Coronary Lumen and Plaque Segmentation from CTA Using Higher-Order Shape Prior

Yoshiro Kitamura[1,2], Yuanzhong Li[1], Wataru Ito[1], and Hiroshi Ishikawa[2]

[1] Imaging Technology Center, Fujifilm Corporation, Tokyo, Japan
[2] Department of Computer Science and Engineering, Waseda University, Tokyo, Japan

Abstract. We propose a novel segmentation method based on multi-label graph cuts utilizing higher-order potentials to impose shape priors. Each higher-order potential is defined with respect to a candidate shape, and takes a low value if and only if most of the voxels inside the shape are foreground and most of those outside are background. We apply this technique to coronary lumen and plaque segmentation in CT angiography, exploiting the prior knowledge that the vessel walls tend to be tubular, whereas calcified plaques are more likely globular. We use the Hessian analysis to detect the candidate shapes and introduce corresponding higher-order terms into the energy. Since each higher-order term has any effect only when its highly specific condition is met, we can add many of them at possible locations and sizes without severe side effects. We show the effectiveness of the method by testing it on the standardized evaluation framework presented at MICCAI segmentation challenge 2012. The method achieved values comparable to the best in each of the sensitivity and positive predictive value, placing it at the top in average rank.

Keywords: Shape prior, Graph cuts, Multiple labels, Higher order function, Coronary stenosis.

1 Introduction

Cardiovascular disease is one of the largest causes of death in developed countries. Thanks to the recent progress of multi-detector CT techniques, 3D CT angiography has become a standard examination for the disease. Since coronary artery disease (stenosis) is caused by a narrowing of the arteries, accurate lumen segmentation is an important step toward determination of the degree of stenosis. Despite the great amount of past studies, computer aided detection and quantification of coronary stenosis remains a challenging task. A comprehensive analysis and categorization of vessel segmentation techniques can be found in [1] and a quantitative comparison of different algorithms in [2]. Coronary stenosis detection algorithms are classified into two types: classification-based and segmentation-based. Classification-based methods directly detect a plaque (or a stenosis) using extracted features [3], while most of segmentation-based methods aim at accurate lumen segmentation by estimation of healthy vessel diameter to detect stenosis [4,10,12]. On the other hand, early accumulation of plaques in coronary arteries is associated with compensatory enlargement of

P. Golland et al. (Eds.): MICCAI 2014, Part I, LNCS 8673, pp. 339–347, 2014.

vessel walls (positive remodeling). Therefore, segmenting vessel walls as well as lumen is a key to accurate assessment of stenosis grade. For this purpose, several methods such as the model-guided level-sets [5] have been proposed.

In this paper, we present a novel segmentation method based on the Conditional Random Field (CRF) framework. As the plaque tissue and surrounding tissues have nearly the same intensities, typical graph-cut methods utilizing first-order potentials tend to fail in distinguishing them. To overcome this challenge, we utilize higher-order functions that can model more complex structures than mere continuity. Kohli et al. [7] presented a framework utilizing higher-order functions that can combine multiple segmentations generated using unsupervised segmentation algorithms. Recently, Kadoury et al. [8] presented a segmentation approach in which higher-order potentials ensure regional consistency. In contrast, we take advantage of the prior knowledge regarding the shape of the vessel walls and calcified plaques.

Even when arteries are narrowed with (calcified or soft) plaques, their walls tend to be tubular, while calcified plaques are more likely globular. We use the Hessian analysis to detect such candidate shapes and introduce higher-order terms to encourage segmentation along them. Such a higher-order term takes a low value if and *only if* most of the voxels on the border stick together inside or outside. Since each higher-order term has an effect only when its highly specific condition is met, we can add many of them at possible locations and sizes without worrying about side effects.

We show the effectiveness of the method by testing it on the standardized evaluation framework presented at MICCAI segmentation challenge workshop in 2012.

