Elimination of clavicle shadows to help automatic lung nodule detection on chest radiographs

G. Simkó[1], G. Orbán[1], P. Máday[1], G. Horváth[1]

[11] Dept. of Measurement and Information Systems,
Budapest University of Technology and Economics, Budapest, Hungary

Abstract— **Lung nodule detection is one of the most important goals of large scale screening of chest radiographs. The success of nodule detection can be increased if it is possible to suppress the bony structures from the chest radiographs. While one possible way to do this is to use dual-energy imaging, most of the commercial X-ray machines do not offer this technology. Finding an alternative approach is an important task. This paper proposes a new solution for it, where first the contours of bones –especially the contours of clavicles– are detected then the shadow of the bones is removed. The performance of the clavicle suppression algorithm is then evaluated by the use of a nodule detection system. The algorithms were tested on images from the JSRT database. The nodule detection algorithm was looking for up to 20 suspicious areas in the original images. According to our first results in the images with suppressed clavicle shadows 20 of 380 false positives were removed, while all true positives were preserved. 69 of the original 380 false positives were around the clavicle, which would mean 24% decrease in the concerned region.**

Keywords— **medical image segmentation, clavicle detection, bone shadow elimination, nodule detection.**

I. Introduction

Lung nodule detection is one of the most important goals of radiograph-based large scale chest screening. Its success can be significantly increased if it is possible to suppress the bony structures – the rib system and the clavicles – from the chest radiographs. One possible way to suppress bony structure is to use dual-energy imaging [1]. As the majority of commercial X-ray machines do not offer this technology, finding alternative approaches still remains an important task. This paper proposes a novel solution to get a "bone-free" image. First the contours of bones – especially the clavicles – are detected than the shadows of the bones are removed. The second goal of the paper is to show how bone shadow elimination can improve the performance of nodule detection algorithms.

The tests were evaluated on such images, where the nodules are under or in the neighbourhood of the clavicles. The main aim of this test was to see whether clavicle shadow removal will change true positive-false positive hit ratios.

II. Clavicle detection

Complete clavicle detection is probably an impossible task with the exclusive use of normal chest radiograph images. The medial epiphysis of clavicle produces an almost invisible, very dim shadow. This is the reason why we decided to try detecting only the diaphysis of the bone instead of detecting the whole bone.

Our clavicle detection is performed in three steps: preprocessing of the image, approximate curve fitting and pixel precise detection. The first step of the clavicle detection is preprocessing of the X-ray images: the image is smoothed by a Gaussian filter, then the edges are enhanced by multiplication with the gradient magnitudes. At the same time, we suppress incorrect gradient directions: we create four different gradient images according to the four clavicle edges.

After preprocessing we try to estimate the position of the clavicles. As the contours of clavicles have very little curvature, we try to approximate them with straight lines. For this purpose we apply Radon transformation as it is a fast algorithm and works well even on noisy images.

The local maxima of the transformed image will specify lines through edges of the original image. However as ribs have similar characteristics as clavicles have, special care must be taken to find the latter. In most situations the clavicles have sharper edges than ribs and the clavicle edges are located in the upper part of the image.

We developed an iterative method that relies on these facts: the algorithm searches for edge pairs that belong to the highest Radon maxima and are in the most upper corner of the lung. We also filter the results with a simple criterion: only non-crossing edges are accepted if they are in an appropriate distance from each other. After finding the approximate contour lines, we try to determine the exact shape of the bones. Active contour modeling fits this task very well [2], that solves the task by the minimization of an energy function representing the criteria mentioned before.

Consider the following equations:

$$c(s) = [x(s), y(s)] \qquad s \in [0,1]$$

$$E = E_{\text{int}} + E_{\text{ext}} = \int_0^1 \frac{1}{2}(\alpha \dot{c}^2 + \beta \ddot{c}^2) + E_{\text{ext}} ds \tag{1}$$

where $c(s)$ defines the 2D parametric curve – representing the clavicle contour –, α and β are weight parameters, $\dot{c}(s)$ and $\ddot{c}(s)$ are the tension and rigidity of the curve and $E_{\text{ext}} = -\nabla I(x,y)^2$ is determined from the $\nabla I(x,y)$ gradient image. The aim is to find $c(s)$ that minimizes E.

