

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

Maxime Descoteaux Louis Collins Kaleem Siddiqi *

May 3, 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 paper 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 [8] 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 [27] can be applied

*Maxime Descoteaux and Kaleem Siddiqi are with the Center for Intelligent Machines and the School of Computer Science, McGill University. Louis Collins is with the McConnell Brain Imaging Center and the Departments of Neurology and Neurosurgery at the Montréal Neurological Institute, McGill University. A preliminary version of this paper appears in the proceedings of the Computer Vision Approaches to Medical Image Analysis Workshop at ECCV, 2004.

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.

Keywords: Vessel segmentation, geometric flows, MRI, multi-scale analysis, validation.

1 Introduction

A three-dimensional (3D) representation of cerebral vasculature can be extremely important in pre-surgical planning, image-guided neurosurgery and clinical analysis. It is unfortunately often the case that in order to obtain such representations from an MRI volume 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, 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 from a single brain makes large scale clinical studies of vasculature infeasible. As a consequence, the computer vision and image analysis community has paid significant attention to automating the extraction of vessels or vessel centerlines. Several methods have been shown to give promising results on 2D projection angiography and 3D CT and MR angiography [28, 14, 3, 22, 20, 19, 27]. However, few techniques currently exist for the automatic extraction of vessel boundaries in more standard MRI volumes such as the proton-density (PD) weighted data set in Figure 1(a). Here it is




Figure 1: A mid-sagittal slice of a proton density (PD) weighted MRI volume, a time of flight (TOF) MR angiogram and a phase contrast (PC) MR angiogram of the same subject (one of the authors) acquired at our institute.

clear that a signal decrease is present in the vascular regions (the spaghetti-like structures), but the contrast between blood vessel and surrounding tissue is not as great when compared to the angiographic sequences (1(b) and 1(c)). Hence, the problem of recovering vessels from image intensity contrast alone on PD-weighted images is a challenge and requires shape information to constrain the segmentation. If successful, such a procedure could result in a vascular model that could be used in surgical planning while eliminating the need for an additional scan thus saving time during image acquisition and easing the burden on the patient as well as reducing the amount of time required to segment and prepare data for use in planning.

In this paper 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.* [24] where Frangi’s vesseness measure [8] 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 vesseness 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 [27] to be applied to recover vessel bound-

aries. 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 angiography and TOF angiography volumes, all obtained for the same subject (Figure 1).

The paper is outlined as follows. In Section 2 we review relevant background literature on the modeling of tubular structures, vessel segmentation and centerline extraction. We then develop our multi-scale geometric flow by incorporating Frangi’s vesseness measure [8] in the flux maximizing flow algorithm of [27] in Section 3. We present qualitative and quantitative validation results in Section 4. We then conclude with a discussion of the results and present directions for future work in Section 5.

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 also refer to the recent article by Aylward and Bullitt [3].

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} [18, 25, 8, 3, 14, 16, 24, 22]. 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 shape information. In particular, an Eigen value analysis extracts extremum changes in the normal vector to the iso-intensity level set passing through a particular point. At locations centered within tubular structures the

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 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.

smallest Eigen value of \mathbf{H} is close to zero (reflecting the low curvature along the direction of the vessel) and the two other Eigen values are high and are close to one another, reflecting the fact that the cross-section of the vessel is approximately circular. The associated Eigen vectors span the vessel direction and the cross-sectional plane, respectively. 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 1. Two prominent approaches for capturing vessel-like or tube-like structures based on the Hessian are the techniques proposed in [16, 8].

First, Krissian *et al.* propose a model-based approach to detecting tubular structures [16]. 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 [15].

Second, Frangi *et al.* propose a vesselness measure which incorporates information from

all three Eigen values and has an intuitive geometric interpretation [8]. This method is close in spirit to previous work by Lorenz *et al.* [18] and Sato *et al.* [25]. 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 1, it can be seen that R_B is non zero only for blob-like and noisy structures. The R_A ratio differentiates sheet-like from tube-like structures. Finally, S , the Frobenius norm, is used to ensure that random noise effects are suppressed from the response. 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

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

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.

