Principal Curves: a Technique for Preliminary Carotid Lumen Segmentation and Stenosis Grading

Release 1.00

Wilbur C. K. Wong, Ronald W. K. So and Albert C. S. Chung

July 17, 2009

Lo Kwee-Seong Medical Image Analysis Laboratory
The Hong Kong University of Science and Technology
Clear Water Bay, Hong Kong

Abstract

A vessel lumen centerline extraction method is presented. Its preprocessing and initialization for the application to 3D CT angiography (CTA) image volumes are also described. This approach is based upon the theory of principal curves to finding lumen axis. Meanwhile, it also estimates the lumen local width with an 1D intensity model that depicts an ideal lumen cross-sectional profile. We tested the method on 46 CTA datasets supplied with the Carotid Lumen Segmentation and Stenosis Grading Challenge (CLS) of the "3D Segmentation in the Clinic: a Grade Challenge", the 3rd MICCAI Workshop. The evaluation of 31 testing datasets suggests that our method can provide a good basis for post-processing if high accuracy of segmentation and stenosis grading is desired.

Contents

1 Introduction

This paper describes the preprocessing, initialization and core algorithm of a vessel lumen centerline extraction method [9] participated in the Carotid Lumen Segmentation and Stenosis Grading Challenge (CLS) of the 3rd MICCAI Workshop in the series "3D Segmentation in the Clinic: a Grade Challenge". This challenge pursues two main goals: (1) carotid bifurcation lumen segmentation and (2) internal carotid artery stenosis grading. By assuming circular lumen cross-section, both the segmentation and grading can be derived from the outputs of the presented method. Our method exploits principal curves to extract lumen centerlines from angiography. Principal curve is a generalization of the first principle component [5]. A principal curve is a smooth path that passes through the "middle" of a set of points. By transforming image voxels to a scalar-valued (image intensity) point cloud, we use principal curves to find lumen trajectories. We define an 1D intensity model that depicts the cross-sectional profile of an ideal lumen. Our method identifies lumen centerlines at locations where there are high positively correlation with the intensity model. The radius parameter of the model gives a good estimation of the lumen local width. The rest of the paper is presented as follows. Data and inputs provided by the CLS organizer are described in Section 2. The preprocessing, initialization and core algorithm to find the principal curves are elaborated in Section 3. The implementation to obtain lumen segmentation and stenosis grading is given in Section 4. Finally, results are presented in Section 5 followed by a conclusion.

2 Data and Inputs

There are totally 46 datasets (excluding the on-site datasets) provided by the CLS organizer, 15 of them are for training and the rest 31 are for testing. They were acquired with three different protocols viz Erasmus MC (30 in total, 9 training and 21 testing), Hadassah (8 in total, 3 training and 5 testing) and Louis Pradel (8 in total, 3 training and 5 testing). The reconstructed resolution of the Erasmus MC and Louis Pradel protocols is higher than the real in-plane resolution 0.6×0.6 mm. The former protocol doubles its acquisition resolution giving reconstructions with in-plane pixel sizes of $0.23 - 0.26 \times 0.23 - 0.26$ mm. The latter protocol gives reconstructed resolution comparable to the real in-plane one. As such, we downsampled every image slice from the Erasmus MC protocol by a factor of two to save computational power without losing accuracy. No downsampling was performed in datasets from the Hadassah and Louis Pradel protocols.

For each dataset, there are three points supplied as the only input. These three points mark the start location of the common carotid artery (CCA) segment of interest, and the end locations of the internal carotid artery (ICA) and external carotid artery (ECA) segments of interest. No extra location is used in our method. We used a default set of parameter setting for data preprocessing and the algorithm to find principal curves. In exceptional cases where the outputs are undesired, such as the centerline extracted is off the vessel trajectory and poor estimation of the local radius, we change the parameter setting (for details see 4.4) to get better results.

