thinning	2 Cannot limit to extracting tubes	Neurosurgical Planning - Non-tubular tubes possible - Modality limited by ROI seg.
Global Thresholding	Post-Process: Global thresholding	3	1 No small vessels & centerline requires thinning	1 Cannot limit to extracting tubes	Coarse Surgical Planning - Non-tubular tubes possible - Modality limited by ROI seg.
Hand Labeling	Explicit: Hand labeling	0	0 Poor localization and connectivity	0 None possible	None - Arbitrary tubes possible - Arbitrary Modality

TABLE I

A REVIEW OF THE CENTERLINE EXTRACTION CAPABILITIES OF A RANGE OF PREVIOUSLY PRESENTED VESSEL MODELING METHODS. THIS REVIEW INCLUDES METHOD PRESENTED IN THIS PAPER (FIRST ROW). NUMERIC SCALE: 0=UNSATISFACTORY, 1=GOOD, 2=BETTER, 3=BEST. NOTE: THIS ANALYSIS IS BIASED AGAINST IMPLICIT AND POST-PROCESS CENTERLINE GENERATION METHODS - THOSE METHODS WERE NOT INTENDED FOR CENTERLINE MODELING, AND THIS ANALYSIS DOES NOT REFLECT THEIR TUBULAR-OBJECT MODELING CAPABILITIES.

III. CENTERLINE EXTRACTION

To introduce the dynamic-scale traversal method we begin by specifying our criteria for a point to be considered on a height ridge. We then detail the importance of scale to ridge formation and traversal. Subsequently we present the default centerline modeling method: given a seed point x' and scale σ' , a point on the local ridge is found, and the full extent of the ridge is traversed by a step-maximize procedure. We then present our dynamic-scale enhancements which involve multi-scale heuristics and tube width estimation to implement an optimal-scale step-maximize ridge traversal procedure.

A. Ridge Criteria

Without loss of generality, we assume that tubes-of-interest appear brighter than the background, that is, their centerlines are intensity ridges not valleys. Via a sign-change, the following methods can be applied to track valleys. Via

blurring, filtering, and other image processing operations these methods can be applied to extract tubular objects that do not have central extrema, are defined by texture, or are differentiated from their surround by other image features. In this paper we focus on the processing of three-dimensional data, but these methods have also been applied to two and four-dimensional data – our notation guides the application of this method to N-dimensional data. We define:

I	an N-dimensional volume of data
x	a point in \mathfrak{R}^N on a 1D ridge in I
σ	the scale at which measures in I at x are made; a scale at which the ridge exists
∇I	the gradient vector of I at x
\vec{t}	the ridge's tangent direction at x

We are specifically interested in maximum convexity height ridges, and so the basis directions normal to the ridge are defined using the eigenvectors of the Hessian matrix at x . We define:

$\vec{v}_1, \vec{v}_2, \vec{v}_3$	the eigenvectors of the Hessian of I at x
$\alpha_1, \alpha_2, \alpha_3$	the eigenvalues of $\vec{v}_1, \vec{v}_2, \vec{v}_3$ with vectors and values ordered such that $\alpha_1 \leq \alpha_2 \leq \alpha_3$

Two conditions must hold for x to be on a maximum-convexity height ridge (see [11], [16] for additional detail):

One, the point x must be a ridge point. Second-derivative information distinguishes ridges from valleys, saddles, planes, and spheres. The Hessian matrix captures second-derivative information at a point in an image. For a point to be on a 1D ridge of an ND surface, N-1 of eigenvectors of the Hessian of I at x must have negative eigenvalues. Given ordered eigenvalues, we test this condition by verifying

$$\alpha_1 \leq \alpha_2 < 0 \quad (1)$$

By this condition, the directions \vec{v}_1 and \vec{v}_2 are assumed to be the directions normal to the ridge (this assumption is relaxed in the dynamic-scale implementation - Section III-F).

Two, the point x must be on an (N-1)-dimensional extreme, i.e., the projection of the gradient at x onto the N-1 directions normal to the ridge must be equal to zero

$$\vec{v}_1 \bullet \nabla I = 0 \text{ and } \vec{v}_2 \bullet \nabla I = 0 \quad (2)$$

Following the work by Eberly [11], to define the ridge's tangent direction we avoid third-derivative calculations and instead approximate the tangent \vec{t} of the ridge at x by \vec{v}_3 the remaining (the maximum eigenvalued) eigenvector of the Hessian of I at x .

We also define a third condition to assure that the ridge is central to a tubular object that has a nearly circular cross-section. We have found this third condition to be essential for proper ridge termination. As a tube becomes dim and its borders diffuse relative to its background (as occurs where blood vessels feed capillary beds), a ridge may still exist but the tubular structure becomes poorly defined and often appears more elliptical than circularly symmetric. We test this condition by the ratio of the eigenvalues in the directions normal to the ridge.

$$\frac{\alpha_2}{\alpha_1} \geq 1 - \epsilon \quad (3)$$

This ratio is one when the tube's cross-section is symmetric. We allow for deviations from perfect circular symmetry via ϵ .

In this manner, there are four parameters in our ridge criteria (Eqs. 1-3): how to test for equal-to-zero in Equ. 2; the value of ϵ in Equ. 3; the location x of the ridge point being tested; and the scale σ at which the ridge at x exists, i.e., the scale at which the measures at x are made.

For all experiments and applications presented in this paper, as the test for equal to zero in Equ. 2 we used

$$\sum_{i=1}^{N-1} (\nabla I \bullet \vec{v}_i)^2 < 0.0001 \quad (4)$$

As formulated, this test enforces orthogonality at places of large gradient, but as the gradient magnitude goes toward zero, less orthogonal directions are tolerated. Specifically, all data is normalized to be between 0 and 1. If the scaled gradient is 10% of the intensity range near a ridge, then ∇I and \vec{v}_i are said to be orthogonal if their angle is bounded by $\text{invcos}(\pm 0.0001/0.1) = 0.5\pi \pm 0.003$, yet, if the scaled gradient is 0.1% of the intensity range, orthogonality is relaxed to $\text{invcos}(\pm 0.0001/0.001) = 0.5\pi \pm 0.319$. This reflects the fact that as a tube becomes dim and “flat”, it is more difficult to know gradient and normal directions with confidence; however, because of the flatness of the ridge, less accurate gradient and normal directions are required to stay “on” the ridge. To counter-balance this effect in extreme conditions (without introducing an application dependent threshold), Equ. 3 was developed.

In Equ. 3, for all experiments and applications in this paper, we use 0.5 for ϵ . This allows cross-sectional intensities to deviate from circular and resemble an ellipse with a 2:1 ratio between the lengths of the major and minor axis.

The focus of this work is determining the remaining ridge parameters, x and σ . The default implementation uses a fixed σ while localizing x . The dynamic-scale implementation dynamically optimizes σ while localizing x . The importance of scale is discussed next.

