

Multi-scale Nodule Detection in Chest Radiographs

Arnold M.R. Schilham, Bram van Ginneken, and Marco Loog

Image Sciences Institute, University Medical Center Utrecht, The Netherlands
{arnold,bram,marco}@isi.uu.nl, <http://www.isi.uu.nl>

Abstract. Early detection is the most promising way to enhance a patient's chance for survival of lung cancer. In this work, a novel computer algorithm for nodule detection in chest radiographs is presented that takes into account the wide size range for lung nodules through the use of multi-scale image processing techniques. The method consists of: *i*) Lung field segmentation with an Active Shape Model [1]; *ii*) Nodule candidate detection by Lindeberg's multi-scale blob detector [2] and quadratic classification; *iii*) Blob segmentation by multi-scale edge focusing; *iv*) k Nearest neighbor classification. Experiments on the complete JSRT database [3] show that by accepting on average 2 false positives per image, 50.6% of all nodules are detected. For 10 false positives, this increases to 69.5%.

1 Introduction

Lung cancer is the second most common cancer among both men and women. With a mortality of over 90% it is the leading cause for cancer deaths for both sexes [4]. Early detection of lung cancer (i.e. nodules) on a chest x-ray or computer tomography (CT) image, may increase the patient's chance of survival. Most chest exams are radiographs, in which it can be very hard to detect lung nodules. In a lung cancer screening program for heavy smokers it was found that 90% of peripheral lung cancers were visible on chest radiographs on retrospective examination [5]. Thus there is a large incentive to improve the detection of nodules in thorax x-ray images.

In this paper, a novel approach for a computer aided diagnosis (CAD) scheme for detecting pulmonary nodules in chest radiographs is presented. The algorithm isolates possible nodule locations in the image, which could then be offered to a radiologist for a final verdict.

By now several approaches for nodule detection have been published (see [6] and the references therein and e.g. [7,8,9]) and one commercial product has been approved by the FDA (Deus Technologies, Rockville, MD). It is of crucial importance to assess the performance and relative assets of the various methods. This can be done when researchers report results on common, publicly available databases. The JSRT database [3] is such a database. For the experiments in this paper we used the full JSRT database, without any selection of cases. This database contains 93 normal cases and 154 chest radiographs with a proven

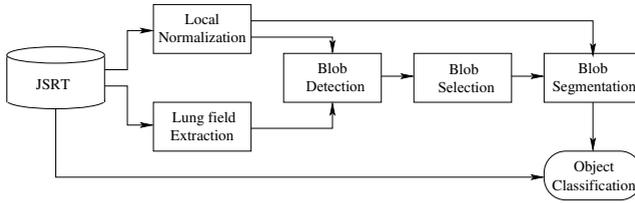


Fig. 1. The nodule detection scheme.

lung nodule (100 malignant ones). The diameters and positions of the nodules are provided, and the nodules are subdivided in five categories based on their degree of subtlety. We show results in the form of FROC curves [10] on the complete database and for each subtlety category separately.

In Section 2 we proceed step by step through the detection procedure. Section 3 presents the results obtained on the JSRT database. Finally, in Section 4 we discuss the usefulness of our results and various ways for further enhancement thereof.

2 Methods

Our nodule classification encompasses the following steps (see Fig. 1), which will be explained in more detail below:

1. A thorax image is selected from the JSRT database and subsampled to 1024 by 1024 pixels.
2. The lung field is extracted with an Active Shape Model (Section 2.1).
3. The contrast of the image is locally enhanced with a local normalization (LN) filter as a preprocessing step for the blob detector (Section 2.2).
4. Lindeberg's multi-scale blob detector is used to find blobs in the lung fields in the LN images (Section 2.3).
5. With a simple classification scheme based on features that result from the blob detector the number of nodule candidates is reduced (Section 2.4).
6. The regions of interest are segmented from the background by multi-scale edge focusing on rays cast radially through the detected blob centers (Section 2.5).
7. Finally, using the segmentation, a selected set of features, and a k nearest neighbor (k NN) classifier determine the probability that a candidate represents a nodule (Section 2.6).

