

**Tubular Objects in 3D Medical Images:
Automated Extraction and Sample Applications**

Stephen R. Aylward¹, Elizabeth Bullitt², Charles Chung¹, Stephen Pizer¹

¹Department of Radiology

²Department of Surgery

Medical Image Display and Analysis Group

University of North Carolina, Chapel Hill

Supported in part by P01CA47982 NIH-NCI, R01CA67812 NIH-NCI, and an Intel equipment award

Abstract

In this paper we present and evaluate a novel technique for generating representations of tubular objects in 3D medical data. Tubular objects are abundant in medical images, e.g., vessels, bones, ducts, spinal cords, and bowels. Tubes can be characterized as smoothly varying, yet possibly branching, structures in 3D that have nearly circular cross sections.

While other techniques have been suggested for segmenting tubular objects, our method rapidly generates accurate and consistent tubular representations with minimal user interaction by exploiting the geometry of tubes. Specifically, tubular objects defined via contrast are special in that blurring produces a central intensity ridge that well approximates the objects' central skeleton. Our method operates by traversing those central skeletons. Once extracted, those central skeletons also serve to stabilize a width estimation process.

Our method is also advantageous because the form of the representations it generates (i.e., central skeleton and widths). These representations enable abstract reasoning and hence a variety of clinical applications. These representations can be easily joined, split, quantitatively evaluated, and viewed in isolation. As examples, we illustrate the formation and manipulation of vascular and bronchial trees for surgical planning, and we highlight our work regarding the use of our representations for multi-modal, multi-dimensional registration.

Keywords

Vessels, Ridges, Segmentation, Registration

1. Introduction

Segmentation remains a critical issue for medical imaging. Generic, clinically effective (i.e., easy-to-use, broadly applicable, consistent, sensitive, and specific) algorithms have eluded researchers. When segmenting an object in a 3D image from a particular modality, expert anatomical and image segmentation processing knowledge must be applied. Expert anatomical knowledge is required to handle uncertainty (e.g., the limitations of segmentation due to image noise can be handled if an object model is available). Expert segmentation knowledge is required to understand and adjust the parameters of most segmentation techniques (e.g., to determine scale, boundary curvature constraints, or energy function weightings) and to weight the pros and cons of the large number of incommensurate segmentation techniques that are available (e.g., to decide between landmark, connected component, or edge based techniques). Both forms of knowledge can be incorporated into a segmentation tool, but that incorporation has proven to be quite difficult and very task dependent. Simpler (i.e., model-free) segmentation tools can be more broadly applied, but then the user must provide the anatomical expertise. In most clinical environments, where expert anatomical knowledge is prevalent, semi-automated, generic, hand-contouring techniques with very few parameters are frequently chosen to segment images. Such contouring techniques are effective for some tasks, but more often, they are too laborious and have too high of inter and intra user variance. Additionally, the representations formed via contouring are too low-level to facilitate abstract reasoning (e.g., task specific visualization, quantification, or simulation). The limitations of model-based and model-free segmentation systems are particularly evident when segmenting tubular objects in medical images.

Tubular objects are prevalent in medical images. Tubular objects include vessels, bowels, ducts, nerves, and select bones. Such objects are visible in magnetic resonance images (MRI) and x-ray computed tomography images (CT), MR and CT angiograms (MRA and CTA), and confocal microscopy images. The types of tubes that we are interested in, including those listed, have three basic geometric properties:

- A tube has a central skeleton that forms a smooth 1D curve in 3D.

- Orthogonal to that central skeleton, a tube has a nearly circular, smoothly varying cross-section.
- Tubes can branch.

We assume that the tubes are delineated from their background by contrast (versus texture) differences.

Our goal was to develop an effective tubular object segmentation method. Efficiency must be judged with respect to the segmentation process and the representations it produces. The segmentation must address the issues of ease-of-use, range-of-use, consistency, sensitivity and specificity. The representations formed must be appropriately abstract to facilitate reasoning not just visualization. For example, dependent on the task, significant information lies in the shape of each tube, how the image data deviates from tubular, the arrangement (e.g., tree structure) of a set of tubes and the arrangement of tubes across multiple dimensions and modalities.

In general, model-based and model-free segmentation tools are ineffective for tubular object extraction from medical images. The global structure of most tubular objects (e.g., the arrangement of the bowel or the tree of kidney vasculature) varies significantly from patient to patient and does not adhere to a global model; thereby rendering model-based methods ineffective. Model-free segmentation methods such as hand contouring, thresholding, or connective components are generally ineffective due to the amount of user interaction required and their sensitivity to spatial variations in the images and to variations in the appearance of the objects (e.g., the cross sectional intensity variations in vessels captured via time-of-flight MRA). Furthermore, the inability of model-free techniques to support abstract reasoning is particularly detrimental for tubular structures where automated quantification (e.g., stenosis or blockage detection) is critical or when a physical system is being simulated (e.g., intracranial embolization simulation or multimodal/multidimensional registration).