2 Method

In this paper, we address the problem of segmentation of coronary lumen and plaques in order to quantify the stenosis grades of the coronary arteries. We assume that the centerlines of coronary arteries have been extracted and a stacked volume has been reconstructed from each centerline by generating multi-planar reconstructed (MPR) images perpendicular to the direction along the centerline. Our goal is to detect the double boundaries of vessel wall and lumen from the stacked volume (Fig. 1). Since a lumen boundary is always inside the vessel wall, we solve this as an ordered three-class labeling problem, where the three classes are {1: *Background*, 2: *Plaque*, 3: *Lumen*}. In the following, we first describe the general multi-label CRF energy setup. Then we explain the higher-order potentials in more specifics. Finally, we describe the unary and pairwise potentials as well as our stenosis detection algorithm.

2.1 Multi-label CRFs with Higher Order Potentials

First, we describe the multi-label Conditional Random Field (CRF) with higher-order potentials. We consider minimizing a function of the form:

$$E(\mathbf{y}) = \sum_{a \in V} \theta_a(y_a) + \sum_{a \in V, b \in N_a} \theta_{ab}(y_a, y_b) + \sum_{c \in C} \theta_c(\mathbf{y}_c), \quad (1)$$

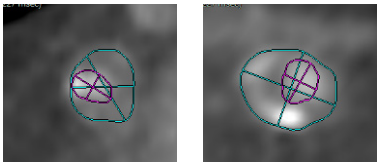

(a) (b)

Fig. 1. The sectional images of the vessel with (a) soft plaque, and (b) calcified plaque. The cyan and magenta lines represent the vessel-wall and the lumen boundaries, respectively.

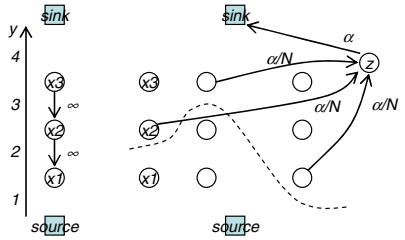

(a) (b)

Fig. 2. (a) Encoding of a multi label to binary variables. (b) The graph construction for higher order multi-label function. z is an auxiliary variable. This construction penalizes any labeling above the dashed line.

where **y** is the vector of variables y_i indexed by the set V of voxels taking values in a label set $L = \{1, 2, \ldots, l\}$ with a linear order and N_a is the set of a's neighbors. The functions θ_a and θ_{ab} are the potentials for the label y_a and the label pair (y_a, y_b), respectively, while the higher order potential $\theta_c(\mathbf{y}_c)$ depends on the variables y_i for voxels in its *clique* $c \subset V$. According to [6], the multi-label variables y_i can be encoded by $l-1$ binary variables $\{x_{1, \ldots,} x_{l-1}\}$. For instance, {1, 2, 3} can be mapped to the labelings of two binary labels $x_1x_2 = \{11, 01, 00\}$. In order to ensure a bijective encoding, very high costs are assigned to the unused labelings such as $x_1x_2 = \{10\}$ (see Fig. 2 (a)). Next, multi-label higher-order functions with the form:

$$\theta_c(\mathbf{y}_c) = \begin{cases} \alpha & \text{if } \exists i \in c : y_i < l_i, \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

can be converted to submodular quadratic pseudo Boolean functions, where α is a positive weight and the threshold label l_i can vary from voxel to voxel. This potential encourages every variable y_i in the clique to take the value equal to or more than l_i. Similarly, there is a family of functions that encourage the variables y_i to take the label less than l_i. Although eq. (2) is a "hard" function that takes two values depending on whether or not *all* the variables satisfy the condition, we can introduce the idea of the "robust P^n Potts model" [7] into this and let the function to take gradually larger value as the number of voxels that violate the condition increases. With a positive integer N, such a potential is formulated as:

$$\theta_c(\mathbf{y}_c) = \min\left\{\alpha, \sum_{i \in c : y_i < l_i} \frac{\alpha}{N}\right\}, \tag{3}$$

which gives an increasing penalty as more voxels violate the condition, until it saturates at N violating voxels. The graph construction in this case is shown in Fig. 2 (b).

2.2 Shape Constraints Using Higher Order Potentials

We now introduce a novel use of higher-order functions as shape prior in segmentation. The novelty is in the choice of the voxels in the clique c for $\theta_c(\mathbf{y}_c)$, which we call the shape term. The idea is to add the higher-order terms such as eq. (3) to encourage: a) all of the voxels inside a candidate shape to be foreground and, b) all of the voxels outside the shape to be background. For lumen and plaque segmentation, we utilize the prior knowledge that the vessel walls tend to be tubular, even with plaques, while calcified plaques are liable to be globular. Note that severely calcified plaques are not always globular. However, they tend to be brighter than other structures, enough to be distinguished based on the intensities (i.e. by the unary term).

