

Simultaneous Segmentation and Anatomical Labeling of the Cerebral Vasculature

David Robben¹ *, Engin Türetken², Stefan Sunaert³, Vincent Thijs⁴, Guy Wilms³, Pascal Fua², Frederik Maes¹, and Paul Suetens¹

¹ iMinds - Medical Image Computing (ESAT/PSI), KU Leuven, Belgium

² CVLab, EPFL, Lausanne, Switzerland

³ Department of Radiology, University Hospitals Leuven, KU Leuven, Belgium

⁴ Department of Neurology, University Hospitals Leuven, KU Leuven, Belgium
david.robben@esat.kuleuven.be

Abstract. We present a novel algorithm for the simultaneous segmentation and anatomical labeling of the cerebral vasculature. The method first constructs an overcomplete graph capturing the vasculature. It then selects and labels the subset of edges that most likely represents the true vasculature. Unlike existing approaches that first attempt to obtain a good segmentation and then perform labeling, we jointly optimize for both by simultaneously taking into account the image evidence and the prior knowledge about the geometry and connectivity of the vasculature. This results in an Integer Program (IP), which we solve optimally using a branch-and-cut algorithm. We evaluate our approach on a public dataset of 50 cerebral MRA images, and demonstrate that it compares favorably against state-of-the-art methods.

Keywords: Cerebral Vasculature, Segmentation, Reconstruction, Anatomical Labeling, Circle of Willis, Integer Programming

1 Introduction

Automated segmentation and anatomical labeling of blood vessels is an important problem with many practical applications. In clinical settings, it can give an interventional radiologist extra guidance when navigating through a patient's vasculature, or it can allow automatic quantification of specific vessel segments. In a research context, it can be used to detect patterns in the vasculature that may be correlated to the incidence of vascular pathologies.

In this work, we focus on the cerebral vasculature and more specifically on the Circle of Willis (CoW) as well as its adjacent vessels. The CoW is a circle of arteries in the skull base that connects the left and right side of the anterior cerebral circulation with the posterior cerebral circulation (Fig. 1). It is supplied with blood via three large arteries, namely the left and right ICA and VBA.

* David Robben is supported by a Ph.D. fellowship of the Research Foundation - Flanders (FWO).




Fig. 1: Configuration of the Circle of Willis for two different subjects as segmented from MRA. The colors indicate the anatomical names of the vessel segments. The left CoW is complete, while the right misses several segments.

Although the CoW has a very characteristic morphology, it is highly variable: less than half of the population has a complete circle, while in the majority of cases, one or more arteries are missing.

Most existing approaches to anatomical labeling of the vasculature pose the problem in a graph-based setting, in which the vertices represent furcations and the edges the branches of the segmented vasculature. For example, Robben et al. [6] label the CoW by matching bifurcations in this graph to a probabilistic atlas, taking into account both unary potentials of the bifurcations and also pairwise potentials between them. Bogunović et al. [1] also label the CoW by matching bifurcations to an atlas. They use the bifurcation properties and have several reference graphs to model the topology of the bifurcations. The method is evaluated on ground truth segmentations as it requires topologically correct segmentations. Mori et al. [5] label the bronchial branches, which are tree-like with no loops, in an edge matching approach. A trained classifier gives a probability to each possible pair of branch and label. The solution is the global optimal assignment of labels taking into account several topological constraints. All these approaches rely on a pre-existing segmentation in terms of a graph of potential blood vessels. They account for the fact that vasculature is not a random set of tubular structures but an organ with specific connectivity patterns. However they fail to exploit this knowledge to improve the segmentations. Finally, Lu et al. [4] segment and label three non-branching coronary arteries by generating many possible segmentations and selecting – based solely on geometry – for each label the most likely.

By contrast, in this paper, we propose to perform the segmentation and labeling jointly. To the best of our knowledge, we are the first to propose such a simultaneous model-based approach for vascular structures. Not only does this approach yield better results than state-of-the-art methods [7] but it is also very generic and could equally well be applied to other curvilinear structures.




Fig. 2: (a) Maximum intensity projection of the overcomplete graph capturing the cerebral vasculature overlaid on the MRA image. (b) Illustration of the edge and edge pair labels.

2 Methodology

Following of Türetken et al. [7], we first compute a 4D scale-space *tubularity* volume, in which the last dimension stands for the vessel radius, or scale. The values represent the likelihood that there is a vessel of that radius centered at the voxel in the image. We sample local maxima of this volume at regular intervals (5 mm) and treat the samples as vertices V_I in a directed graph $G_I = (V_I, E_I)$, which contains the vasculature (Fig. 2a). The edges E_I represent tubular paths and are obtained by connecting pairs of samples that are within a certain distance from each other, using the Fast Marching Algorithm [3] in the scale space. We assume that G_I is overcomplete such that its edges cover all the vessels in the image, but it also contains spurious branches that are not part of the vasculature.

