TB detection using modified Local Binary Pattern features

Joshua Leibstein[#], Andre Nel[*]

[#]*School of Electrical Engineering, University of Johannesburg*

[1]joshua.leibstein@gmail.com

[*]*School of Mechanical and Industrial Engineering, University of Johannesburg*

[2]andren@uj.ac.za

Abstract—**This paper explores a computer-aided detection scheme to aid radiologists in making a higher percentage of correct diagnoses when analysing chest radiographs. The approach undertaken in the detection process is to use several proprietary image processing algorithms to adjust, segment and classify a radiograph. Firstly, a Difference of Gaussian (DoG) energy normalisation method is applied to the image. By doing this, the effect of differing equipment and calibrations is normalised. Thereafter, the lung area is detected using Active Shape Models (ASMs). Once identified, the lungs are analysed using Local Binary Patterns (LBPs). This technique is combined with a probability measure that makes use of the the locations of known abnormalities in the training dataset. The results of the segmentation when compared to ground truth masks achieves an overlap segmentation accuracy of 87,598±3,986%. The challenges faced during classification are also discussed.**

***Keywords*— Computer vision, image processing, biomedical imaging, segmentation, texture analysis**

I. Introduction

South Africa is estimated to have the third highest incidence of tuberculosis (TB) in the world according to the World Health Organisation (WHO) [1]. The high incidence of the disease in this region is sustained through a large number of new infections each year, with approximately one percent of the South African population developing tuberculosis infections annually [2].

People infected with HIV are around 30 times more likely to develop tuberculosis than the rest of the civilian population. Tuberculosis is also the main cause of HIV-related death [1]. Mineral miners exhibit the highest prevalence of the disease of any working population in sub-Saharan Africa [3]. Due to prolonged exposure to silica dust and the resulting silicosis, gold miners are especially susceptible to the development of pulmonary tuberculosis during their lifetimes [4].

These affected civilian and mining populations generate large volumes of chest radiographs that must be dealt with efficiently and effectively by radiographers and radiologists in order to alleviate the burden of this epidemic. The South African National AIDS Council (SANAC) has set a goal of reducing the number of new infections and deaths from tuberculosis by fifty percent between 2012 and 2016 [2]. To fulfil this goal every possible approach to early detection and diagnosis needs to be evaluated.

This paper seeks to address those issues by providing an image processing (IP) based tuberculosis detection scheme to assist radiologists in detecting the disease when analysing chest radiographs.

An overview of the problem and the classification complexities involved are provided in Section II of this paper. In Section III, a detailed description of the proposed work flow used during classification is given. The experiments performed and the results of these experiments are explained in Section IV. Section V concludes the paper and Section VI provides suggestions for future work.

II. The Problem

Radiographs captured from separate sources can demonstrate different characteristics that complicate the further classification of such images when using analysis techniques such as machine learning. Factors influencing an image's characteristics include user settings, capturing equipment used, post-processing and analogue digitisation. Poorly trained or overworked radiographers may also neglect to give the correct instructions to the patient when performing a chest X-ray (CXR). This can result in a "collapsed" chest radiograph. These radiographs are more difficult to assess and are usually due to the patient slouching or not raising their arms correctly during capture. An indication of the variance that can be observed in two radiographs of the same person is shown in Figure 1.

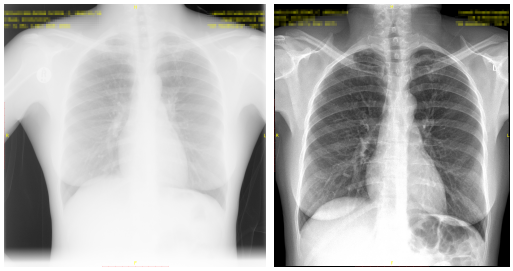

Figure 1: Two radiographs captured from the same patient

These radiographs were captured at the same clinic, 6 weeks apart, by separate radiographers using different equipment. Large variances between CXRs are particularly noticeable in cases of mine worker radiographs taken in clinics associated with rural communities. It is this level of variance that is often the limitation to the application of automated tools in TB identification.

