Computer-Assisted Radiology and Surgery

Original Investigations

Automatic Detection of Calcifications in the Aorta from CT Scans of the Abdomen[1]

3D Computer-Aided Diagnosis

Ivana Isgum, MSc, Bram van Ginneken, PhD, Marco Olree, MD

Rationale and Objectives. Automated detection and quantification of arterial calcifications can facilitate epidemiologic research and, eventually, the use of full-body calcium scoring in clinical practice. An automatic computerized method to detect calcifications in CT scans is presented.

Materials and Methods. Forty abdominal CT scans have been randomly selected from clinical practice. They all contained contrast material and belonged to one of four categories: containing "no," "small," "moderate," or "large" amounts of arterial calcification. There were ten scans in each category. The experiments were restricted to the vertical range from the point where the superior mesenteric artery branches off of the descending aorta until the first bifurcation of the iliac arteries. The automatic method starts by extracting all connected objects above 220 Hounsfield units (HU) from the scan. These objects include all calcifications, as well as bony structures and contrast material. To distinguish calcifications from non-calcifications, a number of features are calculated for each object. These features are based on the object's size, location, shape characteristics, and surrounding structures. Subsequently a classification of each object is performed in two stages. First the probability that an object represents a calcification is computed assuming a multivariate Gaussian distribution for the calcifications. Objects with low probability are discarded. The remaining objects are then classified into calcifications and non-calcifications using a 5-nearest-neighbor classifier and sequential forward feature selection. Based on the total volume of calcifications determined by the system, the scan is assigned to one of the four categories mentioned above.

Results. The 40 scans contained a total of 249 calcifications as determined by a human observer. The method detected 209 calcifications (sensitivity 83.9%) at the expense of on average 1.0 false-positive object per scan. The correct category label was assigned to 30 scans and only 2 scans were off by more than one category. Most incorrect classifications can be attributed to the presence of contrast material in the scans.

Conclusion. It is possible to identify the majority of arterial calcifications in abdominal CT scans in a completely automatic fashion with few false positive objects, even if the scans contain contrast material.

Key Words. Calcifications; computer-aided diagnosis; abdominal CT.

RATIONALE AND OBJECTIVES

Recent studies have shown that the presence and extent of vascular calcifications is associated with a higher risk for a variety of diseases. Vliegenthart and co-workers (1) observed a strong and graded association between the amount of coronary calcification and the presence of myocardial infarction in an elderly population. Kondos and co-workers (2) determined an association between coronary artery calcium and cardiac events in an initially asymptomatic middle-aged population. Pohle and co-

Acad Radiol 2004; 11:247–257

[1] From the Image Sciences Institute, University Medical Center Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands. (I.I., B.v.G) and the Department of Radiology (M.O.), University Medical Center Utrecht, Utrecht, The Netherlands. Received August 15, 2003; revision requested September 29; revision received October 3; revision accepted October 10. **Address correspondence to** I.I. e-mail: ivana@isi.uu.nl

doi:10.1016/S1076-6332(03)00673-1

workers (3) showed that nearly all young patients with myocardial infarction have calcifications in their coronary arteries. In healthy controls, calcifications are significantly less common. De Leeuw and co-workers (4) showed a relation between aortic atherosclerosis and white matter lesions in the brain, which are associated with cognitive impairment and dementia. Aortic atherosclerosis was considered present when radiographs showed calcified deposits in the abdominal aorta. Using the same measurement technique, Witteman and co-workers (5) found that a decline in diastolic blood pressure, which is related to mortality from coronary heart disease, is associated with the presence of calcifications in the aorta.

For further study of the relation between calcifications and risk for disease, it is necessary to conduct large-scale epidemiologic studies. Computerized methods that detect calcifications automatically would facilitate such research. Another reason to develop such methods is the advent of new generations of multi-detector row CT scanners that can perform fast acquisitions of the complete body. From these scans a total body calcium score can be determined. This requires processing datasets of over 1500 slices, which makes computer assistance mandatory. In this work a method is presented for automatic detection of calcifications. It is applied to abdominal CT and CTA scans in which calcifications in the aorta and iliac arteries are detected.

In clinical practice, the presence of calcifications can be assessed visually or, more commonly, a region of interest is selected and all clusters in the region above a certain threshold are determined. Usually the threshold value is 130 Hounsfield units (HU). Typically, small clusters above the threshold are ignored because they are considered to represent noise instead of true loci of calcium. From the larger clusters the Agatson score (6) or other measures, such as calcium mass, which may be more reproducible (7), can be determined.

The most difficult part of automating this procedure is determining the region of interest. For example, a fully automatic segmentation of the aorta and iliac arteries in CT scans with variable slice thickness and in which contrast may be present or absent is a very difficult task. Published vessel segmentation methods are usually interactive or put constraints on the types of scans to which they can be applied. Therefore we adopted a different approach that circumvents the segmentation problem. Our method starts with thresholding the complete scan. In this way many objects are obtained. They include all calcifications, but these are a small minority among objects representing bony structures or contrast material. A set of features is computed for each object and these features are used to determine which objects are calcifications. The purpose of this study is to show that with this approach it is possible to correctly identify a large majority of calcifications in CT scans with and without contrast material. We also investigate if we can automatically categorize the complete scan based on its total amount of calcification.

MATERIALS AND METHODS

Data

For this research, data from an ongoing clinical study of the University Medical Center Utrecht, The Netherlands, was used. In that study the presence of calcifications in the abdominal aorta is linked to risk factors for atherosclerotic disease. The patients included were randomly selected from clinical practice. They were examined for various indications unrelated to vascular disease, and scans were acquired with different protocols. Contrast was administered either intravenously, orally, or both in all examinations. Scans were acquired with Philips Tomoscan AV scanners (Philips, Best, The Netherlands). Data was reconstructed to 512 $\times$ 512 matrices. Slice thickness varied from 5–7 mm and slices were reconstructed every 3–5 mm. The in-plane resolution varied from 0.58 mm–0.74 mm, depending on patient size.

