A pattern recognition approach to automated coronary calcium scoring

I. Išgum, B. van Ginneken
Image Sciences Institute
University Medical Center Utrecht
Utrecht, The Netherlands
{ivana, bram}@isi.uu.nl

M. Prokop
Department of Radiology
University Medical Center Utrecht
Utrecht, The Netherlands
M.Prokop@azu.nl

Abstract

An automated method for coronary calcification detection is presented. First the heart region is extracted, in which objects potentially representing calcifications are obtained by thresholding. Besides coronary calcifications, the set of objects includes other heart calcifications, bony structures and noise. For each object, features describing its size, shape, position and appearance are computed. Several classifiers and classification strategies are evaluated. Best results are obtained with a specifically designed sequence of kNN classifiers that employ sequential forward feature selection. First obvious non-calcifications are removed, then calcifications are distinguished from non-calcifications and a final classifier discerns coronary calcifications from other cardiac calcifications. In 14 CT scans containing 61 coronary calcifications, 46 (75%) are detected at the expense of on average 0.9 false positive objects per scan.

1 Introduction

Coronary artery disease is a leading cause of death in the western world [5]. The amount of calcification present in the coronary arteries – expressed in terms of a coronary calcium score [7] – is related to the presence of coronary artery disease [6]. Clinical studies have shown that a high coronary calcium score is a predictor of cardiac events in both symptomatic and asymptomatic patients [7]. Multislice CT scan, obtained with ECG gating can be used for coronary calcium scoring [2, 6]. Several commercial tools are available for interactive segmentation of coronary calcification: areas of high density in the heart are shown and a user has to select those that represent coronary calcium.

We present a method for automatic detection of calcifications in the coronary arteries. From the obtained results the coronary calcium score can be computed directly. The first strategy for coronary calcification detection to spring to mind may be to segment the coronary arteries and extract all
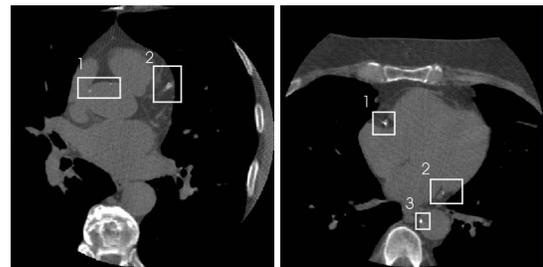
Figure 1. 2D axial views of calcifications in the heart. Left: calcifications in (1) the ascending aorta, (2) the left anterior descending coronary artery. Right: Calcifications in (1) the right coronary artery, (2) left circumflex coronary artery, (3) the descending aorta.

objects in the vessel wall above a fixed threshold. However, to our knowledge no automatic method for segmentation of the non-contrast enhanced coronary artery tree exists as of today. Published methods require the placing of starting points and work on contrast scans in which the arteries are enhanced. Furthermore, non-coronary calcifications in the heart may be very close to the coronaries; to exclude those, the coronary segmentation would have to be extraordinarily precise. Therefore, we use a different approach based on feature extraction and classification. After a segmentation of the heart, candidate objects are extracted by thresholding. For each object a number of features is computed. From that moment, the detection of coronary calcium is a pattern recognition problem. We compare various methods to perform the classification.

2 Materials

For this study 14 CT scans of the heart were used. They were acquired with a 16-detector CT scanner (Mx8000IDT,

Philips, Best, the Netherlands) using prospective ECG gating. No contrast material was present in the scans. Data was reconstructed to 512 x 512 matrices and the in-plane resolution varied from 0.390 mm to 0.429 mm. The slice thickness was 1.5 mm. The scans were cropped in the z-direction to include only slices of the heart, from the bifurcation of the pulmonary artery until the apex.

A gold standard was set by careful inspection of each scan and manually placing a point in each calcified region in the heart, followed by 3-dimensional region growing of the neighboring voxels above 130 HU. This is the generally used threshold value for calcifications [2, 7]. Not all marked calcifications are coronary calcifications. Calcification also occurs in the aorta, in the heart valves and valvular rings, and in the pericardial muscle.

The scans contained 61 coronary calcifications and 36 other calcifications in total; 6 scans contained no coronary calcifications. Figure 1 shows examples of calcifications in the heart.