B. Scale: The Critical Parameter

The critical parameter of our ridge criteria is the scale at which the intensity, gradient, and Hessian of I at x are calculated. This scale serves three purposes. One, it designates the tube-of-interest. Without the specification of a scale, for example, it may be unclear if the user seeks to define a blood vessel within an arm, the arm, or the entire torso of the person imaged. Two, scale provides insensitivity to image noise. Measures made at larger scales are less effected by high frequency variations in the data (Fig. 1). Three, scaled measures such as blurring may be necessary to create a central intensity ridge in a tubular object. Intensity ridges inherently exist along the central tracks of many tubular objects in medical images. This is true, for example, for small vessels in time-of-flight (ToF) MRA images. However, blurring may be necessary to create intensity ridges along the centers of other tubes. For example, in contrast CT data, the intensity profile of a large vessel may be flat, and in ToF MRA large vessels may appear brighter near their edges than along their middles because of laminar flow. Intensity ridges will persist along the centerlines of tubular objects for a wide range of scales. However, too much blurring smooths the curves of an intensity ridge and therefore degrades the accuracy of an extracted centerline. This trade-off is specifically explored in our Monte Carlo experiments. An excellent study of the effects of scale is presented in Lindeberg [16]. Given its importance, ridge traversal initialization (and ease of automation) must consider x' and σ' together.




Fig. 1. A wide range of scales creates stable centerlines along tubes in 3D. (A) a noisy image containing a tube of uniform intensity, (B) when viewed as a height surface, a centerline doesn't exist, (C) after blurring using a Gaussian with $\sigma = 20$, the tube become more obvious, (D) the height-surface ridge that is the tube's centerline is obvious.

C. Specifying a Tube

Appropriate heuristics for automatically designating tubes-of-interest for extraction are usually application dependent, but, in general, such heuristics are complicated if extraction initiation has many parameters or if its initial parameter values strongly affect the quality or quantity of the centerline model generated. As discussed in Section III-B, the minimum parameters necessary for designating a tube-of-interest are

x' the “seed point” that designates the image-space location of the tube-of-interest

σ' the “seed scale” that designates the approximate size (location in scale-space) of the tube-of-interest

To designate a tube for our method, only very approximate initial values for these minimal parameters are needed. Our current implementation requires ∇I at x' at scale σ' be directed towards the central track of interest. That is, gradient ascent from x should lead to the centerline of interest. Assuming a measurement aperture extends $\pm 3\sigma$, this requirement is generally satisfied in practice if x' occurs within $3\sigma'$ of the tube of interest and σ' is within the broad range of scales that removes erroneous local extremes resulting from image noise and yet does not significantly integrate surrounding objects into the local calculations.

The simplicity of these requirements facilitates automation. Regarding x , given our assumption that the tubes are brighter than the background, the identification of a point x' on a tube can usually be achieved via thresholding based on intensity or local contrast. While global thresholding is not sufficient for segmenting the entire extent of a tube (Table I), finding a point on a tube via thresholding is generally quite feasible for most applications. Gradient ascent from such a point will lead to the tube's centerline. Regarding σ' , assuming the tube's cross-sectional intensity profile is initially flat or ridge-like, as long as the extent of the Gaussian kernel is wider than the tube, a central extreme will exist until a neighboring object is encountered. The smallest appropriate scale may also be dictated by image noise or by the existence of structured intensity variations within the tube. Additionally, appropriate scales for σ' may not be appropriate scales for ridge traversal since using overly large scales during traversal will smooth the centerline as illustrated in the Analysis section. Demonstrations of applications involving automated seed-point selection and the dynamic-scale traversal method are given at the end of the Analysis section. First, however, we must focus on going from a seed-point to an actual point on a ridge.

D. Finding a Ridge

Given a seed point and seed scale, the remainder of the centerline extraction process is automatic. The first step is finding a point on the centerline local to the seed point. We define:

x_0 an initial point on the intensity ridge that is the centerline of the tube-of-interest

σ_0 a scale appropriate for traversing the ridge at x_0

From our initial seed point x' , making measures at scale σ' , we perform a line search to reach the point that is locally maximal in the gradient direction, i.e., the point is a 1D maximum. From that 1D maximum, we expand our search to find a local (N-1)D maximum. We use the N-1 most negative eigenvalued eigenvectors of the Hessian at the 1D maximum as the basis of our search space – these directions are assumed to be normal to the ridge by Equ. 1. The ridge criteria are tested at that local (N-1)D maximum. If they are satisfied, that local maximum becomes our initial ridge point x_0 . If they fail, the N-1 most negative eigenvalued eigenvectors of the Hessian at that point are used to find a new local (N-1)D maximum. If the ridge criteria at that new (N-1)D maximum also fail, our method reports that a local ridge cannot be found. Otherwise, that new local maximum is used as our initial ridge point x_0 . In the default ridge-following implementation (presented next), σ_0 is set equal to σ' . The dynamic-scale ridge-following method (Section III-F) estimates the radius of the object at x_0 and uses that estimate for σ_0 .

E. Following a Ridge

Our ridge traversal method uses an iterative step-maximize procedure. Given an initial ridge point x_0 , the ridge will extend in the positive \vec{t}_0 and negative $-\vec{t}_0$ ridge tangent directions. Each ridge direction is traversed independently.

For the default implementation, at the i^{th} point x_i traversed on a ridge the approximate tangent direction \vec{t}_i is defined as \vec{v}_3 the maximum eigenvalued eigenvector of the Hessian at x_i . The direction of ridge traversal is maintained by multiplying \vec{v}_3 by the sign of the dot-product of \vec{v}_3 and the previous tangent direction \vec{t}_{i-1} .

$$\vec{t}_i = \text{sign}(\vec{v}_3 \bullet \vec{t}_{i-1}) \vec{v}_3 \quad (5)$$

The approximate normal directions at x_i (defined as \vec{v}_1 and \vec{v}_2 in the default implementation) specify an (N-1)-Dimensional plane that the local ridge passes through. Under the assumption of smoothness, if that normal plane is shifted by a small amount β in the tangent direction (see Fig. 2), the ridge should continue to pass through that shifted normal plane – the ridge will exist as a local (N-1)D maximum in that shifted normal plane. That local maximum is determined, and the ridge criteria are tested at that local maximum. If the criteria are met, that local maximum becomes the next ridge point x_{i+1} . Otherwise, the ridge traversal process terminates in that direction and, if not previously done, the ridge extending in the direction $-\vec{t}_0$ from x_0 is traversed.

To assure the continuity of the ridge points extracted, we define two criteria that must be met at each shifted normal plane's local maximum for that maximum to be accepted as a ridge point, x_{i+1} .

One, to terminate traversal when perpendicular branches are encounter, the next tangent direction \vec{t}_{i+1} must point in the “same direction” as the current tangent direction \vec{t}_i . We have chosen 0.7 as the minimum acceptable $\vec{t}_i \bullet \vec{t}_{i+1}$ value indicating the “same direction.” Minor experimentation has demonstrated that this threshold is not critical. This




Fig. 2. An illustration of the step-maximize procedure used to iteratively traverse a centerline. Eigenvectors of the local Hessian matrix approximates the tangent and normal directions. The shifted normal plane bounds the search for the next ridge point.

threshold is used in all experiments and applications in this paper.

Two, to terminate traversal when a large spatial discontinuity is encountered, the (N-1)D maximum found in the shifted normal plane x_{i+1} must be “close” to the current ridge point x_i . We assume that centerline resolution at the inner scale of the data is desired. We therefore do not allow the Euclidean distance between any two continuous ridge points to be greater than one voxel.

The primary parameter of this iterative traversal process is β , the distance the normal plane is shifted per step. Appropriate values for β (and “same direction” and “close”) are dictated by the assumed smoothness of the tube. For our implementations, β is 0.2 voxels. For some applications in which the tubes are generally much larger than the inner scale of the data, it may be beneficial to use an adaptive step size β_i that is proportional to the local scale of the tube σ_i (“same direction” and “close” thresholds should also be adapted).