In the aforementioned clinical study, patients are assigned one of four categories based on the amount of calcification encountered. The category labels are: “none,” “small,” “moderate,” and “large,” and they are defined as follows:

“none”—no visible calcifications

“small”—fewer than 5 locations of calcium, each with a maximum of 5 mm diameter in the axial slice

“moderate”—up to15 locations of calcifications, each with a maximum of 15 mm diameter in the axial slice

“large”—otherwise

Following the above guidelines and based on the visual inspection, a radiologist (M.O.) associated a category label with each scan.

For this study, we randomly selected 10 scans of each category. Only those slices from the point where the superior mesenteric artery is branching off the descending aorta until the first bifurcation of the iliac arteries are taken into account (Figure 1).

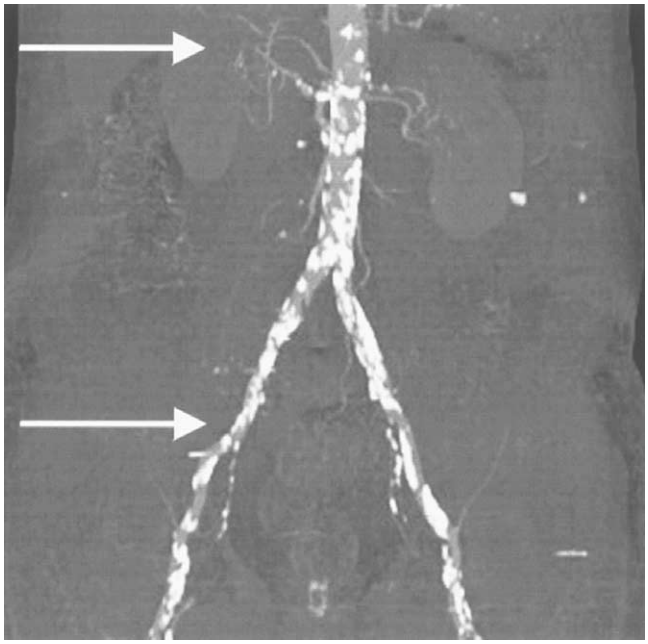

Figure 1. Maximum intensity projection from an abdominal CT scan displays calcifications in the aorta and iliac arteries. The arrows indicate the vertical range from the superior mesenteric artery to the first bifurcation of the iliac arteries. This range was used in the experiments.

References Standard

A calcification is commonly defined to be a volume of density of more than 130 HU (8). Because of the administered contrast in the scans, intensity values in the aorta are higher than 130 HU and that value is therefore not a suitable threshold level to extract calcifications. A higher value of 220 HU was chosen. This threshold level corresponds well with a visual assessment of the size of a calcification. Only clusters of three and more connected voxels and less than 2000 connected voxels (using 26-connectivity in 3 dimensions) above the threshold of 220 HU were taken into consideration. Clusters smaller than three voxels usually represent noise, and the calcifications are objects much smaller than 2000 voxels. Clusters extracted in this way will either represent calcifications or bony structures such as ribs, spine, or contrast material. Figure 2 shows an example of a slice from a dataset and all extracted clusters.

To determine a reference standard, a medical student marked all clusters that represented calcifications in each of the 40 scans in the database. The student received training from a radiologist (M.O.) and in case of doubt, consulted the radiologist.

Method Overview

The method consists of the following three steps:

1) Extracting candidate objects;
2) Calculating features for each candidate object;
3) Classification of candidate objects into calcifications and non-calcifications.

Based on the result of the third step, the scan is categorized as containing "none," "small," "moderate," or "large" amounts of calcifications.

The candidate objects are extracted by thresholding at 220 HU, as explained above, followed by component labeling in 3D using 26-connectivity (9), and discarding objects smaller than 3 and larger than 2000 voxels in size. The features that are computed for each object are described in detail in the next section. Once an object is represented by a feature vector, it can be classified as calcification or non-calcification using any classifier from statistical pattern recognition (10). We adopted a two-stage classification strategy explained below.

Object Features

For each candidate object, a total number of 18 features is calculated. They can be divided in two groups:

(i) features computed directly from the candidate object and
(ii) contextual features extracted from the object's surrounding.

Features computed from the candidate objects.—These include spatial, intensity level, size, and shape features.

Spatial features describe the location of an object in the scan. Calcifications are located in the vessels positioned around the center of the body while other candidate objects may be located anywhere in the scan. Some of them, for example ribs, will always be far from the body center. Therefore it seems useful to consider the spatial position of the center of mass of an object as a feature. It is calculated in a local coordinate frame to correct for varying patient size and position within the field of view. The body is extracted by thresholding, a box is placed around it and the lower-left corner of the box represents the position (0, 0, 0) and the upper right corner the position (1, 1, 1).

Intensity features are the average and maximum HU of the object. These values are usually higher for calcifications than for other objects.