The segmentation method presented in this paper is effective in terms both of its segmentation process and the representations it generates. Our segmentation technique achieves its efficiency by integrating knowledge about the geometry of tubes into the segmentation process and into the form of the representations it generates. We therefore refer to this class of segmentation techniques as geometry-based methods. Specifically, our method

exploits the fact that tubular objects, defined via contrast, are special in that blurring produces a central intensity ridge that well approximates the objects' central skeleton. Our method operates by traversing those central skeletons. Once extracted, those central skeletons also serve to stabilize a width estimation process which is simplified by exploiting the cross-sectional characteristics (e.g., nearly circular and smoothly varying) of tubes. In this manner, the segmentation process is invariant to rotation, translation, and absolute image intensities and is insensitive to image noise (range-of-use). It requires minimal user-interaction and runs on standard PC hardware (ease-of-use); it operates with little dependence on user expertise (consistent); and it has been statistically judged to provide accurate (sensitive and specific) representations. The representations generated by our system take the form of central skeletons and their widths. This form enables joining, splitting, and shape quantification as well as the selective visualization of tubular objects. For example, via joining and splitting we can depict vessel trees and branches and simulate the flow of a contrast bolus [Bullitt 1997]. Additionally, the quantification and selective visualization of tubular segments aids in the identification and characterization of stenosis and bowel obstructions. Furthermore, the central skeletons have been proven to provide an excellent basis for the registration of multi-modal / multi-dimensional images (e.g., MRA / DSA registration).

In this paper we present our extraction algorithm (Section 3) and an evaluation of its ease-of-use, range-of-use, and consistency (Section 4). Measures of sensitivity and specificity are applications specific; statistical evaluations of our methods performance for neurosurgical planning and interventional radiology have been previously published [Bullitt 1999]. We will demonstrate the effectiveness of the form of our representations (Section 5). Specific neurosurgical applications of the representation's form are detailed in [Bullitt 1997; Bullitt 1999].

2. Background

A significant amount of research [Bullitt 1999; Dachman 1998; Frangi 1998; Masutani 1998; Park 1998; Wilson 1997; Gerig 1993] has focused on representing tubular objects for diagnosis (e.g., stenosis quantification and

virtual fly-through), surgical planning (e.g., aneurysm clipping), and multi-modal and multi-dimensional registration (e.g., 3D MRA fusion with 2D x-ray angiograms). Methods previously presented include thresholding, region growing, morphological, adaptive thresholding, kernel, image filtering, and fuzzy-logic techniques.

In general, threshold techniques are not feasible for MRAs due to large-scale intensity inhomogeneities and are not feasible for MRAs and CTAs due to flow artifacts (e.g., contrast dissipation and laminar flow effects that produce irregular cross-sectional intensities; see the study by Hoogeveen [Hoogeveen 1998]). Region growing methods appear quite promising. They are robust in that they do not make assumptions about the cross-sectional shape or continuity of tubes, but those assumptions can help when dealing with uncertainty (i.e., image noise); those methods appear to have difficulties with small vessels, especially for "one-voxel" vessels whose diameter is approximately the resolution of the imaging device. With the resolution of MRA on the order of 1 mm and with the increased use of 1 mm and smaller intravenous devices, the accurate representation of such vessels is critical. Kernel-based techniques tend to be over-constrained; for example, most implementations are unable to handle a wide range of vessel sizes. Image filtering and fuzzy-logic methods can provide quite excellent visualizations, but these techniques were never devised for producing representations for reasoning about the tubes contained within an image. Many of the above listed techniques also require extensive user-interaction during initialization, are dependent on the setting of several parameters based on the modality or object being processed, and/or require significant computational power.

The segmentation technique presented in this paper is closely related to the methods of core extraction developed within our group [Pizer 1996]. The method presented here, however, differs in both theory and application. Essential to core extraction is multi-scale medialness maximization. Multi-scale medialness maximization *simultaneously* solves for the middle (i.e., central skeleton) and the widths of objects in medical images; our method solves for these structures *sequentially*, at a single scale, and only for tubes. Cores have a theoretic advantage since multi-scale theory and simultaneous maximization ensures important invariance and insensitivity

properties, yet core extraction can be quite computationally expensive. In practice, for tubular objects, the method presented in this paper has demonstrated similar insensitivity and yet requires significantly less computation power.

3. Methods

The strength of our technique comes from its incorporation of tubular object geometry into the extraction process. Additionally, the extraction process utilizes this geometry in a manner that is invariant to rotation, translation, scale, and contrast. Specifically, our method operates in three steps:

- (1) **Tube-of-Interest Designation:** To designate the tubular object of interest, the user must select a point within that tube and estimate the tube's width at that point (Section 3.1). We have found our method to be quite insensitive to these specifications (Section 4). Alternatively, we have devised two automated, task specific methods for extracting the entire set of tubes or a specific tree-of-tubes from 3D medical images (Section 3.4).
- (2) **Central-Skeleton Extraction:** Our method automatically extracts the full extent of the central skeleton of the designated tube of interest. Tubular geometry and the delineation of tubes via contrast ensure that simply blurring an image creates a central intensity ridge (or valley) along any tube. Our method employs an invariant height-ridge (or valley) traversal technique to extract that central intensity ridge. (Section 3.2)
- (3) **Width Estimation:** The width of the tube of interest is estimated at points along its central skeleton. Our method takes advantage of the fact that vessels have nearly circular cross sections. The width of a tube about a central skeleton point is proportional to the scale that produces a maximal response from a cylindrical medialness measure. (Section 3.3)

3.1. Tube-of-Interest Designation

To designate a tube of interest so that it can be extracted from a medical image, it is necessary and sufficient to designate the corresponding central skeleton and an approximate width at the designation point. In this manner, our process is easy-to-use in that the minimal information / user-interaction is required to specify a tube-of-interest. Even this minimal requirement is reduced when the automated extraction methods presented in Section 3.4 are considered.

