Codimension-Two Geodesic Active Contours for the Segmentation of Tubular Structures

Liana M. Lorigo[1], Olivier Faugeras[1,2], W. Eric L. Grimson[1],
Renaud Keriven[3], Ron Kikinis[4], Arya Nabavi[4], Carl-Fredrik Westin[4]

[1]MIT Artificial Intelligence Laboratory, Cambridge, MA USA
[2]INRIA, Sophia Antipolis, France
[3]Cermics, ENPC, Champs sur Marne, France
[4]Harvard Medical School, Brigham & Women's Hospital, Boston MA, USA

liana@ai.mit.edu

Abstract

Curve evolution schemes for segmentation, implemented with level set methods, have become an important approach in computer vision. Previous work has modeled evolving contours which are curves in 2D or surfaces in 3D. Our objective is to explore recent mathematical work enabling the evolution of manifolds of higher co-dimension. We consider 1D curves in 3D (codimension-two) for the application of automatically segmenting blood vessels in volumetric magnetic resonance angiography (MRA) images. This paper describes the theoretical foundations of our system, CURVES, *then provides segmentation results compared against segmentations obtained interactively by a neurosurgeon. Segmentations of bronchi in lung computed tomography (CT) scans are also presented. The new experiments, comparisons to manual segmentations, and sample comparison to the use of a codimension-one regularization force are the primary contributions of this report.*

1. Introduction

Curve-shortening flow is the evolution of a curve over time to minimize some distance metric. When this distance metric is based on image properties, it can be used for segmentation. The idea of *geodesic active contours* (or *geometric active contours*) is to define the metric so that indicators of the object boundary, such as large intensity gradients, have a very small "distance" [4, 12, 3]. The minimization will attract the curve to such image areas, thereby segmenting the image, while preserving properties of the curve such as smoothness and connectivity. Geodesic active contours can also be viewed as a more mathematically sophisticated variant of classical snakes which evolve a curve explicitly according to image data and regularization constraints [11]. Further, geodesic active contours are implemented with level set methods [24, 19] which are based on recent results in differential geometry [7, 5]. This implementation technique provides for topological flexibility and increased numerical stability. Geodesic active contours have been extended to evolve surfaces in 3D for the segmentation of 3D imagery such as medical datasets [4, 26].

A limitation of the method was that its theoretical foundations applied only to hypersurfaces. Recent work in differential geometry, however, developed the equations necessary to evolve arbitrary dimensional manifolds in arbitrary dimensional space [1]. Subsequent work developed and analyzed a diffusion-generated motion scheme for codimension-two curves [22]. We have developed the first application of 1D geodesic active contours in 3D (codimension-two), based on [1]. Our system, CURVES, has been applied to automatic segmentation of blood vessels in medical images [17]. It evolves an initial boundary estimate toward the true structures in the image using a codimension-two regularization force which is based on the the vessel centerlines instead of on the vessel surface. This novel regularization force enables the segmentation of very thin structures which would not be possible with a traditional codimension one evolution. This paper reviews CURVES and its underlying theory, then presents novel experiments and validation comparisons. In particular, we compare CURVES segmentations of cerebral vessels to segmentations obtained manually by a neurosurgeon. We also show the advantage of a codimension-two regularization force over the traditional mean curvature force for the segmentation of lung CT data.

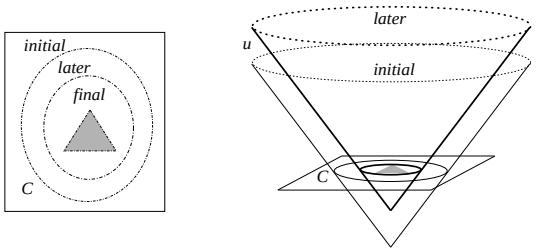

(a) (b)

Figure 1. Simple segmentation example: (a) Evolving curve. (b) Level set implementation of curve evolution.

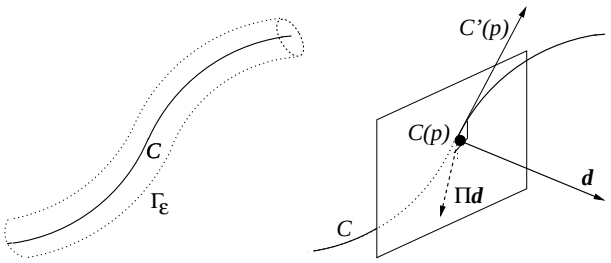

(a) (b)

Figure 2. Codimension-two curve: (a) Tubular isolevel set Γ_ε of C. (b) The tangent to C at p, the normal plane, the external vector $\vec{d}$, and its projection onto the normal plane.