To mitigate these factors, the image must be processed using some form of image normalisation. Applying the correct normalisation method can increase segmentation and detection accuracy, especially when analysing data captured from more than a single dataset [5] or from more than one form of source equipment. To address this issue, the next section describes the proposed workflow used to pre-process and classify an unnormalised radiograph.

III. Proposed Workflow

There are three main research areas that can be identified from the literature with regards to computer-aided radiographic analysis: Normalisation, segmentation and analysis [6]. Processing in each of these areas should take place to achieve an accurate analysis of the presented radiographic material. A diagram describing the order of processing is shown in Figure 2.

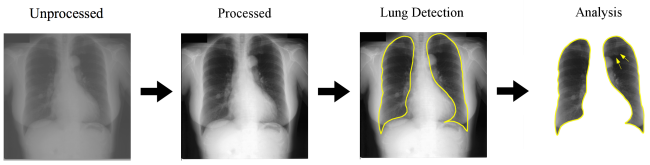

Figure 2: An overview of the detection process

The workflow used in this paper is similar to other screening systems, such as CAD4TB by Delft Imaging Systems [7], but is differentiated by the classification approach used during analysis. The datasets used, as well as the normalisation, segmentation and classification techniques implemented in this paper are described in the following sections.

A. Datasets

The radiographs used in this study were drawn from three different sources:

- Japanese Society of Radiological Technology (JSRT) database [8]
- Perinatal HIV Research Unit (PHRU) digital database
- PHRU analogue database

B. Pre-processing

Since Computer-Aided Diagnosis (CAD) usually relies on training data to accurately analyse medical images, it follows that it should derive the parameters for normalising an input image from its training set. The multi-band energy normalisation procedure described by Philipsen et al. [9] follows this approach and this is the normalisation step used in this paper. The suggested method captures the energy values from each of these frequency bands for an entire dataset. The image is then reconstituted by weighting each of these bands using its corresponding energy value and summing them together. By doing this each frequency band will contain the same energy and the reconstituted image will be normalised.

To start the frequency band separation process, a Gaussian blur is applied to the input image with $\sigma = 1$. The output of this operation is subtracted from the input image. This results in a high frequency image I_1. This process is repeated for $n - 1$ iterations. For each iteration the blurred image from the previous step is used as the input image and the σ used for the Gaussian blur is increased according to Equation 1:

$$\sigma_i = 2^{i-1} \tag{1}$$

where i is the current iteration.

On the n^{th} iteration the blurred image is not subtracted from the input image but rather used as the output for this iteration. The n^{th} output image contains all the remaining low frequency information of the original input image. The output of the whole process is the original input image divided into n frequency sub-bands $I_1 \cdots I_n$.

The frequency separation process is applied to each radiograph in the training set, using $n = 6$ frequency bands as described in [9]. To illustrate the output of this process, the frequency bands of a sample radiograph are shown in Figure 3 (The bands have been equalised for better visual clarity).

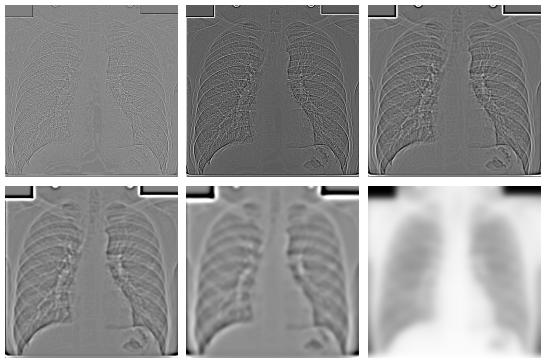

Figure 3: Frequency sub-bands

Once an image has been processed, each resulting sub-band can then be weighted and summed together using Equation 2:

$$I = \sum_{i=1}^{n} \lambda_i I_i \tag{2}$$

If each weighting λ_i is set to 1, the reconstituted image I will be the original input image.