3.1.1. Designating the Central Skeleton

To designate a central skeleton, that central skeleton must *exist* as a 1D intensity extreme (ridge or valley) and the same extreme must be consistently *selected* despite a broad range of user-interaction skill (i.e., clinical effectiveness: ease-of-use). The concept of a 1D extreme is best illustrated by viewing an ND image as an ND surface in (N+1)D. Consider the 2D image of a slice from an MRA (Figure 1). That slice can be viewed as a 2D surface in 3D when intensity is mapped to height. Intensity ridges and valleys are evident (Figure 2). The same concept applies to 3D images that form 3D surfaces in 4D; 1D extrema similarly exist.




Figure 1. An MRA slice as a 2D image




Figure 2. An image slice can be viewed as a 2D surface in 3D with corresponding 1D extrema (ridges and valleys)

Concerning the *existence* of a central skeleton, as is evident in Figure 2, an intensity extreme inherently exists along the center of certain tubular objects in medical images. Such inherent extrema result, for example, from the flow properties of blood through vessels or as a result of the topographic reconstruction process. In those

situations in which such an extreme is not inherent, blurring a 3D image using a Gaussian kernel will produce an intensity ridge (or valley) along the central skeleton of a tubular object defined via contrast. The creation of such extrema via blurring is illustrated in Figure 3 and 4. The stability of the extrema given such blurring is discussed in Section 3.2.




Figure 3. (A) A 2D image containing a uniform intensity stripe; Gaussian noise has been added; (B) A rendering of the height-surface (mapping intensity to height) of the image in 1A; no central skeleton / ridge is present




Figure 4. (A) Blurring the image in Figure 3A produces a central skeleton along the stripe; (B) The intensity ridge is best visualized via a rendering of 1A's height surface. Our method traverses such 1D ridges in 3D (not in 2D as shown here) to extract the central skeleton of tubular objects.

Concerning the *selection* of a particular extreme (i.e., central skeleton), we use a height-ridge-specific search process to reach the extreme nearest a user-specified mouse-click in a 2D slice of a 3D volume. To explain this process, a definition for a 1D-in-3D height ridge must be given:

Define: σ as the scale at which all ridge measures are taken

I as the intensity at $\mathbf{x} \in \mathfrak{R}^3$ (i.e., at scale σ)

\mathbf{H} as the Hessian matrix of I at \mathbf{x}

\mathbf{v}_i and a_i as the eigen-vectors and associated eigen-values of \mathbf{H}

where $a_1 < a_2 < a_3$

If \mathbf{x} is on a 1D in 3D ridge, then

\mathbf{v}_3 approximates the ridge's tangent direction at \mathbf{x}

\mathbf{v}_1 and \mathbf{v}_2 approximate the ridge's normal directions

For \mathbf{x} to be classified as being on a 1D in 3D ridge, it must be true that:

$a_2 / (a_1^2 + a_2^2 + a_3^2)^{1/2} \sim 0.5$ "Cross-sectional intensity is nearly circular"

$\mathbf{v}_1 \cdot \nabla I \sim 0$ "The point is an extreme in the directions normal

$\mathbf{v}_2 \cdot \nabla I \sim 0$ to the ridge"

Using the eigen-vectors of the Hessian to define the principle directions of an extreme makes this height ridge definition a "maximum-curvature" height-ridge definition. Note that the explicit tangent and normal directions can be calculated, but they require third-order information. Such high-order information is costly to generate, e.g., requiring fourth-order spline approximations to have C^1 continuity of tangent vectors. As a result, in practice, the approximations provided by the eigen-vectors of the Hessian are used. [Eberly, 1997]

The one parameter of the designation process is σ , i.e., the scale at which all measures are made. This parameter, however, is well defined by the width of the object being extracted which the user must already specify.

Therefore, no additional burden is placed on the user. The insensitivity of the extraction process to this parameter is discussed in Section 4.

Given this height-ridge definition and an initial user-supplied starting point \mathbf{y}_0 , a sequence of line-searches in the gradient and maximum curvature directions is used to reach a local extreme point and thereby specify the central skeleton of interest. Specifically,

- 1) A line-search in the gradient direction at \mathbf{y}_0 leads to \mathbf{y}_1 , the local maximum in intensity.
 - This location will be on a 2D height ridge in 3D.
 - This location may be close to but not necessarily on a 1D-height ridge in 3D.
- 2) The eigen-vectors of the Hessian at \mathbf{y}_1 are calculated to approximate the ridge's normal directions, and a gradient ascent in the subspace defined by those normal directions leads to the point \mathbf{y}_2 .
 - This is our first candidate ridge point.
 - If it meets all ridge criteria, it designates the central skeleton of interest. Otherwise...
- 3) The Hessian at \mathbf{y}_2 is calculated to better approximate the ridge's normal directions, and a gradient ascent in that normal subspace is performed to produce \mathbf{y}_3 .
 - If \mathbf{y}_3 is not a valid ridge point, the user must supply another stimulation point.

