

A Geometric Flow for Segmenting Vasculature in Proton-Density Weighted MRI

Maxime Descoteaux^a D. Louis Collins^b Kaleem Siddiqi^{c,*}

^a*Odyssée Project Team, INRIA, Sophia-Antipolis, France*

^b*McConnell Brain Imaging Centre, Montreal Neurological Institute, McGill University*

^c*Center for Intelligent Machines & School of Computer Science, McGill University*

Abstract

Modern neurosurgery takes advantage of magnetic resonance images (MRI) of a patient's cerebral anatomy and vasculature for planning before surgery and guidance during the procedure. Dual echo acquisitions are often performed that yield proton density (PD) and T2-weighted images to evaluate edema near a tumor or lesion. In this paper we develop a novel geometric flow for segmenting vasculature in PD images, which can also be applied to the easier cases of MR angiography data or Gadolinium enhanced MRI. Obtaining vasculature from PD data is of clinical interest since the acquisition of such images is widespread, the scanning process is non-invasive, and the availability of vessel segmentation methods could obviate the need for an additional angiographic or contrast-based sequence during preoperative imaging. The key idea is to first apply Frangi's vesselness measure [1] to find putative centerlines of tubular structures along with their estimated radii. This measure is then distributed to create a vector field which allows the flux maximizing flow algorithm of [2] to be applied to recover vessel boundaries. 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 validate the approach quantitatively on a single-subject data set consisting of PD, phase contrast (PC) angiography and time of flight (TOF) angiography volumes, with an expert segmented version of the TOF volume viewed as the ground truth. We then validate the approach quantitatively on 19 PD data sets from a new digital brain phantom, with semi-automatically obtained labels from the corresponding angiography volumes viewed as ground truth. A significant finding is that both for the single-subject and multi-subject studies, 90% or more of the vasculature in the ground truth segmentation is recovered from the automatic segmentation of the other volumes.

Key words: vessel segmentation, geometric flows, proton-density MRI, validation.

1 Introduction


A three-dimensional (3D) representation of cerebral vasculature can be extremely important in pre-surgical planning and image-guided neurosurgery. For example, a vessel-free path must be identified when planning for tumor biopsies or stereotaxically placed depth electrodes for the evaluation of epilepsy since access to the surgical target is achieved through a small burr-hole in the skull. In open craniotomies, the surgical approach may be constrained by the location of major vessels. 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. Even when using angiographic sequences, 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.

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 [2–10]. 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). The rationale for using PD is two-fold. First, by taking advantage of the “black-blood” contrast phenomenon by which there is a signal decrease in vascular regions (see [11,12]), it may in fact be possible to obtain a more complete vessel tree when compared to time-of-flight (TOF) or phase-contrast (PC) angiography. This translates to a more efficient use of PD images that are already acquired in a dual echo sequence in the imaging protocol for patients at our institute. Second, magnetic resonance angiography (MRA) acquisitions can be time consuming, and since the imaging protocol includes two T1 acquisitions (with and without Gadolinium), a T2/PD acquisition and can include diffusion MRI and fMRI, any possible time savings during acquisition is appreciated by the patient and by the medical staff.

The problem of segmenting vasculature from PD is non-trivial because the contrast between blood vessel and surrounding tissue is not as strong as it is in angiographic sequences (see Figs. 1(b) and 1(c)). Furthermore, in contrast to angiographic data, in PD images vessels can have similar contrasts at their boundaries to the boundaries of other anatomical structures. Hence, the prob-

* Corresponding author. Address: McGill University, School of Computer Science, Rm 318, McConnell Engineering, 3480 University Street, Montréal, QC H3A 2A7, Canada. Tel: (514) 398-3371. Fax: (514) 398-7348.

Email address: siddiqi@cim.mcgill.ca (Kaleem Siddiqi).



(a) PD

(b) TOF

(c) PC

Fig. 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.

lem of recovering vessels from image intensity contrast alone on PD-weighted images requires shape information to constrain the segmentation. If successful, after careful validation with radiographic angiography techniques, such a procedure could result in a vascular model that could be used in surgical planning. This would eliminate the need for an additional angiographic sequence or a contrast-based scan and would thus ease the burden on the patient as well as reduce the amount of time required to segment and prepare data for use in planning.

When applied to data from normal subjects, a technique that builds a 3D representation of the vasculature tree can also be used to study anatomical variability in the normal population and this information may be used to build a normative database for comparison with diseased cases. Dose constraints prohibit the use of techniques like radiographic or CT-based angiography on normal subjects. While MRI-based imaging methods of the vasculature are promising, TOF and PC techniques have limitations in resolution, and in detection of small vessels and may be associated with long acquisition times. In addition, contrast agents such as Gadolinium cannot be used on normal subjects. Since the technique proposed here is applicable to PD-weighted images, one could take advantage of existing databases of normal subjects that already contain T1, T2 and PD-weighted data (e.g., [13,14]) to build models of vessel-tree variability.

In this article we develop an 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. We clarify that our aim is to detect and identify the existence of vessel segments with two main applications in mind. The first is to analyze vessel morphology in normal subjects, where eventually we would like to characterize the variability in branching patterns. The second goal is to aid in surgical planning for patients that are candidates for cranial tumor surgery. In both cases, our interest is focused on where the blood vessel lies with respect to neighboring vessels and tissues, and not in the

exact boundary of the vessel itself. In the case of surgery, knowing that the vessel lies in the surgical path is important, and the surgeon will plan to take another path to avoid the vessel or will decide to continue, but perhaps maintain a safety margin of 5-10mm to the vessel. Knowing the exact boundary of the vessel could certainly be interesting to evaluate stenosis or to compare vessel sizes between subjects, and has been the focus of recent articles [9,10]. These surface evolution approaches, as we explain in Section 2, though related to ours, are not applicable to automatic full brain segmentation from PD.

The method we develop is motivated by the technique of [15], where Frangi's vesselness measure [1] is thresholded to find centerlines and then tubular fits to vessel boundaries are then obtained using a form of connected component analysis and a generalized cylinder model. However, rather than threshold the vesselness measure, which yields incomplete vessel trees, we extend it to yield a vector field which is locally normal to putative vessel boundaries, in their vicinity. This in turn allows the flux maximizing geometric flow of [2] to be applied to recover vessel surfaces. This flow has a formal motivation, is topologically adaptive due to its implementation using level set methods, and finally is computationally efficient. Furthermore, connections between this flow and optimal edge integrators related to zero-crossings of a Laplacian have been developed in [16].

The main contributions of this article are:

- (1) We combine a vesselness measure [1] with a geometric flow [2] for vessel segmentation. Whereas this has been done in the context of carotid artery segmentation [9] our goal is develop a method which can be applied to automatic full brain segmentation from PD images, which are routinely acquired in neurosurgical applications.
- (2) We demonstrate qualitative results on simulated data, magnetic resonance angiography (MRA) data, as well as on the more challenging cases of Gadolinium enhanced MRI and PD weighted MRI volumes.
- (3) We validate the approach quantitatively in three ways. First, we evaluate the automatic segmentations of numerically simulated tubes of varying radius and curvature. Second, we compare the automatic segmentations of PD, PC angiography and TOF angiography volumes, all obtained for the same subject (Figure 1), with one another and with a ground truth segmentation obtained by an expert on the TOF volume. Third, we carry out a multi-subject study where segmentations obtained on PD volumes are compared with ground truth vessel data for 19 subjects from a recent digital brain phantom [17].