The energy for a sub-band is expressed as its standard deviation e, as described by Equation 3:

$$e = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (p_i - \mu)^2} \tag{3}$$

Given an unnormalised sample radiograph, the energy value e_i is found for each sub-band in the sample and the final weighting is calculated as follows:

$$\lambda_i = \frac{r_i}{e_i} \tag{4}$$

The sample image is then reconstituted using Equation 2 and the calculated weightings found using Equation 4. The output of this operation results in an energy normalised radiograph. Once normalised, processed radiographs are segmented for further processing.

C. Segmentation

Deformable model-based segmentation makes use of a "shape model" obtained from manual segmentations which

are defined by "landmarks". These landmarks are chosen by a human observer and they are used to train the system that constrains segmentations to shapes that fall within the shape model. One of the most popular deformable model-based methods used in chest radiograph segmentation is the Active Shape Model (ASM) [10]. ASMs have been used to segment lungs for texture analysis [11] and a variant of the ASM has been shown to segment lungs with a good degree of accuracy [12]. Due to the robust nature of the ASM, it is the chosen segmentation method used in this paper.

ASMs require a supervised training set, where the boundaries of the lungs for several images have been identified manually. The corresponding boundary points in each image should be placed in a similar position and each image should have the same number of boundary points. These points are said to form a feature vector and this feature vector describes the shape of the lungs in each image.

When performing an ASM segmentation, the first t largest eigenvalues are extracted from a covariance matrix generated from the model shape vectors. The corresponding eigenvectors are then placed into a matrix

$$\mathbf{\Phi} = [\phi_1, \phi_2, \cdots, \phi_t] \tag{5}$$

where $\phi_1, \phi_2, \cdots, \phi_t$ are the eigenvectors corresponding to the descending eigenvalues $\lambda_1, \lambda_2, \cdots, \lambda_t$.

The lung shape can be approximated from the first t modes using the mean shape and a weighted sum of these deviations, as described by Equation 6:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \mathbf{\Phi b} \tag{6}$$

where $\bar{\mathbf{x}}$ is the mean shape and $\mathbf{b} = [b_1, b_2, \cdots, b_t]^T$ is a vector of weights for the model. $\mathbf{b}$ is calculated by using the following equation:

$$\mathbf{b} = \mathbf{\Phi}^T (\bar{\mathbf{x}} - \mathbf{x}) \tag{7}$$

New shapes are generated by varying the weights of the elements in vector $\mathbf{b}$ within a suitable constraint. This constraint is defined by Equation 8:

$$-N\sqrt{\lambda_i} \leq b_i \leq N\sqrt{\lambda_i} \tag{8}$$

where b_i is an element in $\mathbf{b}$, λ_i is the i^{th} eigenvalue and N is a constant.

In order to perform a segmentation, separate manually segmented feature vectors are used which delineate the lungs for m training images. Each feature vector must have a size of n. Using these feature vectors, models are created from profiles sampled through each vector point. In this implementation, the profiles sample pixels perpendicularly to the gradient of the contour at each point using Bresenham's line algorithm. k pixels are sampled from either side of the corresponding boundary point, resulting in each profile having a length of $s = 2k + 1$.

Once sampled, the first derivative, or gradient, of each profile is found. To create a gradient profile, each element at position i in the greyscale profile is replaced with the difference between that element and the element at $i - 1$. The gradient profiles are then normalised by dividing each value in the profile by the sum of the absolute values of the profile. A gradient profile is represented by a set of values $\mathbf{g} = g_1 \cdots g_s$ and after normalisation $\sum_{i=1}^{s} g_i = 1$.