The approximate line determines the initial $c(s)$ and because clavicles have a nearly horizontal position, we may simplify the problem: we regard $x(s)$ to be fixed, and $y(s)$ is the only variable in which E is minimized. Moreover clavicles tend to be concave upward. Considering these constraints we constructed a new energy function:

$$E = E_{\text{int}} + E_{\text{ext}} = \int_0^1 \frac{1}{2}\beta(\dot{c}^2 \cdot penalty(\ddot{c}(s), s)) + E_{\text{ext}}\, ds \qquad (2)$$

Here $\beta(s)$ is the same weight parameter as before, it may depend on s, and $penalty(\ddot{c}(s), s)$ is a new weight parameter controlling the concavity of the curve:

$$penalty(\ddot{c}(s), s) = \begin{cases} 1, & \text{if} \quad \ddot{c}(s) \le 0 \\ \text{concavity\_penalty(s)}, & \text{else} \end{cases} \qquad (3)$$

So penalty is 1, if $c(s)$ is concave downward, else it punishes upward concavity. By using dynamic programming we can easily solve these equations [3].

III. Bone suppression

In the previous section we showed an algorithm for detecting clavicle edges. Using this information we suppress the shadow of the bones from the chest radiographs. We present a general algorithm for this task, which works well not only for clavicles but also for ribs. Clavicle suppression is tested using the contours from the previous step, while rib suppression is tested with manually marked contours.

The algorithm works in gradient image space. In the end the resulting images need to be transformed back into the normal image space. This means that we need an invertible gradient operator: the central difference operator in our case. Conceptually the main steps of the algorithm are:

1. Calculating gradient image.
2. Building statistical bone model.
3. Recalculation: image from modified gradient.

The essence of the algorithm is the creation of a statistical edge model, it does not use any a-priori information or machine learning. The gradients are evaluated along the bone edge and the model is built from the edge surroundings.

Based on the parametric curve equation (1), we may define its upper and lower neighborhood as:

$$J(s,d) = [x(s), y(s) + d] \qquad (4)$$

where d means the vertical shift from the bone curve (it may be negative, which means a shift toward the inner part of the bone), the other parameters are the same as in (1).

Applying an appropriate smoothing kernel along J we get the edge model: provided the density of bone does not change too fast, the smoothed gradient will average to the bone model. The reason behind this property is that considering an area of gradients, there will be the gradient from bones and fluctuations due to noise and lung tissues. Modeling the fluctuations with Gaussian noise, they cancel out each other, if the area is large enough. Hence the average gradient calculated on an adequate-size area is the gradient resulted from the underlying bones. Of course if the area is too large, there will be several bones and massive tissues degrading the results. This effect is maximally reduced by averaging only in the neighborhood of the observed bone.

A 2D Gaussian kernel with $\sigma_x = 45$ and $\sigma_y = 0.75$ is used. This means a Gaussian filter with an effective window size of approximately 269×7 pixels. It is very important to notice that this size means that an average nodule (1mm-20mm, 7-140 pixel in a 2000×2000-pixel image) does not affect the bone model, hence it remains on the output image.

After constructing the bone model, we subtract it from the original image, hence removing the shadow of the bone.

IV. Nodule detection

The major task of lung CAD systems is the localization of cancerous regions on chest radiographs. The suspicious regions on the input images are marked, decreasing the load on the diagnosing specialist, who is able to pay more attention to these areas throughout the examination. However, it is also quite important to avoid marking regions unnecessary, in order to avoid confusing results.

By suppressing the shadows of clavicles, the characteristics of underlying tissues become more observable.

The algorithm improves the recognition rate in two different ways. First of all, the amount of true positives is increased. By suppressing the disturbing bony structure, the image features calculated in the clavicle region better characterize the tissues underneath. The less biased features calculated can increase the performance of the classifier.

Additionally the algorithm helps reducing the number of false positives. On some images the shadow of the overlapping parts of the clavicles and the ribs, may seem similar to that of the nodules and the prefilter may mark these parts as suspicious regions. In the worst case the classifier may fail to recognize this situation that results in a false positive marking. By considering the known anatomical properties, it is possible to remove the misleading signs of the clavicle region that prevents this scenario.

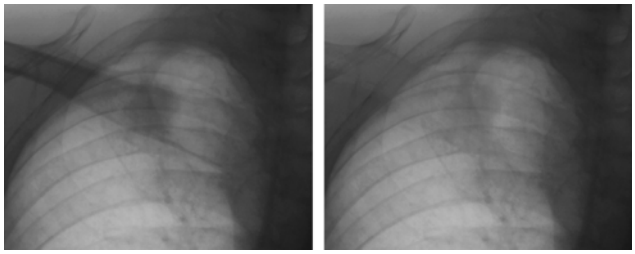

Fig. 1 Clavicle suppression. Notice the nodule in the middle.

