

Fast Automatic Detection of Calcified Coronary Lesions in 3D Cardiac CT Images

Sushil Mittal^{1,2}, Yefeng Zheng², Bogdan Georgescu², Fernando Vega-Higuera³,
Shaohua Kevin Zhou², Peter Meer¹, and Dorin Comaniciu²

¹ Electrical and Computer Engineering Department, Rutgers University, USA

² Siemens Corporate Research, USA

³ Computed Tomography, Siemens Healthcare, Germany

Abstract. Even with the recent advances in multidetector computed tomography (MDCT) imaging techniques, detection of calcified coronary lesions remains a highly tedious task. Noise, blooming and motion artifacts etc. add to its complication. We propose a novel learning-based, fully automatic algorithm for detection of calcified lesions in contrast-enhanced CT data. We compare and evaluate the performance of two supervised learning methods. Both these methods use rotation invariant features that are extracted along the centerline of the coronary. Our approach is quite robust to the estimates of the centerline and works well in practice. We are able to achieve average detection times of 0.67 and 0.82 seconds per volume using the two methods.

1 Introduction

Coronary artery disease (CAD) is one of the leading causes of death in the western world [1]. It is the abnormal constriction of the coronaries usually caused by gradual building up of plaque in the walls of the arteries. Even with the recent advances in multidetector computed tomography (MDCT) imaging techniques, conventional invasive coronary angiography is still the standard procedure to diagnose CAD. The reason can be attributed to the fact that analysis of MDCT for CAD is a tedious task requiring a very high degree of accuracy and clinical expertise. However, given the high amount of risk and cost involved in the invasive coronary angiography procedure, an automatic method for lesion detection in MDCT data would be highly welcomed.

Coronary lesions can be broadly classified into two categories – calcified lesions and soft plaques. In a contrast-enhanced CT, calcified lesions can be spotted as small, bright regions while soft plaques are usually, low contrast lesions. Both types of lesions extend from the endothelium lining in the walls of the coronary towards the center of the coronary, blocking the lumen. Due to significant difference in the visual appearance of the two types of lesions, there does not seem to be an obvious single way to detect both the types.

Over the last few years, many studies have shown the utility of contrast-enhanced CT in detection of coronary lesions [2], [3], [4]. Achenbach [5], in his review article has summarized the results of many such clinical studies conducted,

comparing the accuracy of contrast-enhanced CT with that of the conventional invasive angiography. Recently, various image processing based methods have been proposed to help the detection of lesions. Among a few semi-automatic methods proposed, Rinck et al. [6] and Wesarg et al. [7] proposed methods that make use of user-specified seed points to make various anatomical measurements of the coronaries which are then used to detect calcified lesions. Toumoulin et al. [8] presented a level set based approach to improve the detection accuracy of the vessel contours and calcifications. A few fully-automatic methods have also been proposed. Işgum et al. [9] proposed a method that used a two-stage classifier to distinguish coronary lesions among candidate lesions in non-contrast-enhanced CT images. Saur et al. [10] proposed a rule-based technique that involved extraction and registration of lesion candidates in both contrast-enhanced and non-contrast-enhanced images. Kurkure et al. [11] proposed a hierarchical supervised learning method for classification of candidate regions while Brunner et al. [12] proposed an unsupervised classification algorithm to distinguish between arterial and non-arterial lesions. Both [11] and [12] used non-contrast CT scans. More recently, Tessmann et al. [13] proposed a learning-based detection technique which used AdaBoost to train a strong classifier with various local and global features. Their method claimed to be able to detect both calcified and non-calcified types of lesions.

We propose a learning based method for automatic detection of calcified lesions. A cylindrical sampling pattern for feature extraction with the axis of the cylinder aligned to the coronary centerline, is used. We then extract rotation invariant features along the entire length of the cylinder at varying radii. These features are used to train two classifiers - a probabilistic boosting tree (PBT) based classifier [14] and a random forests (RF) based classifier [15]. The performance of the two classifiers is compared in the results section. We perform lesion detection over three main coronary arteries with an average detection time of 0.67 seconds per volume using PBT and 0.82 seconds per volume using RF. Please note that except [13], all the other fully automatic detection techniques mentioned above, use a simple thresholding on the intensity values to first extract candidate lesions. Since we only perform detection along the coronary centerlines, we do not need to generate candidate regions. Due to the same reason, it is not easy to directly compare our results (detection rates, sensitivity, specificity etc.) to those of the above mentioned methods. Given the general and simple nature of our approach as compared to other methods, our results still reflect state-of-the-art performance.