The assumption is that the profiles are distributed as a multivariate Gaussian, which is a Gaussian distribution generalised to higher dimensions. An n -dimension multivariate Gaussian can be described by an n-dimensional mean vector and an n-by-n dimension covariance matrix. The mean profile $\bar{\mathbf{g}}$ and the covariance matrix $\mathbf{S}_\mathrm{g}$ are therefore calculated for each vector point from the gradient profiles $\mathbf{g}_1 \cdots \mathbf{g}_\mathrm{n}$ in each training image.

Once the system has been trained, the mean shape is used as an initial shape vector for an input sample image. New profiles are sampled from each point in this mean shape vector. The sample points for these profiles are iteratively shifted along "whiskers" defined for each point in the feature vector.

The Mahalanobis distance between the new profiles and the profile model is used to fit the model. The Mahalanobis distance is used since it takes into account the covariance of the profiles and is invariant to differences in scale and correlation of the distance variables. Given a profile $\mathbf{g}_\mathrm{i}$ from the sample image, the Mahalanobis distance $f(\mathbf{g}_\mathrm{i})$ is calculated using Equation 9:

$$f(\mathbf{g}_\mathrm{i}) = (\mathbf{g}_\mathrm{i} - \bar{\mathbf{g}})^T \mathbf{S}_\mathrm{g}^{-1} (\mathbf{g}_\mathrm{i} - \bar{\mathbf{g}}) \tag{9}$$

Each whisker is then iteratively searched for the profile $\mathbf{g_i}$ that minimises the Mahalanobis distance. The profile that minimises the Mahalanobis distance will have the highest statistical similarity to the profile model for that point in the vector. A new boundary point is located at the centre of each of these minimising profiles. This process is applied at multiple resolutions, by adjusting the pixel step size p of the sampling profiles and scale σ of the images. At the original resolution $\sigma = 0$ and $p = 1$. At the next level $\sigma = 1$ and $p = 2$. Thereafter σ and p are doubled for each sampling resolution. For each resolution, coarse to fine, the whiskers are searched and the centre of the profile with the smallest Mahalanobis distance is used as the starting point for the iterations at the next resolution. The detected points at the finest resolution are used to plot an accurate segmentation vector that delineates the lungs in a sample radiograph.

The accuracy of segmenting images in the JSRT dataset using this process, when compared to their ground truth segmentations, is discussed in the results section of this paper.

Once the lungs of the images have been segmented, they are then analysed for the presence of abnormalities. The following sections describe the implemented classification approach.

D. Subdivision

Before a radiograph can be classified, the mean lung shape of the training set must be found. This shape is divided into regions using a process similar to the one demonstrated in [11]. In total, the lungs are subdivided into 42 regions at 3 different scales. The first scale has 24 regions, the second 12 regions and the third has 6 regions. The regions at each scale are equal in area. This is applied to the mean shape of the 154 lung nodule images in the JSRT database, with the lung boundary landmarks being provided by the SCR database [13]. The result of this subdivision is shown in Figure 4.

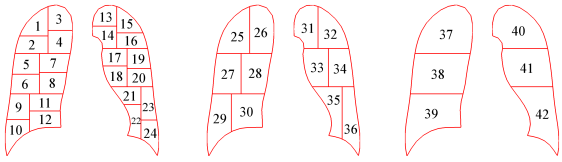

Figure 4: Mean lung shape subdivision

Once subdivided, a mask representing these regions is warped onto a sample lung shape using radial basis function (RBF) interpolation.

E. Radial basis function interpolation

Interpolating multi-dimensional scattered data is often achieved using an RBF. The characteristic feature of an RBF is that its response monotonically decreases or increases as the distance from a central point is increased. The RBF interpolant is given by Equation 10:

$$s(\mathbf{x}) = \sum_{i=1}^{n} w_i \, \phi(\| \mathbf{x} - \mathbf{x}_i \|) \qquad (10)$$

where n is the number of points in a shape, w_i is a weighting, $\mathbf{x}$ is a shape vector and $\phi(r), r \geq 0$, is a radial basis function.