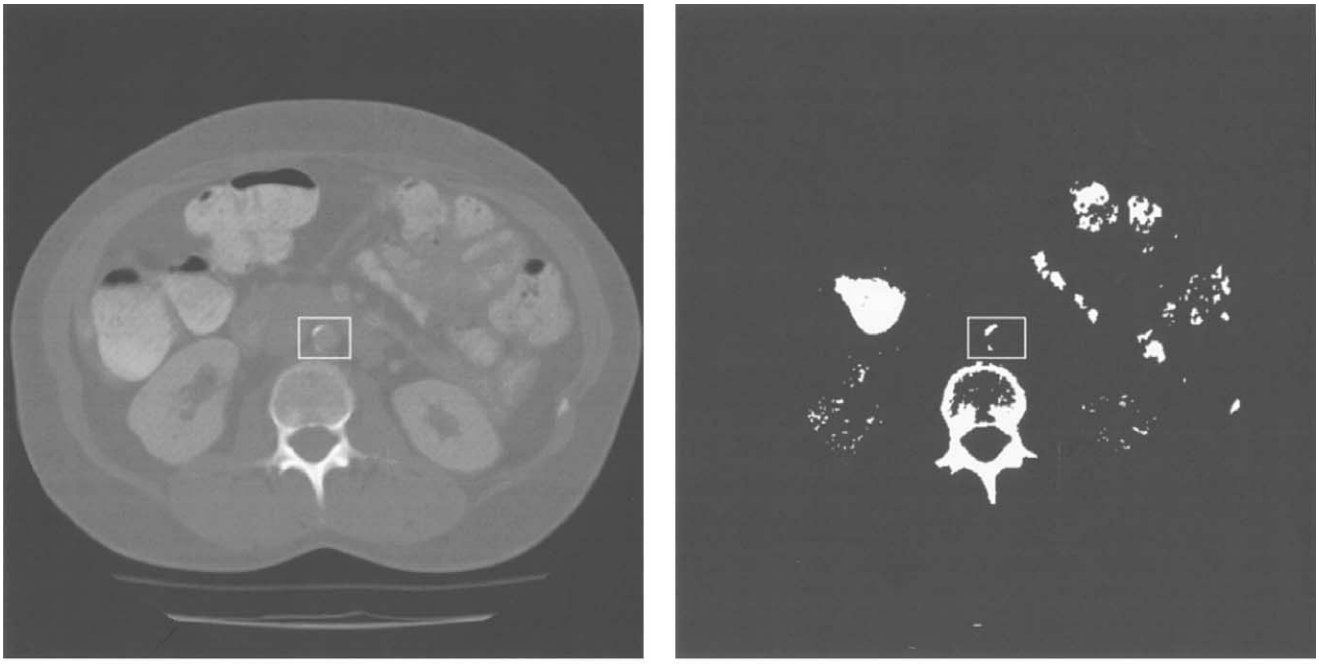

a. b.

Figure 2. **(a)** A slice from CT scan. **(b)** A 2D view of 3D candidate objects extracted by thresholding at 220 HU. Objects represent either calcifications (inside the square on both images), bony structures (a rib and spine in this slice) or contrast material (in this slice a large amount of contrast material in the bowels).

The object's volume (given by the number of voxels) is used as a size feature. The volume of objects can discriminate calcifications from bony structures that are typically much larger. Another size feature is the area of the object in the slice containing its center of mass. This slice will be referred to as the *central slice* in the remainder of this article. The volume of a calcification can be similar to the volume of objects containing contrast material or bony structures, but because of shape differences, the area in the central slice is often different. Unlike other features extracted directly from the candidate objects, which are calculated in three dimensions, this is a 2D feature.

Shape features calculated are the maximum and average distance from the object's center of mass to the border and its compactness. Features related to shape can discriminate calcifications from bony structures that generally have a different shape. Compactness is a feature that measures how spherical an object is. It is calculated as

$$C = \frac{s^3}{36\pi V^2} \quad (1)$$

where s is the object's surface area (calculated as the number of voxels on the boundary of the object), and V is its volume. For a sphere, $C = 1$. The more the object deviates from a sphere, the higher C will be.

Contextual features.—Calcifications are usually located in the vessel wall. In the data used in this study, almost all calcifications are in the abdominal aorta or the iliac arteries. The direction of these vessels is mainly perpendicular to the axial plane. It is therefore to be expected that if one examines the area around a candidate object in the central slice, it will look like the cross-section of a vessel only when the object is a calcification. For the other objects, their surroundings will look different. We have developed two algorithms to determine the cross-section of the vessel in the central slice. In the first approach the cross-section of the vessel is determined based on the intensity values around the object, and the computed feature is a measure of its circularity. In the other approach, a circle is fitted around the vessel and features are extracted based on the gray values inside the fitted circle. In the former case it is not guaranteed that a circular structure matching the vessel will be found, while in the latter one, it is guaranteed that a circle will be fitted, but it might not contain a circular structure.

The first approach to detect the cross-section of the vessel is based on the assumption that the intensity values inside the vessel are higher than in its surroundings. The initial threshold level of 220 HU used to extract candidate objects is lowered in small steps. At a certain level, the

calcification will grow to the vessel's cross-section. At every step the compactness of the grown object is calculated according to

$$cp = \frac{perim^2}{4\pi A} \quad (2)$$

where *perim* is the perimeter of the object's border and *A* is the cross-sectional area of the object. For a circle, $cp = 1$. Less-circular objects have higher values for *cp*. The value of *cp* closest to one, for all thresholds considered, is used as a feature. As the cross-section of the vessel is circular, we expect this *cp* to be close to 1 for a calcification. On the other hand, in the case of a non-calcification, a candidate object can grow to any shape and it is expected that its compactness will be lower. This is illustrated in Figure 3. The algorithm describing the procedure can be found in Appendix A.

