Lung Nodule Detection in CT Scans

M. Antonelli, G. Frosini, B. Lazzerini, F. Marcelloni

Abstract—In this paper we describe a computer-aided diagnosis (CAD) system for automated detection of pulmonary nodules in computed-tomography (CT) images. After extracting the pulmonary parenchyma using a combination of image processing techniques, a region growing method is applied to detect nodules based on 3D geometric features.

We applied the CAD system to CT scans collected in a screening program for lung cancer detection. Each scan consists of a sequence of about 300 slices stored in DICOM (Digital Imaging and COmmunications in Medicine) format. All malignant nodules were detected and a low false-positive detection rate was achieved.

Keywords—computer assisted diagnosis, medical image segmentation, shape recognition.

I. INTRODUCTION

LUNG cancer is the leading cause of death in the world mainly because patients with lung cancer often do not develop symptoms until the cancer is in an advanced stage [1]. Periodic screening tests with CT of populations at risk for lung cancer could dramatically reduce the death rate [2, 3]. However, a major problem in the screening for lung cancer with CT is the fact that the objects to recognize in the image slices, such as nodules, blood vessels and bronchi, often have similar characteristics in terms of density, shape and size. To help solve this problem we have developed a computer-aided diagnosis (CAD) system for automated detection and diagnosis of pulmonary nodules on CT images. First the system extracts the lung regions, then it detects and classifies the nodules.

We used a database consisting of eight CT scans, each with about 300 slices. A slice is a 512×512 pixels matrix with slice thickness of 1.25 mm. All the malignant nodules were correctly detected and recognized, even in the case in which nodules were next to blood vessels or the pleura. Further, we found a very low false-positive number. This is remarkable if we consider that false positives just origin from the presence of a few s-shaped blood vessels not distinguishable from nodules attached to the pleura.

II. EXTRACTION OF PULMONARY PARENCHYMA

In order to extract the pulmonary parenchyma, we perform the following actions. First we remove the background (i.e., the pixels with the same grey level as the lungs but located outside the chest) from the image. To this aim, due to the high similarity between the grey levels of the lungs and the image background (see Fig. 1), we cannot simply apply a thresholding technique. Instead we use an ad-hoc operator which, starting from the four corners of the image, moves along the four directions identifying as background pixels those pixels whose grey level is within a pre-fixed range.

Fig. 1 A DICOM image.

When the operator finds a pixel value outside the range (*out pixel*), it analyses a few more pixels in the scan direction. If such pixels have values within the range, both these pixels and the out pixel are marked as background pixels and the scan goes on. Otherwise, the scan is interrupted along the direction under examination and the successive row/column is analysed in a similar way. The image produced by this operator (shown in Fig. 2) is converted to a binary image by means of a thresholding technique that uses either a static or a dynamic threshold depending on the lung zone the slice belongs to.

Fig. 2 The image after background removal.

More precisely, after discarding the first few initial and last

slices of a CT scan that do not contain lung images, we classify the remaining slices into three groups corresponding, respectively, to the upper, middle and lower parts of the lung volume. In the first and last groups the lungs image represents a smaller percentage of the slice than in the second group. We therefore apply a different thresholding technique depending on the group the given slice belongs to: we use a dynamic threshold, determined through an iterative procedure [4], for the slices of the lung middle part and a threshold determined empirically for all other slices. Thanks to the thresholding operator each slice is transformed into a binary image where the foreground, i.e., the object we are interested in, consists of the lungs and all the rest is background (see Fig. 3).
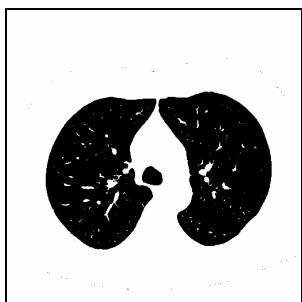

Fig. 3 The image after thresholding.

Afterwards, we apply the morphological opening and closing operators so as to improve the image and border definition; this allows us, in particular, to enhance the separation between distinct regions and to fill the gaps in the borders. The resulting image is shown in Fig. 4.
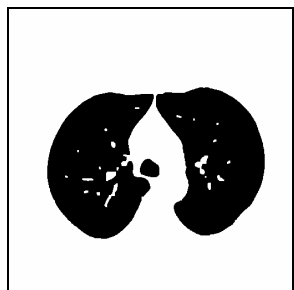

Fig. 4 The image after the opening and closing operators.