Since the sample shape $\mathbf{s} = [s(\mathbf{x}_1), s(\mathbf{x}_2), \cdots, s(\mathbf{x}_n)]^T$ and mean shape $\mathbf{x}$ are known, the weightings $\mathbf{w}$ are obtained by solving a linear system containing the square matrix $\mathbf{A}$, resulting from radial basis functions $\phi(\| \mathbf{x} - \mathbf{x}_i \|)$. The linear system therefore has the form $\mathbf{Aw} = \mathbf{s}$.

The resulting weightings can then be used to interpolate any point within a sample lung shape onto the mean lung shape using Equation 10. An illustration of interpolating the mean shape onto a sample radiograph is shown in Figure 5.

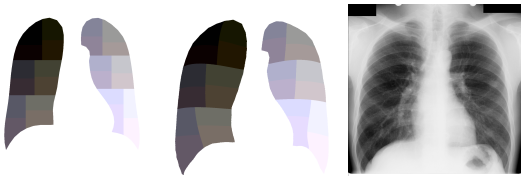

Figure 5: Interpolation of the mean shape onto a sample lung segmentation

After interpolating the mask onto a sample radiograph, each region defined by Figure 4 can be classified as normal or abnormal. To do this a training set is constructed for each region using all of a region's abnormal feature vectors and the same number of randomly selected normal feature vectors from the training set. To extract a feature vector from a region in a radiograph, Local Binary Patterns (LBPs) are implemented.

F. Local Binary Patterns

LBPs were first introduced by Ojala et al. [14] and have proven to be an effective descriptor in the classification of arbitrary textures [15].

To process an image with the LBP operator it must first be converted to a greyscale image. The LBP operator is a grey-level descriptor and does not take the RGB colour values of an image into account. This makes the LBP especially suitable to chest radiographs since they do not contain any colour information.

After the image has been converted to greyscale, the LBP operator is then applied to each pixel within the image. LBPs use local neighbourhoods of a set size around each pixel in an image. These neighbourhoods generate a binary number for each neighbouring pixel based on that pixel's value in comparison with the centre pixel value. The centre pixel is then replaced with this binary value in the output image.

Once the LBP operator has been applied to every pixel in the image, a feature vector describing the textural properties of the image is then obtained from a histogram of the LBP values. During classification this feature vector acts as a statistical representation of the texture and this is what is used to differentiate it from other textures.

The $LBP_{P,R}^{riu^2}$ descriptor used in this research is scalable, rotation invariant and exhibits reduced dimensionality by using uniform patterns. Uniform patterns are the most common LBPs present in observed textures. Uniform patterns are allocated their own individual bins in a histogram while the $P + 1$ (non-uniform) patterns are grouped into a single bin due to their low individual frequency. The result of processing a radiograph with $LBP_{P,R}^{riu^2}$ at three scales is shown in Figure 6.

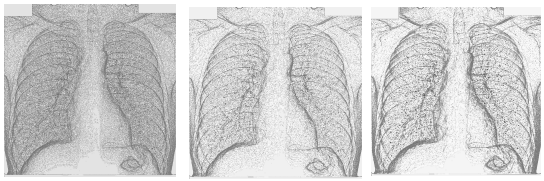

Figure 6: The result of processing a radiograph with $LBP_{8,1}^{riu^2}$, $LBP_{16,2}^{riu^2}$ and $LBP_{24,3}^{riu^2}$

In the above images, the histogram bins have been spaced at equal intervals over the greyscale range.

Once a feature vector has been generated for a region in a sample radiograph, the k-Nearest Neighbours are extracted from the training set containing both the normal and the abnormal feature vectors. If the vector to be classified is in

the training set, then it is not included by leaving out the zero distance neighbour. If a region does not have enough abnormal samples then abnormal samples are randomly included from other regions to build the training set. A sample is then classified as the class (Normal or abnormal) with the majority vote of the k-Nearest Neighbours.

