

Computer-Aided Detection of Pulmonary Pathology in Pediatric Chest Radiographs

André Mouton¹, Richard D. Pitcher², and Tania S. Douglas¹

¹ MRC/UCT Medical Imaging Research Unit, Department of Human Biology,
University of Cape Town, South Africa

² Department of Radiology, Tygerberg Hospital, Cape Town, South Africa
tania@ieee.org

Abstract. A scheme for triaging pulmonary abnormalities in pediatric chest radiographs for specialist interpretation would be useful in resource-poor settings, especially those with a high tuberculosis burden. We assess computer-aided detection of pulmonary pathology in pediatric digital chest X-ray images. The method comprises four phases suggested in the literature: lung field segmentation, lung field subdivision, feature extraction and classification. The output of the system is a probability map for each image, giving an indication of the degree of abnormality of every region in the lung fields; the maps may be used as a visual tool for identifying those cases that need further attention. The system is evaluated on a set of anterior-posterior chest images obtained using a linear slot-scanning digital X-ray machine. The classification results produced an area under the ROC of 0.782, averaged over all regions.

Keywords: classification, computer-aided diagnosis, tuberculosis, pulmonary abnormality, chest radiograph.

1 Introduction

Chest radiography remains the most common radiological examination, in both adults and children, and in both well-resourced and poorly resourced settings; poorly resourced settings, however have limited recourse to specialist radiological services. The chest X-ray is often the main diagnostic tool for tuberculosis (TB) in children in endemic settings [1]. The manner in which PTB manifests itself radiographically varies with the age of the patient as well as the stage of the TB infection (primary or post-primary). Children generally develop primary TB and display associated radiological features.

Chest X-rays are extremely difficult to interpret. Poor contrast and complex backgrounds comprised of superimposed anatomical structures overlapping with regions of interest make the detection of abnormalities difficult, even for experienced radiologists [2]. Thus a diagnostic tool which assists in triaging pediatric chest radiographs for specialist interpretation would be useful.

Although much research has been conducted into the use of computer-aided textural analysis for the detection of interstitial lung diseases in chest radiographs, research into

computer-aided methods for the detection and diagnosis of general pulmonary pathology in chest radiographs, especially in children, is limited. Explicit application to a pediatric population, has, to the authors' knowledge, not been reported. Pulmonary abnormality would include interstitial disease (of the supporting structures of the lung including the alveolar wall) and alveolar disease, which involves the alveolar airspace itself and manifests by diffuse, homogeneous opacification. We apply methods developed for the detection of adult TB and interstitial disease, to the detection of pulmonary abnormality in pediatric chest radiographs; we assess the utility of this framework with reference to the results obtained by others in adults.

2 Methods

Schemes for the detection of interstitial lung disease (ILD) have a common structure. Multiple regions of interest (ROIs) are manually or automatically selected within the lung fields. Features, in the form of textural measurements, are then extracted from each of these ROIs. The ROIs are classified using rule-based and/or pattern recognition techniques resulting in soft labels (posterior probabilities) for each of the regions in the lung fields. The results are combined to obtain an overall diagnosis for the image, indicating whether or not it contains any interstitial abnormalities. The approach of this study is guided by that presented in [3].

2.1 Study Data

The data set consisted of 119 anterior-posterior (AP) pediatric chest scans of children, obtained using a linear-slot-scanning digital X-ray machine [4]. The images were obtained as part of a study to evaluate their suitability for visualizing airway abnormalities in children having symptoms of TB. The children were aged between 0 and 5 years at the time the scans were taken. Although the pulmonary changes associated with TB are completely non-specific when viewed in isolation, chest x-rays are an important component of the diagnostic process for TB, especially in children.

To allow for the training of classifiers to distinguish between normal and abnormal regions, the abnormal regions in the images were outlined by the examining radiologist on a computer monitor using a mouse pointer. Of the 119 images, 6 contained no pathology while the remaining 113 images contained, in total, 263 abnormal regions, of which 168 occurred in the right lung field and 95 in the left lung field.

2.2 Lung Field Segmentation

Automatic segmentation of lung fields is considered a required procedure before computerized analysis of a chest radiograph can occur [5].

In this study the lung fields are segmented using the active shape model (ASM) technique [6], which consists of three elements: a global shape component, a multi-resolution grey-level appearance component and a multi-resolution search algorithm. The technique requires a set of annotated training images, in which a set of corresponding points (representing the lung borders) has been marked, from which the model can be built. Statistics about the model's global shape and appearance, in the vicinity of each model point, are captured. This statistical model is used in an

iterative search to locate the lung fields in unseen test images. A semi-automatic ASM initialization is used here, where the user is required to provide a rough starting point for the model search. This is done by clicking a mouse on three landmark points in the lung field to be segmented. The landmarks required are: the bottom lateral corner of the lung; the lung apex; and the bottom medial corner of the lung.