2.1 Lung Field Segmentation

Detection of false positives outside the lungs is avoided by restricting the analysis to the lung fields only. The lung fields are segmented with an Active Shape Model (ASM) [1]. The ASM segmentation scheme requires manually segmented training

images; we used a set of 230 chest radiographs obtained from a tuberculosis screening program. The settings of the ASM scheme are those used in [11]. Note that lung fields in chest radiographs are defined—as usual—as those parts of the lung unobstructed by the heart, the mediastinum and the structures below the diaphragm. However, a significant part of the lungs is actually obscured by these structures. If a nodule is located in this part of the lungs, it will be missed by our detection system.

2.2 Preprocessing

The multi-scale blob detection (Section 2.3) is sensitive to the contrast between a nodule and its surroundings. It is likely to wrongly extract structures whenever there is strong contrast in the image, for example on ridges or near edges. To locally equalize the amount of contrast in the image L , we employ a local normalization (LN): $L_{LN} = (L - \tilde{L}) / \left(\tilde{L}^2 - (\tilde{L})^2 \right)^{1/2}$, where a tilde indicates Gaussian blurring. In Fig. 2(b) the effect of LN preprocessing can be seen. We used a scale $\sigma = 8$ pixels for the Gaussian blurring in LN. These preprocessing operations are used for the multi-scale blob detection and blob segmentation only; subsequent analysis is on the original images.

2.3 Multi-scale Blob Detection

Nodule detection is inherently a multi-scale problem; the nodules in the JSRT database range in diameter from 6 to 60 mm. Therefore, we employ a multi-scale detection scheme to find nodule candidates. A nodule is usually roughly spherical and has a density comparable to water, which is higher than the surrounding lung parenchyma. Consequently, nodules appear as bright, circular blobs in chest radiographs. To find nodule candidates, we used Lindeberg's multi-scale blob detection scheme, using the Laplacian $\Delta L \equiv L_{xx} + L_{yy}$ (with $L_{\alpha\beta}$ short for $\frac{\partial^2 L}{\partial \alpha \partial \beta}$) and gamma-normalized derivatives (see [2]). We subsampled the images to 256 by 256 pixels (i.e. 1.4 mm per pixel) and used a scale range from $\sigma = 1$ to 16 pixels. For each detected blob the position in the image and its corresponding scale (i.e. radius indicating the size of the blob) are stored. It is well possible that blobs overlap; in that case we keep only the strongest blob.

For the analysis in Table 1 of the performance of the blob detector, a nodule was considered to be detected if there was an overlap between the detected blob and the radiologists' circular delineation of the actual nodule. The total number of blobs detected for the whole JSRT database was 33073, i.e. 134 blobs on average per image.

2.4 Features for Blob Selection

The next step is aimed at removing a large part of the detected blobs to minimize further computational processing time, using a small set of 13 features. The first

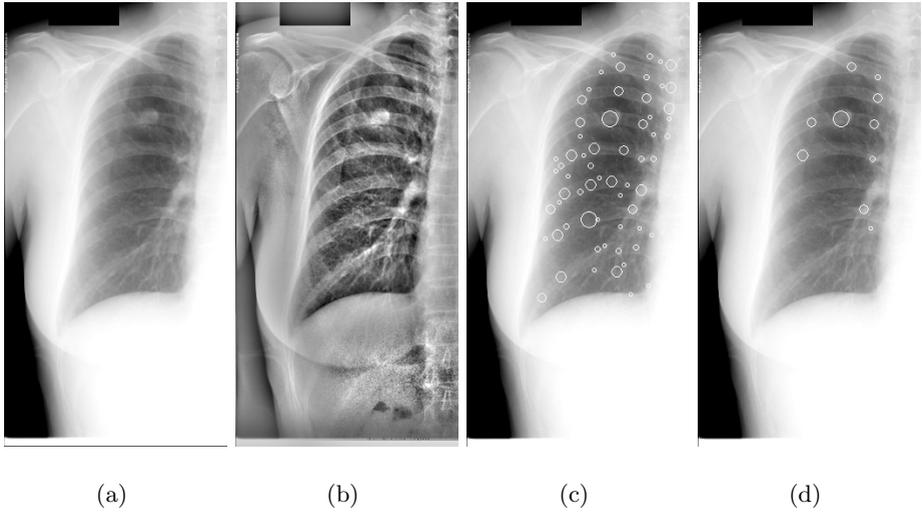


Fig. 2. The various processing steps of the blob detection shown for the right lung of JSRT case JPCLN006. The original image (a) has large contrast variations. After LN, the image shows locally equalized amounts of contrast (b). Many blobs are detected by the blob detector (c) and only a few are left after blob selection (d). For this case, the actual nodule (the largest blob in (d)) also has the highest probability for being a nodule after the final classification.