Since the vessel wall is represented as the summed regions of *Lumen* and *Plaque* (labels 2 and 3, respectively), we encourage the segmentation to be tubular by setting the "soft" constraint $y_i < 2$ for voxels inside the tube, and $y_i \geq 2$ for voxels outside. Similarly, we encourage $y_i < 2$ and $y_i \geq 3$ for voxels inside the globular shape. Of course, we do not know the position and the size of the candidate shapes in advance. So we place the shape term cliques at various positions and scales, choosing the candidates using the Optimally Oriented Flux (OOF) [9]. The OOF computes the Hessian matrices at the boundaries of spherical regions. As is well known, their eigenvalues can be used to estimate the shape of the object, such as tubular or globular, and the eigenvector corresponding to the smallest eigenvalue indicates the orientation of the tubular object. Since the computation of the OOF is localized at the spherical surface, we can estimate the boundaries of the vessel walls accurately (Fig. 3). In our experiments, we computed the OOF at multiple kernel radii from 0.5 to 4.0mm in steps of scale 1.26. We create a clique out of the voxels just inside the circle regions on the plane orthogonal to the main axis, and create another clique from the voxels just outside the circle. This way, several to hundreds of voxels were selected in a clique, when the resolution of the reconstructed MPR image was 0.25mm.

The advantage of our higher-order terms is that, because each shape term does not affect the energy unless the voxels in its clique satisfy the specific condition, there is not severe side effect by including as many of them as desired, except for the computational cost. If instead we tried to influence the energy using pairwise terms by, for instance, placing many of them along the boundary, each term would independently affect the outcome in a way difficult to control. Also, by using the robust P^n Potts model, we can encourage not only the exact candidate shape but also similar shapes, without increasing auxiliary variables.

2.3 Learning Weights of Higher Order Potentials

Next, we describe how to set the weight α for the shape term. We learn the weight from the reference segmentation data that was manually prepared. Computing the OOF in the reference image, we obtain the objective label y_i at each $i \in c$ and the feature vector at each position and scale. Since the responses of the Hessian filters cannot discriminate complicated structures, to improve the prediction accuracy the

feature vector consists of twelve eigenvalues sampled at the same position but at different scales. We set the shape constraint using the robust P^n Potts model by setting the N in eq. (3) to 20% of the total number of voxels in the clique; thus the potential allows up to 20% of the voxels violating the assumed shape. For each of the two cases whether 80% of the labels are less than the certain label or not, we generate the histograms for the feature values and learn the log likelihood ratio of their probabilities:

$$-\log\left(\Pr_{c\in C}\left(\#\{i\in c \mid y_i < l_i\}\geq N\right)/\Pr_{c\in C}\left(\#\{i\in c \mid y_i < l_i\}< N\right)\right). \tag{4}$$

As a result, we confirmed that the shape terms regarding the vessel walls and calcified plaques were significantly predicted when the feature vectors represented tubular patterns or spherical patterns. The likelihood value corresponding to the feature value for given data is directly used as the weight of the shape term. Note that the weight α of the energy must be positive in order to keep the potential submodular. Then the shape terms are set at where the candidate shapes were detected. Fig. 4 is a visualization of the shape terms and the resulting segmentation.

2.4 Unary and Pairwise Potentials

The unary and pairwise potentials in eq. (1) are set as follows. First, we estimate the intensity distributions of lumen and two types of plaques (soft-plaque and calcified-plaque). The mean and the standard deviation of lumen (I_z^{lumen} and σ_z^{lumen}) and soft-plaque (I_z^{soft} and σ_z^{soft}) at the slice position z along the stacked volume are calculated from the upper sides and lower sides of the histogram of the regions determined by the initial centerline and radii (R_z). The mean value of calcified plaque is given by shifting the center to higher level: $I_z^{\text{calcified}} = I_z^{\text{lumen}} + 3\sigma_z^{\text{lumen}}$. Finally, $I_z^{\text{background}}$ is the mean of low intensity regions around the initial centerline.