3.1.2. Designating the Tube's Width

To complete the designation of the tube of interest, the approximate width of the tube at the initial extreme point must be specified. For example, consider that if such a width is not specified, it cannot be known if the user's indicated starting point within an MRA is meant to initiate the extraction of a vessel, the brain hemisphere, the brain envelope, or the entire head. As stated previously, this initial approximate width is also used to specify the scale used to extract the central skeleton. We have found our system to be quite insensitive to this initial width estimate. For example, every vessel in an MRA of a brain can be extracted using the width estimates ranging from ~0.5mm to ~2mm. This robustness is explored in Section 4.

3.2. Central Skeleton Estimation via Intensity Ridge Extraction

Having designated the tube of interest, the remainder of the extraction process is completely automated. Once a point \mathbf{x} on the skeleton has been identified (Section 3.2), two-steps are repeated to traverse the extent of the skeleton.

- (I) The (approximately) normal plane is shifted along the tangent direction. Assuming the 1D curve is smoothly varying in 3D (i.e., that it is the central skeleton of a tube), that 1D curve will pass through this shifted normal plane.
- (II) The local maximum in intensity in that shifted normal plan is located. This point corresponds to the next point on the central skeleton.

All extraction steps are carried out at a sub-voxel level. The normal plane is shifted by 0.2 voxels per step. We fit an approximating cubic spline to the image data to generate sub-voxel values and first and second derivative information.

Traditionally, this "shift-maximize" process is repeated until a termination criterion is met. Termination criterion include

- Encountering another tube (entering the voxel of a previously extracted central skeleton).
- Experiencing a rapid change in spatial location.
- Experiencing a rapid change in tangent direction.
- Failing to find a valid ridge point.

The latter three termination criterion generally occur when a neighboring object interrupts the search for a local maximum in the shifted normal plane. Often, in these cases, the correct central skeleton can be found by re-flowing in the same normal plane using a slightly smaller image measurement scale σ . We have automated this recovery technique identifying those termination points and then trying step beyond them using a scale $\sigma' = 0.9\sigma$.

We have also greatly improved the effectiveness of the traversal process by extending the notion of a ridge to include connector curves. These curves exist as a continuum of points, between two valid ridges, that deviate from the traditional height ridge definition. We consider one particular type of connector curve, the “M-connector.”

A point is on an M-Connector if the point meets every ridge criterion, but

- $a_1 < a_2 < a_3 \sim < 0$ "The point is maximal in every direction"
- The tangent direction is best approximated by \mathbf{v}_i where $i \neq 3$
"The eigen-vectors have swapped."

We detect the connector curves by comparing the eigen-vectors of the Hessian at a new point \mathbf{x} with the tangent direction used to reach that new point. Specifically, if the eigen-direction with the largest projection onto the previous tangent direction is not the 3rd eigen-vector, then a test for an M-Connector is performed. Such points generally occur at branches. By following the connector curves, the traversal process will continue; one branch, for example, will be taken; and a longer extent of a tubular object will be extracted.

3.3. Tubular Widths Estimation via Maxima of Medialness

Having extracted the extent of the central skeleton of a tubular object, all that remains is an estimation of the tube's width along that track. We define the radius at a point \mathbf{x} on a ridge via

$$r_{\mathbf{x}} = \text{arg-local-max}_r [M(\mathbf{x}, r)]$$

where $M(\mathbf{x}, r)$ is a medialness function. The local maximum for the first ridge point is calculated with respect to the user-supplied initial width estimate which, as previously mentioned, is required to designate a tube of interest (Section 3.1) and which is also used to provide a scale σ for the height ridge calculations (Section 3.3). That local maximum is used to find the local maxima at its neighboring ridge points, those maxima in turn initiate their neighbors' local maxima calculations, and so forth.

Medialness functions respond maximally when applied at the center of a gray-scale object, at a scale (radius r) proportional to the width of the object. Medialness functions are distinguished by the shape of the kernel they use to measure medialness.

We have devised a medialness function whose kernel is optimized for the extraction of tubular objects, given that the central skeleton of those objects is already known. Specifically, to measure medialness at a point \mathbf{x} and a scale r , we use the weighted-sum of the responses from a series of cylindrical, Laplacian-of-a-Gaussian medialness kernels applied at scale r along and oriented to the central skeleton about the point \mathbf{x} .

A cylindrical Laplacian-of-a-Gaussian medialness kernel is a 2D Laplacian-of-a-Gaussian in the plane normal to \mathbf{x} with Gaussian fall-off in intensity tangent to that plane. The parameter r defines the scale of the Laplacian and the Gaussian fall-off. This kernel is illustrated in Figures 5 and 6.




Figure 5. (A) A 2D Laplacian-of-a-Gaussian kernel is equivalent to a slice through the normal plane of a 3D cylindrical Laplacian-of-a-Gaussian kernel. (B) The slice in 5A is shown as a height surface.




