

IMAGE ANALYSIS OF RADIOLOGICAL IMAGES
FROM PATIENTS WITH ADVANCED LUNG DISEASE

by

Mekhala Acharya
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
in Partial Fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Bioinformatics and Computational Biology

Committee:

_____	Dr. Jason Kinser, Dissertation Director
_____	Dr. Geraldine Grant, Committee Member
_____	Dr. Donald Seto, Committee Member
_____	Dr. Jeffrey Solka, Committee Member
_____	Dr. James D. Willett, Director, School of Systems Biology
_____	Dr. Timothy L. Born, Associate Dean for Student and Academic Affairs, College of Science
_____	Dr. Vikas Chandhoke, Dean, College of Science

Date: _____ Fall Semester 2012
George Mason University
Fairfax, VA

Image Analysis of Radiological Images from Patients with Advanced Lung Disease

A dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy at George Mason University

By

Mekhala Acharya
Master of Science
George Mason University, 2006
Bachelor of Science
Mangalore University, 2000

Director: Dr. Jason Kinser, Professor
School of Physics and Computational Sciences, College of Science

Fall Semester 2012
George Mason University
Fairfax, VA

Copyright © 2012 by Mekhala Acharya
All Rights Reserved

Dedication

I dedicate this dissertation to my husband Prashanth. I would not have reached this stage in my academic life had it not been for his unwavering support and faith in me. His patience and love for science gave me the encouragement to stay focussed, consistent and dedicated in my aim to complete the dissertation. It was not a question of just completing it, but doing it well and with the utmost satisfaction. I cannot thank him enough for all those hours spent reviewing my material, the discussions and constant encouragement I received. Thank you Prashanth for being there when I needed you most.

Acknowledgments

I would like to thank my advisor Dr Jason Kinser whose encouragement and faith in me made it all possible. There were days of exasperation and frustration, but it was Dr Kinser's cheerful disposition and 'can do' attitude that made this work see the light of the day. I could shrink in the presence of his brilliance, but never once did he make me feel incapable of assimilating the material or producing new work. I always came out with a renewed sense of excitement after the brain-storming sessions. The sheer simplicity with which he conveyed the most complex ideas gave me clarity and purpose which worked wonderfully well especially during the most critical periods. Thank you Dr Kinser. I would especially like to thank the radiologists at Inova Heart and Lung Disease Institute, Dr Steven Nathan and Dr Candy Albano and Registered Nurse, Lori Schlegel. The exhaustive information and feedback provided through data and meetings made the analysis feasible.

I would also like to say a special thank you to Dr Grant, whose initial input and references in understanding the science of this disease gave me the momentum I needed for a great start. Last but not the least I would like to acknowledge my committee members: Dr Seto and Dr Solka whose inputs made this dissertation possible and complete.

Table of Contents

	Page
List of Figures	viii
List of Tables	xii
Abstract	xiii
1 Introduction	1
1.1 Diagnosis and Treatment	1
1.1.1 Differential Diagnosis	1
1.1.2 Epidemiology	4
1.2 Current Challenges in Medical Imaging Modalities of Lung Disease	5
1.2.1 Chest Radiography	5
1.2.2 HRCT Technique	5
1.2.3 Magnetic Resonance Imaging	6
1.2.4 Ultrasound	6
1.3 Disease Quantification using HRCT scans	8
2 Previous Approaches to the Study of IPF	9
2.1 Quantification of Diffuse Lung Diseases	9
2.2 CT Density	9
2.2.1 Disadvantages	10
2.3 Texture Analysis	11
3 Research Rationale and Data Collection	15
3.1 Objective	15
3.2 Medical Image Analysis	17
3.3 Materials and Methods	18
3.3.1 Experimental Data-Sets	19
3.3.2 CT Scanning Protocol	19
4 Research Methods	20
4.1 Adaptive Multiple Feature Method	20
4.2 Thresholding or Feature Identification	21
4.2.1 Range Filter	21

4.3	Smoothing or Noise Reduction	26
4.3.1	Adaptive Thresholding	26
4.3.2	Quantitative CT indices	30
4.3.3	Adaptive operation	31
4.4	ROI extraction	35
4.4.1	Co-Occurrence Matrix Measures	35
4.5	Discussion	37
4.6	Backpropagation Neural Network (BPNN)	38
4.6.1	Radial Basis Function (RBF)	40
4.6.2	Implementation and Analysis	41
4.7	Gabor Filter	42
4.7.1	Goal	45
4.7.2	Implementation	47
4.8	Discussion	47
5	Pulse Images	49
5.1	PCNN Details	49
5.1.1	Original Model	49
5.1.2	Implementing in Python	51
5.1.3	Spiking Behavior	53
5.1.4	Collective Behavior	57
5.1.5	Neural Connections	58
5.2	The ICM	62
5.2.1	Minimum Requirements	63
5.2.2	ICM Theory	65
5.2.3	Connections in the ICM	66
5.2.4	Interference	68
5.2.5	Curvature Flow	70
5.2.6	Centripetal Autowaves	71
5.3	Summary	74
5.4	Comparing ICM with Gabor	75
5.5	Segmentation of IPF Images with Pulse Images	76
5.5.1	System Overview	77
5.5.2	Pulse Image Generator	77
5.5.3	ICM in Texture Analysis	78
5.5.4	Pulsation in Cycles	80
5.5.5	Pulse Streams for Classification	82

5.5.6	Analysis of Pulse Streams	83
5.6	Associative Memory	84
5.6.1	FAAM Demonstration	84
5.6.2	Training the FAAM	86
5.7	Application to a Single Patient	87
5.7.1	Single Image Response	87
5.7.2	Non-Training Image Response	87
5.8	Analysis	89
5.9	Multiple Patients	90
5.9.1	Normalization	90
5.9.2	Training and Recall	90
5.9.3	Analysis	92
5.10	Discussion	92
5.11	Enhancements and Conclusions	95
6	Comparisons and Conclusions	100
6.1	Gold Standards and Image Processing	101
6.2	Volumetrics for IPF images	101
6.3	PFT Measurements	103
6.4	Materials and Methods	106
6.4.1	Statistical Analysis	106
6.5	Results	106
6.6	Discussion	111
7	Summary	112
7.1	Previous Approaches	112
7.2	Proposed Algorithms	113
7.2.1	Adaptive multiple Feature Method	113
7.2.2	Neural Networks	113
7.2.3	Gabor Filter	113
7.3	PCNN and IPF	114
7.4	Volumetric Analysis and Gold standards	114
7.5	Conclusion	115
	Bibliography	116

List of Figures

Figure	Page	
1.1	Classification of interstitial lung disease (ILD). COP, cryptogenic organizing pneumonia; CTD, connective tissue disease; HPS, Hermansky-Pudlak syndrome; IBD, inflammatory bowel disease; IIP, idiopathic interstitial pneumonia; IPF, idiopathic pulmonary fibrosis; LAM, lymphangioleiomyomatosis; LIP, lymphocytic interstitial pneumonia; NSIP, nonspecific interstitial pneumonitis; PAP, pulmonary alveolar proteinosis [1].	2
1.2	Patient with advanced IPF. Chest radiograph shows extensive bilateral reticular abnormality, more marked on the right than on the left. The images show peripheral and basal reticular abnormality, honeycombing and traction bronchiectasis.	2
1.3	Subpleural, lower lobe honeycombing and traction bronchiectasis.	3
1.4	IPF incidence is comparable to that of many other serious diseases.	4
1.5	HRCT scans of healthy subject and patient with IPF. Details such as the smaller bronchi are easier to appreciate on the HRCT scan.	7
3.1	Paradigm of Image Analysis with associated structure of Knowledge-based models that will be used at different stages of processing.	18
4.1	Characteristic honeycombing seen in IPF.	21
4.2	The value of element B(2,4) was calculated from A(2,4). The Range Filter uses a 3-by-3 neighborhood and the neighborhoods or different shapes and sizes can be specified.	23
4.3	The structuring element SE creates a flat, disk-shaped structuring element, where R specifies the radius. R must be a nonnegative integer, in this case a value = 3.	23
4.4	Honeycomb Detection using Range Filter.	24
4.5	Range filter with line element and Range filter with pair element.	25
4.6	Original Image.	27
4.7	Adaptive filter with Exponential distribution.	27

4.8	Original Image.	28
4.9	Adaptive filter with Rayleigh distribution.	29
4.10	A CT density histogram from a fibrotic lung is less skewed, less kurtotic and have increased mean lung attenuation compared to healthy lungs.	32
4.11	CT density histogram from normal lung shows a sharp peak and is skewed to the left in comparison to the Gaussian normal distribution.	33
4.12	Block diagram of the Adaptive Multistage Nonlinear Filter (AMNF).	34
4.13	Texture quantified using texture metrics such as statistical moments (variance, kurtosis) and co-occurrence matrices.	37
4.14	A basic computational neural element or Perceptron for classification.	38
4.15	Feedforward backpropagation network with one hidden layer.	39
4.16	Gabor filter composition for 1D signals:(a) sinusoid, (b) a Gaussian kernel, (c) the corresponding Gabor filter.	43
4.17	Gabor filter composition: (a) 2D sinusoid, (b) a Gaussian kernel, (c) the corresponding Gabor filter.	44
4.18	Example of Gabor filter oriented at 66° with the x -axis	45
4.19	Gabor filters with different frequencies and orientations.	46
4.20	Segmentation result after Gabor filter is applied.	47
5.1	Schematic representation of a PCNN processing element [2].	50
5.2	An input and selected pulse images from the PCNN [2].	52
5.3	An example of the progression of the states of a single neuron. See the text for explanation of L , U , T and F [2].	53
5.4	A typical PCNN example.	54
5.5	Collective behavior of F [2].	57
5.6	Collective behavior of L [2].	58
5.7	Collective behavior of Θ [2].	59
5.8	An example input [2].	60
5.9	Output pulse images [2].	61
5.10	Outputs of a PCNN with an asymmetric kernel, as discussed in the text. These outputs should be compared to those shown in Figure 5.11 [2].	61
5.11	Outputs of a PCNN with an on-center/off-surround kernel [2].	63
5.12	An input image [2].	64

5.13	Neuron that fired in the first 10 iterations for systems with $P = 1$, $P = 2$, and $P = 3$ [2].	64
5.14	An input image and a few of the pulse outputs from the ICM [2].	66
5.15	An original image and several selected pulse images from the PCNN [2]. . .	67
5.16	Results from the ICM [2].	67
5.17	A target (flower) pasted on a background [2].	69
5.18	The signature of the flower without a background (G.plt) and the signature of the flower with a background (Gb.plt) [2].	69
5.19	The propagation of curvature flow boundaries [2].	70
5.20	The progression of an autowave from an initial shape [2].	71
5.21	The signatures of the flower and the flower with a background using the centripetal autowave model [2].	74
5.22	PCA on extracted vectors from Gabor filter correlations.	75
5.23	PCA on extracted vectors from pulse images.	76
5.24	Original image of a patient with IPF.	77
5.25	Overview of the system with a training track (upper path) and a query track (lower path).	78
5.26	Two pulse images from the first cycle.	79
5.27	Four pulse images from the second cycle.	81
5.28	Four pulse images from the third cycle.	82
5.29	The x -axis represents the ICM iteration number, n , and the y -axis represents the pulse activity for the neurons	83
5.30	Demonstration of the FAAM.	85
5.31	Learning curves from a typical FAAM for three different regions in the same patient.	86
5.32	The response from image 14 from both FAAMs.	88
5.33	The response from image 11 from both FAAMs.	88
5.34	Area metrics of classification. The top most graph depicts the total area of the lung for each slice. The lower most graph depicts the area classified as fibrotic and the graph above it depicts the area classified as healthy. The second plot depicts the summation of the healthy and diseased regions. . . .	89
5.35	Image and Histogram before normalization.	91
5.36	The image is divided into different regions(back1, back2, organ, tissue, bone, vascular, diseased and healthy).	92

5.37	Training scans from multiple patients.	93
5.38	Demonstration of the FAAM from multiple patients.	94
5.39	Honeycombing and Vascular regions from Patient a0017.	96
5.40	Honeycombing and Vascular regions from Patient d0034.	97
5.41	Honeycombing and Vascular regions.	98
5.42	Honeycombing and Vascular regions from Patient p0023	99
6.1	Diagnostic approach to IPF.	100
6.2	Area metrics of classification. The top most graph depicts the total area of the lung for each slice. The graph in red depicts the area classified as fibrotic, green depicts vascular and the graph in blue depicts the area classified as healthy.	102
6.3	Lung volumes and capacities are depicted on a volume-time graph.	105
6.4	Correlation graphs between sum of vascular and healthy (tot) against PFT measures.	109
6.5	Correlation graphs between honeycomb regions (pcthon) and PFT measures.	110

List of Tables

Table	Page
6.1 Volumetric Percentage by Patient.	103
6.2 Pulmonary function measurements in patients with IPF.	105
6.3 Correlation Coefficients (r) and p -value in patients with IPF. Note: pctvas, pcthealthy, pcthon: Volumetric measures computed from regional classification for vascular, healthy and honeycomb regions respectively, tot = vascular+healthy measure of the lung.	107
6.4 Coefficients of correlation (r -values) between volumetric measures and PFT measures. Note: pctvas, pcthealthy, pcthon: Volumetric measures computed from regional classification for vascular, healthy and honeycomb regions respectively, tot = vascular+healthy measure of the lung.	107

Abstract

IMAGE ANALYSIS OF RADIOLOGICAL IMAGES FROM PATIENTS WITH ADVANCED LUNG DISEASE

Mekhala Acharya, Ph.D.

George Mason University, 2012

Dissertation Director: Dr. Jason Kinser

There are diagnostic challenges in the evaluation of a patient with a known or suspected diagnosis of interstitial lung disease [ILD] because of the extensive possibilities of diverse potential diagnoses with similar symptoms. High-Resolution Computed Tomography (HRCT) has changed the diagnostic evaluation of patients with ILD and is particularly useful in the diagnosis of idiopathic pulmonary fibrosis (IPF) [3]. According to NIH, IPF is a disease in which over a period of time the lung tissue becomes thickened, stiff and scarred. The characteristic HRCT findings of IPF are reticular abnormality and honeycombing with basal and peripheral predominance [4] and the radiographic pattern differs with the stage of the disease.

The quantification of disease by CT is important to indicate prognosis and to evaluate progression of the disease or response to treatment. It is difficult to convey the complex textural information offered by a CT scan hindered by the lack of user friendly technology for image analysis. Automated tools are presented which extract information from the CT images and isolate visual evidence of the disease from healthy lung tissue. Each CT image is converted to a set of pulse images, which through collective synchronization of pixels extract pertinent information of the diseased regions.

In spite of the obvious difference in contrast, volume and texture, the healthy and diseased regions are distinguished and classified using pulse images. The technique used is successful in classifying the healthy and diseased portions of the lung. The goal is to train adequate and varied stages of IPF images and to be able to extract sufficiently enough information from test images. The algorithm was tested on HRCT scans procured through Inova Fairfax Hospital, Department of Radiology. Two expert radiological reviewers compared the initial results of the segmentation algorithm with the manual segmentation of the original scans. Comparison revealed agreement regarding the presence or absence of honeycombing. Algorithms and results for the analysis of patients with IPF and healthy patients are presented.

The absence of gold standards in image processing makes quantification challenging for early stage images of IPF and blinded images. Thus medical image processing validation often cannot rely on availability of true gold standards. Hence lung volumes derived from Pulmonary Function Tests (PFT) results served as established clinical parameters and were used as "gold standards" [5]. The results of the segmentation were compared with the measurements of the pulmonary function tests. The relationship between image segmentation results and the PFT results were calculated using linear regression analysis and Pearson's product moment correlation. Volumetric measurements of honeycomb, vascular and normal regions are found to correlate with results of PFTs in patients with IPF. The greatest correlation was between honeycomb regions and forced vital capacity (FVC). The healthy and honeycomb regions correlated negatively with PFT measure diffusing capacity (DLco). Results demonstrate that the segmentation of IPF images using PCNN techniques are useful in extracting quantitative information.

Keywords: Idiopathic pulmonary fibrosis, high resolution computed tomography, Pulse coupled neural network, Intersecting cortical model, Fast analog associative memory

Chapter 1: Introduction

Many acute and chronic lung diseases with variable degrees of pulmonary inflammation and fibrosis are collectively referred to as interstitial lung diseases (ILDs) or diffuse parenchymal lung diseases [1]. Idiopathic pulmonary fibrosis (IPF), also known as cryptogenic fibrosing alveolitis, is a distinct interstitial lung disease that is limited to the lung, has an unknown etiology, and is associated with a histological pattern of usual interstitial pneumonia [3,6]. IPF is the most common ILD accounting for 25 to 50 percent of the diffuse lung diseases [7]. Figure 1.1 gives the current classification of ILDs according to the American Thoracic Society [1].

1.1 Diagnosis and Treatment

The diagnosis of IPF is based on clinical, radiographic and histopathologic evaluations [1]. IPF often presents clinically with progressive dyspnea, restricted lung function, and impaired gas exchange. A patchy pattern of peripheral, sub-pleural and predominantly lower lobe reticular opacities combined with honeycombing, traction bronchiectasis, and the absence of significant ground glass together constitute the classic radiographic features of IPF as seen in Figure 1.2 and Figure 1.3. Another feature of the disease is its pathologic heterogeneity within the lungs. In the absence of clinical and radiographic features, a surgical biopsy is recommended for definitive diagnosis of IPF [8].

1.1.1 Differential Diagnosis

Usual interstitial pneumonia (UIP) is the histopathological pattern that identifies patients with IPF [1]. In the absence of surgical lung biopsy, the diagnosis of IPF remains uncertain. Once diagnosed, IPF carries a bleak prognosis, with 5-year survival rates estimated at 30

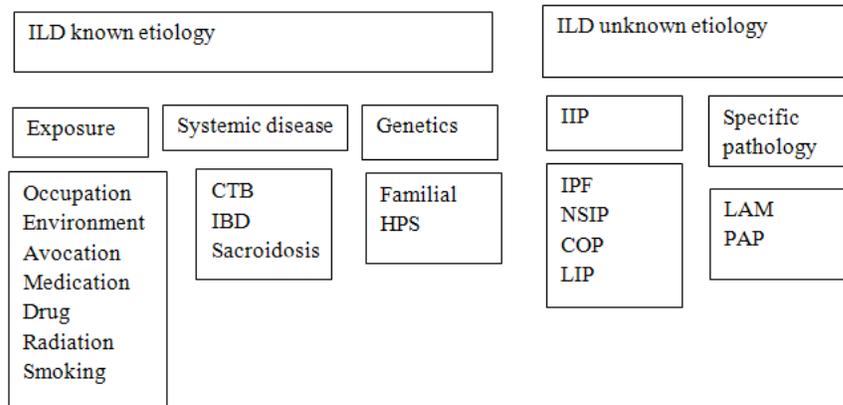


Figure 1.1: Classification of interstitial lung disease (ILD). COP, cryptogenic organizing pneumonia; CTD, connective tissue disease; HPS, Hermansky-Pudlak syndrome; IBD, inflammatory bowel disease; IIP, idiopathic interstitial pneumonia; IPF, idiopathic pulmonary fibrosis; LAM, lymphangiomyomatosis; LIP, lymphocytic interstitial pneumonia; NSIP, nonspecific interstitial pneumonitis; PAP, pulmonary alveolar proteinosis [1].

