Recognizing Emphysema - A Neural Network Approach

Ola Friman, Magnus Borga, Mikael Lundberg, Ulf Tylén[†], Hans Knutsson

Dept. of Biomedical Engineering, Linköping University and [†]Dept. of Radiology, Göteborg University, Sweden

Abstract

An accurate and fully automatic method for detecting and quantifying emphysema in CT-images is presented. The method is based on an image preprocessing step followed by a neural network classifier trained to separate true emphysema from artifacts. The proposed approach is shown to be superior to an established method when applied on real patient data.

1. Introduction

Emphysema is a common chronic respiratory disorder characterized by destruction of lung tissue, leading to reduced oxygen extraction and hence limitations in physical activities. The main cause of emphysema is tobacco smoking, although there may be other factors that may contribute to the development of the disease. In terms of rehabilitation and health care costs, emphysema is among the top five diseases in the western world today. From these perspectives, it is of vital importance to develop methods for diagnozing emphysema, both for clinical and research use.

Computed Tomography (CT) has emerged as an important tool for early assessment of emphysema. CT imaging produces high resolution images where pixel intensity is interpreted as tissue density. Figure 1 shows a CT image of the chest. The lungs contain mostly air (low density) and are thus seen as the two darker objects within the body. Intermixed with the lungs are bright blood vessels with higher tissue density. Since emphysema implies lack of tissue, density is lower in damaged regions in the lung. This is reflected in the CT image as areas with lower pixel intensities. Hence, emphysema is seen as darker regions within the (already dark) lungs. The objective is to quantify the level of damage by estimating the area of destroyed tissue in the CT images. A visual evaluation systematically overestimate the percentage of damaged lung area. In addition, reproducibility is poor both between examiners and between assessments made by the same examiner at different occasions [2]. Thus, there is a need for an objective and accurate method for detecting and quantifying emphysema. Since emphysematous areas mainly are characterized by lower intensity in the images, a commonly applied method is simply to count pixels surviving a suitably selected intensity threshold [4, 1]. Thresholding CT images is facilitated by the use of the normalized Hounsfield scale where 0 Hounsfield units corresponds to the density of water while -1000 corresponds to air. On account of this scale, a fixed threshold can be determined and used for all images. Even though such a method of course yields reproducible results the accuracy is unsatisfactory, particularly when the degree of emphysema is low. This is due to noise and a number of artifacts also present in the CT images, which deteriorate the possibility of finding emphysema by thresholding alone. The approach pursued here combines image processing and neural networks into a detection method which provides the desired accuracy and reproducibility. The proposed method can be divided into three separate steps; an image preprocessing step where potential emphysematous regions are detected, a feature extraction step and finally a classification step which separates true emphysema from image artifacts. The procedure is applicable on both 2D images and 3D volumes.

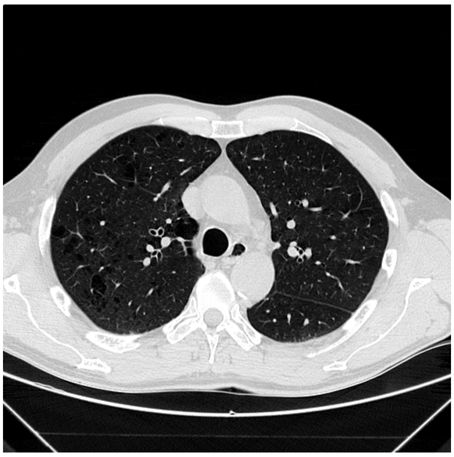

Figure 1. A CT image of a chest.

2. Image preprocessing

The aim of the image preprocessing step is to extract a set of potential emphysematous regions in the CT image. The preprocessing essentially consists of four parts; automatic segmentation of the lungs, an intensity correction, a moderate smoothing and finally a thresholding. Due to the limited space available, these steps are only briefly described.

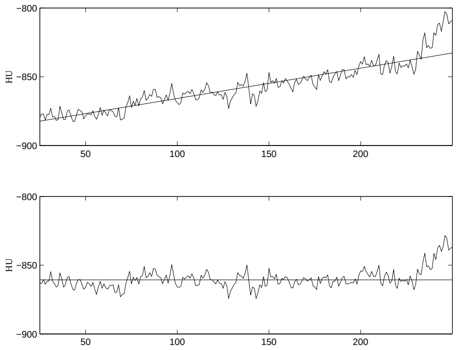

Figure 2. Illustration of the tissue density gradient and its correction. The horizontal axis indicates vertical position in the lung.

