Lung Disease Detection Using Frequency Spectrum Analysis

Ching Ming Jimmy Wang Mamatha Rudrapatna Arcot Sowmya

School of Computer Science and Engineering,
University of New South Wales, Sydney,
NSW 2052, Australia

{cmjw514, mamathar, sowmya}@cse.unsw.edu.au

Abstract

This paper introduces a novel approach for automatic detection of a type of diffuse lung pattern, known as honeycombing, from high resolution computed tomography (HRCT) scans of the lung. The algorithm, which is based on frequency spectrum analysis of the HRCT image, assesses and combines outputs from a pre-defined Gabor filter bank to form a preliminary lesion detection mask. Several morphological filters are then employed to remove noise from the detection mask. The algorithm is applied to a total of 352 images and the outputs are validated against lung HRCT images marked by 2 qualified radiologists. The algorithm achieved a sensitivity of 87.5% and a specificity of 84.4%, which compares favorably with other approaches.

1. Introduction

Computed tomography (CT) scans, alternatively known as CAT scans, are computer rendered x-ray scans which provide a cross-sectional view of the screened object. They are a standard diagnosis technique for lung diseases. Examination of CT images is performed by qualified radiologists. However, this is a time consuming and weary task as a continuous scan of the lung can produces up to 500 – 600 images for an individual patient. As most of the slices are similar in appearance, the task of diagnosis can be tedious and can potentially reduce diagnostic accuracy. To improve the efficiency, automated disease detection algorithms may be developed. In this paper, we tackled a specific type of lung lesion known as honeycombing which belongs to the diffuse lung opacity disease category. Honeycombing refers to the characteristic appearance of pulmonary fibrosis.

Diffuse opacity lung diseases, as the name suggests, spread to large portions of the lung and form distinctive patterns. Such lesions are often studied using texture analysis methods, a classical research area in image processing. Amongst the entire diffuse opacity lung disease family, honeycombing is by and large the hardest to detect automatically. This has been shown separately by Uppaluri et al [1] and Doi et al [2]. They developed multi-class diffuse lung disease classification systems using different approaches and reported that honeycombing is the least detectable pattern in both studies. Figure 1 shows a honeycombing region in contrast to normal lung tissue.

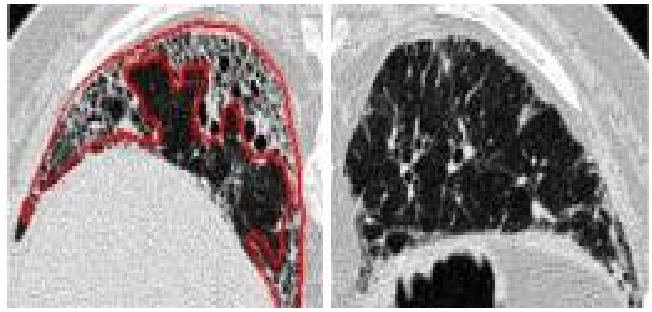

Figure 1: Top part of the lung on the left is affected by honeycombing while the right image shows a normal lung.

Most of the currently available honeycombing detection algorithms are based on statistical learning [1], [2]. In this paper, we take a different approach and base our algorithm on frequency spectrum analysis. In particular, Gabor filters are employed to assist the frequency extraction process. Some of the advantages of our algorithm including fast processing time with each slice taking only around 20 seconds to process on a standard home personal computer. Also, no training is needed for this algorithm, in contrast to statistical learning machines which are notorious for their long and exhausting training phase. Our experiments have shown that this approach yields comparative results (if not better) to statistical methods.

2. Related Work

Mitani et al [3] performed texture classification on diseased lung patterns. The aim was to classify different disease patterns from images containing only homogeneous lesion textures. The team created their test sets by manually cutting the diseased regions out from CT images which formed a new image with block size 128×128. The textural images were divided into 5 classes, namely reticular, ground glass, nodular, consolidation and normal lung tissue. Textural features were extracted using an intensity histogram and a fixed scale Gabor filter bank with 4 orientations. They reported that the usage of intensity histogram together with Gabor filtered output dramatically improved the classification results compared to only Gabor filtered outputs and they achieved a classification accuracy of 90%.

