

Detection, Grading and Classification of Coronary Stenoses in Computed Tomography Angiography

B. Michael Kelm¹, Sushil Mittal^{2,3,*}, Yefeng Zheng², Alexey Tsymbal¹,
Dominik Bernhardt⁴, Fernando Vega-Higuera⁴, S. Kevin Zhou²,
Peter Meer³, and Dorin Comaniciu²

¹ Image Analytics and Informatics, Corporate Technology, Siemens AG, Erlangen, Germany

² Image Analytics and Informatics, Siemens Corporate Research, Princeton, NJ, USA

³ Electrical and Computer Engineering, Rutgers University, NJ, USA

⁴ Computed Tomography, Healthcare Sector, Siemens AG, Forchheim, Germany

Abstract. Recently conducted clinical studies prove the utility of Coronary Computed Tomography Angiography (CCTA) as a viable alternative to invasive angiography for the detection of Coronary Artery Disease (CAD). This has led to the development of several algorithms for automatic detection and grading of coronary stenoses. However, most of these methods focus on detecting calcified plaques only. A few methods that can also detect and grade non-calcified plaques require substantial user involvement. In this paper, we propose a fast and fully automatic system that is capable of detecting, grading and classifying coronary stenoses in CCTA caused by *all* types of plaques. We propose a four-step approach including a learning-based centerline verification step and a lumen cross-section estimation step using random regression forests. We show state-of-the-art performance of our method in experiments conducted on a set of 229 CCTA volumes. With an average processing time of 1.8 seconds per case after centerline extraction, our method is significantly faster than competing approaches.

1 Introduction

According to the American Heart Association, Coronary Artery Disease (CAD) is a leading cause of death in the western world. Every year, about six million patients in the United States emergency departments are examined for acute chest pain [5]. The current diagnostic standard is conventional invasive angiography which involves a very high amount of risk and cost. New generations of cardiac Computed Tomography (CT) scanners enable the acquisition of Coronary CT Angiography (CCTA) images with unprecedented quality [2]. In the review article [1], Achenbach has summarized the results of many clinical studies, comparing contrast-enhanced CCTA with conventional invasive angiography. These results prove CCTA a viable alternative with very high negative predictive value. However, reading CCTA images requires substantial experience and only well-trained physicians are able to interpret CCTA reliably [12]. An automated system that can rule out clinically-relevant stenoses (grade > 50%) in the coronary arteries could be used as a second reader in the absence of an expert physician in the emergency department.

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

Detection and grading of coronary stenoses in CCTA is very challenging due to varying image quality. In an endeavor to reduce the radiation dose during CT scans, often images with relatively low signal-to-noise ratio are acquired [2]. Motion artifacts are frequently encountered despite the routine use of beta blockade to reduce the heart rate. Gated or modulated acquisition protocols may result in stair-case reconstruction artifacts which further complicate the analysis. Finally, coronary arteries follow long and treacherous paths, extending over only a few voxels in diameter in the distal parts. Even experts sometimes struggle to give a correct diagnosis due to these challenges [12].

In this work, we present an algorithm for detecting, grading and classifying severe (i.e. clinically relevant) stenoses along automatically extracted centerlines of the coronaries. The contribution of our work is three-fold. First, we introduce a novel centerline verification step (Section 3.1). Second, a novel regression approach replaces the lumen segmentation that is needed for grade estimation (Section 3.2). Finally, in contrast to previous work, we propose a complete end-to-end and fully automated system that works on *all* types of plaques.

2 Related Work

Vessel Tracing. Manual tracing of coronary centerlines in 3D cardiac CT volumes is a highly tedious task. Many algorithms for automatic tracing of centerlines have been proposed, the most important of which are reviewed in [8]. Owing to the importance of the problem, the MICCAI association also organized a competition for automatic and semi-automatic coronary artery tracking [10]. In our work, the centerlines were traced using the state-of-the-art method of [6]. The method uses multi-scale *medialness* filters in a graph-based algorithm to extract centerlines by computing minimum-cost paths.

