VESSEL SEGMENTATION USING 3D ELASTICA REGULARIZATION

Noha Youssry El-Zehiry and Leo Grady

Siemens Corporate Research, Princeton, NJ, USA.

ABSTRACT

Vascular diseases are among the most important health problems. Vessel segmentation is a very critical task for stenosis measurement and simulation, diagnosis and treatment planning. However, vessel segmentation is much more challenging than blob-like object segmentation due to the thin elongated anatomy of the blood vessels, which can easily appear disconnected in the acquired images due to noise and occlusion. In this paper, we present a generic vessel segmentation approach that extracts the vessels by globally minimizing the surface curvature. The low curvature model enforces surface continuity and prevents the formation of false positives (leakages) and false negatives (holes). We present two contributions: First, we introduce a generic 3D vessel segmentation model by penalizing the boundary surface curvature. Second, we introduce an attraction force as a generalization of the boundary length in the elastica model, which guarantees a complete global solution and avoids shrinkage bias of length regularization. Our results will illustrate that the approach works efficiently across different acquisition modalities and for different applications.

Index Terms— Segmentation, Vessel Segmentation, Combinatorial Optimization, Curvature, Graph Methods

1. INTRODUCTION

Vessel segmentation approaches vary widely according to the prior information used in the segmentation, e.g.; appearance models, geometric models, and hybrid models (see [1] for extended reviews). Despite this wide variety of vessel extraction schemes, there is not a single model that can work efficiently across modalities. Each vessel segmentation approach seems to be customized to handle a specific application in a particular modality. Here, we highlight the drawbacks associated with the state-of-the-art vessel extraction schemes as reviewed by Lesage *et al.* [1] and propose a generic model that overcome these drawbacks. *Centerline based methods* aim at extracting the vessel centerline without providing any surface information. *Region growing* extraction techniques extend the vessels by testing the neighbors of a given voxel against a predefined inclusion criterion. They suffer from topological problems that result in holes (false negatives) and leakages (false positives) in the final segmentation. *Active contours* represent another major category of vessel extraction schemes; Parametric active contours are proven to be efficient in 2D segmentation but the parameterizations get very complicated in 3D. On the other hand, the formulation of geometric active contours (level sets formulation) can simply be extended to 3D but at a very high computational cost. Thin structures can also be segmented by imposing directional propagation information such as flux based approaches [2]. Although these approaches help eliminate leakage and premature stopping, they can not bridge gaps in the vessels. Additionally, the Euler-Lagrange optimization schemes, used to minimize active contour models, converge easily to local minima. Hence, we conclude that a *generic* vessel segmentation model should satisfy the following criteria: 1) **Topology preservation**: The algorithm is respective of the vessel topology, *i.e.*, preserve continuity of the surface and prevent the formation of false negatives (holes) and false positives (leakages). 2) **Global optimization**: The algorithm can be optimized globally to capture the global solution and eliminate the need of tailored initialization. 3) **Reusability**: The algorithm is reusable across different modalities.

In order to provide such a *generic* model, we build on the recent work in [3] to propose a segmentation method which generates the vessel segmentation as the *global optimization* of a model that penalizes boundary curvature. The advantage of curvature optimization (shown theoretically by Mumford [4]) is that, in the absence of reliable image data, it encourages the segmentation of straight lines, which can be used to bridge gaps. Additionally, leaking is suppressed because the leak points often create unnecessary curvature, which is penalized. Curvature minimization has been proven in practice to provide good vessel segmentation (e.g., [5, 6]), but previous work has not been able to provide a global optimization.

The algorithm developed in [3] was demonstrated to have the ability to bridge regions of a vessel for which the image information was incomplete as a result of stenosis, occlusion, etc. In this work, we extend this method in two important ways. First, we extend the formulation of the method to 3D, which is necessary in order to apply the method in several imaging modalities. However, this 3D formulation has the problem that the optimization had more difficulty producing a global optimum (see Section 2). Additionally, when leaking is a concern, we want the flexibility to deemphasize the intensity model and also to fix background seeds in confounding structures. In the original formulation of [3], the reduction of influence for the intensity model in favor of seeds would also cause the optimization method to fail. Therefore, the second contribution of this work is to decompose the boundary length term in the elastica model [4] into foreground and background *attraction forces* yielding a generalized elastica

model. This decomposition allows independent control over foreground and background voxels that yields a solution with the foreground voxels intact and without causing the shrinkage bias associated with length regularization. More importantly, it enables the Quadratic Psuedo Boolean Optimization (QPBO) and Quadratic Psuedo Boolean Optimization with Probing (QPBOP) to provide a complete global solution. Elimination of unary terms enables the algorithm to segment, with aid of seeds, objects the share the same intensity profile.