3 Method

3.1 Heart segmentation

The purpose of the heart segmentation algorithm described here is to obtain a volume of interest from which candidate objects that may represent coronary calcium can be extracted. It is therefore essential that the segmented volume contains the complete heart, while a high accuracy of the segmentation is not the most important goal.

In the majority of slices the heart is partially surrounded by lung tissue with low attenuation. Therefore the original image was thresholded to obtain a lung mask image containing all voxels of intensity values below -200 HU.

The points on the heart border were detected in the following way: A landmark point was manually placed at the bifurcation of the pulmonary artery (see Figure 2). From there, rays were cast in radial directions at 10 degrees angular offset. During casting, the gray level derivative along the ray is computed at a scale of 1 pixel. The point with the most negative derivative is remembered. The algorithm has two stopping criteria: (1) if a point is further from the mask than d pixels and the derivative is bigger than a predefined constant c_1; (2) if the distance to the mask equals zero. Additionally, a point is to be ignored if (3) the derivative is smaller than a predefined constant c_2 and closer to the mask than d pixels. The point on the heart border is the one with the most negative derivative of all the observed points on the ray. Criterion (1) ensures that the process stops when moving from the heart tissue to the bones, usually sternum. Criterion (2) prevents the ray from entering the lungs or area outside the body. Finally, criterion (3) ensures to ignore points on the border of big calcifications close to the
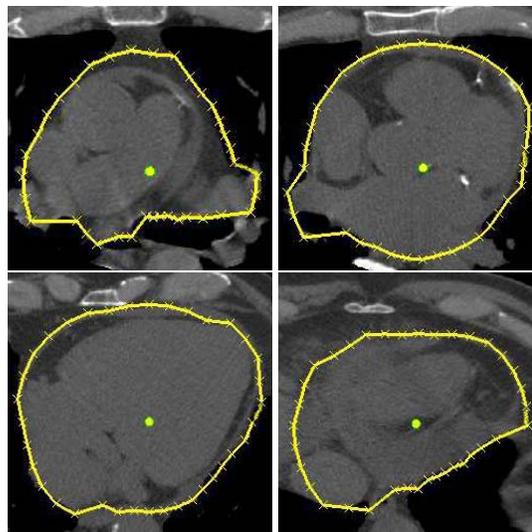

Figure 2. Example of heart segmentation.

heart border. Because they can have a high intensity value, the derivative on their border will be negative and regularly more negative than that on the heart border. The following settings were used: $c_1 = 80$, $c_2 = 150$, $d = 40$ pixels.

After all border points in the slice have been found, the border is smoothed by determining a new distance of each point from the origin of the ray. This new distance is the median radius of 5 points: the observed one and the two closest points on each side. The final result is copied to the next slice and a new center of mass is computed an used for ray casting. In this next slice, points are adjusted following the same procedure but allowing only a restricted displacement from the copied starting position. The minimum and the maximum distance of a new point from the center are: (1) when the starting distance is small, the new one can be at minimum 10% shorter and at maximum 10 pixels longer; (2) otherwise, the new ray can be at minimum 10% shorter and at maximum 10% longer than the starting one. Figure 2 shows examples of segmentation results.

3.2 Candidate and feature extraction

The segmented heart volume was thresholded at 130 HU, and clusters of connected voxels are considered candidate objects that may represent calcifications. Objects with less than 3 or more than 3000 voxels were discarded because they always represented either noise and bony structures partially included in our heart segmentation.

Properties that may distinguish coronary calcifications from other objects include their size, shape, spatial position, appearance (intensity value, and variations) and the appearance of their surroundings. For each candidate object a rich

IEEE
COMPUTER
SOCIETY

set of 66 features is computed that describes these properties.

The coordinates of the object's center of mass were calculated in a local coordinate system with the bifurcation of the pulmonary artery as origin and a scaling factor depending on the heart size. Intensity values were sampled from a 3 by 3 by 3 grid in a cube with center point at the object's center of mass and side length four times the maximum radius of the object. More appearance feature were derived from the image first and second order derivatives in x and y directions calculated at scales of 1, 3 and 5 pixels.