Lumen Segmentation. Estimating the cross-sectional area (or the radius) of the vessel lumen along the centerline is a key feature to detecting and grading coronary stenoses. But methods relying on exact segmentation of the lumen are slow owing to high computational complexity of the segmentation algorithms. A comprehensive review of all major lumen segmentation algorithms is provided in [8]. Given a centerline, we propose a novel, automatic regression-based method to directly estimate the vessel radius, which is significantly faster than a segmentation-based approach.

Stenosis Detection. In the past, a variety of algorithms have been proposed for detection of coronary plaques in CCTA. However, most of this work focuses on the detection of calcified plaques only, e.g., see [14,11]. Fewer methods have been proposed for fully automatic detection of non-calcified plaques, which are usually harder to detect and grade with high confidence. For example, the methods proposed in [13] and [15] need substantial user input in order to localize and grade coronary stenoses. A learning based method was proposed in [14] which could detect both calcified and non-calcified plaques. However, this approach does not consider stenosis grade and thus also reports non-severe lesions which are clinically irrelevant. To the best of our knowledge, the only system that currently analyzes CCTA for coronary artery disease in a fully automated way by reporting location, type and severity of coronary lesions is Rcadia's COR

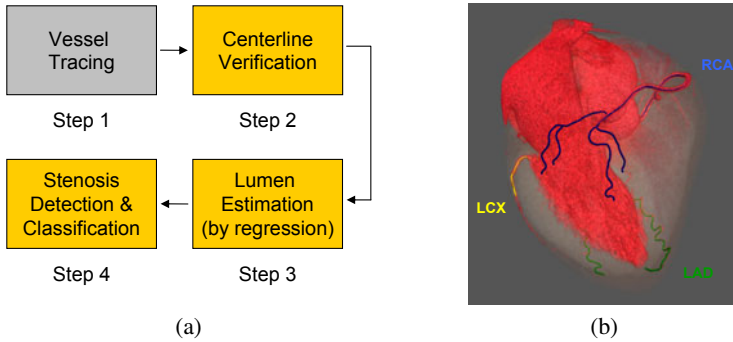


Fig. 1. (a) Overview of the four-step approach for stenosis detection, grading and classification. (b) Example vessel tree showing the three major coronary arteries: left anterior descending (LAD), left circumflex (LCX) and right coronary artery (RCA).

Analyzer. Being a commercial product, however, not much detail is known about the approach. Recent studies evaluating the system were published in [3] and [7].

3 Methods

Given a CCTA image volume, the proposed system can automatically detect and classify coronary stenoses using the general four-step approach sketched in Fig. 1 a. In the *first* step, centerlines of the three major coronary arteries, i.e. left anterior descending (LAD), left circumflex (LCX) and right coronary artery (RCA) (Fig. 1 b), along with their branches, are automatically extracted using the method of [6]. The left main (LM) coronary artery is processed as the common part of LAD and LCX. Using a learning-based classification approach, the *second* step verifies the accuracy of the extracted centerlines and removes parts of the vessel tree belonging to non-coronary regions. This ensures that the subsequent stenosis detection is only performed along the actual vessel, thus producing more stable results. The *third* step employs a learning-based regression approach to locally estimate the cross-sectional area of the vessel lumen (which may be narrowed by coronary plaques), along the extracted centerlines. In the *fourth* step, candidate stenoses are extracted for each individual segment of the vessel tree as max-min-max triples of a baseline-corrected and smoothed radius curve along with a grade estimate. Each candidate stenosis is either discarded (e.g. grade < 50%) or accepted and classified as calcified, non-calcified or mixed stenosis. This decision is based on image features as well as features of the candidate stenosis such as length, cross-sectional area/radius and distance to the distal end of the vessel.