We use Brentd’s line-search method [23] to find local extremes. We use cubic splines to calculate sub-voxel values and first and second derivatives. Also, the image is not blurred as a pre-processing step. Scaled calculations are made as necessary during the traversal process – this is more efficient assuming that the tubes are relatively sparse in the image, and this is required if the dynamic-scale implementation is used. More details on ridge traversal speed and complete details on the dynamic-scale implementation follow.

F. Dynamic-Scale Enhancements for Improved Traversal Speed, Accuracy, and Automation

Extraction speed is one of the strengths of the default implementation (see Section IV). We have already mentioned that making scaled measures during processing is one source of its speed. Another major requirement for speed is that an image’s voxels should be isotropic. Most of our medical data is acquired with anisotropic voxels (isotropic in x and y and thicker in the z dimension). The only preprocessing that the data must undergo prior to centerline extraction is the use of cubic spline interpolation to super-sample in the z dimension to create data with isotropic voxels. Without isotropic data, our system requires nearly five times as long to run since every derivative, eigenvector, etc. calculation must be normalized to compensate the anisotropy.

Regarding accuracy and automation, two adaptations have been made: heuristics have been added to recover from local discontinuities and tube radius estimation has been integrated into the step-maximize process to adapt measurement

aperture size σ_i to provide optimal-scale traversal. We measure how these enhancements benefit speed, accuracy, and automation using Monte Carlo experiments in Section IV.

F.1 Heuristics for Recovery from Local Discontinuities

Local discontinuities can occur in three situations: when branchpoints are encountered, when a tube’s intensity approaches that of the background, or when an intensity ridge is corrupted by image noise. We use four recovery heuristics to try to continue ridge traversal past these discontinuities: (a) smoothing tangent directions during traversal, (b) reducing step-size β at places of high curvature, (c) detecting vector swapping at singularities, and (d) perturbing the image data to remove singularities.

(a) Tangent direction smoothing as a ridge is traversed helps reduce the effect of ridge degradation due to image noise. Smoothing is accomplished by averaging the local Hessian’s implied tangent direction \vec{v}_3^{\rightarrow} at x_i with the previous tangent $\vec{t}_{i-1}^{\rightarrow}$.

$$\vec{t}_i^{\rightarrow} = \text{sign}(\vec{v}_3^{\rightarrow} \bullet \vec{t}_{i-1}^{\rightarrow}) \frac{\vec{v}_3^{\rightarrow} + \vec{t}_{i-1}^{\rightarrow}}{\|\vec{v}_3^{\rightarrow} + \vec{t}_{i-1}^{\rightarrow}\|} \quad (6)$$

(b) At places of high curvature the stepsize used to shift the normal plane β is reduced by 10%. As straighter sections are once again encountered, the stepsize is increased by 10%. The stepsize is always limited to being between 0.1 and 0.6 voxel. Using intra-cranial MRA, we observed that $\vec{t}_{i-1}^{\rightarrow} \bullet \vec{t}_i^{\rightarrow}$ was consistently less than 0.9 at curves. We chose that as our threshold for high curvature. Although we used intra-cranial MRA data to guide our selection of a threshold, that threshold has remained fixed and was used in all experiments in this paper regardless of the organ or modality.

(c) Discontinuities can exist as singularities at branchpoints as well as at other locations along a ridge. At these points, the image data contains intensities that more closely resemble a sphere than a tube. All three eigenvalues of the local Hessian become negative and the ordering of the eigenvectors with respect to the previous ridge point’s tangent and normal directions may temporarily change. We detect eigenvector swapping and re-orient them by finding which eigenvector best matches the previous tangent direction.

$$k = \text{argmax}_{j \in \{1..3\}} (\vec{v}_j^{\rightarrow} \bullet \vec{t}_{i-1}^{\rightarrow}) \quad (7)$$

We then use that best-matching eigenvector as the tangent direction in a modified Equ 6

$$\vec{t}_i^{\rightarrow} = \text{sign}(\vec{v}_k^{\rightarrow} \bullet \vec{t}_{i-1}^{\rightarrow}) \frac{\vec{v}_k^{\rightarrow} + \vec{t}_{i-1}^{\rightarrow}}{\|\vec{v}_k^{\rightarrow} + \vec{t}_{i-1}^{\rightarrow}\|} \quad (8)$$

(d) Another method for handling singularities and other discontinuities is based on the fact that such structures are non-generic in the image data. Because they are non-generic, local perturbations to the image data will cause those points to be displaced or destroyed. Therefore, if after the application of the above recovery techniques, a valid ridge point still cannot be found, image perturbation is employed, i.e., σ is slightly decreased, β is slightly increased, and then the step-maximize process is retried. If a second discontinuity is encountered within the next two voxels, a second recovery is not attempted.




Fig. 3. A rendition of the kernel used to estimate a tubular object’s approximate width. Kernel is aligned normal to the centerline; inner circles represent positive spheres; outer circles represent negative spheres; spheres bound the radius being tested; sphere radius is proportional to the radius being tested.

These recovery enhancements address accuracy and automation. Specifically, if longer extents of the centerlines of tubular objects can be extracted from each seed, fewer seeds will be needed to generate a more complete coverage of the region of interest. Next, we present optimal-scale selection which also addresses accuracy and automation.

F.2 Optimal-Scale Selection

Because of scale’s influence on ridge traversal, the dynamic-scale implementation seeks optimal scale height ridges; i.e., the scale used during traversal is dynamically adjusted based on the radius of the object. For all dynamic-scale applications and experiments presented in this paper, the scale is set equal to the estimated local radius of the tube. Application-specific knowledge of image noise and centerline tortuosity can be used to specify alternative mappings from object radius to scale, but we have found such tuning unnecessary.

We estimate the local radius of the tube by finding a local maximum of a medialness function at x_i .

$$\sigma_i = \operatorname{argmax}_\rho (M(x_i, \vec{t}_i, \rho)) \quad (9)$$

Our medialness function $M(\cdot)$ uses an adaptive convolution kernel formed by a ring of boundary operators centered about x , aligned with the normal plane at x , and at a distance ρ from x . Each boundary operator is a pair of spherical operators of size 0.25ρ aligned radially about x at a distance of 0.75ρ and 1.25ρ and with the inner sphere having a positive sign and the outer sphere having a negative sign. The convolution kernel is adaptive in that, to increase the stability of the medialness response given image noise and neighboring objects (including branches), the radial direction that produces the weakest convolution response is eliminated from the kernel’s convolution sum.

We perform only a coarse estimation of the object’s local width during ridge traversal. The tube’s local width is estimated only every ten voxels, and the local maximum search uses a tolerance of 0.1 voxels for σ_i . As a result, the speed of traversal is minimally affected. In the next section, the change in speed as well as the improvements in accuracy

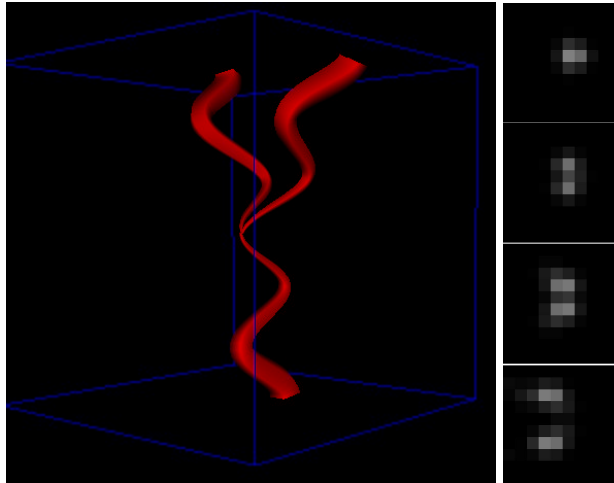


Fig. 4. *Left:* A rendering of the mathematically defined surface of the tortuous, branching object used in the simulated-data Monte Carlo experiments. *Right: Top-to-Bottom:* Slices 49-52 that involve the branch have been blurred by 0.5 voxel, enlarged, and cropped to a 10x10 pixel region.

and automation due to these enhancements are quantified by Monte Carlo experiments.