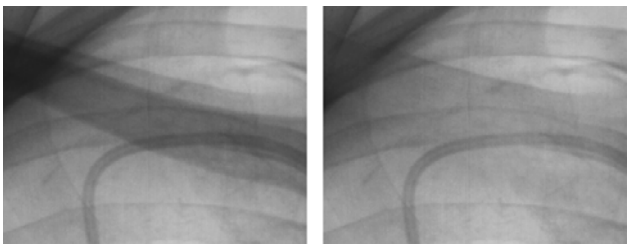

Fig. 2 The cable of pacemaker is clearly visible after suppression.

The nodule detection system contains three major steps. [4] The first is the segmentation of the lung area. The surrounding parts have characteristics different from the inner lung areas and the classifier trained for the inner part may provide unsatisfactory results here. In this paper the lung contour is regarded to be given.

The aim of the second step is the selection of candidate regions, which may contain cancerous areas. As the classifier algorithm in the third step is responsible for the selection of the true positive areas, the aim of the second step is to mark all the suspicious areas.

As cancerous regions on radiographs appear as rounded regions with smooth boundary, the main requirements for the prefilter is to find each object on the image with these properties. Additional requirements need to be considered however, as the nodule sizes may vary significantly, and the cancerous regions are often really hard to observe due to their weak contrast to the background.

One approach to fulfill these requirements is to apply a filter from the Convergence Index (CI) family, such as Sliding Band Filter [5]. CI filters apply a criterion related to the directional properties of the image gradients, and are able to operate contrast insensitively. Sliding Band filters additionally are able to recognize nodules with a wide variety of sizes due to the adaptive selection of the region boundaries applied. The candidate regions are selected by thresholding.

The third step of the process is the application of a classifier. The classification of candidate regions is based on the selected image features describing the properties of the region of interest (ROI) by characterizing it in different scales. The local properties are described by textural features, based mainly on statistical measures, while the geometrical features, like the characteristics of the contour of the ROI, represent the image properties of larger scale. The selection of the appropriate set of features was done by a feature relevance determination method described in [6].

The selected classifier is a Support Vector Machine (SVM), as SVMs are well known for their ability to deal with high dimensional input data, and for their robustness.

The classifier is trained on samples from the Japanese Society of Radiological Technology (JSRT) database, as it serves as a standard benchmark dataset. Separate experiments tested clavicle detection and suppression, and the improvement on nodule detection. We used two sets of chest radiographs for test purposes: the standard JSRT database and 108 additional images.

After viewing the ground truth data linked with JSRT database, we found that it contains some imprecise clavicle data, so its use would produce false evaluation. Moreover the medial joint of clavicle is in the most situations invisible, so it is easy to see that using “radiologist’s marking” exact ground truth information can not be acquired.

During clavicle detection tests we evaluated the algorithm on 108 images by visual inspection using the help of a pulmonologist. We made 3 classes of results: pixel-precise hit, false hit and partly-precise hit (on some regions it is pixel-precise, elsewhere it makes several pixel errors).As it can be seen from Table I. the results are quite impressive.

Table 1 clavicle detection on 108 images – 216 bones

Pixel Precise	False	Partly-precise
175	11	30
81%	5%	14%

Concerning clavicle suppression visual inspection remained the only possible way to evaluate the results in the absence of dual-energy ground truth data. The results are promising, but later exact evaluation is needed.

V. Results

To evaluate the performance of our integrated clavicle removal–nodule detection system a subset of the JSRT database was selected and the results were compared to ground truth results. We selected those 20 images from the database, where the nodules are around a clavicle. We also ran test on the first 120 images, to be able to make more accurate statistics on false positive rate.

We inspected the results both at the end of the candidate selection state and on the output of the classifier. The latter can show performance gain in real life applications, but the former can illustrate the effect better. The threshold of the candidate selector was set to select the first 20 candidates

for the original images. Fig. 3 shows an example image after the candidate selection state.

Concerning the hit rate after ROI selection we obtained minor improvements. The real nodules were always among the 20 candidates on each of the inspected 20 images. In 9 cases the performance improved, while in 6 cases we found the results to be slightly worse then in the original case.