2.2 Vessel Segmentation and Centerline Extraction Methods

Statistical Methods Wilson and Noble [28] propose a statistical approach for segmenting blood vessels from TOF angiography data, such as that shown in Figure 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

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

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 of Figure 1(a) or to Gadolinium enhanced MR volumes.

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(c). Aylward and Bullitt [3] 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. This idea is also the basis of work by Koller *et al.* for the multi-scale detection and traversal of curvilinear structures in intensity images [14]. 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 relate the problem of finding centerline paths to that of finding paths of least action in 3D intensity images [5]. 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 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. Beyond the requirement that

the user must specify the starting and end points of a particular path, the algorithm requires little user interaction. The flow is implemented using fast marching schemes and is hence computationally efficient.

Wink *et al.* have recently presented an approach to centerline extraction, applied in the context of vessel tracking, which combines features of the above two approaches [22]. More specifically, they use Frangi’s vesselness measure (Eq. 1) to characterize putative vessel centerline locations [8]. 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. Their method has been validated qualitatively in the presence of stenoses and imaging artifacts.

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.* [11]. These models have also been extended to handle changes in topology due to the splitting and merging of contours [21]. In the context of geometric flows for segmenting vasculature using level-set techniques, there are two recent approaches which are relevant to the development here. First, Lorigo *et al.* propose a regularization of a geometric flow in 3D using the curvature of a 3D curve [19]. This approach is grounded in the recent level set theory developed for mean curvature flows in arbitrary co-dimension [2]. 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 [2], 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 $\nabla\mathcal{I}$ 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 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 [27]. With S an evolving surface and $\vec{\mathcal{V}}$ the vector field, this flow is given by

$$\mathcal{S}_t = \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}} \quad (2)$$

where $\vec{\mathcal{N}}$ is the unit inward normal to each point on S . The motivation behind this flow is that it evolves a surface to a configuration where its normals are aligned with the vector field. 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.


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 section 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 [23].

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 [27] can be applied to recover them.

Returning now to Frangi’s vesselness measure (Eq. 1), 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 using derivatives of Lindeberg’s γ -parametrized normalized Gaussian kernels [17]. 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 [8]. In practice we have found these parameter settings to yield stable results over a wide range of image modalities. 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 2 for a synthetic branching structure and a synthetic helix. 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 which correspond to the estimated vessel orientation. 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.



(a) branching structure



(b) helix

Figure 2: A synthetic branching structure and a synthetic helix. For each structure the red vectors indicate the estimated vessel orientation at locations where the multi-scale vesselness measure (Eq. 1) is high.

This allows us to lift many of the restrictions on the flow pointed out in Section 2.2, because an explicit model of a tubular structure is now incorporated along with an appropriate notion of scale.

3.1 Distributing the vesselness measure to vessel boundaries

The key idea is to distribute the vesselness measure, which is concentrated at centerlines, to the vessel boundaries which are implied. At each voxel (x, y, z) where the vesselness measure is a local maximum in a $3 \times 3 \times 3$ neighborhood 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. If (x, y, z) is taken to be the origin $(0, 0, 0)$ and the xy plane is taken to coincide