2. Vascular Segmentation in Medical Images

In neurological surgery the vasculature is of utmost importance. Direct visualization of images acquired with current imaging modalities, however, cannot provide a spatial representation of small vessels. These vessels, and their branches which show considerable variations, are most important in planning and performing neurosurgical procedures. In planning they provide information on where the lesion draws its blood supply, and where it drains. This is of particular interest in vascular malformations. During surgery the vessels serve as landmarks, and guidelines to the lesion. The more minute the information is, the more precise are navigation and localization of computer guided procedures.

For these reasons we consider the segmentation of volumetric vasculature images, such as the magnetic resonance angiography (MRA) images pictured in section 5, with a focus on segmenting the small vessels. Areas of flowing blood appear bright in this imaging modality. The images are displayed in *maximum intensity projection* in which the stack of slices 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. The approach of simply thresholding the raw data is commonly used for segmentation but incorrectly labels bright noise regions as vessel and cannot recover very small vessels which may not appear connected in the volumetric image.

Several more sophisticated approaches have been explored. Multiscale approaches have involved convolving the image with gaussian filters at multiple scales and analyzing the eigenvalues of the hessian matrix at each voxel in the image to determine the local shape of the structures in the image [23, 8, 16, 14, 15]. Anisotropic diffusion has also been applied [13, 20]. Another multiscale approach based on medial axes uses the fact that the centerlines of the vessels appear brightest to detect these centerlines as intensity ridges of the image [2]. A differential geometry approach treats the 3D MRA image as a hypersurface of 4D space in which extrema of curvature correspond to vessel centerlines [21]. A statistical approach based on the expectation maximization (EM) algorithm has been proposed as a fully automatic MRA segmentation method [25, 6]. Deformable models have also been applied to 3D vascular segmentation [4, 18].

3. Background

The CURVES algorithm is an extension of geodesic active contours research, also using a level set implementation. The extension was enabled by theoretical work on level set methods for arbitrary co-dimensional manifolds.

3.1. Geodesic Active Contours

The task of finding the curve that best fits the object boundary is posed [4, 3, 12] as a minimization problem over all closed planar curves $C(p) : [0,1] \to \mathbb{R}^2$. The objective function is

$$\oint_0^1 g(|\nabla I(C(p))|)|C'(p)|dp$$

where $I : [0,a] \times [0,b] \to [0,\infty)$ is the image and $g : [0,\infty) \to \mathbb{R}^+$ is a strictly decreasing function such that $g(r) \to 0$ as $r \to \infty$, such as $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$.

To minimize this objective function by steepest descent, consider C to be a function of time t as well as spatial parameter p. The Euler-Lagrange equations yield the curve evolution equation

$$\vec{C}_t = g\kappa\vec{N} - (\nabla g \cdot \vec{N})\vec{N} \qquad (1)$$

where κ is the Euclidean curvature and $\vec{N}$ is the unit inward normal. In the absence of image gradients, this equation

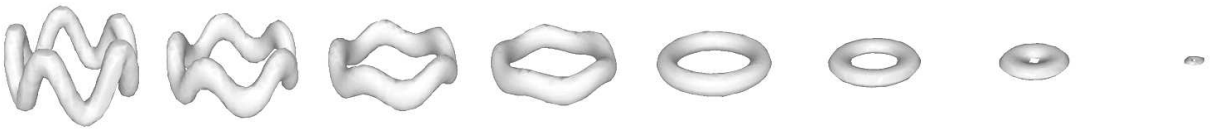

Figure 3. Tubular object evolving to smooth the underlying curve, as in CURVES. Notice the bumps are first smoothed out until the shape approximates a torus, then the torus shrinks to a point.

causes the curve to shrink according to its curvature; the presence of image gradients causes the curve to stop on the object boundary (Figure 1a).