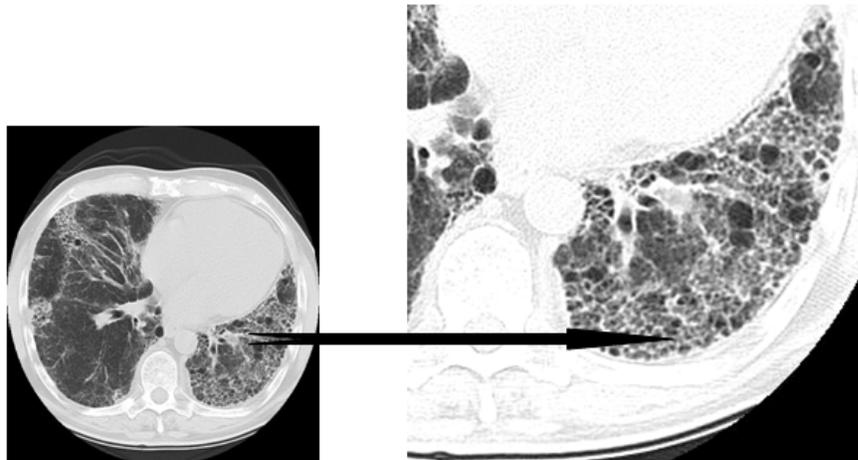


Figure 1.2: Patient with advanced IPF. Chest radiograph shows extensive bilateral reticular abnormality, more marked on the right than on the left. The images show peripheral and basal reticular abnormality, honeycombing and traction bronchiectasis.

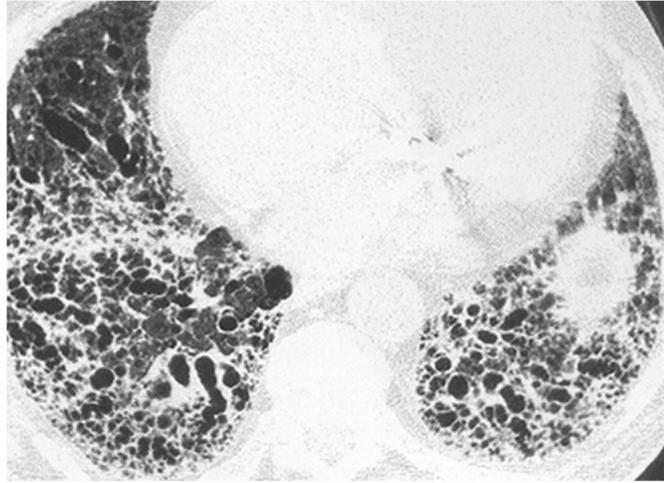


Figure 1.3: Subpleural, lower lobe honeycombing and traction bronchiectasis.

to 50 percent [1]. The course of IPF is unpredictable and patients with well-maintained lung function can be at significant risk of death. Some patients can have rapid progression of their disease and in some such instances this can be due to an acute exacerbation, which is itself a poorly characterized complication of the disease. There are no proven effective medical therapies for IPF. The only recourse for patients is lung transplant, however only a few patients with the disease are appropriate candidates for this.

There are diagnostic challenges in the evaluation of a patient with a known or suspected diagnosis of diffuse infiltrative lung disease (interstitial lung disease [ILD]) because of the extensive list of diverse potential diagnoses. The general availability of high-resolution computerized tomography (HRCT) has allowed this imaging technique to become a standard tool for the evaluation of patients with diffuse lung disease. Therefore it is necessary to develop tools to extract information from these images which are useful in narrowing the broad differential diagnosis.

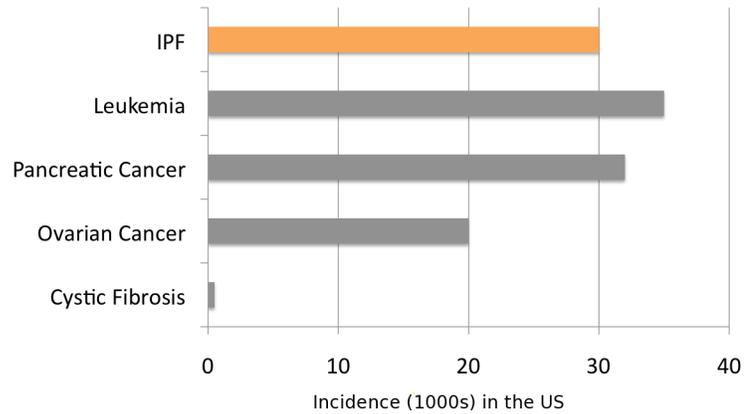


Figure 1.4: IPF incidence is comparable to that of many other serious diseases.

1.1.2 Epidemiology

Surprisingly little is known about the epidemiology of IPF in the United States. Once considered a relatively rare disease, IPF is now recognized as the most common interstitial lung disease [9]. The incidence of IPF is greater than that of ovarian cancer, similar to those of pancreatic cancer and of all leukemias combined, and nearly 30 times that of cystic fibrosis [10–12]. IPF is a fatal disease with an average 5-year mortality higher than that of many difficult-to-treat cancers [7, 10, 13]. Disease awareness is among the critical challenges for patients with IPF, their caregivers, and their healthcare teams. Greater disease awareness can facilitate earlier diagnosis and frequent evaluation, both of which are recommended to optimize the management of IPF [1, 14]. The annual incidence of IPF in the United States is over 30,000 [11]. IPF occurs more commonly in males than in females. The US prevalence is approximately 83,000 [11]. 73% of patients with IPF are greater than 65 years of age [11].

1.2 Current Challenges in Medical Imaging Modalities of Lung Disease

Medical imaging systems include a wide variety of technologies, methods and applications designed to provide information about the structure and function of living beings. The term is used most frequently in reference to modalities associated with diagnostic radiology, for example computed tomography (CT), ultrasound, projection radiography and magnetic resonance imaging (MRI). It can, however, also be used more generically to include systems that obtain images of other types, including those visualized during clinical procedures such as endoscopy or derived from pathologic specimens.

1.2.1 Chest Radiography

Conventional chest radiography is often the first or only imaging procedure performed in the patient with suspected or known lung disease. When a confident diagnosis of IPF is made on the basis of the chest radiograph, it is correct in 48 to 87 percent of cases [15,16]. Given the indolent course of IPF, the clinical utility and optimal timing of follow-up chest radiographs are unclear. Limitations include the fact that it is relatively insensitive for early ILD and gives little if any useful information regarding the severity of inflammation in ILD.

1.2.2 HRCT Technique

HRCT allows detailed evaluation of the lung parenchyma by using 1 to 2 mm thick slices which images only 10 percent of the lung parenchyma [3] with a reconstruction algorithm that maximizes spatial resolution [1]. This has resulted in more accurate differential diagnosis and determination of extent of fibrosis. The accuracy of a confident diagnosis of UIP made on HRCT by a trained observer appears to be about 90 percent [16–18]. However because a confident diagnosis of IPF is made by HRCT in two thirds of patients with IPF, one-third of the cases are missed by relying on HRCT alone [19]. Less experienced observers

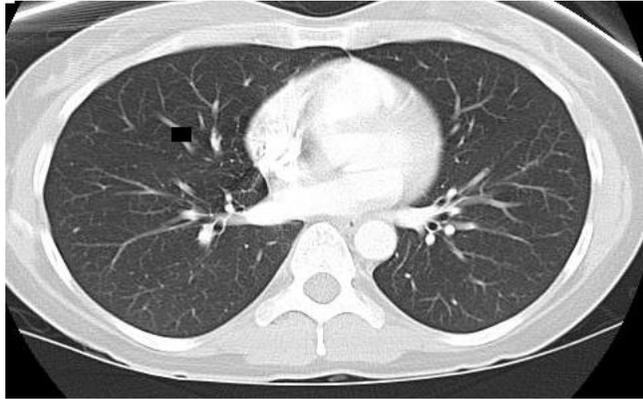
are substantially less accurate than experienced observers [17]. But the quantitation is also challenging for several reasons. It is difficult and potentially misleading to use a single number to convey the complex textural information offered by the CT scan. There is moderate inter-observer agreement in scoring pattern type and disease extent in patients with IPF [20]. All existing methods of CT quantification of ILD are dependent on patient's inspiration, which determines the volume of the lung and lung density. Spirometric control of the level of inspiration on CT is the logical solution and not widely used in the United States [21]. A significant problem with quantitative CT techniques is the lack of an automated method for segmenting the lung and deriving quantitative indices. In spite of these drawbacks, HRCT features alone are sufficient to make the diagnosis of UIP/IPF in 50 to 60 percent of cases selected based on the clinical suspicion of IPF. As seen in Figure 1.5 there is a stark contrast between the HRCT images of the healthy subject as compared to the HRCT scan from the IPF patient. IPF patient shows extensive basal reticular abnormality compatible with lung fibrosis.

1.2.3 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) has the ability to provide a combination of high resolution, excellent soft tissue contrast and high signal to noise ratio [22]. Its an excellent diagnostic tool allowing for morphological evaluation of lung parenchyma and cardiac function. However evaluation of lung parenchyma using MRI is not yet common in patients with severe IPF and with reduced parenchymal structure. The role of MRI for the assessment of interstitial, inflammatory and obstructive disease is not yet determined. In general MRI is less available, will take more time and will cost more than CT [23].

1.2.4 Ultrasound

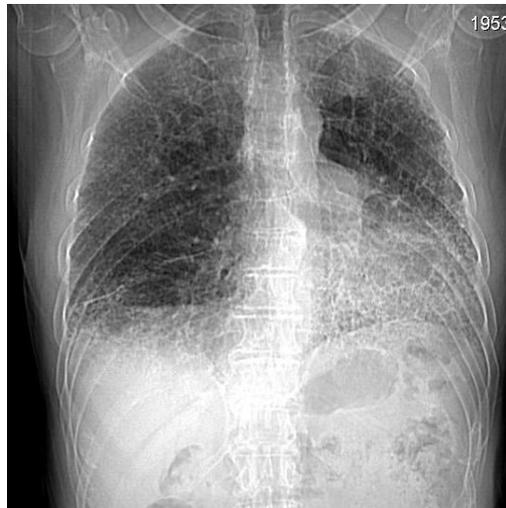
The value of Thoracic Sonography (TS) in ILD was studied by Angelika Ressig [24]. TS is cost-effective, non-invasive, broadly available method that requires neither ionizing radiation nor a contrast medium. However HRCT is the most sensitive and specific radiologic



(a) Healthy Subject



(b) IPF Patient.



(c) IPF Patient.

Figure 1.5: HRCT scans of healthy subject and patient with IPF. Details such as the smaller bronchi are easier to appreciate on the HRCT scan.

technique in comparison. TS is thus considered as a complementary method.

1.3 Disease Quantification using HRCT scans

Quantification of disease by CT is important to indicate prognosis and to evaluate progression of the disease or response to treatment. CT quantification may be visual or digital. Most studies of diffuse lung disease use a visual approach asking observers to assign a numerical grade to the severity of the disease on a four or five point scale [20,25,26]. However such simple scoring systems may not have the resolution to detect subtle changes in extent of disease on serial scans. A wide range of inter-observer and intra-observer correlations (r values) are reported indicating substantial variation in visual assessment scores. Digital methods of quantification of disease extent rely on a standard scale of CT attenuation, measured in Hounsfield unit (HU). The mean attenuation of normal lung is around -820 to -860 HU, on HRCT, but this is critically dependent on the volume of inspired air [27]. For digital quantification, the lung must be first isolated (segmented) from the surrounding structure.

The advantages of visual quantification include ease of performance, moderate reproducibility with experience and training and sensitivity to patterns of disease. Advantages of digital quantification include observer independence, excellent reproducibility and potentially greater sensitivity to subtle changes [4].

Chapter 2: Previous Approaches to the Study of IPF

Precise clinical measurement is an important part of the science of medicine. However radiologists have been slow to provide quantitative measurements of disease extent. In the management of patients with diffuse lung diseases, it is important to define the extent of the disease, particularly if the patient is being monitored for progression of disease or response to treatment. In this chapter, some of the earlier approaches in the quantitative analysis of IPF are explored.

2.1 Quantification of Diffuse Lung Diseases

Quantitative techniques have an important role in improving our understanding of the pathophysiology of lung diseases. Physiologic evaluation provides a global measurement of lung function, including contributions from normal and abnormal lung parenchyma.

Quantitation of radiologic abnormalities is challenging for several reasons. Radiology is a descriptive visual speciality in which the subtleties of image appearance are important. It is difficult and potentially misleading to use a single number to convey the complex textural information offered by a chest radiograph or the chest CT scan. Visual integration of the extent of abnormality is a complex task subject to substantial inter observer variation. Use of digital data from chest radiography and chest CT has been hindered by the lack of user-friendly technology for image analysis and by the difficulty of validating the system used.

2.2 CT Density

A number of authors have used measures of CT density to evaluate the extent of diffuse parenchymal lung diseases such as IPF [28–35]. CT scans were obtained from IPF and

control patients and lung volumes were estimated from measurements of voxel size, and X-ray attenuation values of each voxel. Quantitative estimates of lung structure were obtained from biopsies obtained from diseased and normal CT regions using stereologic methods. CT density was used to calculate the proportion of tissue and air, and this value was used to correct the biopsy specimens to the level of inflation during the CT scan. The data shows that IPF is associated with a reduction in airspace volume with no change in tissue volume or weight compared with control lungs. Thus Coxson et al. [28] concluded that the ratio of air to tissue in the lung, estimated from CT density, correlated with that found in histopathology.

Wollmer et al. [29] found a significant increase in CT lung density in 33 asbestos workers, only 3 of whom had evidence of asbestosis in chest radiographs. In excess of detailed morphologic information, CT provides good estimates of regional lung density [30,31]. A paper by Reuter et al. [32] demonstrated that subjects with asbestos exposure and normal or near-normal chest radiographs had substantially higher lung density than age-matched normal subjects reinforcing the concept that quantitative CT may be able to detect lung disease earlier than analysis of CT images.

Lynch et al. [33] used quantitative CT indexes such as skewness, kurtosis and mean lung attenuation to correlate IPF with pulmonary abnormality. One hundred and forty four subjects participated in a double blind placebo controlled study and strong correlation was found between thin-section CT histograms with pulmonary function tests of patients with IPF.

2.2.1 Disadvantages

However all existing methods of CT quantification of ILD are less than ideal. Visual estimation techniques are subjective and show only moderate reproducibility. All available techniques are dependent on the depth of the patient's inspiration, which determines the overall volume of the lung and the lung density. Quantitative methods based on CT density histogram are particularly susceptible to variation with differing depths of inspiration.

Spirometric control of the level of inspiration on CT is the logical solution [34]; though available in Europe, it is not widely used in the United States. The optimal lung volume for the measurement of lung density in patients with lung fibrosis is unclear. A study of normal subjects by Verschakelen et al. [35] showed that scanning at low volumes accentuated the problem of increased density in the dependent lung. This might potentially obscure density changes due to early lung fibrosis. However a study by Beinert et al. [36] showed that optimal discrimination between patients with lung fibrosis and normals occurred at a level of inspiration of 20% or 50% of vital capacity. The impact of the reconstruction algorithm on density measurements is unclear.

A significant problem with quantitative CT technology is the lack of user-friendly automated method for segmenting the lung and deriving quantitative indices. It seems particularly surprising that none of the major CT manufacturers provide a user-friendly package for performing quantitative analysis of lung images [4].

2.3 Texture Analysis

Texture based methods examine the relationship between the attenuation of pixels within a certain distance of each other. Previously published computer-based methods for evaluating CT images of the lung [37–42] have used information from only one first-order textural feature such as density and therefore do not take full advantage of the complexity of the lung parenchyma as shown and recorded in digital form by HRCT.

Uppaluri et al. [43] developed an adaptive multiple feature method (AMFM) to assess 22 independent texture features to classify six different tissue patterns in a CT image. The features with the greatest ability to discriminate between the different patterns were chosen and a Bayesian classifier was evaluated on six images and manually compared by three radiologists. The method used an objective assessment of global and regional changes in pulmonary parenchyma to detect emphysema. This computerized method used a combination of statistical and fractal texture features for characterization of lung tissues based upon high resolution computed tomography (HRCT) scans. The AMFM is a texture-based method

that combines statistical texture measures and a fractal measure. In total, 17 measures of texture were used. Included in these measures were the grey level distribution measures, run-length measures, co-occurrence matrix measures, and a geometric fractal dimension (GFD).

A feature selection program based on the divergence measure was used to select an optimal subset of features to best discriminate the tissues under consideration. A database of HRCT slices of known classification, called the training set, was used to train a Bayesian classifier. A test set of HRCT slices, distinct from the training set, was used to compute the accuracy of the method in tissue characterization.

HRCT images obtained from normal subjects and those with emphysema, IPF, and sarcoidosis were compared in a global analysis. This analysis was performed using the whole single image slice of the lung field as the region of interest (ROI) to extract the texture features. Two subject groups, three subject groups, and all four of the subject groups were compared. Regional analysis was then performed allowing for comparisons of normal and diseased lungs using the AMFM. Although this method performed well, it was found AMFM did better with global disease categories rather with basic lung patterns such as honeycombing, ground glass, nodular, etc.

Uchiyama et al. [44] proposed a scheme in which the lung was segmented and divided into 32×32 contiguous regions of interest (ROI). An artificial neural network was used to classify each ROI. In a study by Delorme et al., a similar approach was undertaken wherein they trained a computer system to establish rules for recognition of image features such as normal parenchyma ground glass attenuation and fibrosis [45].

In his paper, Uchiyama et al. [44] used a database consisting of HRCT images from patients which included normal and abnormal slices related to six different patterns, i.e., ground-glass opacities, reticular and linear opacities, nodular opacities, honeycombing, emphysematous change, and consolidation. The areas with a specific pattern, which three radiologists marked independently and consistently as the same patterns, were used as gold standard for specific abnormal opacities in this study. The lungs were first segmented from

the background in each slice by use of a morphological filter and a thresholding technique, and then divided into many contiguous ROIs with a 32×32 matrix.

Six physical measurements which were determined in each ROI included the mean and the standard deviation of the CT value, air density components, nodular components, line components, and multilocular components. Artificial neural networks (ANNs) were employed for distinguishing between seven different patterns which included normals and six patterns associated with diffuse lung disease. Uchiyama et al. [44] studied the use of CAD and reported a sensitivity of 99.2% for identifying any abnormal lung patterns (ground-glass opacities, reticular or linear opacities, nodular opacities, honeycombing, emphysematous changes, or consolidation) . Although these early results suggested that CAD may eventually be able to assist radiologists in their assessment of diffuse lung disease, the necessary software tools are still in the theoretical development stage.

Fractal analysis have also been applied to CT patterns [46]. They computed generalized fractal dimensions for HRCT images to investigate their value in the discrimination and quantification of idiopathic pulmonary fibrosis (IPF) from normal lung parenchyma. The probability distribution that was based on the pixel value in each image was used to compute capacity, information, and higher fractal dimensions for a series of 52 HRCT slices obtained from four patients. Qualitative classification of normal, mild, moderate, and severe IPF cases was achieved. The generalizability of the model was tested by predicting the extent of IPF for each patient from a regression model computed with the remaining slices in the database.