We then select a subgraph in G_I and anatomically label its edges such that it most likely represents the true vasculature. This is done jointly by optimizing a global objective function that captures both the image evidence and the prior knowledge about the geometry and connectivity of the labeled arteries.

In contrast to earlier approaches that attempt to sample their vertices from the true furcations of the vasculature, which are very hard to obtain accurately, the vertices in our graphs do not need to coincide with the furcations. As a consequence, an edge can belong to multiple anatomical segments as illustrated in Fig. 2b. However, we infer the position of the furcations once the optimal solution is obtained by merging its overlapping edges that occupy the same 3D space.

Notation Given an image I , let $G = (V, E)$ be the directed graph obtained by adding a virtual vertex v_v to G_I such that $V = \{v_i\} = V_I \cup \{v_v\}$ and $E = \{e_{ij} = (v_i, v_j)\} = E_I \cup \{(v_v, v_i) | v_i \in V_I\}$. Let also $S = \{s_i\}$ be the set of the anatomical labels for distinct segments of the vasculature extended with a

void label (na) for unnamed vessel segments (as illustrated in Fig. 1). This is necessary since we are interested in segmenting the whole cerebral vasculature rather than only the labeled segments of the CoW.

We formulate our problem in terms of consecutive edge pairs in G , since it allows us to capture more global appearance and geometry information, and it gives rise to a linear objective function and constraints [7]. Let $F = \{e_{ijk} = (e_{ij}, e_{jk})\}$ be the set of consecutive directed edge pairs and $\hat{L} = \{(s_1, \dots, s_n) | 1 \leq n \leq 4, \forall i : s_i \in S\}$, the set of the edge pair labels, where an edge pair label is a tuple of segment labels. We define $\mathbf{X} = \{X_{ijk}^l\}$ to be the vector of binary random variables, each representing the (non-)existence of a vessel segment along the directed edge pair e_{ijk} with label l . Their realizations are denoted by the vector of binary variables $\mathbf{x} = \{x_{ijk}^l\}$. In the following, we pose the joint segmentation and labeling problem as an integer program (IP) over \mathbf{x} subject to a set of constraints. We solve the resulting IPs to provable optimality (with a solution gap of $1e^{-4}$) using the branch-and-cut procedure of the Gurobi Optimizer¹. The optimization took on average 20 mins. per image on a single core.

Objective Function We formulate the problem as a maximum likelihood (ML) inference over the binary variables \mathbf{x} :

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} P(I, G | \mathbf{X} = \mathbf{x}) \quad (1)$$

where \mathcal{X} denotes the set of feasible solutions that satisfy the constraints described later in this section. We decompose $\mathbf{X} = \{X_{ijk}^l\}$ into two sets of random variables $\mathbf{T} = \{T_{ijk}\}$ and $\mathbf{L} = \{L_{ijk}\}$, with the binary variable T_{ijk} representing whether the edge pair e_{ijk} belongs to the underlying vasculature and L_{ijk} representing the label $l \in \hat{L}$ of the edge pair. We write:

$$P(I, G | \mathbf{X} = \mathbf{x}) \propto P(\mathbf{T} = \mathbf{t} | I, G) \frac{P(\mathbf{L} = \mathbf{l} | \mathbf{T} = \mathbf{t}, I, G)}{P(\mathbf{T} = \mathbf{t}, \mathbf{L} = \mathbf{l})}. \quad (2)$$

We omit here the derivation of the first term $P(\mathbf{T} = \mathbf{t} | I, G)$ as it is given in [7]. To derive the second term, we assume conditional independence of the image evidence given X_{ijk}^l , a uniform distribution $P(L_{ijk} | T_{ijk} = 0)$ – if the edge is not part of the solution, the choice of label is arbitrary – and a uniform prior distribution over \mathbf{T} . As we show in the supplementary material², this yields:

$$\frac{P(\mathbf{L} = \mathbf{l} | \mathbf{T} = \mathbf{t}, I, G)}{P(\mathbf{T} = \mathbf{t}, \mathbf{L} = \mathbf{l})} \propto \prod_{e_{ijk} \in E} \prod_{e_{ij} \in E} \prod_{l \in \hat{L}} \left[\frac{P(L_{ijk} = l | I, G, T_{ijk} = 1)}{P(L_{ijk} = l | T_{ijk} = 1)} \right]^{x_{ijk}^l}. \quad (3)$$