IV. ANALYSIS

In this section, we quantify and compare the speed (time to extract a centerline), accuracy (difference from ideal centerline), and automation (ability to extract long centerlines and pass branchpoints) of the default and dynamic-scale implementations. We first present a simulated dataset and two experiments that illustrate the influence of scale on ridge traversal accuracy and that quantify the speed of the extraction methods. We then describe and present results from a set of Monte Carlo experiments involving the simulated data to quantify the accuracy and automation of the methods. Using those results, Null Hypotheses are tested to quantify the improvements offered by the dynamic-scale implementation. Finally, a Monte Carlo experiment is conducted using MRA data to detail the consistency of the dynamic-scale method on clinical data.

A. Simulated Data and Random Seed Points

We created a 100x100x100 voxel, isotropic dataset containing a tortuous, branching, tubular object. The object's centerline passes through 350 voxels - an ideal traversal of this centerline results in the extraction of 800 centerline points using $\beta = 0.2$. At its ends, the tube's radius is 4.0 voxels. In the middle, at the branchpoint, the radius is 0.5 voxels. A rendering of the explicitly generated surface of the object is shown in Fig. 4A. The background intensity is 100, and the object's cross-sectional intensities have a parabolic profile ranging from 150 at the object's edge to 200 at the object's middle (see Fig. 5). This profile is typical to contrast MRA and ToF MRA for small vessels [15].

Two different types of noise are present in this data. First, the discrete sampling of the continuous object produces errors. This is illustrated in the slice sections shown in Fig. 4. Second, intensity noise was added to the data. Specifically, three variants of the image were created by adding Gaussian noise having standard deviations η of 20, 40, and 80 (Fig. 6). The $\eta=20$ data is representative of the noise level in MR and CT data. The $\eta=40$ data more closely resembles the




Fig. 5. The cross-sectional intensities of the simulated tubular object have a parabolic profile with intensities ranging from 150 to 200. The background intensity is 100. This profile matches that of small vessels in MRA data.

noise characteristics of ultrasound data. The $\eta=80$ data was chosen to explore the methods' performance on worst-case data.

There are three important details regarding these data. First, the data does not simulate a specific anatomic structure. Our centerline extraction method is general. The methods work with bronchi in CT; vessels in MRA, CT, and ultrasound; nerves in confocal microscopy; and bones in CT. Second, the branchpoint profile is arguably optimally difficult for ridge traversal methods. Each fork in the branch is identical in theory and in the data. This branching configuration creates a singularity. In general, singularities do not exist at branchpoints in human anatomy and are infrequent in medical images of that anatomy. Singularities will only exist if both branches have identical intensities using scaled measurements; this is improbable. Flow tends to favor one branch, thereby making the branches' ToF MRA responses different in terms of width and intensity. The length and branch-point crossing results presented in this paper are worst-case. Third, the $\eta=80$ data represents an *extremely* challenging situation – the object's median intensity is less than one standard deviation from the background's mean intensity. Our experience suggests that the accuracy results we are reporting from the $\eta=80$ image are well beyond any worst-case numbers for ridge traversal for any clinically acceptable MRA, CT, or ultrasound data.

Regarding the random selection of seed points x' for extracting tubes from this data, it is our experience that it is generally easy to find a point on a tube using simple, albeit application-specific, heuristics. The efficacy of such seed-point selection heuristics is demonstrated on clinical data in Section IV-F. Therefore, for the experiments in this paper involving the simulated data, seed points were chosen using a uniform distribution along the extent and within the local radius of the ideal tube. Nevertheless, to demonstrate our method's insensitivity to distance from the tube (as another indication of its ease of automation), in the experiments involving the MRA data, the seed points were chosen using a uniform distribution along but extending twice the radius from the tube's centerline. The influence of ridge traversal's other initial parameter, scale σ' , is illustrated next.



Fig. 6. A slice (x -vs- z at $y=50$) along the center of the tube in the simulated data with three different levels of Gaussian noise added. Noise standard deviations are (*Left*) $\eta=20$, (*Middle*) $\eta=40$, and (*Right*) $\eta=80$. The tube's noise-free, cross-sectional intensity profile is shown in Fig. 5.




Fig. 7. Average x -coordinate for each z -coordinate for 500 centerline extractions at three different scales. Larger scales straighten the centerlines at curves ($z=33$ and $z=67$) and are affected by neighboring objects such as branches ($z = 50$).

B. Scale: The Critical Parameter Revisited

As mentioned in Section III-B, scale is the critical parameter of ridge traversal. One aspect of its importance is well illustrated in Fig. 7. This plot was formed using the default implementation and the $\eta=20$ data by averaging the x -coordinate for each z -coordinate for 500 random extractions performed at each σ' value. We conclude that using a larger scale straightens the curves in a centerline and increases the influence of neighboring objects (at the branchpoint the presence of two tubes causes the $\sigma' = 4$ centerline to be deflected). Next we evaluate the influence of scale on the speed of each method.




Fig. 8. The default processing method’s time to traverse 100 points (~ 20 voxels) of a centerline varies from 0.1 to 0.8 seconds depending on the scale used (x-axis). The dynamic-scale method’s time is consistently ~ 0.33 despite different initial scale values.

C. Speed

The speeds of the default and dynamic-scale implementations were quantified in terms of average number of seconds to traverse 100 ridge points (approximately 20 voxels given $\beta = 0.2$). Results are given in Fig. 8. These averages were calculated by performing 200 extractions for each method at each σ' using the $\eta=20$ data. These extractions were performed on a Dell computer running Microsoft Windows 2000 with a 733 Mhz Pentium III processor and 128 megabytes of memory.

Two conclusions are indicated by Fig. 8. First, as the scale used for fixed-scale traversal increases, the dynamic-scale enhancements actually reduce extraction time by using smaller scale measurements when appropriate. Second, both methods are fast. Most importantly, the dynamic-scale method requires ~ 0.33 seconds to follow a 20 voxel extent of the centerline of a 0.5 to 4.0 voxel radius tube. Given the prevalence of 0.5 to 2 voxel radius vessels (1-4 mm diameter) in intracranial MRA images (1x1x1 mm voxel size), it is quite easy to process an intracranial MRA dataset in 15 minutes. The potential increase in speed offered by the dynamic-scale method becomes even more meaningful when the improvements in accuracy and automation afforded by its enhancements are considered.

D. Monte Carlo Experiments using Simulated Data

To quantify the accuracy and automation capabilities of the default and dynamic-scale methods, Monte Carlo experiments were conducted using the simulated data. For each method, 200 random starting points x' were evaluated using three initial scale values ($\sigma'=1.0, 2.0,$ and 4.0) and three different levels of image noise ($\eta=20, 40,$ and 80). Due to using this broad range of initial scales, centerline extractions were initiated with $\sigma' = 4.0$ at seed points where the tube's radius was actually 0.5 , thereby overestimating the actual radius by a factor of eight. Similarly, centerlines were sought where the tube's radius was 4.0 , but instead σ' was set to 1.0 .