In the segmentation step the lungs in the CT image are located and extracted. This is a quite straightforward operation since there are a number of features characterizing the lungs, such as pixel intensity, size and position in the image. Figure 3a shows a segmented lung.

Next, a correction for an intensity nonuniformity over the lungs must be applied. Usually, the CT images are acquired while the patient is lying on the back in the scanner. The small amount of fluid normally present within the lungs then becomes unevenly distributed with higher concentrations in the lower parts. Thus, tissue density becomes slightly higher in the lower parts of the image compared to the upper parts. Figure 2 shows the mean intensity of each row in a typical image. As is seen, the density varies linearly with position in the lungs. If not corrected for, this nonuniformity will of course introduce problems when subjecting the image to a threshold. Inevitably, more spurious regions will be detected in the upper somewhat darker parts of the lungs. Therefore, a line is fit to the observations in Fig. 2 and the pixel intensities are levelled out over the lungs.

The next preprocessing step is to smooth the image slightly. The rationale for smoothing the image is to avoid spurious errors caused by noise and to gather low intensity regions into homogeneous areas. A plain smoothing will however blur the edges of the segmented lungs and the interpretation of pixel values according to the Hounsfield scale will no longer be valid. Therefore, a filtering technique termed *Normalized Convolution*, [5], is used instead. Normalized Convolution is the preferred filtering method when there are uncertain or missing data present. Indeed, pixels outside the segmented lungs are not interesting in this case and can be seen as missing. Denote by I the original CT image and by f a Gaussian filter kernel. An ordinary smoothing is then obtained by $I_{smoothed} = I * f$, where $*$ is a convolution, see Fig. 3b. To perform Normalized Convolution a so-called *certainty* c is also introduced, reflecting how certain we are that each pixel value in the image I is correct. In the present context, all pixel values inside the segmented lungs are considered known and are assigned certainty 1, while pixel values outside the lungs are considered unknown and are therefore assigned certainty 0 (i.e. c is here binary, but this need not be the case in general). A smoothed image is subsequently obtained by

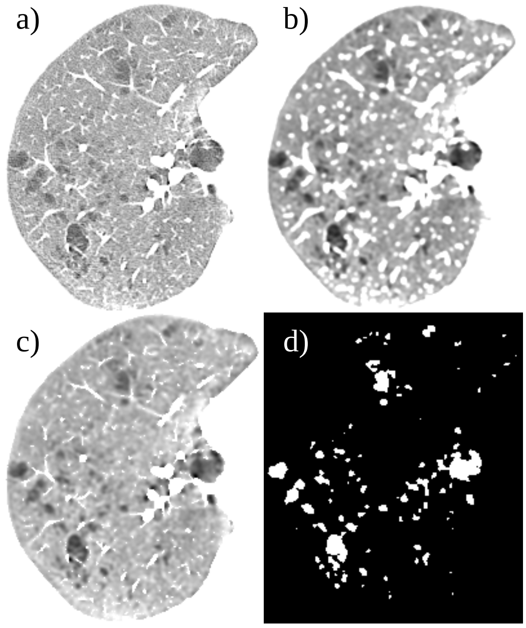

Figure 3. a) Original segmented lung. b) Smoothed lung. c) Smoothing obtained by Normalized Convolution. d) After threshold operation.

$$I_{smoothed} = \frac{c\,I * f}{c * f}, \qquad (1)$$

where multiplications and divisions are taken pointwise, see Fig. 3c. By using Normalized Convolution the Hounsfield scale is preserved also at the edges of the segmented lungs. Hence it is still possible to subject the image to a fixed threshold in order to find potential emphysematous regions, and this is the final step in the image preprocessing, see Fig. 3d.

3. Feature extraction

Each region detected in the preceding image preprocessing step is considered as belonging to one of two classes; truly emphysematous or artifactual. It is possible to define a set of heuristic rules according to which the detected regions can be dichotomized into these two classes. For example, an oblong region originating from a vessel is a typical artifactual pattern and should thus be removed. However, a more general approach is to present a number of manually classified regions to a neural network, which by supervised learning can derive a better set of “rules” to recognize emphysematous regions.

