

Precise Lumen Segmentation in Coronary Computed Tomography Angiography

Felix Lugauer^{1,2*}, Yefeng Zheng³, Joachim Hornegger¹, and B. Michael Kelm²

¹ Pattern Recognition Lab, University of Erlangen-Nuremberg, Germany

² Imaging and Computer Vision, Siemens AG, Corporate Technology, Erlangen, Germany

³ Imaging and Computer Vision, Siemens Corporate Research, Princeton, USA

Abstract. Coronary computed tomography angiography (CCTA) allows for non-invasive identification and grading of stenoses by evaluating the degree of narrowing of the blood-filled vessel lumen. Recently, methods have been proposed that simulate coronary blood flow using computational fluid dynamics (CFD) to compute the fractional flow reserve non-invasively. Both grading and CFD rely on a precise segmentation of the vessel lumen from CCTA. We propose a novel, model-guided segmentation approach based on a Markov random field formulation with convex priors which assures the preservation of the tubular structure of the coronary lumen. Allowing for various robust smoothness terms, the approach yields very accurate lumen segmentations even in the presence of calcified and non-calcified plaques. Evaluations on the public Rotterdam segmentation challenge demonstrate the robustness and accuracy of our method: on standardized tests with multi-vendor CCTA from 30 symptomatic patients, we achieve superior accuracies as compared to both state-of-the-art methods and medical experts.

Keywords: CCTA, lumen segmentation, Markov random field, tubular surface

1 Introduction

Coronary artery disease (CAD) is a leading cause of death in the western world according to the American Heart Association [2]. CAD is indicated by the build-up of coronary plaque which is accompanied by an inflammatory process in the vessel wall. It may result in a local narrowing of the lumen, known as stenosis, which in turn may cause an ischemic heart failure.

The diagnostic standard in current clinical practice is invasive coronary angiography (ICA) which requires catheterization. Besides the degree of stenosis, pressure differences across lesions can be measured under induced hyperemia. Fractional flow reserve (FFR) is computed as the pressure ratio and has been shown to be indicative for ischemia-causing stenoses. However, the procedure is costly and involves considerable risks and inconvenience for the patient.

*The author has been with Siemens Corporate Technology for this work.

Alternatively, with coronary computed tomography angiography (CCTA) a volumetric image of the contrasted coronary vessels is acquired which allows for a non-invasive identification and analysis of coronary stenoses. Usually, stenoses with grades (degree of anatomical obstruction) above 50% are considered hemodynamically relevant. Recently, the non-invasive computation of FFR has been proposed as an alternative to ICA. This is achieved by simulating coronary blood flow using computational fluid dynamics (CFD) based on lumen geometry extracted from CCTA. Latest studies show that CFD provides a better accuracy in identifying ischemia-causing stenoses than the anatomical grade [8]. Both grading and CFD heavily depend on the accuracy of the extracted coronary lumen. As manual segmentation is very time-consuming and prone to high variability among medical experts, automatic and robust methods are desirable.

There exists a substantial body of research on vessel extraction and segmentation algorithms [5]. For brevity, we only mention related methods and approaches that have been applied to coronary artery segmentation using the standardized evaluation framework proposed by Kirişli *et al.* [4]. Previous approaches tend to focus on healthy vessels and often fail to correctly segment the lumen in diseased vessels, *i.e.* in the presence of calcified, non-calcified and mixed plaques. This can be partially accounted for by explicitly modeling or suppressing calcified plaque before segmentation [7,9,10]. Many of the proposed approaches also involve post-processing and refinement of the segmentation results for fixing artifacts. Shazad *et al.* [10], for example, propose a voxel-based graph-cut segmentation followed by a radial resampling and smoothing which patches non-tubular segmentation results (*e.g.* dissected lumen segmentations) and artifacts due to the voxel-level accuracy. A shape prior as included in the model-based level set approach proposed by Wang *et al.* [12] helps to avoid leakages and allows for a robust lumen extraction even in low contrast (ambiguous) regions. Similarly, Mohr *et al.* [9] employ a level-set approach based on results from a tissue classification and calcium segmentation step. Lugauer *et al.* [7] apply boundary detection and calcium removal within the segmentation approach of Li *et al.* [6] which strictly enforces topological constraints. However, this closed-set formulation only allows to apply a particular (“ ϵ -insensitive”) smoothness prior, which is not well-suited for suppressing the high amount of noise encountered in coronary lumen segmentation. Post-smoothing of the segmentation results is thus suggested in [7].