The extent of IPF was predicted well within the 90% confidence interval given by the model. Rodriguez et al. [46] also showed that the automatic discrimination and quantification of IPF from normal lung parenchyma was possible using generalized fractal dimensions on HRCT images. However, they did not attempt to compare the results of the segmentation with measurements such as pulmonary function tests, which is important for the treatment option and the patient prognosis.

Clustering techniques based on first order histogram analysis have been applied by

researchers. In [47], researchers have applied rule based scheme for chest radiographic image analysis in combination with artificial neural network to detect and classify interstitial abnormalities.

Katsuragawa et al. [47] describe a computerized method to quantify ILD by using physical texture measures obtained from an analysis of power spectrum of lung textures in digital chest radiographs. He and his colleagues used Fourier transformation to determine the frequency power spectrum of the radiographic texture of the lung on the ILO standard radiographs [48]. The most detailed and most widely used system for quantification of diffuse lung diseases on chest radiography is the International Labour Organization (ILO) classification system for pneumoconiosis.

They found an increasing profusion of opacities was associated with increased root mean square variation in the pixel values, while an increasing size of opacities was associated with a shift of the first moment of the power spectrum toward lower frequencies. The use of these two measures can, therefore, provide a measure of the profusion and size of the radiographic opacities. These authors have progressively enhanced their analysis algorithm, which has resulted in substantial decrease in false-positive detection of interstitial abnormalities. However the system has not been used to quantify the extent of abnormality in subjects with diffuse lung disease.

Chapter 3: Research Rationale and Data Collection

The goal of this study is to quantify features that appear in radiological images, specifically chest HRCT scans, and to measure their inter relationships. The aim is to understand and measure the features that exist in radiological images. Of specific interest is the development of algorithms to detect evidence of early onset of IPF. In this regard it is the peripheral area of the lung and mostly the bases that are first involved in the disease process than the more central areas.

The ideal quantitative technique would be simple to perform, reproducible and observer independent and offer a valid measurement of disease extent, by comparison with physiologic or pathologic indices. The purpose of this dissertation is to use these criteria to evaluate methods for scoring abnormalities.

3.1 Objective

This goal will be addressed through several major steps.

1. Feature Identification.

Several researchers have identified the visible features of IPF through a manual process. At first these features will be collected and parsed manually.

2. Feature Extraction.

Segmentation is important for feature extraction. The first step in automatic disease pattern detection is to segment the lungs from the background [49]. Next, through a manual down selection, features deemed significant will be isolated using appropriate algorithms. Concurrently, methods will be developed to measure the strength of the feature (size, intensity, etc) and to measure the confidence of detection.

3. Feature Significance

Once the image has been segmented, shape and texture quantification techniques will be applied to address different aspects of the data. The extracted features will be compared to predefined disease patterns. In this step the ability to extract features of significance will be reviewed.

4. Discrimination Ability of Features.

In this step, first-order algorithms of discrimination will be applied to the features extracted from the images. The goal is to determine the ability of each feature to be associated with diagnosis. It is expected that some features will be strongly associated with IPF images and others will have a weaker association. These associations will be measured and the ability of individual features will be quantified.

5. Onset Detection.

In this step the features on early onset images will be compared to normal and later term images. Features from previous steps will be considered and measured. The goal is to identify those features which are evident in early onset images.

6. Higher Order Considerations.

In this step combinations of features will be considered in higher order fashions. The goal is to detect the interplay of different data domains within this set of images.

3.2 Medical Image Analysis

Figure 3.1 shows a paradigm of image analysis with associated structure of knowledge-based models that can be used at different stages of processing. The knowledge of physical constraints and tissue properties can be very useful in imaging and image reconstruction.

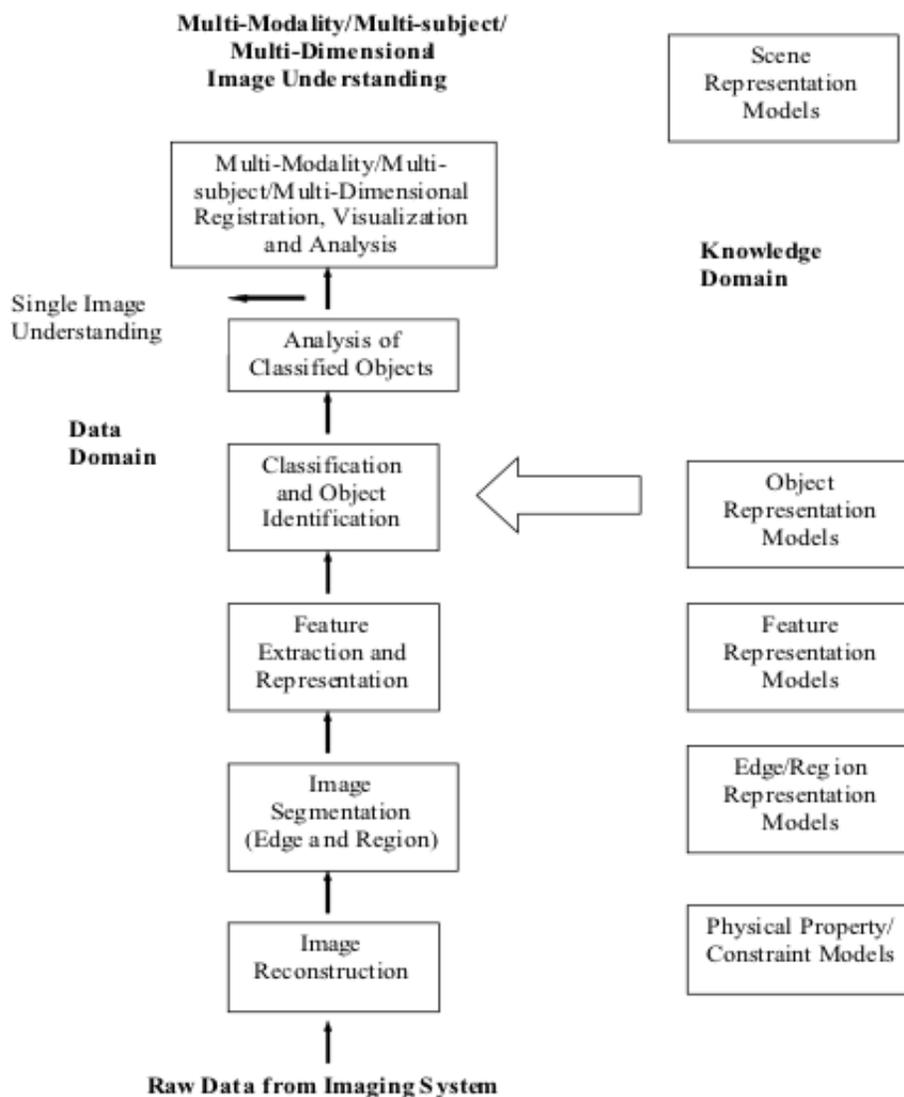


Figure 3.1: Paradigm of Image Analysis with associated structure of Knowledge-based models that will be used at different stages of processing.

3.3 Materials and Methods

This study was approved by the Institutional Review Board at George Mason University. The manuscript was reviewed and approved according to the procedures outlined by the

Lung Tissue Research Consortium. Further information on the Consortium is available on the website [www.ltrcrepublic.com].

3.3.1 Experimental Data-Sets

Radiological images have been provided by Inova. The principal investigator is Dr Steven D. Nathan, MD from the Advanced Lung Disease Program, Inova Heart and Lung Institute. The images have been classified into healthy and diseased. The study population is divided into normal, IPF, early onset and other lung diseases. Three CDs of images, two containing healthy images and one having IPF images have been provided at present. The IPF CD contains two studies with three time series each. Study 1 contains 110 HRCT scans and Study 2 contains 60 HRCT scans. The healthy CD contains 64 chest scans.

Since this is a retrospective study, a standardized imaging protocol was not utilized. Each patient included in the sampling (n=48) will have had one of three chest CT imaging protocols used depending on what was ordered for the particular clinical indication. The types are as follows:

- HRCT (high resolution CT without contrast)
- Routine Chest (with or without contrast)
- CT Angio (with contrast, rule out pulmonary embolism)

3.3.2 CT Scanning Protocol

All patients had HRCT studies of the thorax. Collimation was 1 mm for all but 1 study, which had a collimation of 1.5 mm. Images were acquired at 1- or 2-cm intervals. Expiratory images were obtained in 29 cases. All images were reconstructed with a high spatial frequency reconstruction algorithm and photographed at an appropriate window (1500 1750 Hounsfield units [HU]) and level (2500 to 2700 HU) settings. The average time between imaging and biopsy was 9 months.

Chapter 4: Research Methods

IPF is the most common ILD (Interstitial Lung Disease) accounting for 25 to 50 percent of the diffuse lung diseases. [7]. Quantization of IPF images is challenging because it is difficult and potentially misleading to use a single metric to convey the complex textural information offered by the CT scan. In this chapter three different approaches have been proposed. The corresponding analysis, results and drawbacks have also been presented.

4.1 Adaptive Multiple Feature Method

Texture based methods examined the relationship between the attenuation of pixels within a certain distance of each other. Uppaluri et al [50] developed an adaptive multiple feature method (AMFM) to assess 22 independent texture features to classify six different tissue patterns in a CT image. The features with the greatest ability to discriminate between the different patterns were chosen and a Bayesian classifier was evaluated on six images and manually compared by three radiologists.

A computationally efficient method is presented which automatically segments the ROI (region of interest) using a contrast enhancement function and adaptive thresholding. The goal is addressed through several major steps:

- threshold or feature identification,
- smoothing or noise reduction and
- ROI extraction.

4.2 Thresholding or Feature Identification

The scans contained low-contrast fuzzy regions. As a result the histogram modes corresponding to the different types of regions in the image overlapped making segmentation by thresholding difficult. Although applying a global threshold was computationally fast and simple, it failed because there was low contrast between the object and the back-ground.

Figure 4.1 shows a typical lung scan for a patient with IPF. The lung on the left shows a tremendous amount of a honeycomb structure. Honeycombing is defined as a cluster or row of cysts [4]. Honeycomb cysts are very small (less than 5 mm in diameter) and are often subpleural. As areas of honeycombing increases honeycomb cysts also increases in size in patients with IPF. Whereas the middle portion of the lung on the right shows a smoother texture indicative of healthier tissue.

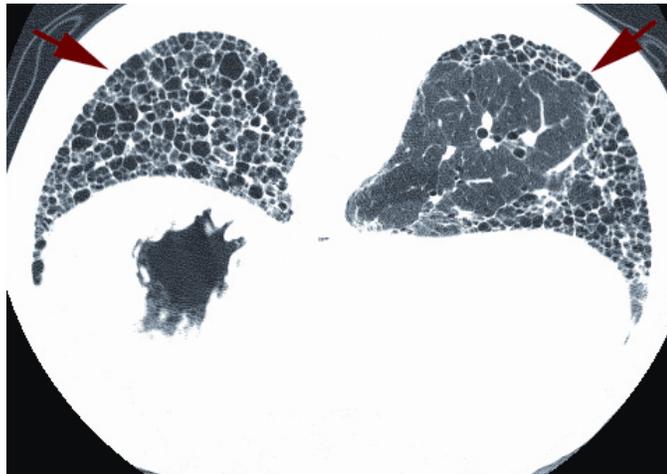


Figure 4.1: Characteristic honeycombing seen in IPF.

4.2.1 Range Filter

Local thresholds were determined by examining the image intensities in the neighborhood of each pixel. A texture segmentation algorithm was implemented using a range filter which returned an array of pixels. This was accomplished using a native Matlab function called

‘rangefilt’. The syntax is shown below:

```
J = rangefilt(I) and J = rangefilt(I, NHOOD)
```

`J = rangefilt(I)` returns the array `J`, where each output pixel contains the range value (maximum value - minimum value) of the 3×3 neighborhood around the corresponding pixel in the input image `I`. The image `I` can have any dimension. The output image `J` is the same size as the input image `I`.

`J = rangefilt(I, NHOOD)` performs range filtering of the input image `I` where user specifies the neighborhood in the parameter ‘`NHOOD`’. `NHOOD` is a multidimensional array of zeros and ones where the nonzero elements specify the neighborhood for the range filtering operation. `NHOOD`’s size must be odd in each dimension. By default, ‘rangefilt’ uses a 3-by-3 matrix of logical ones. ‘rangefilt’ determines the center element of the neighborhood.

Each output pixel contained the range value (maximum-minimum) of the 3×3 neighborhood around the corresponding pixel in the input image as shown in Figure 4.2. The range filter performed the dual step of thresholding and noise reduction. An automated gray-level remapping method for enhancement of details in chest region CT images were used by Davis et al [51]. The objective was to reduce dark line artifacts in CT images and to provide an aid to the automatic detection of the lung region. This formed the basis for adaptive thresholding.

In addition, to specify neighborhoods of various shapes, such as a disk, use the ‘strel’ function to create a structuring element object and then use the native Matlab function ‘getnhood’ method to extract the neighborhood from the structuring element object.

The filter used morphological operations of dilation and erosion to determine the maximum and minimum values in the specified neighborhood. Additional filters were applied and compared in order to select an effective and optimal smoothing filter for noise reduction. Three low-pass filters: pairs, disk and lines with the ability to smooth the image and reduce the noise were evaluated and the corresponding results are seen in Figures 4.4 and 4.5. The Range filter with disk element performed well in the dual step of thresholding and noise reduction. Once again Matlab native function ‘strel’ was used to create the disk

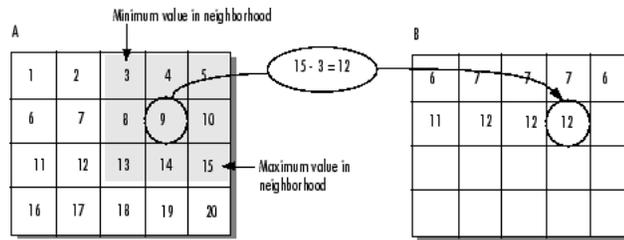


Figure 4.2: The value of element B(2,4) was calculated from A(2,4). The Range Filter uses a 3-by-3 neighborhood and the neighborhoods of different shapes and sizes can be specified.

element. The syntax is shown below.

`SE = strel('disk', R, N)` creates a flat, disk-shaped structuring element, where R specifies the radius as seen in Figure 4.3. R must be a nonnegative integer. N must be 0, 4, 6, or 8. When N is greater than 0, the disk-shaped structuring element is approximated by a sequence of N periodic-line structuring elements. When N equals 0, no approximation is used, and the structuring element members consist of all pixels whose centers are no greater than R away from the origin. If N is not specified, the default value is 4. The range filter performed the dual step of thresholding and noise reduction.

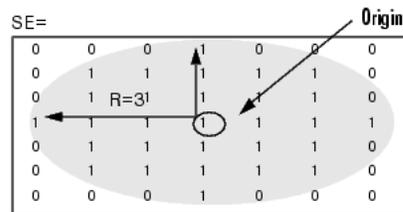
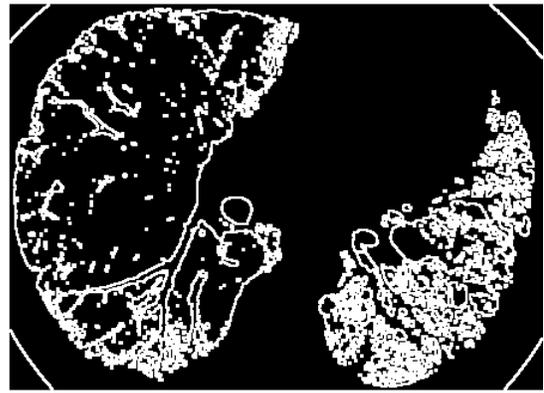


Figure 4.3: The structuring element SE creates a flat, disk-shaped structuring element, where R specifies the radius. R must be a nonnegative integer, in this case a value = 3.

Images were contrast enhanced using a Matlab function 'adapthiseq'. Next, adaptive thresholding was completed and texture metrics such as first order statistics, co-occurrence matrix and autocorrelation function to remove background inhomogeneities were applied.



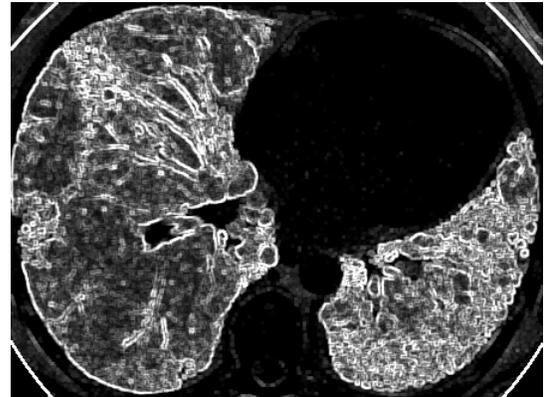
(a) Subject with Honeycombing



(b) After applying Range Filter with DISK element

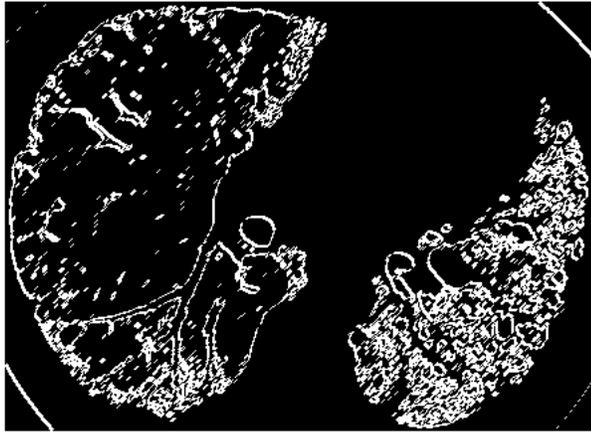


(c) Subject with Honeycombing

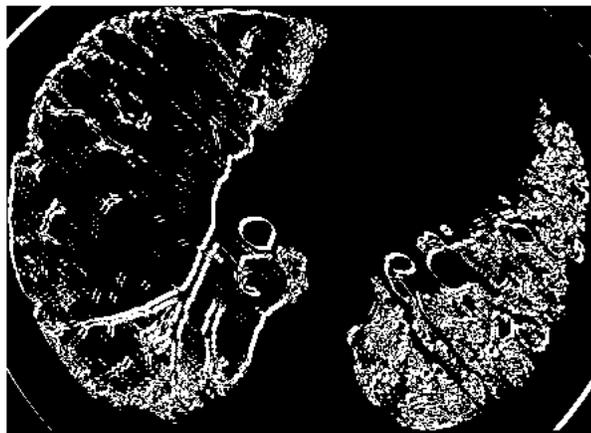


(d) After applying RangeFilter with DISK element

Figure 4.4: Honeycomb Detection using Range Filter.



(a) Line Element



(b) Pair Element

Figure 4.5: Range filter with line element and Range filter with pair element.

This is further explained in detail in the next section. Regions of interest (ROI) were delineated over which quantitative texture features are computed. Each pixel or region was assigned to one or several quantitative texture feature metrics. The information collected in the classification step was used to determine the category into which pixel or region belongs. Within a defined region, its texture was categorized as healthy or diseased.

