

A New Approach for Tubular Structure Modeling and Segmentation Using Graph-Based Techniques

Jack H. Noble and Benoit M. Dawant

Dept. of Electrical Engineering and Computer Science, Vanderbilt University
Nashville, TN 37235, USA
{jack.h.noble, benoit.dawant}@vanderbilt.edu

Abstract. In this work, a new approach for tubular structure segmentation is presented. This approach consists of two parts: (1) automatic model construction from manually segmented exemplars and (2) segmentation of structures in unknown images using these models. The segmentation problem is solved by finding an optimal path in a high-dimensional graph. The graph is designed with novel structures that permit the incorporation of prior information from the model into the optimization process and account for several weaknesses of traditional graph-based approaches. The generality of the approach is demonstrated by testing it on four challenging segmentation tasks: the optic pathways, the facial nerve, the chorda tympani, and the carotid artery. In all four cases, excellent agreement between automatic and manual segmentations is achieved.

Keywords: Tubular segmentation, optimal paths, Dijkstra's algorithm.

1 Introduction

Tubular structure localization has been studied extensively over the past few decades (see Lesage et al. [1] for an extensive review). One common segmentation approach is to use anatomy-specific models constructed from a training set. For instance, de Bruijne et al. [2] use specialized active shape models to localize abdominal aortic aneurysms. In the work presented herein, a new algorithm is presented that can be used to segment a wide variety of tubular structures. The method uses a novel approach to build anatomy-specific models and only requires a small number of sample segmentations for training. The segmentation problem is solved by finding an optimal path in a high-dimensional graph. The graph is designed with novel structures that permit the incorporation of prior information from the model into the optimization and account for several weaknesses of traditional graph-based approaches.

One common approach for tubular structure extraction is to treat it as the problem of finding an optimal path in a graph, $G=\{V,E\}$ [3]. Typically, when used for tubular structure segmentation, V , the set of nodes, corresponds to the image voxels; E the set of edges, connects neighboring voxels; and the costs associated with the edges are defined by some heuristically derived function. This basic approach is widely used because it is computationally efficient and finds a global optimum. The algorithm also

has four identifiable limitations: (1) It is difficult to incorporate local and/or training-based prior knowledge into the segmentation framework. Thus, the SOI (Structure Of Interest) must be enhanced prior to segmentation. This is not always possible as many tubular structures have intensity characteristics that vary along their length and borders that lack contrast. (2) The algorithm identifies a curve, rather than the 3D structure. Further processing must be applied to identify the full structure (an alternative is to add a 4th dimension corresponding to width as done in [4]). (3) The algorithm is biased with respect to length, i.e., because the algorithm finds the path composed of edges with the minimum cumulative costs, it tends to favor shorter, straight paths, which can result either in segmentation failure or in erroneous shortcutting at structure corners. (4) The method requires manual definition and tuning of cost functions.

The algorithm presented in this paper uses an unconventional optimal path-based extraction scheme. Segmentation is performed on a unique high-dimensional graph that is designed to incorporate information from a novel tubular structure model. This approach addresses limitations of the traditional 3D formulation.

2 Methods

Although other features could be used, in this work intensity and shape-based features that are used for training and segmentation include: (1) the curve orientation, measured using finite differences, which provides a local prior on the curve's shape; (2) intensity values along radii normal to the curve at distances of 0, 0.2, 0.9, and 1.1 times the width of the structure at every $\pi/4$ radians; and (3) intensity derivatives on those radii at 0.1 and 1.0 times the structure width. The radii directions chosen when building the model are stored and used to compute the features on the same radii in target images for segmentation.

2.1 The Tubular Structure Model and Its Training

In our approach, a large number of features are used, but techniques will be used to weigh the features that are locally the most discriminant. The approach relies on separating the SOI into a set of segments, which will be referred to as "localities" $\{l\}$. Each locality will correspond to the same anatomic substructures across samples. The separation of the SOI into localities is done automatically as part of the training algorithm. A locality l of a particular training sample s can range from one to any number of voxels in length. Figure 1 shows a synthetic 2D example of a SOI with 2 training samples divided into 5 localities. As shown, the localities correspond to similar regions across each training structure. The model training procedure is completed in four steps: (1) point correspondence across training centerlines is established, (2) the boundaries of the localities are determined, (3) model features are computed, and (4) weights for each feature are established.