Unary potentials: The unary potentials are given by

$$\theta_{\text{a}}(y_i = Lumen) \propto 1 - \exp\left(-\frac{(I_i - I_z^{\text{lumen}})^2}{2(\sigma_z^{\text{lumen}})^2}\right)\cdot h(D_i), \tag{5}$$

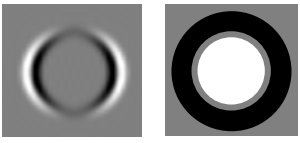

Fig. 3. (LEFT) A sectional plane of OOF filter kernel. (RIGHT) The higher-order potentials set according to the filter response. The voxels in the white and black regions are encouraged to be foreground and background, respectively.

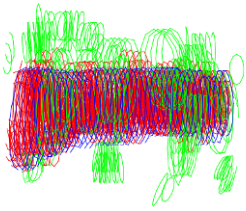

Fig. 4. The visualization of the shape terms. The red circles represent the active terms, i.e., most voxels inside them were labeled as plaque or lumen in the optimization. The green circles represent inactive terms. The blue lines correspond to the vessel-wall boundaries after optimiza-

where I_i is the intensity at voxel y_i, D_i is the distance of the voxel from the centerline, and $h(D_i)$ is 0 if $D_i < R_z$ and 1 otherwise. Similarly, the likelihood of the *Plaque*, $\theta_a(y_i = Plaque)$ takes a small value when I_i is close to I_z^{soft} or $I_z^{\text{calcified}}$. Finally, $\theta_a(y_i = Background)$ is set to take a small number at voxels with $D_i >> R_z$.

Pairwise Potentials: The pairwise potentials are defined as functions of the image gradient and the absolute intensity at each pair of 18-neighboring voxels:

$$\theta_{ij}(y_i, y_j) \propto \begin{cases} N(I_i - I_j, \sigma_g^2) \cdot (1 - N(I_i - (I_z^{\text{plaque}} + I_z^{\text{lumen}})/2), \sigma_z^2)/D_i & \text{, if } y_i < 3 \text{ and } y_j = 3 \\ N(I_i - I_j, \sigma_g^2) \cdot (1 - N(I_i - (I_z^{\text{background}} + I_z^{\text{plaque}})/2), \sigma_z^2)/D_i & \text{, if } y_i < 2 \text{ and } y_j \geq 2 \\ 0 & \text{, otherwise.} \end{cases} \quad (6)$$

I_z^{plaque} is I_z^{soft} if $I_i < I_z^{\text{lumen}}$ and $I_z^{\text{calcified}}$ otherwise. These potentials encourage label changes at voxels with strong image gradient or intensity close to the middle of the estimated distributions of each class. The weights between unary, pairwise and higher order potentials in eq. (1) were determined heuristically.

2.5 Stenosis Detection and Quantification Algorithm

We implemented automatic stenosis detection and quantification algorithm for verification with the standardized evaluation framework presented at MICCAI segmentation challenge workshop in 2012 [2]. The evaluation is done in three categories: 1) Stenosis detection, 2) Stenosis quantification, and 3) Lumen segmentation. For categories 1 and 2, participants are required to report 3D positions of detected stenoses and their stenosis grades that represent the diameter reduction between 0% (healthy) to 100% (occluded). Category 1 evaluates the sensitivity and positive predictive value (PPV) of the algorithm only for severe stenoses with diameter reductions more than 50%. Category 2 evaluates the differences of the estimated stenosis grades and the reference, which have been prepared for Quantitative coronary angiography (QCA) and CT angiography (CTA) by three observers.

We reported these measures in the following steps. We used extraction results obtained by the method in [11] as the initial input of coronary centerlines. After the lumen and plaque segmentation, short axis lengths of lumen (d_{lumen}) and vessel wall (d_{wall}) are computed at every position z along the centerline. Next, we calculated the ratio of the two measures: d_{lumen} / d_{wall} for CTA stenosis grades. On the other hand, QCA stenosis grades are defined as the ratio of minimal (narrowing) diameter and estimated normal lumen diameter (d_{normal}). We estimated d_{normal} as the average value of d_{lumen} that were sampled from 50 mm length along the centerline. For the category 3, the boundaries of lumen obtained by the proposed method were evaluated.