Six measures were made on each extracted centerline during the Monte Carlo runs. The first four of these measures quantify the accuracy of each method. The last two measures quantify how well the methods can be automated. The measures are

1. *Average Error*: Mean distance between a point on the extracted ridge and its closest ideal centerline point.
2. *Maximum Error*: Maximum distance between each point on the extracted ridge and the closest ideal centerline point.
3. *Percent of Points Within 0.5 Voxel*: Percent of points on the extracted ridge within 0.5 voxel of their closest ideal centerline point.
4. *Percent of Points Within 1 Voxel*: Percent of points on the extracted ridge within 1 voxel of their closest ideal centerline point.
5. *Average Percent Length*: Percent of points in the extracted ridge. Extraction using a noise-free image produced 800 points, so this number is reported as an percentage out of 800 . This number is strongly correlated with the percent of branch-points crossed (measure 6).
6. *Percent of Branch-Points Crossed*: Percent of the 200 extractions that spanned the z -dimension from slice 47 to slice 53 – these slices bound the branchpoint.

We calculated the average and standard deviations for the first five of these measures over the 200 runs for each three-way combination: η , method, σ' . The sixth measure is a count over each set of 200 runs. Results are given in Figs. 9-11.

D.1 Discussion of Accuracy

To compare the accuracy of the different methods, we tested the Null Hypothesis that there was no difference between the mean performance of the default and the dynamic-scale implementations. Levels of significance $P_{sig}(t > t_0)$ were calculated. Differences between means were considered significant if $P_{sig}(\cdot) < 0.01$. We also calculated the minimum difference between the means that can be detected with a power of 0.99 and a significance level of 0.01 (i.e., $t_{\beta(1)=0.01}$ and $t_{\alpha=0.01}$) using pooled estimated variance at each level of image noise [24], [6]. Variances were pooled within each level of image noise since performance varied more as a function of image noise than as a function of method or seed-scale. Results are given in Table II.

There are three general conclusions that can be drawn from these experiments. First, both methods are extremely accurate. For the $\eta = 20$ and the $\eta = 40$ data, average error was less than one voxel, maximum error was about two voxels, about 90% of the extracted points were within one voxel of the ideal, and for the $\eta = 20$ data on average over 90% of the ideal centerline was traversed by the dynamic-scale method (it consistently crossed the difficult, inner-scale branch point nearly 90% of the time). Two, for every measure except maximum error the dynamic-scale enhancements

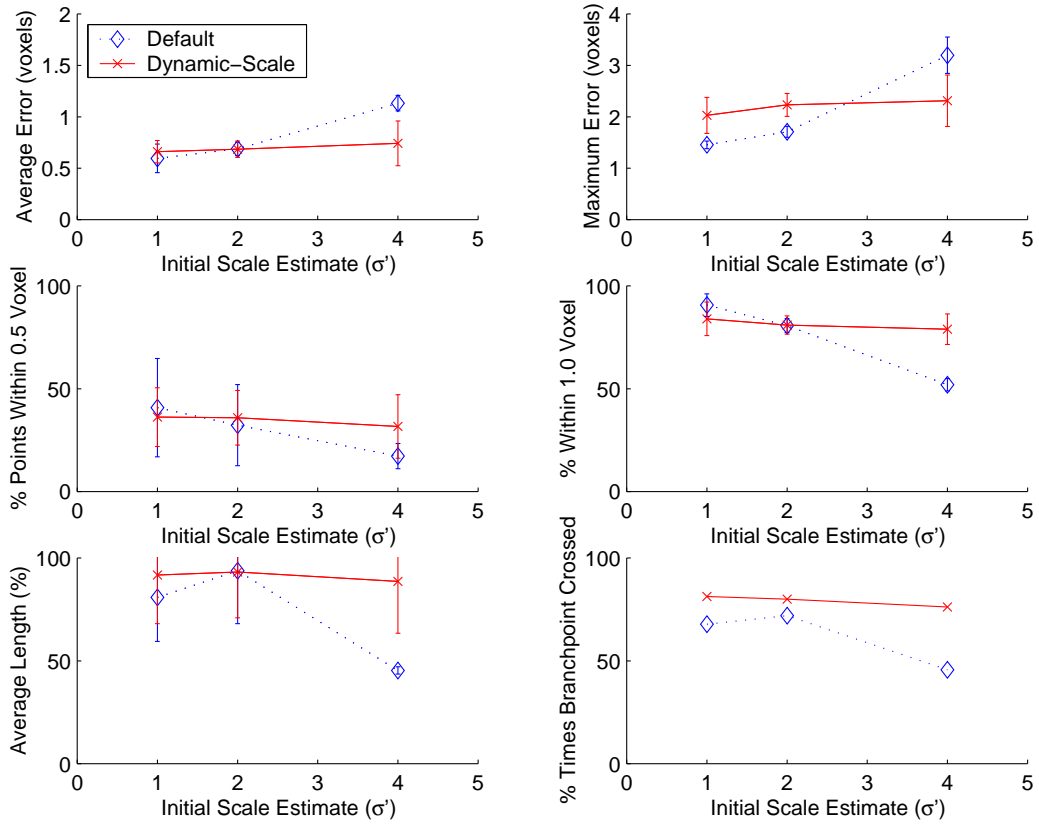


Fig. 9. Means and standard deviations of the six Monte Carlo measures for the $\eta=20$ image (Fig. 6A).

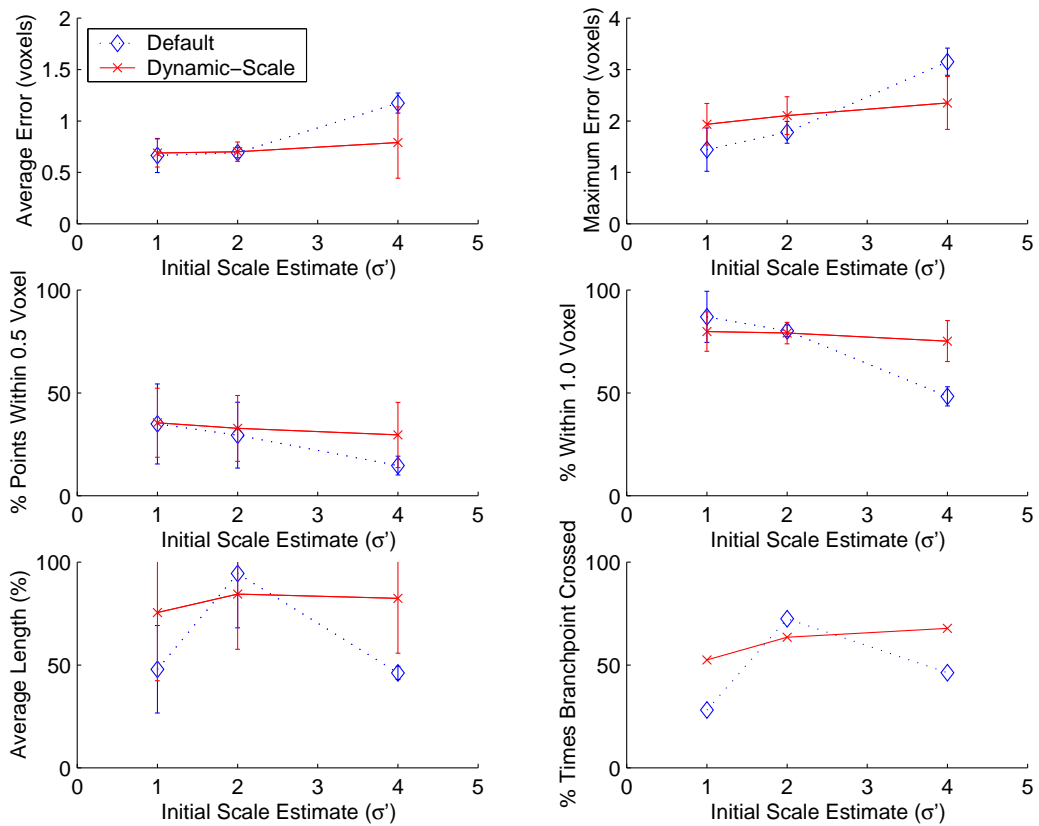


Fig. 10. Means and standard deviations of the six Monte Carlo measures for the $\eta=40$ image (Fig. 6B).