IV. Results

A. *Normalisation*

Using the JSRT dataset as the training set, the effect of normalising a subset from the PHRU dataset is illustrated by Figure 7 and Figure 8.

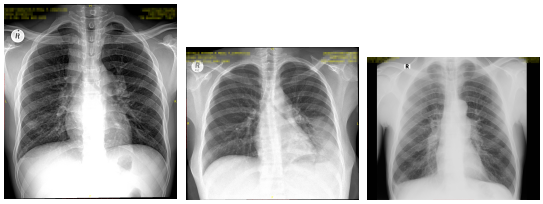

Figure 7: Three radiographs from the PHRU set

It is observed from Figure 7 that images from the PHRU database have varying dimensions, so these images are padded to be square before normalisation. In addition to the differing image sizes, the samples selected also illustrate the dissimilar image characteristics of the radiographs in this dataset.

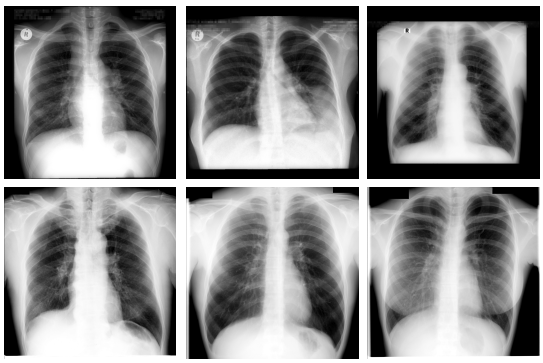

Figure 8: Energy normalised PHRU radiographs (Row 1) and reference radiographs from the JSRT set (Row 2)

The images in first row of Figure 8 are the result of applying the energy normalisation process to the selection of PHRU radiographs in Figure 7. For comparison purposes, radiographs from JSRT database are provided in the second row of Figure 8. To aid visual observation, the grey values in the central 70% of each image in Figure 8 have been used to normalise the radiographs between 0 and 255. It is observed that the normalised images have similar characteristics to the JSRT images and are not easily identified as images from a different dataset.

B. *Lung Segmentation*

Once the PHRU set has been normalised, the images are segmented before analysis. The JSRT set is used as the training set. The JSRT images are also segmented and their segmentations compared to the manual segmentations from the SCR set. A comparison of three segmented radiographs from the JSRT dataset, along with their respective ground truths, is shown in Figure 9.

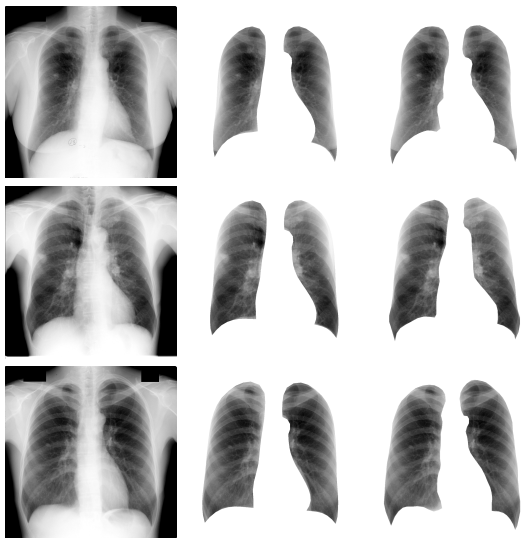

Figure 9: JSRT inputs (Column 1), ASM segmentations (Column 2) and SCR ground truths (Column 3)

To statistically compare segmentations an "overlap" measure is used, defined by:

$$\Omega = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP} + \mathrm{FN}}$$

where TP (True positive) is the correctly classified lung area, FP (False positive) is the area incorrectly classified as lung and FN (False negative) is the area incorrectly classified as non-lung.

