



McGill School of
Computer Science

A Multi-Scale Geometric Flow for Segmenting Vasculature in MRI: Theory and Validation

Maxime Descoteaux

School of Computer Science,
McGill University, Montréal

June 2004

A Thesis submitted to McGill University in partial fulfilment
of the requirements for the degree of Master of Science

© Maxime Descoteaux, 2004

Abstract

Often in neurosurgical planning a dual echo acquisition is performed that yields proton density (PD) and T2-weighted images to evaluate edema near a tumor or lesion. The development of vessel segmentation algorithms for PD images is of general interest since this type of acquisition is widespread and is entirely non-invasive. Whereas vessels are signaled by black blood contrast in such images, extracting them is a challenge because other anatomical structures also yield similar contrasts at their boundaries.

In this thesis we present a novel multi-scale geometric flow for segmenting vasculature from PD images which can also be applied to the easier cases of MR angiography data or Gadolinium enhanced MRI. The key idea is to first apply Frangi's vesselness measure [Frangi et al. (1998)] to find putative centerlines of tubular structures along with their estimated radii. This multi-scale measure is then distributed to create a vector field which is orthogonal to vessel boundaries so that the flux maximizing flow algorithm of Vasilevskiy and Siddiqi (2002) can be applied to recover them. We carry out a qualitative validation of the approach on PD, MR angiography and Gadolinium enhanced MRI volumes and suggest a new way to visualize the segmentations in 2D with masked projections. We also validate the approach quantitatively on a data set consisting of PD, phase contrast (PC) angiography and time of flight (TOF) angiography volumes, all obtained for the same subject. A significant finding is that over 80% of the vasculature recovered in the angiographic data sets is also recovered from the PD volume. Furthermore, over 25% of the vasculature recovered from the PD volume is *not* detectable in the TOF angiographic data.

Thus, the technique can be used not only to improve upon results obtained from angiographic data but also as an alternative when such data is not available.

Résumé

Il est souvent nécessaire, pour une bonne planification neurochirurgicale, d'obtenir une imagerie par résonance magnétique standard (IRM) de type densité protonique (PD) pour tenter d'évaluer l'oedème au pourtour d'une tumeur ou d'une lésion. Comme le sang qui circule dans les vaisseaux donne une absence de signal (noir) et que d'autres structures anatomiques voisines du cerveau ont un signal semblable, l'extraction des vaisseaux représente un défi de taille qui demande une contrainte de forme tubulaire précise. Le développement d'algorithmes de segmentation des vaisseaux sanguins pour l'imagerie par résonance magnétique de type PD est donc d'un très grand intérêt puisque ce type d'imagerie médicale non invasive est couramment utilisée.

Dans cette thèse, nous présentons un flot géométrique multi-échelle pour segmenter les vaisseaux sanguins automatiquement à partir des images acquises en PD. L'algorithme peut également être appliqué aux données angio IRM où il y a rehaussement par un produit de contraste (gadolinium). L'idée principale est d'appliquer la mesure de vaisseaux proposée par Frangi afin de trouver les lignes centrales principales des structures tubulaires et d'en estimer leur diamètre. Par la suite, cette mesure est distribuée pour créer un champ vectoriel qui est perpendiculaire aux parois des vaisseaux de sorte que le débit par flux maximal de Vasilevskyi et Siddiqi 2002 puisse être appliqué. Nous avons effectué une validation qualitative de cette approche sur des résonances magnétiques de type PD, des données angiographiques et des volumes IRM rehaussés par gadolinium. De plus, nous suggérons une nouvelle façon de visualiser les segmentations en 2D en utilisant des projections masquées par les segmentations 3D obtenues à l'aide de notre technique. Nous avons également validé l'approche quantitativement sur une IRM de type PD, une angiographie par contraste de phase (PC) et une angiographie par temps de vol (TOF), toutes obtenues chez le même sujet. Les résultats démontrent que 80% des vaisseaux mis en évidence par angiographie le sont également par

l'IRM de type PD. De plus, 25% des vaisseaux mis en évidence à partir du volume PD ne l'étaient pas à partir des données angiographiques par temps de vol (TOF).

Ainsi, la technique développée peut être employée non seulement pour améliorer les résultats obtenus à partir de l'angio IRM mais également comme alternative quand de telles données ne sont pas disponibles.

Acknowledgments

First and foremost, I would like to thank my supervisor, Prof. Kaleem Siddiqi, for introducing me to shape analysis and the subject of geometric flows. He has been a constant presence throughout my two years in the Master's program and was able to keep me motivated, challenged and productive at all times. His assistance and support go well beyond the expectation of a Master's thesis supervisor. He has given me all the tools to continue in research and academia.

I am indebted to Prof. Louis Collins for his collaboration in the image acquisition and quantitative validation of our approach. He has introduced me to many interesting problems in medical imaging. I also thank Bruce Pike, Simon Drouin, Ingerid Reinertsen and others at the Montreal Neurological Institute for helpful discussions and inspiring conversations on medical imaging.

I am grateful to my "shape analysis" colleagues past and present who shared lab spaces with me: Alexander Vasilevskiy, Abeer Ghuneim, Pavel Dimitrov, Carlos Phillips, Peter Savadjev, Scott McCloskey and especially Sylvain Bouix for his expertise, help and advice on many aspects of computer vision. We had a great work environment and many good laughs.

I am thankful to McGill Athletics for their great intramural sports and squash and tennis varsity programs. Special thanks to all my friends, roommates, teammates and Kaleem for being part of these extra-curricular activities. Without them, I would not have been able to keep my sanity.

A huge "thank you" to Melanie for her support and understanding from stopping me from becoming a workaholic. She was an important inspiration for all my accomplishments. I am most indebted to my family: H  l  ne, Claude, Pierre-Olivier, Andr  e, Andr  , Louise, Ir  ne who have given me unconditional trust and support since the beginning of my scientific career.

Contributions of Authors

Parts of this thesis have appeared in peer reviewed conferences or have been submitted to an archival journal for review. These parts have involved collaborations between myself, Kaleem Siddiqi and Louis Collins which proceeded in three phases.

First, the algorithm for vessel segmentation was designed in close collaboration with Kaleem Siddiqi. This lead to our first publication at the workshop on Computer Vision Approaches in Medical Image Analysis (CVAMIA) held in conjunction with the 8th European Conference on Computer Vision (ECCV) 2004. Chapter 3 of the thesis is based on this publication. All the implementation and programming with the MNI medical libraries was performed by myself.

Second, a careful quantitative validation protocol was designed with Louis Collins at the MNI. We were able to acquire three different image data sets, a proton density (PD) weighted MRI, a phase contrast (PC) angiogram and a time of flight (TOF) angiogram of my brain. This allowed us to compare and quantify the performance of our algorithm on a standard MRI. This work lead to the paper accepted at the Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2004 conference. This quantitative cross validation is part of Chapter 4.

Finally, Kaleem Siddiqi, Louis Collins and myself worked on a submission to an archival journal article combining the two previous publications and new qualitative cross validation results. The manuscript was submitted to a special issue on vascular imaging of the IEEE transactions on Medical Imaging and is currently under review.

Contents

1	Introduction	1
1.1	Imaging Modalities	3
1.2	Problem Statement	5
1.3	Method Overview	6
1.4	Contributions	7
1.5	Organization	8
2	Background	9
2.1	Modeling Vasculature using the Hessian	10
2.2	Vessel Segmentation and Centerline Extraction	17
2.2.1	Statistical Methods	17
2.2.2	Centerline Extraction	17
2.2.3	Geometric Flows	19
3	A Multi-Scale Geometric Flow	23
3.1	Introducing a Tubular Model	24
3.2	Extending the Vesselness Measure to Boundaries	28
3.3	The Multi-Scale Geometric Flow	31
3.4	Algorithms and Implementation Details	33
4	Validation	41
4.1	Image Acquisition	42

4.2	Qualitative Results	43
4.3	Quantitative Results	48
5	Discussion and Conclusions	53
5.1	Discussion	54
5.1.1	Why use a Geometric Flow?	54
5.1.2	Quantitative Validation and Ground Truth Data	55
5.2	Contributions and Summary	56
5.3	Future Work	58
	Bibliography	64

List of Figures

1.1	A proton density (PD) weighted MRI, time of flight (TOF) MRA, phase contrast (PC) MRA, and Gadolinium enhanced MRI	4
2.1	Basic differential geometry of surfaces	11
2.2	Direction of the normal and principal curvatures at a point on a surface	13
2.3	Eigen vector directions of the Hessian matrix at locations centered within tubular structures	14
2.4	Flux maximizing flow illustration for a 2D curve	20
3.1	Sagittal slice of the PD weighted MRI data set with its corresponding vesselness measure	25
3.2	Local vessel orientation on tubular examples	27
3.3	Distributing the vesselness measure to the implied boundaries	29
3.4	Approximate intensity profiles of $I, \nabla I, \phi, \nabla \phi$	32
4.1	Segmentation on a MRA example using three initial seeds	44
4.2	Automatic segmentation when re-sampling the original data set . . .	45
4.3	Segmentation on a Gadolinium enhanced MRI volume	46
4.4	Intensity projections of the original data sets, the vesselness measure and the segmented vessels on PC, TOF, and PD volumes	47
4.5	Pair-wise vessel extraction comparison obtained on the PC, TOF and PD volumes	49

List of Algorithms

3.1	Vesselness computation	34
3.2	Construction of the extended vector field	35
3.3	Level set based geometric flow	36

List of Tables

2.1	A classification of local structures	15
3.1	Behavior of $I, \nabla I, \phi$, and $\nabla \phi$ at vessel boundaries.	33
4.1	A pair-wise comparison between vessels extracted from different modalities	50

Chapter 1

Introduction

The visualization and quantification of cerebral vasculature can be extremely important in pre-surgical planning, image-guided neurosurgery and clinical analysis. A common approach is to use a maximum intensity projection (MIP) where three-dimensional (3D) data is projected onto a 2D plane by choosing the maximal intensity value along that projection direction. A major drawback of this method is that background artifacts and other tissues may occlude vascular structures of low contrast and small width. Thus, it is desirable to extract the vasculature tree before it is visualized. In this thesis, we define *segmentation* as the process of labeling 3D voxels as “vessel” or “non-vessel” points. Once the data is segmented and we have a 3D volumetric representation, the visualization and further analysis of the complex human vasculature is greatly simplified.

It is unfortunately often the case that in order to obtain such representations from a medical data set, an expert has to interact with the data manually, in a slice-by-slice fashion, while coloring regions of interest and connecting them using image processing operations. This process is extremely laborious and is prone to human error. Since a technician preparing data for surgical planning has a limited amount of time, there is a trade-off between the number of manually segmented structures and the quality of the segmentations. In addition, the significant amount of time required to properly segment the vasculature (e.g. from a single brain MRI) makes large scale clinical studies of vasculature infeasible. Another simplistic approach to vessel extraction is thresholding the original data set. Here, depending on the image modality, all voxels with intensity above or below a threshold are labeled as *vessel* and the others as *non-vessel*. However, due to non-homogeneous intensity distribution in medical data sets, a conservative threshold typically does not capture small and low contrast vessels and an aggressive threshold selection incorrectly labels many non-vessel or background voxels. As a consequence, the computer vision and image analysis community has paid significant attention to automating the extraction of vessels or vessel centerlines.

1.1 Imaging Modalities

Several methods in the computer vision literature have been shown to give promising results on 2D projection angiography, 3D Computed Tomography and Magnetic Resonance Angiography (CTA and MRA). These image modalities are currently the most widely used acquisition techniques when one seeks to extract the brain vasculature. While these imaging approaches acquire high-contrast and high resolution volumes and are designed to image blood vessels, each has its own limitations and can be invasive due to contrast agent injection and radiation.

In CTA, X-ray contrast material is injected directly into the blood stream through a catheter. Tomographic images are then generated by collecting 1D X-ray signals of an object at many angles. Then, a cross-sectional image is reconstructed representing the attenuation coefficient of the X-ray beam in that slice. This process is repeated over many planes to construct 3D volumetric data. A similar method is Computed Rotational Angiography (CRA) which produces 3D data sets by acquiring projection radiographs from many angles around the patient, followed by a reconstruction procedure using CT algorithms. These angiograms are generally more accurate than standard MRA acquisition. However, the contrast agent injection and radiation dose given to the patient are major drawbacks.

Magnetic Resonance Imaging (MRI) of the blood vessels is referred to as Magnetic Resonance Angiography (MRA). MRI is a largely noninvasive technique which utilizes the properties of magnetism of the hydrogen atoms in our body to create nondestructive, three-dimensional, internal images of the soft tissues of the body, including the brain. In the context of blood vessel acquisition, there are three widely used methods. The first is Phase Contrast (PC) angiography, in which contrast is determined by tissue motion. Static tissue yields no signal, and is therefore black, as in Figure 1.1(c). The second is Time Of Flight (TOF) angiography, where vessel brightness is proportional to blood flow velocity. However, complex flow or turbulence can cause signal loss in the vessels in such data, as seen in 1.1(b). In




FIGURE 1.1: A mid-sagittal slice of a proton density (PD) weighted MRI volume (a), a time of flight (TOF) MR angiogram (b) and a phase contrast (PC) MR angiogram (c) of the same subject acquired at the Montreal Neurological Institute. A Gadolinium enhanced MRI acquired on a patient with a brain tumor is shown in (d).

these data sets, the vessel/non-vessel contrast is sharp *only* at vessel boundaries. Hence, a simple thresholding of these volumes typically yield a crude estimate of the vascular structure. This makes the segmentation problem easier. The third angiographic-like image acquisition used to highlight vessels is Gadolinium enhanced MRI, seen in Figure 1.1(d). It is the most invasive angiographic technique and is the method currently used in almost all neurosurgical cases involving brain tumors. Gadolinium is a contrast agent injected into patients to alter the signal intensity of soft tissues as well as the blood pool. Hence, Gadolinium enhances blood vessels but is also absorbed by non-vessel surrounding tissues resulting in several bright/dark contrast changes in the data. This makes the segmentation problem a significant challenge. Standard algorithms designed to work on MRA and CTA

typically fail on such volumes.