Point correspondence across centerlines is necessary to determine the localities that correspond to the same distinct regions across the sample SOIs, and it is extracted using a graph-based optimization. One training volume is chosen as the reference.

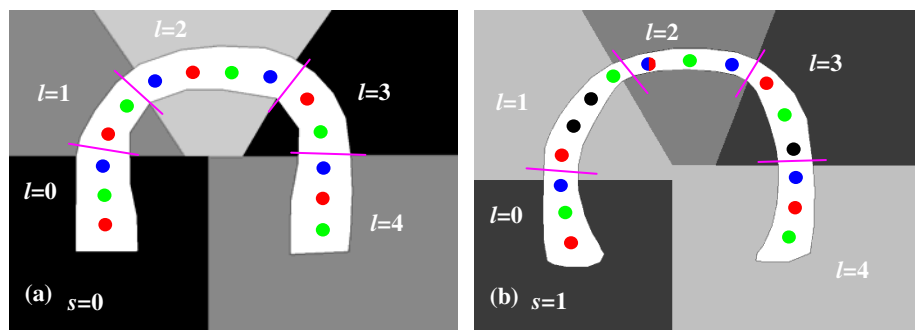


Fig. 1. Synthetic 2D example of a structure model. (a) and (b) are the reference and training SOIs. RGB values of centerline points indicate correspondence across volumes. Locality divisions are indicated by the purple lines.

The search determines one voxel along each training centerline to be associated with each reference centerline voxel. The cost function used in the search is optimized when local features between the corresponding voxels are best matched. The color-coded voxels in the centerlines in Figure 1 represent example correspondences. In (b), black voxels indicate no association with a reference voxel, and multiple colors indicate multiple associations. The main purpose of this step is to provide a way to establish a correspondence between locality boundaries across training samples. Thus, a one-to-one correspondence between centerline points is not necessary.

Although this could be automated, the number of localities for each SOI is chosen manually. However, the positions of the locality boundaries are determined automatically using a gradient descent technique. The cost function used favors smaller intra-locality and larger inter-locality variance of model features across all the training samples, and it includes a regularization term that favors boundaries spread evenly across the SOI. An example of optimal locality selection is shown in Figure 1. Once the localities are determined, the mean and variance of each feature for each locality/sample are computed.

The next step involves determining the features that are the most discriminant in each locality. A desirable feature is one that has a low value inside the SOI relative to the background. To identify these, features are first computed across the training images and the resulting feature images are convolved with a Gaussian-based valley detecting filter with scale equal to the structure width. A high positive response indicates a desirable feature. A low or negative response indicates an undesirable one. Weights for each feature are thus set to the response of the filter if it is positive and set to zero otherwise. After this step the creation of the models is completed. These consist of a sequence of structure samples, each sectioned into localities representing similar anatomical regions; and statistical values and importance weightings for a set of features within each sample/locality.

2.2 Structure Segmentation

Once the structure model is constructed, it can be used to segment new images using an optimal path finding algorithm. This algorithm requires three pieces of information: (1) the curve's starting and ending voxels, (2) a defined graph structure, and (3) a defined cost function. In this work, the starting and ending voxels were chosen manually (this could be automated in future work using, for instance, atlas-based methods as was done in [5,7]).

To achieve our goal, a novel graph structure has been implemented that is unique in two ways: (1) Nodes are associated with specific training samples/localities. This allows costs to be computed using the model's features associated with each sample/locality. Also, this permits reconstruction of the full 3D structure from the optimal path by assuming, since the features are width dependent, that the width of the structure at a node in the optimal path should match the width of the sample in the locality indicated by the node. (2) Novel graph structures are implemented such that the only paths that exist between the seed and end-node obey hard constraints on the minimum/maximum length of the structure. The length of a path in a graph can be constrained by creating an abstract dimension P in the graph. The seed lies at $p=0$ in P , and a set of candidate endnodes are chosen at $p=[l_{\min}, l_{\min+1}, \dots, l_{\max}]$ in P . The graph is constructed such that all neighbors of the seed are at $p=1$, all of their neighbors are at $p=2$, and in general all neighbors of node v_i at p_i lie at p_i+1 . It follows from this formulation that v_i is only reachable from the seed in p_i steps, i.e., by paths that contain p_i edges. An optimal path algorithm will find the path from the seed to the candidate endnode that results in the path of minimal cost, and because the graph is structured so that the only paths that exist between the seed and endnodes contain a bounded number of edges, it is guaranteed that the optimal path will have length on the interval $[l_{\min}, l_{\max}]$. This concept can be extended to apply length constraints on each individual locality. This ensures that the lengths of localities in the optimal path, and the length of the path itself, are within specified bounds. The bounds on the length of each locality are chosen as the range of the lengths of each locality in the training set.