4.3 Smoothing or Noise Reduction

Noise reduction is an essential part of image processing, and there are many algorithms for the noise reduction process. The noise reduction process can make images smooth; on the other hand it makes the edges blurred because it is hard to distinguish the edges from the noise. In smoothing, the data points of a signal are modified so individual points (presumably because of noise) are reduced, and points that are lower than the adjacent points are increased leading to a smoother signal.

There are several sources of noise which affect image quality. Some of which are due to inherent properties of the imaging device such as spatial resolution, energy resolution, non-uniformity, or distortions. Other degrading factors are dependent on the patient and organ localization. A large patient will increase the influence of scattered photons. An organ deep in the body will be overlapped by other tissues, which will increase the background registrations. Patient and organ movements will also degrade the image quality [52].

4.3.1 Adaptive Thresholding

Adaptive thresholding takes into account spatial variations in illumination. The goal is to develop a contrast enhancement function to enhance the IPF features while suppressing the noise. For this purpose an adaptive neighborhood structure is used. This is defined as set of two neighborhoods: inner and outer [53]. Depending on the disease extent, lung features are of arbitrary shape. Hence a variable shaped feature adaptive neighborhood criterion is needed that adapts the arbitrary shape and size of the local features to obtain the ‘center‘ and ‘surround‘ regions [53]. The center consists of pixels forming the feature

and the surround consists of pixels forming the background for that feature. These regions can be used to compute the local contrast for the centered pixel. The adaptive image-enhancement methods using first-order derivative information can be used to enhance specific IPF features such as honeycomb cysts and traction bronchiectasis. The method operates on small regions in the image, called 'tiles', rather than the entire image.

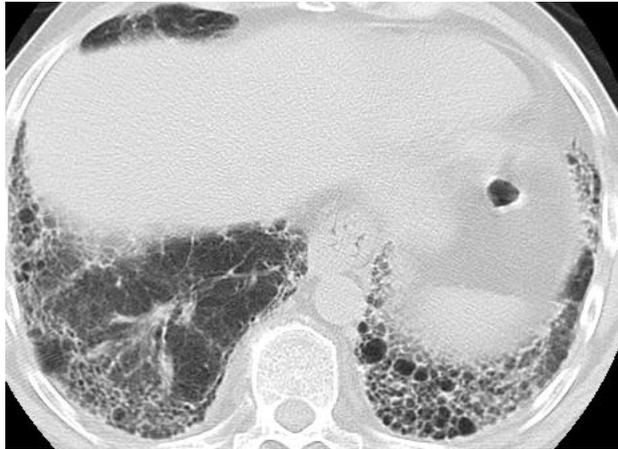


Figure 4.6: Original Image.

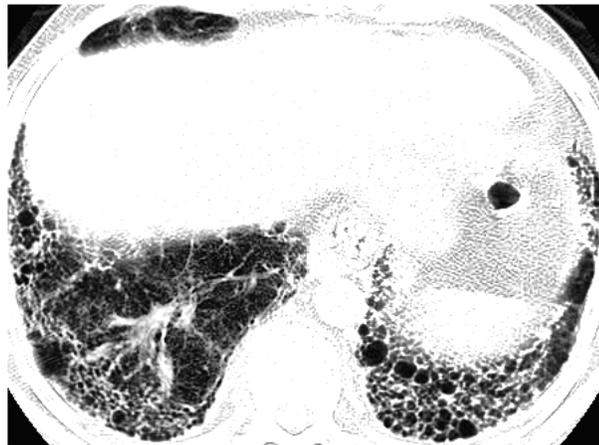


Figure 4.7: Adaptive filter with Exponential distribution.

The key step in the image processing chain proposed is the computation of local contrast information based on adaptive neighborhoods. For each pixel $I(x, y)$ the first criterion for classifying the neighborhood pixels as ‘center‘ or ‘surround‘ (feature and background) is gray value similarity. Any neighborhood pixel $I(x + \delta x, y + \delta Y)$ is classified as ‘center‘ if its intensity meets the condition:

$$I(x, y) - T \leq I(x + \delta x, y + \delta y) \leq I(x, y) + T \quad (4.1)$$

where T is the threshold value. Any neighborhood pixel is classified as ‘surround‘ if it does not meet the intensity criterion. Pixels outside the threshold window that are not connected to the center count as unclassified. For example if the window is 5×5 pixels, the only pixels counted are those for which $\delta x = \pm 2$ or $\delta y = \pm 2$.



Figure 4.8: Original Image.

Zuiderveld [54] described a contrast limited adaptive histogram equalization method (CLAHE) which operated on small regions in the image. A Matlab function ‘adaphiseq‘ enhances the contrast of the gray-scale image transforming the values using CLAHE. The method operates on small regions in the image, called ‘tiles‘, rather than the entire image. Each tile’s contrast is enhanced, so that the histogram of the output region approximately

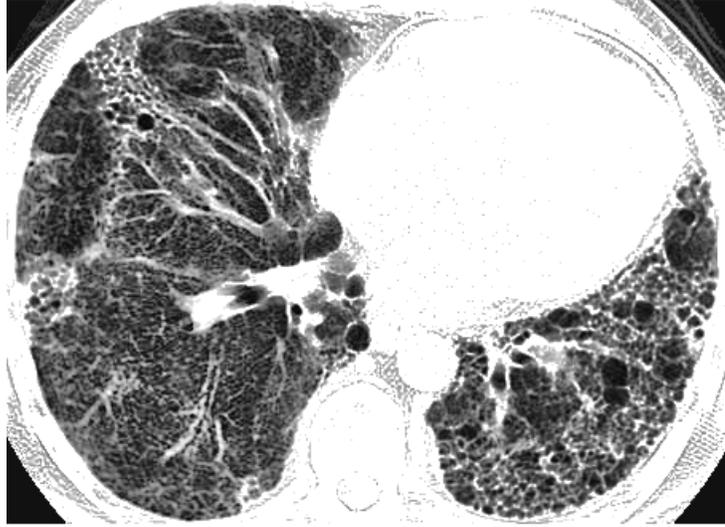


Figure 4.9: Adaptive filter with Rayleigh distribution.

matches the histogram specified by the ‘Distribution‘ parameter as seen in Figures 4.7 and 4.9. The ‘Distribution‘ parameter is one of the input parameters to the ‘adapthiseq‘ function which specifies the distribution that CLAHE uses as a basis for creating the contrast transform function.

The Figure 4.12 shows CLAHE applied for the whole image. An image $g(i, j)$ consists of two parts, low frequency $g_l(i, j)$ and high frequency $g_h(i, j)$ which can be expressed as

$$g(i, j) = g_l(i, j) + g_h(i, j). \quad (4.2)$$

The low-frequency part may be dominant in homogeneous regions, whereas the high-frequency part may be dominant in edge regions. The two-component model allows for different treatment of the components, and it can be used for adaptive image filtering and enhancement [55]. The high frequency part may be weighted with a signal dependent weighting factor to achieve enhancement. A two-component model is suitable not only for noise suppression, but also for many other enhancement operations such as statistical differencing and contrast enhancement. In adaptive multistage nonlinear filtering (AMNF),

the first stage includes multiple linear and non-linear filters to filter the input image. In the second stage, for each pixel, the output of only one filter is selected using an adaptive criterion.

The structure of AMNF is shown in Figure 4.12 and its output is:

$$g_{AMNF}(i, j) = g_{AF}(i, j) + b(i, j)[g(i, j) - g_{AF}(i, j)], \quad (4.3)$$

where $b(i, j)$ is a signal dependent weighting factor that is a measure of the local signal activity and $g_{AF}(i, j)$, is the output from the second stage of the filter in Figure 4.12. The value of $b(i, j)$ is obtained from local statistics around processed pixel as

$$b(i, j) = \sigma_f^2(i, j) / [\sigma_f^2(i, j) + \sigma_n^2(i, j)], \quad (4.4)$$

where σ_f^2 is the signal variance and σ_n^2 is the noise variance. In the flat regions of the input image, the signal-to-noise ratio is small, so $b(i, j)$ becomes small and $g_{AMNF}(i, j)$ approaches $g_{AF}(i, j)$. On the other hand, the signal-to-noise ratio is large around the edges in the image, so $b(i, j) \rightarrow 1$, and $g_{AMNF}(i, j) \rightarrow g(i, j)$. The operation of the filter therefore should preserve the edges in the image.

4.3.2 Quantitative CT indices

Concurrently, computer derived indexes such as mean lung attenuation, skewness and kurtosis were obtained from the frequency histogram. These are especially attractive because they clearly scored the disease extent in patients with IPF. The first four moments of texture (mean, variance, skewness, and kurtosis) do show promise in separating the two types of lung regions.

This is seen in the histogram analysis done for the healthy subject and the IPF patient. In healthy subjects the first order histogram of CT attenuation is sharply peaked (kurtotic) and skewed to the left in comparison to the Gaussian normal distribution as seen in Figure

4.11. Histogram from patients with IPF are less skewed, less kurtotic and have increased mean lung attenuation compared with those from patients with normal lungs as seen in Figure 4.10.

The histograms of the healthy and honeycomb regions overlap and therefore segmentation of the two regions by a simple threshold is not plausible. Also quantitative methods based on CT density histogram are particularly susceptible to variation with differing depths of inspiration [34]. Statistical moments do not contain spatial information because they are derived from the image histogram.

Mean, variance, skewness and kurtosis can be found using the equations given below.

$$\text{Mean } \mu_1 = 1/N \sum (x_i), \quad (4.5)$$

$$\text{Variance} = 1/N \sum (x_i - \mu_1)^2, \quad (4.6)$$

$$\text{Skewness} = 1/N \sum (x_i - \mu_1/\sigma)^3, \quad (4.7)$$

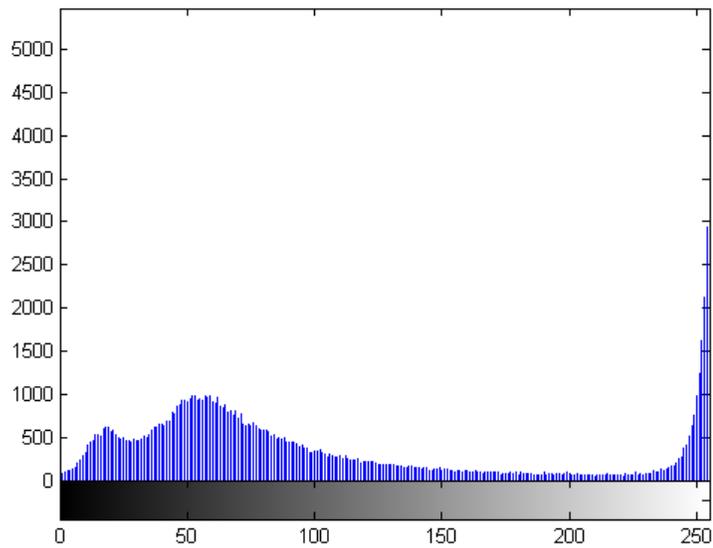
$$\text{Kurtosis} = 1/N \sum (x_i - \mu_1/\sigma)^4 - 3. \quad (4.8)$$

4.3.3 Adaptive operation

In order to achieve better adaptive properties, five different filters with different window sizes are selected according to the value of $b(i, j)$ as shown in Figure 4.12. Linear filters smooth the edges, average details with noise, and decrease greatly the spatial resolution. As a consequence, the nonlinear filter with a small window (e.g. 3×3) can be used in suspicious areas containing honeycombing and calcified lesions, while a linear filter with a large window (e.g. 7×7 or 9×9) can be used in areas without suspicious features to achieve noise removal and background smoothing.

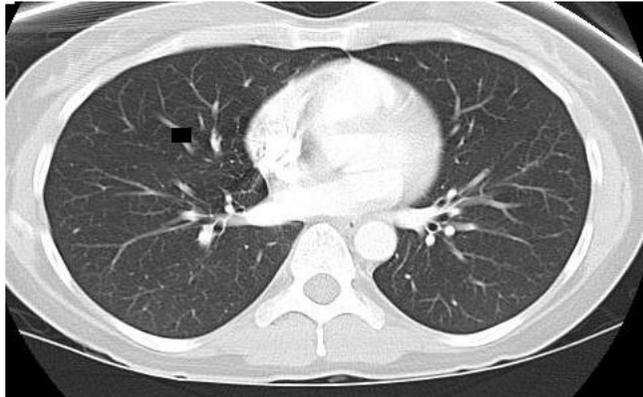


(a) IPF Subject

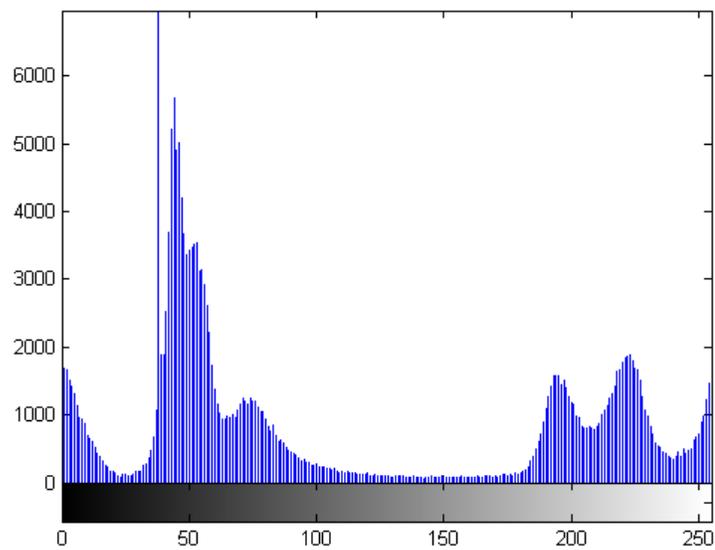


(b) Corresponding Density Histogram

Figure 4.10: A CT density histogram from a fibrotic lung is less skewed, less kurtotic and have increased mean lung attenuation compared to healthy lungs.



(a) Healthy Subject



(b) Corresponding Density Histogram

Figure 4.11: CT density histogram from normal lung shows a sharp peak and is skewed to the left in comparison to the Gaussian normal distribution.

The tree structured filter (TSF) seen in Figure 4.12 is based on the central weighted median filter (CWMF), which provides a selectable compromise between noise removal and edge preservation in the operation of the conventional median filter. The TSF is a multi-stage non-linear filter that consists of multiple CWMFs organized in a tree structure.

During the enhancement process, the filter will be automatically selected by the weighting factor $b(i, j)$. The outputs of the five filters in the first stage will be subject to a binary weighting that will select one filter out of five.

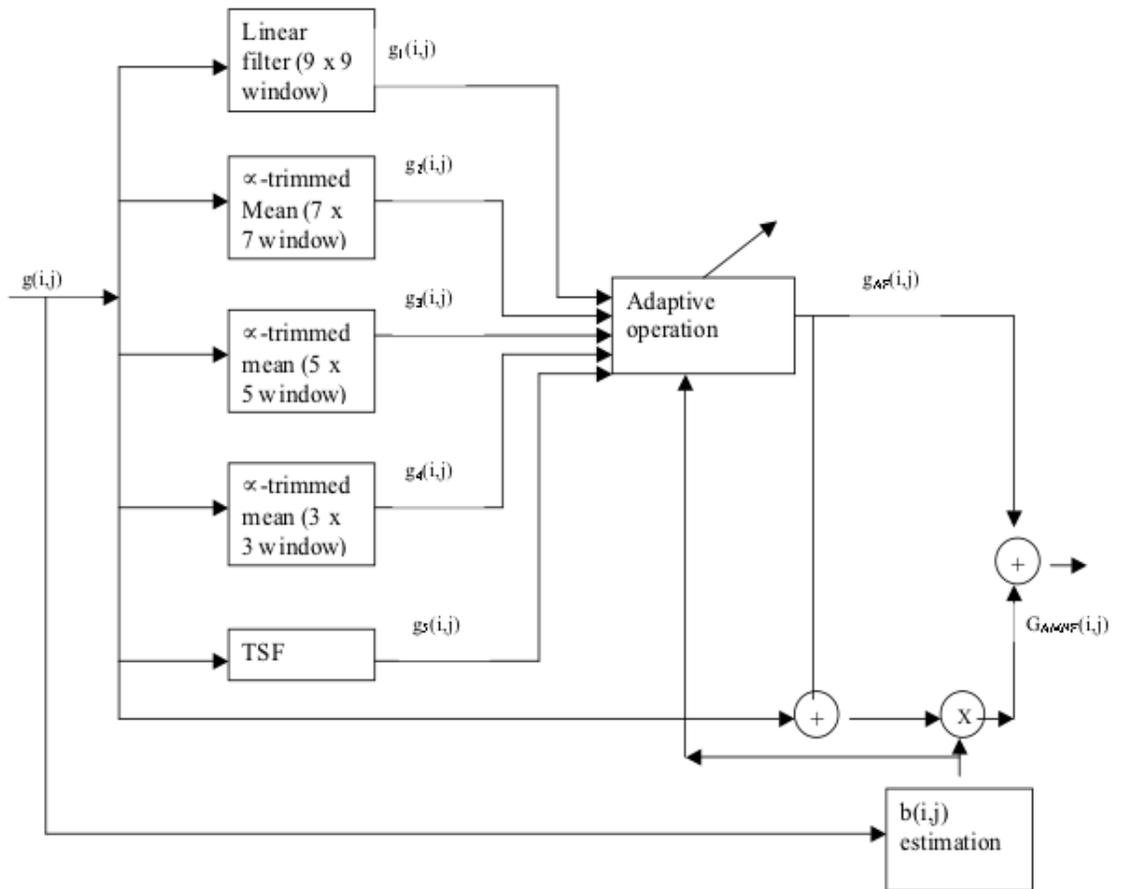


Figure 4.12: Block diagram of the Adaptive Multistage Nonlinear Filter (AMNF).

Representing the outputs of the five filters as the vector g_j , this adaptive process can be written with the vector notation [56]:

$$g_l = c^T g_j, 1 \leq j \leq m. \quad (4.9)$$

The elements of c are set adaptively according to the weighting factor $b(i,j)$.

$$[c_1, c_2, c_3, c_4, c_5]^T = \begin{cases} [1, 0, 0, 0, 0]^T & b(i, j) \leq \tau_1 \\ [0, 1, 0, 0, 0]^T & \tau_1 < b(i, j) \leq \tau_2 \\ [0, 0, 0, 1, 0]^T & \tau_3 < b(i, j) \leq \tau_4 \\ [0, 0, 0, 0, 1]^T & b(i, j) > \tau_4 \end{cases}. \quad (4.10)$$

The threshold τ_1 to τ_4 , will be set by the user.

4.4 ROI extraction

Regions of interest (ROIs) are portions of images, either selected graphically or selected by other means such as thresholding. Once the HRCT images have been thresholded and smoothed, the next logical step is to extract the regions of interest or ROI. These are the diseased or healthy portions of the lung which help quantify the disease extent and aid in the radiological diagnosis of the disease.

4.4.1 Co-Occurrence Matrix Measures

Texture analysis is a multistep process. The first step is usually an enhancement step to emphasize the texture and reduce the influence of unwanted processes. The enhancement step includes noise reduction, background removal, and possibly, histogram equalization. The presence of noise may create a pseudo texture that has the potential to dominate and

confound texture classification [57]. Background inhomogeneities may affect any intensity-based operations, such as thresholding, and they may also confound some texture metrics.

