

Improving Accuracy in Coronary Lumen Segmentation via Explicit Calcium Exclusion, Learning-based Ray Detection and Surface Optimization

Felix Lugauer^{a,b,*}, Jingdan Zhang^c, Yefeng Zheng^c,
Joachim Hornegger^a, and B. Michael Kelm^b

^a Pattern Recognition Lab, Department of Computer Science,
University of Erlangen-Nürnberg, Erlangen, Germany

^b Imaging and Computer Vision, Siemens AG, Corporate Technology, Erlangen, Germany

^c Imaging and Computer Vision, Siemens Corporation, Corporate Technology, Princeton, USA

ABSTRACT

Invasive cardiac angiography (catheterization) is still the standard in clinical practice for diagnosing coronary artery disease (CAD) but it involves a high amount of risk and cost. New generations of CT scanners can acquire high-quality images of coronary arteries which allow for an accurate identification and delineation of stenoses. Recently, computational flow dynamics (CFD) simulation has been applied to coronary blood flow using geometric lumen models extracted from CT angiography (CTA). The computed pressure drop at stenoses proved to be indicative for ischemia-causing lesions, leading to non-invasive Fractional Flow Reserve derived from CTA (FFRCT). Since the diagnostic value of non-invasive procedures for diagnosing CAD relies on an accurate extraction of the lumen, a precise segmentation of the coronary arteries is crucial. As manual segmentation is tedious, time-consuming and subjective, automatic procedures are desirable. We present a novel fully-automatic method to accurately segment the lumen of coronary arteries in the presence of calcified and non-calcified plaque. Our segmentation framework is based on three main steps: boundary detection, calcium exclusion and surface optimization. A learning-based boundary detector enables a robust lumen contour detection via dense ray-casting. The exclusion of calcified plaque is assured through a novel calcium exclusion technique which allows us to accurately capture stenoses of diseased arteries. The boundary detection results are incorporated into a closed set formulation whose minimization yields an optimized lumen surface. On standardized tests with clinical data, a segmentation accuracy is achieved which is comparable to clinical experts and superior to current automatic methods.

Keywords: lumen segmentation, calcium suppression, coronary artery segmentation, CAD, CTA

1. INTRODUCTION

Cardiovascular diseases, among which coronary artery disease (CAD) is most common, are the leading cause of death worldwide according to the World Health Organization. The standard procedure for diagnosing CAD is catheterization which is invasive and involves high risk and costs. However, non-invasive computed tomography angiography (CTA) of the coronary arteries can effectively rule out obstructive CAD. Newer generations of CT acquisition devices enable a diagnostically conclusive accuracy similar to catheterization. Even in cases of severe coronary calcification, sensitivity and specificity of CTA in comparison to invasive angiography for significant stenosis are high [1]. Furthermore, automatic systems have been developed that enable the detection and grading of coronary stenoses based on accurate lumen estimations [2]. Recent studies also show that the blood pressure analysis of coronary arteries is a better predictor for ischemia-causing lesions than anatomically visible lumen obstructions. The Fractional Flow Reserve (FFR), which is the blood pressure decline at a coronary stenosis, is typically measured by an invasive and costly catheterization. Only recently

*Further author information:

E-mail: felix.lugauer@i5.informatik.uni-erlangen.de




Figure 1: Coronary artery segmentation from CTA as obtained with our proposed method.

it is derived with similar accuracy from simulations on segmented coronary lumen from CTA data (FFRCT) [3, 4]. Thus, both the automatic detection of vessel obstructions as well as the blood flow simulation require an accurate segmentation of the coronary arteries (see Fig. 1). However, a manual segmentation is very time consuming (up to an hour per patient) making automatic segmentation desirable.

Extracting the lumen robustly and with high precision is a challenging task as even the main coronary arteries have diameters of only a few millimeters. In addition, the contrast of the vessel lumen, which is visually enhanced by contrast dye, is in between the contrast of non-calcified and calcified plaque which complicates an accurate distinction. The segmentation of diseased coronary lumen obtained by existing segmentation algorithms tends to be misled by calcifications leading to over- or under-estimation of lesions. However, ensuring that plaque is properly excluded from the segmentation is of paramount importance in order to allow for a diagnostically conclusive stenosis detection. Therefore, we developed a new fully-automatic lumen segmentation algorithm that is able to extract the lumen surface reliably and accurately. A novel calcium exclusion approach ensures a correct segmentation in the presence of calcified tissue.