2 Data Preparation

We worked with scans obtained from 165 patients. The slice thickness was 0.5 *mm* with *x-y* resolution typically between 0.3–0.4 *mm*. Each scan typically consisted of around 200 – 300 slices. There were a total of 355 calcified lesions. In all of these volumes, both the coronary centerlines and the calcified lesions were annotated manually for training and evaluation purposes. We analyze three main

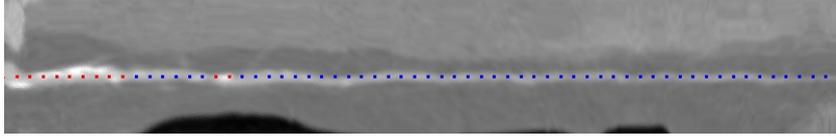


Fig. 1. Stretched CPR view of an LAD coronary artery with centerline and lesion annotations. Note that the control points are not always exactly in the center of the vessel lumen. Control points lying in the calcified and non-calcified regions are shown in red and blue colors respectively.

coronary arteries for presence of calcified lesions – left anterior descending artery (LAD), left circumflex artery (LCX) and right coronary artery (RCA). The left main coronary artery (LM) was always annotated as a part of the LAD artery. To generate positive and negative samples for the training purpose we represented the centerline of each coronary artery by a set of control points. The feature extraction was done around each control point and will be described in the next section. Since the lengths of these arteries vary considerably from one scan to another, for the sake of consistency, we fixed the distance between every two consecutive control points as 1 mm . There are many approaches proposed to automatically extract centerlines for coronary arteries. A comprehensive overview of segmenting out branched vessels in medical images is given in [16]. In this paper, we focus on lesion detection assuming that rough estimates of the centerlines are given. As stated above the centerline annotation was done manually. Most of the control points were not annotated exactly along the center of the lumen, however sufficient care was taken to make sure that almost all of them lie inside the outer walls of the coronary artery. This annotation scheme makes our lesion detection algorithm robust to the inaccuracy of the given centerline. Figure 1 shows a stretched CPR view of an LAD coronary artery with our annotation scheme. For training purposes, we label each control point along the centerline based on whether it lies in a calcified or a non-calcified region. Control points in calcified regions are represented using red dots while those in non-calcified regions are represented using blue dots.

3 Feature Extraction

For any supervised learning algorithm to work effectively, the selected features should sufficiently capture the characteristic properties of the underlying classes of the data. Coronary lesions have no specific shape, size or location along the centerline. The selected sampling pattern should therefore be invariant to such changes. It seems quite reasonable to choose a cylindrical sampling pattern with its axis aligned to the centerline of the coronary. The sampling points are selected around the axis of the cylinder at fixed radii. The length of the cylinder should be carefully chosen. It should be small enough to exploit the locally cylindrical shape of the coronary artery. At the same time, it should be large enough so that there is sufficient overlap between the sampling patterns of any

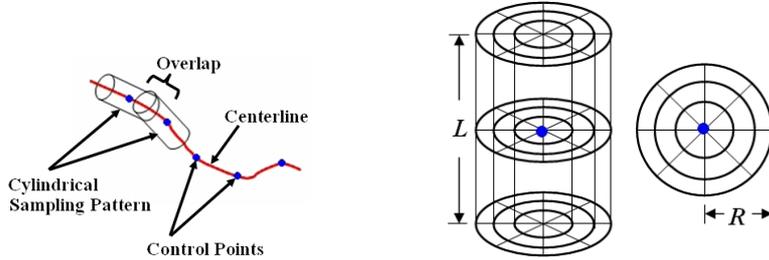


Fig. 2. Cylindrical sampling pattern. *Left:* Overlapping patterns with axes of the cylinders aligned with the centerline of the coronary. *Right:* A cylindrical sampling pattern of length L and radius R . The blue point in the center of the cylinder represents the control point C .