The second approach to detecting the vessel cross-section is by fitting a circle to the image gradient. If a candidate object is a calcification, the area inside the vessel cross-section is brighter than the area surrounding the vessel. The mean intensity value in the circle will be similar to the HU of blood or higher if the vessel contains contrast material. Also, there is a little variation in the intensity values inside the vessel compared with its surroundings. On the other hand, when a candidate object is not a calcification, it is expected that there will be more variation in the intensity values inside the circle. This is illustrated in Figure 4. Seven features are computed based on the intensity levels inside the fitted circle, inside a ring around that circle, and inside the object. The width of the ring is equal to half the radius of the fitted circle. The algorithm describing how the circle is fitted is presented in Appendix B. Table 1 enumerates the complete set of features computed for each object.

Classification

The classification of objects is performed in two stages.

In the first stage the goal is to reduce the number of candidate objects while retaining most of the calcifications in the data set. From the feature vectors of the calcifications in the train set, the mean $\bar{c}$ and covariance matrix S_c are computed. Now for each object c_i the Mahalonobis distance (10) can be computed:

$$f(c_i) = (c_i - \bar{c})S_c^{-1}(c_i - \bar{c}). \quad (3)$$

The Mahalonobis distance is inversely proportional to the probability that the object is a calcification, assuming that the calcification feature vectors follow a Gaussian distribution. A threshold on the Mahalonobis distance was determined such that 99% of the calcifications from the train set would be retained. This was used to discard objects from train and test set: only when the Mahalonobis distance for an object is below the threshold is the object retained.

In the second stage the goal was to classify the remaining objects as calcifications or non-calcifications. A 5-Nearest Neighbor classifier (10) was used. In the train set for every feature a scaling was computed which normalized the feature to zero mean and unit variance. Subsequently sequential forward selection (SFS) (11) was performed. This means that features are added in a stepwise fashion to give the best result on the train set. The samples in the test set were scaled with the same scaling factors as determined for the train set. Using the features selected for the train set and using the classifier trained on that set, objects in the test scan are classified.

RESULTS

The 40 scans contained 1756 slices in the vertical range of interest. The region of interest contained between 163 and 2583 candidate objects. In total 12,838 candidate objects were detected, of which 249 were identified as calcifications by the human observer who set the reference standard. The non-calcifications were bony structures, such as ribs or spine, and contrast material in bowels and vessels.

In both classification stages, leave-one-out experiments were performed, meaning that each scan was used as a test set while the remaining scans were used for training.

After the first stage of classification where objects were eliminated, the number of objects to be classified ranged from 62–1381 per scan, totaling 11,510 objects, among which the total number of calcifications was 242. Thus, seven calcifications (2.8%) were discarded and also 1328 non-calcification objects (10.3%).

The average results and standard deviations of the second classification stage per scan were: accuracy 99.2 $\pm$ 1.2%, sensitivity 85.6 $\pm$ 22.8% and specificity 99.5 $\pm$ 0.9%. The average sensitivity is computed by excluding scans without calcifications because there sensitivity is not defined.

In the second stage of classification, on average 10 features were selected. The most often selected features

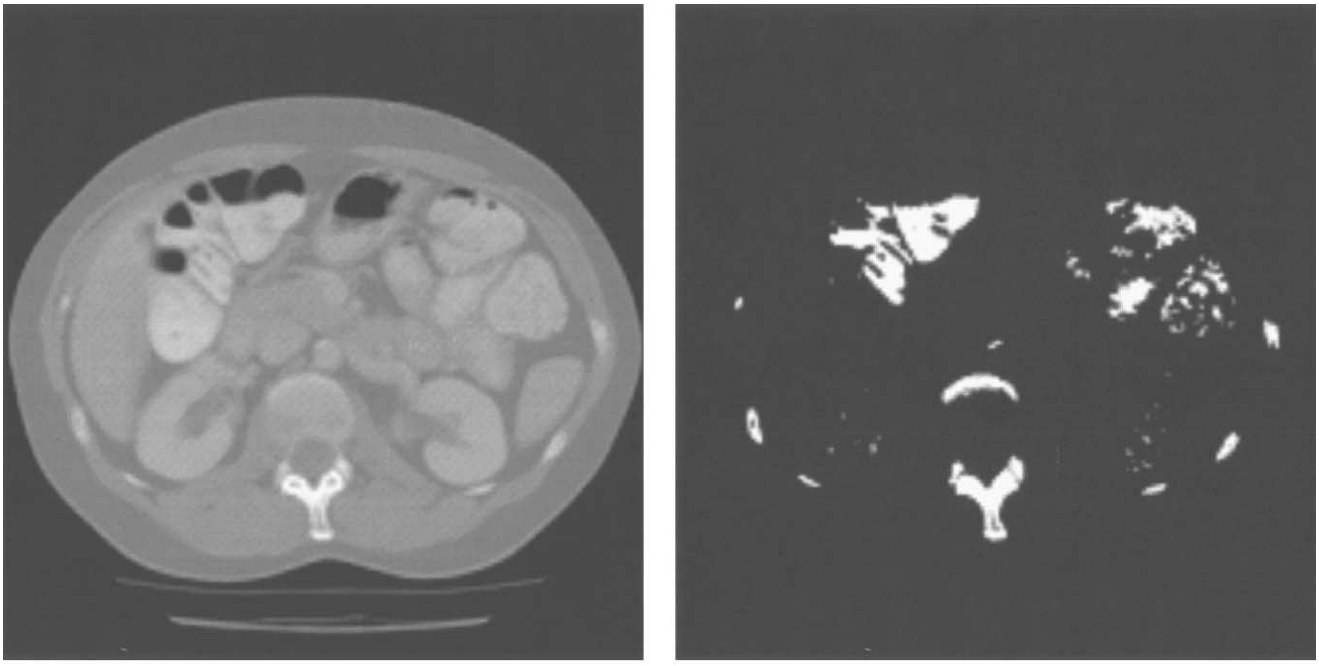

a. b.