2. RELATED WORK

There is a vast amount of literature on the segmentation of coronary arteries. Thus, we refer to the survey paper of Lesage *et al.* [5] for a comprehensive overview on numerous lumen segmentation techniques including a variety of vessel models and features. Within scope of our work are recent algorithms whose performance is assessed by standardized tests [6] and, thus, can directly be compared to our proposed method (see Sec. 4). The best-ranking methods of the coronary lumen segmentation challenge are level-set [7, 8] and graph-cut [9] based algorithms. The semi-automatic algorithm by Shahzad *et al.* [9] implements a graph-cut approach using probabilities derived from sampling around the centerline to generate the lumen segmentation. The typical graph-cut-induced problem of generating surfaces that are under-segmented is mitigated by robust kernel regression. However, the lack of a special calcium treatment reflects in an inferior accuracy for diseased arteries (compare Sec. 4.3). Wang *et al.* [8] use level-sets guided by a vessel-model that alternates between evolving the level-set and re-estimating the vessel diameter in order to update the vessel model. As the contour propagation is regulated on the basis of a single radius per cross-section, deformed or sharply changing contours that occur in the presence of plaque might not be well captured. The best-performing challenge algorithm uses level-sets with additional tissue classification. Similar to our approach, Mohr *et al.* [7] use an explicit calcium exclusion based on a data-driven intensity threshold.

3. LUMEN SEGMENTATION FRAMEWORK

Our segmentation algorithm yields the 3D lumen surface of the coronary arteries in three main stages plus a pre-processing where a set of previously traced arterial centerlines [10] serves as input (see Fig. 2). A region of interest around the initial centerline path (warped volume) is extracted in a pre-processing step (Sec. 3.1). In the first stage, potential lumen boundary positions are detected by casting rays in a cylindrical fashion. The boundary likelihood is estimated along each ray using a robust machine learning algorithm (Sec. 3.2). Second, a new calcium exclusion approach is applied to the potential boundary position. These are analyzed for calcified-plaque along the ray direction and are excluded from

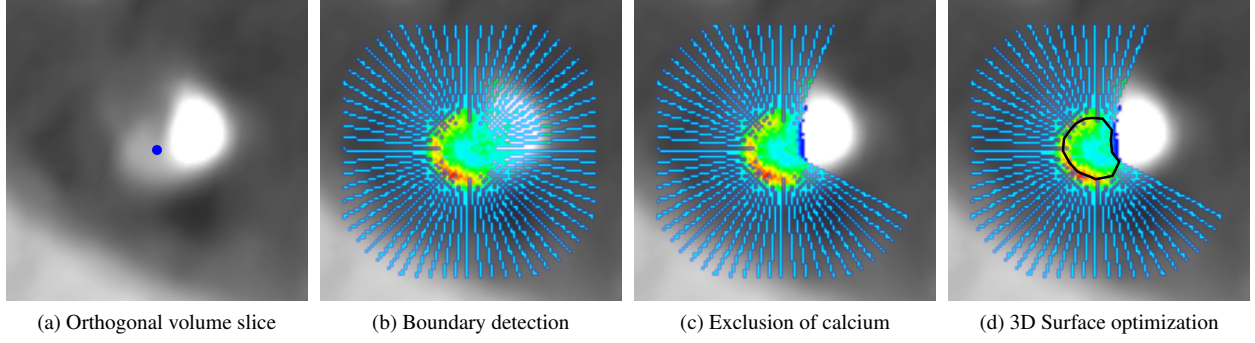


Figure 2: Orthogonal slices are extracted along consecutive centerline points (blue dot) (a). A cascade of classifiers evaluates potential boundary positions radially (b). The boundary likelihood is visualized as a color-coded overlay within blue (0) and red (1). Positions of calcified plaque are identified and used to eliminate incorrect boundary responses (c). The lumen surface is computed by solving a 3D graph-based closed set optimization on the boundary probabilities (d).

further consideration if plaque is found (Sec. 3.3). Third, with respect to these potential boundary candidates and their probabilities, an optimal lumen surface —constrained by predefined smoothness parameters— is computed by applying the method of Li *et al.* [11] followed by a refinement step (Sec. 3.4).