2. METHODS

We begin this section with a review of the 2D curvature optimization framework presented in [3] before proceeding to our 3D extension of this approach, addition of an attraction force and details on the optimization.

The continuous formulation of Mumford's Elastica model [4] is defined for curve $\mathcal{C}$ as

$$E(\mathcal{C}) = \int_{\mathcal{C}} (a + b\kappa^2) ds \quad a, b > 0. \tag{1}$$

where κ denotes the scalar curvature and ds represents the arc length element. When $a = 0$, the model reduces to the integral of the boundary squared curvature $E(\mathcal{C}) = \int_{\mathcal{C}} \kappa^2 ds$.

The use of combinatorial optimization by [3] to minimize the elastica model prompted the discrete formulation of the curvature on a graph. A graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ consists of a set of vertices $v \in \mathcal{V}$ and a set of edges $e \in \mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. An edge incident to vertices v_i and v_j is denoted e_{ij}. In our formulation, each voxel is identified with a node, v_i. A weighted graph is a graph in which every edge e_{ij} is assigned a weight w_{ij}. An edge cut is a set of edges that separates the graph into two sets, $\mathcal{S} \subseteq \mathcal{V}$ and $\overline{\mathcal{S}}$, which may be represented by a binary indicator vector x, such that, $x_i = 1$ if $v_i \in \mathcal{S}$ and 0, otherwise. The cost of the cut represented by any x is given by

$$\text{Cut}(x) = \sum_{e_{ij}} w_{ij} |x_i - x_j|. \tag{2}$$

In [3], we have introduced the discrete representation of curvature on the primal graph. In this formulation, if two edges incident on a node v_i, e_{ij} and e_{ik}, are cut then the cut is penalized with value $w_{ijk} = \frac{\alpha^p}{\min(||\overrightarrow{e_{ij}}||, ||\overrightarrow{e_ik}||)}$, where α is the angle between the edges. This cut penalty is then exactly decomposed into three edge weights

$$E_{(x_i, x_j, x_k)} = w_{ij}|x_i - x_j| + w_{ik}|x_i - x_k| - w_{jk}|x_j - x_k|, \tag{3}$$

where $w_{ij} = w_{ik} = w_{jk} = \frac{1}{2} w_{ijk}$. Despite the negative weights, it was shown in [3] that QBPOP was able to find a minimum cut in most circumstances. Notice that although the curvature clique was designed to penalize the cut of both edges e_{ij} and e_{ik}, the decomposition to pairwise interactions add an edge e_{jk} with negative weight. We denote the set of effective edges with nonzero weights as $\mathcal{E}^* \supseteq \mathcal{E}$.

Extension to 3D: The graph-based formulation presented in [3] and reviewed above associates cut costs with the curvature for the boundary on a dual graph. Unfortunately, many graphs of interest in 2D (e.g., an 8-connected lattice) are nonplanar and therefore have no dual. However, in [3] it was shown that the formulation may also be applied to nonplanar graphs by simply computing angles between the edges incident on a node, applying the decomposition in (3).

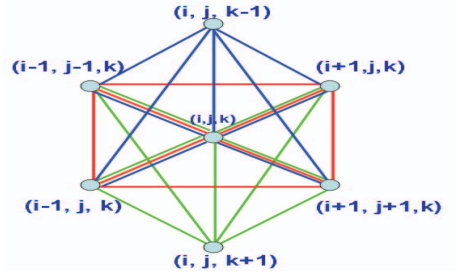


Fig. 1. Six point neighborhood used in evaluating the 3D curvature.

Extending this method to 3D follows the same approach. Since a 6-connected lattice, depicted in Figure 1, has a dual complex [7], it is straightforward to interpret this curvature formulation in terms of penalizing the corners formed on the 3D surface of the dual complex. However, boundary optimization on a 6-connected lattice is well-known to produce undesirable gridding artifacts. Consequently, it would be desirable to extend this method to apply to 3D lattices of higher-order connectivity, such as a 26-connected lattice. Unfortunately, since the dual representation is not as clear in this case, it is quite difficult to determine which pairs of edges should be penalized. Here, we performed this extension by penalizing 8 planar cliques (same cliques of the 8-connected lattice in 2D [3]) and 16 cliques; eight to each of the upper and lower planes. The added cliques represent a higher resolution partitioning of the unit sphere which should yield a smoother surface.