3.3 Classification

For the classification step, any method from pattern recognition theory may be used. We experimented with three different classifiers: a linear and quadratic discriminant classifier (LDC and QDC, respectively) and a k-nearest-neighbor (kNN) classifier [1]. For the kNN classifier, all features were scaled to unit variance.

The simplest approach is to apply these classifiers directly to the candidate objects with the feature set described above. We note, however, that the dimensionality of the feature space is fairly large, and it is unlikely that every feature is equally useful in detecting coronary calcifications. Therefore an optimal set of at most 8 features was determined on the training set using sequential forward selection [4]. Computationally more demanding floating selection schemes [4] did not yield better results. The optimal value for k in the kNN classifier was determined on the training set for the range $k = 1...9$. For feature selection k was fixed at 9.

Finally, we designed a classifier that is intended to cope with two characteristics we observed in our data. First, the prior probability for negatives in our data set is over 99%. Many of these negatives can be easily classified correctly. Second, the negatives contain both calcifications (that are not in the coronaries) and other objects. It is likely that the calcifications outside the coronaries are confused with coronary calcium. To cope with the first observation, we start with a simple classifier that eliminates those objects that have a very high probability of being negative from the training set. To keep the classifier simple, only two features are selected. For LDC and QDC a posterior probability threshold of 0.99 was used, for kNN we set $k = 101$ and discarded those samples with 101 negative neighbors. To deal with the two types of calcifications - positives and negatives - in our data set, a second classifier is trained that considers all calcifications in the training set as positive. At most 8 features are selected and the settings for k are the same as above. Only the samples considered positive are retained in the training set. A last classifier is trained, again with the same feature selection procedure and settings for k, which treats only coronary calcifications as positive. This classification strategy thus is a sequence of three classifications. Because the first classifier is simple and discards most samples, both training and testing of this classifier is fast compared to the single classifier with feature selection.

4 Experiments & results

From the 14 scans, 4253 candidate objects were extracted. Among these were all 61 coronary calcifications and 36 other calcifications.

Because the number of scans in our database is limited, the method is evaluated in a leave-one-scan-out basis, that is, when the system is tested on the n-th scan, all but the n-th scan are used for training (so 14 classification systems are trained with 13 training scans each, and for each system an independent feature selection procedure is performed). Table 1 lists the result for each classification strategy (no feature selection, feature selection, and the sequence of three classifiers) and each type of classifier (LDC, QDC, kNN). Results are given in terms of sensitivity and average number of false positives per scan, and, combining these two criteria in a single figure of merit, the average number of errors per scan (false positives and false negatives).

For each strategy kNN obtains the lowest error rate. The results obtained with the sequential classification strategy and the kNN classifier are substantially better than those of any other system.

Figure 3 shows examples of false negative and false positive objects.

5 Discussion & Conclusions

The method contains three main steps; heart segmentation, feature extraction and classification.

Table 1. Results of coronary calcification detection in 14 CT scans, for each classification strategy and each classifier.

strategy	type	sensitivity	FP/scan	error/scan
no feature selection	LDC	0.64	2.6	4.1
	QDC	0.61	4.4	6.1
	kNN	0.49	1.2	3.4
feature selection	LDC	0.48	2.1	4.4
	QDC	0.70	2.2	3.5
	kNN	0.44	0.9	3.4
sequence classifier	LDC	0.46	2.0	4.4
	QDC	0.49	1.3	3.5
	kNN	0.75	0.9	1.9
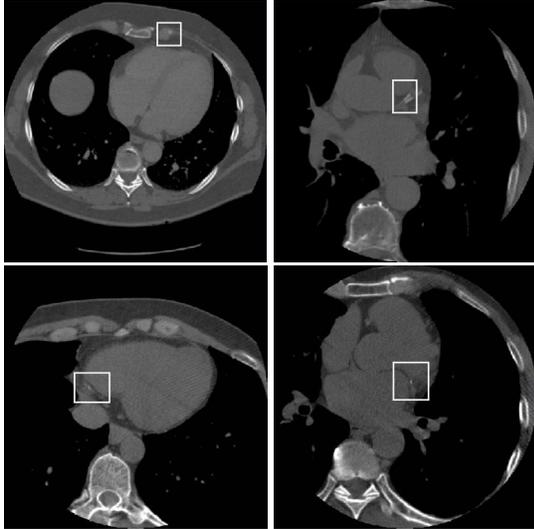
Figure 3. Examples of misclassified objects. Top: false positives (left: sternum, right: stent). Bottom: false negatives (missed coronary calcifications).