3.1 Centerline Verification

To avoid the risk of missing a potentially diseased coronary vessel, any automatic centerline extraction algorithm must be highly sensitive, especially in low-contrast vessel regions, which is true for the employed algorithm [6], for example. Thus, the method should be able to trace centerlines of the vessels occluded by non-calcified plaques, that

often exhibit the same contrast as the vessel wall. Inevitably, such algorithms are prone to various types of errors in tracing and centerlines may get wrongly traced into other arteries or into non-coronary regions like veins, heart chambers, etc. Figure 2 (left) shows Curved Planar Reformation (CPR) views of some more noticeable examples for erroneous tracing results of the employed algorithm. Table 1 summarizes the errors in tracings for 229 volumes with 1472 traced centerlines in the distal vessel regions.

Table 1. Summary of errors in tracing coronary centerlines in the distal vessel regions. A total of 229 volumes were processed and 1472 centerlines were traced.

Error in tracing	$\geq 5\text{mm}$	$\geq 10\text{mm}$	$\geq 15\text{mm}$
# vessels affected	259 (17.6%)	226 (15.4%)	210 (14.3%)
# volumes affected	131 (57.2%)	116 (50.66%)	107 (46.7%)

Clearly, erroneously traced centerlines are susceptible to many false alarms while detecting lesions. Therefore, we propose to use a learning based algorithm for automatic detection of non-coronary regions along the extracted centerlines. Similar to [11], a cylindrical sampling pattern for feature extraction, with its axis aligned to the coronary centerline, is employed. We then extract altogether 171 rotation invariant features along the entire length of the cylinder at varying radii. These features are used to train a random forests (RF) classifier [4] (100 trees, 14 randomly selected features, stratified sampling of 8000 examples with replacement). Given an unseen volume, the RF classifier outputs a probability that a given point on the centerline belongs to a non-vessel region. After removing outliers using a median filter of width 11, points with scores higher than a fixed threshold are excluded from further analysis. In our experiments this threshold was determined as 0.7 to yield a specificity of at least 98%, so that at most 2% of the points are erroneously discarded (see Fig. 2, right).

3.2 Lumen Estimation

Instead of segmenting the lumen and computing the lumen cross-sectional area along the vessel centerlines we propose using a non-linear regression approach to directly estimate the cross-sectional area from local image features.

Alternatively for estimating the cross-sectional area we trained for and estimated the radius r of an equivalent circle with the same area. Thus, a function for the radius $r(\mathbf{x}|p)$ is estimated that depends on the feature vector \mathbf{x} and a set of parameters p that are learned from a training set $T = \{(\mathbf{x}_i, r_i)\}_{i=1}^N$. The training set is constructed from semi-manual lumen segmentations of coronary arteries by computing cross-sections and the corresponding radii r_i at altogether N centerline points. At the same points, rotation-invariant features \mathbf{x}_i are extracted according to [11]. A regression function is learned by minimizing the squared loss $L(p) = \sum_{i=1}^N (r(\mathbf{x}_i|p) - r_i)^2$ on the training set T with respect to the parameters p . For this purpose an ensemble of randomized regression trees, i.e. a random regression forest was employed [4]. Formally, each regression tree models the dependent variable (the lumen radius) as a piecewise constant function, which results in the following model for the ensemble:

$$r(\mathbf{x}|p) = \frac{1}{T} \sum_{t=1}^T \sum_{m=1}^{M^t} r_m^t I(\mathbf{x} \in \mathbf{R}_m^t) \quad (1)$$