Very few techniques currently exist for the automatic extraction of vessel boundaries in more standard anatomical MRI volumes such as the proton-density (PD) weighted data set in Figure 1.1(a) and the Gadolinium enhanced MRI. In PD, it is clear that a signal decrease is present in the vascular regions (the spaghetti-like structures), but there are several other bright/dark contrast change at boundaries of non-vessel structures (between gray and white matter, cerebellum area). Also, the contrast between blood vessel and surrounding tissue is not as great when compared to the angiographic sequences (1.1(b) and 1.1(c)). Hence, the problem of recovering vessels from image intensity contrast alone is a challenge and requires shape information to constrain the segmentation.

1.2 Problem Statement

The goal of this thesis is to solve the segmentation problem on common clinical MRI. In particular, our aim is to automatically classify as much of the vascular structure as possible. The vessel extraction must be accurate and competitive (similar or better) to vessel segmentation achieved from the easier cases of MRA and CTA and the algorithm must be able to extract vessels of variable widths and contrast.

If successful, such a procedure could be used in surgical planning while eliminating the need for an additional scan. This would save time during image acquisition and would ease the burden on the patient as well as reduce the amount of time required to segment and prepare data for use in planning. The 3D vessel structure from our approach could be used as the basis for registration between different non-angiographic modalities. One such application is the registration between intra-operative ultrasound and pre-operative Gadolinium enhanced MRI or PD weighted MRI, to estimate brain shift during brain tumor surgery [Reinertsen et al. (2004)]. Finally, the method could be useful for visualization of the vascu-

lar networks of different organs such as the brain, the liver and the lungs. With a true three-dimensional representation, the complex spatial relationships between the vasculature and surrounding anatomical structures could be made explicit. A user could interact with the derived model, depending upon the task at hand, and could visualize it from arbitrary viewing directions. This is very important in minimally invasive neurological surgery. Typically, a needle is inserted in the scalp of the patient to access the region that is operated on, such as a tumor. However, the neurosurgeon does not see the tip of the needle and all navigation is guided by images. It is thus of utmost importance to have a precise visualization of the location of blood vessels in order to avoid puncturing them.

1.3 Method Overview

We introduce a novel algorithm for vessel segmentation which is designed for the case of PD images, but can be applied as well to angiographic data or Gadolinium enhanced MRI volumes. The algorithm is motivated in part by the approach of Ostergaard et al. (2000) where Frangi's vesselness measure [Frangi et al. (1998)] is thresholded to find centerlines. In this technique, tubular fits to vessel boundaries are then obtained using a form of connected component analysis and a generalized cylinder model. This latter step typically yields results that are disconnected. In our approach, rather than threshold the vesselness measure, we extend it to yield a vector field which is locally normal to putative vessel boundaries. This in turn allows the flux maximizing geometric flow of Vasilevskiy and Siddiqi (2002) to be applied to recover vessel boundaries. This flow has a formal motivation, is topologically adaptive due to its implementation using level set methods, and finally is computationally efficient. We show qualitative results on magnetic resonance angiography (MRA) data, as well as on the more challenging cases of Gadolinium enhanced MRI and proton density (PD) weighted MRI volumes. We also validate the approach quantitatively by comparing the segmentations from PD, PC angiogra-

phy and TOF angiography volumes, all obtained for the same subject (Figure 1.1).

1.4 Contributions

In this thesis, we propose a three step algorithm for blood vessel segmentation. We first introduce a tubular structure model incorporating local vessel centerline orientation and width. Then, we extend this measure to the implied vessel contours to finally apply a flux maximizing geometric flow. The main contributions can be summarized as follows:

1. We describe a new multi-scale geometric flow which can extract vasculature from standard MRI. The approach is able to segment blood vessels on several image modalities, including MRA, Gadolinium enhanced MRI, and PD weighted MRI.
2. We propose a 2D visualization of the vasculature by intensity projections (MIPs) of the original volume masked by the binary segmentation obtained by our algorithm.
3. We carry out a qualitative comparison of the vessel extraction on PD, PC and TOF volumes obtained from the same subject. This suggests that the PD segmentation improves upon results obtained from TOF angiography and is very similar to that obtained from PC angiography.
4. We perform a careful quantitative validation confirming our qualitative observations. In particular, we note that 80% and 89% of the PC and TOF data respectively, is accounted for by the PD segmentation. Moreover, 26% of the PD reconstruction is not present in the TOF vessel extraction.

1.5 Organization

The thesis is outlined as follows. In Chapter 2 we review relevant background literature on the modeling of tubular structures, vessel segmentation and center-line extraction. We then develop our multi-scale geometric flow by incorporating Frangi's vesselness measure [Frangi et al. (1998)] in the flux maximizing flow algorithm of Vasilevskiy and Siddiqi (2002) in Chapter 3. We present qualitative and quantitative validation results in Chapter 4. We then conclude with a discussion of the results and present directions for future work in Chapter 5.

Chapter 2

Background

We now review the use of the Hessian as a descriptor for modeling tubular structures and then provide an overview of vessel segmentation and centerline extraction methods in the literature. This overview is necessarily not exhaustive; it is based on a selection of representative techniques. For a more thorough discussion of the relative strengths and weaknesses of such approaches we encourage the reader to refer to the recent article of Aylward and Bullitt (2002). Also, we refer to relatively standard differential geometry definitions, propositions and theorems that are described in greater detail in any standard differential geometry text such as DoCarmo (1976).

2.1 Modeling Vasculature using the Hessian

Several multi-scale approaches to modeling tubular structures in intensity images have been based on properties of the Eigen values of the Hessian matrix \mathbf{H} [Lorenz et al. (1997); Sato et al. (1998); Frangi et al. (1998); Aylward and Bullitt (2002); Koller et al. (1995); Krissian et al. (2000); Ostergaard et al. (2000); Wink et al. (2004)]. For a function $f(x_1, x_2, \dots, x_n)$, the Hessian is given by the Jacobian of the derivatives $\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \dots, \frac{\partial f}{\partial x_n}$. This matrix encodes important local shape information. To understand why this is so, we must review some basic concepts in differential geometry of surfaces. Referring to Figure 2.1, we look at how rapidly a surface \mathcal{S} pulls away from the tangent plane \mathcal{T}_p in a neighborhood of a point $p \in \mathcal{S}$. This is the same as measuring the rate of change, $d\mathcal{N}_p$, of the unit normal vector field \mathcal{N}_p on a neighborhood of p . It can be shown that this differential $d\mathcal{N}_p$ is a self-adjoint linear map [DoCarmo (1976)] giving rise to the *second fundamental form* II_p of a surface \mathcal{S} at a point p . To see how the Hessian operator appears in this shape analysis of surfaces, we consider surfaces given as the graph of a differentiable function $z = h(x, y)$. Such graphs are common in the computer vision and in the active contour literature. For example, the intensity values of a 2D image are often regarded as a height surface $z = I(x, y)$. Most importantly, it is known that locally, any surface is the




FIGURE 2.1: Basic differential geometry of surfaces. Local surface representation of a regular surface S . T_p is the tangent plane of the surface at a point $p \in S$, π_N is the plane of the normal vector \mathcal{N}_p to the surface at point p .

graph of a differentiable function [DoCarmo (1976) (cf. Prop 3, Sec. 2-2)]. That is, given a point $p \in S$, one can choose the coordinate axis of \mathbf{R}^3 so that the z -axis is along the normal of the surface (\mathcal{N}_p) and the xy plane agrees with T_p . Thus, a neighborhood of $p \in S$ can be represented in the form $z = h(x, y)$. If the surface is parametrized as $(x, y, h(x, y))$, a simple computation [DoCarmo (1976) (Sec. 3-3 ex.5)] shows that the unit normal field is given by

$$\mathcal{N}(x, y) = \frac{(-h_x, -h_y, 1)}{\sqrt{h_x^2 + h_y^2 + 1}}$$

Thus, the second fundamental form of such a surface at a point p applied to a vector $(x, y) \in \mathcal{T}_p$ becomes,

$$\begin{aligned}
 II_p(x, y) &= -\langle d\mathcal{N}_p((x, y)), (x, y) \rangle \\
 &= h_{xx}x^2 + 2h_{xy}xy + h_{yy}y^2 \\
 &= (x \ y) \begin{pmatrix} h_{xx} & h_{xy} \\ h_{xy} & h_{yy} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}
 \end{aligned} \tag{2.1}$$

In this case, the Hessian of h is the second fundamental form of \mathcal{S} at p .

In general, the second fundamental form II_p has some very important geometric properties in the tangent plane \mathcal{T}_p for any surface. In particular, the value of II_p for a unit vector $v \in \mathcal{T}_p$ is equal to the normal curvature of a regular curve passing through p and tangent to v . In fact, we can show that all curves lying on a surface \mathcal{S} and having the same tangent line along v at point $p \in \mathcal{S}$ have the same normal curvature. This allows one to speak of the normal curvature in a particular direction v at p . We are usually interested in the extreme values, *maximum* (κ_1) and *minimum* (κ_2), of the normal curvature. These are called *principal curvatures*. A nice theorem states that there exists an orthonormal basis $\{e_1, e_2\}$ of \mathcal{T}_p such that $d\mathcal{N}_p(e_1) = -\kappa_1 e_1$ and $d\mathcal{N}_p(e_2) = -\kappa_2 e_2$, [DoCarmo (1976)]. Thus, the normal curvature can always be expressed as a linear combination of the minimum and maximum curvatures. In our context, this means that the Eigen values of the Hessian matrix give the principal curvatures (κ_1, κ_2), and the corresponding Eigen vectors (e_1, e_2) span the tangent plane \mathcal{T}_p (Figure 2.2). Hence, the Eigen value decomposition of the second fundamental form (or the Hessian matrix) is all one needs to locally describe the shape of a 2D surface.

In computer vision, one often works with three-dimensional (3D) images. Hence, we must extend the 2D differential geometry to 3D iso-intensity surfaces present in the image data, $\mathcal{I}(x, y, z)$. In this thesis, we model blood vessels as closed tubular iso-surfaces as is popular in the medical imaging literature [Koller et al. (1995);




FIGURE 2.2: Illustration of the direction of the gradient or normal \mathcal{N} , the direction of the minimum and maximum curvatures of a surface \mathcal{S} at a point p .

Frangi et al. (1998); Krissian et al. (2000)]. In particular, at any given voxel in a 3D image \mathcal{I} , we wish to know if we are *inside*, *on* or *outside* a tubular structure implied by the data. To do so, we must explore the variations of intensity in small regions. A common approach to analyzing local shape behavior in the neighborhood $\Delta\bar{x}$ of a voxel \bar{x} of an image \mathcal{I} is to consider its Taylor expansion. Neglecting terms of degree higher than two we obtain

$$I(\bar{x} + \Delta\bar{x}) \approx I(\bar{x}) + \Delta\bar{x}^T \nabla I(\bar{x}) + \Delta\bar{x}^T \mathbf{H}(\bar{x}) \Delta\bar{x}$$

where

$$\mathbf{H} = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix} \quad \text{and} \quad \nabla I = \begin{pmatrix} I_x \\ I_y \\ I_z \end{pmatrix}$$

The vector of first derivatives of the image is the gradient vector $\nabla \mathcal{I}$ and it gives the normal vector (\mathcal{N} in Figure 2.2) to the implied iso-intensity surface. The Hessian




FIGURE 2.3: At locations centered within tubular structures the Eigen vector (e_1) corresponding to the smallest Eigen value (λ_1) of the Hessian matrix is along the vessel direction and the other Eigen vectors $\{e_1, e_2\}$ span the cross-sectional plane.

matrix \mathbf{H} looks at how this normal vector varies in all three directions. Intuitively, when on a tubular iso-surface, the normal vector variation along the tube is small due to low curvature whereas the variation is important in the other two orthogonal directions due to high curvature of the cross-section. Hence, in the tubular case, one expects a row of \mathbf{H} to be composed of zero or close to zero entries (low curvature in that direction) and have the other two rows equal or almost equal (high curvature of circular or almost circular cross-section). In this case, the Eigen decomposition of the Hessian matrix \mathbf{H} , which seeks for vectors $\vec{e} \in \mathbf{R}^3$ and scalars λ such that

$$\mathbf{H} \vec{e} = \lambda \vec{e}$$

gives a zero or close to zero Eigen value and two other equal or almost equal Eigen values with high magnitude. The associated Eigen vectors form a coordinate frame giving the minimum and maximum curvature directions in the tangent plane to the iso-surface at that point and the direction of the normal vector. This is illustrated in Figure 2.3.