The article is organized 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

eigenvalue 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 eigenvalues of the Hessian matrix. Here, we assume that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. The sign of the highest eigenvalues 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.

incorporating Frangi’s vesselness measure [1] in the flux maximizing flow algorithm of [2] in Section 3. We present qualitative and quantitative validation results in Section 4. We 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.

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 eigenvalues of the Hessian matrix \mathbf{H} [18,19,1,5,4,20,15,6]. 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 eigenvalue 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 smallest eigenvalue of \mathbf{H} is close to zero (reflecting the low curvature along the direction of the vessel) and the two other eigenvalues are high and are close to one another, reflecting the fact that the cross-section of the vessel is approximately circular. The associated eigenvectors span the vessel direction and the cross-sectional plane, respectively. The eigenvalue 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 captur-

ing vessel-like or tube-like structures based on the Hessian are the techniques proposed in [20,1].

First, Krissian *et al.* propose a model-based approach to detecting tubular structures [20]. An eigenvalue 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 eigenvalues 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 [21].

Second, Frangi *et al.* propose a vesselness measure which incorporates information from all three eigenvalues and has an intuitive geometric interpretation [1]. This method is close in spirit to previous work by Lorenz *et al.* [18] and Sato *et al.* [19]. 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

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

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 eigenvector associated with the smallest eigenvalue of the Hessian gives its local orientation.

¹ The 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 $\lambda_3 > 0$.

2.2.1 Statistical Methods

Wilson and Noble [3] 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 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.

A more recent method in this class is that of Tyrrell et al. [22], where 3D vasculature is modeled using superellipsoids. These primitives are fit to image data using a robust estimation framework to handle artifacts and noise. Applications to vessel centerline tracing and branchpoint extraction are demonstrated on synthetic imagery.

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(c). Aylward and Bullitt [5] present a centerline tracking approach which is based on a characterization of intensity ridges in 3D data sets. The eigenvectors 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 [4]. 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 [23]. 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 [6]. More specifically, they use Frangi’s vesselness measure (Eq. 1) to characterize putative vessel centerline locations [1]. 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.

More recently, Sofka and Stewart have addressed the problem of vessel centerline extraction using a combination of matched filter responses, vessel boundary measures and vessel confidence measures [24]. This combination is shown to demonstrate potential improvements to centerline detection based on the Hessian. Their algorithm has been applied to finding narrow, low contrast vessels in retinal images with the goal of eliminating false positives.

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.* [25]. These models have also been extended to handle changes in topology due to the splitting and merging of contours [26]. In the context of geometric flows for segmenting vasculature using level-set techniques, there are a few 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 [8]. This approach is grounded in the recent level set theory developed for mean curvature flows in arbitrary co-dimension [27]. 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 eigenvalue of a particular matrix [27], 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 [2]. 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. 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.

A method which extends the domain of application of geometric flows to contrast enhanced MRA is that of [9]. Building on the work of [1] and [23] a vesselness map is first created by applying Frangi's vesselness measure to the input data. A centerline is then determined between two user selected seed points as a minimal cost path between them. This centerline is used to initialize an outward snake model [25] implemented using level set ideas, with the image-based speed term being a function of a full-width-half max criterion to determine vessel width. This choice of speed term has been specifically designed for segmentation of the internal carotid artery, with accurate stenoses characterization and vessel width estimation in mind. It requires a quantity

$i - max$ to be determined separately in cross-sectional planes to the centerline and hence cannot be applied, without further modification, to automatic full brain segmentation in PD.

A distinct geometric flow-based approach with application to vessel tracking in the presence of branching configurations, is that of [10]. The method incorporates a standard active surface model, with the exponential form of [28] to yield a speed term inversely proportional to edge strength. The key novelty is to extract a curve skeleton and to constrain its topology, which is assumed to be known in advance, to model vessel segments or vessel bifurcations. This constrained topology is then used to reinitialize the evolution. In this fashion segmentation using the flow and skeletal-constrained re-initialization alternate. This method has been successfully applied to vessel axis tracking in the carotid artery using CT angiography data sets with controlled seeding. However, it too is not designed for automatic full brain segmentation of vasculature from PD data.

In the following section we build on the above developments by using Frangi's vesselness measure as the basis of an extension mechanism by which vessel boundaries are approximately modeled. We then use conservative seeding to initialize an outward geometric flow to recover vessels. Our goals are different than those of [9,10], in particular, we assume unconstrained topologies and aim to efficiently recover full brain vasculature from PD data. We also show that the method can be applied to other modalities including angiographic data and Gadolinium enhanced MRI volumes.

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 signal 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 (in their approximate vicinity) so that the flux maximizing flow algorithm of [2] can be applied to recover them. In spirit this idea is similar to the basis of the technique of [9], but we do not explicitly extract individual centerlines between user defined seed points.

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

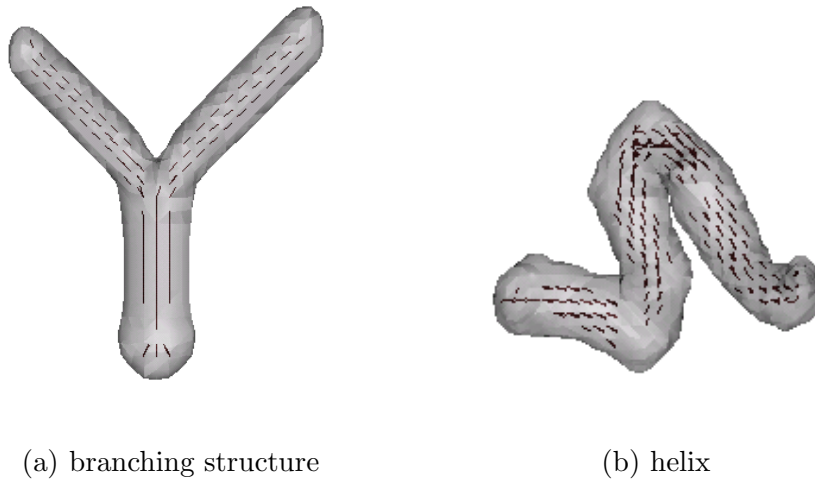


Fig. 2. A synthetic branching structure and a synthetic helix. For each structure the black vectors indicate the estimated vessel orientation at locations where the multi-scale vesselness measure (Eq. 1) is high.

the entries which comprise the Hessian matrix by using derivatives of Lindeberg’s γ -parametrized normalized Gaussian kernels [29]. Vesselness responses are then computed at different scales and the scale associated with the maximum response is selected. This scale gives an approximation to the radius of the vessel and the eigenvector associated with the smallest eigenvalue its local orientation. Implementation details are discussed in Section 3.3.

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 eigenvectors which correspond to the estimated vessel orientation. With the exception of severe bends where the local geometry is “blob-like” (Figure 2(b)) it is apparent that locations of high vesselness are close to the expected centerlines, and that the estimated vessel orientation at such 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. 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 (vessel radius).

