

Geometric Flows for Segmenting Vasculature in MRI: Theory and Validation

Maxime Descoteaux¹, Louis Collins², and Kaleem Siddiqi¹

McGill University, Montréal, QC, Canada

¹School of Computer Science & Centre For Intelligent Machines

{mdesco, siddiqi}@cim.mcgill.ca

²McConnell Brain Imaging Centre, Montréal Neurological Institute

{louis}@bic.mni.mcgill.ca

Abstract. Often in neurosurgical planning a dual spin 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 noninvasive. 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 standard MRI which can also be applied to the easier cases of angiography data. We first apply Frangi’s vesselness measure [3] to find putative centerlines of tubular structures along with their estimated radii. This measure is then distributed to create a vector field which is orthogonal to vessel boundaries so that the flux maximizing geometric flow algorithm of [14] can be used to recover them. We perform a quantitative cross validation on PD, phase contrast (PC) angiography and time of flight (TOF) angiography volumes, all obtained for the same subject. A significant finding is that whereas more than 80% of the vasculature recovered from the angiographic data is also recovered from the PD volume, over 25% of the vasculature recovered from the PD volume is *not* detected in the TOF data. Thus, the technique can be used not only to improve upon the results obtained from angiographic data but also as an alternative when such data is not available.

1 Introduction

A three-dimensional (3D) representation of vasculature can be extremely important in image-guided neurosurgery, pre-surgical planning 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 it manually, coloring regions of interest and connecting them using image processing operations. This process is extremely laborious, is prone to human error and makes large scale clinical studies of vasculature infeasible. In computer vision there has been a significant amount of work towards automating the extraction of vessels or vessel centerlines, typical examples of which include [10, 6, 5, 1]. However, most of these methods have been developed for 2D projection angiography or 3D CT and MR angiography, several of which require the injection of contrast agents. To our knowledge no method currently exists for the automatic extraction of vessel boundaries in standard



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. The spaghetti-like structures in the PD correspond to vasculature, which is more easily detected in the angiograms where there are sharp bright to dark contrast changes *only* at vessel boundaries.

MRI volumes, such as the proton density (PD) image in Figure 1. Here it is clear that a signal decrease is present in the vascular regions (the spaghetti-like structures), but the contrast between blood vessels and surrounding tissue is not limited to vessel boundaries as it is in the corresponding angiographic sequences. The problem of recovering vessels from image intensity contrast alone on PD 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.

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 volumes. The algorithm is motivated in part by the approach in [12] where Frangi’s vesselness measure [3] is thresholded to find centerlines. However, rather than threshold this 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 [14] to be applied to recover vessel boundaries. This approach allows for a type of local integration, while affording the ability to connect branching structures. The flow has a formal motivation, is topologically adaptive due to its implementation using level set methods, and finally is computationally efficient. We perform a quantitative comparison of segmentations from PD, PC and TOF volumes, all obtained for the same subject (Figure 1).

2 A Multi-Scale Geometric Flow For Segmenting Vasculature

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 $\lambda_1, \lambda_2, \lambda_3$ of the Hessian matrix \mathbf{H} [8, 13, 3]. These methods exploit the fact that at locations centered within tubular structures the 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 being equal because the cross-section of the vessel is approximately circular. The corresponding Eigen vectors span the vessel direction and the cross-sectional plane.




Fig. 2. (1) The superimposed vectors on the synthetic helix indicate the estimated vessel orientation at locations where the multi-scale vesselness measure (Eq. 1) is high. (2) 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. 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.

We choose to focus here on Frangi’s vesselness measure [3] because it incorporates information from all three Eigen values. With the Eigen values sorted such that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, three quantities are defined to differentiate blood vessels from other structures: $R_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2 \lambda_3|}}$, $R_A = \frac{|\lambda_2|}{|\lambda_3|}$, and $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$. R_B is non zero only for blob-like structures, the R_A ratio differentiates sheet-like objects from other structures and 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 $G(\sigma)$ at that scale¹. The vesselness measure is then defined by

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

This measure is designed to be maximum along the centerlines of tubular structures and close to zero outside vessel-like regions.