3.1 Volumetric Warping

First, a warped and resampled version of the original CTA data, also known as orthogonal image stack, is generated to simplify the subsequent segmentation. In order to get a homogeneous slice distance and thus avoid image distortion, the provided centerline tree is resampled (inter-slice resolution *e.g.* 0.1 mm) using bi-cubic spline interpolation (see Fig. 3a). The initial centerline that describes the path of the coronary arteries should not pass through calcified plaque in order to guarantee a proper function of our algorithm. If this is not the case for the provided centerlines, a correction via orthogonal shifts of the centerline points towards the lumen center can be applied to cure this problem. Then, for each point of the centerline, a slice is extracted orthogonal to the vessel path using bi-linear interpolation. An interpolation consistent with inter-slice resolution guarantees the generation of an isotropic warped volume which is basically a parallel stack of the extracted slices (see Fig. 3b). Note, that the size of the extracted slice should be adjusted such that it covers at least an area that is big enough to show a cross-section through a vessel with maximum diameter. Thus, the slice can safely be restricted to 15 mm x 15 mm.




Figure 3: A warped volume (b) is generated by extracting orthogonal slices along a resampled centerline tree (a). Thus, the centerline becomes a straight line in the warped space simplifying the subsequent segmentation.




Figure 4: (a) Cylindrical coordinate system for representing a tubular surface in the warped space. (b) Candidate boundary points for slice z of the warped volume along the normal of the initial circular contour. The example shows $K = 8$ equi-angular rays with each $R = 5$ equi-distant candidate points.

As a result, the centerline becomes a straight line in the warped image space and runs through the center of each slice of the warped volume. Another advantage of this transformation step is that the lumen segmentation can now be performed in a cylindrical coordinate system which is suitable for tubular structures such as blood vessels.

3.2 Boundary Detection

Next, potential candidates for the lumen boundary are sampled cylindrically and evaluated using a machine learning approach. For effectively representing the sought tubular surface in the warped image volume we use a cylindrical coordinate system (Fig. 4a). Hereby, a slice in the warped volume is expressed by the coordinate z , whereas the angle k and the radial distance r determine a point in the cross-sectional plane (*i.e.* slice) in polar coordinates. Boundary candidate points are then generated for an equidistant selection of slices $z \in [1, Z]$ of the warped volume, *e.g.*, for every 5th slice. In each such slice, R equidistant points along K rays are generated as depicted in Fig. 4b. This way, $Z \times K \times R$ directed candidate boundary points are generated.

Each such directed candidate boundary point is then evaluated for its likelihood to lie on the lumen boundary. There are various applicable models such as weighted intensity differences, gradient magnitudes or probabilities determined by previously trained classifiers. Primarily, they should yield a higher score at positions that are close to the true lumen boundary. For reasons of robustness, we recommend using boundary classifiers as described in the following subsection.

For convenience, the obtained likelihood values can be stored in a volume $l(z, k, r)$ of dimension $Z \times K \times R$.

Learning-based Ray Detection

Within the context of machine learning, boundary detection is commonly formulated as a binary classification problem. Given adequate ground-truth annotations, a binary classifier can be trained to predict whether a candidate position is part of the lumen border or not. While any binary classifier can be used, we tested the *probabilistic boosting tree* (PBT) [12] and the *random forest* (RF) [13] classifiers, which both gave similar results.

In this approach, the prediction for a point at location (z, k, r) is based on low-level image features $F(z, k, r)$ which are extracted from a local neighborhood. For each point in question, an image feature sampling pattern is defined by the local orientation of the ray connecting the center point and the potential boundary point as shown in Fig. 5. These so-called steerable features [14], which are the combination of low-level image features such as gradients computed at each position of the sampling pattern, yield a discriminative high-dimensional feature vector \vec{x} . Thus, the prediction and training of the boundary classifier is based on the evaluation of \vec{x} for each candidate position.