Doi et al [2] developed an automated computerized method of detecting diffuse lung opacity diseases. Their data set contained ground glass, honeycombing, reticular and linear opacity, consolidation, emphysema and nodular opacity. Doi's team first employed an algorithm to segment the lung region out from the image and divide it into 16×16 blocks. Twenty or more statistical features were extracted from each block and a Bayesian classifier is adopted to classify between normal lung tissue and different opacities. A sensitivity of 89.7% (26/29) for honeycombing detection was reported by the team.

Uppaluri et al [1] used AMFM(Adaptive Multiple Feature Method) to identify diffuse lung opacity diseases. Their data set contained ground glass, honeycombing, reticular and linear opacity, consolidation, emphysema and nodular opacity. The group first isolated the lung region from the background and divided it into 31×31 pixel blocks. Higher order statistical features were extracted from each of the blocks and Adaptive Multiple Feature Method (AMFM), a texture classification algorithm based on statistical learning, was used to classify each of the blocks into different disease patterns. The group reported a sensitivity of 82.5% (33/40) and a specificity of 99.5% for honeycombing detection.

3. Overview

In our approach, we moved away from statistical texture analysis and machine learning completely. Instead we employed frequency spectrum analysis with Gabor filterbank and followed by post processing of the results. Specifically, the lower and higher frequency spectrum of the input image is extracted by a Gabor filterbank. The filtered output are combined together after thresholding. Two types of binary filters, a region-based filter and an area-thresholding filter, are used to remove noise in the detector output. Figure 2 outlines the basic building blocks of our proposed algorithm.

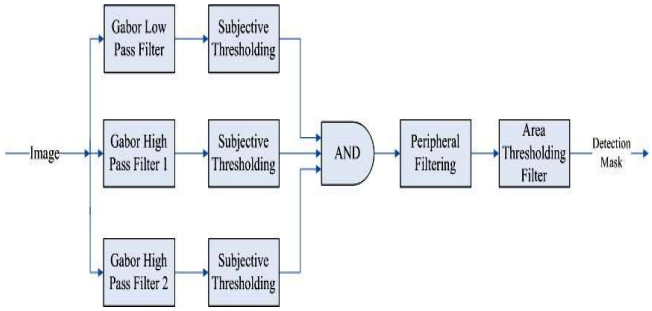

Figure 2: Outline of the honeycombing detection algorithm.

4. Gabor Filters

For any signal processing application, it is always desirable to have filters with sharp cut offs. However in practice such filters, best represented by rectangular filters (all or nothing filters), inject noise into the filtered output in the spatial domain due to their sinc-like Fourier transform counterpart. Noise level decreases as the cut off of the filter smooths out and therefore a trade off between noise level and filter cut offs has to be made. This is known as the joint spatial-frequency error and it is formally proven by J. Daugman [4] that 2D Gabor filters minimize this trade off. In this paper, this family of filters is used to obtain parts of the image frequency spectrum. 2D Gabor filters and their Fourier transform can be written as

$$g(x,y) = \frac{1}{2\pi\sigma_x\sigma_y} exp\left[-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + j2\pi(xu_0 + yv_0)\right)\right] \quad (1)$$

$$\hat{g}(u,v) = exp\left[-\left(\frac{(u-u_0)^2}{2\sigma_u^2} + \frac{(v-v_0)^2}{2\sigma_v^2}\right)\right] \quad (2)$$

where $\sigma_u = 1/2\pi\sigma_x$ and $\sigma_v = 1/2\pi\sigma_y$ and u_0, v_0 are the center frequency of the 2D filter. Gabor filters are also linear shift invariant and orientation selective filters, which makes them very useful in image processing domains.

5. Frequency Masking for Honeycombing Detection

Honeycombing patterns have large internal intensity variations with a mean around one-third the maximum pixel intensity as determined experimentally. This is in contrast to

normal lung parenchyma which have constant low intensity. Low pass filtering of an image containing honeycombing would allow the lesion areas to stand out as the lower frequency spectrum encodes local average intensity information. As mentioned in section 4, the lower frequency spectrum is extracted using Gabor filters which have the following form:

$$\hat{g}_L(u,v) = exp\left[\frac{u^2 + v2}{2\sigma_L^2}\right] \quad (3)$$

This filter is often known as a Gaussian filter, which is a special case of Gabor filters.

However, a different family of lung diseases, known as increased opacity diseases, cause a constant increment of pixel intensity within the lesion area in contrast to normal lung tissue. Therefore, lesions such as ground glass opacity, result in similar local mean intensity to honeycombing and is therefore unable to be differentiated when subjected to low pass filtering. An image of ground glass is shown in figure 3.