Some texture metrics are more robust against background inhomogeneities. These include first order statistics, co-occurrence matrix and correlation function [58]. ROIs are delineated over which quantitative texture features are calculated. Each pixel or region is assigned to one or several quantitative texture feature metrics as seen in Figures 4.13. The information collected in the classification step is used to determine the category into which pixel or region belongs. Within a defined region, its texture can be categorized as healthy or diseased as seen in Figure 4.13.

The Gray-Level Co-Occurrence Matrix (GLCM) is a method for extracting second order statistical texture features [59]. The GLCM functions in Matlab software are used to characterize the texture of an image by calculating how often pairs of pixel with specific values and in a specified spatial relationship occur in an image, creating a GLCM, and then extracting statistical measures from this matrix. These statistics provide information about the texture of an image. A number of texture features can be extracted from the GLCM. The following two features were used for the purpose of texture analysis: Correlation and Variance.

$$\text{CORR} = \sum_{i=0}^{G-1} \sum_{j=0}^{G-1} \frac{\{i \times j\} \times P(i, j) - \{\mu_x \times \mu_y\}}{\sigma_x \times \sigma + y}. \quad (4.11)$$

Correlation is a measure of gray level linear dependence between the pixels at the specified positions relative to each other.

$$\text{VAR} = \sum_{i=0}^{G-1} \sum_{j=0}^{G-1} (i - \mu)^2 P(i, j). \quad (4.12)$$

Variance puts relatively high weights on the elements that differ from the average value of $P(i, j)$.

G is the number of gray levels in the image. The gray level co-occurrence matrix $P(i, j)$ contains the second order statistical probability values for changes between the gray level intensities i and j . μ is the mean value of P . $\mu_x, \mu_y, \sigma_x, \sigma_y$ are means and standard deviations of P_x and P_y respectively.

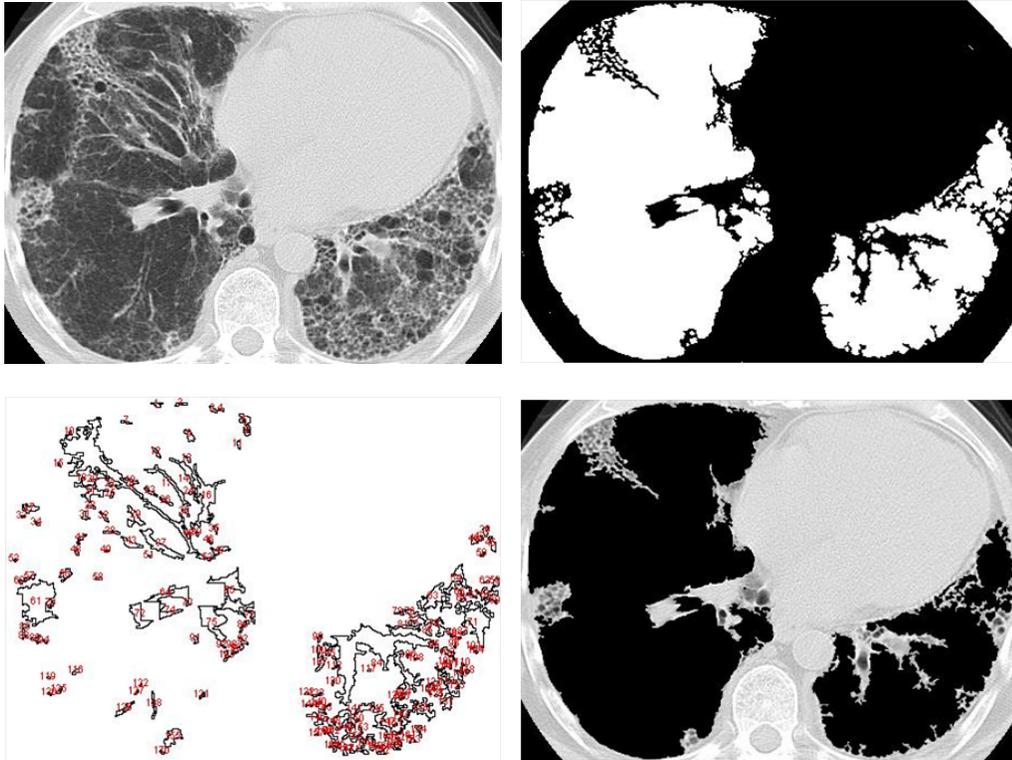


Figure 4.13: Texture quantified using texture metrics such as statistical moments (variance, kurtosis) and co-occurrence matrices.

4.5 Discussion

An adaptive thresholding algorithm has been presented which utilized quantitative CT indexes to correlate IPF with pulmonary abnormality. Simulation results demonstrated that the algorithm performed well in identified IPF images as seen in Figures 4.7, 4.9 and 4.13. However the absence of gold standards made quantification challenging for early stage

images of IPF and blinded images. Although, statistical moments such as variance and kurtosis gave us good segmentation results, complex textural information such as in IPF can be better analyzed in the Fourier domain. Spectral measures obtained from Fourier transform of the image can quantify texture, particularly when repetitive patterns are present.

4.6 Backpropagation Neural Network (BPNN)

The architecture of choice for Neural Network will be a ‘perceptron’ as shown in Figure 4.14. A weighted sum of the input constitute the argument of a non-linear activation function such as a Sigmoidal function. The resulting threshold value of the activation function is the output of the neural element. For the learning process a backpropagation algorithm will be chosen.

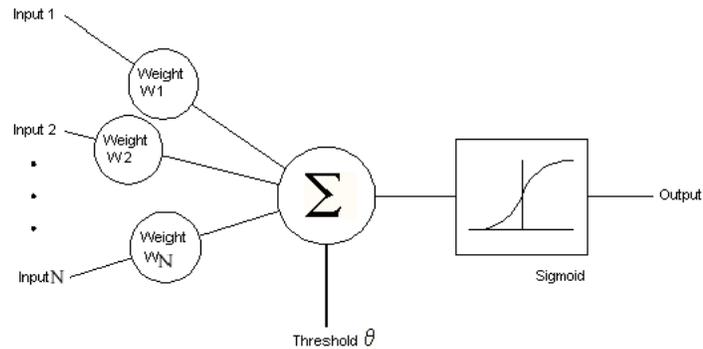


Figure 4.14: A basic computational neural element or Perceptron for classification.

Consider Figure 4.15 which is a feed forward back propagation network with one hidden layer. For our research, we would need a multi-layer feed forward network with L layers of N neural elements (perceptrons) in each layer such that

$$y^k = F(W^k y^{k-1}) \quad (4.13)$$

for $k=1,2,\dots,L$, where y^k is the output of the k th layer neural elements with $k = 0$ representing the input layer and W^k is the weight matrix for the k th layer.

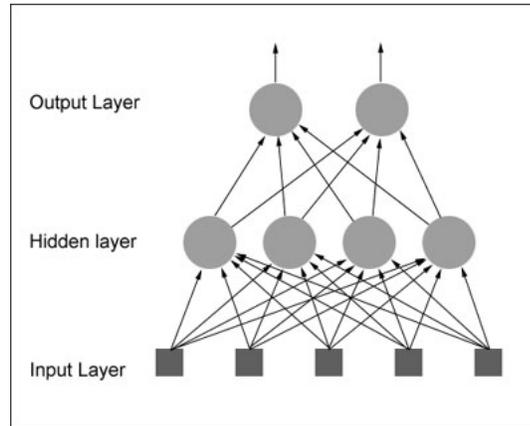


Figure 4.15: Feedforward backpropagation network with one hidden layer.

This method will need to be preceded by some preprocessing of the image data. This would entail identification of regions of interest (ROI) using pixel classification methods like histogram analysis. The features would be normalized to the unit interval $[0,1]$ and then split into the training and test set. The training examples would represent a reasonably complete statistical distribution of the input data. The BPNN follows a gradient descent algorithm.

1. Assign random weights in the range of $[-1,+1]$ to all weights w_{ij}^k .
2. For each classified pattern pair $\{y^0, y^L\}$ in the training set, do the following steps:
 - Compute the output values of each neural element using the current weight matrix.
 - Find the error $e^{(k)}$ between the computed output vector and the desired output vector for the classified pattern pair.

- Adjust the weight matrix using the change $\Delta W^{(k)}$ computed as

$$\Delta W^{(k)} = \alpha e^{(k)} [y^{(k-1)}] \quad (4.14)$$

for all layers $k = 1, \dots, L$ where α is the learning rate that can be set between 0 and 1.

3. Repeat step 2 for all classified pattern pairs in the training set until the error vector for training example is sufficiently low or zero.

4.6.1 Radial Basis Function (RBF)

The basic RBF contains an input layer for the image, a hidden layer for the RBF and an output layer. The input layer starts with n input neurons with activations $x_i, i \in (1..n)$, the activation pattern of the input layer is represented by an n -dimensional vector x (input image) is the 'Feature space' R^n [56]

The hidden layer provides the radial basis function processing. The activation is propagated to the N neurons of the hidden layer by directed connections with weights. The output of each RBF unit is then weighted and summed such that the weight $w_j \in R^n, j \in (1..n)$ are computed as a set of vectors that represent the dataset in the feature space. The RBF (a Gaussian function) is applied to the Euclidean distance between the input vector and its own centroid (virtual position w_j of the hidden layer neuron j). The activation of the hidden layer neuron should decline with increasing distance between the vector x and its centroid. The output of each RBF unit is then weighted and summed through a linear combiner to provide the final output of the network.

Proposed Algorithm for Segmenting Honeycomb Features in IPF

1. Specific features with information about texture is needed. First order texture features characterized by the gray level histogram are mean, variance, skewness and kurtosis are obtained.

2. Second order features such as contrast and local inhomogeneities [60] are obtained using second order derivative functions such as laplachian and Gaussian function.
3. The contrast value of a pixel can be computed from the largest difference in gray values between a pixel and its 8-connected neighbors. Alternatively, adaptive neighborhood processing can be used to obtain the contrast value of the central pixel.
4. The feature vector consists of gray values of the center pixel and its neighbors obtained with the contrast value of the center pixel.
5. All features are normalized to unit variance to be used in the RBF network classifier.
6. A training set of example images is prepared with a large number of pixels that are manually classified as belonging to the background or honeycomb region.
7. Using representative features of all pixels, feature vectors are prepared and scaled for fuzzy clustering
8. The centers of the clusters become locations of RBF that are used in the RBF network.
9. The RBF network classifier is used for training feature vectors and classifying pixels into two classes: background and honeycomb structure.

4.6.2 Implementation and Analysis

The Neural Network was implemented using the Neural Network Toolbox from Matlab. The multilayer feedforward neural network is the workhorse of the Neural Network Toolbox software.

The lack of any improvement in segmentation quality was accompanied by a considerable increment in human intervention: manual pressing time for interactive labeling of the training data was approximately 20 to 30 minutes per image. The second problem related to the amount of input data. For achieving a high and reliable performance for non-training cases, a large number of training cases are commonly required. If an artificial neural network (ANN) is trained with only a small number of cases, the generalization ability (performance

for non-training cases) will be lower (e.g., the ANN may fit only the training cases). An unblinded number of 16 IPF patients was insufficient to train and evaluate the quality of the segmentation results.

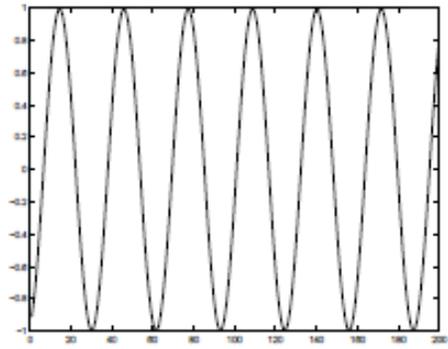
4.7 Gabor Filter

The honeycomb target has objects with certain textures which can be described in terms of size but not shape. These targets are a cluster or row of cysts [4]. Honeycomb cysts are very small (less than 5 mm in diameter) and are often subpleural. As areas of honeycombing increases honeycomb cysts also increase in size in patients with IPF. Whereas the middle portion of the lung on the left shows a smoother texture indicative of healthier tissue.

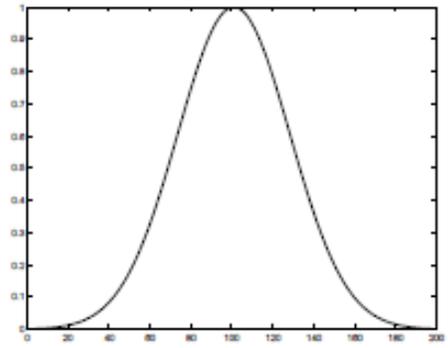
2-D Gabor filter has been recognized to be a very useful tool in computer vision and image processing, especially for texture analysis. The increasing research on Gabor analysis is motivated by biological findings. Numerous papers have been published on Gabor analysis since Gabor proposed the 1-D Gabor function [61–66]. A Gabor filter is obtained by modulating a sinusoid with a Gaussian. For the case of one dimensional (1D) signals, a 1D sinusoid is modulated with a Gaussian. This filter will therefore respond to some frequency but only in a localized part of the signal. Figure 4.16 shows the gabor filter composition for a 1D signal. This comprises of the sinusoid, Gaussian kernel and the corresponding Gabor filter. Let $g(x, y, \theta, \phi)$ be the function defining a Gabor filter centered at the origin with θ as the spatial frequency and ϕ as the orientation. Gabor filters can be viewed as:

$$g(x, y, \theta, \phi) = \exp\left(-\frac{x^2 + y^2}{\sigma^2}\right) \exp(2\pi\theta i(x \cos \phi + y \sin \phi)) \quad (4.15)$$

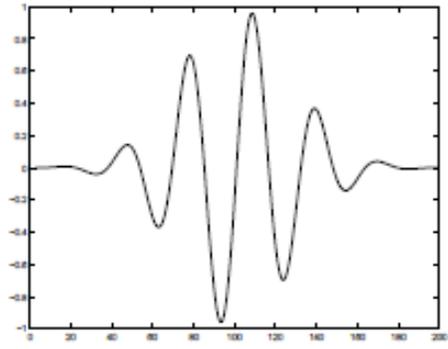
It has been shown that σ , the standard deviation of the Gaussian kernel depends upon the spatial frequency to measured, i.e. θ . The relationship between σ and θ is $\sigma = 0.5 \times 1/\theta$ where $1/\theta$ represents the wavelength of the cosine factor and is inversely proportional to frequency.



(a)



(b)



(c)

Figure 4.16: Gabor filter composition for 1D signals:(a) sinusoid, (b) a Gaussian kernel, (c) the corresponding Gabor filter.

A Gabor filter is a linear filter obtained by modulation of a sinusoidal function with a Gaussian function of standard deviation ' σ '. The application of this filter to the images, seen as a two-dimensional signal, is as follows: a sinusoidal function with orientation of 30° with the x -axis multiplied by a Gaussian kernel. The Figure 4.17 shows a) the sine function, in b) the Gaussian kernel and c) the result.

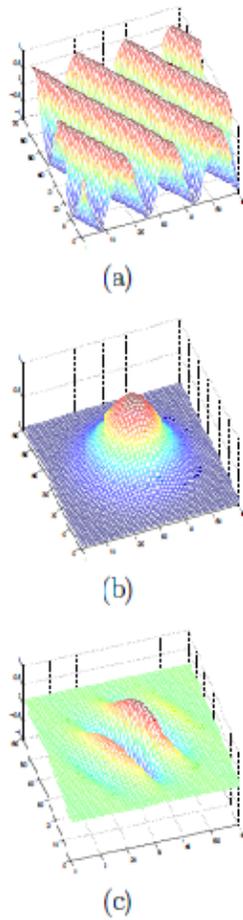


Figure 4.17: Gabor filter composition: (a) 2D sinusoid, (b) a Gaussian kernel, (c) the corresponding Gabor filter.

The response of a Gabor filter to an image is obtained by a 2D convolution operation.

Let $I(x, y)$ denote the image and $G(x, y, \theta, \phi)$ denote the response of a Gabor filter with frequency θ and orientation ϕ to an image at point (x, y) on the image plane. $G(\cdot)$ is obtained as

$$G(x, y, \theta, \phi) = \int \int I(p, q)g(x - p, y - q, \theta, \phi) dp dq \quad (4.16)$$

Fixed dimensions $n \times n$, is calculated by a Gabor filter with the above equation, the result is displayed as two images representing the real part (left) and the imaginary part of the filter (right) as seen in Figure 4.18. Also shown in Figure 4.19 are Gabor filters with different frequencies and orientations. The segmentation result after the filter is applied is shown in 4.20. The diseased portion of the lung is seen clearly towards the periphery. It is also identified in the left upper section of the lung.

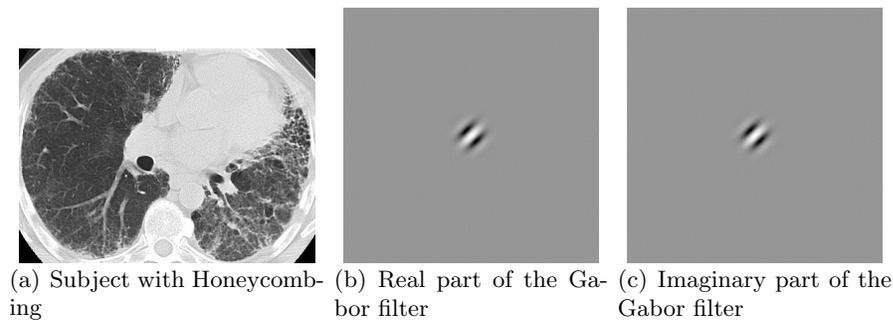


Figure 4.18: Example of Gabor filter oriented at 66° with the x -axis .

4.7.1 Goal

This section discusses invariant texture segmentation based on multichannel analysis. The traditional Gabor filter is modified into a circular symmetric version. A very important property of this new version is that it is rotation invariant. Texture images are decomposed into several channel outputs. K-means clustering algorithm is employed for pixel classification. The goal will be addressed through several major steps.