3 Methods

3.1 Preprocessing

The goal of image preprocessing is to prepare a suitable input for our core algorithm to find prinicipal curves that depict the lumen centerlines. Furthermore, irrelevant regions are excluded to speed the later processing. Image slices are resampled to a resolution comparable to their real in-plane resolution, if necessary. Each dataset is cropped to a rectanglar region at the carotid artery bifurcation. This region includes the CCA, ICA and ECA segments of interest. More preciously, we define the region with an axis-aligned bounding box. This box is determined based upon the three input points and their centriod. We found that the carotid artery is oriented along the Z-axis in all the datasets with lower Z-coordinate in CCA start point and higher Z-coordinate in ICA/ECA end point. As such, we use these two coordinates as the box boundary in the Z-direction. A small value (4 mm, comparable to the radius of CCA lumen cross-section) is subtracted/added to the coordinates to widen the boundary. The in-plane (X-Y plane) rectangular region of the box is centered at the centriod (only the X- and Y-coordinates are considered) of the input points. The size of this region is a fraction of the slice dimensions. We used $1/2$ (i.e., 256×256 mm before the downsampling) for the Erasmus MC datasets and $1/4$ (i.e., 128×128 mm) for the Louis Pradel and Hadassah datasets. A smaller fraction is used in the latter two sets because their field of view is larger. In case, any of those input points lay outside the bounding box, we enlarge the box such that they are included.

The algorithm to find principal curves assumes the vessel lumens are one of the brightest structures in the images. It requires also the intensity values range from 0 and 1. To satisfy these requirements, we clamp the image intensity to get rid of irrelevant structures, such as, calcification, bone and soft tissue, followed by an intensity rescaling. Perhaps, due to imaging artifacts, we noticed that the image intensity (an offset of the Hounsfield Unit) is not consistent with the composition and nature of the imaged tissues across different datasets acquired from the same protocol. We employ Otsu multiple-thresholding [6] to determine the lower and upper clamp intensity values. Intensity less than the lower clamp value (L) corresponds to the air and dim soft tissue, while intensity higher than the upper clamp value (U) corresponds to the bone and calcification. We use 4 thresholds and 128 histogram bins in the Otsu algorithm. Empirically, we found that the second and the last thresholds (ascending thresholds) are good candidates of the two clamp values. Pixels having intensity smaller than the lower clamp value are limited to L, and those having intensity greater than U are set to $\max(L-(U-L),0)$. Thus, after the intensity rescaling, bone and calcification pixels have 0 value, air and dim soft tissue pixels have values close to 0.5, and vessel lumen pixels have values > 0.5.

3.2 Initialization

Our lumen centerline extraction algorithm requires the initial prinicipal curves roughy follow the trajectory of the lumen of interest. The given three input points defining the start and distal ends of the carotid artery bifurcation are thus no sufficient. To get a proper initialization, we first find two polylines (a polyline is a continuous line composed of multiple line segments) that depicts the vessel segment from CCA to ICA and to ECA. We employ Frangi *et al.* vesselness measures [3] and the minimum cost path (MCP)-based vessel tracking method proposed by Wink *et al.* [8] to find a path between the given start and end points. The found path is a sequence of adjacent voxels, which is too complex if it is transformed to a polyline (by mapping voxel centers to vertices of the line segmens). Therefore, we use an algorithm [2] to further simplify it.

The bifuration location is then determined by investigating the distance apart of the two polylines. We take fine samples (1000 locations per line segment) along the polyline representing the CCA-ICA lumen, and calculate the shortest Euclidean distance from the counterpart polyline. Scanning from the CCA start point, we look for a sample that has the shortest distance greater than a threshold (we use 4 mm). This sample location is our junction point. The initial principal curves are then formed by three polylines joint together at the junction point. This initial model is therefore represented as an undirected graph, in which the polylines' points are the vertices and their line segments are the edge of the graph. Each polyline is orginated from one of the input points and ends at the junction point. The polylines are obtained from the MCP-based vessel tracking method followed by a simplifcation as mentioned previously. In rare case where the determined junction point is too distal from the bifurcation, one of the polylines (either the one from the junction point to ICA or ECA) runs against the Z-direction (Figure 4). To correct this, we replace the junction point by the vertex having the lowest Z-coordinates amongst those two polylines. Then we reconstruct the graph again.

