

Co-Dimension 2 Geodesic Active Contours for MRA Segmentation

Liana M. Lorigo¹, Olivier Faugeras^{1,2}, W.E.L. Grimson¹, Renaud Keriven³,
Ron Kikinis⁴, Carl-Fredrik Westin⁴

¹ MIT Artificial Intelligence Laboratory, Cambridge MA, USA
liana@ai.mit.edu

² INRIA, Sophia Antipolis, France

³ Cermics, ENPC, France

⁴ Harvard Medical School, Brigham & Women’s Hospital, Boston MA, USA

Abstract. Automatic and semi-automatic magnetic resonance angiography (MRA) segmentation techniques can potentially save radiologists large amounts of time required for manual segmentation and can facilitate further data analysis. The proposed MRA segmentation method uses a mathematical modeling technique which is well-suited to the complicated curve-like structure of blood vessels. We define the segmentation task as an energy minimization over all 3D curves and use a level set method to search for a solution. Our approach is an extension of previous level set segmentation techniques to higher co-dimension.

1 Introduction

The high-level goal of this research is to develop computer vision techniques for the segmentation of medical images. Automatic and semi-automatic vision techniques can potentially assist clinicians in this task, saving them much of the time required to manually segment large data sets. Specifically, we consider the segmentation of volumetric vasculature images, such as the magnetic resonance angiography (MRA) image pictured in Figure 1.

As shown here, blood vessels appear in MRA images as bright curve-like patterns which may be noisy and have gaps. What is shown is a “maximum intensity projection”. The data is a stack of slices where most areas are dark, but vessels tend to be bright. This stack is collapsed into a single image for viewing by performing a projection through the stack that assigns to each pixel in the projection the brightest voxel over all slices. This image shows projections along three orthogonal axes.

Thresholding is one possible approach to this segmentation problem and works adequately on the larger vessels. The problem arises in detecting the small vessels, and that is the objective of our work. Thresholding cannot be used for the small vessels for several reasons. The voxels may have an intensity that is a combination of the intensities of vessels and background if the vessel is only partially inside the voxel. This sampling artifact is called *partial voluming*. Other imaging conditions can cause some background areas to be as bright as other

vessel areas, complicating threshold selection. Finally, the images are often noisy, and methods using local contextual information can be more robust.

Our method uses the fact that the underlying structures in the image are indeed 3D curves and evolves an initial curve into the curves in the data (the vessels). In particular, we explore techniques based on the concept of *mean curvature flow*, or *curve-shortening flow*, from the field of differential geometry.

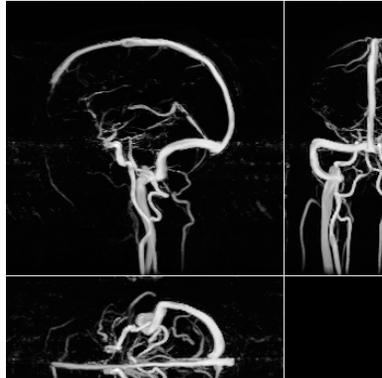


Fig. 1. Maximum intensity projection of a phase-contrast MRA image of blood vessels in the brain.

2 Curvature Evolution Methods

Mean curvature evolution schemes for segmentation, implemented with level set methods, have become an important approach in computer vision [5, 10, 11]. This approach uses partial differential equations to control the evolution. An overview to the superset of techniques using related partial differential equations can be found in [4]. The fundamental concepts from mathematics from which mean curvature schemes derive were explored several years earlier when smooth closed curves in 2D were proven to shrink to a point under mean curvature motion [8, 9]. Evans and Spruck and Chen, Giga, and Goto independently framed mean curvature flow of any hypersurface as a level set problem and proved existence, uniqueness, and stability of viscosity solutions [7, 6]. For application to image segmentation, a vector field was induced on the embedding space, so that the evolution could be controlled by an image gradient field or other image data. The same results of existence, uniqueness, and stability of viscosity solutions were obtained for the modified evolution equations for the case of planar curves, and experiments on real-world images demonstrated the effectiveness of the approach [3, 5].

Curves evolving in the plane became surfaces evolving in space, called *minimal surfaces* [5]. Although the theorem on planar curves shrinking to a point

could not be extended to the case of surfaces evolving in 3D, the existence, uniqueness, and stability results of the level set formalism held analogously to the 2D case. Thus the method was feasible for evolving both curves in 2D and surfaces in 3D. Beyond elegant mathematics, spectacular results on real-world data sets established the method as an important segmentation tool in both domains. One fundamental limitation to these schemes has been that they describe only the flow of hypersurfaces, i.e., surfaces of co-dimension 1.