3.2. Level Set Method for Hypersurfaces

Level set methods increase the dimensionality of the problem from the dimensionality of the evolving manifold to the dimensionality of the embedding space [24, 19]. For the example of planar curves, instead of evolving the one-dimensional curve, the method evolves a two-dimensional surface. Let $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the signed distance function to curve C as in Figure 1b; C is thus the zero level-set of u, and u is an implicit representation of C. Let C_0 be the initial curve. It is shown in [7, 5] that evolving C according to

$$\vec{C}_t = \beta \vec{N}$$

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

$$u_t = \beta |\nabla u|$$

with initial condition $u(\cdot, 0) = u_0(\cdot)$ and $u_0(C_0) = 0$ in the sense that the zero level set of u is identical to the evolving curve for all time. Choosing $\beta = g\kappa - (\nabla g \cdot \vec{N})$ as in Equation 1 gives the behavior illustrated in Figure 1b according to the update equation

$$u_t = g\kappa |\nabla u| + \nabla g \cdot \nabla u.$$

The extension to surfaces in 3D is straightforward and is called *minimal surfaces* [4]. The advantages of the level set 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.

3.3. Level Set Method for Curves in Higher Codimension

For the task of evolving one-dimensional curves in three-dimensional space, however, the above level set relation does not hold. It is applicable only to hypersurfaces, that is, surfaces whose *co-dimension* is one: the *co-dimension* of a manifold is the difference between the dimension of the evolving space and the dimension of the manifold. The examples of a planar curve and a 3D surface have co-dimension one, but 1D curves in three dimensions have co-dimension two. Intuition for why the level set method above no longer holds is that there is not an "inside" and an "outside" to a manifold with co-dimension larger than one, so one cannot create the embedding surface u in the same fashion as for planar curves; a distance function must be everywhere positive, and is thus singular on the curve itself. Recently, however, more general level set equations were found for curvature-based evolution [1].

Let $v : \mathbb{R}^3 \rightarrow [0, \infty)$ be an auxiliary function whose zero level set is identically $C(p) : [0, 1] \rightarrow \mathbb{R}^3$, 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 $P_{\mathbf{q}} = I - \frac{\mathbf{q}\mathbf{q}^T}{|\mathbf{q}|^2}$ where I is the identity matrix 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 for curve-shortening flow is then [1]

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

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

For intuition, let C be some curve in $\mathbb{R}^3$ and v_d be the distance function to C, which satisfies $|\nabla v_d| = 1$. Consider then an isolevel set $\Gamma_\varepsilon = \{x | v_d(x) = \varepsilon\}$ of v_d where ε is small and positive, so Γ_ε is a thin tube around C (Figure 2a). The nonzero eigenvalues of $P_{\nabla v_d} \nabla^2 v_d P_{\nabla v_d}$ are equal to the principal curvatures of this tube. The larger principal curvature depends on ε while the smaller is related to the geometry of C. It is according to C that we want the evolution to proceed; thus, the smaller principal curvature is chosen. For the more general embedding function v, the nonzero eigenvalues are equal to the principal curvatures scaled by $|\nabla v|$, and the intuition for choosing the smaller is the same.

Figure 3 demonstrates the behavior of a shape undergoing this motion, where the smoothing force corresponds to the curvature of the underlying 1D curve. Figure 4 then

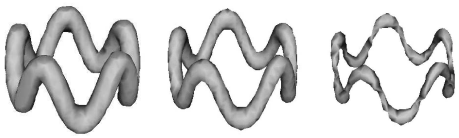

Figure 4. Codimension-one flow. The high curvatures corresponding to the small radius of the tube cause the shape to become thinner until it disappears without perceptibly affecting the geometry of the underlying curve.

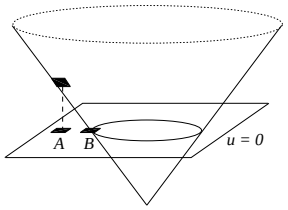

Figure 5. To evolve a point on the distance function, CURVES chooses image information from A instead of B.

compares this behavior to that of traditional (codimension-one) mean curvature flow in which the regularization is based on the mean curvature of the surface.

Now assume there is an underlying vector field driving the evolution, so the desired evolution equation is

$$\vec{C}_t = \kappa\vec{N} - \Pi\vec{d},$$

where Π is the projection operator onto the normal space of C (which is a vector space of dimension 2) and $\vec{d}$ is a given vector field in $\mathbb{R}^3$, (Figure 2b). The evolution equation for the embedding space then becomes [1]

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

4. CURVES

The evolution equation we use follows directly from an energy-minimization problem statement. Beyond that equation, several additional features of the program are incorporated for numerical and application-specific reasons.