Taking the logarithm of Eq. 2 results in an objective function that is linear in the x_{ijk}^l variables:

$$\sum_{e_{ijk} \in F} \sum_{l \in \hat{L}} \left[\log \frac{P(T_{ijk} = 1 | I, G)}{1 - P(T_{ijk} = 1 | I, G)} + \log \frac{P(L_{ijk} = l | I, G, T_{ijk} = 1)}{P(L_{ijk} = l | T_{ijk} = 1)} \right] x_{ijk}^l. \quad (4)$$

¹ <http://www.gurobi.com>

² <http://www.medicalimagingcenter.be/public/MIC/publications/DR201405/>

Bifurcation	Proposed			Bogunović et al.[1]		
	accuracy	precision	recall	accuracy	precision	recall
ICA-OA	99	100	99	-	-	-
ICA-M1	99	99	100	99	100	99
ICA-PCoMA	93	94	96	97	98	98
ACA1-ACoMA	92	93	97	98	97	100
M1-M2	89	89	100	82	82	100
VBA-SCA	95	98	97	-	-	-
VBA-PCA1	94	100	93	96	96	100
PCA1-PCoMA	96	100	94	98	97	100
PCA2-PCA3	89	92	93	-	-	-

Table 1: Comparison of the labeling performance on the ground truth centerlines.

The probabilities $P(T_{ijk} = 1|I, G)$ are obtained using the path classifier introduced in [7], and $P(L_{ijk} = l|I, G, T_{ijk} = 1)$ is obtained from a random forest classifier using geometrical features such as the mean position, direction and radius of the edge pair e_{ijk} .

Constraints Not every \mathbf{x} gives rise to a biologically plausible and feasible solution. For example, we force the active edge pairs in the final solution to be connected to the virtual vertex v_v as in [7], and every edge e_{jk} to have at most one incoming edge pair e_{ijk} as illustrated in Fig. 2b. Furthermore, our algorithm learns from the annotated training data which edge pair labels – and more importantly – configurations of labels are possible in the final solution: labels of overlapping edge pairs should be compatible, some labels can occur only in a furcation, etc. All these constraints are expressed by the linear inequality $W\mathbf{x} \leq \mathbf{b}$, where $W_{ij} \in \{-1, 0, 1\}$ and \mathbf{b} is a binary vector. A more detailed description is given in the supplementary material.

3 Evaluation

In this section, we first evaluate the labeling and segmentation performance of our algorithm separately, each against the state-of-the-art approach for the respective task and then report our combined performance. All experiments are done with a leave-one-image-out cross-validation, using 50 MRA images of the cerebral vasculature from a public dataset [2] together with their ground truth segmentations (as used in and provided by [1]) and anatomical labels manually annotated by an expert. The images are rigidly aligned and cropped to the region that covers the segmentations [1].

Anatomical Labeling on the Ground Truth Centerlines Instead of using an overcomplete graph constructed from the image, we create a graph from the ground truth centerlines. This graph is unlabeled, but contains only valid

Metric	Proposed	Türetken [7]	Bifurcation	acc.	prec.	recall
Prec. overlap	83(5)	84(6)	ICA-OA	60	79	69
Recall overlap	76(5)	77(5)	ICA-M1	91	95	95
Prec. topology	61(39)	44(44)	ICA-PCoMA	68	73	79
Recall topology	70(43)	65(42)	ACA1-ACoMA	55	52	98
			M1-M2	59	59	98
			VBA-SCA	82	95	85
			VBA-PCA1	83	90	90
			PCA1-PCoMA	42	28	52
			PCA2-PCA3	71	73	88

(a)

(b)

Table 2: Results of the simultaneous segmentation and labeling. (a) Comparison of the segmentation performance. Reported numbers are the mean and standard deviation (in parenthesis) over the images. (b) Labeling performance.

edges. By setting $P(T_{ijk} = 1 | I, G) = 1$, we can use our algorithm to only perform labeling. The result is an edge labeled graph, from which we can infer the position of several named bifurcations. The positions are compared with those in the ground truth annotation. If the Euclidean distance is smaller than 2mm, it is considered a true positive. Since we use the same dataset as [1], we can directly compare the performance. Results are given in table 1.

On average, the accuracy is about the same, but on individual PoI, there is quite some difference, especially on the furcation formed by M1-M2. This furcation has many variations and to infer its position, medical experts look at the end positions of its daughter branches. This is done implicitly by our edge labeling algorithm. Finally, it should be noted that the method of Bogunović et al. [1] requires topologically correct segmentations, and uses reference graphs explicitly stating PoI connectivity and order for the entire vasculature. Extending it to a larger number of bifurcations requires a steep increase in the number of reference graphs. For example, inclusion of the left and right VBA-SCA, which can lie either before or after VBA-PCA1 and not necessarily next to each other, would already triple the number of reference graphs from 8 to 24 in their approach.