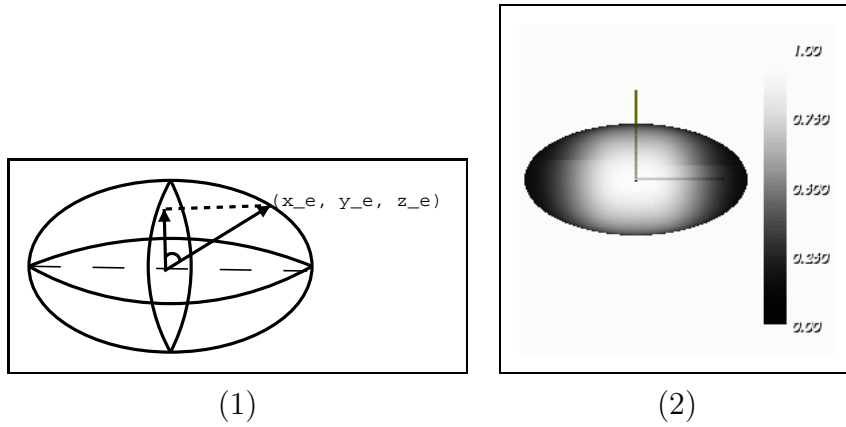


Fig. 3. Distributing the vesselness measure to the implied boundaries. (1) The vector from the center of the ellipsoid to the surface voxel (x_e, y_e, z_e) , as well as its projection onto the cross-sectional plane, taken to be the xy plane. (2) We distribute the vesselness measure to all (x_e, y_e, z_e) on the ellipsoid by scaling it by the magnitude of this projection. The grayscale bar indicates the association between brightness and magnitude.

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 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. It is important to point out that this distribution of the vesselness measure can only locate vessel boundaries approximately since: 1) the ellipsoidal section is only an approximate model of the anatomical ground truth, 2) the model is not well-suited for branching regions (where vessels bifurcate), 3) the estimated radius is approximate, and 4) the voxel center may not coincide exactly with the true centerline of the vessel. By adopting stronger assumptions on

the distribution of intensities along vessel cross-sections, such as the Gaussian profile assumption used in [20], it is possible to model the relationship between centerlines and boundaries more precisely. However, such assumptions are not easily justified in the case of PD images. Our strategy therefore is to compensate for this approximation by combining this ϕ measure with the gradient of the *original* PD image (in which strong gradients are expected to occur at true vessel boundaries). More specifically, an extended vector field is 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 large in the vicinity of 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 of the *original* intensity image, where the gradient magnitude 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

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

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.

² 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.

- (1) The first term $\langle \nabla\phi, \frac{\nabla\mathcal{I}}{\sqrt{|\nabla\mathcal{I}|}} \rangle$ acts like a doublet. $\nabla\phi$ has a zero-crossing at vessel boundaries and $\frac{\nabla\mathcal{I}}{\sqrt{|\nabla\mathcal{I}|}}$ 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 [30–32].
- (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 [33,34]. It is also the basis for several nonlinear geometric scale-spaces such as those studied in [35,36].
- (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 $\vec{\nabla}$.

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 [37].

- (1) To compute 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 [1]. 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. We use Jacobi’s method for symmetric matrices to find the eigenvalues 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 [1]. This is done using the following procedure:

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

$$\{\text{Distribute vesselness over ellipsoid}\}$$

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. The voxels inside the ellipsoidal regions give the seeds used to initialize the geometric flow.

- (3) The derivatives in the doublet term $\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \rangle$ 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. This is essentially because ϕ is by construction relatively smooth since it is constructed by superposition. On the other hand, the gradient of the original intensity image can be sharp and is better captured by an ENO method.
- (4) $\kappa_{\mathcal{I}}$, the mean curvature of each intensity iso-surface is computed using a 3-neighbor central difference scheme for all derivatives:

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

- (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 $\mathcal{F} = \langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \rangle + \phi \operatorname{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right)$, 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 [37].

3.4 Validation of Vesselness Extension

In order to validate the extension of the vesselness measure to boundaries and to evaluate the effectiveness of the speed term driving the geometric flow, we constructed several binary synthetic cylindrical tubes of varying cross-sectional radius and centerline curvatures, with integer values of 1 in the interior and 0 in the background. Each tube was then smoothed using 2 iterations of mean curvature smoothing to simulate partial volume effects at their boundaries. Ground truth surface points were obtained as the 0.5 crossings of each tube, obtained by linear interpolation on the voxel grid. We then added white Gaussian noise to each voxel to simulate noise levels in a typical MR acquisition scan. We used a mean of 0 with a standard deviation equal to 10% of the

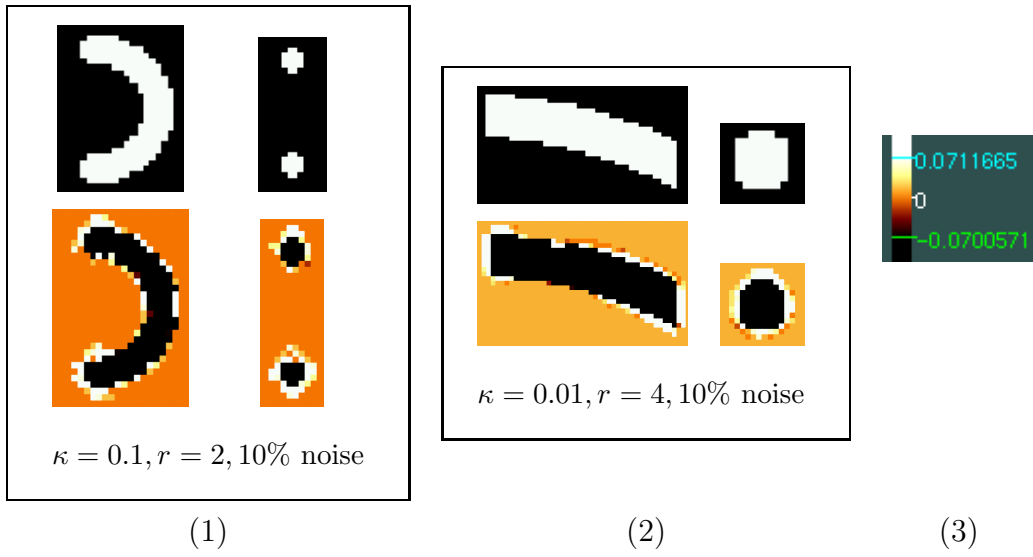


Fig. 4. Distributing the vesselness measure to the implied boundaries. (1) and (2) show views of synthetic tubes of varying radius r and curvature κ (top), along with the numerical speed term after mean curvature smoothing and Gaussian noise addition, as explained in the text. (3) A color-bar to indicate the speed values. Zero-crossings of the speed term correspond to surface points to which the geometric flow would converge.

radius	average distance error, maximum error (voxels)					
0.5	0.29, 0.95	0.48, 1.53	0.68, 1.92	0.47, 1.13	0.51, 1.12	0.47, 1.30
1.0	0.45, 1.14	0.41, 1.15	0.42, 0.78	0.41, 1.10	0.32, 1.17	0.22, 0.51
1.5	0.40, 0.95	0.37, 0.95	0.43, 1.15	0.36, 1.09	0.27, 0.93	0.23, 0.51
2.0	0.43, 1.15	0.37, 1.00	0.40, 1.14	0.34, 1.17	0.31, 0.80	0.21, 0.67
2.5	0.35, 0.87	0.34, 1.15	0.36, 1.16	0.34, 1.09	0.33, 0.91	0.21, 0.68
3.0	0.35, 0.99	0.38, 1.09	0.37, 1.13	0.37, 1.09	0.39, 0.96	0.36, 1.15
3.5	0.34, 1.10	0.34, 1.10	0.33, 1.10	0.34, 1.09	0.38, 1.09	0.39, 1.17
4.0	0.32, 1.08	0.34, 0.95	0.32, 0.95	0.34, 1.05	0.34, 1.07	0.32, 1.13
4.5	0.31, 1.29	0.37, 1.37	0.32, 1.7	0.31, 1.07	0.32, 1.07	0.32, 1.13
5.0	0.30, 1.29	0.31, 1.47	0.30, 1.52	0.29, 1.18	0.30, 1.06	0.29, 1.06
	$\kappa = 0.1$	$\kappa = 0.04$	$\kappa = 0.02$	$\kappa = 0.01$	$\kappa = 0.002$	$\kappa = 0$
	curvature κ					

Table 2

Average and maximum Euclidean distance errors between empirical surface points and their nearest ground truth surface points for a variety of synthetic tubes, with 10% Gaussian noise added (see text).

mean intensity level³. We then followed the steps detailed above to construct the vesselness measure, obtain the ϕ distribution and then build the vector field $\vec{\mathcal{V}}$, using the same parameters throughout. Empirical surface points were defined as the zero-crossings of the speed term $div(\vec{\mathcal{V}})$ in the geometric flow (Eq. 5).