2.3 Lung Field Subdivision

Arguably the greatest challenge in textural analysis of the lungs is the presence of superimposed anatomical structures [5]. Dividing the lung fields into small ROIs allows for each ROI to be classified individually with a classifier trained with features extracted from the specific region only, thereby capturing common anatomical characteristics of particular lung regions in the feature vectors, reducing the effect of background structures [3]. The mean lung shapes are computed from ASM derived from the training data. The boundaries defining the regions are found using a simple search algorithm, with equal area as the objective. The region maps are illustrated in Figure 1; 42 regions of different sizes overlap. These regions are used to automatically position the ROIs in subsequent test images, with segmented lung fields, which are warped to these region maps to define the ROIs. This warping is done using radial basis functions with multiquadrics and interpolation [7].

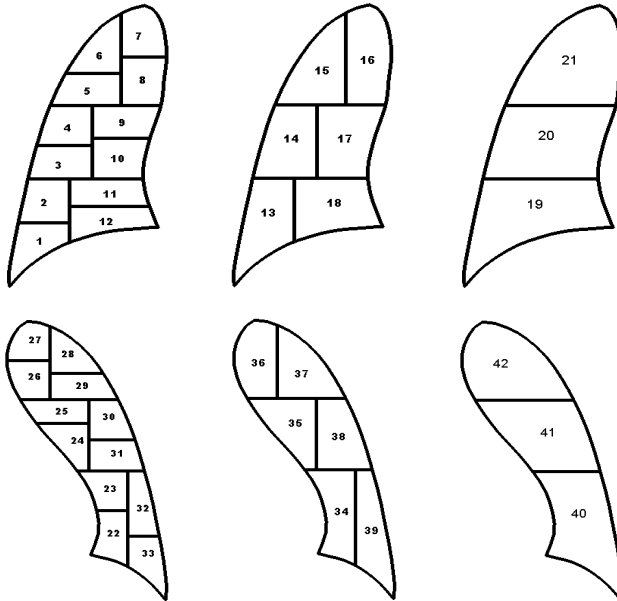


Fig.1. Subdivision of the lung fields into 42 overlapping regions of interest

2.4 Feature Extraction and Classification

Prior to the extraction of image features, several pre-processing steps are completed [8]: the left lung field is flipped about the centre line of the image so that the left lung

will resemble the right lung in orientation; pixel values outside the lungs are mirrored symmetrically with respect to the lung borders to prevent distortions in filter outputs at the lung borders.

A multiscale bank of Gaussian filters is used to extract the texture features [3,9]. The filter bank consists of Gaussian derivatives of orders, 0, 1 and 2 at five scales, $\sigma = 1, 2, 4, 8, 16$. By varying the degree of the filter (i.e. using multiple orders) more of the local image structure can be captured.

For each input image, 30 filtered versions of the image are computed, where the derivative of the image, $L(x)$, computed at an inner scale of σ is denoted by: $L_{n \times ny}$, with n_x and n_y denoting the order of the derivatives in the x and y directions. For each of the 5 scales the following 6 filtered images are computed: $L_{00}, L_{10}, L_{01}, L_{20}, L_{02}, L_{11}$.

The first four central moments (mean, standard deviation, skew and kurtosis) of the pixel intensity distribution from each ROI in the filtered and original images are computed as textural features [3,8].

Training separate classifiers for each of the 42 ROIs, thereby allowing them to be examined independently of one another, incorporates anatomical variation into the description of every region [3]. This is a particularly attractive approach in the present study, where the number of images containing no lung pathology is small; the number of normal regions is, however, much larger.

A kNN classifier is used. A leave-one-out method is used to construct the training and testing sets. The output of the classification process is a soft label for each feature vector in the range $[0, 1]$ where 0 = normal and 1 = abnormal. This allows receiver operating characteristic (ROC) curve analysis.

The per region classification results are used to create probability maps for the chest images where every pixel in the lung fields is replaced by the average of the soft labels of all the regions in which it falls - giving a posterior probability of that pixel being normal or abnormal. These maps are to be used as a visual tool by attracting the attention of a clinician reading the image to a region of abnormality and then allowing the clinician to make a decision regarding further interpretation of the image.

2.5 Performance Evaluation

The student's t -test is used to compare the classification results for the true (manual) and ASM segmentations; such a comparison gives an indication of the performance of the segmentation algorithm.

ROC analysis is performed independently on the 42 regions, resulting in a separate value for the area under the ROC curve (A_z) for each of the regions.

3 Results

The left and right lung fields were segmented independently of one another. The active shape models were trained using 25 images for the right lung field and 20 images for the left lung field.

Using an 11-NN classifier, the final classification results (in terms of area, A_z , under the ROC curve) of the system gave a per region score of $A_z = 0.782$ averaged over the 42 regions. No statistically significant difference was found between classifications on regions obtained by manual and ASM segmentation.

Table 1 shows a comparison between the results (in terms of per region Az scores) presented in this study and those presented in [3], where a similar scheme was used for the detection of TB in a screening database and interstitial disease in images obtained from daily clinical practice. Our database is similar to the TB database in [3], in that it contains a range of abnormalities encountered in chest screening, and we expect their results on the TB database to be a suitable basis for comparison with ours. The abnormalities identified in our images were primarily alveolar in origin, apart from a small group with nodular opacities, which could be either interstitial or alveolar.