1. Create a set of Gabor filters.

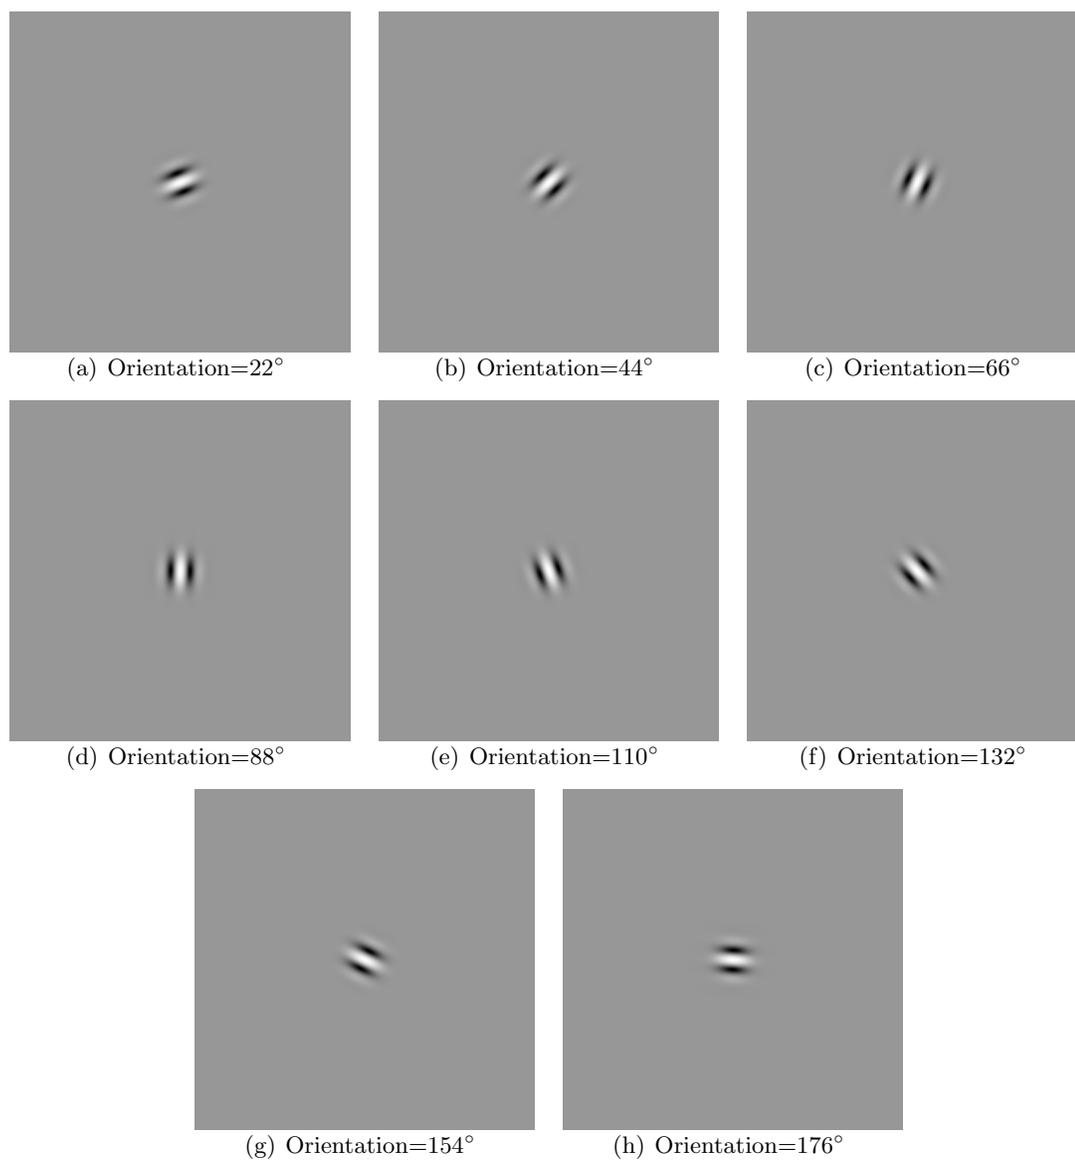


Figure 4.19: Gabor filters with different frequencies and orientations.

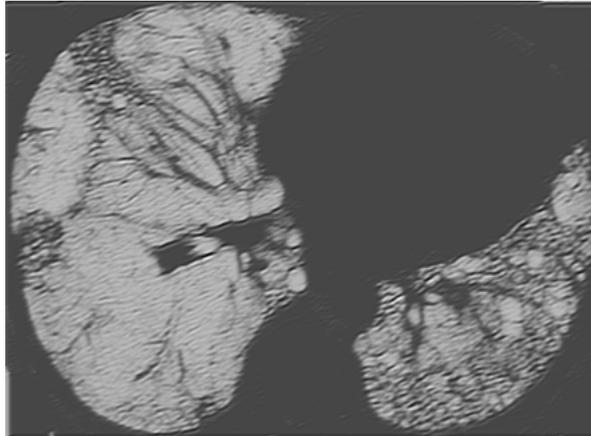


Figure 4.20: Segmentation result after Gabor filter is applied.

2. Correlate the filters with the image.
3. Smooth the correlations.
4. Extract selected pixels for training.
5. Train a K-means system for recognition of the pixels in the image.
6. Classify all pixels in the image.

4.7.2 Implementation

The Gabor Filter has been implemented using ImageJ, a Java image processing program. Two-dimensional Gabor functions were proposed by Daugman et al [67] to model the spatial summation properties (of the receptive fields) of simple cells in the visual cortex. They are widely used in image processing, computer vision, neuroscience and psychophysics. The parameterization used in 4.15 follows references[68–73].

4.8 Discussion

The lessons from the earlier approaches such as Adaptive thresholding led to the conclusion that complex textural information such as in IPF can be better analyzed in the Fourier

domain. Spectral measures obtained from Fourier transform of the image can quantify texture, particularly when repetitive patterns are present. Keeping this in mind other texture based approaches such as Gabor filter have been explored to improve upon the segmentation and develop algorithms to detect evidence of early onset of IPF. Gabor filter is a very useful tool for texture analysis in both domains and hence combines the advantages of both filters. However the disadvantage of image dependence unexpectedly resulted in image enhancements. Certain configuration settings for wavelength, orientation, phase offset or aspect ratio worked well with one set of images but not with another.

One of the possible reasons in the difficulty in parameter selection for Gabor filter design could be attributed to the nonorthogonal characteristic of Gabor filters [67, 74]. It was challenging to identify identical texture with less or more resolution than the original image presented to the Gabor filter as the same texture. It is known that a set of Gabor filters represents a complete but non orthogonal basis of a two-dimensional signal [75, 76]. Non-orthogonality implies a certain degree of redundancy among the features that can be extracted from the various filters. In general the greater the superposition zone between the adjacent filter, the higher the information redundancy [77]. The effects of filter orthogonality on pattern classification is not completely understood. The most commonly adopted approach, in classification based on Gabor filters, is to minimize non-orthogonality by designing Gabor filters that can overcome this redundancy [76, 78–80]. However it has been noticed that filter orthogonality would make the response of the filter bank unstable under rotation and translation, with negative effects on classification [81].

Chapter 5: Pulse Images

In this section two digital models will be presented. The first is the Pulse-Coupled Neural Network (PCNN) which for many years was the standard for many image processing applications [2, 82, 83]. The PCNN is based solely on the Eckhorn model [84] but there are many other cortical models that exist. These models all have a common mathematical foundation, but beyond the common foundation each also had unique terms. Since the goal here is to build image processing routines and not to exactly simulate the biological system a new model was constructed. This model contained the common foundation without the extra terms and is therefore viewed as the intersection of the several cortical models, and it is named the Intersecting Cortical Model (ICM) [85].

5.1 PCNN Details

The PCNN differs from the Eckhorn model to account for digitization[86]. The orientation, scale and skewness of the inputs did not alter the results in prototypes of PCNN. A wide range of images were run through the PCNN implementation chosen to determine the basis of its working mechanisms and runtime parameters, to ensure its efficiency as an image processing engine for general purpose HRCT images.

5.1.1 Original Model

A PCNN neuron shown in Figure 5.1 contains two main compartments: the Feeding and Linking compartments. Each of these communicates with neighboring neurons through the synaptic weights \mathbf{M} and \mathbf{W} respectively. Each retains its previous state but with a decay factor. Only the Feeding compartment receives the input stimulus, \mathbf{S} . The values of these two compartments are determined by,

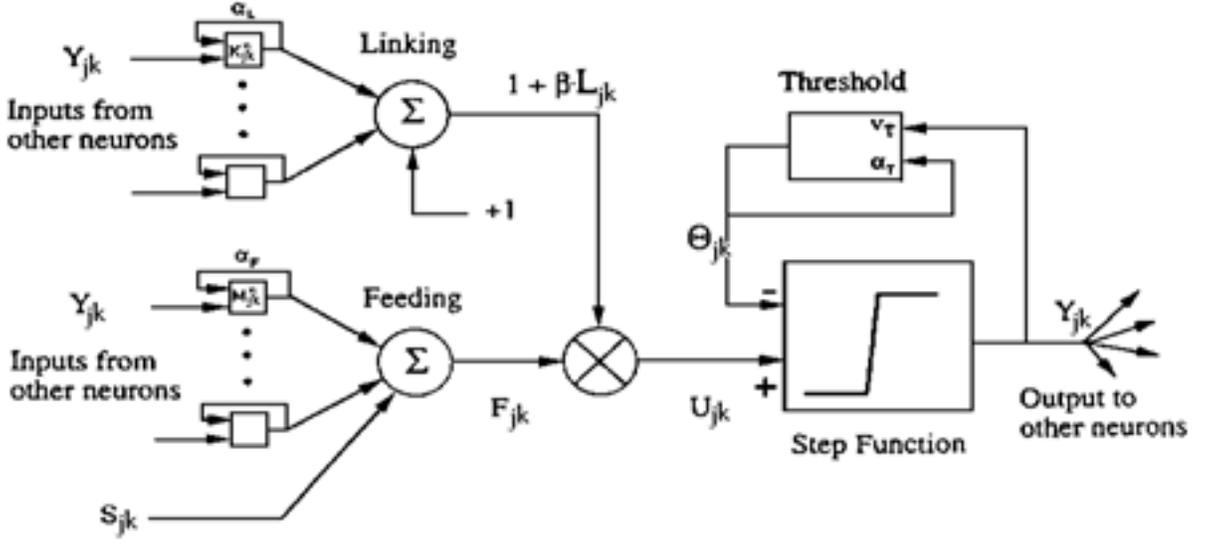


Figure 5.1: Schematic representation of a PCNN processing element [2].

$$F_{ij}[n] = e^{\alpha_F \delta n} F_{ij}[n-1] + S_{ij} + V_F \sum_{kl} M_{ijkl} Y_{kl}[n-1], \quad (5.1)$$

$$L_{ij}[n] = e^{\alpha_L \delta n} L_{ij}[n-1] + V_L \sum_{kl} W_{ijkl} Y_{kl}[n-1], \quad (5.2)$$

where F_{ij} is the Feeding compartment of the (i, j) neuron embedded in a two-dimensional array of neurons, and L_{ij} is the corresponding Linking compartment. Y_{kl} are the outputs of neurons from a previous iteration $[n-1]$. Both compartments have a memory of the previous state, which decays in time by the exponent term. The constants V_F and V_L are normalizing constants. If the receptive fields of \mathbf{M} and \mathbf{W} change then these constants are used to scale the resultant correlation to prevent saturation.

The state of these two compartments are combined in a second order fashion to create the internal state of the neuron, \mathbf{U} . The combination is controlled by the linking strength,

β . The internal activity is calculated by,

$$U_{ij}[n] = F_{ij}[n] (1 + \beta L_{ij}[n]). \quad (5.3)$$

The internal state of the neuron is compared to a dynamic threshold, Θ , to produce the output, \mathbf{Y} , by

$$Y_{ij}[n] = \begin{cases} 1 & \text{if } U_{ij}[n] > \Theta_{ij}[n-1] \\ 0 & \text{Otherwise} \end{cases}. \quad (5.4)$$

The threshold is dynamic in that when the neuron fires ($Y_{ij} > \Theta_{ij}$) the threshold then significantly increases its value. This value then decays until the neuron fires again. This process is described by,

$$\Theta_{ij}[n] = e^{\alpha \Theta \delta n} \Theta_{ij}[n-1] + V_{\Theta} Y_{ij}[n], \quad (5.5)$$

where V_{Θ} is a large constant that is generally more than an order of magnitude greater than the average value of U_{ij} .

5.1.2 Implementing in Python

Consider a case where the input is a 360×360 frame and centered in this frame is a solid ‘T’. The results are shown in Figure 5.2 with the original being shown in Figure 5.2(a) shows how the pulsed images are generated from the original image at each iteration of pulsing (advancing n).

A combination of neurons that pulse at $n = 0$ not pulsing again for many iterations, and the neighboring pixels being encouraged to pulse through the connections \mathbf{M} and \mathbf{W} causes the decay of the initial pulse and the waves traveling thru the matrix of neurons.

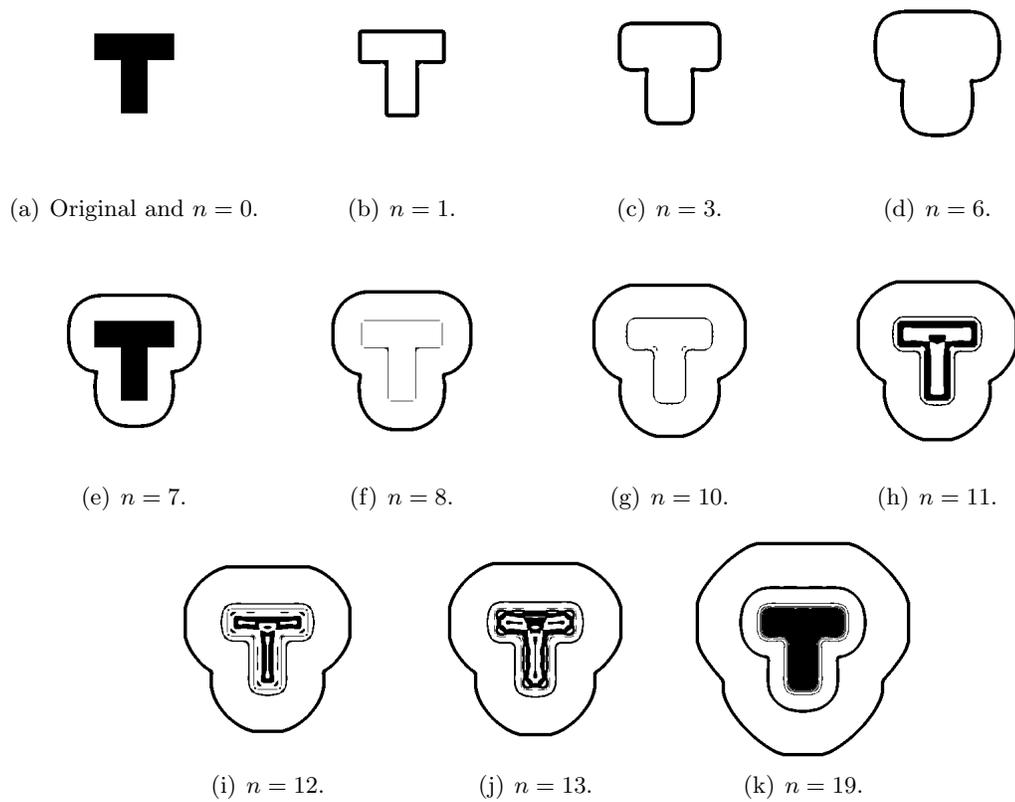


Figure 5.2: An input and selected pulse images from the PCNN [2].

5.1.3 Spiking Behavior

For a single neuron, the input stimulus, \mathbf{S} , and the stimulus from neighbors determines the output. The internal activity rises until it becomes larger than the threshold value. Then the neuron fires and the threshold sharply increases then begins its decay until once again the internal activity becomes larger than the threshold. This causes the pulsing nature of the PCNN. Figure 5.3 shows the activity within a single neuron embedded in a 2D array as it progresses in time.

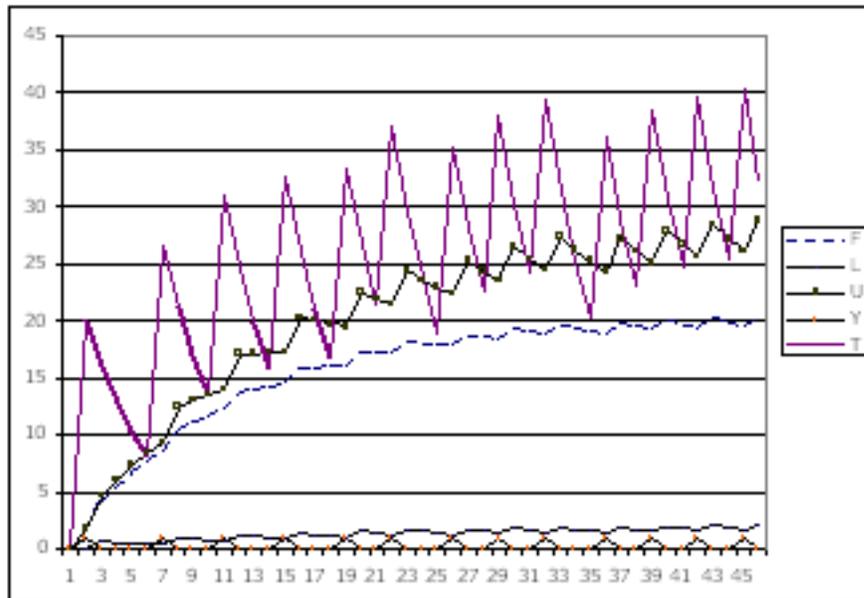


Figure 5.3: An example of the progression of the states of a single neuron. See the text for explanation of \mathbf{L} , \mathbf{U} , \mathbf{T} and \mathbf{F} [2].

The pulses also trigger communications to neighboring neurons when the output of the neuron is high. Equations (5.1) and (5.2) explain this in detail. Consider three neurons A, B, and C that are linearly arranged with B between A and C. For this example, only A is receiving an input stimulus. At $n = 0$, the A neuron pulses sending a large signal to B. At $n = 1$, B receives the large signal, pulses, and then sends a signal to both A and C. At $n = 2$, the A neuron still has a rather large threshold value and therefore the stimulus is not

large enough to pulse the neuron. Similarly, neuron B is turned off by its threshold. On the other hand, C has a low threshold value and will pulse. Thus, a pulse sequence progresses from A to C.

This process is the beginning of the *autowave* nature of the PCNN. Basically, when a neuron (or group of neurons) fires, an autowave emanates from that perimeter of the group. Autowaves are defined as normal propagating waves that do not reflect or refract. In other words, when two waves collide they do not pass through each other. Autowaves are being discovered in many aspects of nature and are creating a significant amount of scientific research [87, 88]. The PCNN, however, does not necessarily produce a pure autowave and alteration of some of the PCNN parameters can alter the behavior of the waves.

Consider the image in Figure 5.4. The original input consists of two ‘T’s. The intensity of each ‘T’ is constant, but the intensities of each ‘T’ differ slightly. At $n = 04$ the neurons that receive stimulus from either of the ‘T’s will pulse in step $n = 16$ (denoted as black). As the iterations progress, the autowaves emanate from the original pulse regions. At $n = 10$ it is seen that the two waves did not pass through each other. At $n = 12$ the more intense ‘T’ again pulses. Thus, small differences in intensity are clearly visible in the pulsing patterns.

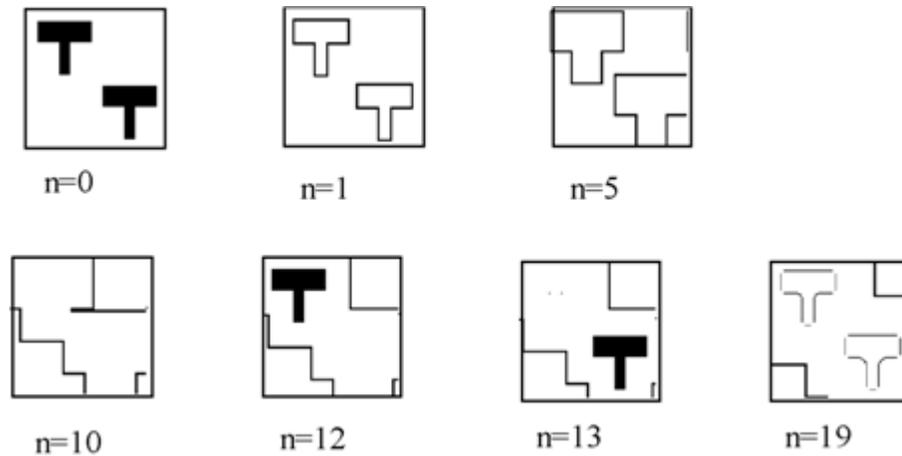


Figure 5.4: A typical PCNN example.