Simultaneous Segmentation and Labeling We compare the segmentation quality of our method with the method of Türetken et al. [7], which can be thought of as our algorithm without the labeling step. The comparison is performed based on two criteria: overlap and topology of the solutions. The former is computed by measuring the distance between the centerline points of the ground truth and the solution. We consider a point in the solution to be a true (false) positive if its distance to the closest ground truth point is less (more) than half the ground truth’s radius at that point.

The topology criterion, on the other hand, is a more global measure since it captures connectivity of the vasculature. We compute it by first finding all

the paths that extend between the start points of the VBA or ICAs segments in the ground truth. Note that such a path visits multiple anatomical segments in the CoW. We only compute the paths between the VBA and ICAs segments because they are the brain supplying arteries, and therefore, the origin of the blood supply to the entire vasculature. We then enumerate all the paths between these three segments in the obtained solution and find their corresponding paths in the ground truth if they exist. We consider a path in the solution to be a true positive if the overlap precision and recall of its centerline points (as described earlier in this section) are both above 0.7.

Table 2a gives the precision and recall values over the evaluated images for the two measures. The overlap values of the two methods are relatively similar but still significantly better than thresholding and thinning, as used in [6] (precision: 73(7), recall 64(6)). In terms of the topology criterion, our approach clearly outperforms that of [7], which is also illustrated in the close-ups of Fig. 3.

The labeling results are given in Table 2b, which are calculated as described in Section 3. The results show that the additional difficulty of segmenting the image causes the labeling accuracy to drop. Finally, Fig. 4 shows automated segmentation and labeling results of our algorithm for two complete vasculatures.

4 Conclusion

To the best our knowledge, we presented the first algorithm for simultaneous segmentation and anatomical labeling of the vasculature. Our probabilistic formulation results in an integer program, which we solved optimally on the evaluated dataset of 50 images. We demonstrated that our approach compares favorably against specialized state-of-the-art algorithms that address the segmentation and labeling problems separately. In future work, we will evaluate the segmentation quality more extensively on different structure types and imaging modalities.




Fig. 3: The centerlines of the ACA1 and AComA according to the ground truth (left), the method of Türetken et al. (middle), which contains spurious branches, and the proposed method (right).




Fig. 4: Two automatically segmented and labeled vasculatures. As in Fig. 1, the colors indicate the anatomical labels of the vessel segments.

Acknowledgement: This work was supported in part by the EU ERC project MicroNano. The MR brain images from healthy volunteers used in this paper were collected and made available by the CASILab at The University of North Carolina at Chapel Hill and were distributed by the MIDAS Data Server at Kitware Inc. We thank the authors of [1] for sharing their centerline delineations.

References

1. Bogunovic, H., Pozo, J.M., Cardenes, R., Roman, L.S., Frangi, A.F.: Anatomical labeling of the Circle of Willis using maximum a posteriori probability estimation. *IEEE Transactions on Medical Imaging* 32(9), 1587–99 (Sep 2013)
2. Bullitt, E., Zeng, D., Gerig, G., Aylward, S., Joshi, S., Smith, J.K., Lin, W., Ewend, M.G.: Vessel tortuosity and brain tumor malignancy: A blinded study. *Academic Radiology* 12(10), 1232–40 (Oct 2005)
3. Li, H., Yezzi, A.: Vessels as 4-D curves: global minimal 4-D paths to extract 3-D tubular surfaces and centerlines. *IEEE Transactions on Medical Imaging* 26(9), 1213–23 (Sep 2007)
4. Lu, L., Bi, J., Yu, S., Peng, Z., Krishnan, A., Zhou, X.S.: Hierarchical Learning for Tubular Structure Parsing in Medical Imaging : A Study on Coronary Arteries Using 3D CT Angiography. In: *ICCV*. pp. 2021–2028 (2009)
5. Mori, K., Ota, S., Deguchi, D., Kitasaka, T., Suenaga, Y., Iwano, S., Hasegawa, Y., Takabatake, H., Mori, M., Natori, H.: Automated anatomical labeling of bronchial branches extracted from CT datasets based on machine learning and combination optimization and its application to bronchoscope guidance. In: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI. LNCS*, vol. 5762, pp. 707–14. Springer Berlin Heidelberg (Jan 2009)
6. Robben, D., Sunaert, S., Thijs, V., Wilms, G., Maes, F., Suetens, P.: Anatomical labeling of the Circle of Willis using maximum a posteriori graph matching. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI. LNCS*, vol. 8149, pp. 566–73. Springer Berlin Heidelberg (Jan 2013)
7. Türetken, E., Benmansour, F., Andres, B., Pfister, H., Fua, P.: Reconstructing loopy curvilinear structures using integer programming. In: *CVPR*. pp. 1822–1829 (2013)