For an optimal accuracy, all prior information should be considered by the segmentation approach and post-processing should not be required. We propose a method that addresses the above mentioned issues by (1) explicitly accounting for calcified-plaque, (2) enforcing tubular structure, and (3) providing robust surface regularization suitable for coronary lumen segmentation by adopting a general Markov random field (MRF) [3] formulation to tubular segmentation.

2 Method

A set of centerlines (CTLs) approximately describing the curve of the coronary artery tree initializes the proposed surface extraction method. The lumen ge-

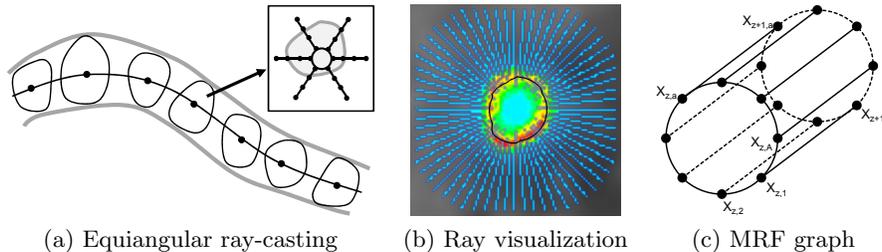


Fig. 1: (a) Radial candidate positions along rays in cross-sectional slices orthogonal to the given centerline. (b) Associated boundary probabilities for rays within a slice. (c) Cylindrical element of a tubular MRF graph where each node represents the selection of a candidate position along a certain ray and edges implement smoothness priors.

ometry is modeled on the basis of generalized cylinders with radially varying contours along the given centerline for each vessel segment. A discrete number of radial candidate positions along equiangular rays in orthogonal slices (Fig. 1a) are constructed and represented as random variables of a corresponding MRF graph (Fig. 1c). Using learning-based boundary detectors (Fig. 1b) and robust semi-convex regularization terms, this formulation allows for an optimal surface segmentation within polynomial time using well-known min-cut/max-flow solvers. We first describe the learning-based boundary detection model, followed by details on the MRF-based optimal surface generation.

2.1 Detection Model

The probability of a lumen wall is estimated in every cross-sectional slice for a vessel segment comprising Z slices using a fixed number of A equiangular rays with R equally-spaced (δ_r) radial candidate positions. While fixed terms based on image gradient and structural information could be used [6], we propose a learning-based probabilistic estimate based on steerable features. These are a collection of low-level image features sampled on a ray-oriented pattern [13]. Probabilistic boosting trees [11] are trained on manually segmented lumen contours by bootstrap aggregation and tested at every of the $Z \cdot A \cdot R$ candidate positions to estimate the probability $P_B(z, a, r)$ of lying on the lumen boundary. In order to improve the detection of boundaries from diseased lumen, the same amount of training samples were randomly selected from healthy and diseased (calcified or soft-plaque) lumen tissue. A calcium removal step, as proposed in [7], ensures the exclusion of calcified plaque from the segmented lumen by modifying the boundary probabilities. Essentially, every ray is analyzed for intersections with calcium and, if found, boundary probabilities are exponentially damped at these positions in outbound direction.

Since vessel segments were processed in parallel for performance reasons, enforcing matching contours at vessel furcation points would require a tree topology scheduled processing which has been omitted for simplicity.

2.2 Optimal Surface Generation

Finding the lumen surface is now cast as a combinatorial optimization problem: for every ray an optimal candidate position has to be selected. The trade-off between image-based likelihood (*w.r.t.* the boundary probabilities) and the prior assumption that the true lumen surface is smoothly varying, is best expressed with a first order Markov random field.