To reduce false detections, the following rules were considered in the stenosis detection. 1) The segments in the coronary tree were labeled according to the model matching method in [11]. Then, clinically relevant 18 segments defined by the Society of Cardiovascular Computed Tomography (SCCT) were evaluated. 2) Only z positions along the centerline from the ostia till 100mm length point were evaluated, excluding around branching points.

3 Results and Discussion

The summary of the three evaluation categories, comparing our method with other top participants for each metric as of May 2014 are shown in Table 1, 2 and 3, respectively. The processing time was at most 10 seconds per a centerline, depending on its length, on a quad-core 3.4GHz PC. As can be seen in Table 1, the method achieved a sensitivity of 51.1% and a PPV of 33.3% compared to CTA reference, which are almost equal to the best ones for each measure. And it achieved the best quantification performance compared to the CTA reference with kappa value: 0.32. Fig. 5 shows an example of the segmentation result compared to the reference. The boundaries obtained by the proposed method are close to observer's one. On the other hand, the quantification performance compared to the QCA was not very good. Our method underestimated the QCA stenosis grades because the normal lumen diameter calculated as an average value of local lumen diameters did not show significant decrease in some cases. As is discussed in [2], robust regression may improve the accuracy of QCA-based stenosis quantification.

An advantage of our method in stenosis detection is accurate detection of the vessel wall boundaries. Fig. 6 shows a comparison of segmentation results obtained with and without the shape terms. It can be seen that the shape terms improve the segmentation accuracy of vessel wall boundary, and successfully segment calcified regions. The method can improve stenosis quantification accuracy by avoiding the difficulty of estimating the normal diameter, and can also assess the positive remodeling effect accurately. A limitation of the method derives from the Hessian-based features, which cannot distinguish very wide variety of structures. For example, the shape constraints do not work well around branching points. The method also tends to over or under-segment where surrounding structures form a tubular pattern.

Table 1. Stenosis detection results compared to other top participants for each metric

Method	QCA Sensitivity		QCA PPV		CTA Sensitivity		CTA PPV		Avg. rank
	%	rank	%	rank	%	rank	%	rank	
Proposed method	35.7	12	32.3	6	51.1	6	**33.3**	**5**	**7.2**
Lugauer et al. [12]	60.7	6	24.6	7	46.8	9	25.0	10	8.0
Cetin et al. [3]	53.6	9	19.2	10	**53.2**	**5**	26.0	8	8.0
Wang. et al. [5]	25.0	14	**50.0**	**3**	10.6	17	**33.3**	**5**	9.8
Eslami et al. [10]	**67.9**	**4**	9.4	16	51.1	6	4	16	10.5

Table 2. Stenosis quantification results compared to other top participants for each metric

Method	QCA Avg. Abs. diff.		QCA R.M.S. diff.		CTA Weighted Kappa		Avg. rank
	%	rank	%	rank	K	rank	
Proposed method	30.8	5	36.6	5	**0.32**	**5**	**5.0**
Wang et al. [5]	**28.8**	**1**	**33.7**	**1**	0.18	10	5.5
Lugauer et al. [12]	49.0	11	55.1	12	0.30	6	8.8

Table 3. Lumen segmentation results compared to other top participants for each metric

Method	DICE *diseased*		DICE *healthy*		MSD *diseased*		MSD *healthy*		MAXSD *diseased*		MAXSD *healthy*		Avg. rank
	%	rank	%	rank	%	rank	%	Rank	%	rank	%	rank	
Lugauer et al. [12]	**0.74**	**3.6**	0.73	**4.3**	**0.35**	**3.9**	0.55	**4.6**	2.99	4.2	3.73	4.2	**4.2**
Mohr et al. [4]	0.70	4.9	0.73	4.5	0.40	5.5	**0.39**	4.7	**2.68**	**3.6**	**2.75**	**2.7**	**4.2**
Proposed method	0.71	5.2	**0.75**	4.6	0.42	6.2	0.41	5.4	3.17	5.4	3.69	4.3	5.1

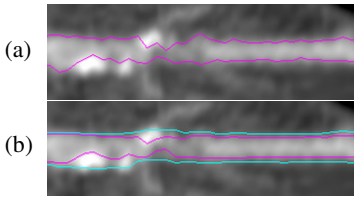

Fig. 5. The results in segment LAD7 in training dataset #08. (a) manually annotated reference of observer #1, (b) the result obtained by the proposed method.