The Eigen value analysis can be extended to differentiate *tube-like*, *blob-like*, *sheet-like*, and *noise-like* structures from one another as summarized in Table 2.1. Sheet-like or plate-like structures are encountered in data sets with flat bones, skin

Eigen value conditions	local structure	examples
$\lambda_1 \approx 0, \lambda_2 \approx \lambda_3 \gg 0$	tube-like	vessel, bronchus
$\lambda_1 \approx \lambda_2 \approx 0, \lambda_3 \gg 0$	sheet-like	cortex, skin
$\lambda_1 \approx \lambda_2 \approx \lambda_3 \gg 0$	blob-like	nodule
$\lambda_1 \approx \lambda_2 \approx \lambda_3 \approx 0$	noise-like	noise

TABLE 2.1: A classification of local structures based on the Eigen values of the Hessian matrix. Here, we assume that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. The sign of the highest Eigen values generally indicate whether the local structure is dark on a bright background or bright on a dark background. A positive sign corresponds to a dark structure on a bright background which is the case for PD weighted MRI volumes.

or cortex, where two of the Eigen values are close to zero due to small normal changes in the plane corresponding to the flattened shape. For sphere-like or blob-like structures, it is expected that all three Eigen values are high and almost equal because of the isotropy of the data in all three directions. This local structure is often detected at branch points and at very high curvature sections of blood vessels, as pointed out later in Figure 3.2. Finally, close to zero Eigen values represent locations with the absence of structure. This is often the case for points in the background or in noisy parts of the data. In this thesis, we are interested in tube-like structures for the task of segmenting vasculature. Two prominent approaches for capturing vessel-like or tube-like structures based on the Hessian are the techniques proposed in Krissian et al. (2000) and Frangi et al. (1998).

First, Krissian et al. (2000) propose a model-based approach to detecting tubular structures. An Eigen value decomposition of the Hessian matrix is carried out analytically for each assumed model that is fit to the image data. They report that whereas this analysis provides a good descriptor at the center of a vessel, its quality decreases at locations close to vessel boundaries. Hence, they define a vessel detector which combines the highest two Eigen values of the Hessian matrix and a gradient term which is known to play a significant role at vessel boundaries. They have recently demonstrated the robustness of this operator in the context of segmenting the aorta in low contrast 3D ultrasound images [Krissian et al. (2003)].

Second, Frangi et al. (1998) propose a vesselness measure which incorporates information from all three Eigen values and has an intuitive geometric interpretation. This method is close in spirit to previous work by Lorenz et al. (1997) and Sato et al. (1998). Three quantities are defined to differentiate blood vessels from other structures:

$$R_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}} \quad R_A = \frac{|\lambda_2|}{|\lambda_3|} \quad S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}.$$

From Table 2.1, it can be seen that R_B is non zero only for blob-like structures. The R_A ratio differentiates sheet-like from other structures because it is zero only for sheet points. Finally, S , the Frobenius norm, is used to ensure that random noise effects are suppressed from the response. For all non noise-like structures, this measure is high because at least one of the Eigen values is significant. For a particular scale σ the intensity image is first convolved by a Gaussian at that scale, $G(\sigma)$, and the following vesselness response function, $V(\sigma)$, is computed:¹

$$V(\sigma) = \begin{cases} 0 & \text{if } \lambda_2 < 0 \text{ or } \lambda_3 < 0 \\ (1 - \exp\left(-\frac{R_A^2}{2\alpha^2}\right)) \exp\left(-\frac{R_B^2}{2\beta^2}\right) (1 - \exp\left(-\frac{S^2}{2c^2}\right)). & \end{cases} \quad (2.2)$$

This measure is designed to be maximum along the centerlines of tubular structures and close to zero outside vessel-like regions. The scale σ associated with the maximum vesselness response provides an estimate of the width of the tubular structure centered at a particular location and the Eigen vector associated with the smallest Eigen value of the Hessian gives its local orientation. This is illustrated in Figure 3.2 and will be further explained later when we develop our approach.

¹The vesselness expression is given for the case of a dark tubular structure on a brighter background (as in a PD volume). In the case of angiographic data, the signs in condition 1 must be changed, i.e., $V(\sigma) = 0$ if $\lambda_2 > 0$ or $\lambda_3 > 0$.

2.2 Vessel Segmentation and Centerline Extraction

In this section, we review three classes of vessel segmentation algorithms relevant in the context of this thesis. We first look at a statistical approach and then investigate ridge traversal and centerline extraction methods. Finally, we provide a discussion on other geometric flows existing for blood vessel segmentation from angiographic data sets.

2.2.1 Statistical Methods

Wilson and Noble (1997) propose a statistical approach for segmenting blood vessels from TOF angiography data, such as that shown in Figure 1.1(b). They introduce a mixture of three probability distributions which is based on physical properties of blood and brain tissues. Vessel labels are assumed to arise from a uniform distribution and two Gaussian distributions are used to model other structures, one for tissue outside the head and another for eyes, skin, bone and brain tissue. The parameters of these models are estimated using a classical expectation maximization (EM) algorithm. The vasculature tree is then obtained following a thresholding procedure that is sensitive to signal to noise ratio and intensity contrast between vessel and non-vessel structure in the data. It is important to point out that this method does not employ a multi-scale analysis and also has no explicit model for tubular structures. Hence, it cannot be applied to non-angiographic data sets such as the PD volume or to Gadolinium enhanced MR volumes of Figure 1.1(a) and Figure 1.1(d).

2.2.2 Centerline Extraction

Another class of methods attempts to find centerlines of tubular structures as they are manifest directly in intensity (MR or CT) images, such as those in Figure 1.1(c). Aylward and Bullitt (2002) present a centerline tracking approach which is based

on a characterization of intensity ridges in 3D data sets. The Eigen vectors of the Hessian matrix are used to estimate the local orientation of vessels and a normal plane is iteratively updated to follow the vessel's cross-section, as illustrated locally in Figure 2.3. This idea is also the basis of work by Koller et al. (1995) for the multi-scale detection and traversal of curvilinear structures in intensity images. Aylward and Bullitt pay particular attention to the validation of their method, demonstrating its robustness under parameter changes, changes in scale and simulated image acquisition noise. The method is an iterative one, where the centerline is continuously extended in the estimated direction of its local orientation. As we shall later see, this local Hessian analysis is similar to the one used in our geometric flow based approach. However, rather than traverse the ridge at a single scale and compute vessel widths using a multi-scale analysis, we use multi-scale orientation and scale estimates directly to propagate information from centerlines to vessel boundaries.

Deschamps and Cohen (2001) relate the problem of finding centerline paths to that of finding paths of least action in 3D intensity images. This leads to a form of the well-known Eikonal equation where a front is propagated in the image with a speed determined by a scalar potential that depends upon location in the medium. The minimal path is extracted using a simple steepest gradient descent. The framework aims to infer the boundaries of tubular structures in a first stage, using a standard surface evolution method. The potential function is then designed to take into account a Euclidean distance function from the boundary, so that the minimal paths are centered. The flow is implemented using fast marching schemes and is hence computationally efficient. The algorithm requires little user interaction but the user must specify the starting and end points of a particular path. As a consequence, the method is only applicable to small portions of blood vessels since it extracts only a single path at a time. The major drawback is that the technique cannot, in its current form, handle bifurcations and multiple trajectories

naturally, making global segmentation of vasculature quite difficult.

Wink et al. (2004) have recently presented an approach to centerline extraction, applied in the context of vessel tracking, which combines features of the above two approaches. More specifically, they use Frangi's vesselness measure (Eq. 2.2) to characterize putative vessel centerline locations [Frangi et al. (1998)]. They then formulate the problem of finding paths between user selected points as a minimum cost path problem which they solve computationally using wavefront propagation. The interesting feature of their approach is that they do a 3D search for a minimum path in the space of multi-scale responses. Standard methods usually take a maximum projection of the multi-scale responses to perform a 2D search. To achieve this 3D search across scales, they incorporate a term that controls how easily the scale of the vessel can be changed in the path tracking. Their method has been validated qualitatively in the presence of stenoses, vessel crossings, several proximal vessels and imaging artifacts. However, the technique is again a local one and it does not handle the branching of vessels naturally.

2.2.3 Geometric Flows

There is a long history on the use of deformable models for segmentation in the computer vision literature, motivated in large part by the classical parametric snakes introduced by Kass et al. (1987). These models have also been extended to handle changes in topology due to the splitting and merging of contours [McInerney and Terzopoulos (2000)]. In the context of geometric flows for shape modeling, the first curve and surface evolution models were developed independently by Malladi et al. (1993, 1994, 1995) and Caselles et al. (1993). This work lead to two recent vasculature segmentation approaches which are relevant to the development here. First, Lorigo et al. (2001) propose a regularization of a geometric flow in 3D using the curvature of a 3D curve. This approach is grounded in the recent level set theory developed for mean curvature flows in arbitrary co-dimension based on work




FIGURE 2.4: Illustration of the flux maximizing flow of Vasilevskiy and Siddiqi (2002) for a 2D curve placed in a vector field. The curve evolves as to increase the inward flux through its boundary as fast as possible. The resting flux maximizing configuration is one where the inward normals to the curve are everywhere aligned with the direction of the vector field. The figure is adapted from Vasilevskiy and Siddiqi (2002).

by Ambrosio and Soner (1996). It yields the flow

$$\psi_t = \lambda(\nabla\psi, \nabla^2\psi) + \rho \langle \nabla\psi, \nabla\mathcal{I} \rangle \frac{g'}{g} \nabla\psi \cdot \mathbf{H} \frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|}.$$

Here ψ is an embedding surface whose zero level set is the evolving 3D curve, λ is the smaller nonzero Eigen value of a particular matrix [Ambrosio and Soner (1996)], g is an image-dependent weighting factor, \mathcal{I} is the intensity image and \mathbf{H} is the determinant of its Hessian matrix. For numerical simulations the evolution of the curve is depicted by the evolution of an ϵ -level set. Without the multiplicative factor $\rho \langle \nabla\psi, \nabla\mathcal{I} \rangle$ the evolution equation is a gradient flow which minimizes a weighted curvature functional. The multiplicative factor is a heuristic which modifies the flow so that normals to the ϵ -level set align themselves (locally) to the direction of image intensity gradients (the inner product of $\nabla\psi$ and $\frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|}$ is then maximized). The flow is designed to recover vessel boundaries signaled by the gradient in angiography data, while under the influence of a smoothing term driven by the mean curvature of an implied centerline.

Second, Vasilevskiy and Siddiqi (2002) derive the gradient flow which evolves a curve (2D) or a surface (3D) so as to increase the inward flux of a fixed (static) vector field through its boundary as fast as possible (Figure 2.4). With S an evolving

surface and $\vec{\mathcal{V}}$ the vector field, this flow is given by

$$\mathcal{S}_t = \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}} \quad (2.3)$$

where $\vec{\mathcal{N}}$ is the unit inward normal (for inward motion) or outward normal (for outward motion) to each point on S . The motivation behind this flow is that it evolves a curve (in 2D) or surface (in 3D) to a configuration where its normals are aligned with the vector field, as seen in Figure 2.4 for the 2D case. In the context of segmenting vasculature in angiographic images, $\vec{\mathcal{V}}$ can be selected to be the gradient of the intensity image which is expected to be orthogonal to vessel boundaries [Vasilevskiy and Siddiqi (2002)].

It is important to point out that both of the above approaches are designed specifically for angiographic data and hence require restrictive assumptions to hold. In particular: 1) both methods are initialized essentially by thresholding such data, and thus would fail when vessel boundaries cannot be identified from contrast alone; 2) neither approach has an explicit term to model tubular structures, but instead relies on the assumption that the gradient of the intensity image yields a quantity that is significant *only* at vessel boundaries; and 3) neither of these methods takes into account explicitly the multi-scale nature of vessel boundaries as they appear in all modalities. In the following chapter we argue that several of the above limitations can be overcome by incorporating a measure of “vesselness”. The result is a modified flow which can be applied to a wide range of modalities, and which also offers computational advantages over other vessel segmentation algorithms due to its implementation using level set techniques [Osher and Sethian (1988)].

Chapter 3

A Multi-Scale Geometric Flow for Segmenting Vasculature

The approach we develop proceeds in two steps. First, we apply Frangi's vesselness measure to find putative centerlines of tubular structures along with their estimated radii. Second, this multi-scale measure is distributed to create a vector field which is orthogonal to vessel boundaries so that the flux maximizing flow algorithm of Vasilevskiy and Siddiqi (2002) can be applied to recover them.