Optimization: In the previous sections, curvature regularization was formulated into the problem of finding a minimum cut on a graph in which some of the edge weights were negative. Unfortunately, the negative edge weights introduced by the third term of (3) cause the min-cut problem to be nonsubmodular [8], *i.e.*, straightforward max-flow/min-cut algorithms will not yield a minimum cut. However, it was shown in [3] that the Quadratic Pseudo Boolean Optimization (QPBO) and Quadratic Pseudo Boolean Optimization with Probing (QPBOP)[8] frequently offered a solution to the optimization problem that is complete and optimal.

In our extension of this work to vessel segmentation, we encounter two difficulties with this optimization approach. 1) In our experience the structure of the negative weights encountered in the 3D construction more often leads to incomplete solutions from QBPO. 2) Intensity models are often unable to distinguish vessels from other proximal structures (e.g., two touching vessels). Therefore, in these circumstances we want the ability to modify the construction to remove the intensity model and instead supply a seed in the confounding structure. Unfortunately, without an intensity model (unary term), then QBPO will be completely unable to label the voxels. To resolve this problem, it is possible to add the length term of the Elastica to decrease the non submodularity of the energy function. However, it is well-known that the length regularization introduces a shrinkage bias so we decompose the length term into foreground and background *attraction energies* in order to

produce a complete labeling of the vessel and control the foreground and background independently to avoid shrinkage.

Attraction Energy: The second term in the elastica energy of (1) is the boundary length term. The boundary length term corresponds to the minimum cut term in graph-based methods, which identifies the boundary length with a cut using edge weights which may be weighted to reflect Euclidean boundary length [9]. Therefore, combining our formulation of discrete curvature with boundary length, we obtain a discrete formulation of the elastica model as:

$$E(x) = \lambda \sum_{e_{ij} \in \mathcal{E}^*} w_{ij}|x_i - x_j| + \mu \sum_{e_{ij} \in \mathcal{E}^*} w^*_{ij}|x_i - x_j|, \quad (4)$$

where w_{ij} could be set to $w_{ij} = 1$ or to reflect the Euclidean boundary length (as in [9]) and the weights w^*_{ij} are the curvature penalties in (3). We may generalize the elastica model by decomposing the length term into a **foreground attraction force** and a **background attraction force**. The minimization of the length term expressed by $\sum_{e_{ij} \in \mathcal{E}^*} w_{ij}|x_i - x_j|$ is equivalent to the maximization of the function

$$\sum_{e_{ij} \in \mathcal{E}^*} w_{ij} \left(x_i x_j + (1 - x_i)(1 - x_j)\right). \quad (5)$$

Foreground and background attraction forces are analogous to inflation and deflation *balloon forces* presented in [10], with the difference that the attraction forces operate on pairs of voxels. While an inflation force encourages every individual voxel to be labeled foreground, a foreground attraction force encourages attraction between pairs of voxels by enforcing neighboring voxels to have the same foreground label. Consequently, we may consider a generalized elastica model that consists of the curvature term with two attraction forces. This generalized discrete elastica model is written as

$$\begin{aligned} E_{\text{elastica}}(x) = &-\lambda_1 \sum_{e_{ij} \in \mathcal{E}^*} w_{ij} x_i x_j - \\ &\lambda_2 \sum_{e_{ij} \in \mathcal{E}^*} w_{ij}(1 - x_i)(1 - x_j) + \mu \sum_{e_{ij} \in \mathcal{E}^*} w^*_{ij}|x_i - x_j| \end{aligned} \quad (6)$$

where λ_1 and λ_2 may be independently controlled to weight the foreground or background attraction forces.

A key value of the attraction force is that it allows for an optimization of the curvature energy even if the data (unary) terms are removed. Specifically, the graph construction represents the negative attraction energy by adding an edge e_{12} and an edge $e_2\mathcal{T}$ with positive weights. The addition of positive weights changes the sign of some of the negative weights introduced by the curvature term. These sign changes affect the optimization problem by strongly decreasing the number of non submodular terms in the energy.

Summary: The segmentation problem is modeled as the solution, x, which minimizes $E(x) = E_{\text{data}}(x) + E_{\text{elastica}}$.

The data term is typically instantiated by an intensity model for the object (vessel). In our experiments, we employ a very simple Chan-Vese data term [11] which models the foreground and background each with a single intensity.

$$E_{\text{data}}(x) = \sum_{v_i \in \mathcal{V}} x_i (g_i - \mu_F)^2 + \sum_{v_i \in \mathcal{V}} (1 - x_i)(g_i - \mu_B)^2, \quad (7)$$

where μ_F represents the expected foreground intensity and μ_B represents the background intensity and the elastica energy E_{elastica} is defined in (6). We intentionally chose a simple data model in order to highlight the contribution of the other energy terms. In practice, the algorithm could be customized for a specific application by replacing this data term with a data term that models the appearance of the target vessel.