six features are derived directly from the Hessian matrix H (i.e. all (gamma-normalized) second order derivatives) and its largest and smallest eigenvector (λ_1 and λ_2 , respectively) at the location and scale σ of each detected blob: $L_{xx}, L_{xy}, L_{yy}, \det H = \lambda_1 \times \lambda_2, \lambda_1, \lambda_2$. We add the scale σ and two position features (x, y) , determined in a local coordinate system with the center of mass of the lung fields as the origin and the standard deviations in the x - and y -directions giving the unit lengths along the axes. Finally, four features are computed from the original image: the mean and standard deviation of the intensity values inside the blob and in a band of width σ around the blob.

A quadratic classifier with these 13 features was trained on all blobs. We put a threshold on the posterior probability to retain on average twenty candidates per image (5028 nodule candidates in total). See Fig. 2(d) for a typical result of the blob selection. The penalty for reducing the number of candidates by a factor 6.6 is the loss of 16% of the nodules (see Table 1).

2.5 Blob Segmentation

We developed a technique to improve the separation of the nodule from the surrounding background. The key ingredient of our segmentation scheme is scale-space edge focusing of rays cast through the detected blob centers, akin to the

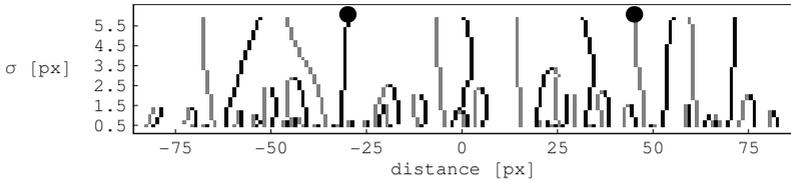


Fig. 3. The scale-space edge focusing technique. The dots indicate the strongest left (black) and right (grey) edges at the largest scale (σ). These are followed down to the smallest scale to find their corresponding locations in the original image. Positions (horizontal axis) are relative to the blob center as detected by the blob detector.

technique used in [12] for the segmentation of follicles in 3-D ultrasound. We start with constructing a Gaussian scale space of the data sampled along the cast rays. The extent of the lines is restricted to three times the diameter found by the blob detector. Since a blob is brighter than its surroundings, we search for edges of positive (negative) gradient on the left-hand (right-hand) side of the line (‘left’ and ‘right’ edges, respectively).

Generally we encounter several left and right edges (see Fig. 3), so we pick the strongest edges (we identify ‘strength’ here with the magnitude of the gradient at the given scale space location) to mark the boundaries of the blob. These edges are traced down in scale space to the smallest scale level to find the corresponding edge location in the image. In our application, the scale space spans a scale range $\sigma = 0.5\text{--}6$ pixels with increments of $1/6$ pixel. We cast 30 rays in a homogeneous orientation distribution, resulting in 60 boundary points per nodule candidate.

The images are preprocessed before ray casting with two filters. Firstly, LN is used to normalize the contrast level, such that contrast between bright (positive values) and dark spots (negative values) is comparable throughout the image. Secondly, we employ a filter operation to promote the edges between bright and dark areas over edges within bright or dark regions: $L' = -a_1 L_{LN}$ for $L_{LN} < 0$ and $L' = a_2 L_{LN}$ for $L_{LN} \geq 0$, with L_{LN} denoting the pixel intensities in the LN image, and a_1 and a_2 positive constants. We find satisfactory results for $a_2/a_1 \geq 10$ (typically a_1 is set to 1 and a_2 to 50).

After edge focusing, two post-processing steps are applied. First, the one-dimensional lists of radii is median filtered to remove the inevitable outliers. Next, the radii are allowed to grow back to the nearest edge (with the correct sign) if it lies within 10 pixels of the filtered location.