3.1 Introducing a Tubular Model

Returning to Frangi's vesselness measure (Eq. 2.2), a subtlety arises when a multi-scale analysis is employed. The difficulty is that one has to compare the results of the response function at different scales, while the intensity and its derivatives are decreasing functions of scale. Hence, each individual response function must be suitably normalized before the comparison can be done. Fortunately, this can be done quite efficiently by directly computing the entries which comprise the Hessian matrix by convolving the image, \mathcal{I} , with second order derivatives of Lindeberg's γ -parametrized normalized Gaussian kernels [Lindeberg (1998)]. This is a general heuristic principle stating that local maxima over scales of combinations of *γ -normalized derivatives*,

$$\partial_{x,\gamma-norm} = t^{\frac{\gamma}{2}} \partial_x$$

serve as useful indicators reflecting the spatial extent of corresponding image structures. In this expression, the scale parameter used is the evolving time, t , in the classical *Heat Equation*,

$$\mathcal{I}_t = \Delta \mathcal{I} = \text{div}(\nabla \mathcal{I})$$

This comes from the fact that convolving an image \mathcal{I} with a Gaussian kernel at scale σ is equivalent to evolving every point of \mathcal{I} according to the Heat equation




FIGURE 3.1: A mid-sagittal slice of the original PD weighted MRI volume with its corresponding vesselness map. Note how the vesselness measure is now an angiographic-like volume similar to the PC data set of 1.1(c).

for $t = \sigma^2$ iterations [Hummel (1986)]. Thus,

$$\begin{aligned} \partial_{x,\gamma-norm} &= t^{\frac{\gamma}{2}} \partial_x \\ &= (\sigma^2)^{\frac{\gamma}{2}} \partial_x \\ &= \sigma^\gamma \partial_x. \end{aligned}$$

It is standard to choose $\gamma = 1$ when gamma is a fixed value. Otherwise, it is adaptively changed to maximize some quantity as in Lindeberg (1998) and Krissian et al. (2000) for automatic detection of edges or special features. Hence, anytime we compute a derivative, we scale its value by σ . We compute every entry of the Hessian matrix using

$$\mathcal{I}_{uv}(\bar{x}, \sigma) \equiv \sigma^2 \frac{\partial^2 G(\bar{x}, \sigma)}{\partial u \partial v} * \mathcal{I}(\bar{x}) \quad \text{where} \quad G(\bar{x}, \sigma) = \frac{1}{\sqrt{(2\pi\sigma^2)}} \exp \frac{-\langle \bar{x}, \bar{x} \rangle}{2\sigma^2}$$

In our implementation of the vesselness measure, we set the parameters α, β and c to 0.5, 0.5 and half the maximum Frobenius norm respectively, as suggested in Frangi et al. (1998). In practice we have found these parameter settings to yield stable results over a wide range of image modalities. Figure 3.1 shows a mid-sagittal slice of the original PD weighted MRI volume with its corresponding ves-

selness map. Note how the vesselness measure is an “angiographic-like” volume with strong intensity in vessel regions and almost null values outside vasculature, similar to the phase contrast angiography of Figure 1.1(c). At each voxel we compute vesselness responses using ten log scale increments between $\sigma = 0.2$ and $\sigma = 2.5$ (in our data the maximum radius of a vessel is 2.5 voxels) and select the maximum vesselness response along with its scale. The chosen scale gives the estimated radius of the vessel and the Eigen vector associated with the smallest Eigen value its local orientation.

This process is illustrated in Figure 3.2 for a synthetic branching structure, a synthetic helix and cropped portions of a PD weighted MRI and MRA. The gray surface coincides with a particular level set of the vesselness measure, which quickly drops to zero away from centerline locations. Within this surface locations of high vesselness are indicated by overlaying the Eigen vectors associated with the lowest Eigen values, which correspond to the estimated vessel orientation. Note that some vessel regions do not have vectors overlaid. This is because they have a low vesselness value at these locations. This is evident at the branch point of Figure 3.2(a) and the high curvature part of Figure 3.2(b), where at these voxels, the local shape is more blob-like. However, neighboring points have high vesselness measures with correct associated vessel orientations. Furthermore, it is apparent that locations of high vesselness are close to the expected centerlines, and that the estimated vessel orientation at these locations is accurate. This information along with the estimated radius of associated vessels can be used to construct an appropriate vector field to drive the flux maximizing geometric flow, as we shall now see. The use of a geometric flow is a simple and natural way to handle branching structures while integrating the vesselness information with some amount of local control. This allows us to lift many of the restrictions on the flow pointed out in Section 2.2.3, because an explicit model of a tubular structure is now incorporated along with an appropriate notion of scale.

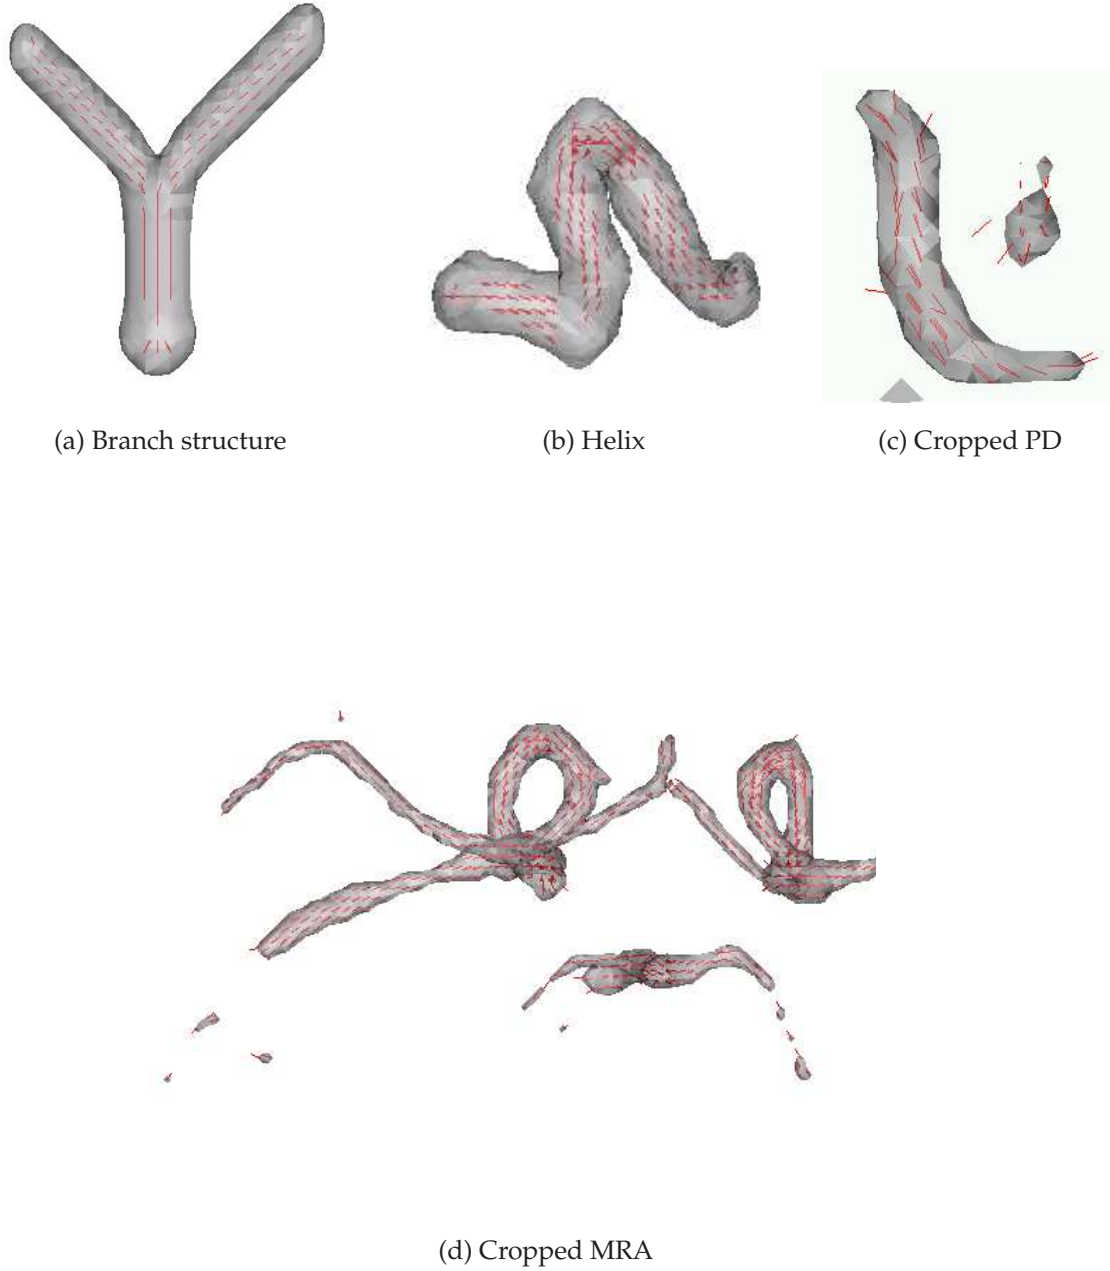


FIGURE 3.2: A synthetic branching structure, a synthetic helix, a cropped region of an MRA and a cropped vessel of a PD weighted MRI. For each structure the red vectors indicate the estimated vessel orientation at locations where the multi-scale vesselness measure (Eq. 2.2) is high. Note that at the branch point of 3.2(a) and the high curvature part of 3.2(b), there are no vectors overlaid. At these locations the local shape is blob-like and hence the vesselness measure is low.

3.2 Extending the Vesselness Measure to Boundaries

As described in Section 2.2.3, in order for an initial surface to evolve so as to align itself to vessel boundaries using the flux maximizing flow equation of 2.3, we need to construct a vector field large in magnitude and orthogonal to vessel contours. Since the vesselness measure is concentrated at centerlines, we need to distribute it to the vessel boundaries which are implied by the local orientation and scale. We consider an ellipsoid with its major axis aligned with the estimated orientation and its two semi-minor axes equal to the estimated radius. In our implementation the semi-major axis length is chosen to be twice that of the semi-minor axes. The vesselness measure is then distributed over every voxel (x_e, y_e, z_e) on the boundary of the ellipsoid by scaling it by the projection of the vector from (x, y, z) to (x_e, y_e, z_e) onto the cross-sectional plane passing through the semi-minor axes, as illustrated in Figure 3.3(a) and Figure 3.3(b). If (x, y, z) is taken to be the origin $(0, 0, 0)$ and the xy plane is taken to coincide with the cross-sectional plane this scale factor works out to be

$$\left\langle (x_e, y_e, z_e), \frac{(x_e, y_e, 0)}{\sqrt{x_e^2 + y_e^2}} \right\rangle = \sqrt{x_e^2 + y_e^2}. \quad (3.1)$$

This distribution cannot be blindly done at all voxels in the vesselness volume because we are only confident in scale and orientation estimates for voxels on vessel centerlines. For instance, it is often the case that tube-like points off the centerline get a significant vesselness measure with an associated scale that is incorrect because the Gaussian kernel responds to the opposite boundary, which is further away. Hence, the estimated scale does not reflect the true vessel width. To resolve this subtlety, we must distribute the measure only from voxels at centerline locations. To find such locations, we adopt a very simple local maximum detection procedure. At each voxel (x, y, z) where the vesselness measure is a local maximum in a $3 \times 3 \times 3$ neighborhood, we perform the vesselness distribution over all voxels on the surface of the associated ellipsoid. This construction is explained in more detail

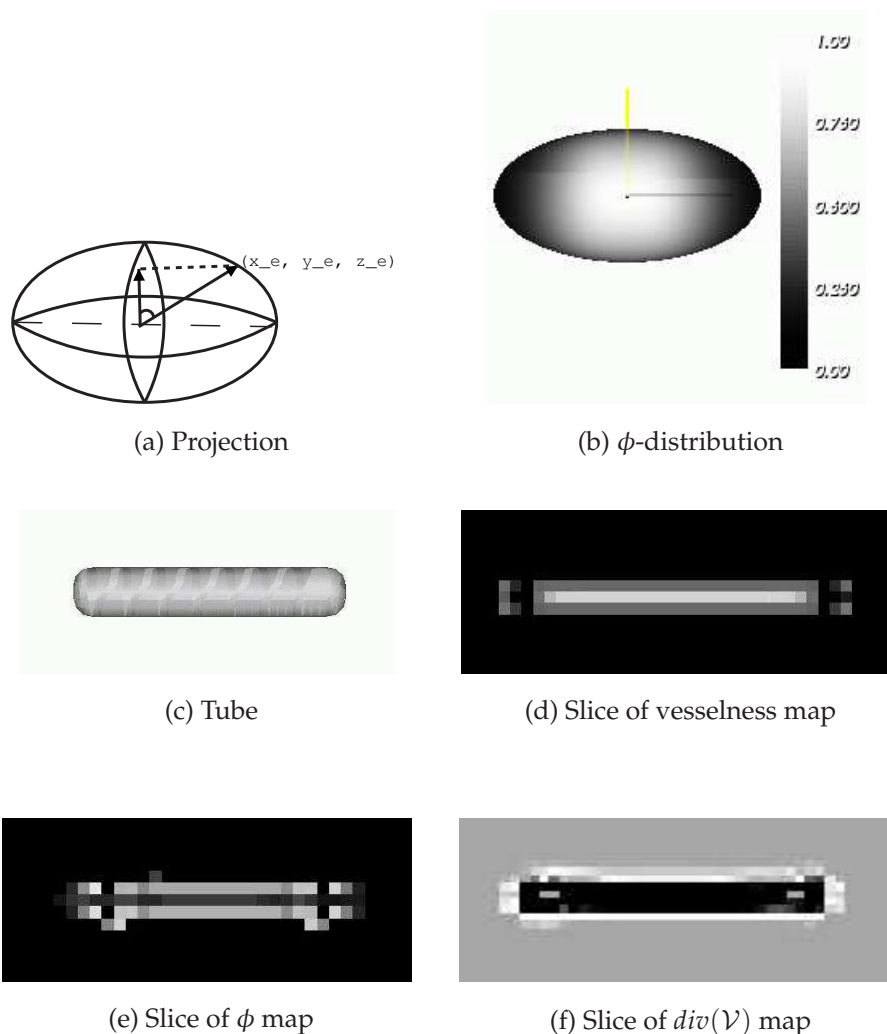


FIGURE 3.3: Distributing the vesselness measure to the implied boundaries. 3.3(a) The vector from the center of the ellipsoid to the surface voxel (x_e, y_e, z_e) , as well as its projection onto the cross-sectional plane, taken to be the xy plane. 3.3(b) We distribute the vesselness measure to all (x_e, y_e, z_e) on the ellipsoid by scaling it by the magnitude of this projection. The color bar indicates the association between brightness and magnitude. 3.3(c) A synthetic tube of radius 2. 3.3(d) A view of the vesselness measure in a slice, with brighter regions indicate stronger intensity. 3.3(e) A view of the ϕ distribution in the same slice. 3.3(f) The divergence of the vector field in Eq. 3.2, with transitions between dark and bright indicating zero-crossings. As expected, we have local maxima of the vesselness measure on the centerline in 3.3(d), local maxima of the ϕ distribution at the boundaries of the tube in 3.3(e) and zero-crossings of the divergence at the boundaries of the tube in 3.3(f).

in Section 3.4. This process of distributing the vesselness measure to the implied boundaries clearly favors voxels in the cross-sectional plane and gradually fades to the ends of the ellipsoid. This is illustrated in 3.3(b), where the surface of the ellipsoid is colored according to the projection value. We define the addition of the extensions carried out independently at all voxels to be the ϕ distribution.