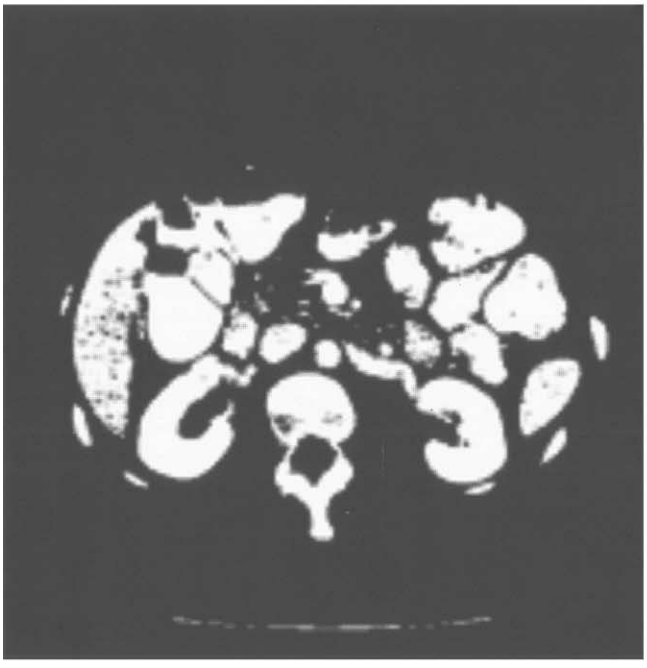

c.

Figure 3. CT slice, shown in **(a)**, thresholded at 220 HU is shown in **(b)**. The goal is to detect the cross-section of the vessel inside which the calcification is located. When the threshold value is lowered, candidate objects grow in size. At some point, the calcification will match the cross-section of the vessel, which has a circular shape. Most of the other objects will not grow to circular structures. Image **(c)** shows a slice thresholded at the value when the calcification has grown to the most circular shape. It can be seen that the grown structure in **(c)** matches the cross-section of the vessel in **(a)**. Compactness of the grown object is calculated as a contextual feature for the given calcification.

were features 4, 1, 2, 13, 14, 15, 17, 10 (in that order) listed in Table 1.

In total 249 calcifications were manually segmented and after both classification stages 209 were detected (true positives), but also 40 non-calcifications were marked as calcifications (false positives). The results per scan category are given in Table 2. The overall result of the system is a sensitivity of 83.9% (=209/249) at the expense of on average 1.0 false-positive object per scan.

The 40 false positives (FPs) were mostly contrast material in bowels, objects inside the spine, and in a few cases, calcifications in renal arteries and contrast in kidneys. Table 3 gives the breakdown of false-positive objects.

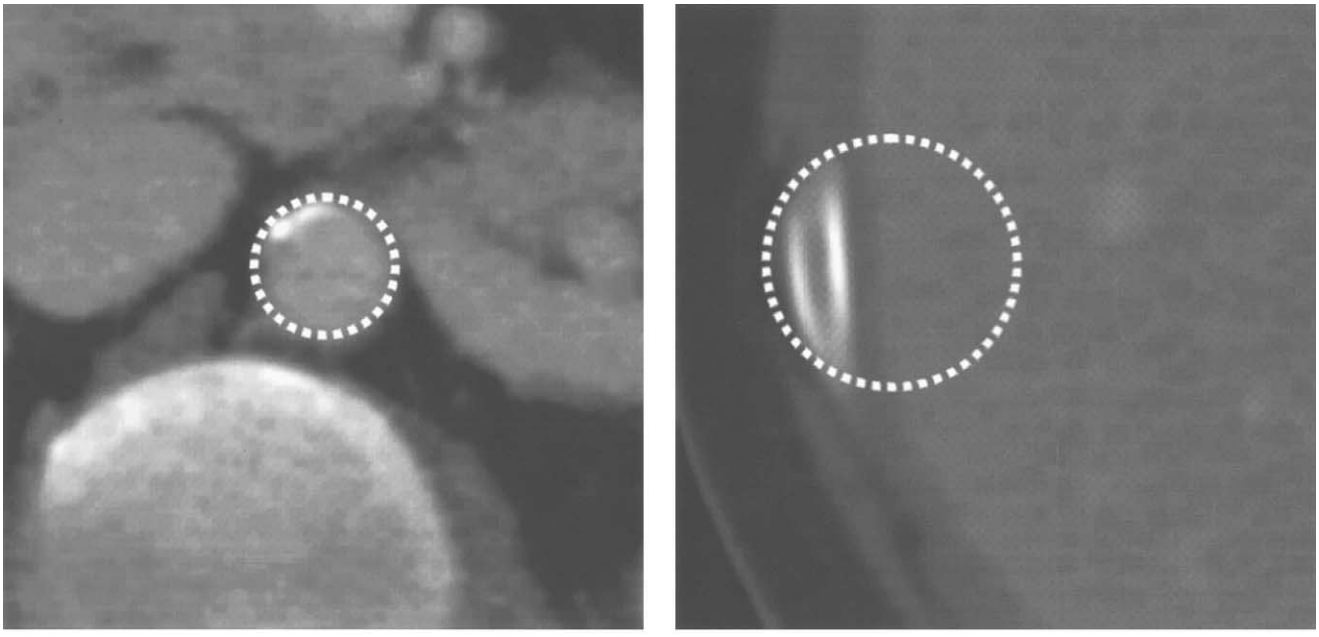

a. b.