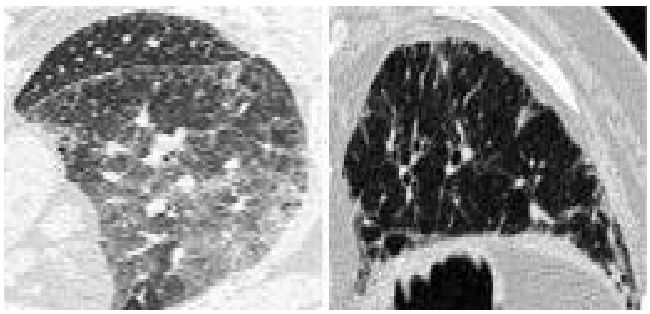

Figure 3: The left image shows an entire lung affected by ground glass while the right image shows a healthy lung.

A solution to this problem can be found by observing that the higher end of the frequency spectrum encodes information about local variations. Therefore, high pass filters should respond strongly to honeycombing patterns but not to increased opacity disease patterns.

The 2D high pass filters are designed by adopting a generic filter bank design strategy discussed by Manjunath and Ma [5]. Filters of the first two scales from a four-scale filterbank are employed as high pass filters. Filtered outputs of different orientations at the same scale are averaged to produce an orientation invariant output. The first two scales of the filter bank layout in the frequency domain is shown in figure 4.

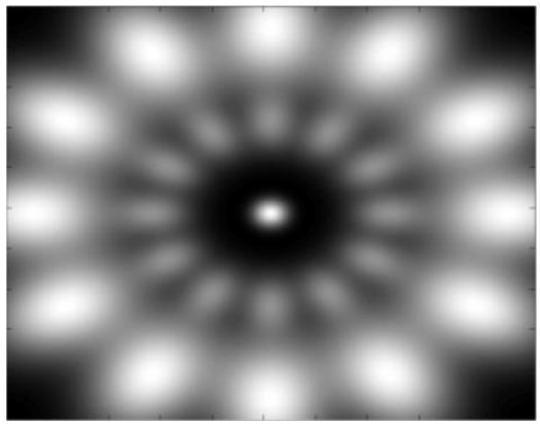

Figure 4: Honeycombing detection filter bank layout.

It is observed that although high pass filters successfully filter out increased opacity disease patterns, they also pick up other artifacts such as large nodules. Large nodules appear as large bright blobs which are similar to rectangular pulses in the signal processing sense. They have large sinc frequency spectrum which extends into the passband of the high frequency filters and hence cannot be filtered out. Although nodules also appear in the low pass filtered output, they can be easily discarded by thresholding, by noticing that the local mean intensity of nodules is around the maximum value while that of honeycombing is just one-third the maximum value.

Furthermore, noise found in the high pass and low pass filtered outputs forms a largely non-overlapping set while the wanted information is contained in both outputs. Therefore, noise can be eliminated by first binarizing the filtered outputs by thresholding and then taking the intersection between the binary masks. Thresholding in this case turns pixel values within the threshold to 1's and the rest to 0's. Figure 5 shows a preliminary honeycombing detection mark obtained by the above process.

It is easy to see that the preliminary mask still contains noise. The next section proposes two different family of filters aimed at removing as much of this noise as possible.

6. Post Processing

It has been suggested by radiologists that honeycombing only occurs in the peripheral regions of the lung [6]. Part of the noise in the detection mask can be removed by creating a peripheral lung mask and intersecting this mask with the preliminary honeycombing detection mask. To create peripheral mask, first the lung regions are extracted using thresholding and morphological operations. The difference

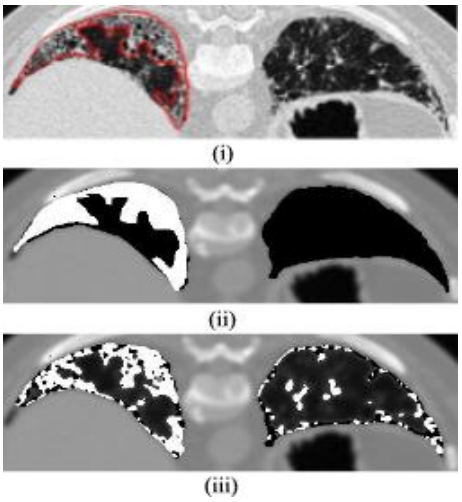

Figure 5: (i) the original lung image; (ii) the honeycombing region labelled by radiologist and (iii) the preliminary honeycombing mask.

image between the convex hull that covers both the lungs and the lung masks yields a mask for mediastinum region (cavity between the lungs where the heart is located). Lung boundary lining the mediastinum is removed using the mediastinum mask. The remaining lung boundary is then dilated to get peripheral mask.