two adjacent control points along the centerline so that no lesion is missed by the feature extraction pattern. Figure 2(a) shows the the cylindrical sampling pattern used with its axis aligned to the centerline of the coronary. Further, since the lesions can potentially occur anywhere around the axis of the cylinder, we choose features that are rotation invariant about the axis. As shown in the right image of Figure 2, suppose a cylindrical pattern of length L and radius R is defined around a control point C . For a point at distance l , $(-L/2 \leq l \leq L/2)$, from C along the axis of the cylinder, we extract the following nine features at a radius r , where $(0 < r \leq R)$ – average, minimum and maximum intensities – $(I_{av}, I_{min}, I_{max})$, gradients along the radial direction – $(G_{av}^r, G_{min}^r, G_{max}^r)$ and gradients along the tangent direction – $(G_{av}^t, G_{min}^t, G_{max}^t)$. We found that the choice of $L = 5$ voxels gives reasonably good overlap between adjacent cylinders and $R = 3$ voxels is sufficient to capture the width of the coronary. Therefore, with $L = 5$ and $R = 3$, we get a $5 \times 3 \times 9 = 135$ dimensional feature vector. Although a similar sampling pattern was also used in [13], the authors did not fully exploit the fact that calcified lesions can occur anywhere around the centerline. Also, in [13] the gradients were computed with respect to the fixed axes of the volume. In our case the gradient based features are computed within the local coordinate system defined by the cylinder. This leads to an increase in robustness of the proposed method.

4 Learning Methods

We use two different supervised learning techniques to automatically detect the calcified lesions along the given centerline. Both learning algorithms make use of the rotation invariant features discussed above. The first technique trains weak classifiers which are then combined into a strong classifier using the probabilistic boosting tree (PBT) [14]. The second technique uses the random forests (RF) method [15] to construct a *forest* of decision trees and then uses the output of all of the individual trees to make a final classification. Both these methods output a probability that a given point along the coronary centerline falls in the calci-

fied region. The threshold over this probability is varied to obtain the receiving operating characteristic (ROC) curves and a suitable operating point can then be selected on the curve. Both these methods have been used successfully on numerous applications in the past. Below, we briefly explain their main ideas for completeness.

4.1 Probabilistic Boosting Tree

Probabilistic boosting tree (PBT) [14] is a binary tree based learning framework in which each node consists of a number of weak classifiers combined into a single strong classifier. In the training stage, the tree is recursively constructed. The entire data is divided into two sets, based on the decision of the root node. Each of these sets is then used to train the left and right sub-trees recursively. In the testing stage, the conditional probability, $p(y|x)$, that a given sample x belongs to the positive (calcified) or negative (non-calcified) class, is computed at each tree node based on the learned classifier, which guides the probability propagation in its sub-trees. The maximum depth of the tree is defined by the user based on the complexity of the dataset. In our experiments this value was set to 6. PBT very nicely combines AdaBoost algorithm [17] with a binary tree based learning method. For further details on PBT, we refer readers to [14]. We use histogram based weak classifiers for PBT. Each histogram is constructed from one single feature with the number of bins fixed. For each node in the PBT, the histograms of all the 135 features are first constructed. The PBT then picks the 20 best performing histograms as weak classifiers to combine into a strong classifier at that node.

4.2 Random Forests

The random forests classifier [15] is an ensemble of many decision trees. It outputs the class that is the mode of the classes output by the individual trees. Alternatively, the outputs of the individual decision trees can also be combined into a probability mass function over various classes. Each individual decision tree in the forest is grown by picking 8 input variables at random out of the total 135. A total of 100 trees are grown.

5 Experiments

The entire data was divided randomly into four subsets which were then used for a 4-fold cross validation. For training of both the classifiers, the control points annotated along the centerline were used. To compensate for the large number of negative samples in comparison to the small number of positive samples, every two consecutive positive control points were linearly interpolated with three additional points. Further, for every positive control point, eight neighboring points in the plane perpendicular to the centerline, were also added to the positive data. These two types of enhancements of the positive data help to avoid overfitting

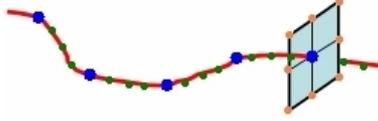


Fig. 3. Additional positive training samples. The blue dots represent the original positive control points, the green dots represent the interpolated points and the orange dots represent the neighboring points in the plane locally normal to the centerline of the coronary.