Altschuler and Grayson studied the problem of curve-shortening flow for 3D curves [1], and Ambrosio and Soner generalized the level set technique to arbitrary manifolds in arbitrary dimension. They provided the analogous results and extended their level set evolution equation to account for an additional vector field induced on the space [2].

We herein present the first implementation of geodesic active contours in 3D, based on Ambrosio and Soner’s work. Specifically, our system uses these techniques for automatic segmentation of blood vessels in MRA images. The dimension of the manifold is 1, and its co-dimension is 2.

3 Mean Curvature Flow

Intuitively, *mean curvature flow* refers to some curve evolving in time so that at each point, the velocity vector normal to the curve is equal to the mean curvature vector. This concept is normally defined for arbitrary generic surfaces, but only curves are necessary for this paper, so we have restricted the definition. More formally, let $C(t)$, $t \geq 0$ be a family of curves in \mathbb{R}^2 or \mathbb{R}^3 , \mathbf{N} the normal for a given orientation. That is, C is a curve, and t represents the “time” parameter or the index into the family of curves, not position. The mean curvature flow equation is then given by the vector equation

$$C_t = \kappa \mathbf{N} \tag{1}$$

with given initial curve $C(0) = C_0$, κ the curvature of the curve, and C_t the time derivative of the curve. Note that since we consider only 1D curves here, as opposed to evolving surfaces, the mean curvature is just the usual curvature of the curve. This motion is also called “curve-shortening flow” since it is the solution, obtained by Euler-Lagrange equations, to the problem of minimizing curve length:

$$\min_C \int |C'(p)| dp$$

where p is the spatial parameter of the curve.

4 Level Set Method for Planar Curves

We give the basic idea of the level set method [12] to evolve a *planar* curve C . Define a function $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ so that C is a level-set of u . We follow the convention that C is, in particular, the zero level set of u , although this choice is not

necessary for the method. The function u is now an implicit representation of the curve C . The advantages of this representation are that it is intrinsic (independent of parameterization) and that it is topologically flexible since different topologies of C are represented by the constant topology of u . Let C_0 be the initial curve.

It is shown in [7] and [6] that evolving C according to

$$C_t = \beta \mathbf{N} \quad (2)$$

with initial condition $C(\cdot, 0) = C_0(\cdot)$ for any function β , is equivalent to evolving u according to

$$u_t = \beta |\nabla u| \quad (3)$$

with initial condition $u(\cdot, 0) = u_0(\cdot)$ and $u_0(C_0) = 0$.

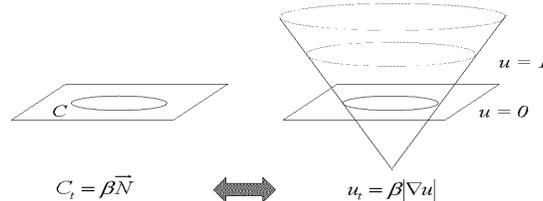


Fig. 2. Level sets of an embedding function u , for a closed curve in \mathbb{R}^2 .

This result is independent of the choice of function u [7, 6]. As customary in the literature, we choose u_0 to be the signed distance function to the curve C (Figure 2).

5 Level Set Method for Curves in Higher Codimension

The level set evolution equations that follow were proven in [2]. They enable us to evolve space curves, with evolution driven by both mean curvature and image information. In the following discussion, C is a curve in 3D.

5.1 Mean Curvature Flow

Let $v : \mathbb{R}^3 \rightarrow [0, \infty)$ be an auxiliary function whose zero level set is identically C , that is smooth near C , and such that ∇v is non-zero outside C . For a nonzero vector $\mathbf{q} \in \mathbb{R}^n$, define

$$P_{\mathbf{q}} = I - \frac{\mathbf{q}\mathbf{q}^T}{|\mathbf{q}|^2}$$

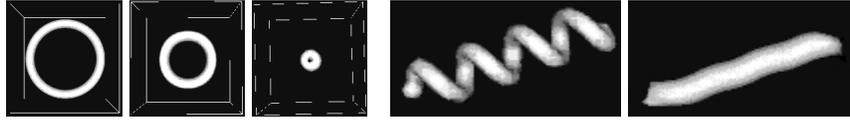


Fig. 3. Evolving curves under mean curvature flow. The first three images show a circle shrinking to a point, and the last two images show a helix shrinking to its axis.