where $I(\cdot)$ denotes the set indicator function. Thus, for each tree t , the feature domain is partitioned into M^t regions $\mathbf{R}_1^t, \mathbf{R}_2^t, \dots, \mathbf{R}_{M^t}^t$, and in each region the radius is modeled by a constant r_m^t [4]. During training of each tree, the sample is successively partitioned by thresholding one of the feature values. The optimal split (feature and threshold) is greedily determined as the one which reduces the squared loss function most. In the leaves (a region \mathbf{R}_M^t), the radius r_m^t is estimated as the average radius of all remaining examples. While it is possible to stop splitting when the variance within a node and/or the number of examples drop below certain values, we use fully grown trees in our experiments, i.e. they were grown until only one example remained in every leaf. For obtaining good generalization performance, it is essential that every tree slightly differs. This is achieved by training each tree on a resampled training sample (at most 4000 examples drawn with replacement) and by randomly selecting a fraction of the features (one third, i.e. 57 of 171) to be considered for each split.

Given the feature vector of an unseen example, each regression tree is evaluated by following its splits from the root to a leaf with an associated radius estimate r_m^t . The random forest averages the estimates of its trees (Eqn. (1)). We used $T = 50$ regression trees in our experiments.

3.3 Stenosis Detection and Classification

Candidate stenosis regions are identified and graded using the estimated lumen radii. The vessel tree is decomposed into disjoint segments so that every segment either starts at an ostium or a vessel bifurcation. Each segment is then analyzed separately.

First, a baseline curve is computed using binomial filtering (Fig. 3, green). It is subtracted from the original radius curve to obtain a de-trended residual curve which is again slightly smoothed (Fig. 3, red). The positions of the local optima are extracted; clearly, local minima and maxima alternate. Every triple (max-min-max) is then regarded as a stenosis candidate for which a grade is estimated by

$$g = 1 - \left(\frac{2r_{min}}{r_{left} + r_{right}} \right)^2, \quad (2)$$

where r_{min} is the minimum radius within, r_{left} the radius at the left (towards the ostium) and r_{right} the radius at the right end of the stenosis candidate (Fig. 3, magenta). At the ostia and bifurcations, the grade is estimated with the alternative formula $g = 1 - (r_{min}/r_{right})^2$ to account for the non-pathologic radius broadening there.

Then it is decided for each stenosis candidate whether it should be discarded or not. In addition to low grade candidates ($< 50\%$), also short (< 0.9 mm) and narrow ($r_{right} < 1.0$ mm) ones are discarded. Candidates close to the distal end of the vessel (< 7.5 mm) are also discarded since the data quality usually gets too low there.

Finally, using probability scores obtained from two classifiers for the detection of calcified and non-calcified plaques similar to [11] (Fig. 3, cyan/orange), each accepted

stenosis candidate is classified into one of three types, “calcified”, “non-calcified” and “mixed” (calcified as well as non-calcified parts).

4 Experimental Results

Training of the system was performed using a total of 229 CCTA volumes that were acquired on several cardiac CT scanners with varying protocols and reconstruction algorithms. The slice distance for these scans varied between 0.3-0.5mm with x - y pixel spacing being between 0.3-0.4mm. Each scan typically consisted of 200-300 slices. For training and evaluation, the data was manually annotated. For automatic vessel tracing, errors were annotated. Coronary plaques were labeled with their type (calcified, non-calcified, mixed) as well as a rough grade (mild, moderate, severe, occluded).

Figure 2 shows cross-validation results for the centerline verification step. Invalid centerline points are reliably recognized for several types of tracing errors (Fig. 2, *Left*). The receiver-operating-characteristic (ROC) curve (Fig. 2, *Right*) shows that high specificities as well as sensitivities are attainable. The inferior performance on the RCA results from the difficulty to distinguish the (invalid) coronary sinus vein from a (valid) artery. For the centerline verification step a high specificity (> 0.98) is desirable in order to ensure that no valid parts of the vessel tree are discarded.