A visualization of the speed term for two example synthetic tubes is shown in Figure 4 (1) and (2). In these sub-figures the zero-crossings of the speed term are indicated by the color bar. We evaluated the accuracy of these locations by computing the average and maximum Euclidean distance errors between each empirical surface voxel and its closest ground truth surface voxel, for each tube. The results, shown in Table 2, indicate that the average error is typically less than 0.4 voxels. We have determined empirically that the maximum errors occur at the two ends of each synthetic tube, which is to be expected since mean curvature smoothing causes the most shrinkage there. These simulation results have the limitation that the synthetic tubes had planar centerlines. True 3D vessels would have corresponding centerlines with torsion which could lead to larger errors.

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. For this data set we obtained an expert segmentation of the TOF volume by voxel painting, which we treat as the ground truth segmentation. Finally, we carry out a multi-subject study where we compare segmentations obtained using our method on PD data, against the corresponding ground truth vasculature from a recent digital brain phantom [17]. These single-subject and multi-subject quantitative comparisons give an indication of how successfully the presence of vessel segments are detected. Whereas they do not directly address the accuracy of vessel boundary segmentation we present the distance errors between ground truth and segmented vessels and demonstrate that they are small.

³ We also tried lower (5%) and higher (15%) noise levels, but found that the average distance errors did not change significantly from those reported in Table 2.

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 (1mm voxels in plane, 2mm thick slices, 50% overlap acquisition, reconstructed on a 3D 1mm isotropic grid to yield 1mm³ 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 5) and a Gadolinium enhanced MRI volume (Figure 6), 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.

For the multi-subject study we used data from 19 subjects taken from the digital brain phantoms of [17]. These MRI volumes were acquired on a 1.5T Siemens Sonata Vision clinical scanner (Siemens Medical Systems, Erlangen, Germany). The data includes T1, T2, PD-weighted scans and a magnetic resonance angiography (MRA) scan. The T1-weighted scans were acquired using a 3D spoiled gradient echo (GRE) sequence (sagittal acquisition, TR=22ms, TE=9.2ms, flip angle 30 degrees, 140-160 slices, 1mm³ voxels). The PD and T2-weighted images were acquired with a dual echo turbo spin echo (TSE) protocol (TR=3300ms, TE=15,104ms, 1mm in-plane, 80 2mm thick slices) . The MRA was acquired using a 3D phase contrast spoiled gradient-echo angiography scan with 176 0.9 mm slices (TR=71ms, TE=8.2ms, flip angle 15 degrees), with 0.47x0.47mm in-plane resolution.

Whereas it could be possible to use high resolution acquisitions such as CT angiography to obtain the ground truth data set, dose limitations precluded this for the 19 normal controls used in the multi-subject study. It was also not possible to rely solely on manual segmentation by painting voxels, because

this would have taken on the order of 7 to 10 hours per subject. As a result we relied on the following strategy to obtain a ground truth “silver” standard for the evaluation of our segmentation results. First the 19 T1, T2, PD and MRA volumes were registered together using a mutual information objective function with a rigid body transformation. (Note that this does not account for possible differing geometric distortion between scans.) The resulting volumes were mapped into stereotaxic space to facilitate comparisons between subjects [38]. Two methods were then used for segmentation, one for surface vessels and one for internal vessels. Surface vessels were segmented by hand on the PD data, using the registered MRA volume for guidance. To obtain the internal vessels we used the technique described in [39] on the MRA volume. These results were then corrected by hand using a process of local thresholding and manual painting and erasing of individual voxels to add false negative voxels and remove false positives. This silver standard is slightly biased in favor of our technique because it too is based on the flux maximizing flow [2]. Nonetheless, this process of semi-automatic segmentation took on the order of 1 to 2 hours per subject, making it possible to include 19 data sets for validation. Furthermore, the multi-subject study is only one part of our overall validation framework.

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 [8,2] 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 5 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 [39]. 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 6 depicts a $133\text{mm} \times 214\text{mm} \times 75\text{mm}$ 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