In Fig. 4 the different stages of the blob segmentation process are displayed for a typical nodule, showing the improvement in isolating the actual nodule.

2.6 Feature Extraction and Classification

In the final phase of the scheme, the probability that each candidate represents an actual nodule is estimated. To this end, a set of features is computed for each candidate, and the k nearest neighbors in feature space among all nodule candidates from different cases in the database are computed. The posterior

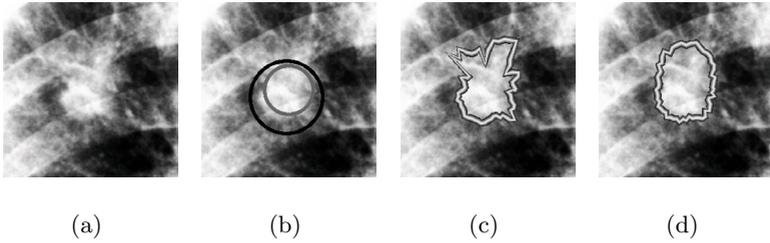


Fig. 4. After histogram equalization, this as ‘fairly obvious’ categorized nodule (case JPCLN014) is clear to see (a). In the JSRT database, the nodule delineation is the black circle, whilst the blob detector finds the grey circle (b). The nodule segmented by the ray casting procedure is shown in (c). After post-processing (median filtering and back growing) the final segmentation of this nodule is given in (d).

probability is n/k where n is the number of actual nodules among the k neighbors. Best results were obtained by using $k = 100$.

Two position features were used, as defined above, and the mean and standard deviation within the segmented area and in a band around the segmented area – defined by doubling the radii that define the segmentation – were computed for a number of filtered versions of the original images. As filters all Gaussian derivatives from 0th to 2nd order for 4 different scales ($\sigma = 1, 2, 4, 8$ pixels) were used. This resulted in a total of 98 features. From these 98 features, a subset of 13 features was selected using sequential forward selection in a leave-one-out manner [13].

3 Results

The performance of the different steps of the nodule detection scheme are summarized in Table 1. Note that 13 nodules out of 154 are obstructed by the heart, mediastinum or structures below the diaphragm. These nodules are outside the analyzed region of the images, and are therefore not detected. Another 25 nodules are missed by the multi-scale blob detection. There might be room for improvement here, but most of these missed nodules are ‘extremely subtle’ or ‘very subtle’ cases which are barely detectable for human observers as well [3].

Table 1 also shows that some objects are lost in the segmentation phase. For these objects, the edge-focusing technique could not find edges of the correct type for all ray orientations. Apparently, these are not blob-like structures completely located within three detected diameters of the found centers. These objects are notably non-circular, none of them were nodules.

The FROC curves [10] of our detection algorithm for the five nodule categories as defined in [3] are shown in Figure 5. Indeed we see that the names of the 5 categories are reflected by the performance of our algorithm. By accepting 2 nodule candidates per image, we already have an overall probability of 51%

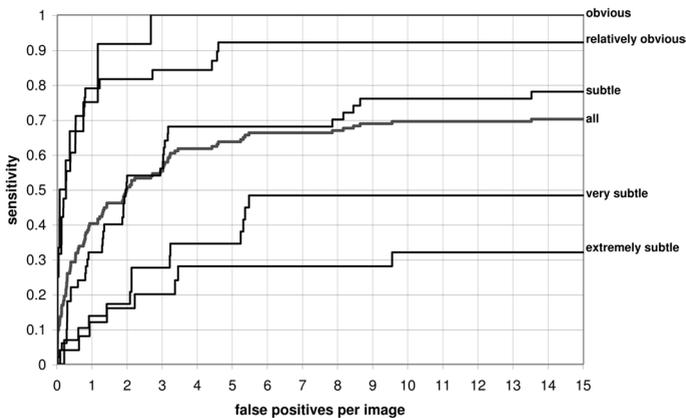


Fig. 5. FROC curves of the system on the complete JSRT database, and per category.

to have located the nodule (if there is one in the image) and by accepting 10 candidates per image this chance increases to 69%. If we ignore the very subtle and extremely subtle nodules, these chances are respectively 69% and 85%.