4.1. Evolution Equation

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

$$\oint_0^1 g(|\nabla I(C(p))|)|C'(p)|dp$$

where $C(p) : [0,1] \rightarrow \mathbb{R}^3$ is the 1D curve, $I : [0,a] \times [0,b] \times [0,c] \rightarrow [0,\infty)$ is the image, and $g : [0,\infty) \rightarrow \mathbb{R}^+$ is a strictly decreasing function such that $g(r) \rightarrow 0$ as $r \rightarrow \infty$ (analogous to [4]). For our current implementation, we use $g(r) = \exp(-r)$ because it works well in practice. The Euler-Lagrange equations give

$$\vec{C}_t = \kappa\vec{N} - \frac{g'}{g}\Pi(\mathbf{H}\frac{\nabla I}{|\nabla I|}),$$

where $\mathbf{H}$ is the Hessian of the intensity function. The auxiliary vector field in the above equation is thus

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

so the equation for the embedding space is

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

4.2. Features

Initial experiments required that the evolving volume be a distance function to the underlying curve; however, it was not clear how to robustly extract the zero level set or even evolve those points since the distance function was singular exactly there. For this reason, we developed the *ε-Level Set Method* which defines a thin tube of radius ε around the initial curve, then evolves that tube instead of the curve. ε does not denote a fixed value here, but means only that the evolving shape is a "tubular" surface of some unspecified and variable nonzero width. Thus, we are now evolving surfaces similar to [4], but that follow the motion of the underlying curve so they do not regularize against the high curvatures found in thin cylindrical structures such as blood vessels. In addition to being more robust, this method better captures the geometry of such structures, which have nonzero diameter. We stress that this is an approximation to evolving the underlying curve but is not equivalent. If we were to constrain the width of the tube to remain constant along the tube, it would be equivalent; however, allowing the image to attract local surface areas independently causes the width to vary, so the tube is no longer as isolevel set of the distance function to its centerline.

Moreover, although early experiments were initialized with a tube of constant width around a pre-specified curve, we obtain improved performance when we start with a tube-like structure derived from the medical data. Current experiments are thus initialized by thresholding the raw data, as seen below in Section 5.1, so ε represents the intuition for using tubes, but never indicates a particular value.

Next, we modified the update equation for the MRA segmentation application. To control the trade-off between fitting the surface to the image data and enforcing the smoothness constraint on the surface, we incorporate an image

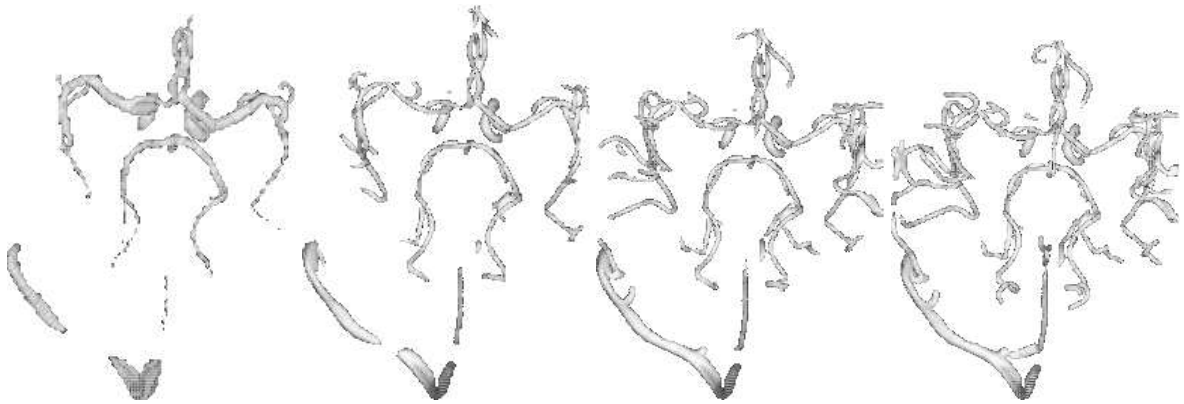

Figure 6. Surface evolution over time: initialization, followed by successive boundary estimates.

weighting term ρ which is set by the user or is pre-set to a default value. Second, because vessels in MRA and bronchi in CT 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

$$v_t = \lambda(\nabla v, \nabla^2 v) + \rho(\nabla v \cdot \nabla I)\frac{g'}{g}\nabla v \cdot \mathbf{H}\frac{\nabla I}{\mid \nabla I \mid} \qquad (2)$$

A third aspect of our system is that we do not propagate the image information off the current object boundary to obtain the values near the boundary, as is customary in level set methods [4, 3, 12, 9]. Instead, we use directly the image information at each point in the image (Figure 5). Our choice has the advantage of enabling attraction to image gradients not on the current boundary, thereby reducing sensitivity to initialization, at the expense of requiring more frequent reinitializations of the distance function and losing the equivalence to the Lagrangian method (updating C directly instead of v) off the boundary.