We must then clean up any left over noise, which primarily comprises boundary artifacts and some of the tiny dots which are often only a few pixels wide. Standard techniques such as closing removes the noise as well as some of the regions of interest (ROIs), which are made up of groups of small unconnected pixels (1's). To avoid the elimination of ROIs, a binary area-thresholding filter was devised. As the honeycombing mask contains 1's and 0's, this filter adds up the total number of 1s inside the window and sets the pixel to 0 if the total is less than a certain threshold. Windowed area-thresholding filter can be formally defined as follows:

$$Im(x,y) = 0 \mid \sum_{q=-N}^{N} \sum_{p=-N}^{N} Im(x+p, y+q) < \gamma \quad (4)$$

where γ is the area threshold and the window size is equal to *2N + 1*. Figure 6 shows the output of the final honeycombing detection mask.

7. Experimental Results

A total of 376 lung HRCT images labelled by qualified radiologists were used to evaluate the performance of the proposed honeycombing detection algorithm (HCDA). 24 images in this set were known to contain honeycombing based

Figure 6: The final honeycombing mask.

on labels. Sensitivity of the algorithm was measured by considering each individual diseased region labelled by the radiologists and counting the result of HCDA to be correct if the algorithm also picked up some part of the labelled region. For example, if there are three different honeycombing regions marked by the doctor in one image and the algorithm picks up two of these regions, then the algorithm has a sensitivity of 66%. Sensitivity is also known as the accuracy of the algorithm. Formally, sensitivity is defined as

$$\text{Sensitivity} = \frac{\#\text{True Positives}}{\#\text{True Positives} + \#\text{False Negatives}} \quad (5)$$

where

- **True Positive(TP):** Regions labelled by the radiologist as honeycombing and recognised by HCDA.
- **False Positive(FP):** Regions recognised by HCDA but not labelled by radiologist as honeycombing.
- **False Negative(FN):** Regions labelled by radiologist as honeycombing but not recognised by HCDA.
- **True Negative(TN):** Regions not labelled or recognised as honeycombing by either the radiologist or HCDA.

There were 40 honeycombing regions within the 24 honeycombing images in the data set and HCDA scored an sensitivity of 87.5% (35/40). Some of the detection results can be found on the last page of this paper.

The specificity of the algorithm, which one could think as one minus the probability of a slice being labelled as honeycombing when there is no honeycombing present, can be calculated as follows:

$$\text{Specificity} = \frac{\#\text{True Negatives}}{\#\text{False Positives} + \#\text{True Negatives}} \quad (6)$$

HCDA was applied to the other 352 images in the database, which contained 17 different disease categories, including ground glass, emphysema, nodular opacity, interstitial thickening and parenchymal bands (but no

Set Name	TN	FP
Bronchiectasis	36	2
Bullae and Blebs	26	7
Consolidation	9	1
Emphysema	57	8
Ground Glass Opacity	26	5
Interface Sign	1	1
Interlobular Septal Thickening.	5	2
Intralobular Intersitial Thickening	31	4
Irregular Lung Opacity	20	4
Large Nodule	19	4
Mosaic Perfusion	8	0
Parenchymal Bands	22	5
Peribronchiovascular Inter. Thickening	9	0
Pleural Plaques	32	1
Small Nodule	16	4
Subplural Interstital Thickening	20	3
Traction Bronchietasis	15	4
Total	352	55

Table 1: Specificity Measurment

honeycombing). The result is shown in Table 1. The results show that HCDA achieved a specificity of 84.4% (297/352).

An additional feature of HCDA is its short processing time. The optimized Gabor filter bank filtering only needs 3 matrix multiplication and 4 Fast Fourier Transforms (FFT), which has complexity $N^2 logN$, where $N \times N$ is the size of the image. The algorithm on average only takes around 20 seconds to process one image with most of the time spent on the lung and peripheral region detection on a standard Pentium 4 2.8GHz home computer.