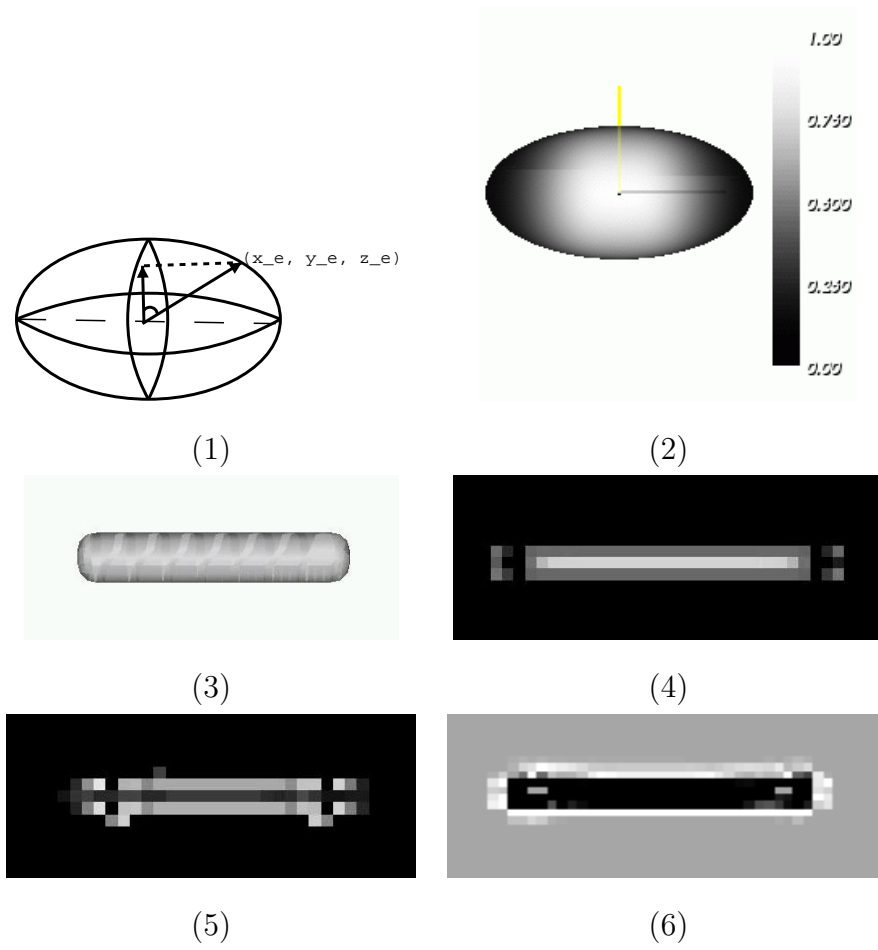


Figure 3: Distributing the vesselness measure to the implied boundaries. (1) The vector from the center of the ellipsoid to the surface voxel (x_e, y_e, z_e) , as well as its projection onto the cross-sectional plane, taken to be the xy plane. (2) We distribute the vesselness measure to all (x_e, y_e, z_e) on the ellipsoid by scaling it by the magnitude of this projection. The colour bar indicates the association between brightness and magnitude. (3) A synthetic tube of radius 2. (4) A view of the vesselness measure in a slice, with brighter regions indicate stronger vesselness. (5) A view of the ϕ distribution in the same slice. (6) The divergence of the vector field in Eq. 4, with transitions between dark and bright indicating zero-crossings. As expected, we have local maxima of the vesselness measure on the centerline in (4), local maxima of the ϕ distribution at the boundaries of the tube in (4) and zero-crossings of the divergence at the boundaries of the tube in (6).

with the cross-sectional plane this scale factor works out to be

$$\left\langle (x_e, y_e, z_e), \frac{(x_e, y_e, 0)}{\sqrt{x_e^2 + y_e^2}} \right\rangle = \sqrt{x_e^2 + y_e^2}. \quad (3)$$

This process of distributing the vesselness measure to the implied boundaries clearly favors voxels in the cross-sectional plane. 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

$$\vec{\mathcal{V}} = \phi \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|}. \quad (4)$$

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 integration². 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.

3.2 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. 1) 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). The surface evolution equation then works out to be

$$\begin{aligned} S_t &= \text{div}(\vec{\mathcal{V}}) \vec{\mathcal{N}} \\ &= \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \text{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right) \right] \vec{\mathcal{N}} \\ &= \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \kappa_{\mathcal{I}} \right] \vec{\mathcal{N}}. \end{aligned} \quad (5)$$

²This 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.

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. $\nabla\phi$ has a zero-crossing at vessel boundaries and $\nabla\mathcal{I}$ does not change sign. Hence, when the evolving surface overshoots the boundary slightly, this term acts to push it back toward the boundary. Such doublet terms have also shown to be beneficial in earlier geometric flows for segmentation [12, 4, 26].
2. The second term behaves like a geometric heat equation since $\kappa_{\mathcal{I}}$ is the mean curvature of the iso-intensity level set of the original intensity image. This equation has been extensively studied in the mathematics literature and has been shown to have remarkable anisotropic smoothing properties [9, 10]. It is also the basis for several nonlinear geometric scale-spaces such as those studied in [1, 13].
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. 5 will evolve toward the closest zero level set of the divergence of the vector field $\vec{\mathcal{V}}$.