Table 1. Comparison of Az values (across regions) with those of a similar study on adults [3]

	Mean	Min	Max
TB [3]	0.676	0.540	0.830
ILD [3]	0.835	0.660	0.930
Current study	0.782	0.655	0.937

The objective of a probability map is to indicate the degree of abnormality of every region in the lung field. The overlays on the lungs in Figures 2 and 3 define this degree of abnormality - a very white region indicates extreme abnormality, while a greyish-white region indicates that some subtle abnormalities could be present.

Figure 2 illustrates successful performance of the system, while Figure 3 illustrates common failures. The most frequent errors occurred in the basal and perihilar regions of the lungs, especially those of the left lung.




Fig 2. Successful probability maps: original image (left) and probability map (right). The healthy lung tissue is accurately depicted by darker grey values. Outlines of abnormal regions, made by a radiologist, are shown; the segmented lung borders are shown.




Fig 3. Common failures of the system. Top row: the probability map for a normal image indicates a greater degree of abnormality in the right perihilar region. Bottom: the image shows a diffuse nodular pattern in both lungs, which is only depicted in the right lung in the probability map.

4 Discussion

Frequent errors occurred in the basal and perihilar regions of the lungs, especially those of the left lung, and may be attributed largely to the sparseness of the data (in terms of number of abnormalities) in those regions of the lungs and the prominence of the vasculature in the lung fields.

Broncho-vascular markings (both normal and increased markings), proved problematic for the system. In many images, where the radiologist indicated increased broncho-vascular markings, the resulting probability maps displayed limited abnormality; while in many images where the radiologist indicated healthy vasculature the resulting probability maps showed abnormality in the regions containing vasculature. This is perhaps an unavoidable shortcoming of the classification system, which, in its current form, defines abnormality based solely on the textural characteristics of the lung tissue, while unhealthy broncho-vascular markings cannot always be distinguished from healthy counterparts using textural appearance alone. Experienced radiologists are able to take into consideration a large variety of factors when making a diagnosis, many of which are not quantifiable and therefore difficult to program. The probability map in Figure 3 (bottom) indicates diffuse abnormality in the right lung field but not in the left lung field for an input image containing a very obvious diffuse nodular pattern in both lung fields. This failure can be attributed to the fact that only three images containing this pattern in both lungs exist in the data set, while

several more exist displaying the pattern in the right lung only. The system is thus very poorly trained on such abnormalities for the left lung and since individual classifiers are trained for every region the results are poor.

The results presented in this study for pulmonary abnormality in children show an improvement over those presented in [3] for adult TB. Since similar methods are used in the two studies, the performance differences must be attributed to the data set. A few scenarios are possible. The abnormalities in this data set could be more obvious and therefore more easily detectable than those in the data set used in [3] – this is supported by the fact that the performance of the system presented in [3] on the interstitial lung disease data set is considerably higher. We used direct digital images, while the images in [3] were digitized from film. The digital X-ray machine we used may produce images which present textural pathologies in the lung fields more clearly. Testing on a larger and more diverse database and on image sets from different imaging systems may provide more clarity on system performance.

Acknowledgments. We thank Dr Rupesh Daya for assistance with data collection, and Dr Bram van Ginneken for guidance on algorithm implementation.

References

1. Marais, B.J., Pai, M.: New approaches and emerging technologies in the diagnosis of childhood tuberculosis. *Paediatric Respiratory Reviews* 8, 124–133 (2007)
2. Katsuragawa, S., Doi, K.: Computer-aided diagnosis in chest radiography. *Computerized Medical Imaging and Graphics* 31, 212–223 (2007)
3. van Ginneken, B., Katsuragawa, S., Romeny, B.M.T., Doi, K., Viergever, M.A.: Automatic detection of abnormalities in chest radiographs using local texture analysis. *IEEE Transactions on Medical Imaging* 21, 139–149 (2002)
4. Pitcher, R.D., Wilde, J.C.H., Douglas, T.S., van As, A.B.: The use of the Statscan digital X-ray unit in paediatric polytrauma. *Pediatric Radiology* 39, 433–437 (2009)
5. van Ginneken, B., Romeny, B.M.T.H., Viergever, M.A.: Computer-aided diagnosis in chest radiography: A survey. *IEEE Transactions on Medical Imaging* 20, 1228–1241 (2001)
6. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. *Computer Vision and Image Understanding* 61, 38–59 (1995)
7. Ruprecht, D., Muller, H.: Image warping with scattered data interpolation. *IEEE Computer Graphics and Applications* 15, 37–43 (1995)
8. Arzhaeva, Y., Tax, D.M.J., van Ginneken, B.: Dissimilarity-based classification in the absence of local ground truth: Application to the diagnostic interpretation of chest radiographs. *Pattern Recognition* 42, 1768–1776 (2009)
9. van Ginneken, B., Romeny, B.M.T.: Multi-scale texture classification from generalized locally orderless images. *Pattern Recognition* 36, 899–911 (2003)