In our implementation we set the parameters α , β and c to 0.5, 0.5 and half the maximum Frobenius norm, respectively, as suggested in [3]. 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 (left) for a synthetic helix. The grey surface coincides with a particular level set of the vesselness measure. Within this surface locations of high vesselness are indicated by overlaying the Eigen vectors which correspond to the estimated vessel orientation.

¹ In practice we directly compute the entries which comprise the Hessian matrix by using derivatives of Lindeberg’s γ -parametrized normalized Gaussian kernels [7]. This allows us to compare responses at different scales.

2.2 Extending the vesselness measure to vessel boundaries

We now construct a vector field which is both large in magnitude at vessel boundaries as well as orthogonal to them, so that the flux maximizing flow of [14] can be applied. Since the vesselness measure is concentrated at centerlines, we need to first distribute it to the vessel boundaries which are implied by the local orientation and scale. At each voxel (x, y, z) where this measure is a local maximum in a $3 \times 3 \times 3$ neighborhood we consider an ellipsoid with its major axis oriented along the estimated orientation and its two other axes equal to the estimated radius. 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, as illustrated in Figure 2 (right). 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 a scalar function ϕ . The extended vector field is given by the product of the normalized gradient of the original intensity image with ϕ :

$$\vec{\mathcal{V}} = \phi \frac{\nabla I}{|\nabla I|}. \quad (2)$$

2.3 The multi-scale geometric flow

We now apply the flux maximizing geometric flow of [14], which is the gradient flow that 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. With S an evolving surface and $\vec{\mathcal{V}}$ the vector field, this flow is given by

$$\mathcal{S}_t = \text{div}(\vec{\mathcal{V}}) \vec{\mathcal{N}} \quad (3)$$

where $\vec{\mathcal{N}}$ is the unit outward 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 their original formulations, the above flow and the related co-dimension 2 flow developed [9] are both quite restrictive since they are designed specifically for angiographic data. They are both initialized essentially by thresholding such data, and thus would fail when vasculature cannot be identified from contrast alone. Furthermore, neither has an explicit term to model tubular structures. Rather, each flow relies on the assumption that the gradient of the intensity image yields a quantity that is significant *only* at vessel boundaries. Finally, neither of these methods takes into account explicitly the multi-scale nature of vessels boundaries as they appear in all modalities.

With regard to the flux maximizing flow, these limitations can be overcome by choosing as the static vector field in Eq. 3 the construction defined in Eq. 2. This vector field embodies two important constraints. First, the magnitude of ϕ is maximum on vessel boundaries and the ellipsoidal extension performs a type of local integration. Second, $\frac{\nabla I}{|\nabla I|}$ captures the direction information of the gradient, which is expected to be high at boundaries of vessels as well as orthogonal to them². The surface evolution equation

² 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. Without this normalization structures such as white and gray matter boundaries could get significant unwanted contributions.

works out to be

$$S_t = \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}} = \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \text{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right) \right] \vec{\mathcal{N}} = \left[\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \kappa_{\mathcal{I}} \right] \vec{\mathcal{N}}, \quad (4)$$

where $\kappa_{\mathcal{I}}$ is the Euclidean mean curvature of the iso-intensity level set of the image. This is a hyperbolic partial differential equation since all terms depend solely on the vector field and not on the evolving surface. The first term $\left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle$ acts like a doublet because at vessel boundaries $\nabla \phi$ has a zero-crossing while $\nabla \mathcal{I}$ does not change sign. Hence, when the evolving surface overshoots the boundary slightly, this term acts to push it back towards the boundary. The second term behaves like a geometric heat equation [4] since $\kappa_{\mathcal{I}}$ is the mean curvature of the iso-intensity level set of the original intensity image. The flow cannot leak in regions outside vessels since by construction both ϕ and $\nabla \phi$ are zero there.