Numeric results for the overall system are provided in Table 2. The “by-lesion” sensitivity quantifies how many of the severe non-calcified stenoses are detected while the “by-vessel” sensitivity quantifies how many of the vessels with severe non-calcified stenoses are identified. Only the “by-vessel” measure allows to compute the specificity and negative predictive value (NPV). For an application as a second reader, sensitivity and NPV are of utmost importance and preferred over specificity, i.e. false positives are acceptable while false negatives are not. With an overall sensitivity of 97.62%, a NPV of 99.77% and a specificity of 67.14% the proposed system performs competitive with




Fig. 2. *Left:* CPR views of wrong centerline tracings obtained using the algorithm of [6]. The first two are partially incorrect and traced into a heart chamber and a vein respectively. The third is entirely wrong and traced into a heart chamber. Points detected outside the coronary are marked with orange ‘+’ sign. *Right:* ROC curves (sensitivity vs. specificity) obtained with 10-fold cross validation using 229 volumes on per vessel point basis.

Table 2. Detection performance on non-calcified plaques (10-fold cross-validation). The system is tuned towards high sensitivity and NPV to cope with its application as a second reader.

		LAD	LCX	RCA	Overall
by-lesion	sensitivity	100.0%	90.0%	95.24%	94.55%
	false positives per volume	0.81	1.03	1.13	2.97
by-vessel	sensitivity	100.00%	93.75%	100.00%	97.62%
	specificity	75.23%	63.16%	62.86%	67.14%
	negative predictive value	100.00%	99.17%	100.00%	99.77%

Rcadia’s system [3,7]. But, while the Rcadia system requires about 10min of processing time per case, on average, our system only needs 1.8 s (up to 3.9 s) after centerline tracking which completes within a minute.

This advantage can mainly be attributed to excellent performance of the lumen regression step, which provides an accurate estimate of the lumen cross-sectional area much faster than a segmentation method. To this end, we also compared our results to a segmentation approach, a learning-based version of the graph cuts approach presented in [9]. While comparable performance was achieved, the overall processing time of 21 s (up to 42 s) after centerline extraction was clearly higher with segmentation approach.

A test of the lumen regression (without retraining) on the eight (training) data sets provided by [10] yielded a bias of 0.18 mm (median 0.19 mm) and a standard deviation of 0.27 mm (median absolute deviation 0.14 mm). Apart from the bias, which can be attributed to systematic annotation differences, these results agree with the 10-fold cross-validation results on our data which yielded a bias of 0.01 mm (median 0.002 mm) and a standard deviation of 0.28 mm (median absolute deviation 0.08 mm).

Fig. 3 shows two examples of patients that have neither been used in training nor the development of the proposed system. In both cases, the system can locate and classify all severe stenotic lesions correctly. Note that the right example shows a second, more distal stenosis for which a grade of about 40% is estimated (cf. graph below the image). It is thus deemed non-severe and therefore not reported.

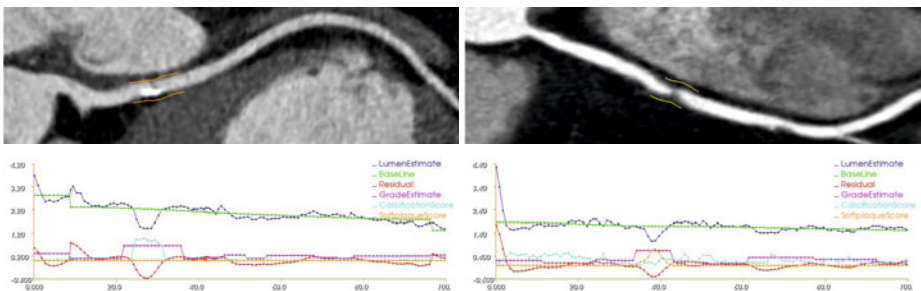


Fig. 3. Two examples of patients from unseen data. The presented system detects severe stenoses caused by both calcified plaques (left) and non-calcified plaques (right). The graph at the bottom shows the lumen estimate (blue), baseline (green), residual (red), grade estimate (magenta), calcification score (cyan) and the score for non-calcified (soft) plaques (orange) along the vessel center line. While the system is highly sensitive it only exhibits a moderate number of false alarms. For the two examples shown above, other severe stenoses were neither present nor detected.