3.3 Implementation Details

Below we review some of the details of the implementation of our multi-scale geometric flow (Eq. 5), which is based on level set methods [23].

1. We compute the Hessian operator over 10 log scales and select the maximum vesselness response as described in Section 2.1. We use Jacobi’s method for symmetric matrices to find the Eigen values of the Hessian.
2. The ϕ distribution in Section 3.1 is carried out from voxels at vessel centerlines since at such locations one has strong confidence in the scale and orientation estimate from

Frangi's vesselness measure [8]. This is done using the following procedure

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

Distribute vesselness over ellipsoid

For most examples we use a vesselness threshold of 0.01 and a percentile of 0.75 and local_max is the maximum vesselness response in a 3x3x3 neighborhood of the voxel.

3. 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.
4. $\kappa_{\mathcal{I}}$, the mean curvature of each intensity iso-surface is computed using a 3-neighbor central difference scheme for all derivatives:

$$\kappa_{\mathcal{I}} = \frac{(\mathcal{I}_{yy} + \mathcal{I}_{zz})\mathcal{I}_x^2 + (\mathcal{I}_{xx} + \mathcal{I}_{zz})\mathcal{I}_y^2 + (\mathcal{I}_{xx} + \mathcal{I}_{yy})\mathcal{I}_z^2 - 2(\mathcal{I}_x\mathcal{I}_y\mathcal{I}_{xy} - \mathcal{I}_x\mathcal{I}_z\mathcal{I}_{xz} - \mathcal{I}_y\mathcal{I}_z\mathcal{I}_{yz})}{(\mathcal{I}_x^2 + \mathcal{I}_y^2 + \mathcal{I}_z^2)^{\frac{3}{2}}} \quad (6)$$

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

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

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. 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.

6. The evolving surface S is obtained as the zero level set of the Ψ function. The numerical derivatives used to estimate $||\nabla \Psi||$ must be computed with up-winding in the proper direction as described in [23].

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 PD, TOF and PC volumes, all obtained for the same subject.

4.1 Image acquisition

We acquired four different volumes from the same subject (one of the authors) 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. We supplemented these three data sets with an MRA volume (Figure 4) and a Gadolinium enhanced MRI volume (Figure 5), 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.