MRF Formulation. Let $G = (V, E)$ denote an undirected graph with a set of vertices V and undirected edges $E \subset V \times V$. Each vertex $v \in V$ is associated with a multivariate random variable $X_v = X_{z,a} = r$ where $r \in [1 \dots R]$ describes its state while a and z denote the angular position and the slice of a certain ray. Edges are constructed according to the generalized cylinder model within and across contour slices as depicted in Fig. 1c. For computational efficiency, we propose to use a first order MRF energy formulation as in Ishikawa [3]

$$E(X) = \sum_{(u,v) \in E} \gamma_{uv} g(X_u - X_v) + \sum_{v \in V} h(v, X_v), \quad (1)$$

where $g(\cdot)$ is a convex function of the label difference of vertex u and v . The first sum represents the smoothness prior, whereas the second sum, the data term, incorporates the observations (*i.e.* the boundary likelihood). The constant edge weight factors γ_{uv} have been used to trade between intra-slice γ_a and inter-slice smoothness γ_z . Energy functions of this form can be minimized exactly within polynomial time by a *min-cut/max-flow* algorithm [3]. As convex priors that penalize label differences $d = X_u - X_v$ across edges, we considered

$$g(d) = \beta |d| \quad (\text{L1-norm}) \quad (2)$$

$$g(d) = \beta \begin{cases} 0 & |d| \leq \alpha \\ (|d| - \alpha) & |d| > \alpha \end{cases} \quad (\epsilon\text{-insensitive for } \beta \gg 1) \quad (3)$$

$$g(d) = \beta \begin{cases} (d/\alpha)^2 & |d| \leq \alpha \\ 2|d/\alpha| - 1 & |d| > \alpha \end{cases} \quad (\text{Huber}) \quad (4)$$

with threshold and slope parameters α and β (see Fig. 4a). While any kind of convex function could be used, the above three are especially well suited due to their robust nature and computational advantages which will be discussed later.

Flow Graph Construction. As shown in [3], a transformation of the undirected graph G into a directed graph H (flow graph) results in an s - t min-cut problem which can be solved exactly by min-cut/max-flow algorithms. The flow graph $H = (V_H, E_H)$ is derived from G such that for any vertex $v \in V$, $R - 1$ vertices $V_{z,a,r} \in V_H$ are created. Two special vertices s, t mark the source and the sink of the flow graph summing up to $Z \cdot A \cdot (R - 1) + 2$ vertices in total. H consists of directed edges $(u, v) \in E_H$ (from vertex u to v) with positive capacity $c(u, v)$.

The graph is constructed in a way such that every s - t cut in H (a cut separating source and sink) corresponds to a configuration (variable assignment) $X \in \mathcal{X}$ of the MRF and the cost of the cut—sum of all edge capacities $c(u, v)$ in the cut—is the cost of this configuration according to (1). Note that, differing from [3], we omit a layer of superfluous vertices in our description (by merging data- and constraint-edges of the first of R layers which connects the source and $Z \cdot A$ vertices).

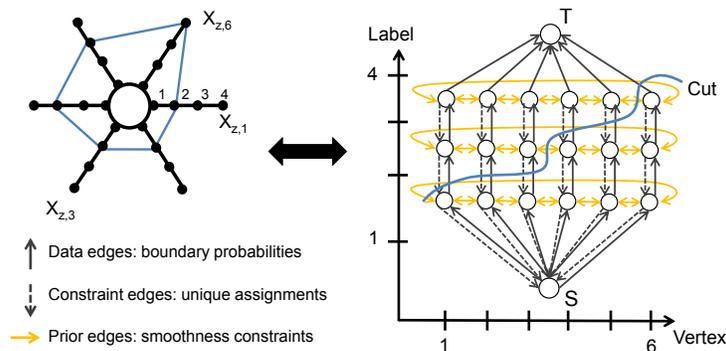


Fig. 2: Exemplary 2D flow graph for a single slice ($a = 6$ rays, $r = 4$ radial candidates) using the L1-norm prior(2). The cut represents an MRF configuration and, thus, uniquely describes a particular contour ($x_{z,1-3} = 2, x_{z,4-5} = 3, x_{z,6} = 4$). By analogy, a cut plane in the 3D flow graph defines a vessel surface geometry.