as the projector onto the plane normal to \mathbf{q} . Further define $\lambda(\nabla v(x, t), \nabla^2 v(x, t))$ as the smaller nonzero eigenvalue of $P_{\nabla v} \nabla^2 v P_{\nabla v}$. The level set evolution equation is then

$$v_t = \lambda(\nabla v(x, t), \nabla^2 v(x, t)).$$

That is, this evolution is equivalent to evolving C according to $C_t = \kappa \mathbf{N}$ in the sense that C is the zero level set of v throughout the evolution.

Figure 3 demonstrates this evolution. As discussed above, a circle shrinks to a point under mean curvature motion. Under this motion, a helix evolves into its axis.

5.2 Incorporation of Vector Field

This section discusses the situation where there is an underlying vector field driving the evolution, in combination with the curvature term. Assume the desired evolution equation is of the form

$$C_t = \kappa \mathbf{N} - \Pi \mathbf{d},$$

where Π is the projection operator onto the normal space of C (which is a vector space of dimension 2) and \mathbf{d} is a given vector field in \mathbb{R}^3 . The evolution equation for the embedding space then becomes

$$v_t = \lambda(\nabla v, \nabla^2 v) + \nabla v \cdot \mathbf{d}.$$

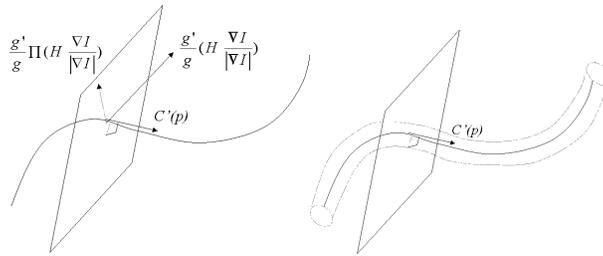


Fig. 4. (a) The tangent to C at p , the normal plane, the image-based vector, and its projection onto the normal plane. (b) ε -level set method.

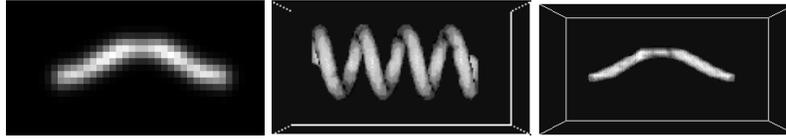


Fig. 5. Evolving helix under mean curvature flow with additional vector field: target curve, initial level set, level set after evolution with endpoints constrained.

5.3 3D Image Segmentation

For the case of 1D structures in 3D images, we wish to minimize

$$\int_0^1 g(|\nabla I(C(p))|) |C'(p)| dp$$

where $C(p) : [0, 1] \rightarrow \mathfrak{R}^3$ is the 1D curve, $I : [0, a] \times [0, b] \times [0, c] \rightarrow [0, \infty)$ is the image, and $g : [0, \infty) \rightarrow \mathfrak{R}^+$ is a strictly decreasing function such that $g(r) \rightarrow 0$ as $r \rightarrow \infty$ (analogous to [5]). For our current implementation, we use $g(r) = \exp(-r)$ because it works well in practice. Another common choice is $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$. By computing the Euler-Lagrange equations, we find that the curve evolution equation is

$$C_t = \kappa \mathbf{N} - \frac{g'}{g} \mathbf{H} \frac{\nabla I}{|\nabla I|}, \quad (4)$$

where \mathbf{H} is the Hessian of the intensity function. The second term in the above equation is illustrated in Figure 4(a). That is,

$$\mathbf{d} = \frac{g'}{g} \mathbf{H} \frac{\nabla I}{|\nabla I|},$$

so the equation for the embedding space is

$$v_t = \lambda(\nabla v(x, t), \nabla^2 v(x, t)) + \frac{g'}{g} \nabla v(x, t) \cdot \mathbf{H} \frac{\nabla I}{|\nabla I|} \quad (5)$$

Thus, Ambrosio and Soner's work has provided the basis for the use of mean curvature flow and level set methods to segment 1D structures in 3D. Figure 5 illustrates how underlying image information can attract the evolving tube. The underlying volumetric image data is shown, as a maximum intensity projection, in the first image. This volume was generated by drawing a cosine curve in the volume, then smoothing with a Gaussian filter. The second image shows the initial curve, a helix. The result of the evolution is shown in the rightmost image.