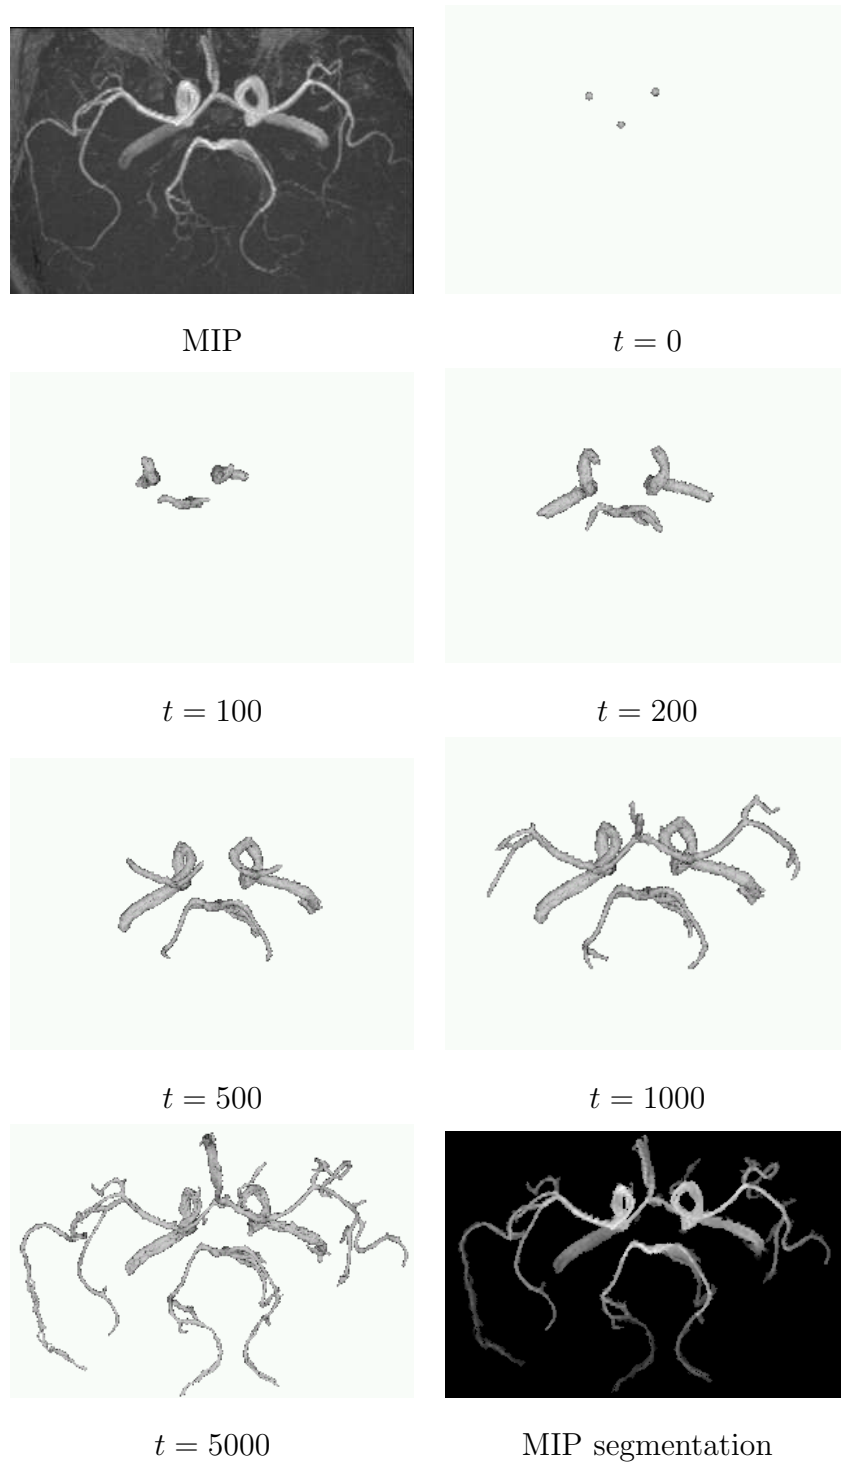


Figure 4: An illustration of the multi-scale geometric flow on a $68 \times 256 \times 256$ MRA image. An MIP of the data is shown at the top left and the other images depict different stages of the evolution from three seeds. The bottom right figure depicts an MIP of the input MRA data masked by the binary segmentation.





Figure 5: 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)).


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 3.3. We should point out that whereas prior geometric flow based methods [19, 27] could be applied to the angiographic volumes, they would fail entirely on both the Gadolinium enhanced MRI volume and the PD data set where high contrast regions are not limited to vessel boundaries.


Figure 4 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.



PC




vesselness of PC



PC masked by segmentation



TOF




vesselness of TOF




TOF masked by segmentation



PD



vesselness of PD



PD masked by segmentation
(reversed contrast)

Figure 6: Transverse views of (maximum or minimum) intensity projections of the PC, TOF and PD data sets, the associated vesselness measures and the segmentations obtained by the multi-scale geometric flow. See the associated text for a discussion.

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 [6]. 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.

Figure 5 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 integration scheme is able to connect a section of the supra-callosal arteries which has very low contrast in the original Gadolinium data set.

Finally, Figure 6 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. 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 modalities. 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 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.

4.3 Quantitative Results


Figure 7 compares the segmentations obtained on the PC, TOF and PD volumes (Figure 6) 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 7 indicate that the TOF reconstruction is missing a large number of vessel labels when compared to the PC and PD reconstructions.

We now carry out a quantitative analysis of these segmentation results by computing a number of statistics between each pair of modalities, treating one as the ground truth data set and the the other as the test data set. These comparisons are shown in Table 2 and include the following measures:


1. The kappa coefficient defined by

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


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 [7].



PC (truth) vs PD (test)



TOF (truth) vs PD (test)



PC (truth) vs TOF (test)

Figure 7: Each row shows a pair-wise 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.

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

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 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 (this is essentially the set of red voxels in the figure 7).

It is clear from Table 2 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, a vessel extracted from the different data sets is indeed similar.