Figure 6. (A) A slice 2D along a 3D cylindrical Laplacian-of-a-Gaussian. (B) The slice in 6A is shown as a height surface.

The kernel's response is measured at the points \mathbf{x}_{-2t} \mathbf{x}_{-t} \mathbf{x} \mathbf{x}_{+t} \mathbf{x}_{+2t} where t is chosen so as to space the kernels by approximately one voxel. The distance (αt) of a kernel's center from the central kernel's center (\mathbf{x}) weights that kernel's response.

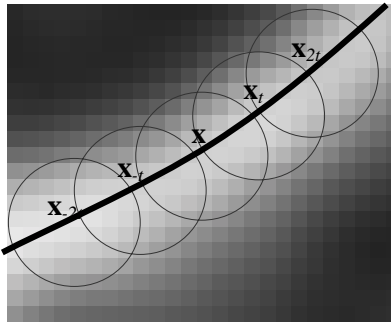


Figure 7. The responses from multiple kernels are weighted and summed to estimate the width at a point \mathbf{x} on the central skeleton.

This multi-kernel approach to width estimation provides improved performance by exploiting the spatial consistency of tubular objects common in medical images. There is significant benefit to this approach. Two of the most important benefits are that (1) it covers a large extent of the tube, thereby providing additional insensitivity to image noise; and (2) it fits the effective kernel to the spatial curvature of the ridge, thereby reducing assumptions about the shape of the tube.

Together, the central skeleton and widths well describe a tubular object in a medical image. Examples of the types of visualization that are possible using this method are given in Section 5.

3.4. Towards Complete Automation

We have devised two automated methods for extracting *multiple* tubular objects from an image. These techniques differ in the assumptions they make regarding the relative configurations of the tubes, and thereby they differ in the processing time they require to analyze a volume. The first technique (Section 3.4.1) makes no assumption regarding how the tubes are interconnected. The second technique (Section 3.4.2) assumes the tubes form a tree and thereby significantly reduces the processing time. We refer to these processes as the extract-all and the extract-tree processes respectively. It should be noted that these methods have not demonstrated the broad applicability of the semi-automated (i.e., point-and-click) method that we have presented; in general, these processes need to be tuned for each anatomical tube group and modality to which they are applied. Work on these automated methods is continuing.

3.4.1. Extract-All

This auto-extraction process involves a full search of the image volume. At each point in the search, the local maximum in the gradient direction is identified. That point is then used to stimulate the tube extraction process (y_0). We place tighter restrictions on the height-ridge criteria (e.g., test for equality to zero) for the initial point on each potential skeleton in order to reduce the extraction of short, ill-defined tubular objects. Yet, vessels are not missed since valid tubes will meet the tighter ridge criteria at some point along their extent, if not at every point. We do not allow the same tube to be extracted twice (see termination criterion, Section 3).

This auto-extraction technique has proven useful in the extraction of vessels of the lung and in the processing of certain confocal microscopy images. Consider the extraction of all vessels within a subspace of a chest CTA. In

Figure 8, the user specified a rectangular region within the shown slice. All vessels passing through that region were then automatically extracted. Note the long segments that were extracted and the multiple branching points.




Figure 8. (A) Our method automatically extracts all vessels passing through a portion (indicated by the box) of a slice. The intersections between that slice and the extracted representations are overlaid. Pixel resolution gives the illusion of the representations overestimating the widths. (B) A 3D rendering of the tubular representations generated via the search of the region in 8A; the slice in 8A is transparently blended with the 3D rendering.

The speed of this method depends most heavily on the number of search points that it must consider. The exhaustive search of a $256 \times 256 \times 48$ voxel volume can take several hours. Various extensions of this technique which address this issue are possible. The most viable speedup works by limiting the stimulation points searched to those points above a threshold value. This constraint is also helpful if extraneous vessels are being extracted from within a particularly noisy image. This additional constraint greatly reduces the processing time, however, as a threshold, must be adjusted for each image processed.

3.4.2. Extract-Tree

The second method limits the auto-extraction search volume to the space just beyond surface of each tube. The user must specify one tube to be extracted. All tubes that abut that tube will then be automatically extracted. Once

again, stricter ridge criteria are employed to eliminate the extraction of ill-defined tubes. This automated extraction technique combined with a threshold constraint has proven useful for generating representations of kidney vessel trees. It still takes over an hour to process such a tree, but no user-interaction is required during that time.

4. Speed, Consistency, Sensitivity, and Specificity of Tubular Object Extraction

By requiring the minimal interaction to designate a tube-of-interest, the time required to extract the normal vasculature captured by an MRA is near minimal for a semi-automated method. Specifically, with little training, a user can use our semi-automated (point-and-click) method to extract such vasculature in about 30 minutes on a 266 Mhz Pentium II laptop computer. The process generally requires about 60-70 point-and-clicks. The only parameter adjustment potentially required is an increase in object scale (width estimate) for the extraction of the carotids. Because of the consistency, sensitivity, and specificity of our method, a person with little anatomical knowledge can be quickly trained to perform such extractions. Additionally, in many situations, only a subset of the tubular objects in an image (e.g., one hemisphere of the intra-cranial vasculature) needs to be extracted. We therefore consider 30 minutes to be acceptable for a pre-operative planning system. Even this nominal time can be significantly reduced if the one of the automated methods (Section 3.4) can be employed.