Figure 4. Seven features are calculated after a circle has been fitted around the candidate object. In the case of a calcification, we want to find the cross-section of the vessel inside which the calcification is located. An example is shown in image **(a)**. The calcification is located in the upper left area of the circle, which encloses the cross-section of the vessel. The cross-section of the vessel has little variation in the intensity values and it is brighter than the area around it. When a circle is fitted around a non-calcification, for example around a rib illustrated in image **(b)**, there is more variation in intensity values inside the circle.

The 40 false negatives (FNs) can be divided into two groups: large calcifications of an unusual shape, often around a bifurcation of the aorta (25%), and small calcifications with intensity values similar to the intensity values in the aorta (75%).

Figure 5 shows 2D views of typical examples of FPs and FNs.

Further, we wanted to investigate if the method could be used to assign the category labels "none," "small," "moderate," and "large" amounts of calcifications to a scan, as was done in the clinical study from which the data originated. To determine the thresholds on the number of voxels for each category, we related the number of voxels in the reference standard to the category assigned by the study. Using these threshold values we labeled scans according to the amount of calcification detected by the computer method. Out of 40 scans, 30 scans were correctly labeled and only 2 cases were more than one category off. The results are given in Table 4. The percentage of agreement between the method and the reference standard is 75% and the weighted kappa statistic (12) is 0.76. To compare the performance of the computer method with the performance of a human observer, one of the authors (I.I.) manually marked all calcifications in 20 out of 40 scans. The observer correctly marked all but

Table 1
Contextual Features

Features computed from the object
1. x-coordinate of the object's center of mass
2. y-coordinate of the object's center of mass
3. z-coordinate of the object's center of mass
4. Average intensity inside the object
5. Maximum intensity inside the object
6. Average distance from the object's border to its center of mass
7. Maximum distance from the object's border to its center of mass
8. Volume
9. Compactness
10. Area in the central slice
Contextual features
11. Average intensity inside the detected circle
12. Standard deviation of the intensity inside the detected circle
13. Average intensity inside the ring around the detected circle
14. Standard deviation of the intensity inside the ring of the detected circle
15. Difference of features 11 and 13
16. Difference of features 12 and 14
17. Difference of intensities inside the detected circle and inside the object
18. Compactness of the grown candidate object

Table 2
Overall Classification Results, and Results Per Scan Category

	Number of scans*	TP*	FP*	FN*
Large amounts of calcification	10	133 (119.7)	5 (6.6)	24 (247.2)
Moderate amounts of calcification	10	58 (78.2)	11 (97.4)	12 (29.5)
Small amounts of calcification	10	18 (47.3)	11 (62.9)	4 (105)
No calcifications	10	0 (-)	13 (27.8)	0 (-)
All	40	209 (101.9)	40 (54.0)	40 (158.2)

*The numbers refer to the number of objects; the values in parentheses represent the average volume in voxels.
Abbreviations: TP, true positives; FP, false positives; FN, false negatives.

1 of the calcifications and only marked 2 objects incorrectly. Clearly a human observer performs significantly better than the computer method.

To evaluate whether contextual features, two-stage classification, and feature selection are essential for good performance of the method, the additional experiments without these components were performed. Results are listed in Table 5 and the following questions were answered:

Does the use of contextual features improve the system? — Yes. From the first two rows in Table 5 it can be seen that adding contextual features results in better performance in terms of FNs and FPs as well as in volume of misclassified objects.

Does two-stage classification improve the system? — Somewhat. The second and third rows in Table 5 show similar results in terms of FNs and FPs, but the FP volume decreased by a factor of three. However, it seems likely that other classification strategies might yield similar performance. It is important to realize that for assigning category labels to a scan we used the detected *volume,* and users of the system may also be primarily interested in volume measurements, while the system works on a *per object* basis.

Does feature selection improve the system? — Somewhat. Comparing the last two rows in Table 5, it can be seen that feature selection decreased the number of FPs by 16%. Simply using all features also yields good results though, and it was not possible to identify only a very small set of features that yielded similar performance. Again, other classification techniques — which do not use sequential forward feature selection — might lead to similar performance.

Table 3
Breakdown of FP Objects into Contrast Material, Bony Structures and Calcifications in the Renal Arteries

	False positive objects*
Contrast material	27 (72.6)
Bony structures	10 (4.7)
Other calcifications	3 (50.3)

*Values in parentheses are the average volume in voxels.

DISCUSSION

The presented method has a high sensitivity and low false-positive rate, but the performance of a second observer shows that the task is not extremely difficult and there is undoubtedly room for improvement for the computerized method. From the analysis of false positive and false negative objects it is clear that a near-perfect result would be obtained if no contrast material were mistaken for calcifications and when the relatively rare large calcifications of an unusual shape were correctly detected. In that case only very small low-contrast loci of calcium in the aorta would be missed and the only false positives would be very small clusters of bony material in the spine.

Obviously the simplest way to avoid the incorrect classification of contrast material is to apply the method to scans without contrast. Another possibility would be to first exclude the bowel region from the scans or to remove contrasted colonic material and then apply our method. Some methods for that purpose have been published (13,14), and it would be interesting to investigate if they are sufficiently robust to be used in this application. A third possibility is to use the method in an interactive setting. A user could quickly remove incorrectly detected objects from a processed scan.