3.3 Principal Curves to Depict Lumen Centerlines

A principal curve is a smooth path passes through the 'middle' of a point cloud. In this work, we treat voxels as a set of points in a 3-D space. Each point has an associated scalar value, its voxel intensity value. We assume vessel lumens are bright objects on a homogeneous background. As such, we want the principal curve to go through high valued points whereas whose peripheral points have low intensity. Such curve depicts the lumen centerline. To achieve this, we exploit an intensity model and search for locations where the point intensity is positively correlated to the intensity model. We use a Butterworth-shaped function as the model. It is a 1-D function taking a distance value as one of the parameters. This distance value corresponds to the shortest Eucidean distance from a point to the curve. The model gives the highest value at distance zero. The return value drops to a background value when the distance is greater than a threshold. Mathematically, this function is expressed as $h(d,B,H,R,n) = B + (H-B) / \left(1 + (d/R)^{2n}\right)$, where d is the distance value, B is the background value, H is the value at $d = 0$, R denotes the lumen width and n controls the sharpness of the lumen boundary. Other models are possible, such as Gaussian and rectangle. However, ours is a good complement of them (Figure 4). Since our model is in 1-D, we further assume the normal cross-section of the lumen is circular.

3.4 Polygonal Line Algorithm

Our algorithm to find principal curves is based on the polygonal line algorithm [5], it is an iterative algorithm. Starting from the initial principal curves represented as an undirected graph, the algorithm adds new vertices to the existing edges progressively while aligns the edges and vertices to locations with high correlation to the intensity model. The algorithm consists of three core steps: (1) data point projection, (2) vertex optimization and (3) intensity model optimization. In the first step, the scalar valued points are projected to the nearest graph edge. This is to prepare a list of points that is relevant to the intensity model correlation calculation. Alignment of the edges and vertices to the high correlation locations is performed in the second step. This is done by changing the vertex coordinates while minimizing the sum of squared difference (SSD) between the projected point scalar values and the intensity suggested by the model on a per edge basis. We call this SSD the data term. Apart from this term, a regularization term is also used to keep the polylines smooth. In the last step, with the optimized graph vertices and the projected points, we optimize the intensity model parameters such that the data term per edge is minimized. This is an essential step to estimate the lumen radius, the model parameter R gives such estimation. After these three steps and before the algorithm goes to the next iteration, new vertices are added to line segments that are long and having large SSD. These line segments are good candidates for furhter refinement because their projected points show low correlation to the intensity model. The algorithm keeps iterating until the averge edge length of the undirected graph is less than a threshold. Our initial intensity model parameters are $B = 0.5$, $H = 0.7$, $n = 4$, $R = 2$ for the graph edges correspond to the CCA segment and $R = 1$ for ICA/ECA segment. Details of the algorithm can be found in [9].

4 Implementation

4.1 Lumen Segmentation

The CLS organizer requires the participants to provide their segmentation results as a floating point partial volume segmentation. In this segmentation, the voxel intensity ranges from 0 to 1. The value 0 denotes a voxel is laid completely outside a lumen, whereas, a 1 implies a complete lumen occupation. The in between values give the volume percentage of such occupation. The outputs from our algorithm that are relevant to the lumen segmentation are the lumen centerlines and their associated radii. We use a brute force method to compose the segmentation result. We visit every voxel in the volume and check if all the 8 voxel corners are completely inside/outside the lumen. This is done by determining if the shortest Euclidean distance from a corner to the lumen centerlines is greater than the associated radius. Greater denotes outside and smaller denotes inside, a value of 0 and 1 is assigned to the voxel, respectively. As far as we find some corners are outside the lumen and some are laid inside, the voxel is subdivided into 8 even volume smaller voxels. The lumen occuption test is then conducted amongst these subvoxels. Further subdivision may be needed, the maximum subdivision level in this work is 3. This gives an accuracy up to $1/2^{3\times3} \times 100\% = 0.2\%$ volume.