5 Conclusion

An automatic method for the detection and classification of stenotic lesions in coronary computed tomography angiography is proposed. The centerline verification step helps the system cope with tracing errors and vessels with low data quality. Lumen cross-sectional area is accurately estimated using a regression approach which is considerably faster than a full segmentation method. A competitive performance is achieved with significantly reduced computational burden as compared to state-of-the-art methods.

References

1. Achenbach, S.: Cardiac CT: State of the art for the detection of coronary arterial stenosis. *J. Cardiovasc. Comput. Tomogr.* 1(1), 3–20 (2007)
2. Achenbach, S., Anders, K., Kalender, W.: Dual-source cardiac computed tomography: Image quality and dose considerations. *Eur. Radiol.* 18, 1188–1198 (2008)
3. Anders, K., Petit, I., Achenbach, S., Pflederer, T.: Diagnostic utility of automated stenosis detection in dual source CT coronary angiography as a stand alone or add-on tool. In: *Proc. SCCT (2010)*
4. Breiman, L.: Random forests. *Mach. Learn.* 45(1), 5–32 (2001)
5. Gallagher, M.J., Raff, G.L.: Use of multislice CT for the evaluation of emergency room patients with chest pain: the so-called “triple rule-out”. *Cath. and Cardiovasc. Interv.* 71(1), 92–99 (2008)
6. Gülsün, M.A., Tek, H.: Robust vessel tree modeling. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 602–611. Springer, Heidelberg (2008)
7. Halpern, E.J., Halpern, D.J.: Diagnosis of coronary stenosis with CT angiography comparison of automated computer diagnosis with expert readings. *Acad. Radiol.* 18, 324–333 (2011)
8. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Med. Image Anal.* 13(6), 819–845 (2009)
9. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal surface segmentation in volumetric images—a graph-theoretic approach. *IEEE Trans. Pattern Anal. Mach. Intell.* 28(1), 119–134 (2006)
10. Metz, C., Schaap, M., van Walsum, T., van der Giessen, A., Weustink, A., Mollet, N., Krestin, G., Niessen, W.: 3D segmentation in the clinic: A grand challenge II – coronary artery tracking. In: *MICCAI Workshop on 3D Segmentation in the Clinic: A Grand Challenge (2008)*
11. Mittal, S., Zheng, Y., Georgescu, B., Vega-Higuera, F., Zhou, S., Meer, P., Comaniciu, D.: Fast automatic detection of calcified coronary lesions in 3D cardiac CT images. In: *Proc. MICCAI Workshop on Machine Learning in Medical Imaging (2010)*
12. Pugliese, F., Hunink, M.G.M., Gruszczynska, K., Alberghina, F., Malag, R., van Pelt, N., Mollet, N.R., Cademartiri, F., Weustink, A.C., Meijboom, W.B., Witteman, C.L.M., de Feyter, P.J., Krestin, G.P.: Learning Curve for Coronary CT Angiography: What Constitutes Sufficient Training? *Radiol.* 251(2), 359–368 (2009)
13. Rinck, D., Krüger, S., Reimann, A., Scheuering, M.: Shape-based segmentation and visualization techniques for evaluation of atherosclerotic plaques in coronary artery disease. In: *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 6141, pp. 61410G–9 (2006)
14. Teßmann, M., Vega-Higuera, F., Fritz, D.: Learning-based detection of stenotic lesions in coronary CT data. In: *Proc. of Vision, Modeling, and Visualization*, pp. 189–198 (2008)
15. Wesarg, S., Khan, M.F., Firle, E.: Localizing calcifications in cardiac CT data sets using a new vessel segmentation approach. *J. of Dig. Imag.* 19(3), 249–257 (2006)