Training is based on available ground-truth annotations (*i.e.* contour or surface describing the lumen border) from which positive and negative samples are extracted. For each orientation (*i.e.* ray), the intersecting point between the ray and the ground-truth annotation is considered as positive and the remaining points on the ray are negative training samples (see Fig. 5b). This yields the locations of the lumen on which positive and negative feature sample sets (T_0, T_1) are calculated. Randomly selected samples are used for training of two cascading classifiers (P_1, P_2) . This is done in a bootstrapping

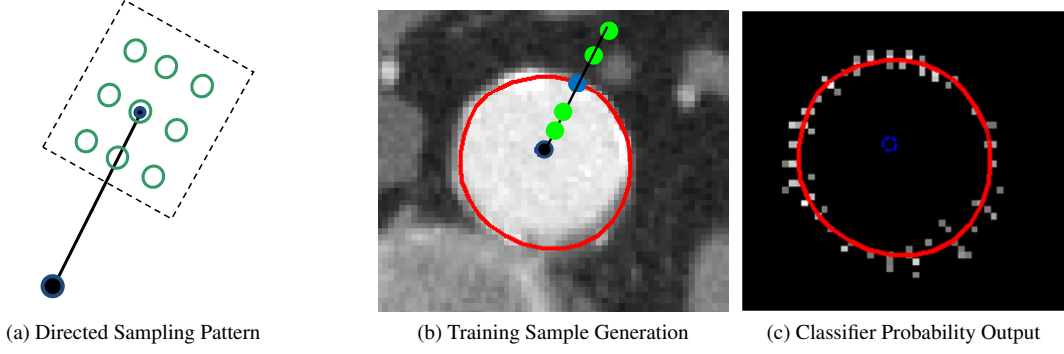


Figure 5: (a) A ray and the associated image feature sampling pattern. (b) Extraction of positive (blue) and negative (green) samples for one ray given an annotated contour (red).

approach: while the first classifier gets to see the full training data set (T_0, T_1) , the second classifier is trained with the same positive samples T_1 but negative samples \hat{T}_0 where only false positives of the first classifier ($P_1 \geq \theta_p = 0.52$) are used. With that technique, we make the second classifier more robust against complex boundaries and conflicting features.

Likewise for detection, a previously trained cascade of classifiers is used to generate the boundary probability P_B (see Fig. 5c). More precisely, the boundary probability P_B for a feature sample \vec{x} is produced by combining the two cascaded classifiers in the following way:

$$P_B(\vec{x}) = \begin{cases} P_2(\vec{x}), & \text{if } P_1(\vec{x}) \geq \theta_p, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

That is, P_B is zero, unless the first classifier meets a certain degree of belief θ_p , then the final probability is that of the second classifier.

3.3 Calcium Exclusion

For an accurate lumen segmentation, the correct handling of calcified plaque is mandatory. Due to the fact that calcium in CTA images is characterized by high intensity values and is, hence, similar to intensities captured inside the lumen, it is often erroneously classified as lumen tissue. Thus, boundaries are often detected between plaque and vessel background rather than lumen and plaque. However, calcified regions in coronary arteries are not part of the blood-flowed lumen and therefore, have to be excluded from the segmentation results. Otherwise the segmentation cannot be used for a reliable analysis of stenoses. Even though we generated the boundary likelihood with a robust machine-learning approach (see Sec. 3.2), the obtained boundary probabilities tend to show false responses for calcified vessels due to high intensity differences at calcium spots.

This issue becomes obvious with the following representative example. As depicted at the bottom of Fig. 6, profiles for the *lumen likelihood* and the *boundary likelihood* can be extracted for a slice z and a ray at angle k along the radial direction r . For that, the lumen likelihood was estimated by computing the normal distribution of the lumen intensity (mean μ_l and variance σ_l measured in HU) from samples extracted along the centerline. For rays passing through calcified plaque (Fig. 6d), these curves are characteristically different from the normal, healthy case (Fig. 6c). While in the healthy case a boundary likelihood profile with only one peak can be expected, the profile through a calcified plaque may exhibit multiple peaks. In the right example in Fig. 6 the most likely boundary candidate point (at about 3.5 mm) would in fact erroneously include the calcified plaque. Thus, a heuristic approach is proposed for removing boundary likelihood peaks that are beyond calcified plaques.