The cost function includes model-based terms but can also include other model-independent, image-based terms, such as vesselness filters (see e.g. Frangi et. al. [6]). In this work, a directed vesselness filter is used, where the filter is aligned with the expected curve orientation given by the model. The analytical form of the segmentation cost function associated with edge e connecting node v_i to v_j is:

$$C(e_{i,j}) = \sum_{l=0}^{L-1} \omega_{s_j,p_j}^l \frac{g^l}{\sigma_{s_j,p_j}^l} + \sum_{k=0}^{K-1} \omega_{s_j,p_j}^{k+L} \frac{(x^k - f_{s_j,p_j}^k)^2}{(\sigma_{s_j,p_j}^k)^2}. \quad (1)$$

The second term is the cost associated with the K model-based features, while the first is associated with the L image-based features. f_{s_j,p_j}^k and σ_{s_j,p_j}^k are the mean and standard deviation of the k th model-based feature associated with sample s_j and locality p_j . x^k is the value of that feature evaluated in the target image. σ_{s_j,p_j}^l is the standard deviation of the l th image-based feature, and g^l is the value of that feature in the target image. The ω 's are the weighting values associated with each of those




Fig. 2. Quantitative segmentation results in mm shown for the (top graph) CH, FN, (bottom graph) OP, and CA

features. The standard deviations are used as normalizing factors to bring all the features into roughly the same scale. The standard deviations and weights for the image-based features are computed with the same approach used for the model-based features, as described in the previous section. With the graph and cost function defined, the algorithm can be used to segment new images.

3 Results

To validate the approach, the algorithm was tested on four structures: the facial nerve (FN), chorda tympani (CH), and carotid artery (CA) in CT, and the optic pathway (OP) in MR/CT. All of these structures pose a significant problem for segmentation algorithms because they have intensity characteristics that vary along their length and borders that lack contrast. The FN and CH are tiny nerves in the ear that range from 5 to 1 voxel in width. The structures lie in pneumatized bone, and surrounding structures have highly variable topology. The CA is large in comparison, and is visible in CT in the temporal bone region. The OPs are composed of the optic nerves, chiasm, and tracts. This structure suffers from lack of contrast in both MR and CT. In this study, one OP is defined as one optic nerve connecting to the contra-lateral tract through the chiasm. Previous studies have shown that localizing both OPs in this fashion and computing the union of the two can lead to an accurate segmentation of the whole optic system [7]. The FN, CH, CA, and OP are divided into 6, 4, 5, and 7 localities, which roughly corresponds to the number of visually distinct regions of the structures. For each SOI, a training set of 7 binary segmentations was acquired for training, and a leave-one-out approach was used for validation. Voxel resolution for the FN and CH images was $0.3 \times 0.3 \times 0.4 \text{ mm}^3$ and was $1.0 \times 1.0 \times 1.5 \text{ mm}^3$ for the CA and OP images. Quantitative error between manual and algorithmic segmentations was measured using symmetric surface distances as shown in mm in Figure 2. In the figure, mean and max surface errors are shown for each structure. The green bar, red




Fig. 3. Renderings of the test case resulting in worst maximum errors for the (top-to-bottom) CH, FN, OP, and CA. Left and middle columns are 3D renderings of the manual and algorithmic segmentations, color-encoded with error distance in mm as indicated by the color-bars. The right column shows contours of the manual (red) and algorithmic (green) segmentations and the automatically generated locality divisions (blue) in the images.

bar, blue rectangle, and black I-bar indicate the median, mean, one standard deviation from the mean, and the range of the errors, respectively. Each dotted-line division in the graph indicates an error of approximately 1 voxel. As seen in the graph, mean errors are sub-voxel, and the mean maximum errors are under 3 voxels for all SOIs.