6 MRA Segmentation System

This section describes our system for segmentation of vessels from MRA using the described level set method. A flowchart is shown in Figure 6. We discuss issues that have arisen in converting the theory above to practice for this application.

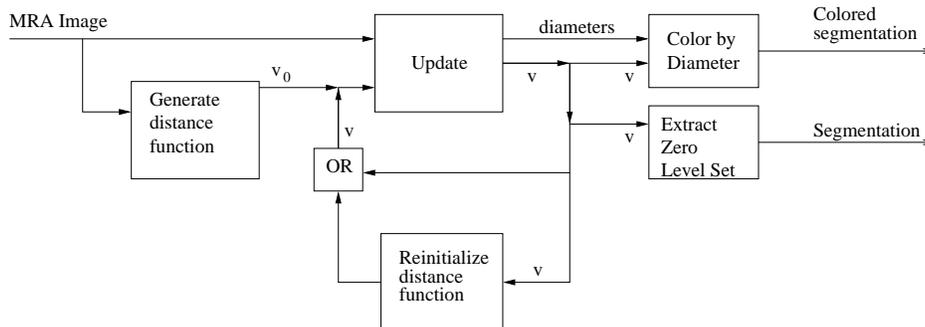


Fig. 6. Overview of segmentation algorithm.

ε -Level Set Method: Since the projection operator $P_{\mathbf{q}}$ is defined only for non-zero vectors q , the method is undefined at $\nabla v = \mathbf{0}$, which is the curve itself, and is numerically unstable near the curve. For this reason, we regard v as a distance function to a “tube” of small radius ε around the curve, instead of extracting the true 1D curve. That is, we evolve the ε -level set instead of evolving the true curve (Figure 4(b)). Note that ε does not denote a fixed value here: we mean simply that the evolving shape is a “tubular” surface of some (unspecified and variable) nonzero width. In addition to being more robust, this method better captures the geometry of blood vessels, which have nonzero diameter.

Banding: Instead of evolving the entire volume, we evolve only the portion of the volume within a narrow band of the zero level set (the current surface). This technique is commonly used in level set methods. Normally, we set the band to include voxels that are up to 6 voxels away from the surface. We have increased this distance up to 12 for some experiments. The advantage of this technique is efficiency, and the disadvantage is that we may miss structures that are outside the band if the potential function g does not have a large enough capture range to attract the segmentation to these structures. This issue can be addressed by ensuring that g is compatible with the band size.

Curvature Instead of Eigenvalues: For computational efficiency and because of numerical instability of the gradient computations and thus the evolution equation near $\nabla v = \mathbf{0}$, we remark that the level sets of the function v flow in the direction of the normal with velocity equal to the sum of their smaller principal curvature and the dot product of ∇v with the image-based vector field \mathbf{d} . Therefore, we compute the smaller curvature directly from v instead of as an eigenvalue of $P_{\nabla v} \nabla^2 v P_{\nabla v}$.

Image Scaling: To control the trade-off between fitting the surface to the image data and enforcing the smoothness constraint on the surface, we add an *image scaling* term *imscale* to Equation 5 to obtain

$$v_t = \lambda(\nabla v(x, t), \nabla^2 v(x, t)) + imscale * \frac{g'}{g} \nabla v(x, t) \cdot \mathbf{H} \frac{\nabla I}{|\nabla I|} \quad (6)$$

imscale is set by the user or can be pre-set to a default value.

Gradient Directionality: Because vessels appear brighter than the background, we weight the image term by the cosine of the angle between the normal to the surface and the gradient in the image. This cosine is given by the dot product of the respective gradients of v and I , so the update equation becomes

$$v_t = \lambda(\nabla v(x, t), \nabla^2 v(x, t)) + imscale * (\nabla v \cdot \nabla I) * \frac{g'}{g} \nabla v(x, t) \cdot \mathbf{H} \frac{\nabla I}{|\nabla I|} \quad (7)$$

For example, if the two vectors point in the same direction, then the brighter region is inside the surface and the darker region is outside; the angle between the vectors is 0, whose cosine is 1, so the image term is fully counted. However, if they point in opposite directions, the negative weighting prevents the evolving vessel walls from being attracted to image gradients that point in the opposite direction.