Potentials of the MRF are incorporated using three types of edges. *Data Edges* implement the observations $h(v, X_v)$ by sink-directed (vertical) edges with capacities equal to the negative log-likelihood of the boundary probabilities, while *Constraint Edges* impose a unique assignment by source-directed edges with infinite capacities. *Prior Edges* incorporate the prior terms $\gamma_{uv}g(X_u - X_v)$ of (1). Their capacities are second-order differences of the convex prior function (2)-(4) and are computed as

$$\text{cap}(d) = \frac{g(d+1) - 2g(d) + g(d-1)}{2}, \quad (5)$$

where $d = r_1 - r_2$ is the radial offset between vertices that represent the variables $X_{z,a_1} = r_1$ and $X_{z,a_2} = r_2$, *i.e.* the edge $(V_{z,a_1,r_1}, V_{z,a_2,r_1})$ in H (see Fig. 2 for an exemplary construction in 2D). Prior edges must be considered in two directions and potentially between all pairs of vertices in two neighboring columns. But since edges with capacity $\text{cap}(d) = 0$ can be omitted, prior terms with linear ranges result in fewer flow graph edges and, ultimately, in a faster min-cut/max-flow computation. The proposed prior functions (2)-(4) exhibit linear ranges (Fig. 4a) of varying extent. That means, the computational complexity of the proposed priors tends to increase from L1-norm over ϵ -insensitive to Huber dependent on the function parametrization (α).

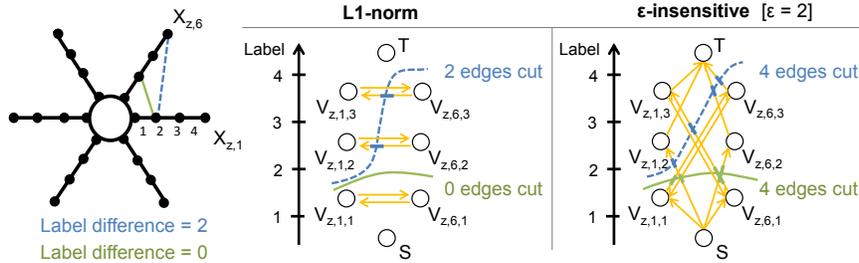


Fig. 3: Prior edges are only added for label differences with non-zero second-order differences of the prior function. Using the L1-norm and ϵ -insensitive functions yields flow graphs with varying complexity. The resulting costs for label variations (difference of 0 and 2) are visualized by two example cuts. For clarity, only prior edges are drawn.

Fig. 3 compares the flow graphs resulting from the L1-norm (2) and the ϵ -insensitive (3) priors for an MRF edge between $X_{z,6}$ and $X_{z,1}$. The L1-norm prior induces costs that grow linearly with the label difference. While there is no penalty difference for the ϵ -insensitive ($\epsilon = 2$) prior between a label difference of $d = 2$ (blue) and $d = 0$ (green), higher label differences would induce higher costs. Thus, choosing the slope parameter β sufficiently large (to dominate any data likelihood term) yields a regularizer like the one implicitly assumed by Li *et al.* [6]. The impact of that choice on the segmentation result is visualized in Fig. 4. While the L1-norm and Huber priors yield smooth, nearly circular contours, applying the ϵ -insensitive prior results in contours which are highly varying and more susceptible to noise (data-affine).

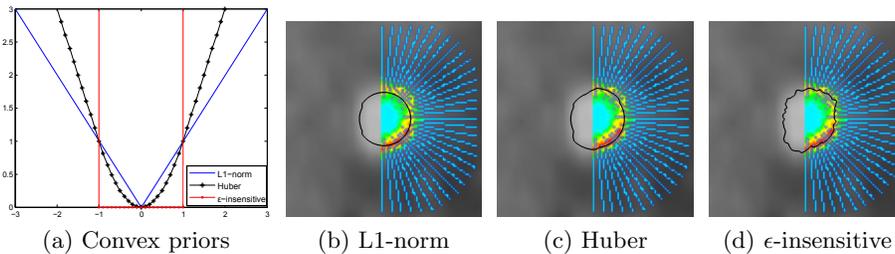


Fig. 4: (a) Different convex penalty functions can trade noise-robust (b) & (c) for data-affine (d) segmentations (right half colored with the boundary likelihood in blue-to-red).