Critical to the utility of the representations generated by our method are the issues of consistency, sensitivity and specificity. In this section we evaluate the consistency of our extraction method; that is, its dependence on the expertise of the user and on the setting of its parameters. Our method's sensitivity and specificity must be judged for each particular task to which its representations are being applied. Such analyses have been previously published in detail [Bullitt 1999; Bullitt 1997]. To summarize, those evaluations have shown that automated software developed by Elizabeth Bullitt can be used to post-process the representations extracted from an MRA in order to form accurate and complete representations of a patient's network of intracranial vasculature. These evaluations were carried out by interventional radiologists who were new to our methods and software. The

automated software makes occasional errors, so an interactive review of the segmentations and the network must be performed. The radiologists, however, judged the tools Dr. Bullitt provided for reviewing and modifying the network to be very acceptable – most review and edit sessions lasted 1-2 minutes.

To evaluate the consistency of our method, two items must be considered: the point-and-click (\mathbf{y}_0) used to designate a central skeleton and the initial width estimate (r) that completes the designation of the tube-of-interest and that determines the scale (σ) at which height-ridge measures are made. The user's expertise influences both of these parameters. These are the only user-adjustable parameters of our system.

Previously, we surmised and experience demonstrated that our method is insensitive to these parameters because it makes use of the geometry of tubes. As mentioned in the introduction, tubes are special in that blurring their images produces a central intensity ridge that also serves to stabilize the width estimation process. Scale-space properties ensure that the location of the extreme remains consistent across all scales above the image noise level. Consider the simple blurring of a 2D image with a central, uniform-intensity disc (tube cross-section) to which noise has been added (Figure 9). A broad range of blurring produces a consistently located extreme in the center of the object. Only the existence of neighboring objects can interfere with the consistency of an extreme through scale. For most tubes, neighboring objects are relatively distant compared to the width of a tube.




Figure 9 A disc of radius 25pixels with Gaussian noise and with Gaussian blurring with $\sigma=5, 10, 15,$ and 25 pixels.

The central extreme remains spatially consistent for a broad range of scales.

To quantify the consistency of our method, we performed a Monte Carlo experiment to evaluate the consistency of its extractions with respect to the location of the stimulation point (i.e., the point-and-click, \mathbf{y}_0) and the initial width estimates. Specifically, in an MRA, we designated a small and tortuous vessel (one containing slice is shown in Figure 10) for repeated extraction. We used our method to extract one "baseline" representation of this vessel. To simplify the comparison of multiple extractions, we limited this and subsequent extractions to a subset of image space. The subspace extends along the vessel for a large extent, i.e., ~ 20 voxels = ~ 20 mm, and in all three dimensions well beyond the vessel's edges. We then randomly generated, at three different initial width estimates (0.5, 1.0, and 2.0 mm), 100 stimulation points (\mathbf{y}_0) distributed uniformly along the vessel segment and extending uniformly $\pm 2r$ in three dimensions about the vessel. We then extracted the 300 tubular objects associated with those stimulation points. Plotting X-vs-Y (ignoring Z although it was also estimated) for the points on the central skeletons from all of those extractions produced Figure 11. That is, Figure 11 shows 300, overlapping central skeletons extracted using different scales and different stimulation points, for an extremely small vessel in an MRA. Occasionally neighboring vessels were extracted as indicated by the branches from the vessel-of-interest, but their extraction is not a failing of our technique – they are accurate extractions given certain stimulation points. Importantly, the maximum difference between the two closest points from any two different extractions of the vessel-of-interest was always much less than $1/10^{\text{th}}$ of a voxel! Our method produces extremely consistent central skeleton estimates - it is insensitive to user-interaction (point-and-click) and parameter settings (scale of image measurement).




Figure 10. One slice that the vessel-of-interest passes through




Figure 11. 300 random extractions (3 scales, 100 per scale) about the vessel-of-interest produce extremely consistent results (shown is a plot of x versus y for all 300 extractions)

For the extractions of the vessel-of-interest, plotting distance-along-ridge-vs-width (i.e., $2r$) produced the plot shown in Figure 12. Along the majority of the extent of the vessel-of-interest, a consistent estimate of the width was generated independent of the stimulation point and initial width estimate. However, at one end of the vessel (near the lower left-hand corner in Figures 11 and 12; not visible in slice shown in Figure 10), the vessel was extremely faint relative to the background and its width approaches the resolution of the scan (1.68 mm. inter-slice spacing). Under these extreme conditions, the width estimates provided by our system degrade, but they are still generally well behaved. We are investigating super-sampling techniques for the medialness kernel convolutions and alternative width estimation functions to improve the width estimates in these small-vessel situations. Our system does not exhibit such degradations for larger vessels.




Figure 12. Plot of distance-along-ridge versus width. Most width estimates

are quite consistent. Arrow indicates degradation in consistency that occurs when the vessel's contrast and width approaches the voxels' resolution.