Qualitative results are shown in Figure 3 for the test case that results in the worst quantitative error. To generate the 2D views, a thin-plate spline transformation was computed that warps the medial axes of the structures to lie in a plane, and then the image data and contours were passed through this transformation. This was done so that a cross section of the entire structure could be viewed in one 2D plane. As the example demonstrates, the primary location of error maxima for the FN occurs at the endpoints of the structure. This is due to errors in endpoint selection. Error maxima for the CA occur at the sharp corners of the structure, where, even though the path length is bounded, the path still may tend to try to shortcut in these regions within the allowable length bounds. Error maxima for the OP occurs at the corners of the chiasm. In this region some shortcutting occurs, and the anatomy is more rectangular in shape, which makes the tubular segmentation model less accurate. The worst errors of about 4 voxels, as demonstrated by the quantitative results and by the examples in Figure 3, are rare and highly localized. On average, results appear to be qualitatively accurate, and excellent agreement is seen between the algorithmic and manually generated contours. To demonstrate the performance of the automatic locality extraction process used to build the model, the locality divisions of the manual segmentations are also shown in the figure. As seen in the figure, the locations of the locality divisions generally correspond to changes in local intensity or curve orientation.

4 Conclusions

In this work, a novel approach for tubular structure segmentation was presented. An optimal path algorithm is used to identify the SOI in a unique high-dimensional graph that incorporates information from a novel anatomy specific tubular structure model. The graph and structure model are used to overcome some limitations inherent to approaches typically used when applying an optimal path finding algorithm. Similarly to the work of Li and Yezzi [3], the graphs presented herein allow a 3D representation of the surface to be recovered from the optimal curve. These graphs also allow local, a priori, intensity and shape information from SOI exemplars to be included and optimized by the graph search. Finally, using a priori knowledge of the approximate length of the SOI, the graphs bound the allowable structure length to remove the bias towards straight curves inherent to traditional methods. While mild shortcutting may still occur, total failure is avoided.

An approach for training the novel structure model was also presented. Using a procedure to find correspondence across exemplars, detect unique structure subsections, and extract and train discriminant features, a model of the SOI is constructed that contains statistical values and importance weightings for a set of intensity and shape features for each subsection of each training exemplar. The information in this model is then incorporated into the segmentation framework. Combining the advantages achieved by the graph structure with the weighting of the features from the model used in the cost function accounts for all four limitations of traditional graph-based approaches listed in Section 1.

For validation, the approach was tested on four challenging anatomical structures. As can be seen in the images of Figure 3, there is poor image contrast at the borders

of these structures and changing intensity profiles along the lengths of the structures. There are also substantial inter-patient variations in the intensity and shape features for each of these structures. All of these factors make the segmentation problem very difficult. However, the results of the presented approach are accurate, suggesting that it is generally applicable for fixed topology, non-branching tubular structures. To the extent of our knowledge, there are only a few existing methods for the segmentation of the FN, CH, and OP. The segmentation results presented in this paper are comparable to those of the most accurate of those existing approaches, which were specifically designed to localize these structures [5,7]. This also indicates that the approach presented in this paper is effective.

The method does require selecting one single parameter: the number of localities of the SOI. However, its selection is intuitive, i.e., it should roughly correspond to the number of visually distinct regions of the SOI. In future work, this will be chosen automatically. One possibility is brute force optimization, i.e., try every possible number of localities between 1 and N on the training set, and use the one that produces the most accurate results.

Acknowledgements. This research has been supported by NIH grants R01EB006193 and F31DC009791. The content is solely the responsibility of the authors and does not necessarily represent the official views of these institutes.

References

1. Lesage, D., Angelini, E., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: Models, features, and extraction schemes. *Medical Image Analysis* 13, 819–845 (2009)
2. de Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J.: Adapting active shape models for 3D segmentation of tubular structures in medical images. In: *Proc. Inf. Process. Med. Imaging*, pp. 136–147. Springer, Heidelberg (2003)
3. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271 (1959)
4. Li, H., Yezzi, A.: Vessels as 4D Curves: Global Minimal 4D Paths to Extract 3D Tubular Surfaces and Centerlines. *IEEE Trans. Med. Imag.* 26(9), 1213–1223 (2007)
5. Noble, J.H., Warren, F.M., Labadie, R.F., Dawant, B.M.: Automatic segmentation of the facial nerve and chorda tympani in CT images using spatially dependent feature values. *Medical Physics* 35(12), 5375–5384 (2008)
6. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A., Delp, S.L. (eds.) *MICCAI 1998*. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
7. Noble, J.H., Dawant, B.M.: An atlas-navigated optimal medial axis and deformable model algorithm (NOMAD) for the segmentation of the optic nerves and chiasm in MR and CT images. *Medical Image Analysis* (2011, in press), doi:10.1016/j.media.2011.05.001