To this end, a number of features are extracted from each detected region. These features should capture the characteristics of a candidate emphysematous area which a human examiner implicitly observe when classifying a region as being either truly damaged or just a spurious artifact. Fundamental for the classification of a separate region are features based on statistical measurements of the pixel intensity values inside the region and the shape of region. For example, a truly damaged region is generally larger and more regular than an artifactual region. However, it is not likely that it is possible to make a classification based on features extracted solely from the region itself, information from the local surrounding is also needed. Therefore, for each candidate region, two areas are defined, see Fig. 4. First the inner region which is the original potentially damaged area. The second region is a 10 pixel wide border area from which information about the surroundings can be extracted. The two

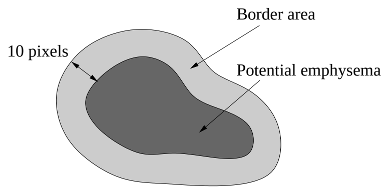

Figure 4. Illustration of a single candidate region and the 10 pixel wide surrounding border.

leftmost columns in Table 1 list the features extracted from each candidate region. The intensity features are calculated

Features		
Intensity	Shape	Global
Mean	Area	Nbr of detected regions
Standard dev.	Eccentrity	Mean intensity of r.
Min	Orientation	Mean area of r.
Max	Compactness	Mean orientation of r.
Contrast		Mean intensity of lung

Table 1. List of features.

for both the inner region and the border region, except for the contrast feature which is a measure of the intensity difference between the two regions. The shape features are calculated for the inner region only. Eccentrity is defined as the ratio between the principal axes of an ellipse fitted around the region and hence measures roundness. Compactness is the ratio between the number of pixels in the region and the number of pixels in the convex hull of the region.

For a human examiner, the general "impression" of the CT image probably influences also the classification of a separate region in the image. Hence, in addition to the local features described above, it may be of importance to extract a set of global features containing information about the overall situation, i.e. features which put the local region under consideration into a context. The global features used are listed to the right in Table 1.

An example of a situation which cannot be resolved by local features only is a patient motion artifact which manifests itself as dark regions oriented in the same direction and which originate from the blood vessels in the image. From a global view this artifact is easily spotted due to the characteristic orientation pattern but it is impossible to detect on a local region level. The mean orientation of the detected regions is therefore included as an indicator of this artifact.

4. Training & Classification

To extract data for training a classifier, 500 CT images from 50 patients were processed as described in the Image preprocessing section. For each image, a set of potentially emphysematous regions were presented and these regions were manually classified as being either true emphysema or as artifacts. This procedure resulted in about 5,000 emphysematous and 50,000 artifactual regions. Even though most artifactual regions are small, these figures indicate the severe problem of artifacts in the detection of emphysema. For each region, the features described in the previous section were extracted. 4,500 examples of each class were used for training a classifier to recognize emphysematous and artifactual regions respectively, and 500 examples from each class were used for evaluating the classification performance.

A number of different classifiers and neural networks were trained and evaluated on this data set. Reported here are the results using a Fisher Linear Discriminant classifier, an ordinary error backpropagation neural network and a Support Vector Machine approach with Gaussian kernels.

5. Results & Discussion

The number of correctly classified test data examples using the evaluated classifiers are reported in Table 2. The nonlinear classifiers achieve somewhat better performance compared to the linear Fisher Discriminant. An error backpropagation network with eight hidden neurons and one output neuron was found to give best classification, 89.4 % correctly classified test examples. An evaluation of the

Fisher Linear Discriminant	85.9 %
Error Backpropagation	89.4 %
Support Vector Machines	88.8 %

Table 2. Detection performance for the evaluated classifiers.

network revealed that the intensity features were the most important followed by the global features, while the shape features (somewhat surprisingly) did not provide much discrimination support. It should be stressed that the training data set was produced by a single examiner and hence suffers from inconsistencies as pointed out in the introduction. The proper way of constructing the training data would be to merge data independently classified by several examiners and remove the discrepancies between these sets. The interference in the feature space between true emphysema and artifacts would then become significantly smaller. Still, all tested classifiers managed well to separate between emphysematous and artifactual regions and are therefore able to weed out most of the artifacts. This is illustrated in Fig. 5 where detected regions before and after the classifier step are shown for two different cases.