In short, the CURVES system takes in a 3D image that contain thin tubular structures, such as an MRA image. An initial segmentation estimate is generated by thresholding the image. That estimate is used to generate an initial distance function, which is iteratively updated according to Equation 2. Convergence is detected automatically when volumetric change in the segmentation is very small over some number of iterations. Further detail was published in [17].

5. Results

We have run CURVES on approximately 30 medical datasets, primarily phase contrast magnetic resonance angiography (PC-MRA), of various resolutions and scanner types. We provide images of several representative segmentations.

We first show successive boundary estimates in a segmentation of a cerebral MRA image to demonstrate the behavior of the algorithm over time, until convergence is reached. We then illustrate the advantage of our system compared to codimension-one surface evolution with an experiment involving the segmentation of bronchi in a computed tomography (CT) image of lung. The codimension-one result was obtained using the mean curvature of the surface as the regularization force, as in previous level set segmentation schemes [4]. Finally, we show CURVES segmentations of cerebral MRA images compared to those obtained with a manual segmentation technique used clinically at our institution.

5.1. Example Evolution

Figure 6 illustrates the behavior of our system over time on a phase contrast MRA image of cerebral vessels. The initial surface is obtained by thresholding the raw dataset, then CURVES evolution produces the subsequent images.

5.2. Bronchi Segmentation, Comparison to Codimension-One

In the lung, typically the artery, vein, and bronchus are running in parallel throughout the bronchial structure. In the CT scan, we see a combination of all of these tissues, but the majority of the signal comes from the vessels. So by segmenting the bright tubular structures, we are extracting the structure of the bronchi, but it is primarily the vessels along those bronchi that are causing the signal. For comparison purposes, we have created a version of the CURVES program which uses the codimension-one regularization force. Figure 7 shows the CURVES segmentation of bronchi in a lung CT dataset followed by its codimension-one segmen-

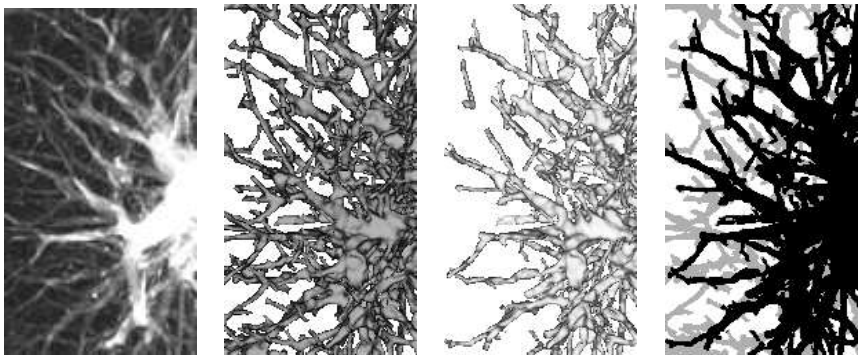

Figure 7. Segmentation of bronchi in CT lung scan: maximum intensity projection of CT dataset, CURVES segmentation, codimension-one segmentation, and overlay image in which the addition structures detected by CURVES are colored gray. Notice that CURVES is better able to capture thin bronchi than is the codimension-one method due to the new regularization force.

tation, for the same parameter settings. The two segmentations are overlaid in the final image, which is colored gray in regions where only CURVES detected structures and black in regions where both methods detected structures. Notice that the codimension-two regularization force does indeed allow the segmentation of more thin structures than does the codimension-one force, as expected.

5.3. Segmentations of Cerebral Vasculature

One specific practical motivation for our work is the use of surface models of cerebral vasculature as an aid in neurosurgical planning and procedure, especially in the context of the image-guided surgery program at our institution [10]. Currently the vessel models are obtained manually as follows. A neurosurgeon interactively chooses a threshold that is used to binarize the MRA dataset: all voxels brighter than that threshold are labeled as vessel, while all others are discarded. A "connectivity" program then partitions the set of labeled voxels into connected components. Each connected component appears in a distinct color on the user interface. The surgeon looks at individual slices and clicks on colored regions that correspond to vasculature. All connected components so chosen are stored as the final manual segmentation. The first drawback of this method is the expert user-interaction required, the second is that the thresholding step implies that all regions of image "noise" which adjoin vasculature are incorrectly labeled as vessel and small thin vessels which may appear broken or disconnected from larger structures will often be omitted. Thus, our goal is to reduce user interaction while increasing the ability to segment thin vessels.