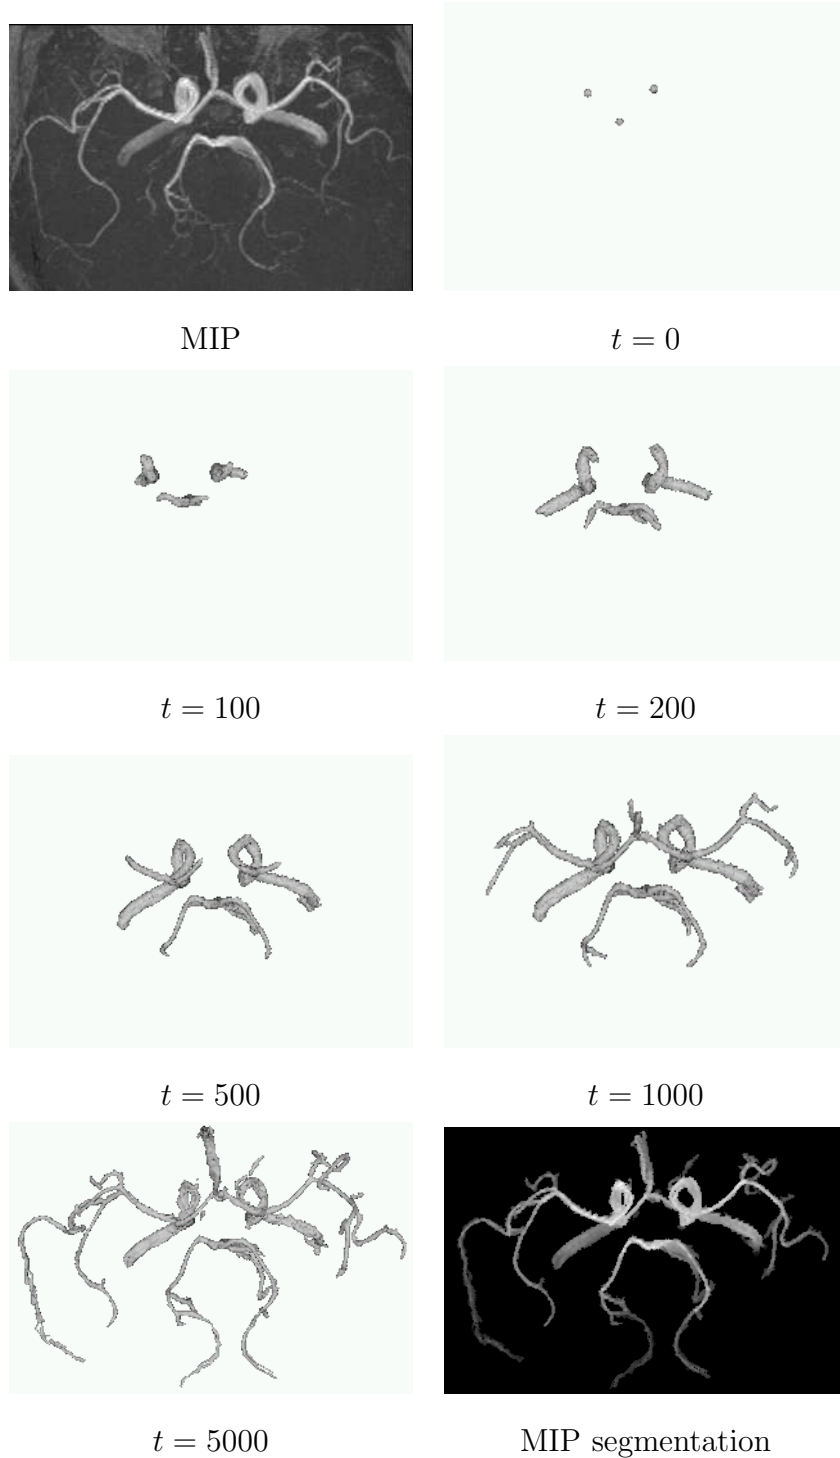


Fig. 5. An illustration of the multi-scale geometric flow on a $256 \times 256 \times 68$ MRA image. An MIP of the data is shown at the top left and the other images depict different stages of the evolution from three seeds. The bottom right figure depicts an MIP of the input MRA data masked by the binary segmentation.

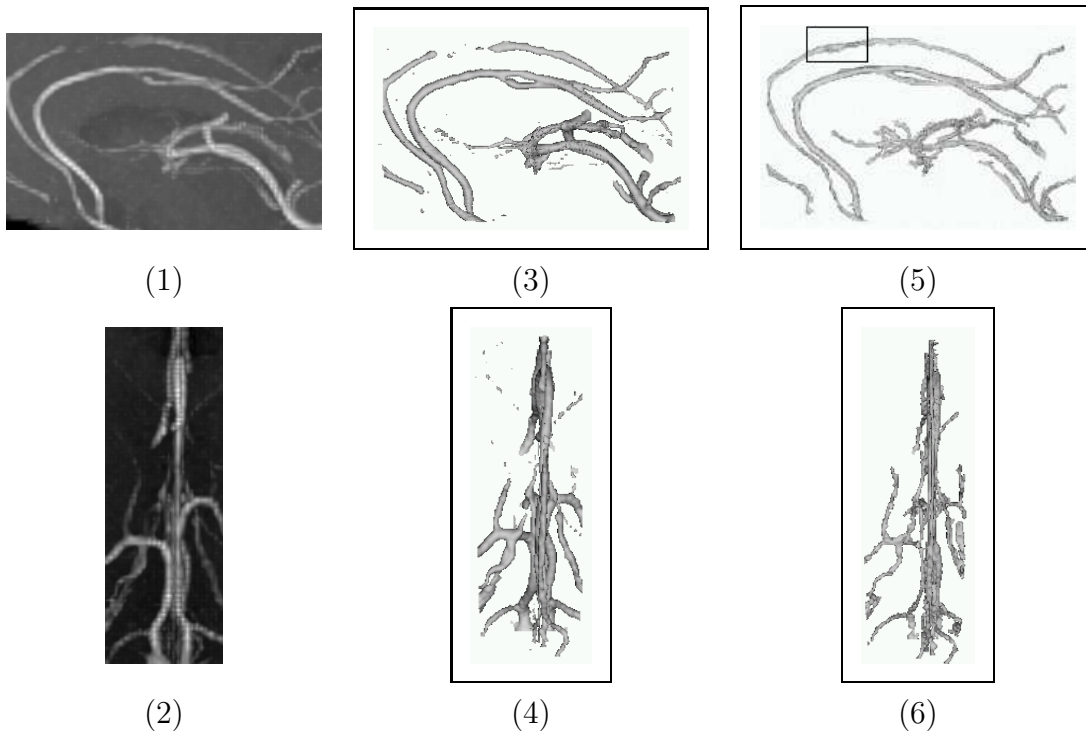


Fig. 6. An illustration of the flow on a 133mm x 214mm x 75mm cropped region of a Gadolinium enhanced MRI. An MIP of the sagittal and transverse views of the data is shown in (1) and (2). Reconstructions obtained by simple thresholding for the same views are shown in (3) and (4). These are clearly sensitive 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)).

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 7 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 x 217 x 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