Reinitializing Volume: As customary in level set segmentation methods, the volume v is periodically reinitialized to be a distance function: the zero level set S is extracted, then each point in the volume is set to be its distance to S . For our implementation, this reinitialization is itself a level set method. To obtain the positive distances, the surface is propagated outward at constant speed of 1, and the distance at each point is determined to be the time at which the surface crossed that point. A second step propagates the surface inward to obtain the negative distances analogously. For some experiments, we have used the Fast Marching Method [12] to implement these steps.

Initial Surface: Figure 7 shows additional detail on the generation of the

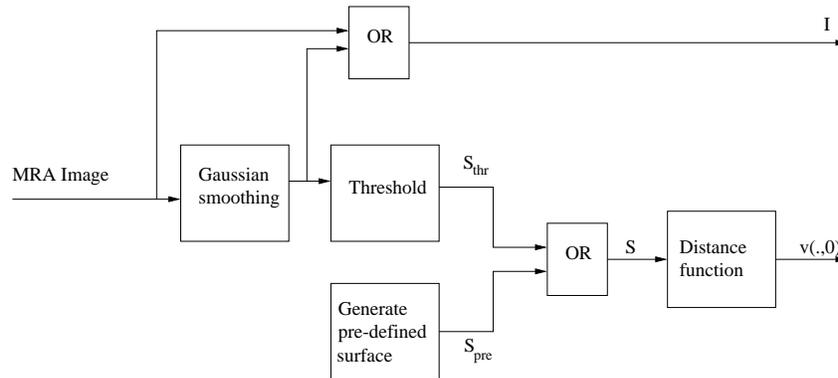


Fig. 7. More detailed illustration of initialization part of algorithm.

initial surface. This initial surface (and thus the initial volume) is normally generated by thresholding the MRA dataset. However, the method does not require that the initial surface be near the target surface but may use any initial surface. Figure 8 illustrates a vertical bar evolving into the segmentation of the first dataset in Figure 10.

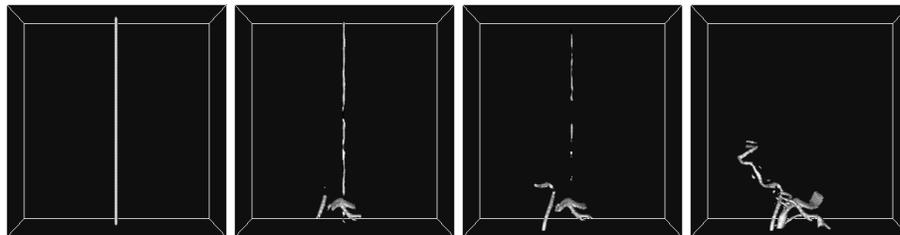


Fig. 8. Illustration of a vertical bar evolving in a segmentation of the first dataset in Figure 10.

Smoothing: As shown in Figure 7, the datasets may be pre-processed to reduce noise. For the results presented here, the raw datasets were convolved with an isotropic Gaussian of $\sigma = 0.5$.

Cleaning: We post-process the segmentations to remove any surface patches whose surface area is less than some threshold (a parameter of the method) to eliminate patches corresponding to noise in the original data.

Vessel Radii Estimation: The larger principal curvature can be useful in measuring the radii of the vessels for a particular application, since radius is the inverse of curvature. This curvature can be easily computed when the smaller principal curvature is computed for the segmentation. We have added the option to color-code our segmentations based on vessel radii, as estimated from the local larger principal curvature of the tubular surface.

7 Results

We demonstrate segmentation results on four datasets, courtesy of the Surgical Planning Laboratory, Brigham and Womens Hospital and Harvard Medical School (Figures 9, 10, and 11). All datasets had an initial resolution of $.9375 \times .9375 \times 1.5 \text{mm}^3$ ($256 \times 256 \times 60$ voxels). The final example only was resampled to $.9375 \times .9375 \times .9375 \text{mm}^3$ ($256 \times 256 \times 96$ voxels) before segmentation; the other segmentations were performed directly on the raw data. The images are not square (256×256) because uninteresting portions were cropped for efficiency. In Figure 10, the initial surface for the segmentation was a surface obtained by thresholding the raw dataset whereas in Figure 9 it was a tube as in Figure 8; *imscale* also varied as discussed below. For comparison, Figure 10 first shows results obtained by thresholding alone. Figure 11 shows an enlargement of a portion of the segmentations and corresponding maximum intensity projection considered in Figure 10.