We have compared CURVES segmentations against manual segmentations constructed as just described for nine cerebral datasets; Figure 8 shows representative comparisons. From left to right, each row shows a maximum intensity projection of the raw dataset, the CURVES segmentation, the manual segmentation, and an overlay image. The overlay image is black where both methods detected vessels, dark gray where only CURVES detected vessels, and light gray where only the manual method detected vessels. The first three rows shows a single dataset and a pair of segmentations from three orthogonal viewpoints. The dataset shown here was acquired on a 1.5T scanner with voxel size of $1.171875 \times 1.171875 \times 0.8\text{mm}^3$ and image size of $256 \times 256 \times 84$ voxels. Notice that CURVES is able to capture more of the thin vessels than the manual procedure which is based on simple thresholding. The final three rows of this figure show comparisons for three more datasets, shown only from an axial viewpoint. The first dataset was acquired using the same protocol as that in the first rows, and the second and third were acquired on a 0.5T scanner with a size of $256 \times 256 \times 60$ voxels and with the same protocol otherwise. Notice that the first of these three datasets exhibits an imaging artifact which causes a bright horizontal smear in the image near the medial cerebral arteries. The manual method must include this smear in the resulting segmentation in order to also include thinner vessels whose intensity is as high as that of the smear. Since CURVES operates on intensity gradients, it is able to capture many of the thinner vessels without mislabeling much of this smear. Color images which more clearly illustrate the comparison are available at http://www.ai.mit.edu/people/liana/3dcurves/multiview.html.

5.4. Acknowledgments

Medical datasets pictured provided by Surgical Planning Laboratory, Brigham & Women's Hospital. L. Lorigo was funded by NSF Contract IIS-9610249, NSF Contract DMS-9872228, and NSF ERC (Johns Hopkins University agreement) 8810-274. A. Nabavi was funded by a grant from the DFG (NA356/1-1). This work was also funded in part by CIMIT and NIH grants P41-RR13218-01 and R01-RR11747-01A. We thank Yoshinobu Sato of Osaka University Medical School for discussions on the algorithm and validation, and we thank Dan Kacher of Brigham & Women's Hospital for acquisition of the cerebral MRA scans used in this study.

References

[1] L. Ambrosio and H. M. Soner. Level set approach to mean curvature flow in arbitrary codimension. *J. of Diff. Geom.*, 43:693–737, 1996.

[2] S. Aylward, S. Pizer, E. Bullitt, and D. Eberly. Intensity ridge and widths for 3d object segmentation and description. In *IEEE Proc. Workshop Mathematical Models Biomedical Image Analysis*, pages 131–138, 1996.

[3] V. Caselles, F. Catte, T. Coll, and F. Dibos. A geometric model for active contours. *Numerische Mathematik*, 66:1–31, 1993.

[4] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int'l Journal Comp. Vision*, 22(1):61–79, 1997.

[5] Y. Chen, Y. Giga, and S. Goto. Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations. *J. Differential Geometry*, 33:749–786, 1991.

[6] A. Chung and J. Noble. Statistical 3d vessel segmentation using a Rician distribution. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 82–89, 1999.

[7] L. Evans and J. Spruck. Motion of level sets by mean curvature: I. *Journal of Differential Geometry*, 33:635–681, 1991.

[8] A. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Vessel enhancement filtering. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 130–137, 1998.

[9] J. Gomes and O. Faugeras. Reconciling distance functions and level sets. In *Proc. Int'l Conf. Scale-Space*, pages 70–81, 1999.

[10] W. Grimson, G. Ettinger, T. Kapur, M. Leventon, W. Wells III, and R. Kikinis. Utilizing segmented MRI data in image-guided surgery. *Int'l Journal of Pattern Recognition and Artificial Intelligence*, 1996.

[11] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int'l Journal Comp. Vision*, 1(4):321–331, 1988.

[12] A. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proc. IEEE Int'l Conf. Comp. Vision*, pages 810–815, 1995.

[13] K. Krissian, G. Malandain, and N. Ayache. Directional anisotropic diffusion applied to segmentation of vessels in 3d images. In *Proc. Int'l Conf. Scale-Space*, pages 345–348, 1997.