5 Conclusion and Discussion

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 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 $\vec{\mathcal{V}}$ (Eq. 4). 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. The results in Figure 6 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. 6. 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 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 (maximum) number of iterations, or when the flow had not hit the narrow band for several (5000) 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 [3, 15] 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.

Acknowledgements This work was supported by grants from NSERC, FQRNT, CFI and CIHR. We thank Bruce Pike, Simon Duchesne and Ingerid Reinertsen for helpful discussions on the validation of our approach.

References

- [1] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel. Axiomes et équations fondamentales du traitement d’images. *C. R. Acad. Sci. Paris*, 315:135–138, 1992.
- [2] L. Ambrosio and H. M. Soner. Level set approach to mean curvature flow in arbitrary codimension. *Journal of Differential Geometry*, 43:693–737, 1996.
- [3] S. R. Aylward and E. Bullitt. Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction. *IEEE Transactions On Medical Imaging*, 21(2):61–75, 2002.
- [4] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *International Conference On Computer Vision*, pages 694–699, 1995.
- [5] T. Deschamps and L. D. Cohen. Fast Extraction of Minimal Paths in 3D Images and Applications to Virtual Endoscopy. *Medical Image Analysis*, 5(4):281–299, December 2001.
- [6] M. Descoteaux, L. Collins, and K. Siddiqi. 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)*, May 2004.
- [7] L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.
- [8] A. Frangi, W. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI’98*, pages 130–137, 1998.
- [9] M. Gage and R. Hamilton. The heat equation shrinking convex plane curves. *Journal of Differential Geometry*, 23:69–96, 1986.

- [10] M. Grayson. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry*, 26:285–314, 1987.
- [11] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1987.
- [12] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *International Conference On Computer Vision*, pages 810–815, 1995.
- [13] B. B. Kimia, A. Tannenbaum, and S. W. Zucker. Shape, shocks, and deformations I: The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15:189–224, 1995.
- [14] T. M. Koller, G. Gerig, G. Székely, and D. Dettwiler. Multiscale detection of curvilinear structures in 2-d and 3-d image data. In *International Conference On Computer Vision*, pages 864–869, 1995.
- [15] K. Krissian, J. Ellsmere, K. Vosburgh, R. Kikinis, and C.-F. Westin. Multiscale segmentation of the aorta in 3d ultrasound images. In *Engineering in Medicine and Biology Society*, pages 638–641, 2003.
- [16] K. Krissian, G. Malandain, and N. Ayache. Model-based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80(2):130–171, November 2000.
- [17] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):77–116, 1998.
- [18] C. Lorenz, I. Carlsen, T. Buzug, C. Fassnacht, and J. Weese. 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*, volume 1205, pages 233–242, 1997.

- [19] L. M. Lorigo, O. D. Faugeras, E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin. Curves: Curve evolution for vessel segmentation. *Medical Image Analysis*, 5:195–206, 2001.
- [20] T. McInerney and D. Terzopoulos. Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Transactions on Medical Imaging*, 18(10):840–850, 1999.
- [21] T. McInerney and D. Terzopoulos. T-snakes: Topology adaptive snakes. *Medical Image Analysis*, 4:73–91, 2000.
- [22] W. O and M. A. V. W. J. Niessen. Multiscale vessel tracking. *IEEE Transactions on Medical Imaging*, 23(1):130–133, 2004.
- [23] S. J. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [24] L. Ostergaard, O. Larsen, G. Goualher, A. Evans, and D. Collins. Extraction of cerebral vasculature from mri. In *9th Danish Conference on Pattern Recognition and Image Analysis*, August 2000.
- [25] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis. 3d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
- [26] K. Siddiqi, Y. B. Lauzière, A. Tannenbaum, and S. W. Zucker. Area and length minimizing flows for shape segmentation. *IEEE Transactions on Image Processing*, 7(3):433–443, 1998.
- [27] A. Vasilevskyi and K. Siddiqi. Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(12):1–14, 2002.
- [28] D. L. Wilson and A. Noble. Segmentation of cerebral vessels and aneurysms from mr angiography data. In *Information Processing in Medical Imaging*, pages 423–428, 1997.