Then we detect the image borders through a tracking algorithm that uses the Sobel operator [5], and reduce the border size to one pixel using the Wang and Zhang thinning algorithm [6]. In this algorithm, a pixel is considered a border pixel if at least one of its neighbours is white (non-border). The algorithm starts building all the border pixels chains (a border chain is a set of connected border pixels), and then for each chain eliminates the excess pixels (i.e., those pixels whose elimination does not influence the border connectness), thus producing an image with one-pixel contours (see Fig. 5).
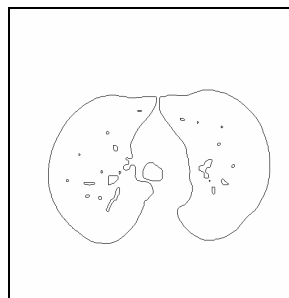

Fig. 5 The image with thinned borders.

The following step is the identification of the border pixel chains that represent the two pulmonary lobes. This step takes both the chain size and the chain location within the slices into account. More precisely, in a slice of either the middle or the lower part of the lung volume the two longest border chains are chosen as pulmonary lobes. On the other hand, for a slice in the upper part of the lung volume, the size of the two lung lobes is comparable with the trachea size: we therefore choose the two longest chains that do not include border pixels with coordinates inside the range in which we expect to find the trachea. Fig. 6 shows a slice with two chains representing the pulmonary lobes.
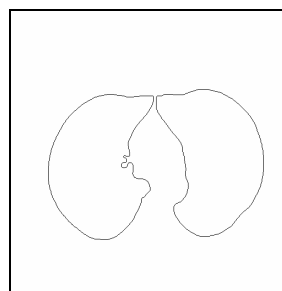

Fig. 6 Two chains representing the pulmonary lobes.

Then we reconstruct the pulmonary lobes borders so as to reinsert the nodules adjacent to the pleura (*juxta-pleura* nodules) previously suppressed by the thresholding operator.

In order to rebuild the pulmonary lobes border parts that have been erroneously eliminated, for each pair of border points, we calculate two distances: the Euclidean distance d_1 and the minimum distance in pixels d_2 between the two points. We then compute the ratio d_2/d_1: if it is greater than a given threshold the two points are candidates for a possible reconstruction. For all candidate pairs a further condition is evaluated before establishing whether the reconstruction has to be done: we verify if the segment connecting the two points is not internal to the lung. If the condition is satisfied, then the border between the pair of points is rebuilt, i.e., the two points are connected by a line (see Fig 7).
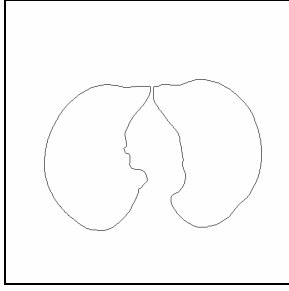
Fig. 7 Reconstructed pulmonary lobes.

Finally, we apply a region filling operator to the pulmonary lobes chains in order to reintroduce the original values of grey levels inside the lungs. The resulting images (see Fig. 8), which contain only the lung regions, are input to the nodule detector.
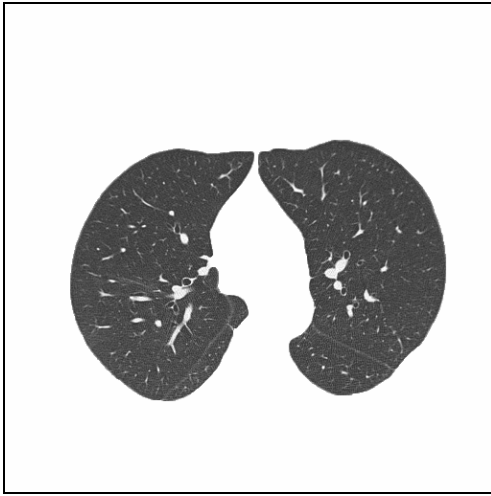

Fig. 8 Restored image.