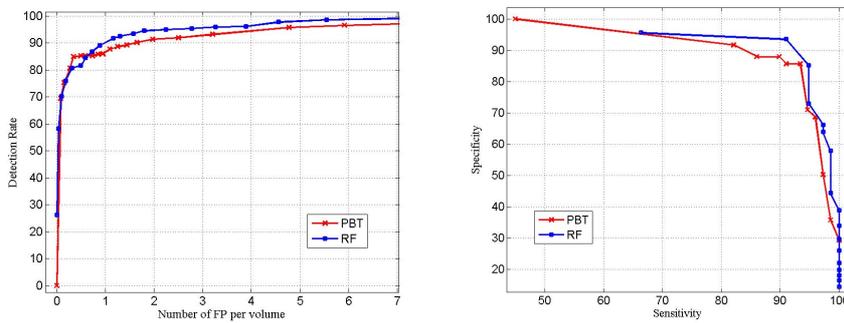


Fig. 4. ROC curves for PBT and RF methods for the average of 4-fold cross validation experiments. *Left:* Number of false positive lesions per scan vs. percentage of correctly detected lesions. *Right:* Sensitivity vs. specificity based on per vessel evaluation scheme.

and compensate for errors in centerline estimation. Figure 3 shows how additional positive training points are obtained from the data using the original positive control points. For each coronary artery, testing was performed on the original set of control points. Figure 4 shows different ROC curves obtained by varying the threshold on the output probabilities of PBT and RF. For lesion based evaluation method, the true detection rate is defined as the percentage of actual lesions detected. For vessel based evaluation method, the sensitivity is defined as the percentage of vessels with lesions that are correctly detected. The corresponding specificity is defined as the percentage of healthy vessels detected correctly as being healthy. We were able to achieve an average detection time of 0.67 and 0.82 seconds per volume using PBT and RF respectively. Figure 5 shows some sample results using both the methods.

Due to significant difference in our detection method compared to other methods, it is very hard to directly compare our results with others. All other methods, due to their employed detection techniques, provide only their *single* best performance values (e.g., best detection/false positive rate on unseen data). On the other hand, our methods output ROCs and a user can then select a suitable operating point on it. For the sake of completeness, we now provide a rough comparison with other methods based on the ROCs in Figure 4. Išgum et al. [9] reported an average detection rate of 73.8% at the cost of 0.1 false positive per scan. By looking at the left ROC in Figure 4, one can observe that we get about

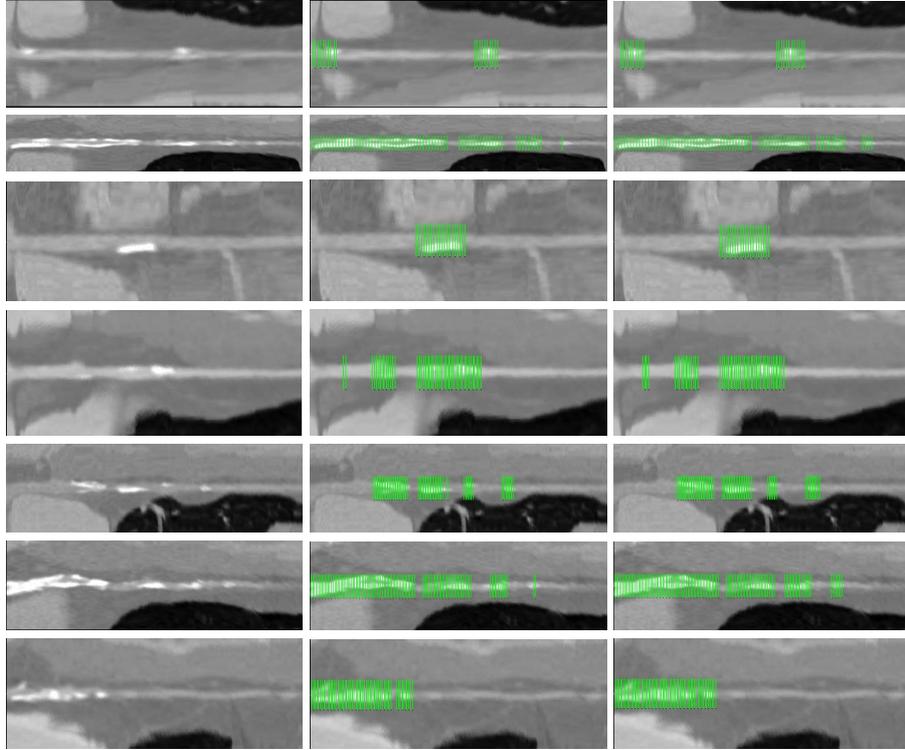


Fig. 5. Sample detection results. Left column shows the original arteries. Middle and right columns show detection results using and probabilistic boosting tree (PBT) and random forests (RF) respectively.