The network also exhibits some synchronizing behavior. In the early iterations segments tend to pulse together. However, as the iterations progress, the segments tend to de-synchronize. Synchronicity occurs by a pulse capture. This occurs when one neuron is close to pulsing ($U_{ij} < \Theta_{ij}$) and its neighbor fires. The additional input from the neighbor will provide an additional input to U_{ij} thus allowing the neuron to fire prematurely. The two neurons, in a sense, synchronize due to their linking communications.

The de-synchronization occurs in more complex images due to signals that are residual in nature. As the network processes the signals, the neurons are affected indirectly from other non-neighboring neurons. This alters their behavior and the synchronicity begins to fail. The beginning of this failure can be seen by comparing $n = 1$ to $n = 19$ in Figure 5.4. Note that the corners of the ‘T’ autowave are missing in $n = 19$. This phenomenon is more noticeable in more complicated images.

Gernster [89] argues that the lack of noise in such a system is responsible for the de-synchronization. However, applications demonstrated in subsequent chapters specifically show the PCNN architecture does not exhibit this link. Synchronization has been explored more thoroughly for similar integrate and fire models [90].

The PCNN has many parameters that can be altered to adjust its behavior. The (global) linking strength, β , in particular, has many interesting properties (in particular effects on segmentation). While this parameter, together with the two weight matrices, scales the feeding and linking inputs, the three potentials, V , scale the internal signals. Finally, the time constants and the offset parameter of the firing threshold are used to adjust the conversions between pulses and magnitudes. The dimension of the convolution kernel directly affects the speed that the autowave travels. The dimension of the kernel allows the neurons to communicate with neurons farther away and thus allows the autowave to advance farther in each iteration.

The pulse behavior of a single neuron is greatly affected by α_{Θ} and V_{Θ} . The α_{Θ} affects the decay of the threshold value and the V_{Θ} affects the height of the threshold increase after the neuron pulses. It is quite possible to force the neuron to enter into a multiple

pulse regime. In this scenario the neuron pulses in consecutive iterations. The autowave created by the PCNN is greatly affected by V_F . Setting V_F to 0 prevents the autowave from entering any region in which the stimulus is also 0. There is a range of V_F values that allows the autowave to travel but only for a limited distance.

There are also architectural changes that can alter the PCNN behavior. One such alteration is quantized linking where the linking values are either 1 or 0 depending on a local condition. In this system the Linking field is computed by,

$$L_{ij}[n] = \begin{cases} 1 & \text{if } \sum_{kl} w_{ijkl} Y_{kl} > \gamma \\ 0 & \text{Otherwise} \end{cases} . \quad (5.6)$$

Quantized linking tends to keep the autowaves clean. In the previous system autowaves traveling along a wide channel have been observed to decay about the edges. In other words a wave front tends to lose its shape near its outer boundaries. Quantized linking has been observed to maintain the wavefronts shape.

In a “real world” image generally all of the neurons receive some stimulus and thus in the initial iteration all neurons will pulse. Then it will take several iterations before they can pulse again. From an image processing perspective the first few iterations are unimportant since all neurons pulse in the first iteration and then non pulse for the next several iterations. An alternative is to initially set the threshold values higher. The first few iterations may not produce any pulses since the thresholds now need to decay. However, the frames with useful information will be produced in earlier iterations than in the “initially 0” scenario. Parodi [91] suggests that the Θ be reset after a few iterations to prevent de-synchronization.

The influence of a dynamic V_Θ was considered by Yamaguchi *et al.* [92] in which a static V_Θ was replaced by,

$$V_\Theta = 1 + a_0 \sin(\omega_0 t), \quad (5.7)$$

where a_0 and ω_0 were the variables that they altered. In a simulation of a 1D network they

demonstrated that certain combinations of a_0 and ω_0 produced a stable case in which two neurons receiving the same input pulse in synchrony, but other combinations produced a chaotic behavior.

5.1.4 Collective Behavior

The PCNN is run again and neural activity is collected for each compartment. This allows for visual inspection of the collective behavior of the neurons. Figure 5.5 shows the activity of the \mathbf{F} compartments of each neuron. At $n = 0$ only the neurons that are receiving a stimulus have a significant value since the only active part of Equation (5.1) is the stimulus \mathbf{S} .

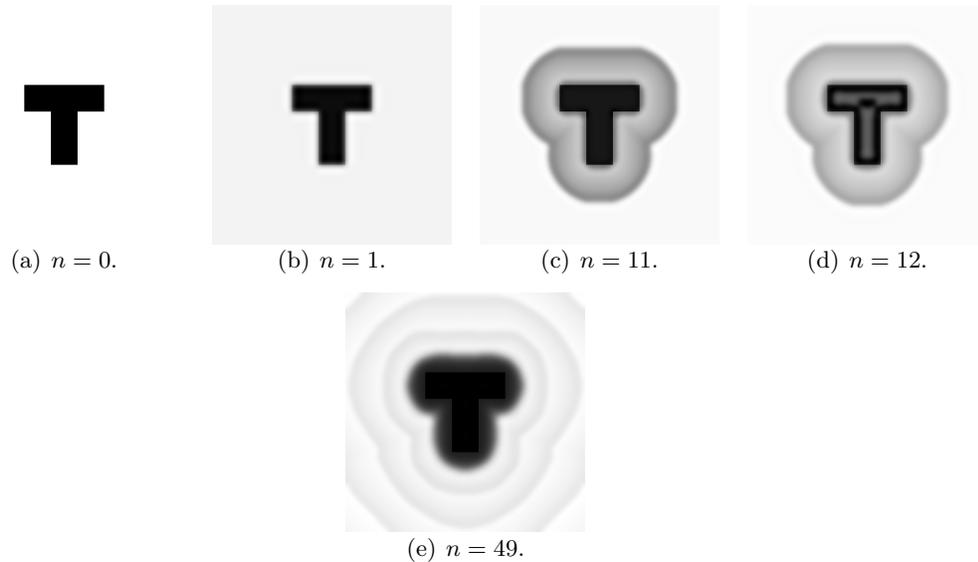


Figure 5.5: Collective behavior of \mathbf{F} [2].

At $n = 1$ the first and third terms in Equation (5.1) begin to contribute to the activity in \mathbf{F} . Basically, the autowave appears and continues to expand up to $n = 11$. In this example, at $n = 12$ the second cycle begins and the values in \mathbf{F} are increased for the neurons that fire. The image shown is scaled and therefore intensities shown on the page between $n = 11$ and $n = 12$ do not have the same magnitude. The last example shown is at $n = 49$ which

shows three cycles of waves expanding.

Figure 5.6 shows the intensities of the \mathbf{L} compartments with $n = 0$ not shown since there is no activity in the first iteration. While the expanding autowave is present it is also noticed that the \mathbf{L} compartment is more sensitive to edges.

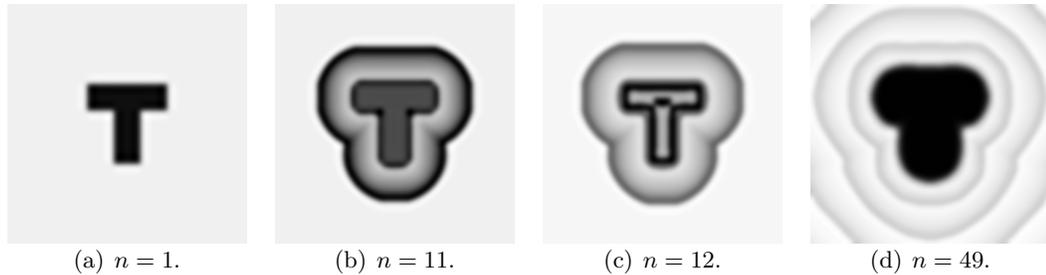


Figure 5.6: Collective behavior of \mathbf{L} [2].

The threshold Θ responses are shown in Figure 5.7. These have larger values when a neuron fires. Clearly, at $n = 12$ the expanding wave is shown to have decreasing values closer to the target since those neurons fired further back in time. The output arrays are shown in Figure 5.2.

5.1.5 Neural Connections

The PCNN contains two convolution kernels \mathbf{M} and \mathbf{W} . The original Eckhorn model used a Gaussian type of interconnections, but when the PCNN is applied to image processing problems these interconnections can be used for altering the behavior of the network.

The few examples shown here all use local interconnections. It is possible to use long range interconnections but two impositions arise. The first is that the computational load increases with the number of interconnections. The second is that PCNN tests to date do not differ in output significantly when using long range interconnections, although long range inhibitory connections of similar models have been proposed in similar cortical models [93].

Subsequent experiments replaced the interconnect pattern with a target pattern in the

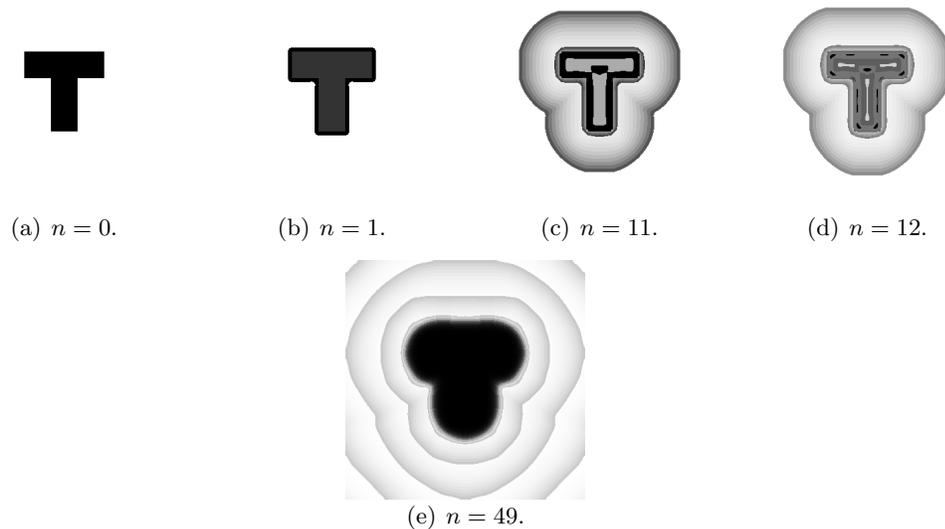


Figure 5.7: Collective behavior of Θ [2].

hope that on-target neurons would pulse more frequently. The matrices \mathbf{M} and \mathbf{W} were similar to the intensity pattern of a target object. In actuality there was very little difference in the output from this system than from the original PCNN. Further investigations revealed the reason for this. Positive interconnections tend to smooth the image and longer-range connections provide even more smoothing. The internal activity of the neuron may be quite altered by a change in interconnections. However, much of this change is nullified since the internal activity is compared to a dynamic threshold. The amount by which the internal activity surpasses the dynamic threshold is not important and thus the effects of longer-range interconnections are reduced.

Manipulations of a small number of interconnections do, however, provide drastic changes in the PCNN. A few examples of these are shown. For these examples we use the input shown in Figure 5.8. This input is a set of two ‘T’s.



Figure 5.8: An example input [2].

The first example computes the convolution kernel by,

$$K_{ij} = \begin{cases} 0 & \text{if } i = m \text{ and } j = m \\ 1/r & \text{Otherwise} \end{cases}, \quad (5.8)$$

where r is the distance from the center element to element ij , and m is half of the linear dimension of K . In this test K was 5×5 . Computationally, the feeding and linking equations are,

$$F_{ij}[n] = e^{\alpha_F \delta_n} F_{ij}[n-1] + S_{ij} + (\mathbf{K} \otimes \mathbf{Y})_{ij}, \quad (5.9)$$

and

$$L_{ij}[n] = e^{\alpha_L \delta_n} L_{ij}[n-1] + (\mathbf{K} \otimes \mathbf{Y})_{ij}. \quad (5.10)$$

The resultant outputs of the PCNN are shown in Figure 5.9. The output first pulses all neurons receiving an input stimulus. Then autowaves are established that expand from the original pulsing neurons. These autowaves are two pixels wide since the kernel extends two elements in any direction from the center. These autowaves expand at the same speed in both vertical and horizontal dimensions again due to the symmetry of the kernel.

Setting the elements of the previous kernel to zero for $i = 0$ and $i = 4$ defines a kernel that is asymmetric. This kernel will cause the autowaves to behave in a slightly different fashion. The results from these tests are shown in Figure 5.10.

The autowave in the vertical direction now travels at half the speed of the one in the

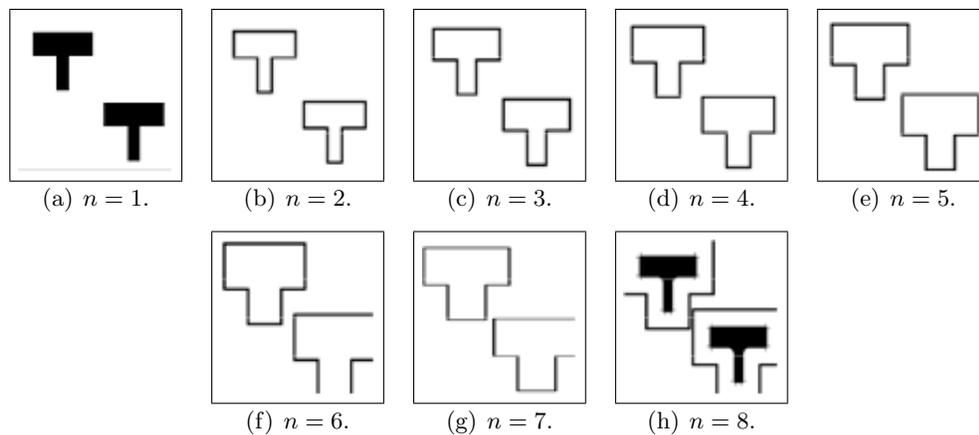


Figure 5.9: Output pulse images [2].

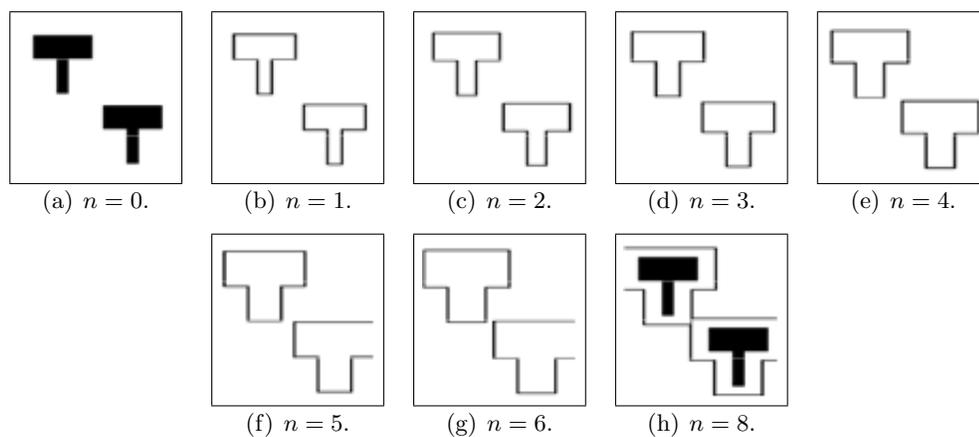


Figure 5.10: Outputs of a PCNN with an asymmetric kernel, as discussed in the text. These outputs should be compared to those shown in Figure 5.11 [2].

horizontal direction. Also the second pulse of the neurons receiving stimulus is delayed a frame. This delay is due to the fact that these neurons were receiving less stimulus from their neighbors. Increasing the values in K could eliminate the delay.

The final test involves altering the original kernel by simply requiring that,

$$K_{ij} = \begin{cases} K_{ij} & \text{if } i = m \text{ and } j = m \\ -K_{ij} & \text{Otherwise} \end{cases}. \quad (5.11)$$

The kernel now has a positive value at the center and negative values surrounding it. This configuration is termed On-Center/Off-Surround. Such configurations of interconnections have been observed in the eye. Furthermore, convolutions with a zero-mean version of this function are quite often used as an “edge enhancer”. Employing this type of function in the PCNN has a very dramatic effect on the outputs as is shown in Figure 5.11. The autowaves created by this system are now dotted lines. This is due to *competition* amongst the neurons since each neuron is now receiving both positive and negative inputs.

5.2 The ICM

The PCNN is a digital model based upon a single biological model. As stated earlier there are several biological models that have been proposed. These models are mathematically similar to the Fitzhugh-Nagumo system in that each neuron consists of coupled oscillators. When the goal is to create image processing applications it is no longer necessary to exactly replicate the biological system. The important contribution of the cortical model is to extract information from the image and there is little concern as to the deviation from any single biological model.

The ICM is a model that attempts to minimize the cost of calculation while maintaining the effectiveness of the cortical model when applied to images. Its foundation is based on the common elements of several biological models.

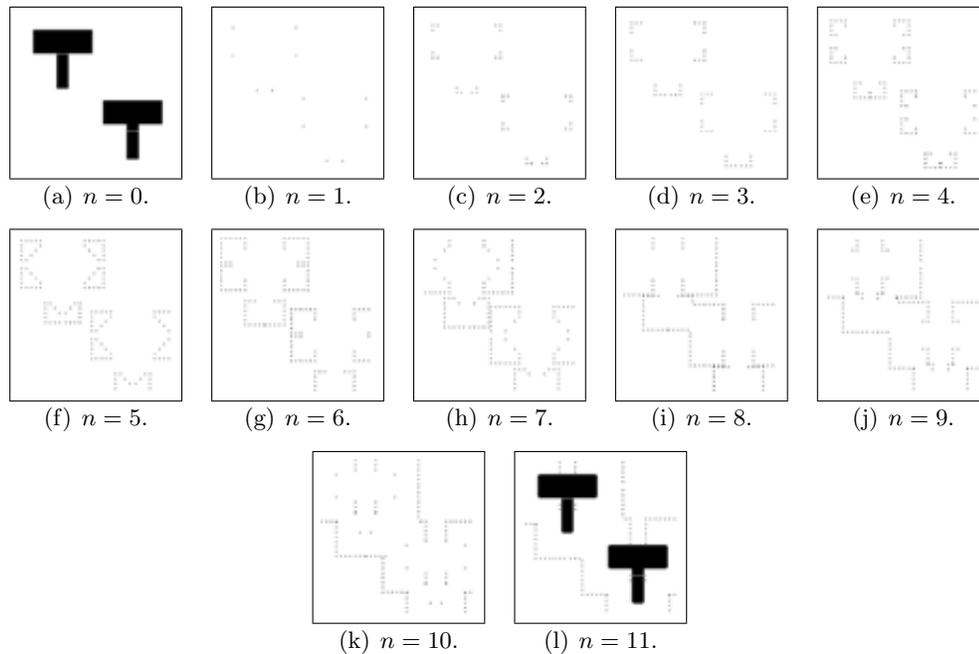


Figure 5.11: Outputs of a PCNN with an on-center/off-surround kernel [2].