III. NODULE DETECTION

Lung nodule detection is a very difficult step in a CAD system development. Actually, in CT lung images, nodules are frequently attached to blood vessels or to the pleura; further their grey tone is so similar to vessels sections that traditional intensity-based methods are inappropriate. Instead, an effective nodule detection algorithm must take both the grey level and the object shape into account. In our CAD system we adopt a method that uses 3D shape information to identify spherical regions with a given grey level.

Following the approach described in [7], the idea is to distinguish spherical (typically nodules) from cylindrical (typically blood vessels) shapes analyzing a shape index (SI), defined in terms of 3D characteristics, extracted from sets of voxels with grey level in the range of the nodule intensity.

More precisely, let $v(x,y,z)$ denote an image voxel (i.e., a 3D point) and $I(x,y,z)$ its intensity. For each voxel we define an infinitely differentiable 3D function $f(x,y,z)$ as the convolution of $I(x,y,z)$ with a smoothing filter $g(x,y,z)$: $f(x,y,z) = I(x,y,z) \oplus g(x,y,z)$. In this case we use a 3D gaussian function as smoothing filter.

To calculate the shape index directly from the isotropic volume we compute the first and second-order partial derivatives utilizing the convolution operator property:

$$f_x = I(x,y,z) \oplus g_x ;\ f_y = I(x,y,z) \oplus g_y ;\ f_z = I(x,y,z) \oplus g_z ;$$

$$f_{xx} = I(x,y,z) \oplus g_{xx} ;\ f_{xy} = I(x,y,z) \oplus g_{xy} ; \text{etc.}$$

An iso-intensity surface can be defined as $S(x,y) = \{(x,y) \in \Re^2 ; I(x,y,\phi(x,y)) = a\}$, where $z = \phi(x,y)$ and a is a constant. If we refer to the smoothed image, we have $S(x,y) = \{(x,y) \in \Re^2 ; f(x,y,\phi(x,y)) = a\}$.

The first and second fundamental forms of differential geometry are:

$$E = S_x \bullet S_x,\ F = S_x \bullet S_y,\ G = S_y \bullet S_y$$

$$L = S_{xx} \bullet Q,\ M = S_{xy} \bullet Q,\ N = S_{yy} \bullet Q$$

where $\bullet$ denotes scalar product, $Q = \dfrac{S_x \times S_y}{\|S_x \times S_y\|}$, $\times$ is the outer product and $\|\|$ is a norm.

Being $\dfrac{\partial \phi}{\partial x} = -\dfrac{\partial f}{\partial x} \Big/ \dfrac{\partial f}{\partial \phi} = -\dfrac{f_x}{f_z}$, we obtain:

$$S_x = \left(1,0,\dfrac{\partial \phi}{\partial x}\right) = \left(1,0,-\dfrac{f_x}{f_z}\right).$$

Therefore it results:

$$E = 1 + \dfrac{f_x^2}{f_z^2},\ F = \dfrac{f_x f_y}{f_z^2},\ G = 1 + \dfrac{f_y^2}{f_z^2}$$

$$L = \dfrac{2 f_x f_z f_{xz} - f_x^2 f_{zz} - f_z^2 f_{xx}}{A^{1/2} f_z^3},$$

$$M = \dfrac{f_x f_z f_{yz} + f_y f_z f_{xz} - f_x f_y f_{zz} - f_z^2 f_{xy}}{A^{1/2} f_z^3},$$

$$N = \dfrac{2 f_y f_z f_{yz} - f_y^2 f_{zz} - f_z^2 f_{yy}}{A^{1/2} f_z^3},$$

with $A = \dfrac{f_x^2 + f_y^2 + f_z^2}{f_z^2}$.