3. RESULTS

This section demonstrates that our method is capable of segmenting a vessel in 3D under challenging conditions. We will present results for three challenging cases in three different modalities (CTA, MRI and US) to demonstrate that:

1. The algorithm works across modalities with no changes.
2. The algorithm can use the curvature regularization and global optimization to connect a vessel in which a signal dropout appears to disconnect the vessel.
3. The algorithm is capable of separating two structures having similar intensity.

We begin by addressing situations in which a signal dropout in the image makes the vessel segmentation challenging. The first case shows an example of a Computed Tomography Angiogram (CTA) acquisition where the Right Coronary Artery suffers a signal drop during descent. Using a simple data model in which μ_F and μ_B in (7) are fixed based on the maximal and minimal intensities in the input volumes, we see in Figure 2 that the curvature regularization is sufficient to connect the vessel under these challenging conditions. A more extreme example is given by Figure 3, which depicts a vessel in a 3D ultrasound . In this case, acquisition resulted in a series of high-intensity blobs which are separated by dark regions. However, by using exactly the same weak intensity model as before (i.e., setting μ_F and μ_B the same as in the CT case) Figure 3 demonstrates that our curvature regularization method is able retrieve the vessel from this series of blobs.

The second major benefit of the algorithm is demonstrated by the ability of the regularization method to separate two structures with a similar intensity, we placed a single seed on one slice to mark a target vessel and a second seed to mark a background vessel. When seeds are incorporated into the segmentation, a foreground seed v_i is set to $x_i = 1$ while a background seed would be set to $x_i = 0$. When seeds are incorporated into the segmentation, a data term is not necessary to avoid the trivial minimum (i.e., $x_i = 1, \forall v_i \in \mathcal{V}$). However, if the data term is removed, then at least one of the attraction terms must be included ($\lambda_1 > 0$ or $\lambda_2 > 0$) to allow the QBPO optimization procedure to find a solution for x. In this scenario, no data term was used. Figure 4 shows an example of this methodology applied to the separation of two blood vessels in an MR acquisition. Needless to say that any data model (alone) would have failed to perform this segmentation because both vessels share the same intensity profile. Moreover, the most common regularization in the literature, length regularization,

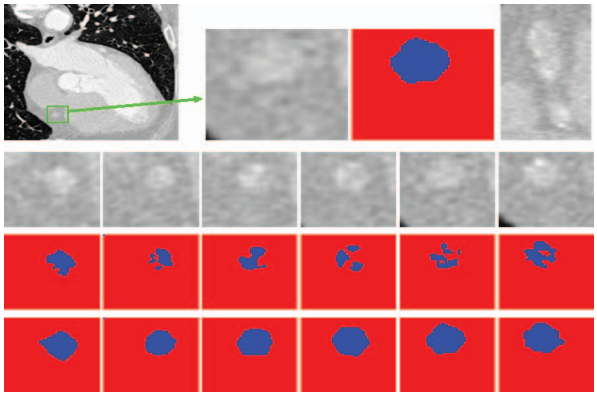


Fig. 2. **Segmentation of the Right Coronary Artery in CTA.** The first image is a coronal slice of a CTA , the second image is zooming on the Right Coronary Artery (RCA), the third image shows our result of the 3D segmentation for the given slice and the fourth image is an orthogonal view that depicts the dropout in the signal. The second row consists of six consecutive slices of the 3D volume cropped around the RCA, the third row exhibits the segmentation result using the data fidelity only and the fourth row depicts the results obtained by data and curvature regularization.

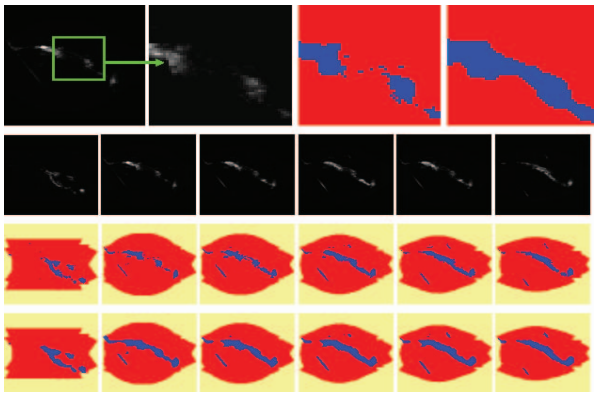


Fig. 3. **Blood vessel segmentation in ultrasound.** Color code: yellow - vowels excluded form the segmentation domain by simple thresholds, Red- Object of interest, Blue- Background. The top row shows a slice of the input volume. The second image is a magnified portion of the first image, the third and fourth images are the segmentation of the second image with data term only and with data term and curvature, respectively. Second row: Six slices of the input volume. Third row: Segmentation using data term only. Fourth row: Segmentation using data term and curvature [μ=15, $\lambda_1 = 10$, $\lambda_2 = 20$].