In our implementation we compute the Hessian operator over 10 log scales and select the maximum vesselness response, using Jacobi's method for symmetric matrices to find the Hessian's Eigen values. The geometric flow in Eq. 4 is then simulated using a first-order in time discretized level-set version of the flow [11]:

$$\Psi_n = \Psi_{n-1} - \Delta t * \mathcal{F} * ||\nabla \Psi_{n-1}||. \quad (5)$$

Here $\mathcal{F} = \left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \text{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right)$, Ψ is the embedding hypersurface typically initialized as a signed Euclidean distance function to the initial seed boundaries, Δt is the step size and the evolving surface S is obtained as the zero level set of the Ψ function.

3 Experiments

3.1 Image acquisition

In order to validate our algorithm we acquired four different volumes from the same subject on a Siemens 1.5 Tesla system at the Montréal Neurological Institute. 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 (Figure 1).

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/background contrast is greatest for the PC data (white on black tissue), intermediate for the PD data (black on grey) and slightly less for the TOF data (white on grey).

3.2 Vessel extraction

Owing to the high resolution of the scans, the volumes lead to significant memory costs when processed in their entirety. Hence we chose to work with a common $259 \times 217 \times 170$ voxel region cropped from each volume, which had vessels of different widths and contrasts in the three modalities. The segmentation process was carried out by computing the vesselness measure, constructing the extended vector field \vec{V} for a selected vesselness threshold and finally applying the geometric flow (Eq. 4), initializing it by placing seeds at locations where the vesselness measure was above a high threshold so that the evolution began from within vessels.

Figure 3 compares the reconstructions obtained on the three different modalities, with transverse views shown in the top row and sagittal views in the bottom row. To allow for slight alignment errors due to the process by which the initial volumes were scaled 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 column 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 column 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. The second and third column of Figure 3 indicate that the TOF reconstruction is missing a large number of vessel labels when compared to the PC and PD reconstructions.

3.3 Cross validation

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 other as the test data set. These comparisons are shown in Table 1 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[2].
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 (for the red voxels). This measure also indicates the degree to which the test data explains the ground truth data, but in terms of an average distance error.

It is clear from Table 1 that the vasculature obtained from the PD volume accounts for more than 80% of that obtained from either of the angiographic sequences. Furthermore a significant proportion (over 25%) of vessel voxels recovered from the PD and PC volume are not seen in the TOF angiographic sequence. The results also indicate that the segmentations are closely aligned, with an average alignment error less than 0.5mm.




Fig. 3. 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.

4 Conclusions

We have presented what to our knowledge is the first multi-scale geometric flow that can be applied for segmenting vasculature in proton density weighted MRI volumes. The key idea is to incorporate a multi-scale vesselness measure in the construction of an appropriate vector field for a geometric flow. We have performed a quantitative cross validation of the flow by applying the algorithm on PD, PC and TOF volumes obtained for the same subject and computing several measures. The results indicate that the vessels segmented from the PD data alone account for 80% and 89% of the vasculature segmented from PC and TOF angiographic data sets respectively, but also that over 25% of the vasculature obtained from the PD data are not recovered from either of the angiographic volumes. This suggests that our algorithm can be used to both improve upon the results obtained from angiographic data but also as a promising alternative when such data is not available, since the PD-weighted MRI data are often acquired when planning brain tumour surgery.

References

1. E. Bullitt, S. Aylward, A. Liu, J. Stone, S. K. Mukherjee, C. Coffey, G. Gerig, and S. M. Pizer. 3d graph description of the intracerebral vasculature from segmented mra and tests of accuracy by comparison with x-ray angiograms. In *Information Processing in Medical Imaging*, pages 308–321, 1999.

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 1. A pair-wise comparison between the different modalities, treating one as the ground truth and the other as the test data. See the text for an explanation of the measures computed.

2. L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.
3. A. Frangi, W. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI'98*, pages 130–137, 1998.
4. M. Grayson. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry*, 26:285–314, 1987.
5. 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.
6. 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.
7. T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):77–116, 1998.
8. 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.
9. 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.
10. T. McInerney and D. Terzopoulos. T-snakes: Topology adaptive snakes. *Medical Image Analysis*, 4:73–91, 2000.
11. 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.
12. 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*, 2000.
13. 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.
14. A. Vasilevskyi and K. Siddiqi. Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(12):1–14, 2002.