3 Experimental Results

The segmentation accuracy of the proposed method was tested using the publicly accessible standardized coronary artery evaluation framework which allows

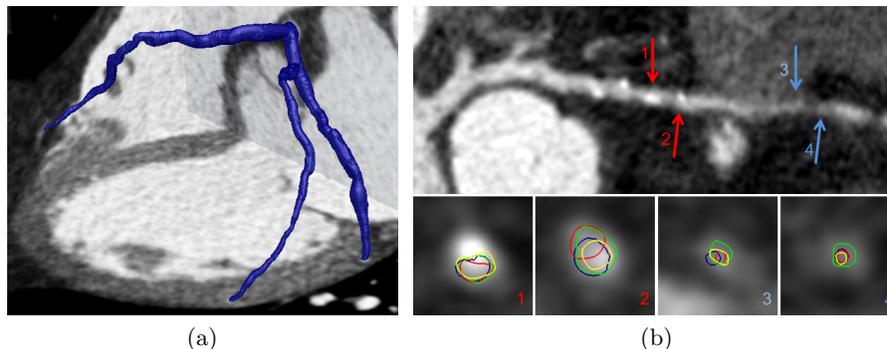


Fig. 5: (a) Surface visualization of the segmented left main coronaries with the LCX and LAD branch of the training dataset #10. (b) The upper part of the LAD is affected by mild mixed plaque while the lower end is narrowed by moderate soft-plaque (CPR, top). (b) Cross-sectional views with expert annotations in green, red and yellow along with our segmentation (blue) through mixed (1,2) and soft plaque (3,4). The proposed method effectively minimizes the variability within expert annotations.

for the comparison with current segmentation algorithms on 48 multi-protocol/-vendor CTA image volumes [4]. Annotations by three medical experts were provided for the first 18 training datasets which were used for training of the boundary detectors. The hyper-parameters were empirically chosen and in particular the Huber function parameters were a trade-off between graph-cut-induced computation time and accuracy *w.r.t.* an independent set of reference annotations: slice distance 0.3 mm, $A = 64$ rays with $R = 50$ candidates at $\delta_r = 0.1$ mm, Huber prior (4) with $\alpha = \beta = 1$ and $\gamma_a = 0.5$, $\gamma_z = 0.1$. The graph-cut solver was implemented in C++ as described in the original work [1]. We used CTLs obtained from the automatic tracking algorithm described by Zheng *et al.* [14].

Qualitative results are presented for pathological data from the training set in Fig. 5 and from the testing dataset in Fig. 6. Fig. 5a shows a surface geometry view of the left coronaries using the proposed segmentation while Fig. 5b shows a CPR of the—by mixed and soft plaque affected—LAD along with cross-sectional views through these lesions. The expert contours which were annotated from the same set of centerlines vary considerably while our contours (blue) rather minimizes the variability towards a consensus segmentation (as all annotations were used for training). In Fig. 6 segmented vessels of two patients with calcified and non-calcified plaques are shown. The lumen is smoothly segmented while calcified plaques are precisely circumscribed.

In Table 1, the maximum/mean surface distance (MAXSD/MSD) and the volumetric overlap (DICE) report the accuracies separately for healthy (H) and diseased (D) arteries averaged over the 30 testing/18 training datasets (hereafter we refer to testing results). The evaluation framework computes these measurements separately *w.r.t.* annotations of the three medical experts and averages

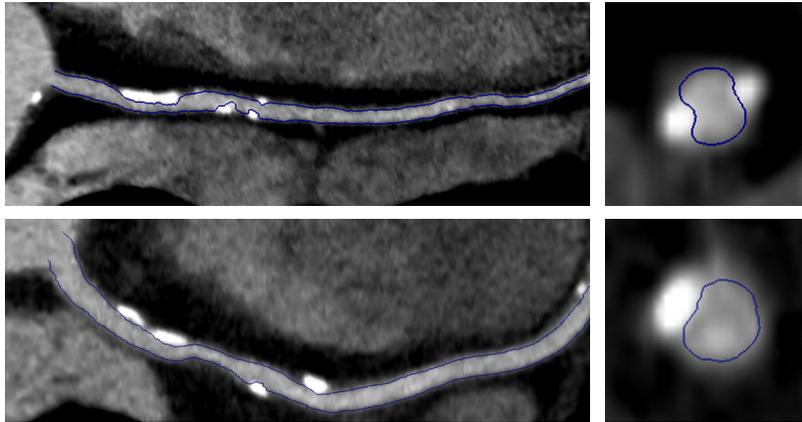


Fig. 6: Segmented lumen of calcified arteries for two patients as obtained by the proposed algorithm visualized from different projections (left: CPR, right: cross-section).