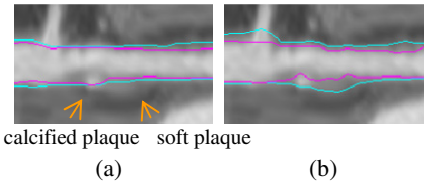

Fig. 6. Segmentation results for the stenosed artery obtained (a) without the shape terms, and (b) with the shape terms. The cyan and magenta lines represent the vessel-wall and the lumen boundaries, respectively.

4 Conclusion

We proposed a novel segmentation method for coronary lumen and vessel wall segmentation from CT angiography images. Our contribution is presenting effective use of higher-order potentials that enable us to introduce shape priors to a CRF energy. The verification tests for stenosis detection and quantification showed the state of the art performance. The method could be applied for general vessel segmentation. Extending its utility would be an interesting future work.

Acknowledgment. In the course of this work, H. Ishikawa was partially supported by the JSPS Grant-in-Aid for Scientific Research (B) #24300075 and the JSPS Grant-in-Aid for Challenging Exploratory Research #25540075.

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(6), 819–845 (2009)
2. Kirişli, H.A., Schaap, M., Metz, C., et al.: Standardized evaluation framework for evaluating coronary artery stenosis detection, stenosis quantification and lumen segmentation algorithms in computed tomography angiography. Medical Image Analysis 17(8), 859–876 (2013)
3. Cetin, S., Unal, G.: Automatic detection of coronary artery stenosis in CTA based on vessel intensity and geometric features. In: Proc. of MICCAI 2012 Workshop 3D Cardiovascular Imaging: A MICCAI Segmentation Challenge (2012)
4. Mohr, B., Masood, S., Plakas, C.: Accurate stenosis detection and quantification in coronary CTA. In: Proc. of MICCAI 2012 Workshop 3D Cardiovascular Imaging: A MICCAI Segmentation Challenge (2012)
5. Wang, C., Moreno, R., Smedby, Ö.: Vessel segmentation using implicit model-guided level sets. In: Proc. of MICCAI 2012 Workshop 3D Cardiovascular Imaging: A MICCAI Segmentation Challenge (2012)

6. Ramalingam, S., Kohli, P., Alahari, K., Torr, P.H.S.: Exact inference in multi-label CRFs with higher order cliques. In: Proc. of CVPR (2008)
7. Kohli, P., Ladicky, L., Torr, P.H.S.: Robust higher order potentials for enforcing label consistency. In: Proc. of CVPR (2008)
8. Kadoury, S., Abi-Jaoudeh, N., Valdes, P.A.: Higher-Order CRF Tumor Segmentation with Discriminant Manifold Potentials. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part I. LNCS, vol. 8149, pp. 719–726. Springer, Heidelberg (2013)
9. Law, M.W.K., Chung, A.C.S.: Three Dimensional Curvilinear Structure Detection Using Optimally Oriented Flux. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part IV. LNCS, vol. 5305, pp. 368–382. Springer, Heidelberg (2008)
10. Eslami, A., Aboee, A., Hodaei, Z., Moghaddam, M.J., Carlier, S., Katouzian, A., Navab, N.: Quantification of coronary arterial stenosis by inflating tubes in CTA images. In: Proc. of MICCAI 2012 Workshop 3D Cardiovascular Imaging: A MICCAI Segmentation Challenge (2012)
11. Kitamura, Y., Li, Y., Ito, W.: Automatic coronary extraction by supervised detection and shape matching. In: Proc. of ISBI, pp. 234–237 (2012)
12. Lugauer, F., Zhang, J., Zheng, Y., Hornegger, J., Kelm, M.: Improving Accuracy in Coronary Lumen Segmentation via Explicit Calcium Exclusion, Learning-based Ray Detection and Surface Optimization. In: Proc. of SPIE MI, 90343U (2014)