8. Concluding Remarks

We presented the development of a computerized honeycombing detection algorithm using frequency spectrum analysis with minimal post-processing. This technique demonstrated high sensitivity and specificity with short processing time. It has been shown by Uppaluri [1] that such values of sensitivity is comparable to the performance of a human expert. Although the proposed algorithm has a lower specificity than some the of currently available algorithms in this field, it is still acceptable because the main purpose of this algorithm is to reduce the number of slices that a radiologist needs to view. A specificity of 84.4% implies that the algorithm is capable of discarding 84 slices out of a hundred slices with high confidence and thus has successfully reduced the number of possible candidates by 6 fold.

Potential future work includes fine tuning of the Gabor high pass filters to suit the honeycombing pattern instead of using filters from a generic filter bank. It is also possible to learn parameter values for thresholds and window sizes by incorporating learning algorithms. The current algorithm is yet to be evaluated by radiologist. We have applied HCDA on 3000 unlabelled lung HRCT images obtained from two hundred patients. The algorithm picked up around 500 images from the set, which correlates well to the specificity which we obtained from our previous labelled data. The 500 images are to be validated by radiologists in the near future.

9. Acknowledgement

This research was partially supported by the Australian Research Council through a Linkage grant (2002 – 2004), with Medical Imaging Australasia as clinical and Philips Medical Systems as industrial partners.

References

[1] 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., Volume 160, Number 2, August 1999, 648 - 654.

[2] K. Doi, Y. Uchiyama, H. Abe, F. Li and J. Shiraishi *Quantitative Copmuter-aided Analysis of Diffuse Lung Disease in High-Resolution Computed Tomography*. Scientific Paper, RSNA 2002.

[3] Y. Mitani, H. Yasuda and S. Kido. *Combining the Gabor and Histogram Features for Classifying Diffuse Lung Opacities in Thin-section CT*. International conference on Pattern Recognition, Vol 1, 53 - 56, 2002

[4] John Daugman. *Uncertainty relation for resolution in space, spatial frequency and orientation optimized by two dimensional visual cortical filters*. J. Opt. Soc. Am. A/Vol.2, No. 7, 1985.

[5] B S. Manjunath, W.Y. Ma. *Texture Features for Browsing and Retreval of Image Data*. IEEE Transactions on pattern analysis and machine intelligence, Vol. 18 No. 8, August 1996.

[6] W. R. Webb, N. L. Muller and D. P. Naidich. *High Resolution CT of the Lung*. Published by Lippincott Williams & Wilkins Publishers, January 2001.

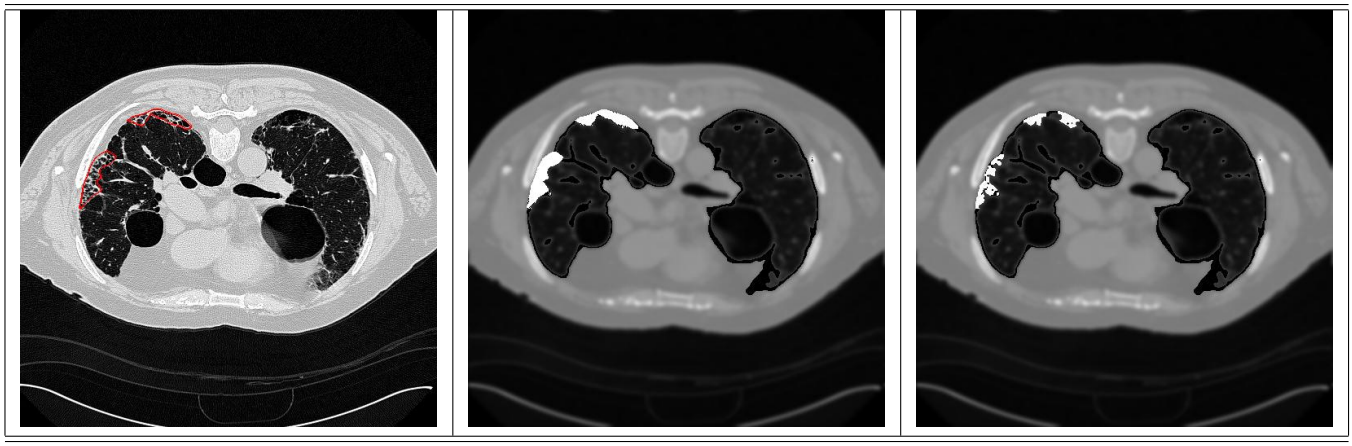

Figure 7: From left: (i) Original Image; (ii) Honeycombing regions (white) labelled by radiologist; (iii) Honeycombing regions (white) labelled by HCDA