To test the accuracy of the segmentations, half of the images in the JSRT database are used to train the system and the other half are then used to test the system. When compared to the ground truth segmentations from the SCR database, the JSRT images achieved an overlap segmentation accuracy of 87,598±3,986%.

C. *Classification*

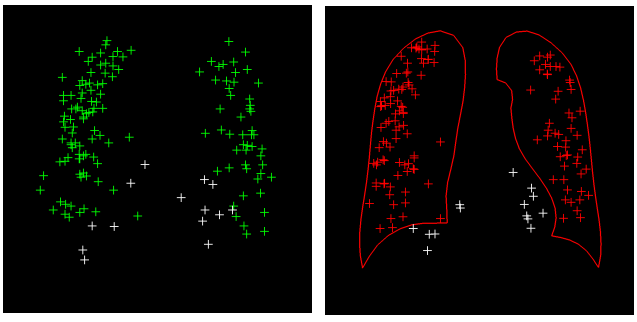

Figure 10: The bounded (Green) and unbounded (White) JSRT abnormalities warped onto the mean shape (Red)

Taking all of the abnormalities in the JSRT lung nodule set (JPCLN), a plot of the original nodule positions and the result of warping these points onto the mean shape using RBF

interpolation is shown in Figure 10. The JSRT set does not only contain lung abnormalities or abnormalities that are bounded by the landmarks in the SCR database. There are 14 unbounded abnormalities in the JPCLN set and these points are indicated by the white crosshairs in Figure 10. It is observed that all of the abnormalities bounded by the SCR landmarks are warped onto the mean shape and that none of the unbounded abnormalities lie within the mean shape after interpolation.

The probabilities generated by these bounded abnormalities will be used as an additional parameter during the classification process. The number of abnormalities per region and the probability of an abnormality occurring in that region is listed in Table 1. The total number of bounded abnormalities in the JSRT set is 140.

TABLE I
THE NUMBER OF ABNORMALITIES AND PROBABILITY PER REGION

R	#	P(A)	R	#	P(A)	R	#	P(A)
1	14	0.1	15	10	0.071	29	11	0.079
2	20	0.143	16	4	0.029	30	7	0.05
3	6	0.043	17	1	0.007	31	2	0.014
4	0	0	18	3	0.021	32	14	0.1
5	14	0.1	19	5	0.036	33	4	0.029
6	13	0.093	20	11	0.079	34	16	0.114
7	0	0	21	3	0.021	35	7	0.05
8	4	0.029	22	4	0.029	36	8	0.057
9	8	0.057	23	7	0.05	37	40	0.286
10	3	0.021	24	1	0.007	38	31	0.221
11	4	0.029	25	34	0.243	39	18	0.129
12	3	0.021	26	6	0.043	40	16	0.114
13	1	0.007	27	27	0.193	41	20	0.143
14	1	0.007	28	4	0.029	42	15	0.107

In preliminary tests, the modified LBP presented in this paper did not accurately detect the subtle abnormalities present in the JSRT database. This is mainly due to the sample size of the regions described in Section III D, which often contain a great deal of texture information which is similar to the sampled non-nodule regions in the training set. Abnormalities present on the edges of more than one region are also problematic due statistical nature of the LBP descriptor. This will be addressed in future research.

V. CONCLUSION

In this paper, a tuberculosis detection workflow was proposed using energy normalisation, ASM segmentation and modified Local Binary Patterns for classification. The normalisation and segmentation steps provided satisfactory results through visual inspection and segmentation accuracy. Probabilities were gathered from the locations of abnormalities in JSRT dataset to generate an additional classification parameter. The classification of the subdivided regions proved problematic due to the relatively large sample sizes used, the subtle abnormalities present in the JSRT dataset and the nature of LBP descriptor. In future work, subsets of these regions will be analysed using overlapping windows to avoid the problems associated with abnormalities that border more than one region. By doing this, the inclusion of healthy tissue in sampled abnormal regions will also be reduced, resulting in a better statistical representation of the abnormal textures during training and analysis.