[14] K. Krissian, G. Malandain, and N. Ayache. Model based detection of tubular structures in 3d images. Technical Report 3736, INRIA, Sophia Antipolis, France, 1999.

[15] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Trousset. Model based multiscale detection of 3d vessels. *Proc. IEEE Conf. Comp. Vision and Pattern Recognition*, pages 722–727, 1998.

[16] C. Lorenz, I.-C. Carlsen, T. Buzug, C. Fassnacht, and J. Weese. A multi-scale line filter with automatic scale selection based on the hessian matrix for medical image segmentation. In *Proc. Int'l Conf. Scale-Space*, pages 152–163, 1997.

[17] L. M. Lorigo, O. Faugeras, W. Grimson, R. Keriven, R. Kikinis, and C.-F. Westin. Co-dimension 2 geodesic active contours for MRA segmentation. In *Int'l Conf. Information Processing in Medical Imaging*, pages 126–139. Springer-Verlag, 1999.

[18] T. McInerney and D. Terzopoulos. Topology adaptive snakes. *Medical Image Analysis*, 1999.

[19] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulation. *Journal of Computational Physics*, 79(1):12–49, 1988.

[20] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(7):629–639, May 1990.

[21] V. Prinet, O. Monga, C. Ge, L. Sheng, and S. Ma. Thin network extraction in 3d images: application to medical angiograms. *Int'l Conf. Pattern Recognition*, 3:386–390, 1996.

[22] S. Ruuth, B. Merriman, J. Xin, and S. Osher. Diffusion-generated motion by mean curvature for filaments. Technical Report 98–47, UCLA Computational and Applied Mathematics, Nov. 1998.

[23] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis. Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.

[24] J. A. Sethian. *Level Set Methods*. Cambridge University Press, 1996.

[25] D. Wilson and J. Noble. Segmentation of cerebral vessels and aneurysms from MR angiography data. In *Int'l Conf. Information Processing in Medical Imaging*, pages 423–428. Springer-Verlag, 1997.

[26] X. Zeng, L. Staib, R. Schultz, and J. Duncan. Segmentation and measurement of the cortex from 3d MR images. In *Proc. Medical Image Conference and Computer Assisted Interventions (MICCAI)*, pages 519–530, 1998.