4 Discussion and Conclusions

We have shown that multi-scale blob detection and blob segmentation by scale-space edge focusing can be used successfully for the detection of pulmonary nodules in chest radiographs. Note that the nodule segmentation methodology proposed here might also prove useful in determining whether a given nodule is benign or malignant [14]. The observer studies in [3] indicated that radiologists are very likely to miss the extremely subtle nodules in the JSRT database (area under the ROC curve $A_z = 0.568$). With only 4 false positives per image, our scheme marks almost 30% of these cases, which is an encouraging result, suggesting that our method could provide a useful clinical tool. However, this remains to be proven in observer studies.

For the object classification, we have used straightforward feature-based classification methods, which may leave room for improvement. For example, it could be beneficial to add shape descriptors to the set of features. Alternatively, the object classification could be based on a statistical model of the nodules and their surrounding structures.

Table 1. Actual nodules and candidates retained after the various steps in the detection scheme.

	Number of candidates	Number of nodules
JSRT Database	–	154 (100%)
Nodules inside the lung fields	–	141 (91.5%)
Blob Detection	33073	133 (86.4%)
Blob Selection	5028	108 (70.1%)
Blob Segmentation	4900	108 (70.1%)

References

1. T.F. Cootes, C.J. Taylor, D. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision and Image Understanding*, 61(1): 38–59, 1995.
2. T. Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30(2): 79–116, 1998.
3. J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *American Journal of Roentgenology*, 174: 71–74, 2000.
4. A. Jemal, A. Thomas, T. Murray, and M. Thun. Cancer statistics 2002. *CA Cancer J Clin*, 52: 23–47, 2002.
5. J.R. Muhm, W.E. Miller, R.S. Fontana, D.R. Sanderson, and M.A. Uhlenhopp. Lung cancer detected during a screening program using four-month chest radiographs. *Radiology*, 148: 609–615, 1983.
6. B. van Ginneken, B.M. ter Haar Romeny, and M.A. Viergever. Computer-aided diagnosis in chest radiography: a survey. *IEEE Transactions on Medical Imaging*, 20(12): 1228–1241, 2001.
7. Q. Li, S. Katsuragawa, and K. Doi. Computer-aided diagnostic scheme for lung nodule detection in digital chest radiographs by use of a multiple-template matching technique. *Medical Physics*, 28: 2070–2076, 2001.
8. B. Keserci and H. Yoshida. Computerized detection of pulmonary nodules in chest radiographs based on morphological features and wavelet snake model. *Medical Image Analysis*, 6: 431–447, 2002.
9. J. Wei, Y. Hagihara, A. Shimizu, and H. Kobatake. Optimal image feature set for detecting lung nodules on chest x-ray images. In *Computer Assisted Radiology and Surgery (CARS 2002)*, pages 706–711, Berlin, 2002. Springer.
10. P.C. Bunch, J.F. Hamilton, G.K. Sanderson, and A.H. Simmons. A free response approach to the measurement and characterization of radiographic-observer performance. *Journal of Applied Photographic Engineering*, 4: 166–171, 1978.
11. B. van Ginneken, S. Katsuragawa, B.M. ter Haar Romeny, K. Doi, and M.A. Viergever. Automatic detection of abnormalities in chest radiographs using local texture analysis. *IEEE Transactions on Medical Imaging*, 21(2): 139–149, 2002.
12. B.M. ter Haar Romeny, B. Titulaer, S. Kalitzin, G. Scheffer, F. Broekmans, E. te Velde, and J. J. Staal. Computer assisted human follicle analysis for fertility prospects with 3D ultrasound. In *IPMI '99*, volume 1613 of *Lecture Notes in Computer Science*, pages 56–69. Springer-Verlag, Heidelberg, 1999.
13. R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. John Wiley & Sons, New York, 2nd edition, 2001.
14. M. Aoyama, Q. Li, S. Katsuragawa, H. MacMahon, and K. Doi. Automated computerized scheme for distinction between benign and malignant solitary pulmonary nodules on chest images. *Medical Physics*, 29: 701–708, 2002.