The extended vector field is now defined as the product of the normalized gradient of the original image with the above ϕ distribution

$$\vec{\mathcal{V}} = \phi \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|}. \quad (3.2)$$

This vector field embodies two important constraints. First, the magnitude of ϕ is maximum on vessel boundaries and the ellipsoidal extension performs a type of local integration. This follows because the local maximum vesselness criterion enforces the condition that the extension is carried out only from locations as close as possible to vessel centerlines. Hence, the maximum value previously on the centerline is translated to the vessel contour. This is demonstrated in Figure 3.3 on a synthetic tubular structure. The vesselness map is maximum along the centerline (3.3(d)) and the ϕ map has maxima distributed to vessel boundaries (3.3(e)). Second, $\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|}$ captures the direction of the gradient, which is expected to be high at boundaries of vessels as well as orthogonal to them. It is important to normalize the gradient of the image so that its magnitude does not dominate the measure in regions of very low vesselness. For example, structures such as white and gray matter boundaries could then get significant unwanted contributions. Figure 3.3(f) shows the divergence of this new vector field of Eq. 3.2.

We have performed a careful numerical validation of the ϕ distribution procedure on synthetic tubes of varying central axis curvature and radius. The vesselness measure, ϕ extension and divergence map were computed as previously explained (illustrated in the example of Figure 3.3). We then found the average and maximum distance error between ground truth surface points and corresponding

extension surface points. The extension surface points are the zero-crossings in the divergence map, which are computed with a simple linear interpolation. We obtained an average distance error of 0.35 voxels and a maximum error of approximately 1 voxel over all the examples. This shows the accuracy of the vessel boundary estimations using the proposed extension.

3.3 The Multi-Scale Geometric Flow

The extended vector field explicitly models the scale at which vessel boundaries occur, due to the multi-scale nature of the vesselness measure $V(\sigma)$ (Eq. 2.2) as well as the expected gradient in the direction normal to vessel boundaries. Thus it is an ideal candidate for the static vector field in the flux maximizing geometric flow (Eq. 2.3). The surface evolution equation then works out to be

$$\begin{aligned} S_t &= \operatorname{div}(\vec{\mathcal{V}}) \vec{\mathcal{N}} \\ &= \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \operatorname{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right) \right] \vec{\mathcal{N}} \\ &= \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \kappa_{\mathcal{I}} \right] \vec{\mathcal{N}}. \end{aligned} \quad (3.3)$$

Here $\kappa_{\mathcal{I}}$ is the Euclidean mean curvature of the iso-intensity level set of the image. Note that this is a hyperbolic partial differential equation since all terms depend solely on the vector field and not on the evolving surface. We now enumerate several properties of this geometric flow.

1. The first term $\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle$ acts like a doublet. To see this, we observe the intensity profiles of $I, \nabla I, \phi$ and $\nabla \phi$ in Figure 3.4. ϕ has a maximum at vessel boundaries which implies that $\nabla \phi$ has a zero-crossing at such locations. Furthermore, I behaves like a smoothed step function at vessel contours which implies that $\nabla \mathcal{I}$ does not change sign there. Therefore, the first term of the evolution Eq. 3.3 is a doublet. Such doublet terms have also shown to be beneficial in earlier geometric flows for segmentation [Kichenassamy et al.

(1995); Caselles et al. (1995); Siddiqi et al. (1998)]. When the evolving surface overshoots the boundary slightly, this term acts to push it back toward the boundary.




FIGURE 3.4: An approximate sketch of the $I, \nabla I, \phi$, and $\nabla \phi$ intensity profiles.

2. The second term is a regularization term since it behaves like the *geometric heat equation*. Here,

$$\kappa_I = \text{div} \left(\frac{\nabla I}{|\nabla I|} \right)$$

is the mean curvature of the iso-intensity level set of the original intensity image. Evolutions driven by such term have been extensively studied in the mathematics literature and have been shown to have remarkable anisotropic

smoothing properties [Gage and Hamilton (1986);Grayson (1987)]. Such terms are also the basis for several nonlinear geometric scale-spaces such as those studied in Alvarez et al. (1992b,a) and Kimia et al. (1990, 1995).

3. Combining both terms, it is clear that the flow cannot leak in regions outside vessels since both ϕ and $\nabla\phi$ are zero there. Hence, when seeds are placed at locations where the vesselness measure $V(\sigma)$ is high the flow given by Eq. 3.3 will evolve toward the closest zero level set of the divergence of the vector field \vec{V} . This will make the evolving surface cling to vessel boundaries.

3.4 Algorithms and Implementation Details

The entire process for extracting vasculature can now be described via three algorithms. First, the vesselness measure is computed using Algorithm 3.1. Second, this measure is used to construct the extended vector field via Algorithm 3.2. Finally, this extended vector field drives the flux maximizing geometric flow for segmentation described in Algorithm 3.3.

Below we review some of the details of the implementation of these algorithms:

1. Typically, a few iterations of mean curvature type smoothing on the original

	At vessel boundaries	inside vessels	outside vessels
I	contrast change (bright to dark for PD and dark to bright for angiography)	roughly constant (step function or Gaussian-like)	large for PD and small for angiography
∇I	local min at one boundary and local max at the other	zero or small	small except at tissue change (ex: between gray and white matter)
ϕ	local max	camel back	zero
$\nabla\phi$	zero-crossing	positive and negative	zero

TABLE 3.1: Behavior of $I, \nabla I, \phi, \nabla\phi$ intensity profiles at vessel contours and inside and outside vessel regions.

Algorithm 3.1: Vesselness computation

Data : \mathcal{I} : 3D input medical data set

Result : Vess: maximum vesselness measure volume

Scale: smallest Eigen value λ_1

E_x : volume containing first component of Eigen vector e_1

E_y : volume containing second component of Eigen vector e_1

E_z : volume containing third component of Eigen vector e_1

for $\sigma = \sigma_{min}$ *to* σ_{max} **do**

for (every voxel $\tilde{x} \in \mathcal{I}$) **do**

 Compute the derivatives of the Hessian matrix as described in Section 3.1;

 Use Jacobi's method to extract the Eigen values $\lambda_1, \lambda_2, \lambda_3$ and the associated Eigen vectors e_1, e_2, e_3 ;

 Sort them such that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$;

 Compute R_a, R_b, S and the corresponding vesselness value $V(\sigma)$ of Eq. 2.2;

 /* keeping maximal response */

if ($V(\sigma) > Vess(\tilde{x})$) **then**

 Vess(\tilde{x}) = $V(\sigma)$;

 Scale(\tilde{x}) = λ_1 ;

$E_x = e_1(x)$;

$E_y = e_1(y)$;

$E_z = e_1(z)$;

image is used as a pre-processing step before segmentation. This is a standard method to remove artifacts such as speckle noise since it smooths along iso-intensity level sets but not across them. However, we have noticed that this process is unnecessary in our implementation as we compute derivatives of the Hessian matrix by convolution with derivatives of Gaussian kernels, which takes care of preliminary smoothing. If a curvature flow is used on top of that, we loose many smaller vessels.

2. In order to favor smaller scales, we use log scale increments when computing derivative entries of the Hessian operator. We then select the maximum

Algorithm 3.2: Construction of the extended vector field

Data : \mathcal{I} : 3D input medical data set
Vess: maximum vesselness volume
Scale: smallest Eigen value λ_1
 E_x : volume containing first component of Eigen vector e_1
 E_y : volume containing second component of Eigen vector e_1
 E_z : volume containing third component of Eigen vector e_1
 ϕ : ϕ -extension volume

Result : \mathcal{F} : speed volume driving the flux maximizing flow, $div(\vec{\mathcal{V}})$

Compute Vess, Scale, E_x , E_y , E_z with Algorithm 3.1;
for (every voxel $\bar{x} \in \mathcal{I}$) **do**
| Initialize $\phi(\bar{x}) = 0$;
for (every voxel $\bar{x} \in \mathcal{I}$) **do**
| /* vesselness extension to vessel boundaries */
| Compute *local_max* variable by finding local maximum of Vess volume
| in a 3x3x3 neighborhood of \bar{x} ;
| **if** ($Vess(\bar{x}) > threshold$) and $\left(\frac{Vess(\bar{x})}{local_max} > percentile\right)$ **then**
| | **for** each \bar{x}_e on the ellipsoid surface of semi-minor length $Scale(\bar{x})$, semi-
| | major length $2*Scale(\bar{x})$ and orientation given by $(E_x(\bar{x}), E_y(\bar{x}), E_z(\bar{x}))$
| | **do**
| | | $\phi(\bar{x}_e) = \text{Distribute}(Vess(\bar{x}))$ as detailed in Section 3.2;
for (every voxel $\bar{x} \in \mathcal{I}$) **do**
| Compute $\frac{\nabla I(\bar{x})}{|\nabla I(\bar{x})|}$;
| Compute $\kappa_{\mathcal{I}}(\bar{x})$ given by Eq. 3.4;
| Compute $\nabla \phi$;
| $\mathcal{F}(\bar{x}) = \left\langle \nabla \phi(\bar{x}), \frac{\nabla I}{|\nabla I|}(\bar{x}) \right\rangle + \kappa_{\mathcal{I}}\phi(\bar{x})$;
return \mathcal{F} ;

vesselness response as described in Section 2.1. We use Jacobi's method for symmetric matrices to find the Eigen values of the Hessian. For a faster multi-scale vesselness volume computation, we have computed this measure over 5 scales without noticeable differences in the vessel extractions.

3. The ϕ distribution in Section 3.2 is carried out from voxels at vessel center-lines since at such locations one has strong confidence in the scale and ori-

Algorithm 3.3: Level set based geometric flow

Data : \mathcal{F} : speed volume driving the flux maximizing geometric flow
 Ψ : volume with the evolving surface embedded as its zero level set
 N : data structure containing points in the narrow band

Result : S : surface representation of the vasculature extracted from \mathcal{I}

Compute vesselness measure and resulting volumes with Algorithm 3.1;

Compute speed function \mathcal{F} using Algorithm 3.2;

/ Surface initialization */*

for (*every voxel* $\bar{x} \in \mathcal{F}$) **do**

if ($Vess(\bar{x}) > initial_threshold$) **then**

$S(\bar{x}) = 1$;

else

$S(\bar{x}) = 0$;

$\Psi = dt(S)$ (dt is the signed Euclidean Distance Transform of Borgefors (1984);

Compute points in the narrow band to the surface and store them in N ;

/ Level set surface evolution equation */*

for ($t = 0$ to $stop_time$) **do**

for (*every voxel* $\bar{x} \in \mathcal{F}$) **do**

if ($\bar{x} \in N$) **then**

 Update Ψ according to the discrete surface evolution Eq. 3.5;

if *narrow band boundary* N *is hit* **then**

for *every* $\bar{x} \in \mathcal{I}$ **do**

if ($\Psi = 0$) **then**

$S(\bar{x}) = 1$;

else

$S(\bar{x}) = 0$;

$\Psi = dt(S)$;

return $S = \Psi(0)$;

entation estimate from Frangi's vesselness measure [Frangi et al. (1998)]. A global thresholding approach is not appropriate as it either misses the smaller structures or allows a lot of non-vessel structures. Hence, we adopt a more

local procedure which is sensitive to both small and large vessels:

$$\text{if } (V(\sigma) > \text{threshold} \ \&\& \ \frac{V(\sigma)}{\text{local_max}} > \text{percentile})$$

Distribute vesselness over ellipsoid

The *threshold* condition is to ensure that we only consider voxels with a significant vesselness measure. The variable *local_max* is the maximum vesselness response in a small neighborhood of a particular voxel. We chose a 3x3x3 neighborhood because we know that if a point has a significant vesselness value, there must be a vessel centerline within at least 3 voxels (the maximum vessel radius is 2.5 voxels in our data). If $\frac{V(\sigma)}{\text{local_max}} > \text{percentile}$, then we have detected a vessel voxel on or very near the center of the vessel. Otherwise, we are either off the centerline or at part of another local structure. For most examples we use a conservative vesselness *threshold* of 0.01 and a *percentile* of 0.75. These parameters give good and stable vessel extractions over all the image modalities tested. In practice, one can choose the *threshold* more aggressively for angiographic data as the difference between vessel and non-vessel regions is then much sharper.