them [4]. The proposed method achieves the best rank¹ compared to all other automatic methods and ranks higher than two of the three medical experts. Among the automatic methods it performs best (results marked bold) on diseased vessel segments for all error measures². Since only the main vessels (LM, LAD, RCA, LCX) were compared with the reference, a labeling mismatch between our applied CTL tracing and the reference causes a large error bias for datasets (20, 28, 30, 31, 32, 41, 47) (can partly be seen in Fig. 7) and deteriorates the MSD and MAXSD error measures in particular for the healthy vessels. Our training results are consistent (only slightly better) with those for testing which indicates that no over-fitting occurred.

While manual segmentation requires substantial expertise and time and is prone to high inter- and intra-user variability, our automatic method yields a robust segmentation in under a minute per patient. The box-and-whiskers plots in Fig. 7 show that our method performs equal or even better than the best medical expert on most datasets. Apparently, the consensus learned from the annotations of the three experts yields superior performance when compared on the testing data. Our method effectively averages individual annotation biases during training and, thus, avoids the considerable inter-user variability seen in the expert annotations.

¹Rank denotes the performance in comparison to all other participants where a rank of 1.0 means that this method yields the best measures for all subjects and vessels.

²Latest results can be found at: <http://coronary.bigr.nl/stenoses/results/results.php>

Table 1: Three error measures are reported separately for diseased (D) and healthy (H) vessels (boldface marks best among automatic methods). Rank lists the overall segmentation ranking compared to all participating methods (*w.r.t.* testing results). Measures are averaged over 30 testing (18 training) datasets (listed testing/training).

Method	Rank	DICE	DICE	MSD	MSD	MAXSD	MAXSD
	avg. ↓	D [%]	H [%]	D [mm]	H [mm]	D [mm]	H [mm]
<i>Expert3</i>	3.5/4.5	79/76	81/80	.23/.24	.21/.23	3.00/3.07	3.45/3.25
Proposed	3.8/4.0	76/75	75/77	.32/.27	.51/.32	2.47/1.96	3.67/2.79
<i>Expert1</i>	4.4/5.4	76/74	77/79	.24/.26	.24/.26	2.87/3.29	3.47/3.61
Lugauer [7]	4.5/5.2	74/72	73/74	.35/.28	.55/.35	2.99/2.02	3.73/2.88
Mohr [9]	4.5/5.6	70/73	73/75	.40/.29	.39/.45	2.68/1.87	2.75/3.73
<i>Expert2</i>	6.1/7.3	65/66	72/73	.34/.31	.27/.25	2.82/2.70	3.26/3.00
Shahzad [10]	6.4/7.9	65/66	68/70	.39/.37	.41/.32	2.73/2.49	3.20/3.04
Wang [12]	6.9/9.0	69/68	69/72	.45/.43	.55/.56	3.94/4.06	6.48/5.23

4 Conclusion

A novel method for segmenting the lumen of the coronary arteries in computed tomography angiography has been proposed. While enforcing a tubular structure of the segmentation, the approach allows for a flexible choice of robust smoothness priors. Combined with a learning-based boundary detection, excellent performance is achieved on public challenge data [4]. Our analysis shows that the automatic segmentation results are as accurate as those from medical experts. While the proposed method is evaluated for coronary vessels only, it is readily applicable to other tubular structures. An extension to evolving arbitrary surfaces under local smoothness constraints is easily possible. The precise segmentation of our new approach will improve automatic stenosis detection and enable an improved non-invasive simulation of coronary blood flow, which is left to be evaluated in future work.

References

1. Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004.
2. A.S. Go et al. Heart disease and stroke statistics–2014 update a report from the american heart association. *Circulation*, 129(3):e28–e292, 2014.
3. H. Ishikawa. Exact optimization for Markov random fields with convex priors. *IEEE PAMI*, 25(10):1333–1336, 2003.
4. H.A. Kirişli, M. Schaap, C.T. Metz, A.S. Dharampal, W.B. Meijboom, 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.
5. D. Lesage, E.D. Angelini, I. Bloch, and G. Funka-Lea. A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical image analysis*, 13(6):819–845, 2009.