would have yielded a cut around one the seeds. This shrinking bias is eliminated by the curvature regularization. Another aspect that worth highlighting is the computational complexity of our algorithm relative to the state-of-the art curvature optimization schemes (that are mostly local). The most recent curvature based segmentation approach [12] reported a time varying from 10 minutes to 3.5 hours for the segmentation of a 2D image, without even guaranteeing optimality, while our approach performs volumetric segmentation of a small volume in less than a minute on a similar machine. For example, The segmentation of an ultrasound data set of size $128 \times 128 \times 15$ was performed in 25 seconds.

4. CONCLUSION

The paper presented a generic curvature based vessel segmentation model that extracts vessels across data acquisition modalities without any modifications. Moreover, our model is

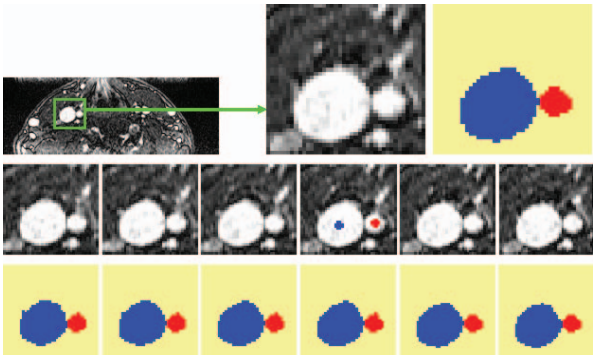


Fig. 4. **Separation of vessels in MR**. Top row shows a sample of the input slices in the first image. Second and third images are a magnified portion of the image and its segmentation, respectively. Second row: six consecutive input slices with seeds. Third row: segmentation result.

globally optimized and does not require any initialization manipulation to provide the desired results. Curvature has been shown to provide a good mechanism for segmenting vessels in both theory [4] and practice [5, 6]. Here we showed how the 2D method for global optimization of curvature in [3] could be extended to 3D with the addition of an attraction force (which also allows us to employ seeds rather than a data term). This global optimization approach avoids the difficulties of traditional vessel-following methods by considering the data as a whole, which allows it to connect regions of a vessel which are disjoint as a result of noise or pathology. Future work will address more sophisticated data models and customization of our method to specific problems in vessel segmentation.

5. REFERENCES

[1] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea, "A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes," *Med. Imag. Anal.*, vol. 13, no. 6, 2009.

[2] M. Descoteaux, L. Collins, and K. Siddiqi, "A multi-scale geometric flow for segmenting vasculature in MRI," in *MICCAI*, 2004.

[3] N. El-Zehiry and L. Grady, "Fast global optimization of curvature," in *CVPR*, June 2010, pp. 3257 –3264.

[4] David Mumford, "Elastica and computer vision," *Algebraic Geometry and Its Applications*, pp. 491–506, 1994.

[5] K. Ram, Y. Babu, and J. Sivaswamy, "Curvature orientation histograms for detection and matching of vascular landmarks in retinal images," in *SPIE*, 2009, vol. 7259, p. 52.

[6] S. Garg, J. Sivaswamy, and S. Chandra, "Unsupervised curvature-based retinal vessel segmentation," in *ISBI*, 2007.

[7] L. Grady, "Minimal surfaces extend shortest path segmentation methods to 3D," *PAMI*, vol. 32, no. 2, pp. 321–334, Feb. 2010.

[8] V. Kolmogorov and C. Rother, "Minimizing nonsubmodular functions with graph cuts: A review," *PAMI*, vol. 29, no. 7, pp. 1274–1279, 2007.

[9] Y. Boykov and V. Kolmogorov, "Computing geodesics and minimal surfaces via graph cuts," in *ICCV*, Oct. 2003, pp. 26–33.

[10] L. D. Cohen, "On active contour models and balloons," *CVGIP: Image understanding*, vol. 53, no. 2, pp. 211–218, 1991.

[11] Tony F. Chan and Luminita A. Vese, "Active contours without edges," *TIP*, vol. 10, no. 2, pp. 266–277, 2001.

[12] T. Schoenemann, F. Kahl, and D. Cremers, "Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation," in *ICCV*, Kyoto, Japan, 2009.