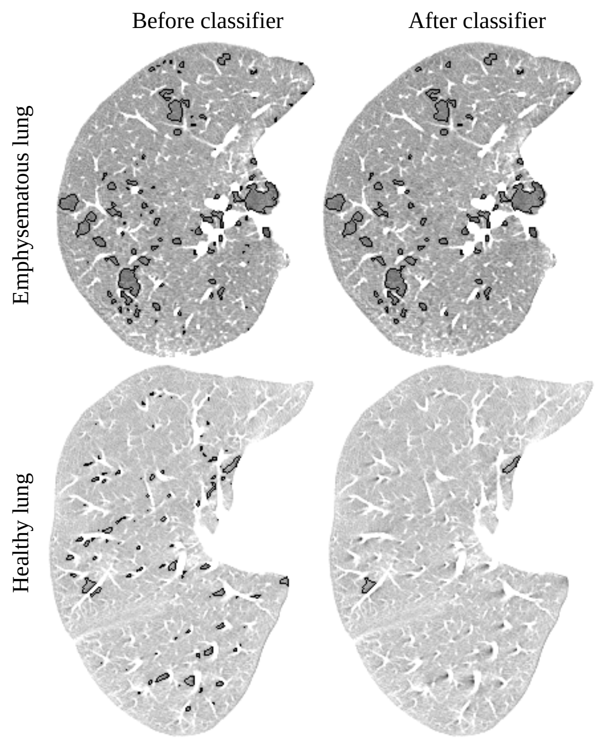

Figure 5. Regions detected before and after the classifier step in an emphysematous lung and in a normal lung with some artifacts.

As the primary goal is to estimate the area of destroyed lung tissue, a set of 20 images of healthy and damaged lungs were used for evaluating the new method in this respect. The commonly used Density Mask method [4], which employs a direct thresholding of the original CT image as described in the Introduction, was also applied on these images. The result is reported in Fig. 6. The new method is able to identify healthy and diseased lungs with good accuracy while the Density Mask method performs very poorly. Again, these results emphasize the importance of recognizing the artifacts when detecting emphysema in CT images. Elaborate extensions of the density mask method, e.g. [3], often fails in this respect. In a work by Uppaluri et al. [6], a Bayesian classifier is used for recognizing several different lung disorders, among them emphysema. However, their approach partitions the lung into 31×31 blocks which subsequently are classified, which immediately implies a loss in accuracy. Moreover, since Uppaluri et al. aim at detecting several different patterns, their preprocessing steps are not tailored for recognizing emphysema. This is important

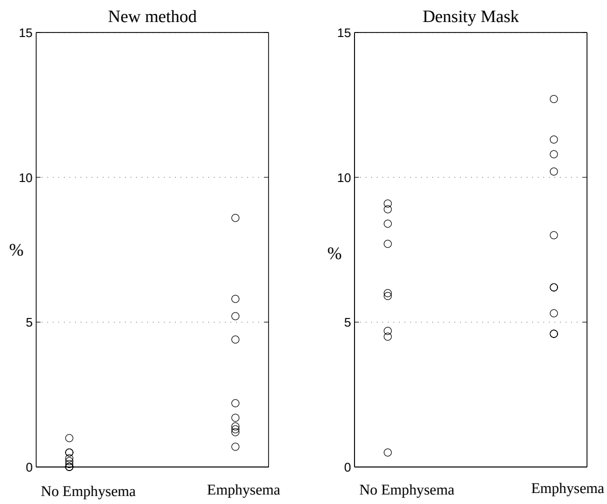

Figure 6. Estimated percentage of damaged lung area in 20 different images using the proposed new method and the traditional Density Mask method.

since the preprocessing step is crucial in all types of learning and neural network problems.

To conclude this paper, the combination of image processing and neural networks has been shown to yield an accurate method for detecting emphysema. As such, the method can facilitate and support physicians in visual screening of the images and also increase knowledge of the development of the disease.

References

[1] A. Bankier. Why and how should emphysema be quantified on CT? 19th annual meeting of the Society of Thoracic Radiology, Thoracic Imaging 2001.

[2] A. Bankier, V. Maertelaer, C. Keyzer, and P. Gevenois. Pulmonary emphysema: Subjective visual grading versus objective quantification with macroscopic morphometry and thinsection CT densitometry. *Radiology*, 211:851–858, 1999.

[3] R. Blechschmidt, R. Werthscützky, and U. Lörcher. Automated CT image evaluation of lung: A morphology-based concept. *IEEE Transactions On Medical Imaging*, 20(5):434–442, 2001.

[4] H. Coxson, R. Rogers, K. Whittall, Y. D'Yachkova, P. Pare, F. Sciurba, and J. Hogg. A quantification of the lung surface area in emphysema using computed tomography. *Am. J. Respir. Crit. Care Med.*, 159:851–856, 1999.

[5] H. Knutsson and C.-F. Westin. Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 515–523, June 1993.

[6] R. Uppaluri, E. Hoffman, M. Sonka, P. Hartley, G. Hunninghake, and G. McLennan. Computer recognition of regional lung disease patterns. *Am. J. Respir. Crit. Care Med.*, 160:648–654, 1999.