4.2 Stenosis Grading

Stenosis grading in our case is straightforward, the lumen diameter is 2× the estimate radius, and the cross-sectional area equals $\pi\times$ squared of the radius. Since we assume the normal cross-section along the lumen is circular, our diameter and area give only a coarse estimation. Starting from the given CCA point , we take fine samples along the CCA and ICA polylines and interpolated the radius at each sample. The sample having the smallest diameter/area are identified. Then we continue walking towards the given ICA point and calculate the average diameter/area of the reference part. This reference part starts at 20 mm distal of the previously identified sample and has a length of 10 mm. Stenosis grading is then calculated based on the NASCET-like formulae defined by the organizer [4].

4.3 Languages and Libraries

We based our implementation heavily on Python scripts, The Visualization Toolkit (VTK) and Insight Segmentation and Registration Toolkit (ITK). Python scripts were developed to compute the axis-aligned bounding box for volume cropping from the three input points, determine the junction point from the CCA-ICA and CCA-ECA polylines and calculate the stenosis grading. We authored several snippets with `vtkExtractVOI`, `vtkImageResample`, `vtkImageShiftScale`, `itk::ThresholdImageFilter`, `itk::OtsuMultipleThresholdsImageFilter` and `itk::ShortestPathImageFilter` [7] to crop volume of interest, downsample the Erasmus MC datasets, rescale the intensity, clamp the intensity, determine the lower and upper clamp values, and track vessels with the MCP-based method, respectively. We developed our own VTK filters to implement the polygonal line algorithm and the brute force subdivision method to compose the partial volume segmentation result from lumen centerlines and their assoicated radii. All the snippets were packaged as command-line software and glued with several Bash scripts for batch processing on a Linux notebook.

4.4 Parameter Setting

We use a default set of parameters in the preprocessing, initialization and execution of the polygonal line algorithm. In rare cases, the extracted centerlines do not follow the desired lumen trajectories, fine-tuning some of the parameters is required. Generally, there are two reasons for the problem. First, the adjoin vessels are too close. This makes the vesselness measure giving non-maximum responses along the lumen axis. Thus the initialization does not follow roughly the lumen trajectories, providing a poor starting point for the polygonal line algorithm. The default setting that we used in the vesselness measure calculation is two-scale at 1 and 2 mm. The non-maximum responses are entirely due to the large scale. To avoid such poor initialization, we used a single scale at 1.5 mm instead. Smaller scale may be needed if this setting does not help. Second, the prinical curve initialization is over-simplified. The regularization term of the polygonal line algorithm may dominate the vertex optimization, making the output curves shortened at high curvature region (Figure 4). A remedy is to reduce the degree of simplication in the polyline construction. Our default parameter to control the simplification tolerant is 3 mm, we use 0.75 mm to get a complex polyline (Figure 4).

5 Results

5.1 Carotid Lumen Segmentation

We extracted the lumen centerlines and estimated their radii from all the available 46 datasets. Lumen segmentation results were compiled using the method described in Section 4.1. Evaluation of the results from the 31 testing datasets are present in this section. They were prepared by the CLS organizer. The default parameter setting in the preprocessing and initialization was generally sufficient to provide satisfactory results. Only 9 datasets required parameter fine-tuning. Amongst these datasets, our algorithm failed to extract the carotid bifurcation lumen in 4 of them (016, 020, 022, 105). An investigation suggests that the proximate vessels and irrelevant structures (e.g. bones) are causes of the problem. They both lead to bad principal curve initialization and hence the segmentation. In order to have a proper initialization, extra points on the lumen or a manual initialization are needed. Tables 1 and 2 tabulate the evaluation results provided by the organizer. The Dice similarity index (dice), mean surface distance (msd), root mean squared surface distance (rmssd) and maximum surface distance (max) are presented. Detailed formulations of these metrics are described in [4]. Our average Dice similarity index is greater than 0.7 indicating good agreement with the reference standard segmentation [1]. The mean surface distance, on average, is less than the voxel diagonal length 1.24 – 1.31 mm, indicating a fair result can be obtained with the circular cross-section assumption. However, due to this assumption, the average maximum surface distance is relatively high and our segmentations are not as good as those obtained from the human observers (Table 2).