First, the calcified plaque is identified by determining the positions of the intensity profile that are above a certain threshold t_{CAL} . The radial position r_0 which is closest to the center line and which extends up to a position r_1 with a certain minimum length $l_{CAL} = r_1 - r_0$ (e.g. $l_{CAL} = 0.3$ mm) is regarded as the relevant calcified range. Then, the corrected boundary likelihood along the ray is obtained as

$$l_c(z, k, r) = \begin{cases} l(z, k, r) & r < r_0 \\ l(z, k, r) \cdot d_{CAL}^{(r-r_0)/\Delta r} & r \geq r_0 \end{cases} \quad (2)$$

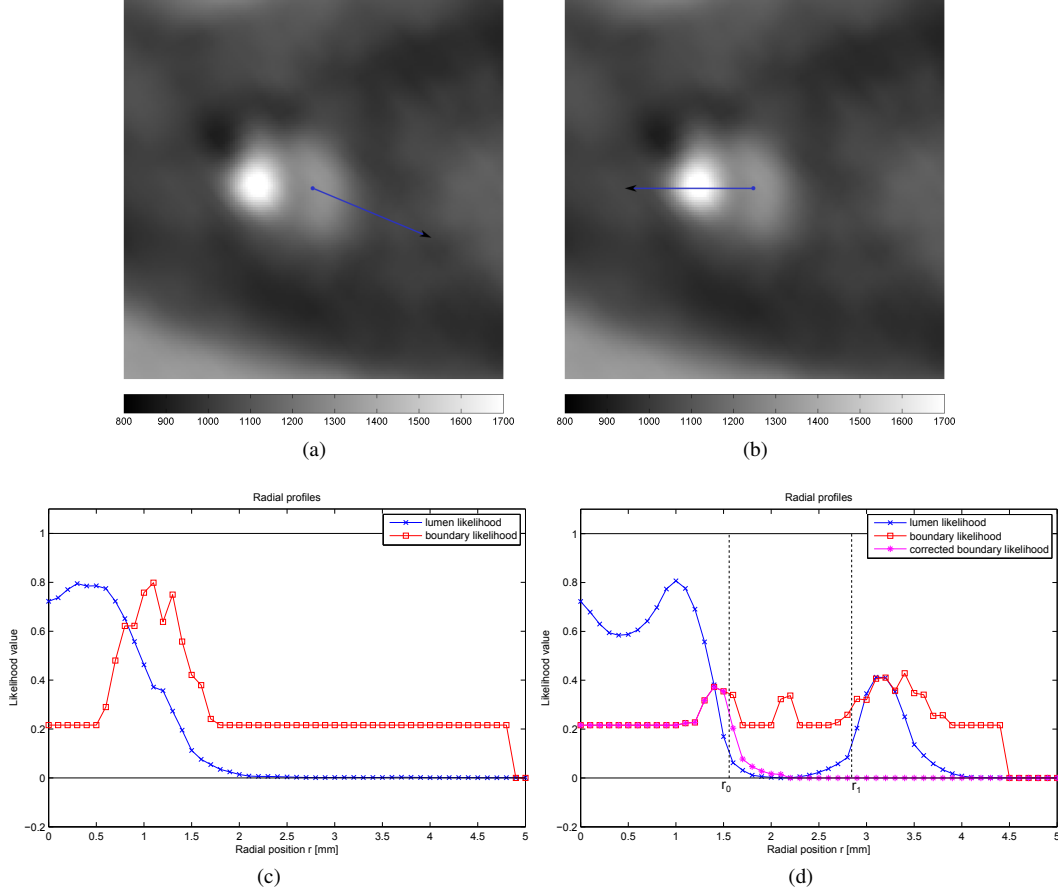


Figure 6: Characteristic ray profiles of the lumen and boundary probabilities are shown for normal (c) and calcified tissue (d) along with the corresponding ray directions in a cross-sectional view. While the boundary profile through healthy tissue has maximum response correctly at about 1.3 mm, the profile through calcium erroneously has multiple peaks with a global maximum at about 3.5 mm which would include the calcified region ($r_0 - r_1$). The calcium corrected profile (magenta) cures this problem by applying a probability damping along the ray for positions that exceed a data-dependent intensity threshold (d).

where Δr is the radial sampling resolution ($\Delta r = 0.1$ mm) and d_{CAL} is a damping factor which determines the speed of decay of the boundary likelihood profile proportional to distance from the starting radius r_0 (e.g. $d_{CAL} = 0.6$). Thus, the boundary response gets exponentially damped along the ray starting from the first relevant calcified position.