The heart segmentation is not very precise, but always included the complete region of interest. The experiment repeated without heart segmentation resulted in many more candidate objects, and an overall sensitivity of 69.9% and a false positive rate of 1.1 objects per scan. False positives that were avoided with the use of heart segmentation were calcifications in the descending aorta and bone in the sternum. The latter type of false positives also occurred with heart segmentation, as incidentally pieces of sternum were included in the segmented heart region. An example can be seen in Fig. 3, top-left. We conclude that it is useful to use the presented heart segmentation method in this context.

The choice of features will be further investigated when a larger database is available. For a more extensive discussion of useful features for calcification detection we refer to [3].

Performing the classification in multiple stages improves performance substantially. The first stage, removal of candidates with a large probability of being negative, is not essential to obtain good performance, but speeds up training and testing significantly. Typically this step eliminates 90% of all samples. The first feature to be selected was always the average intensity inside the candidate object. Five coronary calcifications are misclassified in this stage, but these samples are outliers in feature space and would otherwise have been misclassified in the next stage. In the second and third stage, typically all but one or two of the selected features are different. This indicates that the characteristics that distinguish calcifications from other objects are not the same as those that set apart calcifications inside and outside the coronaries, which motivated the use of this sequential classification strategy. In the second stage, another five coronary calcifications are misclassified. The remaining errors (12 false positives and 5 more false negatives) are made in the final stage.

The results reported here are encouraging, but clearly show that stand-alone automatic coronary calcium scoring is still far from reliable. We believe that to further improve performance a substantially larger database is necessary, in which all varieties of calcifications occur in multiple scans. In the experiments presented, the sequential kNN classifier processed five scans completely correctly, and another five had only one error. The majority of errors (16 out of 27) were made in just two scans, which contained calcifications with characteristics not found in any other scan (for example, all calcifications in the pericardium were contained in one of these scans and three out of total five were misclassified). One scan contained a stent, which was picked up as a false positive (Fig. 3, top-right).

Future work will focus on enlarging the database, investigation of different features, automating the detection of the bifurcation of the pulmonary artery and the heart apex and evaluation of the results in terms of calcium score.

In conclusion, we have presented a system to automatically detect coronary calcifications from CT scans. Using a sequence of three classifiers, we detect 75% of coronary calcifications at the expense of 0.9 false positive responses per scan. The method provides the first steps towards (semi-)automated coronary calcium scoring.

References

[1] R.O. Duda, P.E. Hart, D.G. Stork. *Pattern Classification, 2nd ed*. Wiley, New York, 2001.

[2] L.R. van Hoe, K.G. de Meerleer, P.P. Leyman, P.K. Vanhoenacker. Coronary artery calcium scoring using ECG-gated multidetector CT. *American Journal of Roentgenology*, 181(4):1093–100, 2003.

[3] I. Isgum, B. van Ginneken, M.A. Viergever. Automatic detection of calcifications in the aorta from abdominal CT scans. In *Computer Assisted Radiology and Surgery*, 1037–1042, 2003.

[4] A. Jain, D. Zongker. Feature selection: Evaluation, application, and small sample performance. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 19(2):153–158, 1997.

[5] G.T. Kondos, J.A. Hoff, A. Sevrukov, M.L. Daviglus, D.B. Garside, S.S. Devries, E.V. Chomka, K. Liu. Electron beam tomography coronary artery calcium and cardiac events. *Circulation*, 107(20):2571-2576, 2003.

[6] N. Takahasi, K.T. Bae. Quantification of coronary artery calcium with multi-detector row CT. *Radiology*, 228(1):101-106, 2003.

[7] J.A. Rumberger, L. Kaufman. A Rosetta stone for coronary calcium risk stratification. *American Journal of Roentgenology*, 181(3):743–748, 2003.