The following parameters were used in these experiments; all settings were chosen empirically. For our method, *imscale* varied across the datasets depending on the noise present. A threshold t_{init} was used in Figure 10 to obtain the initial surface from the dataset; such a threshold was obviously not needed in

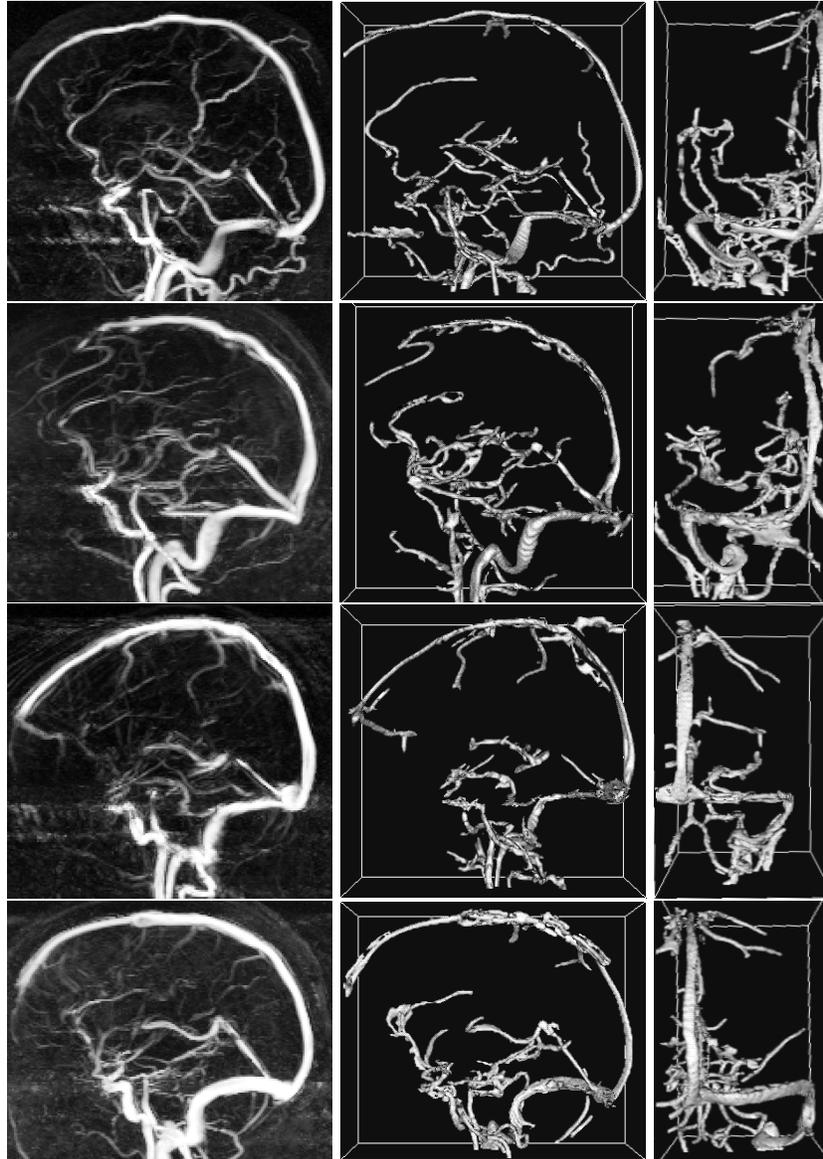


Fig. 9. The first image in each row is the maximum intensity projection of the raw data, and the second and third are the segmentation result from two orthogonal viewpoints. These results are obtained by our method where the initial surface was a vertical bar as showed in Figure 8.