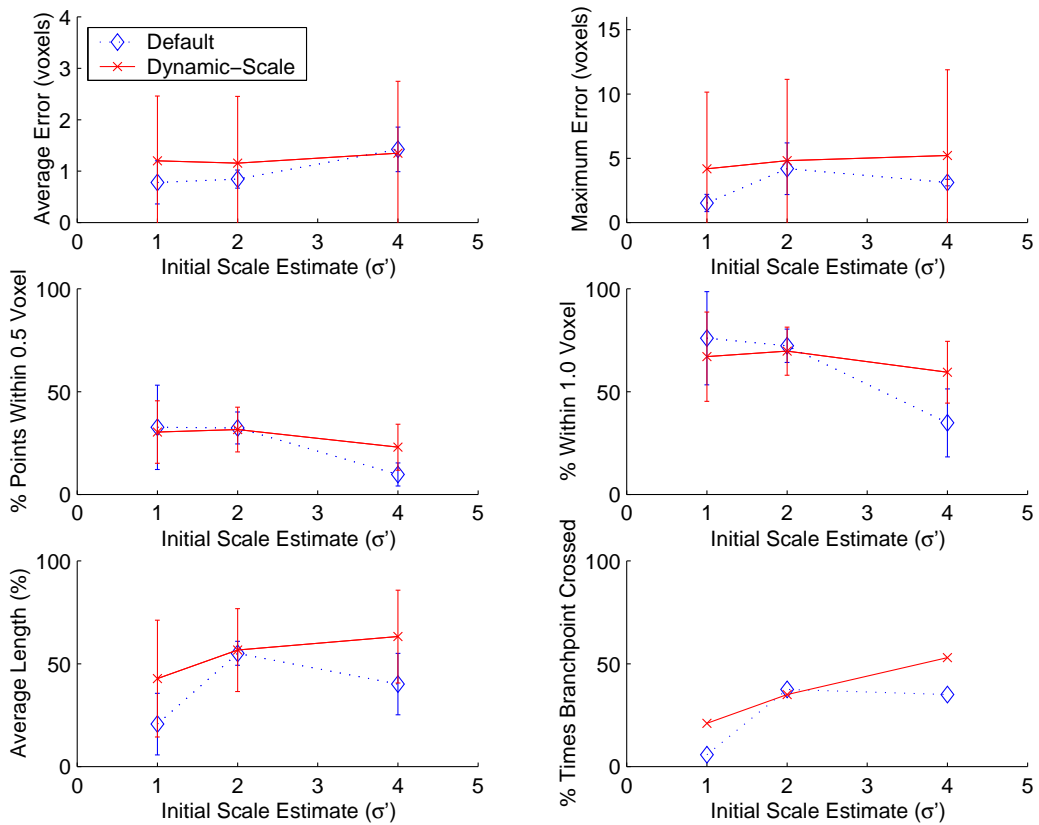


Fig. 11. Means and standard deviations of the six Monte Carlo measures for the $\eta=80$ image (Fig. 6C).

η	Decrease in Average Error	Decrease in Maximum Error	Increase in % within 0.5 voxel	Increase in % within 1.0 voxel	Increase in % Length
20	0.1096	-0.0705 *	4.477 *	6.871	17.873
40	0.1168	-0.0043 *	6.268	6.182 *	17.938 *
80	-0.2184	-1.7948	3.339	4.351 *	15.593 *

TABLE II

Performance change resulting from the enhancements POOLED OVER σ' . POSITIVE NUMBERS INDICATE THAT THE DYNAMIC-SCALE ENHANCEMENTS IMPROVED EXTRACTION ACCURACY. AN * INDICATES THE CHANGE WAS STATISTICALLY SIGNIFICANT WITH $P_{sig}(\cdot) < 0.01$.

improved performance at every noise level. The one exception is at $\eta = 80$; the average error increased but that increase was not statistically significant. Three, the dynamic-scale method did produce higher maximum errors. This is not surprising. It is due to the additional pliability provided by searching for an optimal scale and due to the use of recovery heuristics that increase the likelihood of taking a false step once a traversal mistakenly goes outside of a tube. The increase in maximum error, however, is extremely small - when statistically significant, the increase in maximum error is less than 0.071.

Most surprising about these results is the success of ridge traversal on the extraction of a tube from the extremely noisy $\eta = 80$ data. Average centerline point error was about one voxel with approximately 70% of the extracted




Fig. 12. *Left:* A slice of the $\eta=80$ data is shown. The tube is not apparent. *Middle:* a randomly chosen dynamic-scale Monte Carlo extraction did extract the tube passing thru the slice – its intersection with the slice is overlaid. *Right:* The extracted tube is shown with respect to the ideal tube (shown in wireframe) to illustrate locations and relative magnitude of errors reported. This tube even passed the branch point. This extraction’s max error was 1.7 voxel at the lower curve.

Method	$P_{sig}(\cdot)$ for Average Error	$P_{sig}(\cdot)$ for Maximum Error	$P_{sig}(\cdot)$ for % within 0.5 voxel	$P_{sig}(\cdot)$ for % within 1.0 voxel	$P_{sig}(\cdot)$ for % Length
Default	0.0067 *	1.7e-13 *	0.0706	0.0025 *	1.1e-08 *
Enhanced	0.1016	0.0743	0.1104	0.0397	0.0877

TABLE III

Ease of automation as a lack of statistically significance performance change due to changing seed scale σ' . FOR EACH IMPLEMENTATION THE SIGNIFICANCE OF THE CHANGE IN PERFORMANCE BETWEEN $\sigma' = 2$ & $\sigma' = 1$ AND BETWEEN $\sigma' = 4$ & $\sigma' = 2$ FOR ALL η . AN * INDICATES A STATISTICALLY SIGNIFICANT EFFECT WITH $P_{sig}(\cdot) < 0.01$. EASE OF AUTOMATION IS INDICATED IF MEASURES ARE NOT STATISTICALLY SIGNIFICANTLY DEPENDENT ON σ' .

centerline points within one voxel of the ideal centerline. The branch point was passed less than half of the time, but most extractions covered the full extent of the branch upon which they fell. An example extraction is shown in Fig. 12.