4. The derivatives in the doublet term $\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle$ are computed using central differences for $\nabla \phi$ and a second-order essentially non-oscillatory (ENO) scheme for the normalized gradient of the input image, $\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|}$ [Osher and Shu (1991)]. We choose a central difference scheme when we want a smoother approximation of the derivatives and an ENO scheme for a more precise approximation able to capture sharp changes in intensity. ENO is also computationally more expensive.
5. We have two options to compute this quantity. First, we can use numerical approximations to first compute $\nabla \mathcal{I}$ and obtain a new volume $\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} = A$. Then, we can compute the divergence of this new data set with another

derivative approximation, i.e., $\text{div}(A) = A_x + A_y + A_z$. This approach is less appealing because we need at least three data structures to save the partial volumes and also, it uses numerical approximations at two levels. Our second option is to use the analytic level set expression for the mean curvature of an iso-intensity level set [Osher and Sethian (1988)]. We compute all derivatives using a 3-neighbor central difference scheme,

$$\kappa_{\mathcal{I}} = \frac{\begin{pmatrix} (\mathcal{I}_{yy} + \mathcal{I}_{zz})\mathcal{I}_x^2 + (\mathcal{I}_{xx} + \mathcal{I}_{zz})\mathcal{I}_y^2 + (\mathcal{I}_{xx} + \mathcal{I}_{yy})\mathcal{I}_z^2 \\ -2(\mathcal{I}_x\mathcal{I}_y\mathcal{I}_{xy} - \mathcal{I}_x\mathcal{I}_z\mathcal{I}_{xz} - \mathcal{I}_y\mathcal{I}_z\mathcal{I}_{yz}) \end{pmatrix}}{(\mathcal{I}_x^2 + \mathcal{I}_y^2 + \mathcal{I}_z^2)^{\frac{3}{2}}} \quad (3.4)$$

6. A first-order in time discretized form of the level-set version of the evolution equation is given by

$$\Psi_n = \Psi_{n-1} + \Delta t * \mathcal{F} * \|\nabla \Psi_{n-1}\| \quad (3.5)$$

where $\mathcal{F} = \left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \text{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right)$, Ψ is the embedding hypersurface and Δt is the step size. The evolving surface S is obtained as the zero level set of this Ψ function. The numerical derivatives used to estimate $\|\nabla \Psi\|$ must be computed with up-winding in the proper direction as described in Osher and Sethian (1988). This is now a standard numerical approach for solving partial differential equations of this type since it allows topological changes to occur without any additional computational complexity and can be made efficient using a narrow band implementation. It could be made even more efficient by using a second order in time discretization of the surface evolution equation, since the time step Δt could then be reduced.

7. The narrow band width has an underlying subtlety. There is a trade-off between memory and speed. The smaller the narrow band, the less voxels we have to update at every iteration of the evolution equation. However, when

the narrow band is hit by the evolving surface, we need to reinitialize the binary surface and recompute a new embedding distance function, Ψ . This is computationally expensive and very slow for large data sets because we use the *Euclidean Distance Transform* of Borgefors (1984) to implement Ψ . This uses four *float* measures for every voxel in the data set. Memory allocation quickly becomes a problem when segmenting volumes in the order of 300x300x300 voxels. Hence, we do not want to be hitting the narrow band too often. In our implementation, we have found that a narrow band width of 20 voxels is an effective and computationally efficient choice.

8. Flow algorithms are always challenged by the *initialization* step. Depending on the way the algorithm is used, one can initialize the flow manually or automatically. In this work, we have focused on segmenting as much vasculature as possible automatically. As mentioned in Chapter 2, most existing flows are applied on angiographic data and can be initialized by thresholding the original data set. In the more general case of PD or Gadolinium enhanced MRI, we use the vesselness volume to initialize the surface. We threshold it using 0.1 for standard MRI data sets and use a more aggressive threshold of 0.05 when segmenting angiography data sets. These values give good initial surfaces capturing most of the important vessels. This allows the flow to converge fast to the final segmentation without the need of a constant inflation term to speed up the evolution as necessary in the implementation of Malladi et al. (1993, 1994, 1995) and Caselles et al. (1993). Optimally, we believe a semi-automatic algorithm gives the best results. A user would typically segment automatically as much vasculature as possible in a first step. Then, regions of interest could be selected and seeds could be placed manually to further segment smaller or lower contrast vasculature.
9. The stopping criteria is specified by the user. However, if the narrow band has not been hit in a very long time (5000 iterations), the process is stopped

automatically. In most examples, depending on the initialization used, 15000 iterations is enough to extract most of the vasculature. The algorithm regularly saves intermediate surfaces during the evolution so that a user can restart the segmentation process from a previously saved iteration. In practice, when the original data set is in the order of $300 \times 300 \times 300$, we can compute 10000 iterations per hour on a Pentium IV, 1.5Ghz, 1G RAM machine. The initial computation of the vesselness measure and the vector field needed to drive the flow can be computed in roughly 15 minutes.

Chapter 4

Validation

We now validate our multi-scale geometric flow for extracting vasculature. We first present qualitative segmentation results and masked maximum intensity projections (MIPs) on a variety of modalities. We then carry out a quantitative comparison of the segmentations on a data set consisting of Proton Density (PD) weighted MRI, Time Of Flight (TOF) angiography and Phase Contrast (PC) angiography volumes, all obtained for the same subject.

4.1 Image Acquisition

We acquired four different volumes from the same subject (the author) on a Siemens 1.5 Tesla system at the Montreal Neurological Institute (MNI). We first used a PD/T2-weighted dual turbo spin-echo acquisition with sagittal excitation (2mm thick slices, 50% overlap 1mm³ isotropic voxels, TE = 0.015s TR = 3.3s). Following this, a 3D axial phase-contrast (PC) volume (0.47mm x 0.47mm x 1.5mm resolution, TE = 0.0082s TR = 0.071s) and a 3D axial time-of-flight (TOF) volume (0.43mm x 0.43mm x 1.2 mm resolution, TE = 0.0069s TR = 0.042s) were acquired. Each data set was registered to a standardized coordinate system and re-sampled onto a 0.5mm³ isotropic voxel grid to facilitate processing and comparisons. A mid-sagittal slice of the PD, PC and TOF volumes is depicted in Figure 1.1. We supplemented these three data sets with an MRA volume (Figure 4.1) and a Gadolinium enhanced MRI volume (Figure 4.3), both obtained from the MNI.

In the PC data, contrast is determined by tissue motion. Static tissue yields no signal, and is therefore black. In the TOF data, vessel brightness is proportional to blood flow velocity. However complex flow or turbulence can cause some signal loss in the vessels in such images. In the data presented here, vessel/non-vessel contrast is greatest for the PC data (white on black tissue), intermediate for the PD data (black on gray) and slightly less for the TOF (white on gray). Resolution also affects vessel detectability. In principle the angiographic volumes should be able to show smaller vessels, since they have a higher resolution.

4.2 Qualitative Results

We illustrate our multi-scale geometric flow for segmenting vasculature on a variety of modalities. The same parameters were used throughout, as described in Section 3.4. We should point out that whereas the prior geometric flow based methods of Lorigo et al. (2001) and Vasilevskiy and Siddiqi (2002) could be applied to the angiographic volumes, they would fail entirely on both the Gadolinium enhanced MRI volume and the PD data set. This is because high contrast regions are not limited to vessel boundaries and these techniques do not have an explicit tubular model. Hence, these flows would leak in the gray matter and other non-vessel regions.

Figure 4.1 shows iterations of the flow using three single voxel seeds on an MRA data set obtained from the MNI, as well as an MIP of the data set masked by the final segmentation. In preliminary work we demonstrated that the flow is able to pick up the main vessels automatically when the original 1mm^3 isotropic data is used [Descoteaux et al. (2004a)]. In the current experiment the original data is super-sampled to a 0.5mm^3 resolution. This preprocessing strategy allows us to recover several of the finer vessels which are less than one voxel wide and have low contrast at their boundaries. This is illustrated in Figure 4.2.

Figure 4.3 depicts a $40\text{mm} \times 53\text{mm} \times 91\text{mm}$ region centered on the corpus callosum from a Gadolinium enhanced MRI volume obtained at the MNI. The 1mm^3 isotropic data was super-sampled to a resolution of 0.33mm^3 using a tricubic interpolation kernel, because several vessels in the original data set were less than one voxel wide. In the image one can see the callosal and supra-callosal arteries (the long arching vessels running from left to right). We show an MIP of a sagittal and a transverse view in the left column. A segmentation obtained by thresholding is shown in the middle column. This results in many disconnected vessels as well as artifacts. Our segmentation is shown in the third column and results in the reconstruction of well connected tubular structures. Observe how the local ellipsoidal

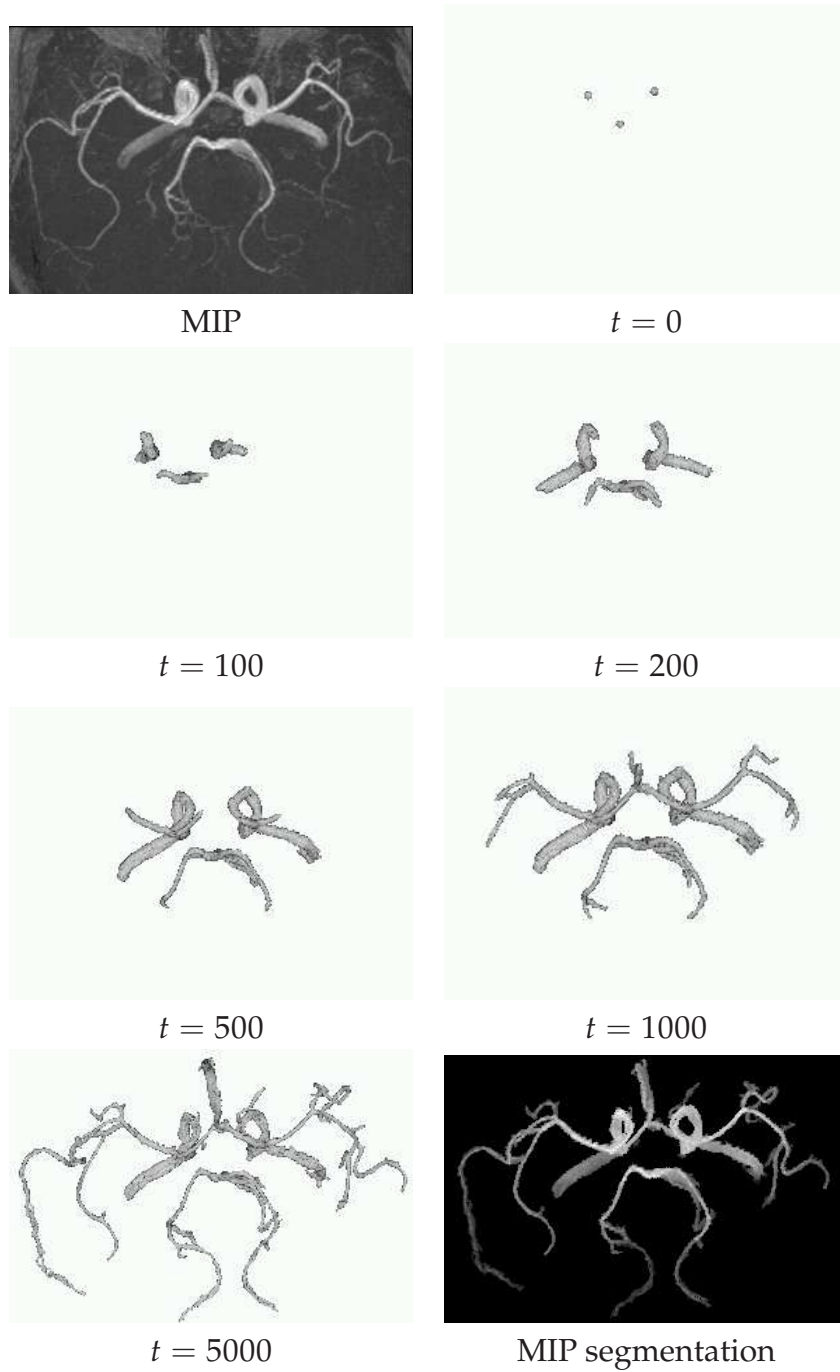


FIGURE 4.1: An illustration of the multi-scale geometric flow on a $68 \times 256 \times 256$ MRA image. An MIP of the data is shown at the top left and the other images depict different stages of the evolution from three seeds. The bottom right figure depicts an MIP of the input MRA data masked by the binary segmentation.