We define the principal curvatures at each voxel v as follows:

$$K_1(v) = H(v) + \sqrt{H^2(v) - K(v)}$$

$$K_2(v) = H(v) - \sqrt{H^2(v) - K(v)}$$

with K and H equal to the Gaussian and the mean curvatures:

$$K = \dfrac{LN - M^2}{EG - F^2},\ H = \dfrac{EN - 2FM + GL}{2(EG - F^2)}.$$

Finally the shape index at each voxel v is defined by:

$$SI(v) = \frac{1}{2} - \frac{1}{\pi} \arctan\left(\frac{K_1(v) + K_2(v)}{K_1(v) - K_2(v)}\right).$$

This index distinguishes spherical from cylindrical shapes: values close to 1 indicate spherical shape, instead values close to 0.75 indicate cylindrical shape. In practice, we considered values in the range [0.8,1] as clue of nodules and values in the range [0.5,0.8) as clue of blood vessels. Testing the method on synthetic images we found that there are some cases in which the SI value is undefined owing to either a particular parameter combination that zeros the denominator in some fraction, or a strong symmetry along the z-axis that zeros the first derivative f_z. To solve this problem we exploited adjusted formulas obtained by axes rotation, thanks to 3D images isotropy.

The gaussian filter dimension was chosen based on a trade-off between the computational time and the desired accuracy: a small filter is unable to recognize different shapes because of the limited observation perspective. In our experiments we found that a value equal to 15×15 is a good trade-off.

We applied a region growing method that utilizes both the grey level and the SI. Starting from a seed point (voxel) which is within the specified ranges with respect to both the grey level and the SI, we expand the region towards the 26 immediate neighbouring voxels. The growing process is interrupted along a direction when an "out of range" voxel condition is found, i.e., the new voxel has either the grey tone or the SI outside the specified ranges. When all the edges of the object under consideration are intercepted, if for all (actually, most) voxels we find SI < 0.8, the object is recognized as a blood vessel, otherwise it is a nodule. For each detected nodule, an estimate of its size is also provided for diagnostic purposes.

The method was validated on both synthetic images and real 3D lung CTs. In the first case it was able to recognize spherical and cylindrical shapes, independently of their position (e.g., sphere attached to cylinder, sphere attached to sphere).

In the real lung CT images the method recognized lung nodules even if they were next to blood vessels or the pleura, and gave a small false-positive number. This happened in correspondence to s-shaped blood vessels not distinguishable from nodules attached to the pleura.

IV. CONCLUSION

We have described a CAD system for automated segmentation and reconstruction of pulmonary parenchyma, and nodule detection. The results achieved by applying the system to a database consisting of eight CT scans with a total of more than 2400 digital images have been judged definitely good by experienced chest radiologists. In particular, as regards nodule detection, all malignant nodules were detected and a very low false-positive detection rate was achieved.

REFERENCES

[1] S. Sone, S. Takashima, F. Li, Z. Yang, T. Honda, Y. Maruyama, M. Hasegawa, T. Yamanda, K. Kubo, K. Hanamura, K. Asakura, "Mass screening for lung cancer with mobile spiral computed tomography scanner", *The Lancet*, vol. 351, April 1998, pp. 1242-1245.

[2] T. Nawa, T. Nakagawa, S. Kusano, Y. Kawasaki, Y. Sugawara, H. Nakata, "Lung cancer screening using low-dose spiral CT: results of baseline and 1-year follow-up studies", *Chest*, vol. 122, n.v1, July 2002, pp. 15-20.

[3] M.J.R. Dalrymple-Hay, N.E. Drury, "Screening for lung cancer", *Journal of the Royal Society of Medicine*, vol. 94, January 2001, pp. 2-5.

[4] S. Hu, E.A. Hoffman, J.M. Reinhardt, "Automatic lung segmentation for accurate quantitation of volumetric x-ray CT images", *IEEE Trans. Medical Imaging*, vol. 20, n. 6, June 2001, pp. 490-498.

[5] K.R. Castleman, "Digital Image Processing", Prentice Hall, Englewood Cliffs, NJ, USA (1996).

[6] P.S.P. Wang, Y.Y. Zhang, "A fast and flexible thinning algorithm", *IEEE Trans. Computers*, vol. 38, n. 5, May 1989, pp. 741-745.

[7] J. Dehmeshki, X. Ye, J. Costello, "Shape based region growing using derivatives of 3D medical images: Application to semi-automated detection of pulmonary modules", *International Conference on Image Processing*, vol. 1, 2003, pp. 1085-1088.