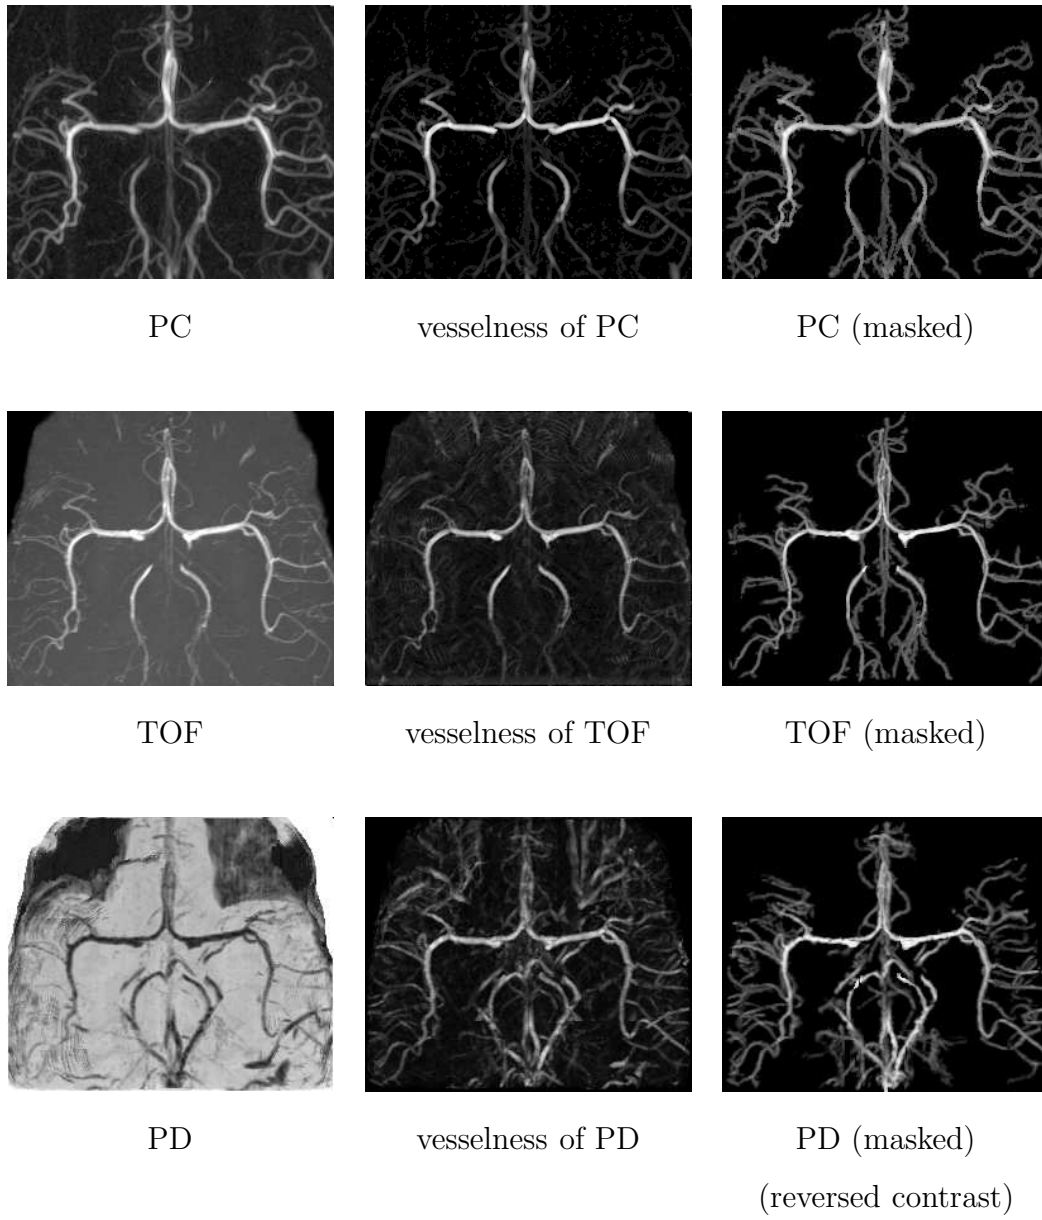


Fig. 7. LEFT: Transverse views of (maximum or minimum) intensity projections of the PC, TOF and PD data sets. MIDDLE: The associated vesselness measures. RIGHT: Visualizations of the left column, but masked by the binary segmentations obtained by the multi-scale geometric flow. It is important to note the large dynamic range of the intensities of the vessel tree in both the original data and in the vesselness maps of the PC, TOF and PD data make it impossible to use a single threshold to extract the complete vessel tree. In the vesselness maps it is apparent that there are many non-vessel structures that have intensities similar to the vessels. This is the main reason that we use conservative thresholding of the distributed vesselness maps to seed the segmentation process, and then let the geometric flow complete the segmentation.

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.

4.3 Quantitative Results

4.3.1 Single Subject Study

Figure 8 compares the segmentations obtained on the PC, TOF and PD volumes (Figure 7) with one another (a cross-validation) and with the ground truth segmentation obtained by an expert (a validation). Transverse views are shown 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.

From the results in the top two rows (cross-validation) it is clear that most of the 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. The results also indicate that the TOF reconstruction is missing some vessel labels when compared to the PC and PD reconstructions. The results in the bottom two rows (validation) indicate a strong agreement between the segmentation results on all three modalities and those obtained by the expert.

We now carry out a quantitative analysis of these segmentations by computing a number of statistics between each pair of results, treating one as the ground truth data set and the the other as the test data set. These comparisons are shown in Table 3 and include the following measures:

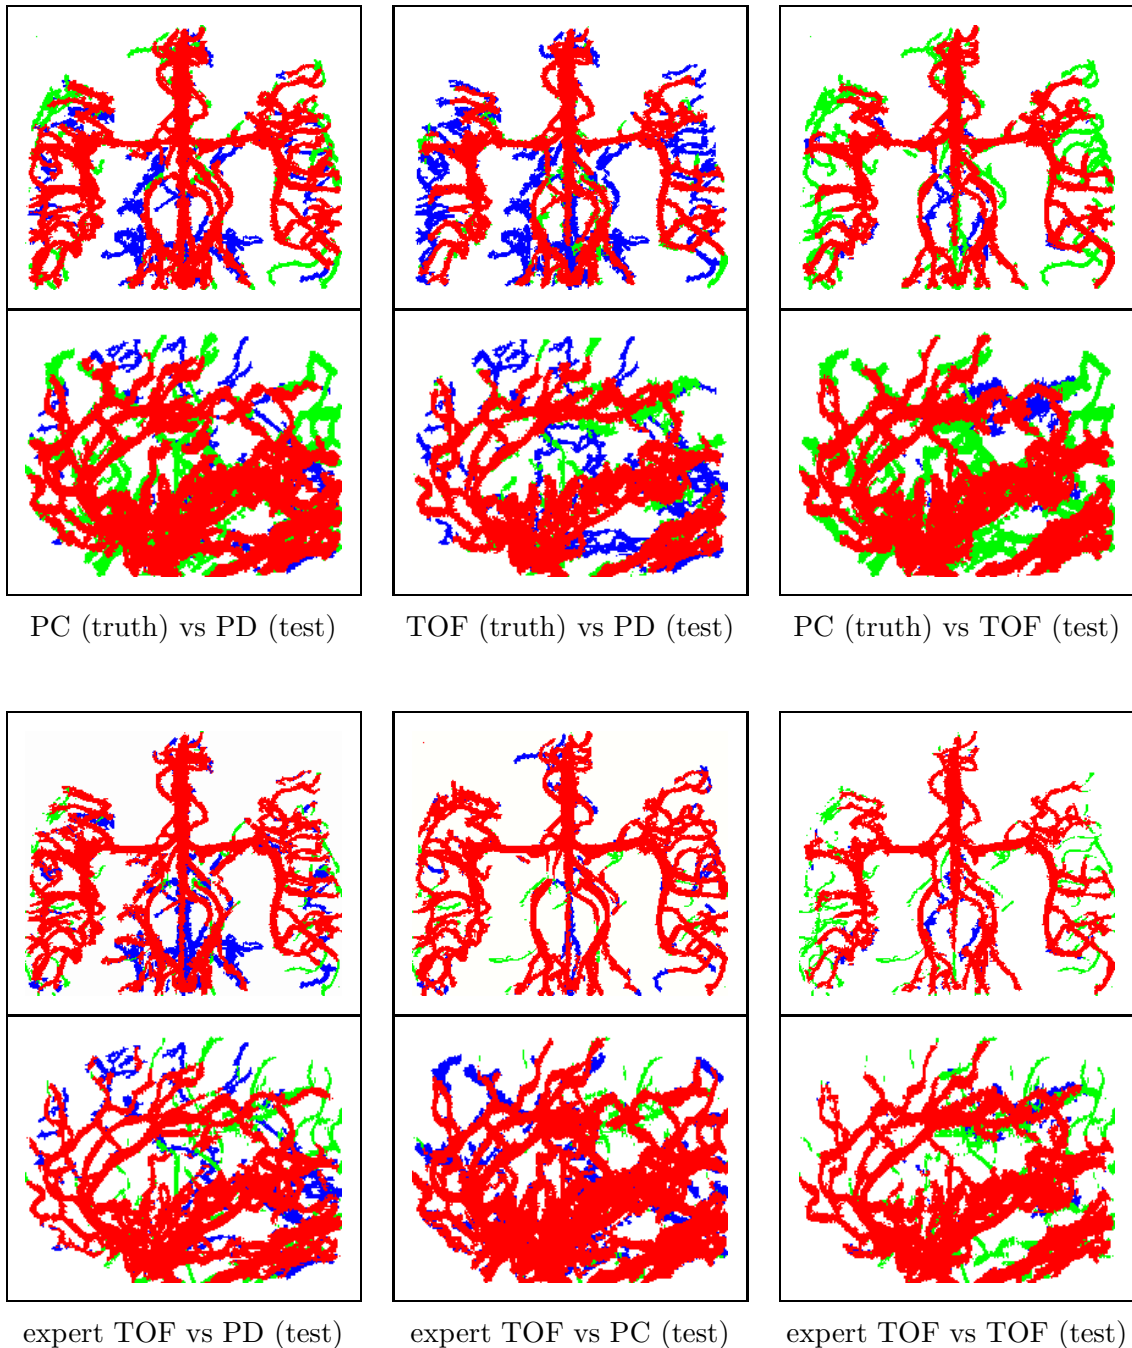


Fig. 8. Each column shows a pair-wise comparison of reconstructions obtained on different modalities, with transverse views in the top row and sagittal views in the bottom row. White labels correspond to the background, red labels to locations common to the ground truth and test data, green labels to locations in the ground truth only and blue labels to locations in the test data only. When more than one label exists along the viewing direction red takes precedence over the other colors and green takes precedence over blue.