Table 1: Summary lumen

Measure	% / mm			rank		
	min.	max.	avg.	min.	max.	avg.
L_dice	54.8%	93.3%	76.8%	4	4	4.00
L_msd	0.21mm	3.74mm	1.16mm	4	4	4.00
L_rmssd	0.30mm	6.11mm	2.07mm	4	4	4.00
L_max	1.05mm	17.16mm	7.59mm	3	4	3.97
Total (lumen)				**3**	**4**	**3.99**

Table 2: Averages lumen

Team	dice		msd		rmssd		max		Total
name	%	rank	mm	rank	mm	rank	mm	rank	rank
media	76.8	4.0	1.16	4.0	2.07	4.0	7.59	4.0	4.0
ObserverA	95.4	1.5	0.10	1.5	0.13	1.6	0.56	1.9	1.6
ObserverB	94.8	2.4	0.11	2.4	0.15	2.3	0.59	1.8	2.2
ObserverC	94.7	2.2	0.11	2.1	0.15	2.1	0.71	2.3	2.2

5.2 Stenosis Grading

We calculated the stenosis grading from the estimated lumen radii. The average absolute differences bewteen the reference standard values are listed in Tables 3 and 4. Our method gives large difference $> 25\%$ in both diameter- and cross-sectional area-based grades. These results may be due to the circular cross-section assumption and the poor segmentations of the 4 problematic datasets.

6 Conclusions

A method that extracts lumen centerlines and estimates lumen local widths from vascular images was described. With the assumption of circular normal cross-section, we compiled carotid bifurcaiton lumen segmentations and computed ICA stenosis grading for the CLS competition of the 3rd MICCAI Workshop. The circular cross-section assumption leads to less satisfactory results obtained from our method than the human observers' one, which is not unexpected. Nonetheless, the results are fair—the Dice similarity index is > 0.7 indicating a good agreement with the reference standard, and the average mean surface distance is smallar than the voxel diagonal length. All these show that the outputs of our method does form a good basis for post-processing in which actual lumen boundary is located if accurate segmentation and stenosis grading is desired.

References

[1] J. J. Bartko. Measurement and reliability: statistical thinking considerations. *Schizophr Bull.*, 17(3):483–489, 1991. 5.1

[2] D. H. Douglas and T. K. Peucker. Algorithms for the reduction of the number of points required to represent a line or its caricature. *The Canadian Cartographer*, 10(2):112–122, 1973. 3.2

Table 3: Summary stenosis

Measure	%			rank		
	min.	max.	avg.	min.	max.	avg.
S_area	0.0%	100.0%	27.5%	1	4	3.29
S_diam	0.0%	100.0%	34.9%	1	4	3.26
Total (stenosis)				**1**	**4**	**3.27**

Table 4: Averages stenosis

Team name	area %	area rank	diam %	diam rank	Total rank
media	27.55	3.3	34.90	3.3	3.3
ObserverA	2.71	1.2	3.61	1.5	1.4
ObserverB	4.55	1.6	5.29	1.8	1.7
ObserverC	5.61	2.2	5.74	2.0	2.1

[3] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Med. Imag. Comput. Comput. Assist. Interv.*, volume 1496 of *Lecture Notes in Comput. Sci.*, pages 130–137. Berlin, Springer-Verlag, 1998. 3.2

[4] R. Hameeteman, M. Zuluaga, L. Joskowicz, M. Freiman, and T. van Walsum. Carotid lumen segmentation and stenosis grading challenge. Online, May 2009. 4.2, 5.1

[5] B. Kégl, A. Krzyzak, T. Linder, and K. Zeger. Learning and design of principal curves. 22(3):281–297, March 2000. 1, 3.4

[6] N. Otsu. A threshold selection method from grey-level histograms. 8:62–66, 1978. 3.1

[7] L. Weizman, M. Freiman, and L. Joskowicz. Implementation of weighted Dijkstra's shortest-path algorithm for n-D images. Insight Journal, January 2009. 4.3

[8] O. Wink, W. J. Niessen, and M. A. Viergever. Multiscale vessel tracking. 23(1):130–133, 2004. 3.2

[9] W. C. K. Wong and A. C. S. Chung. Principal curves for lumen center extraction and flow channel width estimation in 3-D arterial networks: theory, algorithm and validation. *Med. Imag. Anal.*, 2009. In Preparation. 1, 3.4