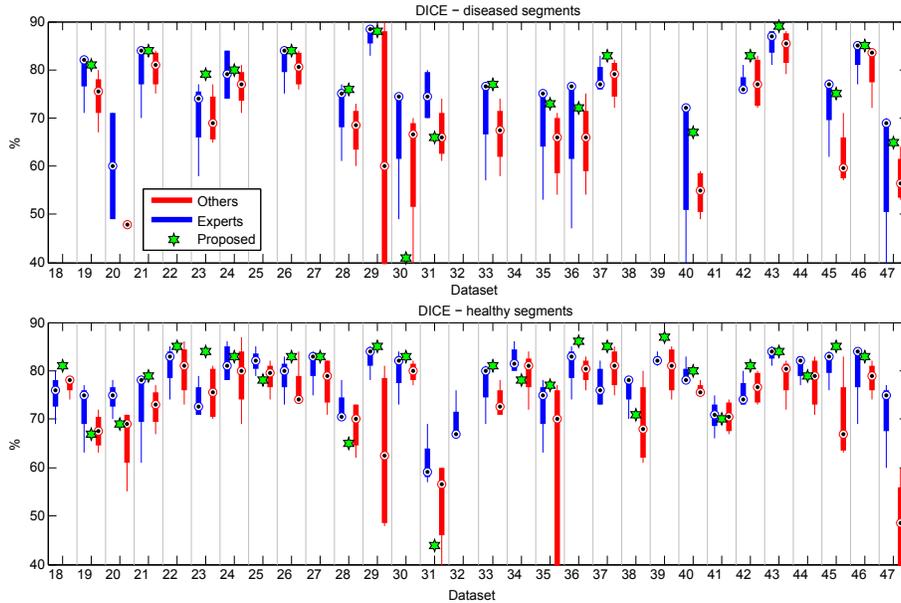


Fig. 7: Inter-user variability boxplots between the three medical experts and the four best-ranking methods (others) in comparison to our results (star) *w.r.t.* DICE measure for each subject of the testing data. Circles indicate medians, the box edges the 25th and 75th percentiles and whiskers extend to outliers.

6. K. Li, X. Wu, D. Chen, and L. Sonka. Optimal surface segmentation in volumetric images—a graph-theoretic approach. *IEEE PAMI*, 28(1):119–34, January 2006.
7. F. Lugauer, J. Zhang, Y. Zheng, J. Hornegger, and B.M. Kelm. Improving accuracy in coronary lumen segmentation via explicit calcium exclusion, learning-based ray detection and surface optimization. In *Proc. SPIE Conf. Medical Imaging*, 2014.
8. M.F.L. Meijs et al. CT fractional flow reserve: the next level in non-invasive cardiac imaging. *Netherlands Heart Journal*, 20(10):410–418, October 2012.
9. B. Mohr, S. Masood, and C. Plakas. Accurate lumen segmentation and stenosis detection and quantification in coronary CTA. In *Proceedings of 3D Cardiovascular Imaging: a MICCAI Segmentation Challenge Workshop*, 2012.
10. R. Shahzad, H. Kirişli, C. Metz, H. Tang, M. Schaap, L. van Vliet, W. Niessen, and T. van Walsum. Automatic segmentation, detection and quantification of coronary artery stenoses on CTA. *Int J Cardiovasc Imaging*, 29(8):1847–1859, 2013.
11. Z. Tu. Probabilistic boosting-tree: learning discriminative models for classification, recognition, and clustering. In *Tenth IEEE International Conference on Computer Vision, ICCV’05*, volume 2, pages 1589–1596. IEEE, 2005.
12. C. Wang, R. Moreno, and Ö. Smedby. Vessel segmentation using implicit model-guided level sets. In *Proceedings of 3D Cardiovascular Imaging: a MICCAI Segmentation Challenge Workshop*, 2012.
13. Y. Zheng, A. Barbu, B. Georgescu, M. Scheuring, and D. Comaniciu. Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features. *Medical Imaging, IEEE Transactions on*, 27(11):1668–1681, 2008.

14. Y. Zheng, H. Tek, and G. Funka-Lea. Robust and accurate coronary artery centerline extraction in CTA by combining model-driven and data-driven approaches. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2013*, pages 74–81. Springer Berlin Heidelberg, 2013.