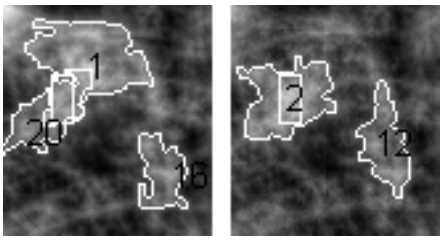

Fig. 3 Candidates marked and ordered on the CI output image. The nodule is marked with a white rectangle: (left) is the original result, (right) is the result after clavicle removal.

The latter 20 and the other set of 120 images were used to examine the change in false positive rate. On the smaller set compared to the original results, 25 false positives disappeared, while 11 appeared in the concerned region. Those new false positives were located mostly at the inner end of the removed clavicle. Because we prefer reducing the number of false positives over increasing recognition rate, we decided to take the intersection of the original candidates and the ones we get after clavicle removal. With this method, we removed a total of 20 false positives which means on average 1 less false positive per image. Concerning the fact that on these images 69 false positives were around the clavicle, the system improved by 29% in the region. The large dataset showed even better results. We managed to eliminate 211 false positives.

Supposing that on other regions of the image, the bones are responsible for false candidates in the same rate, we could expect this 29% increase on the whole image when removing the ribs. We tried this on one image only for illustration. The system produced 16 false positives. Compared to the original amount of 20, it means a 20% performance increase. When automatic rib detection becomes available, an extensive analysis should be made.

On the output of the classifier, we examined the 20 images with nodules near the clavicle, and the set of 120 images. On the former set we found a performance increase both in recognition rate (Fig. 4 left), and in false positive rate. With various SVM parameters, a 6% to 20% increase in recognition rate was observed at identical false positive rates. From another point of view, the false positive rate can be decreased by 16% to 33% while maintaining recognition rate. Also the smallest available false positive rate decreased by 0.25, which means a 6% improvement while recognition rate increased simultaneously. A 3% increase in maximum available recognition rate with less false positives was also observed. On the larger image set, the results were close to our expectations (Fig. 4 right). The false positive rate decreased by 15% to 18% while maintaining recognition rate, and the smallest available false positive rate decreased by 16%. In this case, recognition rate did not change noticeably. For example the highest value increased by 1.5%.

We can notice that on the classifier output the improvements are sometimes better, than we would expect from the result after the candidate selection state. This is probably due to better textural features of nodules originally behind the shadow of the clavicle. We made some tests using the candidates from the clavicle free image, but calculating features on the original images. The results were somewhat worse, proving that clavicle removal will not corrupt the texture of nodules but rather make it more visible.

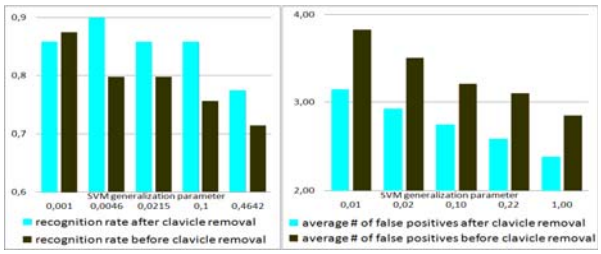

Fig. 4 Classifier performance. Recognition rate on the 20-image set (left). Number of false positives on the 120-image set (right).

VI. Conclusions

In this paper we showed that bone shadows can be eliminated from chest radiographs without using dual energy technology, and based on the clavicle-less images the performance of lung nodule detection can be improved: while keeping the true positive hit rate the number of false positives can be reduced significantly.

References

1 Lehmann et al. (1981). Generalized image combinations in dual KVP digital radiography. Medical Physics 8 (5), pp. 659–667.

2 Kass, M., Witkin, A., Terzopoulos, D. (1987), Snakes: Active Contour Models. International Journal of Computer Vision, 1. (4), pp. 321-331.

3 Amini, A. A., Tehrani, S., Weymouth, T. E. (1988), Using dynamic programming for minimizing the energy of active contours in the presence of hard constraints, Int. Conf. Comput. Vision., II: pp. 95-99.

4 Campadelli, P., .Casiraghi, E., Artioli, D. (2006) A Fully Automated Method for Lung Nodule Detection From Postero-Anterior Chest Radiographs. IEEE Trans. Medical Imaging 25 (12) pp. 1588-1603

5 Wei; Hagihara, Y.; Kobatake, H. (1999). Detection of cancerous tumors on chest X-ray images–candidate detection filter and its evaluation. Proc. of the International Conf. on Image Processing, ICIP 99.

6 Carlos S. Pereira et al. (2007) Evaluation of Contrast Enhancement Filters for Lung Nodule Detection, ICIAR07, pp. 878-888.