The large calcifications that are missed are usually located around the bifurcation of the aorta. They consist of several clusters of calcium that have grown together. The most likely reason for this failure is the fact that large clusters can display many shapes and each type is rare

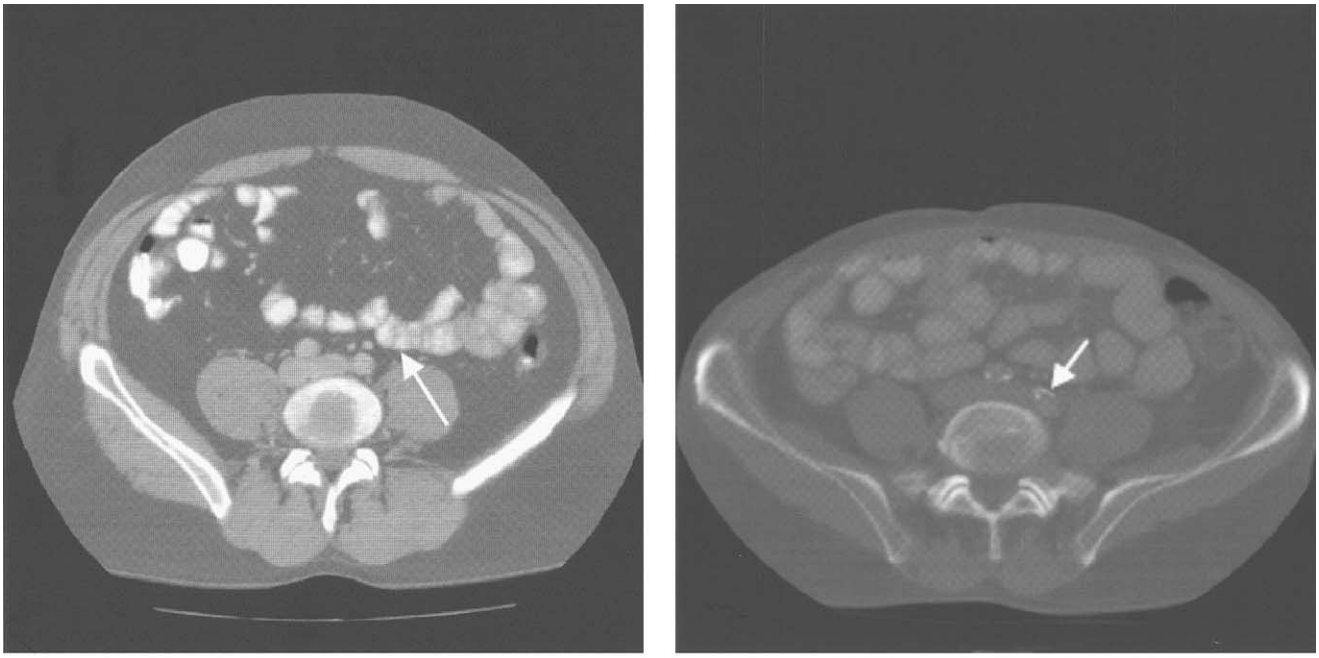

a. b.

Figure 5. A 2D view of typical examples of incorrectly classified objects. Image **(a)** shows a false positive object representing contrast material in bowels. Although it is obvious to a human observer that this cannot be a calcification, its size, location and shape are similar to calcifications in the training set. Image **(b)** shows a 2D view of a large calcification of an unusual shape, just below the bifurcation. Such large interconnected calcified clusters are relatively rare in the training data and therefore these objects are likely to be misclassified.

Table 4
Category Labels Assigned to Scans Based on the Total Volume of Calcifications*

	Classified as Large	Classified as Moderate	Classified as Small	Classified as None
Large (ref.)	9	1	0	0
Moderate (ref.)	3	7	0	0
Small (ref.)	1	1	7	1
None (ref.)	0	1	2	7

*Rows indicate category labels assigned in the reference standard (ref.) and columns show labels assigned by the method.

within our training set. Training the classifier on a larger data set that includes a larger range of examples of such calcifications might improve performance. Another option is to add additional algorithms that focus on larger calcifications. Note that in an interactive setting a human user could also add the missed calcifications.

The system is currently based on a single fixed threshold level of 220 HU. In the literature threshold values ranging from 80–230 HU have been reported (15). It would be interesting to perform an automatic threshold level selection or a segmentation that is not based on thresholding for each object to match its contour more precisely. Such a refinement might also improve the reproducibility of the method, which has not been evaluated yet.

In future work we intend to use the system on other parts of the body as well. One modification is necessary to do this. In general, one cannot assume that the direction of the vessel is perpendicular to the axial plane, as was done here. The direction of the vessel could be estimated using the local Hessian matrix (16).

An advantage of the presented method is that it does not require a segmentation of the vessels of interest. However, if the method would be applied to full body scans or the heart, it would be important to know precisely in which vessels the calcifications are located. For example, in the heart it is essential to be able to distinguish coronary calcium from calcifications in the aortic valves and the aorta (17). It might be possible to infer that location from features extracted from the object's surrounding, by using contextual information as was done in features 11–18 in our method. In that way, the need for a complete segmentation of the vasculature—a daunting task if no intravenous contrast material is administered—could be avoided.

Table 5
Additional Experiments Were Performed With Systems Without Contextual Features, the First Classification Stage, or Feature Selection.

	False negatives*	False positives*
No SFS; features 1–10; only 2nd classification stage;	63 (91.4)	63 (152.5)
No SFS; all features; only 2nd classification stage;	41 (112.2)	49 (143.3)
No SFS; all features; 1st + 2nd classification stages;	40 (130.3)	48 (47.1)
SFS; all features; 1st + 2nd classification stages;	40 (152.8)	40 (53.9)

*Values in parentheses are the average volume in voxels.

CONCLUSION

A method for detecting calcifications in the aorta from CT scans of the abdomen has been presented. Out of 249 calcifications marked in the ground truth, 209 were correctly detected by the method at the expense of on average 1.0 false positive object per scan. Most of these false-positive objects are contrast material. The method is completely automatic and does not require any vessel segmentation.