5. Example Applications

Via the method presented in this paper, a variety of applications are possible. Most of our work has involved the generation of novel visualizations for pre-operative planning and 3D/2D registration for visualization augmentation and intra-operative guidance. In this paper, we illustrate (1) an intracranial arterio-venous malformation as captured by an MRA, (2) a variety of tubular structures in a chest CT, and (3) the small-bowel via CT of an infant. We also discuss (5) our multi-dimensional / multi-modal image fusion research that registers MRA with DSA via the central skeletons of vessels.

5.1. Intracranial Arterio-venous Malformations in MRA

Consider the difficult case of visualizing the complex network of arteries associated with an arterio-venous malformation. Traditionally, maximum intensity projections (MIP) of MRA or digital subtraction angiograms (DSA) would be used to visualize this network. Neither contains sufficient depth information to distinguish overlapping vessels. MIP images are often quite noisy and poorly relate small structures due to occlusion. While DSA images can only illustrate the sub-tree of vessels that are down-stream of the point of release for the contrast bolus and have associated issues concerning radiation and contrast toxicity. Recent research into 3D DSA is promising, but that technique, as with existing vascular imaging modalities, does not directly provide a representation of each vascular segment. As a result, while possibly providing nice visualizations, such methods do not facilitate higher-level reasoning about vessels. That is, without subsequent processing (e.g., the application of the method proposed in this paper), these modalities cannot be directly used to simulate the release of a

contrast bolus, to understand the effect of occlusion, to automatically identify a stenosis, or to visualize a subset of vessels that are relevant to a particular procedure.




Figure 13. Maximum intensity projection of an MRA. Note the aneurysm on the anterior communicating artery and the image noise. Such visualizations do not facilitate the selective viewing of a vascular tree or the characterization of a particular vessel




Figure 14. Digital subtraction angiogram of an arterial venous malformation. High resolution but only 2D and only of a single vessel sub-tree (i.e., beyond the bolus)

Via a central-skeleton and width representation of the vessels in an MRA, a 3D, lighted, surface rendering of the vessels can be generated (Figures 15). On a 266 Mhz Pentium II laptop computer, it required approximately 30 minutes to extract the illustrated tubes. On that same machine, the rendering's point-of-view can be updated at a rate of about one frame per second. Each vessel can be automatically searched for a possible stenosis. Via

software developed by Bullitt, vessel networks can be generated and sub-trees can be viewed in isolation based on their relevance to a surgical procedure, to simulate the passage of a catheter or to generate pseudo-angiograms. Johns Hopkins [REF] is evaluating our representations for use in a patient-specific surgical simulation virtual reality system. Another benefit is the compactness of these representations. Instead of requiring several megabytes of memory (as with MRA), the vessel representations from multiple patients require only a few kilobytes of memory; thus reducing the cost of storage and transmission.




Figure 15. (A) Surface rendering of tubular representations extracted from an MRA of a patient with an arteriovenous malformation; interested translucently into the 3D visualizations is a slice from the MRA volume. (B) Same patient's data: the left-hand image illustrates another slice of the MRA and its intersection with the extracted vessels is overlaid. The right-hand image is a 3D rendering from a different point of view that clearly illustrates the Circle of Willis and an aneurysm.

5.2. Tubular Objects in Chest CT

To illustrate the use of our system on images from other modalities and on other tubular objects in the body, consider the visualization of the vessels, bronchial tubes, sternum, spinal cord, and ribs from a CTA (Figures 16). Such visualizations may aid in the determination of the optimal resection cuts for live-donor partial lung

transplant or aid in the determination a lung's vascular profusion relative to lung volume or bronchial tube extent. Tubes can also be checked for nodules for interstitial tissue diagnosis.




Figure 16. Surface renderings of vessels, bronchial tubes, and ribs, sternum, and spinal cord from a CTA. The 3D spatial distribution of the vessels and bronchia is obvious, simplifying lobe identification.

The exact same code that was used to extract intracranial vessels from an MRA was also used to extract the representations in Figure 16 from the patient's CTA data. The only modifications employed were for the spinal cord and the bronchia; intensity valleys were followed and hollow tube widths were estimated. Such adjustments require little user interaction and occur instantaneously.

5.3. The Bowel via CT

Abdominal CT with contrast is used to detect and localize small bowel obstructions, polyps, and masses. Others have shown the benefit of CT colonography (virtual colonoscopy) for some of these processes [Dachman 1998]. Our center and widths process is ideal of defining the fly-through for such simulations. Figure 17 is a visualization of the bowel from a stillborn infant. The original CT data was acquired using air contrast. We have also applied our method to abdominal CT acquired with barium enema contrast.



Figure 17. A 3D rendering of the bowel of a stillborn infant. Air contrast was used to extend the bowel. The centers and width representation we generate is also ideal for automatically defining the path for virtual colonoscopy.