Data Sets		Validation Measures			
Ground Truth	Test	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
expert TOF	PD	0.88	0.92	0.75	0.36
expert TOF	PC	0.93	0.97	0.51	0.26
expert TOF	TOF	0.92	0.89	0.38	0.19

Table 3

A pair-wise comparison between segmentation results, treating one case as the ground truth and the other as the test data. Cross-validation results are shown in the top six rows and validation results in the bottom three.

- (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 [40].

- (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 (we essentially test all the red voxels in Figure 8).

It is clear from Table 3 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. Finally there is very good agreement between the automatic segmentation of all three data sets and the ground truth segmentation by an expert. This latter ground truth segmentation took on the order of 7 hours for the expert to provide and hence it was infeasible to incorporate segmentations from additional experts for this study. The potential concern about the quality of the ground truth segmentation is mitigated in part by the results of the multiple subject study, which we describe next.

4.3.2 Multiple Subject Study

Table 4 shows pairwise comparisons between the automatic segmentations obtained on PD volumes with the corresponding ground truth data sets for 19 subjects from the digital brain phantom [17], using a distance tolerance threshold of 3 voxels (1.5mm) for correspondence. These results indicate that 92% of the ground truth vessel voxels are also included in the PD segmentation, with an average alignment error that is small (0.55 mm). Part of this misalignment is due to differing geometric distortions in the PD and angiography scans. Furthermore, the average kappa coefficient is 0.79. We also note that in this study on average only 71% of the vessel voxels obtained in the PD volume are also present in the ground truth data. One possible interpretation is that there are more true vessel voxels in the PD segmentation. This appears to be the case in qualitative manual evaluation of the 19 subjects. It can be explained by the fact that TOF angiography is only sensitive to blood flowing at certain speeds.

The pairwise comparison results for the multi-subject study are supplemented by the plot in Fig. 9 which indicates how the kappa and ratio overlap measures change with varying distance thresholds. This plot was generated by computing the average and standard deviation of the kappa and ratio measures for the 19 subjects. As expected, as the tolerance increases, the kappa and ratio values approach 1 and 100% respectively. It is also expected that for low or no tolerance in alignment error, the average kappa and agreement ratio should be low as well. These curves have an almost constant slope between 0 and 3 voxels (0 and 1.5 mm) and then become flat. Therefore, choosing 3 voxels as a tolerance threshold for quantitative cross-validation is reasonable to allow for small alignment errors due to geometric distortions between the PD and angiographic acquisitions.

Validation Measures	Subjects (1-10)									
	1	2	3	4	5	6	7	8	9	10
ratio	0.97	0.92	0.91	0.90	0.96	0.92	0.91	0.94	0.93	0.92
kappa	0.81	0.78	0.79	0.82	0.79	0.84	0.81	0.89	0.83	0.73
alignment (voxels)	0.82	1.25	1.21	1.21	0.96	1.21	1.06	1.06	1.09	1.24
alignment (mm)	0.41	0.63	0.61	0.61	0.48	0.61	0.53	0.53	0.55	0.62
	Subjects (11-19)									
	11	12	13	14	15	16	17	18	19	Avg.
ratio	0.87	0.91	0.96	0.97	0.95	0.92	0.87	0.84	0.95	.92
kappa	0.80	0.81	0.70	0.76	0.75	0.83	0.76	0.80	0.80	.79
alignment (voxels)	1.25	1.11	0.95	0.81	0.92	1.11	1.21	1.39	1.03	1.10
alignment (mm)	0.63	0.56	0.48	0.41	0.46	0.56	0.61	0.70	0.52	.55

Table 4

Pair-wise comparisons between the 19 PD segmentations and the corresponding ground truth segmentation. The average ratio, kappa and alignment measures are shown in the right column in the bottom half of the table.

5 Conclusion and Discussion

We have developed and validated a geometric flow for segmenting vasculature in full brain PD-weighted 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. The qualitative results indicate that a significant amount of vasculature can be recovered by initializing the flow using a few isolated seeds (at least on an MRA volume). 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. 4). In our experience a clinician or researcher familiar with our software, can do this fairly efficiently when focusing on a small region of interest. We have easy-to-use visualization

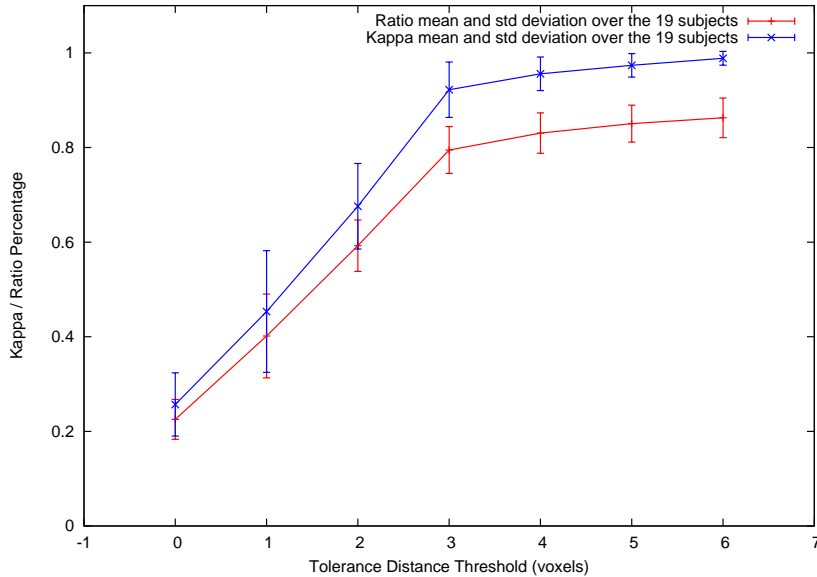


Fig. 9. An examination of the effect of the distance tolerance threshold on the kappa and ratio overlap measures for the multi-subject study. The kappa and ratio measures increase linearly till a tolerance distance of 3 voxels, after which they plateau.

tools that allow one to crop regions, place seeds with a few mouse clicks and re-run the flow on the region from any previously computed intermediate segmentation volume.

We have also 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 7 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. 7. 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.

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 a single subject as well as multi-subject data from the recent digital brain phantom of [17]. The results indicate that for both studies, 90% or more of the vasculature in the ground truth segmentation is recovered from the automatic segmentation of the other volumes, with a small alignment error. We observe also that a significant portion of the vasculature obtained from the PD data is 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 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 [5,21] 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. It could also be interesting to extend the method to handle stenoses as well as branching topologies by incorporating some of the modeling ideas exploited in other surface evolution methods [9,10].

At present, our algorithm is being evaluated as an MR image preprocessing tool for neurosurgical planning, to segment and visualize deep vessels that are used as landmarks during surgery. The procedure reduces the amount of time required for manual preparation and structure segmentation. In addition, the vessel segmentation procedure is being used in a study of *vessel driven brain shift correction* at the MNI, where pre-operative MR images and intra-operative Doppler ultrasound data serve as features in a non-linear registration

paradigm to correct for brain-shift during surgery [41].

Acknowledgments This work was supported by grants from NSERC, FQRNT, CFI and CIHR. We thank Bruce Pike, Simon Drouin and Ingerid Reinertsen for helpful discussions.