REFERENCES

[1] World Health Organisation, "Global Tuberculosis Report 2012," 2012.

[2] South African National AIDS Council, "National Strategic Plan on HIV, STIs and TB 2012 - 2016," 2012.

[3] S. Basu, D. Stuckler, G. Gonsalves, and M. Lurie, "The production of consumption: addressing the impact of mineral mining on tuberculosis in Southern Africa," Globalization and Health, vol. 5, 2009.

[4] B. Girdler-Brown, N. White, R. Ehrlich, and G. Churchyard, "The burden of silicosis, pulmonary tuberculosis and COPD among former Basotho goldminers," American Journal of Industrial Medicine, vol. 51, no. 9, pp. 640–647, 2010.

[5] B. van Ginneken, C. M. Schaefer-Prokop, and M. Prokop, "Computer-aided diagnosis: how to move from the laboratory to the clinic," Radiology, vol. 261, no. 3, pp. 719–732, 2011.

[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, vol. 20, no. 12, pp. 1228–1241, 2001.

[7] L. Hogeweg, C. Mol, P. A. de Jong, R. Dawson, H. Ayles, and B. van Ginneken, "Fusion of local and global detection systems to detect tuberculosis in chest radiographs," Proceedings of the Medical Image Computing and Computer Assisted Intervention, vol. 13, pp. 650–657, 2010.

[8] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K. Komatsu, M. Matsui, H. Fujita et al., "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," AJR, vol. 174, pp. 71–74, 2000.

[9] R. Philipsen, P. Maduskar, L. Hogeweg, and B. van Ginneken, "Normalization of chest radiographs," Proceedings of SPIE, vol. 8670, pp. 86 700G–86 700G–6, 2013.

[10] T. F. Cootes, and C. J. Taylor, "Active shape models: 'smart snakes'," Proceedings of the British Machine Vision Conference, pp. 266–275, 1992.

[11] B. van Ginneken, S. Katsuragawa, B. M. ter Haar Romeny, M. Viergever et al., "Automatic detection of abnormalities in chest radiographs using local texture analysis," IEEE Transactions on Medical Imaging, vol. 21, no. 2, pp. 139–149, 2002.

[12] B. van Ginneken, A. F. Frangi, R. F. Frangi, J. J. Staal, B. M. ter Haar Romeny, and M. A. Viergever, "Active shape model segmentation with optimal features," IEEE Transactions on Medical Imaging, vol. 21, pp. 924–933, 2002.

[13] B. van Ginneken, M. Stegmann, and M. Loog, "Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database," Medical Image Analysis, vol. 10, pp. 19–40, 2006.

[14] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 24, no. 7, pp. 971–987, 2002.

[15] J. Leibstein, A. Findt, and A. Nel, "Efficient texture classification using local binary patterns on a graphics processing unit," in Proceedings of the twenty-first annual symposium of the pattern recognition association of South Africa, pp. 147–152, 2010.

Joshua Leibstein holds the degrees of BSc IT, BSc IT (Hons) and is presently a postgraduate student at the University of Johannesburg. As an undergraduate he won the SA Microsoft Imagine Cup for his research into the detection of TB using texture analysis.

Andre Leon Nel is a professor in the Department of Mechanical Engineering Sciences at the University of Johannesburg. He received his D Ing (Mech) from the Rand Afrikaans University in 1993. He also holds a BA (Class Lang) from the Rand Afrikaans University. He has been on the full time staff at first the Rand Afrikaans University and then the University of Johannesburg since 1984. Since 2010 he has been the Head of School (Mechanical and Industrial Engineering) and is the present lead researcher in the Hypervision Research Laboratory. He has produced publications in a broad range of fields from information theory to legged robotics and the ethical implications of engineering education activities. He is a Member of the IEEE and a reviewer for ISI listed journals.