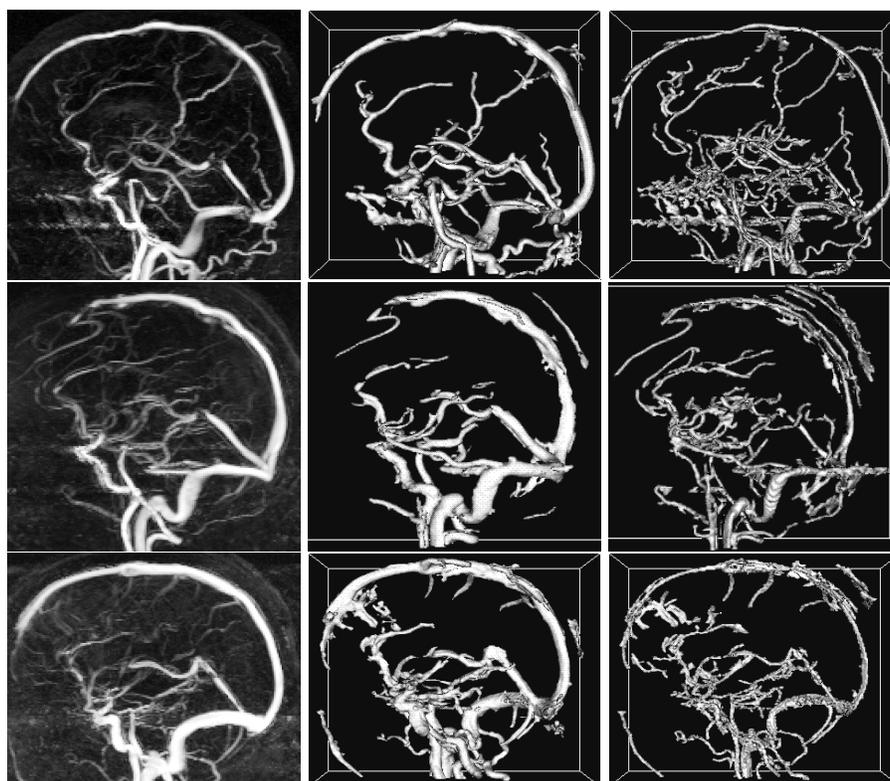


Fig. 10. Results on three datasets are shown. For each image pair, the first image is the maximum intensity projection of the raw data, the second is the segmentation result from thresholding only and the third is the segmentation result using our method.

Figure 9. A cleaning threshold c indicated the minimum surface area of connected components of the surface to be retained in the post-processing “cleaning” step.

For thresholding only, the threshold t_{thresh} was chosen and also the cleaning threshold c . For all datasets, t_{init} was slightly higher than t_{thresh} for the same dataset: although using a lower t_{thresh} alone looks better after the cleaning step, the noise before cleaning worsened our results and led us to use a slightly higher value for initialization.

Recall that obtaining the very small vessels is the goal of this work since the large vessels are easily segmented by thresholding. For this reason, $imscale$ was set fairly high in the experiments in Figure 10 to obtain the small vessels, at the expense of also obtaining many imaging artifacts. A coarser segmentation is obtained in Figure 9 by choosing lower values for $imscale$. Although the results in this figure are only similar to those obtained by simple thresholding, the objective of the demonstration is academic: it shows that we capture the vasculature shape even when the initial guess is meaningless. In practice, better results are obtained using thresholding for initialization.

When considering that the $imscale$ parameter controls the trade-off between noise and small vessels in our method, and when comparing our method to thresholding alone, it is important to note that it would not be possible to similarly lower t_{thresh} to obtain the small vessels (and noise) by thresholding alone. Lowering the threshold obtains large blobs in the volume which do not correspond to vessels. Our method is thus more powerful than thresholding alone.

Finally, we demonstrate the capability to color-code the vasculature surface based on local curvature. Notice (Figure 12) that for a ribbon-like vessel, the flatter sides shows a large radius, and the sharply curved edges show a small radius. In this example, the colorscale is continuous from darkest to lightest intensities, with darkest indicating a radius of curvature $\leq 1\text{mm}$ and lightest indicating a radius of curvature $\geq 2\text{mm}$. The curvatures output by our evolution have been smoothed by a 3x3x3 filter prior to coloring the surface.

8 Future Work

Vessels tend to appear thinner in our segmentations than in those obtained by thresholding. One possible reason is that our method uses gradients instead of intensities, so the vessel wall is found attracted to the strongest gradients, which may be fully inside the bright region indicated by thresholding. A second option is that the underlying mathematics of our algorithm assume that the vessels are 1D curves, not tubular surfaces. We believe that our ε -level set method allows the method to successfully handle tubular surfaces, but have not yet verified this analytically. A final potential reason for the discrepancy is that the segmentations obtained by thresholding may be thicker than the true blood vessels due to noise around the vessels. Future work will involve comparisons to manual segmentations which will provide ground truth to evaluate both methods.