5.4. Multi-Dimensional / Multi-Modal Registration: MRA and DSA

One of the strengths of MRA is the 3D information it provides. One of its weaknesses is its relatively low resolution ($\sim 1\text{mm}$) and its dependence on the speed of flow (i.e., for time-of-flight MRA not for contrast MRA) which can result in aneurysms not being visible because of turbulent flow. One of the strengths of DSA is its high resolution (sub-mm). One of its weaknesses is its 2D format. Dr. Bullitt has devised a method for registering DSA with MRA and via that registration reconstructing orthogonal DSA projections into 3D. Dr. Bullitt's registration method is successful because it uses our method to extract the middles of the vessels from the MRA, and those middles are projected from candidate points-of-view to determine the point of view from which a DSA was acquired. Figure 18 shows the projection of the centers from an MRA onto a DSA after the registration process converged to a solution. Further details are provided in [Bullitt 1998]. Dr Bullitt has also shown that given two nearly orthogonal DSA images registered with an MRA, the MRA serves to disambiguate the vessel correspondences between the two 2D images and thereby enables a high resolution 3D reconstruction from those high resolution images.



Figure 18. The white lines are the vessel centers extracted by our method from a DSA. Those centers follow the vessels in the registered DSA on which they are overlaid. Via such registration, 3D reconstruction from multiple DSA is possible.

9. Conclusion

In this paper we presented a method for extracting tubular objects from medical images. Our system is shown to exploit the geometric properties of tubes so as to provide a clinically effective segmentation of 3D medical data which simplifies the development of a number of clinical tasks including pre-operative planning and intra-operative guidance. We detailed both the user and the algorithmic processes of that method. We used Monte Carlo simulations to demonstrate our method's consistency for a wide range of user-interaction skill and parameter choices. In approximately 30 minutes, our method can be used to extract the intra-cranial vessels captured by an MRA. Automated extraction methods have been devised and were presented. Thereby, our system has demonstrated consistency, speed, and ease-of-use. Its clinical accuracy is detailed in [Bullitt, 1999].

Work is currently focusing on improving the width estimation process for small, faint vessels in MRA and on integrating our methods with the clinical environment. Specifically, we are investigating methods developed by Hoogeveen [Hoogeveen 1997] to provide better width estimates. Clinical applications being investigated include those listed; the current emphasis is on developing application-specific clinical interfaces.

8. References

[Aylward 1996] Aylward, S., Pizer, S., Bullitt, E., and Eberly, D., "Intensity Ridge and Widths for Tubular Object Segmentation and Description," *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 131-138

[Bookstein 1993] Bookstein F, and Green W, "A thin-plate spline for deformations with specified derivatives," *Proceedings of Mathematical Methods in Medical Imaging II*, SPIE 2035, pp. 14-44

[Bullitt 1999] Bullitt E, Aylward S, Liu A, Stone J, Mukherji S, Coffey C, Gerig G, Pizer SM, "3D graph description of the intracerebral vasculature from segmented MRA and tests of accuracy by comparison with x-ray angiograms," *Proceedings of Information Processing in Medical Imaging, Lecture Notes in Computer Science*, In Press

[Bullitt 1997] Bullitt, E., Liu, A., Aylward, S., Soltys, M., Rosenman, J., Pizer, S., "Methods for Displaying Intracerebral Vascular Anatomy," *American Journal of Neuroradiology*, 18, pp. 417-420

[Dachman 1998] Dachman AH, Kyniyoshi JK, Boyle CM, Samara Y, Hoffmann KR, Rubin DT, Hanan I, "CT Colonography with Three-Dimensional Problem Solving for Detection of Colonic Polyps," *American Journal of Radiology*, 171, pp. 989-995

[Eberly 1996] Eberly D, Ridges in Image and Data Analysis, Kluwer Academic Publishers, Dordrecht, 213 pages

[Frangi 1998] Frangi AF, Niessen WJ, Vincken KL, Viergever MA, "Multiscale Vessel Enhancement Filtering," *Proceedings of Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science 1496*, pp. 130-137

[Gerig 1993] Gerig G, Koller T, Szekely G, Brechbuhler C, Kubler O, "Symbolic description of 3-D structures applied to cerebral vessel tree obtained from MR angiography volume data," *Proceedings of Information Processing in Medical Imaging, Lecture Notes in Computer Science 687*, pp. 94-111.

[Hoogeveen 1997] Hoogeveen RM, Bakker CJG, Viergever MA, "Phase derivative analysis in MR angiography: reduced Vessel dependency and improved vessel wall detection in laminar and disturbed flow" *Journal of Magnetic Resonance Imaging*, 7, 1997, p. 321-330 and Hoogeveen RM, Vessel Visualization and Quantification by Magnetic Resonance Angiography, *Dissertation*, Image Sciences Institute, Utrecht University, The Netherlands, 133 pages

[Park 1998] Park W, Hoffman, EA, and Sonka M, "Segmentation of Intrathoracic Airway Trees: A Fuzzy Logic Approach", *IEEE Transactions on Medical Imaging*, 17(4), pp. 489-497

[Masutani 1998] Masutani Y, Schiemann T, Hohne KH, "Vascular Shape Segmentation and Structure Extraction Using a Shape-Based Region-Growing Model," *Proceedings of Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science 1496*, pp. 1242-1249

[Pizer 1996] Pizer SM, Eberly D, Morse BS, Fritsch DS, "Zoom-invariant vision of figural shape: the mathematics of cores," *Computer Vision and Image Understanding*

[Wilson 1997] Wilson DL and Noble JA, "Segmentation of Cerebral Vessels and Aneurysms from MR Angiography Data," *Proceedings of Information Processing in Medical Imaging, Lecture Notes in Computer Science 1230*, pp. 423-428