As depicted in the graph of Fig. 6d, this strategy leads to a boundary likelihood profile with only one peak that describes a boundary which correctly excludes the calcified plaque (see also Fig. 7).

While the calcium intensity threshold can be chosen constant (e.g. $t_{CAL} = 576$ HU), an adaptive threshold based on the lumen intensity distribution improves the effectiveness of the proposed heuristic. To this end we propose using the maximum of a fixed threshold (e.g. $t_f = 576$ HU) and the mean lumen intensity μ_l plus twice its variance σ_l :

$$t_{CAL} = \max(t_f, \mu_l + 2\sigma_l). \quad (3)$$

3.4 Surface Optimization

For surface construction, the lumen segmentation problem is turned into a constrained optimization problem: out of all potential boundaries the best suiting option needs to be selected. This has to be done in respect to decisions for neighboring rays and slices in order to avoid unsteady contours and accomplish a smooth 3D lumen surface.




Figure 7: Calcifications can be excluded from the segmentation by altering erroneous boundary responses along the ray. Please note that the boundary likelihood is visualized color-coded between blue (0) and red (1).




Figure 8: A post-smoothing on the noisy segmentation output yields the refined contours.

We apply the surface optimization scheme of Li *et al.* [11] in order to determine the optimal radial positions for each slice and ray of the warped volume. To this end, negative logarithmic boundary probabilities that are evaluated in cylindrical positions are forming the cost function for Li’s tubular surface optimization algorithm. Thus, for each evaluated candidate position (z, k, r) , a cost-weighted node is added to a 3D directed graph structure. The constructed graph represents the sought tubular surface by unfolding it into a volumetric representation. Then, a minimum closed set formulation on this graph yields the optimal radial position $r_{z,k}^*$ for each slice z and orientation k (determining the sought lumen surface) using s-t cut minimization. The smoothness constraints for this optimization scheme are the maximum tolerated deviations of radii along slice and ray directions, respectively. Surfaces violating those smoothness constraint are infeasible as this is limited by the graph construction with infinite arc weights. Since jumps within the tolerated deviations are not penalized, the obtained solutions are very well adapted to the underlying data. But since coronary CTA images are rather noisy, the resulting surface contours are quite noisy, too. This makes a post-refinement of the surface necessary which is a smoothing of the radius $r_{z,k}^*$ along neighboring rays and slices (see Fig. 8).

4. EVALUATION

4.1 Data

Quantitative and qualitative evaluations have been performed on a number of clinical contrast-enhanced CTA data. For that, two different databases were employed.

The first utilized database \mathcal{D}_1 is provided by the authors of the coronary artery segmentation and stenoses detection evaluation framework¹ and is publicly available for anyone who registers [6]. The database consists of 48 multi-vendor CTA datasets that were acquired from a representative selection of CAD symptomatic patients using several cardiac CT scanners with varying protocols and reconstruction algorithms. Ground-truth lumen annotations (cross-sectional contours) were created for the first 18 datasets by three different medical experts. Those were used for parameter tuning while the

¹<http://coronary.bigr.nl/stenoses/about.php>




Figure 9: Segmented lumen of calcified arteries for two patients (top/bottom) as obtained by the proposed algorithm visualized from different projections (left: CPR, right: cross-section along red arrow in CPR through a stenosis).

remaining 30 datasets were for testing only. Optionally, there were sets of traced centerlines provided that could be used to initialize the segmentation. In our experiments we employed centerlines that were extracted by the combined model-driven and data-driven tracing algorithm described in reference [10].

The second clinical database \mathcal{D}_2 is compiled from a clinical study that targets the investigation of blood flow from CTA datasets. Data was acquired from 24 symptomatic patients on the same cardiac CT scanner. Every scan typically consisted of 200 - 300 slices with an intra-slice resolution ranging between 0.3 - 0.5 mm and an inter-slice resolution of less than 0.4 mm. Each of the CTA volumes was annotated by a medical expert such that for a large set of coronary arteries ground-truth is available.

Fig. 9 shows two representative examples from the publicly available database \mathcal{D}_1 with lumen segmentations of diseased patients using our approach (dataset #11 and #14).