D.2 Discussion of Automation

There are two measures of automation: the percent of times the branchpoint is crossed and the dependence of each method on σ' . The number of branchpoint crossings per method is detailed in the figures. The Monte Carlo experiments also enabled us to test the Null Hypothesis that there is no difference between the measured means of each method for different σ' values (variance pooled across σ' and image noise). Levels of significance and power were calculated. Significance of changing σ' for each measure is reported in Table III.

The results are conclusive: the dynamic-scale enhancements facilitate the automation of centerline extraction via ridge traversal. Nearly 90% of the time for the $\eta = 20$ data and over 50% of the time for the $\eta = 40$ data, the optimally difficult branch point is crossed. More importantly, the default implementation is statistically significantly dependent on the seed scale for nearly every measure of performance, and the dynamic-scale method demonstrates no statistically significant performance dependence on the seed-scale. The strength of this conclusion is obvious from the nearly flat




Fig. 13. One slice of the MRA image containing the vessel-of-interest. The vessel passes branch points, and its width is near the inner scale of the data.

appearance of the dynamic-scale method’s performance in the figures.

Clearly, all of the simulated data experiments support the use of the dynamic-scale implementation. The next section verifies that dynamic-scale method is also applicable to clinical data.

E. Monte Carlo Experiment using MRA Data

As a final evaluation of the dynamic-scale implementation, we performed a Monte Carlo analysis using clinical MRA data. This experiment illustrates the consistency of the extractions (truth is not known). Specifically, a single vessel (see Fig. 13) was extracted using a manually chosen x' with $\sigma'=1.0$. Using that representation, 100 random starting points were chosen along a 20 voxel extent of that tube and within twice its radius. At each starting point, the extraction was repeated with three different σ' values, i.e., $\sigma'=1, 2, \text{ and } 3$. The x,y coordinates of those 300 extractions are shown in Fig. 14.

The dynamic-scale implementation performed extremely well on this test case. The maximum distance between any two closest points on any two of the 300 extractions was 0.1 voxel. At times, branches were extracted instead of the original centerline. That is not a failing of the extraction process, the starting points warranted the extraction of those branches. Those branch-extractions were not factored into our error estimates. They do, however, demonstrate the enhanced method’s ability to cross branchpoints in real-world data.

F. Applications

In this section we present the application of dynamic-scale centerline modeling to five different sets of data. We demonstrate the modeling of lung vessel, ribs, and bronchi from CT data; the modeling and registration of intracranial vascular from two MRA data sets; and modeling vessels of the liver in CT and ultrasound data. The dynamic-scale method and its parameters, as presented in this paper, were not modified to process any of these data. Tube widths were estimated as the scale that produced the local maximum of medialness (Equ. 9) at every point along each extracted




Fig. 14. A plot of the x,y coordinates of 300 dynamic-scale extractions of a vessel in an MRA dataset – 100 random starting points and at three different σ' values - 1, 2, and 3.

centerline. The processing of the CT and MRA data involved automated seed-point selection that exploits our semi-automated organ delineation technique. That method is described next.

F.1 Seed Point Automation via Organ Delineation

To limit the automated extraction of tubes to particular organs, the CT and MRA applications used a semi-automated organ segmentation method based on connected components. From the segmented organs, the seed points x' were then automatically determined. From those points, centerlines and width models were then automatically generated.

Our organ delineation process is as follows:

1. From a user-specified starting point and using lower and upper intensity thresholds specified by the user, the voxels spatially connected to the starting point and having intensities within the thresholds are identified (this is standard connected components).
2. That component is pruned via erosion using a spherical operator to remove sections that are only connected by thin strands to the main component. The main component is subsequently dilated by the same amount to return it to its original borders (minus the clipped regions).
3. The main component is then dilated using a spherical operator to fill-in small holes. The main component is again subsequently eroded by the same amount to return it to its original borders (minus its holes and clipped regions).

We have developed a user-interface that simplifies the specification of the parameters of this process. On a slice of data, the user merely specifies the corners of a box contained by the organ to be delineated. The center of the box is used as the connected-component starting point. The fifth and the 95th percentile of the intensities in that box are

used as the connected component thresholds. The pruning sphere size is application specific (~ 1.2 cm for liver and brain envelope segmentation), and we define the fill sphere size to be the same as the pruning sphere size. For some applications and data, it may be necessary to edit the resulting delineation; however experience indicates that the time required and the inter and intra-user variability for such editing is usually minimal compared to the time and variability of hand-segmentation.

Once an organ is delineated, the seed-point threshold must be specified. Continuing our assumption that the vessels are brighter than the background, the user specifies the threshold in terms of intensity percentile. We have found that using the 99.5 percentile is an effective for threshold for selecting seed-points in intracranial MRA. For the CT data, the threshold was chosen manually. To eliminate the possibility of extracting tubes that abut the delineated organ, seed-points within 6 voxels of the object's boundary are disqualified. More experimentation is needed and more intelligent histogram analysis may be beneficial for the broader application of this approach. Results produced thus far, however, are quite promising. They are presented next.

F.2 Application Illustrations

The application of the dynamic-scale method to five different sets of data is shown in Fig. 15. All extractions used $\sigma' = 1.5$ voxel. Application specifics are as follows.

1. Using non-contrast lung CT data, we modeled a right lung's bronchi, vessels, and ribs as tubular objects. The bronchi, ribs, and three major vessels leading into the lung were extracted manually using one seed-point per object. The vessels within the lung were extracted by segmenting the right lung and then running the extraction process from every point in the lung above a manually chosen threshold.
2. To illustrate the usefulness of centerlines for 3D/3D registration, the vessels from a patient's post-operative MRA data were semi-automatically extracted and rigidly registered with the patient's pre-operative MRA data (the tube-to-image registration method is presented in [2], [3]). The extractions were limited to vessels within the brain envelope in each set of data. In the figure, the vessels from both MRAs are shown. Both sets of vascular models are nearly identical except where non-rigid deformations are present in the data.
3. Contrast CT data captures the portal and hepatic venous networks of the liver. In the figure, the delineated liver and the semi-automatically extracted vessels are shown. Interacting with such visualizations help surgeons evaluate donors for adult-to-adult, living-donor, partial-liver transplants.
4. Vessels appear as tubular voids in 3D ultrasound data. We inverted the intensities in 3D ultrasound data of a liver and then manually selected seed points to extract the vascular representations shown.

Truly the dynamic-scale centerline modeling method is broadly applicable. These clinical data represent a wide range of image noise, tube cross-section intensity profiles, and centerline conspicuities, yet the same implementation and parameters presented in this paper were used to perform all of the extractions shown. Most impressive is the consistency of the extractions generated from the two different sets of MRA data (pre and post operative) from the same patient. The centerline-based registration method demonstrated is one of numerous applications enable by fast, accurate, and automated centerline extraction. For other clinical applications, see the extensive work lead by Dr. Bullitt [7], [8].




Fig. 15. *Top-Left*: These tubular models of a lung's ribs, bronchi, and vessels were extracted from non-contrast CT. A connected component model of the lung defined the region of interest from which lung-vessel seed-points were automatically selected by thresholding (Section IV-F.1). *Top-Right*: These vessels representations were automatically generated from seed-points within semi-automatically defined brain envelopes in pre and post-operative MRA data. The MRA data were aligned by matching the centerlines of the tubes in one image with the data in the other image. Non-rigid deformations near the surgical site are seen as vessel mis-alignments as indicated by arrows. *Bottom-Left*: Using arterial-phase contrast CT, the liver surface has been delineated (shown as dots) and using that delineation, portal and hepatic vessels have been automatically extracted. *Bottom-Right*: 3D ultrasound data of a liver was inverted to make the vessels appear brighter than the background and then seed-points were manually specified to generate this visualization.