FIGURE 4.2: An illustration of the automatic blood vessel segmentation when re-sampling the original data set. Results of column (a) were presented in Descoteaux et al. (2004a) where we worked with the original 1mm³ resolution data. In column (b), we demonstrate the benefits of re-sampling the original data to a 0.5mm³ grid. Although it is computationally more expensive and requires more computer memory, re-sampling allows one to recover smaller vessels automatically.

integration scheme is able to connect a section of the supra-callosal arteries which has very low contrast in the original Gadolinium data set. Other methods that do not have an explicit tubular constraint fail miserably on this modality source as they leak into regions where the the Gadolinium contrast agent is absorbed by non-vessel tissues.

Finally, Figure 4.4 depicts the transverse views of intensity projections of the input data, the vesselness measures and the segmentations of the PC angiography, TOF angiography and PD volumes shown in Figure 1.1. Owing to the large number of short vessels near the surface of the full brain, the 2D visualization of the 3D segmentations poses a challenge since most of the vasculature inside the head is occluded when projecting the data in a certain direction. Hence, we choose to work with a common $259 \times 217 \times 170$ voxel region cropped from the center of each volume, which has vessels of different widths and contrasts in the three modali-




FIGURE 4.3: An illustration of the flow on a 40 mm x 53 mm x 91 mm cropped region of a Gadolinium enhanced MRI. An MIP of the sagittal and transverse views of the data is shown in (1) and (2). Reconstructions obtained by simple thresholding for the same views are shown in (3) and (4). These are clearly sensitive to noise and result in disconnected or missing vessels. The results obtained by our multi-scale geometric flow are shown in (5) and (6). Observe that the flow has connected a section of the callosal arteries which is barely visible in the MIP (see (1),(3),(5)).

ties. In the third column we mask the original volumes with the corresponding binary segmentations obtained by our algorithm, and show a maximum intensity projection (rows 1 and 2) or a minimum intensity projection (row 3). This last result is shown in “reversed” contrast so that it is comparable to the other two. Observe that along each row, the segmentations, vesselness maps and maximum/minimum




FIGURE 4.4: Transverse views of intensity projections (IP) of the PC, TOF (maximum IP) and PD data sets (minimum IP), the associated vesselness measures and the segmentations obtained by the multi-scale geometric flow. Observe that along each row, the segmentations, vesselness maps and maximum/minimum intensity projections agree closely.

intensity projections agree closely, up to some very small vessels. We also note the resemblance between the PC and PD views, where a majority of the vasculature agrees. We carry out a quantitative study of these segmentation results in

the following section. To our knowledge, this is the first segmentation in the literature of a PD weighted MRI obtained using a geometric flow. Movies of the geometric flow on the PC and PD data sets can be found on the author's web page, <http://www.cim.mcgill.ca/~mdesco>.

4.3 Quantitative Results

Figure 4.5 compares the segmentations obtained on the PC, TOF and PD volumes (Figure 4.4) with transverse views in the left column and sagittal views in the right column. To allow for small alignment errors due to geometric distortions between the different acquisitions, we consider two locations to be in common if the Euclidean distance between them is no greater than 3 voxels (1.5 mm). In each figure red labels indicate locations common to the two data sets, green labels indicate locations present in the ground truth data set but not in the test data set and blue labels locations in the test data set which are not in the ground truth data set. It is clear from the first row that most of reconstructed vessels in the PD and PC data agree. The PC reconstruction has some finer vessels apparent in the transverse view where small collaterals branch off the posterior aspects of the middle cerebral artery in the lateral fissure. On the other hand, the PD reconstruction has more vasculature visible in the sagittal view with vessels branching off the callosal and supra-callosal arteries. Finally, the second and third rows of Figure 4.5 indicate that the TOF reconstruction is missing a large number of vessel labels when compared to the PC and PD reconstructions.

We now present a quantitative analysis of these segmentation results, which were presented in preliminary form in Descoteaux et al. (2004b). We compute a number of statistics between each pair of modalities, treating one as the "ground truth" data set and the other as the "test" data set. These comparisons are shown in Table 4.1 and include the following measures:

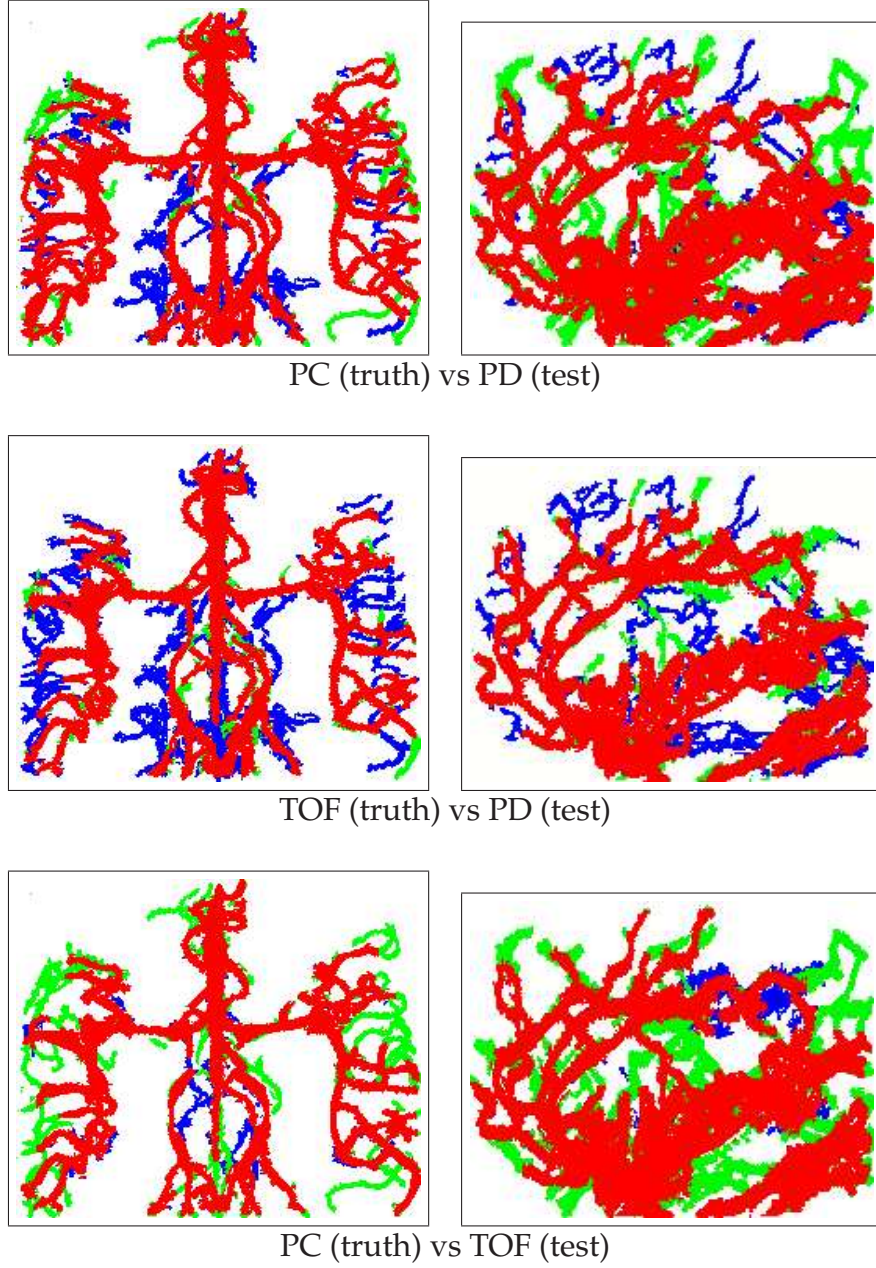


FIGURE 4.5: We consider the angiograms as the “ground truth”. Each row shows a pairwise comparison of reconstructions obtained on different modalities, with transverse views in the left column and sagittal views in the right column. White labels correspond to the background, red labels to locations common to the ground truth and test data, green labels to locations in the ground truth only and blue labels to locations in the test data only.

1. The kappa coefficient defined by

$$\frac{2a}{2a + b + c}$$

Data Sets		Validation Measures			
Ground Truth	Test Data	kappa	ratio	alignment (voxels) (mm)	
PC	PD	0.84	0.80	0.95	0.48
TOF	PD	0.81	0.89	0.66	0.33
PD	PC	0.84	0.89	0.56	0.28
PD	TOF	0.81	0.74	0.60	0.30
PC	TOF	0.81	0.72	0.82	0.41
TOF	PC	0.81	0.94	0.88	0.44

TABLE 4.1: A pair-wise comparison between the different modalities, treating one as the ground truth and the other as the test data.

where a is the number of red voxels, b is the number of green voxels and c the number of blue voxels. This measure tests the degree to which the agreement exceeds chance levels [Dice (1945)]. This measure is commonly used in the medical image analysis community. A kappa coefficient above 60% to 70% is considered as a strong correlations.

2. The ratio

$$\frac{a}{a + b}$$

where a and b are as before. This measure indicates the degree to which the ground truth data is accounted for by the test data.

3. The alignment error, defined by taking the average of the Euclidean distance between each voxel in the ground truth data set and its closest voxel in the test data set. This is done by computing the Euclidean distance transform on the test data and then, at every vessel voxel in the ground truth, adding the corresponding distance value. Recall that this value is the closest Euclidean distance to a vessel structure in the test data. This measure also indicates the degree to which the test data explains the ground truth data, but in terms of an average distance error. In order to avoid measurement bias when an extracted vessel is longer in one segmentation when compared to another, we

do not include voxels whose closest distance is greater than 3 voxels (1.5mm). This is essentially the set of red voxels in Figure 4.5.

It is clear from Table 4.1 that the vasculature obtained from the PD volume accounts for 80% and 89% of that obtained from the PC and TOF angiographic sequences, respectively. Furthermore, whereas 89% of the PD vessel voxels are also found in the PC data, a significant proportion (26%) of PD vessel voxels are not seen in the TOF data. The results also indicate very high alignments between vessel labels in all pair-wise comparisons, which indicates that when segmented, vasculature extracted from the different data sets is indeed similar.

Chapter 5

Discussion and Conclusions

This last and final chapter addresses potential questions concerning presented in this thesis. We also summarize the main contributions and propose a number of directions for future work.

5.1 Discussion

5.1.1 Why use a Geometric Flow?

This question is important and needs to be addressed since one may wonder if the result obtained using our algorithm is almost equivalent to surface obtained by sweeping elliptical disks, whose radii and orientation are determined by multi-scale vesselness responses. This proposition is in fact a direct extension to vessel boundaries of Aylward and Bullitt’s ridge traversal using properties of the Hessian matrix. The theory behind this centerline approach and our method is essentially the same but the geometric flow framework has several important advantages. A flow acts as a local “glue”, i.e., when propagating the surface front it is able collect evidence from neighboring voxels to create a connected surface. Hence, branch points of vascular trees, which locally behave like blobs (Section 3.1, Figure 3.2), are handled naturally. Ridge traversal and centerline techniques need an explicit bifurcation model or a back-tracking method to capture vessel junctions. A flow also allows significantly more control. An expert can interact with the data and segmentation process by stopping it, manually placing seeds and restarting the evolution. This is possible because the flow can adapt to merging surfaces as well as changes in topology. In the end, if one seeks for the 3D centerlines of the blood vessel surfaces, one can use centerline extraction methods such as those in Bouix et al. (2004a,b).

Another question one might ask is why did we chose the flux maximizing flow of Vasilevskiy and Siddiqi (2002) over the flow of Lorigo et al. (2001)? Although the latter has very nice mathematical motivation and an underlying regularization

term involving the Hessian matrix, the flow must be forced to stop at vessel contours by adding a heuristic image gradient term in the surface evolution equation. This is less appealing than the flux maximizing gradient flow of Vasilevskiy and Siddiqi (2002), which is significantly easier to implement.

5.1.2 Quantitative Validation and Ground Truth Data

Although we have carried out a careful qualitative and quantitative cross validation of our method, this falls short of a *true* quantitative validation. This is because we do not have the ground truth segmentation to compare our PD segmentation to. Colleagues have suggested the use of a high quality CT acquisition of a phantom brain to obtain a ground truth 3D representation of the blood vessels. However, this is not useful in our analysis because it is impossible to obtain test data from a phantom brain with similar complexity as a PD weighted MRI of a human brain. Another possibility is to use the virtual brain simulator (<http://www.bic.mni.mcgill.ca/brainweb/>) [Collins et al. (1998)] to generate a virtual angiogram and its corresponding anatomical MRI. However, at this point, the tool can generate T1, T2, and PD MRI composed of only cerebral spinal fluid, gray and white matter tissues but not blood vessels. Hence, the cross validation performed in this thesis is currently the best type of validation we can perform. Our statistical measures suggest that most of the vascular structures in a high quality PC angiography data set can be extracted from a standard clinical PD weighted MRI. One might have doubts on the quality of the angiogram but in fact, the PC MRA used in this comparison was shown to an expert in image acquisition and brain analysis who was impressed by its quality and ability to show the vasculature.

5.2 Contributions and Summary

We have presented what to our knowledge is the first multi-scale geometric flow for segmenting vasculature in standard MRI volumes. Whereas the flow is designed for PD weighted data sets, it can also be applied to a variety of other modalities. We have demonstrated its applicability with both qualitative and quantitative cross validation studies. First, the qualitative results indicate that a significant amount of vasculature can be recovered by initializing the flow using a few isolated seeds. We have also found that a number of finer vessels can also be recovered by super-sampling the data and by placing seeds manually along with an adaptive lowering of the vesselness threshold used in the construction of the extended vector field \vec{V} (Eq. 3.2).