4.2 Accuracy Measures

For measuring the accuracy of the proposed segmentation method, three different error metrics were used in order to analyze the difference between the ground-truth and the computed segmentation. The applied metrics are a variant of the Dice coefficient, the mean surface distance (MSD) and the Hausdorff surface distance (MAXSD). Please note that DICE is measured between 0 - 1 (total overlap) whereas MSD/MAXSD is the mean/maximal distance to the ground-truth in mm. See [6] for definitions and further details.

4.3 Results

The segmentation accuracy averaged over the ground-truth annotations of three medical experts is assessed separately on diseased and healthy arteries of \mathcal{D}_1 . Table 1 shows the results as reported from the online evaluation framework [6] and lists our method along with those of the medical experts and three other (semi-)automatic algorithms (the currently best-ranking teams of the challenge). Fully-automatic methods are marked with an asterisk next to their method name. Results are averaged over 30 (18 training) datasets and sorted according to their overall ranking on testing data (listing order “testing/training”). The overall *avg. ranking* is computed as the average of the ranks obtained for the three accuracy measures DICE, MSD and MAXSD. Among automatic methods, the best achieved result for each accuracy measure is marked boldface. *Expert1-3* denote three medical experts that annotated the lumen surface manually. All accuracies were created by comparing the segmentation with each of the three ground-truth annotations and averaging the results. For more details, please refer to Kirisli *et al.* [6].

In order to rule out potential over-fitting due to training of the boundary detector or parameter tuning we tested two different detector models on the training dataset. One boundary detector model was trained on the 18 training datasets

of \mathcal{D}_1 (**Proposed**) while the other was trained only from annotations of another clinical database \mathcal{D}_2 (**Proposed- \mathcal{D}_2**). A comparison of both models on training data shows that the obtained results are very consistent which indicates that no over-fitting occurred. Hence, similar results can be expected on other sets of data. However, only the first model participated in the challenge where evaluation was performed against unseen ground-truth which is the reason of missing data for **Proposed- \mathcal{D}_2** (*n.a.*).

The following discussion solely considers results on testing data. Our method yields accuracies that are roughly within the range of accuracies achieved by the medical experts. An exception are the noticeably inferior MSD/MAXSD measures for healthy vessels. They seem to be accounted for by a mix-up of the centerlines obtained by the employed tracking algorithm and ground-truth vessels. In comparison to the other automatic methods, our approach goes head to head with the best-ranking system by Mohr *et al.* [7]. While they achieve slightly better results for healthy vessels, our method yields higher accuracies for diseased arteries (see above).

Furthermore, there is a high variability among the reported accuracies achieved by medical experts. Robust automatic methods, in contrast, enable a consistent lumen extraction in a fraction of time which is needed for manual segmentation.

Method	Rank avg. ↓	DICE D [%]	DICE H [%]	MSD D [mm]	MSD H [mm]	MAXSD D [mm]	MAXSD H [mm]
Expert3	3.1/3.1	79/76	81/80	.23/.24	.21/.19	3.00/3.07	3.45/3.25
Proposed*	3.8/3.7	74/72	73/74	.35/.28	.55/.35	2.99/2.02	3.73/2.88
Mohr* [7]	3.8/3.9	70/73	73/75	.40/.29	.39/.45	2.68/1.87	2.75/3.73
Expert1	3.9/3.9	76/74	77/79	.24/.26	.24/.26	2.87/3.29	3.47/3.61
Expert2	5.3/5.1	65/66	72/73	.34/.31	.27/.25	2.82/2.70	3.26/3.00
Shahzad [9]	5.6/5.7	65/66	68/70	.39/.37	.41/.32	2.73/2.49	3.20/3.04
Wang* [8]	6.1/6.3	69/68	69/72	.45/.43	.55/.56	3.94/4.06	6.48/5.23
Proposed-\mathcal{D}_2*	n.a.	n.a./73	n.a./75	n.a./30	n.a./36	n.a./2.21	n.a./3.11

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 (sorted by testing accuracies). Measures are averaged over 30 testing (18 training) datasets (listed testing/training). Asterisks mark fully-automatic methods only. Our method (**Proposed**) is roughly within the range of accuracies yield by medical experts and ranks head to head with the currently best-ranking algorithm. Similar results can be expected on other sets of clinical data which is indicated by the consistency of results on training data between the proposed und the unrivaled method **Proposed- \mathcal{D}_2** (detector model trained on \mathcal{D}_2) .