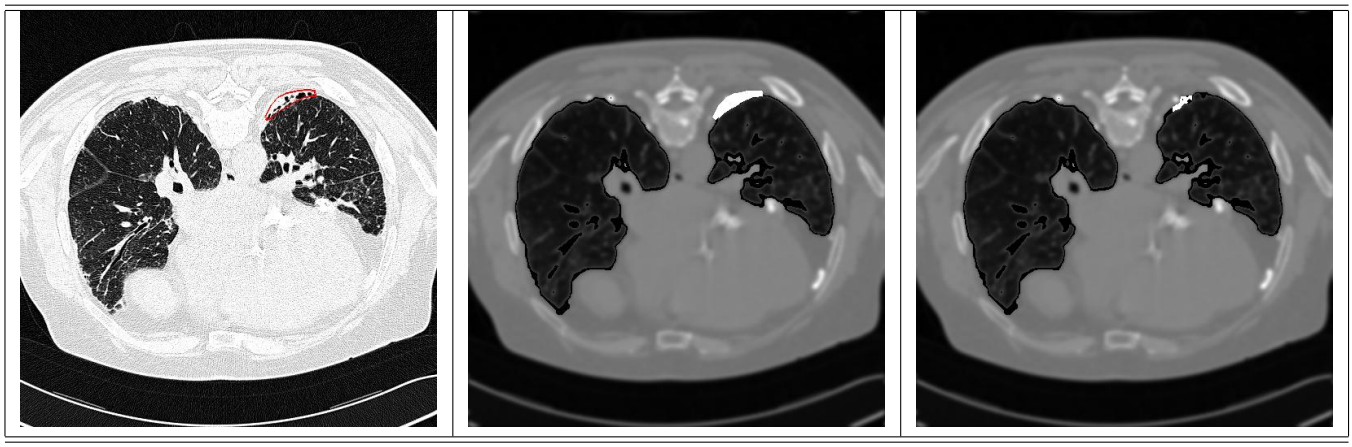

Figure 8: From left: (i) Original Image; (ii) Honeycombing regions (white) labelled by radiologist; (iii) Honeycombing regions (white) labelled by HCDA

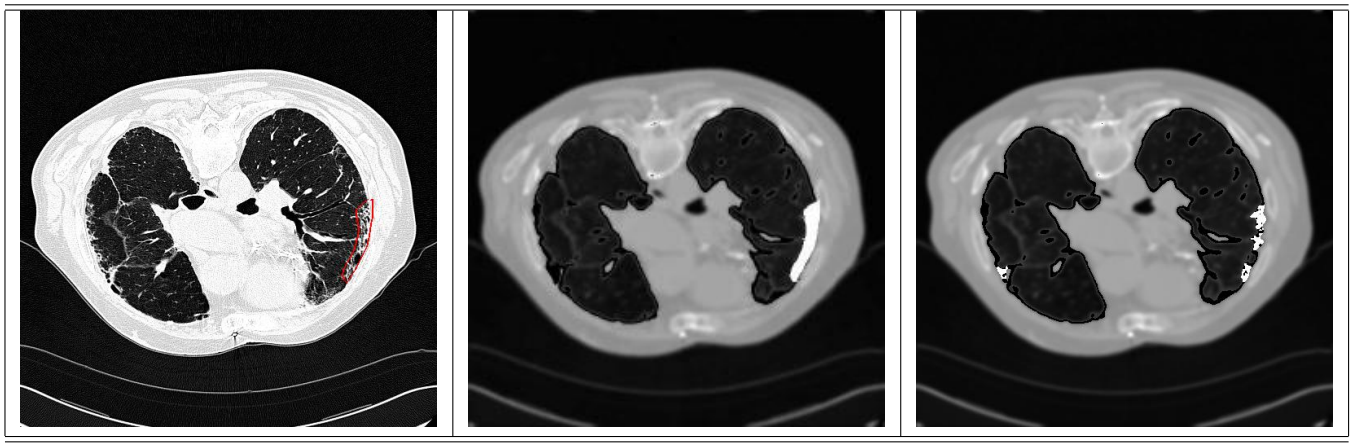

Figure 9: From left: (i) Original Image; (ii) Honeycombing regions (white) labelled by radiologist; (iii) Honeycombing regions (white) labelled by HCDA