REFERENCES

1. Vliegenthart R, Oudkerk M, Song B, et al. Coronary calcification detected by electron-beam computed tomography and myocardial infarction, The Rotterdam Coronary Calcification Study. Eur Heart J 2002; 23:1596–1603.
2. Kondos GT, Hoff JA, Sevrukov A, et al. Electron-beam tomography coronary artery calcium and cardiac events: a 37-month follow-up of 5635 initially asymptomatic low- to intermediate-risk adults. Circulation 2003; 107:2571–2576.
3. Pohle K, Ropers D, Mäffert R, et al. Coronary calcifications in young patients with first, unheralded myocardial infarction: a risk factor matched analysis by electron beam tomography. Heart 2003; 89:625–628.
4. De Leeuw FE, De Groot JC, Oudkerk M, et al. Aortic atherosclerosis at middle age predicts cerebral white matter lesions in the elderly. Stroke 2000; 31:425–9.
5. Witteman JCM, Grobbee DE, Valkenburg HA, et al. J-shaped relation between change in diastolic blood pressure and progression of aortic atherosclerosis. Lancet 1994; 343:504–07.
6. Agatston AS, Janowitz WR, Hildner FJ, et al. Quantification of coronary artery calcium using ultrafast computed tomography. J Am Coll Cardiol 1990; 15:827–832.
7. Hong C, Bae KT, Pilgram TK, et al. Coronary artery calcium quantification at multi-detector row CT: Influence of heart rate and measurement methods on interacquisition variability - initial experience. Radiology 2003; 228:95–100.
8. Gayard P, Garcier JM, Boire JY, et al. Spiral CT quantification of aorto-renal calcification and its use in the detection of atheromatous renal artery stenosis: A study in 42 patients. Cardiovascular and Interventional Radiology 2000; 23:17–21.
9. Jain AK. Fundamentals of digital image processing. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1989.
10. Duda RO, Hart PE, Stork DG. Pattern classification. 2nd ed. New York, NY: Wiley, 2001.
11. Jain A, Zongker D. Feature selection: Evaluation, application, and small sample performance. IEEE Transaction on Pattern Analysis and Machine Intelligence 1997; 19:153–158.
12. Altman DG. Practical Statistics For Medical Research. London: Chapman & Hall, 1991.
13. Masutani Y, Yoshida H, MacEneaney PM, et al. Automated segmentation of colonic walls for computerized detection of polyps in CT colonography. J Comput Assist Tomogr 2001; 25:629–638.
14. Chen D, Liang Z, Wax MR, et al. A Novel approach to extract colon lumen from CT images for virtual colonoscopy. IEEE Transactions on Medical Imaging 2000; 19:1220–6.
15. Hong C, Bae KT, Pilgram TK. Coronary artery calcium: Accuracy and reproducibility of measurements with multi-detector row CT-assessment of effects of different thresholds and quantification methods. Radiology 2003; 227:795–801.
16. Frangi AF, Niessen WJ, Vincken KL, et al. Multiscale vessel enhancement filtering. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI '98: First international Conference, Cambridge, MA, USA, October 1998. Proceedings. Berlin: Springer Verlag 1998; 1496:130-137.
17. Carr JJ. Coronary vascular calcium. RSNA Categorical Course in Diagnostic Radiology: Thoracic Imaging—Chest and Cardiac 2001; 47-55.

APPENDIX A

Feature 18 is calculated by the following algorithm:
find the central slice for the object
initialize threshold level *ThreshLevel* (220 HU)
perform 2D-region growing of candidate object
calculate *size* and *compactness* of the grown object;
bestCompactness = *compactness*
while ((*ThreshLevel* > *ThreshMin)* and (size < *SizeMax))*
ThreshLevel = *ThreshLevel* - *ThreshStep*
perform 2D-region-growing of candidate object
calculate *size* and *compactness* of the grown object
if (*compactness* < *bestCompactness*)
bestCompactness = *compactness*
where the predefined constants *SizeMax* represents the maximum size of the vessel and *ThreshMin*is the minimum HU values expected inside the vessel. We used *ThreshMin* = −100 HU, *SizeMax* = 1500 pixels and *ThreshStep* = 5 HU.

APPENDIX B

The following algorithm describes how the circle is fitted around the candidate object:
determine a rectangular region P_c of size $2r_{max} \times 2r_{max}$ where the center pixel is the center of mass of the candidate object;
for every point in P_c

initialize $F_{best} = 0$;
for every radius r = r_{min}, $r_{min}+1$, . . ., r_{max}
determine a circle with center P_c and radius r;
in the central slice I determine area of the object $A_{object}$and its area outside the circle A_{out};
if ($A_{out} > 0.5\ A_{object}$) reject the circle;
determine a set of equidistant points {P_1, P_2, . . .,P_n} along the circle;
for every point P_i from {P_1, P_2, . . .,P_n}
compute the image gradient F in P_i in the direction of vector V_i pointing from P_i to Pc

$$F = \frac{dI}{dV_i}(P_i)_i$$

if ($F < c_{min}$) reject the circle;
F_{avg} = avg{F};
if ($F_{avg} > F_{best}$)
$F_{best} = F_{avg}$;
$R_{best} = r$;
$Pc_{best} = P_c$;

F is a gradient computed in the point P_i from Gaussian derivatives with scale $\delta = 1$ pixel. We used r_{min}= 7 pixels, r_{max}= 19 pixels, $c_{min} = -10$.

The algorithm contains two rejection criteria. Criterion ($A_{out} > 0.5\ A_{object}$) ensures that the candidate object is inside the fitted circle. The requirement that the directional gradient in each point is above a certain minimum value was found to prevent the circle from fitting to a part of the spine.