V. CONCLUSION

In this paper we define a general method for centerline extraction based on our 1996 work [1] (the “default” implementation) and then offer a set of multi-scale heuristics and optimal-scale measures for handling noise, discontinuities, and singularities in the data (the “dynamic-scale” implementation).

We evaluated these implementations using Monte Carlo experiments with simulated and clinical data. Those experiments evaluated the speed, accuracy, and automation of the methods. The data contained extremely challenging levels of noise and an uncharacteristically difficult branchpoint (local object radius was 0.5 and both branches were identical). All methods were shown to be fast (0.33 second to extract a centerline passing through 20 voxels) and accurate (average error less than 0.5 voxel). The dynamic-scale implementation offered statistically significant improvements in accuracy. The dynamic-scale implementation was shown to be truly supportive automation in that its initial parameter values were not statistically significant in determining its final accuracy (i.e., its average error, maximum error, percent of points within 0.5 voxel of ideal, percent of points within 1 voxel of ideal, and its length).

We then presented demonstrations of the clinical utility of the dynamic-scale centerline modeling method using lung CT, head MRA, liver CT, and liver 3D ultrasound data. Using connected components and morphological operators to segment organs of interest, these applications demonstrated the automated (and manual point-and-click) extraction of vessels, bronchi, and ribs as tubular objects from those organs. We also demonstrated a tube-to-image 3D/3D registration method that exploited these centerline representations.

Future work will focus on the quantification of the performance of various vessel width estimation methods. The methods described in this paper have already been used to support multiple clinical applications, evaluations, and descriptions [7], [8], [9]. Visualizations associated with these and other applications as well as the simulated data used in the Monte Carlo experiments are available on the internet at <http://caddlab.rad.unc.edu/>.

REFERENCES

- [1] S. Aylward, E. Bullitt, S. M. Pizer, and D. Eberly. Intensity ridge and widths for tubular object segmentation and registration. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 131–138, 1996.
- [2] Stephen Aylward, Julien Jomier, Susan Weeks, and Elizabeth Bullitt. Registration and analysis of vascular images. *International Journal of Computer Vision*, In Review:12, 2001.
- [3] Stephen Aylward, Susan Weeks, and Elizabeth Bullitt. Analysis of the parameter space of a metric for registering 3d vascular images. In *MICCAI*, page 8, Utrecht, The Netherlands, October 2001.
- [4] Ingmar Bitter, Arie Kaufman, and Mie Sato. Penalized-distance volumetric skeleton algorithm. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):195–206, July-September 2001.
- [5] Henry Blum and Roger N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10:167–180, 1978.
- [6] G. E. P. Box, W. G. Hunter, and J. S. Hunter. *Statistics for Experiments*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, New York, 1978. Significance: Difference between means.
- [7] E. Bullitt, S. Aylward, E. Bernard, and G. Gerig. Technical report: Computer-assisted visualization of arteriovenous malformations on the home PC. *Neurosurgery*, page Accepted, 2000.
- [8] E. Bullitt, S. Aylward, K. Smith, S. Mukherji, M. Jiroutek, and K. Muller. Symbolic description of intracerebral vessels segmented from MRA and evaluation by comparison with x-ray angiograms. *Medical Image Analysis*, page Accepted, 2000.
- [9] E. Bullitt, A. Liu, S. Aylward, C. Coffey, J. Stone, S. Mukherji, and S. Pizer. Registration of 3d cerebral vessels with 2d digital angiograms: Clinical evaluation. *Academic Radiology*, 6:539–546, 1999.

- [10] Y. Du, D. Parker, and W. Davis. Improved vessel visualization in MR angiography by nonlinear anisotropic filtering. *Journal of Magnetic Resonance Imaging*, 5:353–359, 1995.
- [11] David Eberly. *Ridges in Image and Data Analysis*, volume 7 of *Computational Imaging and Vision*. Kluwer Academic Publishers, Dordrecht, 1996.
- [12] A. F. Frangi, W. J. Niessen, R. M. Hoogeveen, T. Van Walsum, and M. A. Viergever. Model-based quantitation of 3-d magnetic resonance angiographic images. *IEEE Transactions on Medical Imaging*, 18(10):946–956, October 1999.
- [13] Daniel S. Fritsch, David Eberly, Stephen M. Pizer, and Mathew J. McAuliffe. Stimulated cores and their applications in medical imaging. In Y. Bizais, C. Barillot, and R. DiPaola, editors, *IPMI 1995: Information Processing in Medical Imaging*, pages 385–368. Kluwer Series in Computational Imaging and Vision, 1995.
- [14] G. Gerig, T. Koller, G. Szekely, C. Brechbuhler, and O. Kubler. Symbolic description of 3-d structures applied to cerebral vessel tree obtained from MR angiography volume data. In *Information Processing in Medical Imaging '93: Lecture Notes in Computer Science 687*, pages 94–111, 1993.
- [15] R. M. Hoogeveen, C. J. G. Bakker, and M. A. Viergever. Limits to the accuracy of vessel diameter measurement in MR angiography. *Journal of Magnetic Resonance Imaging*, 8, 1998.
- [16] Tony Lindeberg. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Dordrecht, Netherlands, 1994.
- [17] L. M. Lorigo, O. Faugeras, W. E. L. Grimson, R. Keriven, R. Kikinis, and C. F. Westin. Co-dimension 2 geodesic active contours for MRA segmentation. In *Information Processing in Medical Imaging '99: Lecture Notes in Computer Science 1613*, pages 126–139, 1999.
- [18] T. McInerney and D. Terzopoulos. Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Transactions on Medical Imaging*, 18(10):840–850, October 1999.
- [19] B. S. Morse, S. M. Pizer, and A. Liu. Multiscale medial analysis of medical images. In H. H. Barrett and A. F. Gmitro, editors, *Information Processing in Medical Imaging*, 1993.
- [20] M.M. Orkisz, C. Bresson, I.E. Magnin, O. Champin, and P.C. Douek. Improved vessel visualization in MR angiography by nonlinear anisotropic filtering. *Magnetic Resonance Imaging*, 37, 1997.
- [21] S. M. Pizer, D. Eberly, B. S. Morse, and D. S. Fritsch. Zoom-invariant vision of figural shape: The mathematics of cores. In *Computer Vision and Image Understanding*, 1996.
- [22] S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson, and E. L. Chaney. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Transactions on Medical Imaging*, 18(10):851–865, October 1999.
- [23] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, Cambridge, 1990.
- [24] Llya M. Sobol'. *A Primer for the Monte Carlo Method*. CRC Press, Boca Raton, 1994.
- [25] D. L. Wilson and J. A. Noble. An adaptive segmentation algorithm for time-of-flight MRA data. *IEEE Transactions on Medical Imaging*, 18(10):938–945, October 1999.
- [26] D. L. Wilson, D. D. Royston, J. A. Noble, and J. V. Byrne. Determining x-ray projections for coil treatments of intracranial aneurysms. *IEEE Transactions on Medical Imaging*, 18(10):973–980, October 1999.
- [27] P.J. Yim, P.L. Choyke, and R.M. Summers. Gray-scale skeletonization of small vessels in magnetic resonance angiography. *IEEE Transactions on Medical Imaging*, 19(6):568–576, June 2000. "Ordered region growing" - spanning trees. Requires two endpoints or parameters to limit branching rate.