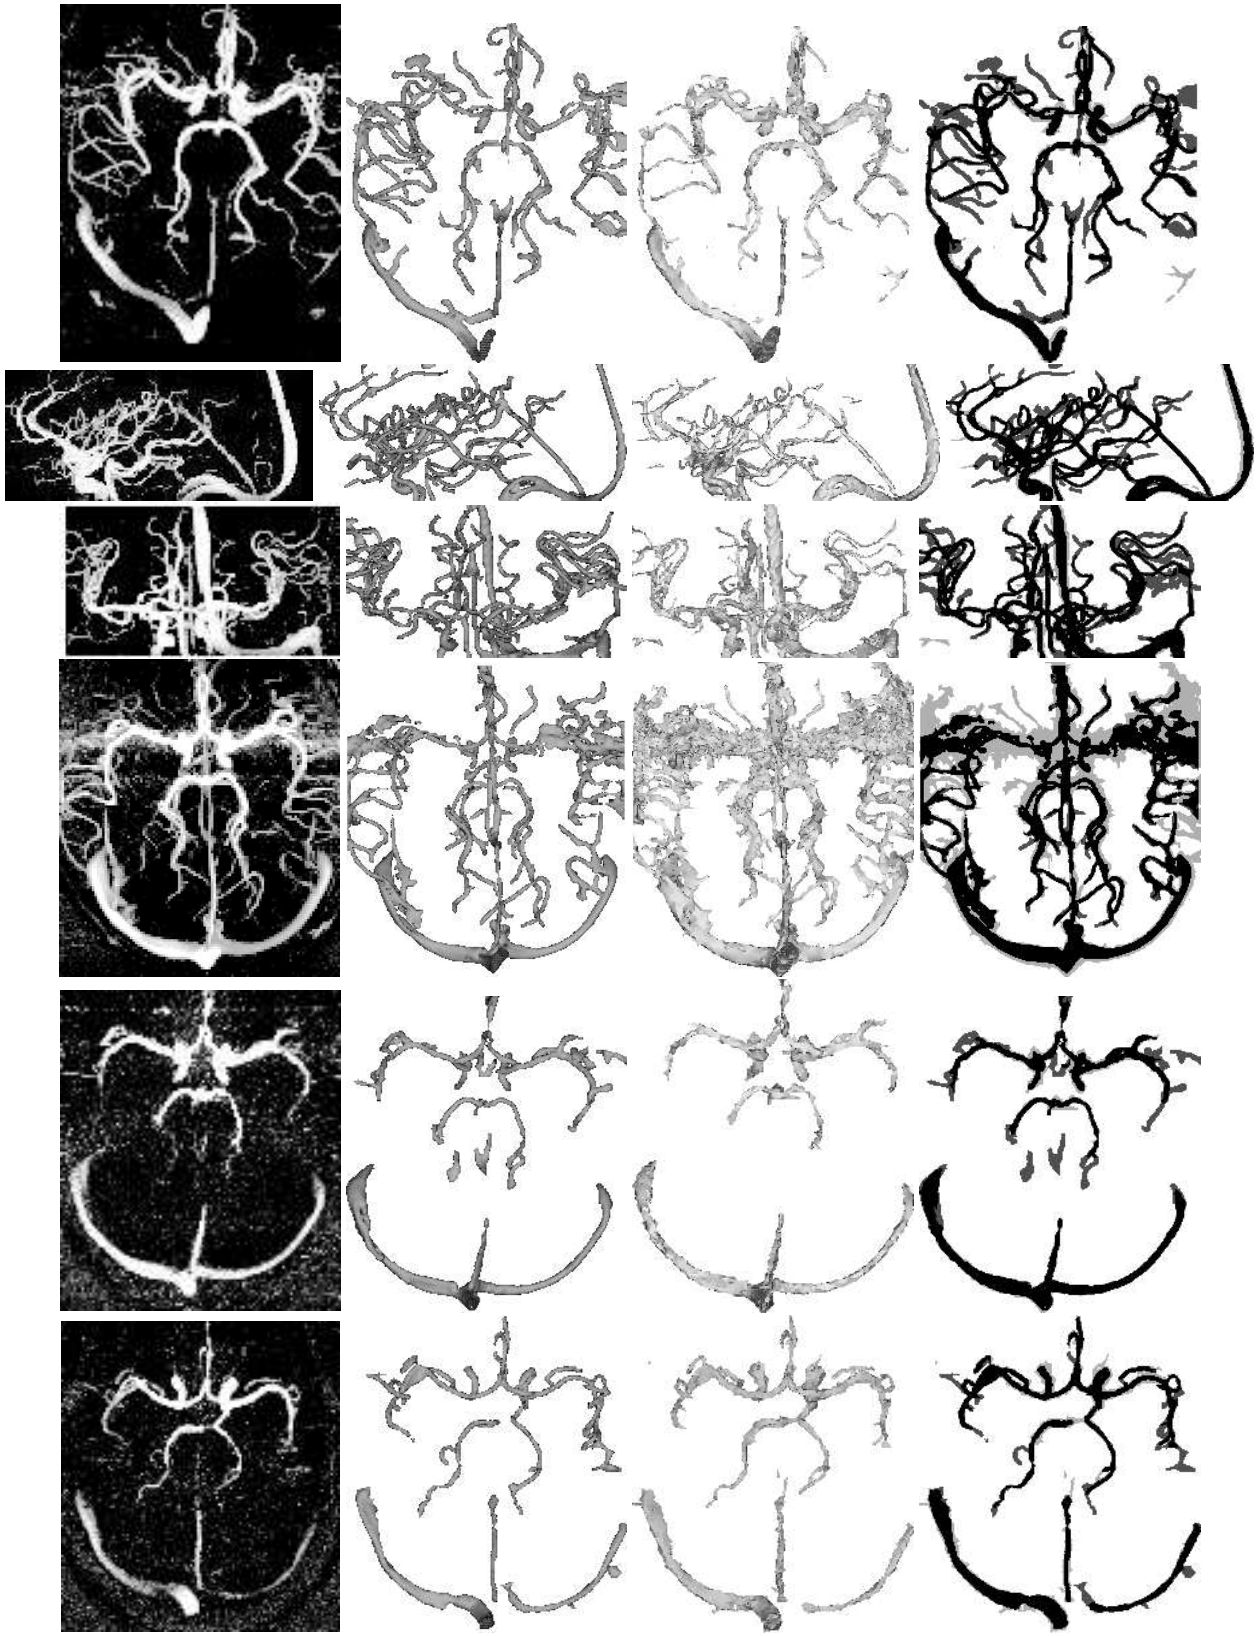

Figure 8. Data, CURVES segmentations, manual segmentations, overlays. Discussion in text.