70% detection rate for the same value of false positives per scan using both the methods. Please note that this operating point is not the most optimal for our methods. For example, we achieve about 81% detection rate at the cost of 0.3 false positives per scan. With less than 1 false positive per scan RF can achieve a detection rate of 90%. Kurkure et al. [11] missed 1.74 lesions per scan at the expense of 5.56 false positives per scan. For the same value of false positives, we missed only 0.024 lesions per scan using RF. Similarly, for the same value of missed lesions, there were only 0.19 false positives per scan using RF.

6 Discussion and Future Work

We presented a novel learning based algorithm to detect calcified coronary lesions in cardiac CT images along the centerlines of the coronaries. We evaluated and compared the performance of two supervised learning methods. The methods are very fast and their performance is shown on a large set of data through 4-fold cross validation experiments. The algorithms assume that the centerlines

are given, but are quite robust to their accuracy. As of now, we only perform detection on the three main coronary vessels. Due to this reason, the detection was poor at the coronary bifurcations. In future, we plan to extend our approach to the entire coronary tree. We can then use specialized detection schemes around the bifurcations to enhance the detection results.

References

1. Lloyd-Jones, D., et al.: Heart disease and stroke statistics – 2009 update. *Circulation* 119(3), 21–181 (2009)
2. Moshage, W.E., Achenbach, S., Seese, B., Bachmann, K., Kirchgeorg, M.: Coronary artery stenoses: three-dimensional imaging with electrocardiographically triggered, contrast agent-enhanced, electron-beam CT. *Radiology* 196, 707–714 (1995)
3. Matsuo, S., Nakamura, Y., Matsumoto, T., Nakae, I., Nagatani, Y., Takazakura, R., Takahashi, M., Murata, K., Horie, M.: Visual assessment of coronary artery stenosis with electrocardiographically-gated multislice computed tomography. *Intl. J. of Cardiovascular Imaging* 20, 61–66 (2004)
4. Reddy, G.P., Chernoff, D.M., Adams, J.R., Higgins, C.B.: Coronary artery stenoses: assessment with contrast-enhanced electron-beam CT and axial reconstructions. *Radiology* 208, 167–172 (1998)
5. Achenbach, S.: Cardiac CT: State of the art for the detection of coronary arterial stenosis. *J. of Cardiovascular C.T.*, 3–20 (2007)
6. 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 (2006) 61410G–9
7. Wesarg, S., Khan, M.F., Firle, E.: Localizing calcifications in cardiac CT data sets using a new vessel segmentation approach. *J. of Digital Imaging* 19(3), 249–257 (2006)
8. Toumoulin, C., Boldak, C., Garreau, M., Boulmier, D.: Coronary characterization in multi-slice computed tomography. *Comp. in Cardiology*, 749–752 (2003)
9. Išgum, I., Rutten, A., Prokop, M., van Ginneken, B.: Detection of coronary calcifications from computed tomography scans for automated risk assessment of coronary artery disease. *Medical Physics* 34(4), 1450–1461 (2007)
10. Saur, S.C., Alkadhi, H., Desbiolles, L., Székely, G., Cattin, P.C.: Automatic detection of calcified coronary plaques in computed tomography data sets. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 170–177. Springer, Heidelberg (2008)
11. Kurkure, U., Chittajallu, D., Brunner, G., Yalamanchili, R., Kakadiaris, I.: Detection of coronary calcifications using supervised hierarchical classification. In: *MICCAI workshop on CVII* (2008)
12. Brunner, G., Kurkure, U., Chittajallu, D., Yalamanchili, R., Kakadiaris, I.: Toward unsupervised classification of calcified arterial lesions. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 144–152. Springer, Heidelberg (2008)

13. Teßmann, M., Vega-Higuera, F., Fritz, D.: Learning-based detection of stenotic lesions in coronary CT data. In: Proc. of Vision, Modeling, and Visualization, Konstanz, Germany, pp. 189–198 (2008)
14. Tu, Z.: Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In: ICCV, Beijing, China, pp. 1589–1596 (2005)
15. Breiman, L.: Random forests. *Machine Learning* 45(1), 5–32 (2001)
16. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. *ACM Computing Surveys* 36(2), 81–121 (2004)
17. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. In: Vitányi, P.M.B. (ed.) EuroCOLT 1995. LNCS, vol. 904, pp. 23–37. Springer, Heidelberg (1995)