References

- [1] A. Frangi, W. Niessen, K. L. Vincken, M. A. Viergever, Multiscale vessel enhancement filtering, in: International Conference On Medical Image Computing and Computer Assisted Intervention, Vol. LNCS 1496, 1998, pp. 130–137.
- [2] A. Vasilevskiy, K. Siddiqi, Flux maximizing geometric flows, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (12) (2002) 1565–1578.
- [3] D. L. Wilson, A. Noble, Segmentation of cerebral vessels and aneurysms from MR angiography data, in: Information Processing in Medical Imaging, 1997, pp. 423–428.
- [4] T. M. Koller, G. Gerig, G. Székely, D. Dettwiler, Multiscale detection of curvilinear structures in 2-d and 3-d image data, in: International Conference On Computer Vision, 1995, pp. 864–869.
- [5] S. R. Aylward, E. Bullitt, Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction, *IEEE Transactions On Medical Imaging* 21 (2) (2002) 61–75.
- [6] O. Wink, M. A. V. Wiro J. Niessen, Multiscale vessel tracking, *IEEE Transactions on Medical Imaging* 23 (1) (2004) 130–133.
- [7] T. McInerney, D. Terzopoulos, Topology adaptive deformable surfaces for medical image volume segmentation, *IEEE Transactions on Medical Imaging* 18 (10) (1999) 840–850.
- [8] L. M. Lorigo, O. D. Faugeras, E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, C.-F. Westin, Curves: Curve evolution for vessel segmentation, *Medical Image Analysis* 5 (2001) 195–206.
- [9] C. M. van Bommel, M. A. Viergever, W. J. Niessen, Semiautomatic segmentation and stenosis quantification of 3d contrast-enhanced mr angiograms of the internal carotid artery, *Magnetic Resonance in Medicine* 51 (2004) 753–760.
- [10] R. Manniesing, M. A. Viergever, W. J. Niessen, Vessel axis tracking using topology constrained surface evolution, *IEEE Transactions on Medical Imaging* 26 (3) (2007) 309–316.

- [11] B. K. Rutt, S. Napel, Magnetic resonance techniques for blood-flow measurement and vascular imaging, *Journal of the Canadian Association of Radiologists* 42 (1) (1991) 21–30.
- [12] J. E. Siebert, J. R. Pernicone, E. J. Potchen, Physical principles and application of magnetic resonance angiography, *Seminars in Ultrasound, CT and MR* 13 (4) (1992) 227–245.
- [13] J. Mazziotta, A. Toga, A. Evans, P. Fox, J. Lancaster, K. Zilles, R. Woods, T. Paus, G. Simpson, B. Pike, C. Holmes, L. Collins, P. Thompson, D. MacDonald, M. Iacoboni, T. Schormann, K. Amunts, N. Palomero-Gallagher, S. Geyer, L. Parsons, K. Narr, N. Kabani, G. Le Goualher, D. Boomsma, T. Cannon, R. Kawashima, B. Mazoyer, A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (icbm), *Philos Trans R Soc Lond B Biol Sci* 356 (1412) (2001) 1293–322.
- [14] A. Evans, the Brain Development Cooperative Group, The nih mri study of normal brain development, *Neuroimage* 30 (1) (2006) 184–202.
- [15] L. R. Ostergaard, O. V. Larsen, G. L. Goualher, A. C. Evans, D. L. Collins, Extraction of cerebral vasculature from MRI, in: 9th Danish Conference on Pattern Recognition and Image Analysis, 2000.
- [16] R. Kimmel, A. M. Bruckstein, On regularized laplacian zero crossings and other optimal edge integrators, *International Journal of Computer Vision* 53 (3) (2003) 225–243.
- [17] B. Aubert-Broche, M. Griffin, G. B. Pike, A. C. Evans, D. L. Collins, Twenty new digital brain phantoms for creation of validation image data bases, *IEEE Transactions on Medical Imaging* 25 (11) (2006) 1410–1416.
- [18] C. Lorenz, I. Carlsen, T. Buzug, C. Fassnacht, 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*, Vol. 1205, 1997, pp. 233–242.
- [19] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, R. Kikinis, 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images, *Medical Image Analysis* 2 (2) (1998) 143–168.
- [20] K. Krissian, G. Malandain, N. Ayache, Model-based detection of tubular structures in 3D images, *Computer Vision and Image Understanding* 80 (2) (2000) 130–171.
- [21] K. Krissian, J. Ellsmere, K. Vosburgh, R. Kikinis, C.-F. Westin, Multiscale segmentation of the aorta in 3D ultrasound images, in: *Engineering in Medicine and Biology Society*, 2003, pp. 638–641.
- [22] J. A. Tyrell, E. di Tomaso, D. Fuja, R. Tong, K. Kozak, R. K. Jain, B. Roysam, Robust 3-d modeling of vasculature imagery using superellipsoids, *IEEE Transactions on Medical Imaging* 26 (2) (2007) 223–237.

- [23] T. Deschamps, L. D. Cohen, Fast Extraction of Minimal Paths in 3D Images and Applications to Virtual Endoscopy, *Medical Image Analysis* 5 (4) (2001) 281–299.
- [24] M. Sofka, C. V. Stewart, Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures, *IEEE Transactions on Medical Imaging* 25 (12) (2006) 1531–1546.
- [25] M. Kass, A. Witkin, D. Terzopoulos, Snakes: Active contour models, *International Journal of Computer Vision* 1 (1987) 321–331.
- [26] T. McInerney, D. Terzopoulos, T-snakes: Topology adaptive snakes, *Medical Image Analysis* 4 (2000) 73–91.
- [27] L. Ambrosio, H. M. Soner, Level set approach to mean curvature flow in arbitrary codimension, *Journal of Differential Geometry* 43 (1996) 693–737.
- [28] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Transactions On Pattern Analysis and Machine Intelligence* 12 (1990) 629–639.
- [29] T. Lindeberg, Edge detection and ridge detection with automatic scale selection, *International Journal of Computer Vision* 30 (2) (1998) 77–116.
- [30] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, A. Yezzi, Gradient flows and geometric active contour models, in: *International Conference On Computer Vision*, 1995, pp. 810–815.
- [31] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, in: *International Conference On Computer Vision*, 1995, pp. 694–699.
- [32] K. Siddiqi, Y. B. Lauzière, A. Tannenbaum, S. W. Zucker, Area and length minimizing flows for shape segmentation, *IEEE Transactions on Image Processing* 7 (3) (1998) 433–443.
- [33] M. Gage, R. Hamilton, The heat equation shrinking convex plane curves, *Journal of Differential Geometry* 23 (1986) 69–96.
- [34] M. Grayson, The heat equation shrinks embedded plane curves to round points, *Journal of Differential Geometry* 26 (1987) 285–314.
- [35] L. Alvarez, F. Guichard, P. L. Lions, J. M. Morel, Axiomes et équations fondamentales du traitement d’images, *C. R. Acad. Sci. Paris* 315 (1992) 135–138.
- [36] B. B. Kimia, A. Tannenbaum, 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 (1995) 189–224.
- [37] S. J. Osher, J. A. Sethian, Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations, *Journal of Computational Physics* 79 (1988) 12–49.

- [38] D. L. Collins, P. Neelin, T. M. Peters, A. C. Evans, Automatic 3d intersubject registration of mr volumetric data in standardized talairach space, *J Comput Assist Tomogr* 18 (2) (1994) 192–205.
- [39] M. Descoteaux, L. Collins, 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)*, Vol. LNCS 3117, 2004, pp. 169–180.
- [40] L. R. Dice, Measures of the amount of ecologic association between species, *Ecology* 26 (3) (1945) 297–302.
- [41] I. Reinertsen, M. Descoteaux, S. Drouin, K.Siddiqi, L. Collins, Vessel driven correction of brain shift, in: *International Conference On Medical Image Computing and Computer Assisted Intervention*, Vol. LNCS 3217(2), 2004, pp. 208–216.