We have proposed a method to visualize vasculature by creating maximum or minimum intensity projections of the original data, but masked by the binary segmentations. These projections are particularly useful for visualizing vasculature in non-angiographic volumes since artifacts due to the brain surface as well as background structures are removed. These are quick to compute over any projection direction and neurosurgeons and radiologists are familiar with them. The results in Figure 4.4 show that the MIPs of the original PC data and the segmented PC data are very similar, indicating that our geometric flow is successful in segmenting all but the very finest vessels. The MIPs of the original TOF and the segmented TOF data are even more similar, although the TOF data contains fewer vessels when compared with the PC volume. Surprisingly, the minimum intensity projection of the PD data also shows a significant number of vessels. This information is greatly enhanced in the vesselness of PD image in the bottom row of Fig. 4.4. The reversed contrast MIP of the masked PD data demonstrates that our vessel segmentation procedure is successful and yields a 2D image which is comparable to the MIP of the segmented PC image and which is almost as informative as the MIP of the original PC. More importantly, the complex spatial relationships between

the vasculature and surrounding anatomical structures can be made explicit since the segmented PD is a true three-dimensional structure. A user can interact with the derived model, depending upon the task at hand, and can visualize it from arbitrary viewing directions.

Second, an important contribution of our work is the quantitative cross validation of the algorithm using a data set comprised of PD, PC and TOF volumes obtained for the same subject. The quantitative results indicate that the vessels segmented from the PD data alone account for over 80% of the vasculature segmented from either of the angiographic data sets, with a very small alignment error. We observe also that 26% of the vasculature obtained from the PD data are not recovered from the angiographic TOF volume. This suggests that our algorithm can be used to improve upon the results obtained from angiographic data but also as a promising alternative when such data is not available, since PD-weighted MRI data are routinely acquired when planning brain tumor surgery.

It is important to point out that all the segmentations were obtained automatically by initializing the flow with a threshold of the vesselness measure and by stopping the surface evolution after a fixed number of iterations, or when the flow had not hit the narrow band for several iterations. In the case of the PD volume, the threshold must be conservative to guarantee that seeds are placed only within vessel regions. It is possible to place seeds less conservatively in the angiographic volumes in which vessels can be identified primarily by contrast. Ideally the algorithm could be semi-automatic to improve the segmentation results. For example, in the event that the automatic reconstruction does not recover some of the finer vessels, these could be later obtained using a finer manual placement of seeds along with an adaptive lowering of the vesselness threshold at such locations.

Finally, it is important to note that the method does depend crucially on the choice of a particular vesselness measure to identify centerlines along with their orientations and associated vessel widths. Whereas our results indicate that Frangi's

vesselness measure is a very promising candidate, other choices have also been proposed in the literature [Aylward and Bullitt (2002); Krissian et al. (2003)] and these would be worth exploring in the context of driving a geometric flow. One issue that must be faced is the normalization of the responses for such operators so that both thin and thick vessels yield quantitatively similar values at expected centerline locations.

5.3 Future Work

It is our hope that our implementation will become a basic image analysis tool for segmenting vasculature in clinical studies. In fact, we have already started using it for *vessel driven brain shift correction* at the Montreal Neurological Institute [Reinertsen et al. (2004)]. Our segmentation algorithm is the basis for the registration of pre-operative MR images and intra-operative Doppler ultrasound data. The vascular tree present in the Gadolinium enhanced MRI is segmented with our algorithm and then the 3D centerline curves are found using the automatic centerline extraction proposed by Bouix et al. (2004a,b). These curves are used as landmarks for registering vessels from the intra-operative ultrasound. It is then possible to find a brain shift estimate.

An accurate segmentation of vasculature from brain MR images is also critical in many other clinical applications. Once segmented, various measures can be used to characterize the vascular tree, such as tortuosity, size and branching, with direct applications in the diagnosis, treatment and follow-up of arterial venous malformations and assessment tumor malignancy. Due to the automatic nature of our vessel segmentation algorithm, one could also analyze large databases of PD/T2 weighted MRIs of healthy subjects and patients with particular diseases.

Finally, there are many other ways one could exploit local shape properties in images. In this thesis, we have only discussed tube-like structures but one can easily define different measures to enhance and detect other structures such as *blobs*

or *sheets*. In particular, we have recently designed a “sheetness” measure to detect sheet-like structures in astrophysics galaxy simulation data using a geometric flow, with promising results.

Bibliography

- Alvarez, L., F. Guichard, P. L. Lions, and J. M. Morel: 1992a, 'Axiomatisation et nouveaux opérateurs de la morphologie mathématique'. *C. R. Acad. Sci. Paris* **315**, 265–268. [33]
- Alvarez, L., F. Guichard, P. L. Lions, and J. M. Morel: 1992b, 'Axiomes et équations fondamentales du traitement d'images'. *C. R. Acad. Sci. Paris* **315**, 135–138. [33]
- Ambrosio, L. and H. M. Soner: 1996, 'Level Set Approach to Mean Curvature Flow in Arbitrary Codimension'. *Journal of Differential Geometry* **43**, 693–737. [20]
- Aylward, S. R. and E. Bullitt: 2002, 'Initialization, Noise, Singularities, and Scale in Height Ridge Traversal for Tubular Object Centerline Extraction'. *IEEE Transactions On Medical Imaging* **21**(2), 61–75. [10, 17, 58]
- Borgefors, G.: 1984, 'Distance Transformations in Arbitrary Dimensions'. *Computer Vision, Graphics, and Image Processing (CVGIP)* **27**, 321–345. [36, 39]
- Bouix, S., A. Tannenbaum, and K. Siddiqi: 2004a, 'Flux Driven Automatic Centerline Extraction'. Technical report, School of Computer Science, McGill University. [54, 58]
- Bouix, S., A. Tannenbaum, and K. Siddiqi: to appear 2004b, 'Flux Driven Automatic Centerline Extraction'. *Medical Image Analysis*. [54, 58]
- Caselles, V., F. Catte, T. Coll, and F. Dibos: 1993, 'A Geometric Model for Active Contours in Image Processing'. *Numerische Mathematik* **66**, 1–31. [19, 39]
- Caselles, V., R. Kimmel, and G. Sapiro: 1995, 'Geodesic Active Contours'. In: *International Conference On Computer Vision*. pp. 694–699. [32]
- Collins, D., A. Zijdenbos, V. Kollokian, J. Sled, N. Kabani, C. Holmes, and A. Evans: 1998, 'Design and Construction of a Realistic Digital Brain Phantom'. *IEEE Transactions on Medical Imaging* **17**, 463–468. [55]
- Deschamps, T. and L. D. Cohen: 2001, 'Fast Extraction of Minimal Paths in 3D Images and Applications to Virtual Endoscopy'. *Medical Image Analysis* **5**(4), 281–299. [18]

- Descoteaux, M., L. Collins, and K. Siddiqi: 2004a, 'A Multi-Scale Geometric Flow for Segmenting Vasculature in MRI'. In: *Computer Vision Approaches to Medical Image Analysis (CVAMIA) and Mathematical Methods in Biomedical Image Analysis (MMBIA)*. [43, 45]
- Descoteaux, M., L. Collins, and K. Siddiqi: 2004b, 'A Multi-Scale Geometric Flow for Segmenting Vasculature in MRI'. In: *Medical Image Computing and Computer Assisted Intervention (MICCAI) 2004*. [48]
- Dice, L. R.: 1945, 'Measures of the amount of ecologic association between species'. *Ecology* **26**(3), 297–302. [50]
- DoCarmo, M. P.: 1976, *Differential Geometry of Curves and Surfaces*. Birkhauser. [10, 11, 12]
- Frangi, A., W. Niessen, K. L. Vincken, and M. A. Viergever: 1998, 'Multiscale Vessel Enhancement Filtering'. In: *MICCAI'98*. pp. 130–137. [iii, 6, 8, 10, 12, 15, 16, 19, 25, 36]
- Gage, M. and R. Hamilton: 1986, 'The Heat Equation Shrinking Convex Plane Curves'. *Journal of Differential Geometry* **23**, 69–96. [33]
- Grayson, M.: 1987, 'The Heat Equation Shrinks Embedded Plane Curves to Round Points'. *Journal of Differential Geometry* **26**, 285–314. [33]
- Hummel, R. A.: 1986, 'Representations Based on Zero-Crossings in Scale-Space'. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 204–209. [25]
- Kass, M., A. Witkin, and D. Terzopoulos: 1987, 'Snakes: Active Contour Models'. *International Journal of Computer Vision* **1**, 321–331. [19]
- Kichenassamy, S., A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi: 1995, 'Gradient Flows and Geometric Active Contour Models'. In: *International Conference On Computer Vision*. pp. 810–815. [31]
- Kimia, B. B., A. Tannenbaum, and S. W. Zucker: 1990, 'Toward a Computational Theory of Shape: An Overview'. In: *European Conference On Computer Vision, Lecture Notes in Computer Science*, Vol. 427. pp. 402–407. [33]
- Kimia, B. B., A. Tannenbaum, and S. W. Zucker: 1995, 'Shape, Shocks, and Deformations I: The Components of Two-Dimensional Shape and the Reaction-Diffusion Space'. *International Journal of Computer Vision* **15**, 189–224. [33]
- Koller, T. M., G. Gerig, G. Székely, and D. Dettwiler: 1995, 'Multiscale Detection of Curvilinear Structures in 2-D and 3-D Image Data'. In: *International Conference On Computer Vision*. pp. 864–869. [10, 12, 18]

- Krissian, K., J. Ellsmere, K. Vosburgh, R. Kikinis, and C.-F. Westin: 2003, 'Multiscale Segmentation of the Aorta in 3D Ultrasound Images'. In: *Engineering in Medicine and Biology Society*. pp. 638–641. [15, 58]
- Krissian, K., G. Malandain, and N. Ayache: 2000, 'Model-based detection of tubular structures in 3D images'. *Computer Vision and Image Understanding* **80**(2), 130–171. [10, 13, 15, 25]
- Lindeberg, T.: 1998, 'Edge Detection and Ridge Detection With Automatic Scale Selection'. *International Journal of Computer Vision* **30**(2), 77–116. [24, 25]
- Lorenz, C., I. Carlsen, T. Buzug, C. Fassnacht, and J. Weese: 1997, 'Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images'. In: *CVRMED-MRCAS'97, Lecture Notes in Computer Science*, Vol. 1205. pp. 233–242. [10, 16]
- Lorigo, L. M., O. D. Faugeras, E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin: 2001, 'CURVES: Curve evolution for vessel segmentation'. *Medical Image Analysis* **5**, 195–206. [19, 43, 54]
- Malladi, R., J. A. Sethian, and B. C. Vemuri: 1993, 'Topology-Independent Shape Modeling Scheme'. In: *Geometric Methods in Computer Vision II, SPIE*, Vol. 2031. pp. 246–258. [19, 39]
- Malladi, R., J. A. Sethian, and B. C. Vemuri: 1994, 'Evolutionary Fronts For Topology-Independent Shape Modeling and Recovery'. In: *European Conference On Computer Vision, Lecture Notes in Computer Science*, Vol. 800. pp. 3–13. [19, 39]
- Malladi, R., J. A. Sethian, and B. C. Vemuri: 1995, 'Shape Modeling with Front Propagation: A Level Set Approach'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(2), 158–175. [19, 39]
- McInerney, T. and D. Terzopoulos: 2000, 'T-snakes: Topology Adaptive Snakes'. *Medical Image Analysis* **4**, 73–91. [19]
- Osher, S. and C.-W. Shu: 1991, 'High-order Essentially Non-Oscillatory Schemes for Hamilton-Jacobi Equations'. *SIAM Journal of Numerical Analysis* **28**, 907–922. [37]
- Osher, S. J. and J. A. Sethian: 1988, 'Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations'. *Journal of Computational Physics* **79**, 12–49. [21, 38]
- Ostergaard, L., O. Larsen, G. Goualher, A. Evans, and D. Collins: 2000, 'Extraction of Cerebral Vasculature from MRI'. In: *9th Danish Conference on Pattern Recognition and Image Analysis*. [6, 10]

- Reinertsen, I., M. Descoteaux, S. Drouin, K. Siddiqi, and L. Collins: 2004, 'Vessel Driven Correction of Brain Shift'. In: *Medical Image Computing and Computer Assisted Intervention (MICCAI) 2004*. [5, 58]
- Sato, Y., S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis: 1998, '3d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images'. *Medical Image Analysis* **2**(2), 143–168. [10, 16]
- Siddiqi, K., Y. B. Lauzière, A. Tannenbaum, and S. W. Zucker: 1998, 'Area and Length Minimizing Flows for Shape Segmentation'. *IEEE Transactions on Image Processing* **7**(3), 433–443. [32]
- Vasilevskiy, A. and K. Siddiqi: 2002, 'Flux Maximizing Geometric Flows'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(12), 1–14. [iii, 6, 8, 20, 21, 24, 43, 54, 55]
- Wilson, D. L. and A. Noble: 1997, 'Segmentation of Cerebral Vessels and Aneurysms from MR Anangiography Data'. In: *Information Processing in Medical Imaging*. pp. 423–428. [17]
- Wink, O., W. J. Niessen, and M. A. Viergever: 2004, 'Multiscale Vessel Tracking'. *IEEE Transactions on Medical Imaging* **23**(1), 130–133. [10, 19]