We also observe a lot of noise in our segmentations of the first and second datasets. As mentioned above, we could obtain much less noise at the expense of

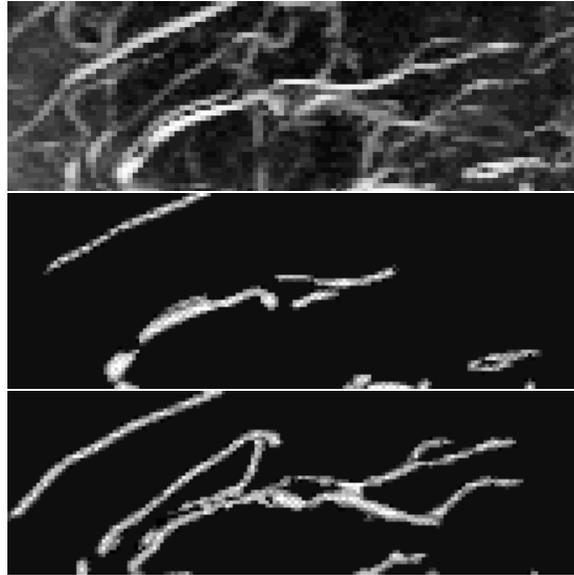


Fig. 11. Enlargement of a portion of the second example from Figure 10. As above, the second image is the segmentation obtained by thresholding alone, and the third image is the result of our method.



Fig. 12. Our method naturally allows estimation of local radii of curvature of the segmented vessels. In this image of a partial segmentation of the first dataset in Figure 10, the colorscale is continuous from darkest to lightest intensities, with darkest indicating a radius of curvature $\leq 1\text{mm}$ and lightest indicating a radius of curvature $\geq 2\text{mm}$.

the thinnest vessels by lowering *imscale*. For the large amounts of noise in these datasets, noise is often indistinguishable from small vessels when only a small local neighborhood is considered, as in our algorithm. To address this problem, one could reduce noise prior to segmentation by filtering or incorporate a more sophisticated image measure into Equation 5.

On the positive side, the segmentation of small vessels that were not obtainable by thresholding encourages us to continue in the development of this algorithm. Although still in preliminary stages, we believe that it has the potential to yield effective segmentations of very thin vessels.

Acknowledgments

This work was supported in part by NSF Contract IIS-9610249, in part by NSF ERC (Johns Hopkins University agreement) 8810-274, and in part by MERL, A Mitsubishi Electric Research Laboratory.

References

1. Altschuler, S. and Grayson, M.: Shortening space curves and flow through singularities. *Journal of Differential Geometry* **35** (1992) 283–298
2. Ambrosio, L. and Soner, H.M.: Level set approach to mean curvature flow in arbitrary codimension. *Journal of Differential Geometry* **43** (1996) 693–737
3. Caselles, V., Catta, F., Coll, T. and Dibos, F.: A geometric model for active contours. *Numerische Mathematik* **66** (1993) 1–31
4. Caselles, V., Morel, J.M., Sapiro, G. and Tannenbaum, A.: Introduction to the special issue on partial differential equations and geometry-driven diffusion in image processing and analysis. *IEEE Transactions on Image Processing* **7(3)** (1998) 269–273
5. Caselles, V., Kimmel, R. and Sapiro, G.: Geodesic active contours. *Int'l Journal of Computer Vision* **22(1)** (1997) 61–79
6. Chen, Y.G., Giga, Y. and Goto, S.: Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations. *Journal of Differential Geometry* **33** (1991) 749–786
7. Evans, L.C. and Spruck, J.: Motion of level sets by mean curvature: I. *Journal of Differential Geometry* **33** (1991) 635–681
8. Gage, M. and Hamilton, R.S.: The heat equation shrinking convex plane curves. *Journal of Differential Geometry* **23** (1986) 69–96
9. Grayson, M.: The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* **26** (1987) 285–314
10. Kichenassamy, A., Kumar, A., Olver, P., Tannenbaum, A. and Yezzi, A.: Gradient flows and geometric active contour models. In *Proc. IEEE Int'l Conf. Computer Vision* (1995) 810–815
11. Sapiro, G.: Vector-valued active contours. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition* (1996) 680–685
12. Sethian, J.A.: *Level Set Methods*. Cambridge University Press (1996)