5. CONCLUSIONS

A cascade of learning-based ray classifiers yields a robust lumen boundary detection whereof the final lumen surface is found by a global optimization scheme. The exclusion of calcified plaque from the segmentation is assured through a novel calcium exclusion technique which allows us to accurately capture stenoses of diseased arteries. The proposed approach is able to fully-automatically extract the coronary lumen of calcified and non-calcified vessels with accuracies almost similar to medical experts. On a standardized testing framework our method ranks better than two of three medical experts and lies level with the so-far best-ranking automatic system.

REFERENCES

- [1] den Dekker, M., de Smet, K., de Bock, G., Tio, R., Oudkerk, M., and Vliegenthart, R., “Diagnostic performance of coronary CT angiography for stenosis detection according to calcium score: systematic review and meta-analysis,” *European radiology* **22**(12), 2688–2698 (2012).
- [2] Kelm, B., Mittal, S., Zheng, Y., Tsybal, A., Bernhardt, D., Vega-Higuera, F., Zhou, S., Meer, P., and Comaniciu, D., “Detection, grading and classification of coronary stenoses in computed tomography angiography,” in *Proceedings of the 14th international conference on Medical image computing and computer-assisted intervention - Volume Part III*, MICCAI’11, 25–32, Springer-Verlag, Berlin, Heidelberg (2011).

- [3] Nakazato, R., Park, H., Berman, D., Gransar, H., Koo, B., Erglis, A., Lin, F., Dunning, A., Budoff, M., Malpeso, J., Leipsic, J., and Min, J., “Non-invasive fractional flow reserve derived from CT angiography (FFRCT) for coronary lesions of intermediate stenosis severity: Results from the DeFACTO study,” *Circulation: Cardiovascular Imaging* **6** (2013).
- [4] Sharma, P., Itu, L., Zheng, X., Kamen, A., Bernhardt, D., Suciu, C., and Comaniciu, D., “A framework for personalization of coronary flow computations during rest and hyperemia,” in [*Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*], 6665–6668, IEEE (2012).
- [5] Lesage, D., Angelini, E., Bloch, I., and Funka-Lea, G., “A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes,” *Medical image analysis* **13**(6), 819–845 (2009).
- [6] Kirişli, H. 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).
- [7] Mohr, B., Masood, S., and Plakas, C., “Accurate lumen segmentation and stenosis detection and quantification in coronary CTA,” in [*Proceedings of 3D Cardiovascular Imaging: a MICCAI Segmentation Challenge Workshop*], (2012).
- [8] Wang, C., Moreno, R., and Smedby, Ö., “Vessel segmentation using implicit model-guided level sets,” in [*Proceedings of 3D Cardiovascular Imaging: a MICCAI Segmentation Challenge Workshop*], (2012).
- [9] Shahzad, R., van Walsum, T., Kirişli, H., Tang, H., Metz, C., Schaap, M., van Vliet, L., and Niessen, W., “Automatic stenoses detection, quantification and lumen segmentation of the coronary arteries using a two point centerline extraction scheme,” in [*Proceedings of 3D Cardiovascular Imaging: a MICCAI Segmentation Challenge Workshop*], (2012).
- [10] Zheng, Y., Tek, H., and Funka-Lea, G., “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*], *Lecture Notes in Computer Science* **8151**, 74–81, Springer Berlin Heidelberg (2013).
- [11] Li, K., Wu, X., Chen, D., and Sonka, L., “Optimal surface segmentation in volumetric images—a graph-theoretic approach,” *IEEE transactions on pattern analysis and machine intelligence* **28**, 119–34 (Jan. 2006).
- [12] Tu, Z., “Probabilistic boosting-tree: learning discriminative models for classification, recognition, and clustering,” in [*Tenth IEEE International Conference on Computer Vision (ICCV’05)*], **2**, 1589–1596, IEEE (2005).
- [13] Breiman, L., “Random forests,” *Machine learning* **45**, 5–32 (2001).
- [14] Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., and Comaniciu, D., “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).