5.2.1 Minimum Requirements

Each neuron must contain at least two coupled oscillators, connections to other neurons, and a nonlinear operation that determines decisively when a neuron pulses. In order to build a system that minimizes the computation it must first be determined which operation creates the highest cost. In the case of the PCNN almost all of the cost of computation stems from the interconnection of the neurons. In many implementations users set $\mathbf{M} = \mathbf{W}$ which would cut the computational needs in half. One method of reducing the costs of computation is to replace the traditional Gaussian type connections.

Another method is to reduce the number of connections. What is the minimum number of neurons required to make an operable system? This question is answered by building a minimal system and then determining if it created autowave communications between the neurons. Consider the input image in Figure 5.12 which contains two basic shapes.

The system that is developed must create autowaves that emanate from these two shapes. So, a model was created that connected each neuron to P other neurons. Each

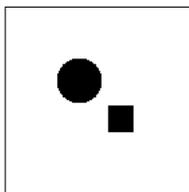


Figure 5.12: An input image [2].

neuron was permanently connected to P random nearest neighbors and the simulation was allowed to run several iterations. The results in Figure 5.13 display the results of three simulations. In the first $P = 1$ and the figure displays which neurons pulsed during the first 10 iterations. After 10 iterations this system stabilized. In other words the autowave stalled and did not expand. In the second test $P = 2$ and again the autowave did not expand. In both of these cases it is believed that the system had insufficient energy to propagate the communications between the neurons. The third test used $P = 3$ and the autowave propagated through the system, although due to the minimal number of connections this propagation was not uniform. In the image it is seen that the autowaves from the two objects did collide only when $P = 3$.

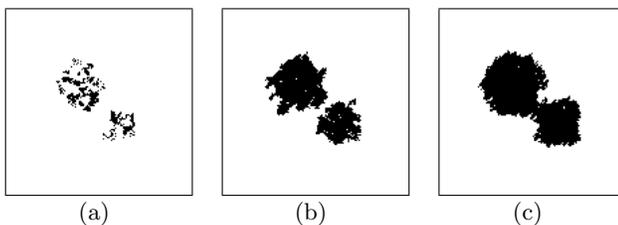


Figure 5.13: Neuron that fired in the first 10 iterations for systems with $P = 1$, $P = 2$, and $P = 3$ [2].

The conclusion is that at least three connections between neurons are needed in order to generate the autowave. However, for image processing applications the imperfect propagation should be avoided as it will artificially discriminate the importance of parts of the

image over others.

Another desire is that the autowaves emanate as a circular wave front rather than a square front. If the system only contained 4 connections per neuron then the wave would propagate in the vertical and horizontal directions but not along the diagonals. The propagation from any solid shape would eventually become a square and this is not desired. Since the input image will be defined as a rectangular array of pixels the creation of a circular autowave will require more neural connections. This circular emanation can be created when each neuron is connected to two layers of nearest neighbors. Thus, $P = 24$ seems to be the minimal system.

5.2.2 ICM Theory

Thus, the minimal system now consists of two coupled oscillators, a small number of connections, and a nonlinear function. This system is described by the following three equations [94],

$$F_{ij}[n + 1] = fF_{ij}[n] + S_{ij} + W\{\mathbf{Y}[n]\}_{ij}, \quad (5.12)$$

$$Y_{ij}[n + 1] = \begin{cases} 1 & \text{if } F_{ij}[n + 1] > \Theta_{ij}[n] \\ 0 & \text{Otherwise} \end{cases}, \quad (5.13)$$

and

$$\Theta_{ij}[n + 1] = g\Theta_{ij}[n] + hY_{ij}[n + 1]. \quad (5.14)$$

Here the input array is \mathbf{S} , the state of the neurons are \mathbf{F} , the outputs are \mathbf{Y} , and the dynamic threshold states are Θ . The scalars f and g are both less than 1.0 and $g < f$ is required to ensure that the threshold eventually falls below the state and the neuron pulses. The scalar h is a large value that dramatically increases the threshold when the neuron fires. The connections between the neurons are described by the function W and for now these are still the $1/r$ type of connections. A typical example is shown in Figure 5.14.

Distinctly the segments inherent in the input image are displayed as pulses. This system

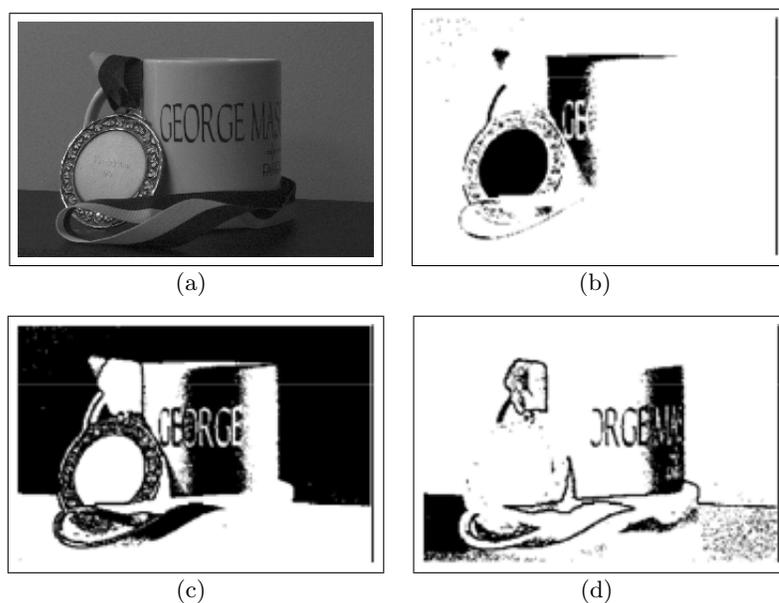


Figure 5.14: An input image and a few of the pulse outputs from the ICM [2].

behaves quite similar to the PCNN and is done so with simpler equations. Comparisons of the PCNN and the ICM operating on the same input are shown in Figures 5.15 and 5.16.

Certainly, the results do have some differences, but it must be remembered that the goal is to develop an image processing system. Thus, the results that are desired from these systems is the extraction of important image information. It is desired to have the pulse images display the segments, edges, and textures that are inherent in the input image.

5.2.3 Connections in the ICM

The function $W()$ manages the connections between neurons as did the \mathbf{M} and \mathbf{W} matrices in the PCNN. However, the PCNN connections were static whereas the ICM connections are dependent upon the pulse activity of the previous iteration. Both the PCNN and ICM do rely on only local connections. Before the function $W()$ is explained the foundation for this dynamic nature of the connections is explained.

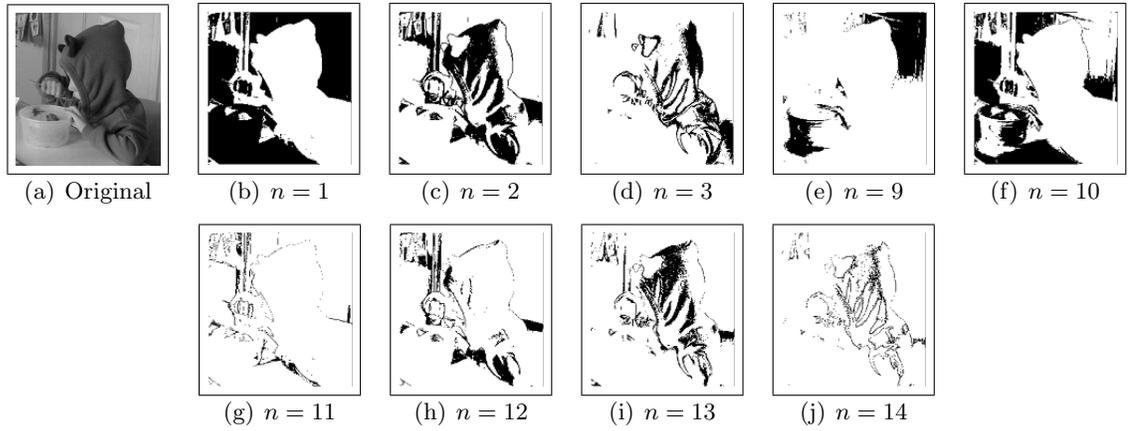


Figure 5.15: An original image and several selected pulse images from the PCNN [2].

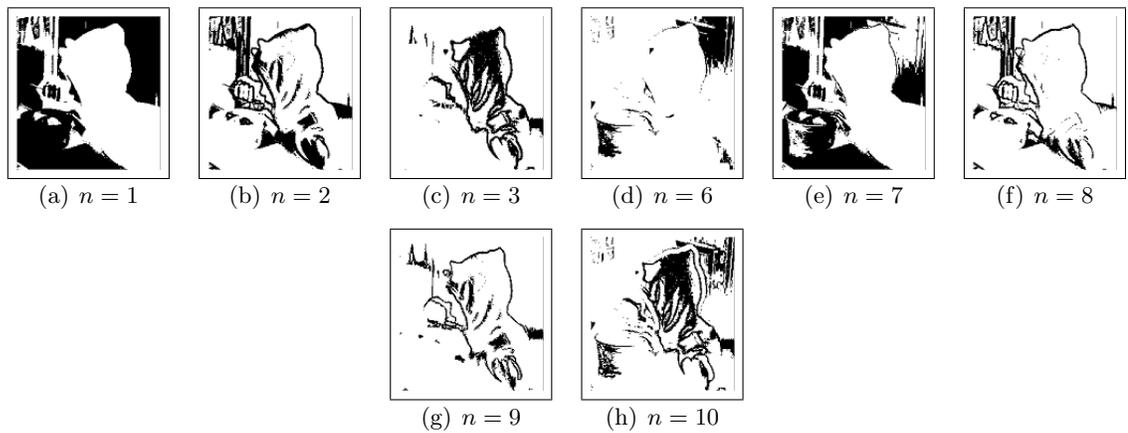


Figure 5.16: Results from the ICM [2].

5.2.4 Interference

In the PCNN model the connections are proportional to $1/r$ and static. This lead to the expanding waves seen in Figure 5.2. The expanding nature of the waves caused an *interference* problem when the PCNN was applied to images with multiple objects. As with wave interference of any kind, the waves expanding in Figure 5.4 are autowaves and so when they collide the wavefronts are annihilated. This means that the presence of the second object interferes with the expanding wave of the first object.

These expanding autowaves are the root cause of interference. Ideally, the autowaves expanding from non-target objects should not alter the autowaves emanating from target objects. But with strongly propagating auto waves, they do. If the non-target object is brighter it will pulse earlier than the target object autowaves, and its autowave can pass through the target region before the target has a change to pulse. The values of the target neurons are drastically altered by the activity generated from non-target neurons. Thus, the pulsing behavior of on-target pixels can be seriously altered by the presence of other objects.

Consider the input image in Figure 5.17 in which the target (a flower) is pasted onto a background. The target was intentionally made to be darker than the background to amplify the interference effect. Two inputs \mathbf{F}_1 and \mathbf{F}_2 were created where \mathbf{F}_1 was the image shown in Figure 5.17 and \mathbf{F}_2 was only the flower without a background. The pulse images from a modified ICM were computed for each input. The modification was to use the connections \mathbf{M} that were also used in the PCNN. The signatures restricted to only the on-target pixels of these two trials are shown in Figure 5.18. As seen the two signatures are very different and thus the pulse activity for the neurons on-target is different solely due to the presence of the background. It would be quite difficult to recognize an object from the neural pulses if those pulses are so susceptible to the content of the background.

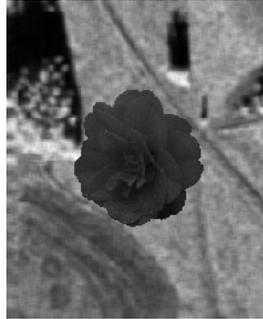


Figure 5.17: A target (flower) pasted on a background [2].

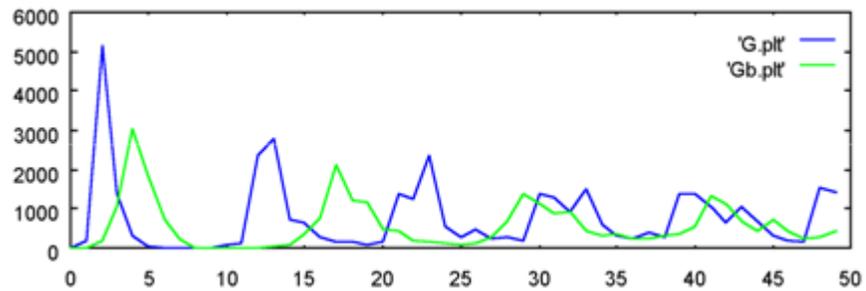


Figure 5.18: The signature of the flower without a background (G.plt) and the signature of the flower with a background (Gb.plt) [2].

5.2.5 Curvature Flow

The solution to the interference effect is based on curvature flow theory [95]. In this scenario the waves propagate towards the centripetal vectors that are perpendicular to the wavefront. Basically, they propagate towards local center of curvatures. For solid 2D objects the curvature flows will become a circle and then collapse to a point [96]. (There is an ongoing debate as to the validity of this statement in dimensions higher than two.)

Such propagation from Malladi & Sethian [95] is shown in Figure 5.19. The initial frame presents a intricate 2D shape. This figure will eventually evolve into a circle and then collapse to a point. There is a strong similarity between this type of propagation and the propagation of autowaves. In both cases the wavefront will evolve to a circle. The difference is that the autowaves will also expand the circumference with each iteration whereas the curvature flow will be about the same size as the original shape.



Figure 5.19: The propagation of curvature flow boundaries [2].

The interference in the ICM that lead to the deleterious behavior in Figure 5.18 was caused when the neural communications of one object interfered with the behavior of another. In other words, the autowaves from the background infringed upon the territory owned by the flower. This stems from the ever expanding nature of the autowaves.

Curvature flow models evolve to the same shape as autowaves but do not have the ever-expanding quality. Thus, the next logical step is to modify the connection function $W()$ to behave more like curvature flow wavefronts.

5.2.6 Centripetal Autowaves

A centripetal autowave follows the mechanics of curvature flow. When a segment pulses its autowave will propagate towards a circle and then collapse. It does not propagate outwards as does the traditional autowave. This results in autowaves developed from two neighboring objects that will have far less interference.

The propagation of a curvature flow boundary is towards the local center of curvature. The boundary, C , is a curve with a curvature vector $\vec{\kappa}$. The evolution of the curve follows,

$$\frac{\partial C}{\partial t} = \vec{\kappa} \cdot \hat{n}, \quad (5.15)$$

where n is normal. In two-dimensional space all shapes become a circle and then collapse to a point. Such a progression is shown in Figure 5.20 where a curve evolves to a circle and then collapses.

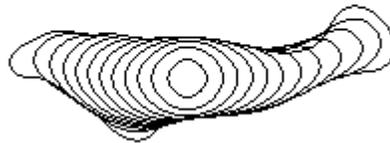


Figure 5.20: The progression of an autowave from an initial shape [2].

The ever-expanding nature of the autowaves leads to the interference and this quality is absent in a curvature flow model. Thus, the logical step is to modify the neural connections to behave as in the curvature flow model. This requires that the connections between the neurons be dependent upon the activation state of the surrounding neurons. However, in creating such connections the problem of interference is virtually eliminated. In this new

scenario neural activity for on-target neurons is the same independent of the presence of other objects. This is a major requirement for the employment of these models as image recognition engines.

The new model will propagate the autowaves towards the local center of curvature and thus obtain the name *centripetal autowaves*. The computation of these connections requires the re-definition of the function $W()$.

Computations for curvature can be cumbersome for large images. Hence, for images, a less computationally intense method is selected. The curves in figure start with the larger, intricate curve and progress towards the circle and then collapse to a point. The neural communications will follow this type of curvature flow progression. Of course, in the ICM there are other influences such as the internal mechanics of the neurons which influence the evolution of the neural communications.

The function $W(A)$ is computed by,

$$W(A) = [[F_{2,A'} \{M(A')\} + F_{1,A'}(A')] < 0.5], \quad (5.16)$$

where,

$$A' = A + [F_{1,A} \{M(A)\} > 0.5]. \quad (5.17)$$

The function $M(A)$ is a smoothing function. The function $F_{1,A} \{X\}$ is a masking function that allows only the pixels originally OFF in A to pass as in,

$$[F_{1,A}]_{ij} = \begin{cases} X_{ij} & \text{if } A_{ij} = 0 \\ 0 & \text{Otherwise} \end{cases}, \quad (5.18)$$

and likewise $F_{2,A}\{X\}$ is the opposing function,

$$[F_{2,A}]_{ij} = \begin{cases} X_{ij} & \text{if } A_{ij} = 1 \\ 0 & \text{Otherwise} \end{cases}, \quad (5.19)$$

The inequalities are passing thresholds as in,

$$[X > d]_{ij} = \begin{cases} 1 & \text{if } x_{ij} \geq d \\ 0 & \text{Otherwise} \end{cases}, \quad (5.20)$$

and

$$[X < d]_{ij} = \begin{cases} 1 & \text{if } x_{ij} \leq d \\ 0 & \text{Otherwise} \end{cases}, \quad (5.21)$$

This system works by basically noting that a smoothed version of the original segment produces larger values in the off-pulse region and lower values in the on-pulse region in the same areas that the front is to propagate. The nonlinear function isolates these desirable pixels and adjusts the communication wavefront accordingly.

The centripetal autowave signatures of the same two images used to generate the results in Figure 5.18 are shown in Figure 5.21. It is easy to see that the background no longer interferes with the object signature. The behavior of the on-target neurons are now almost independent of the other objects in the scene. This quality is necessary for image applications.

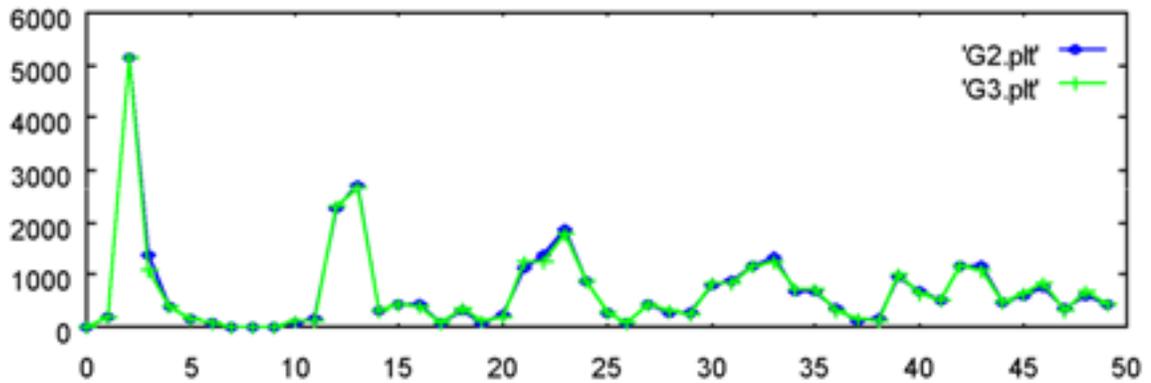


Figure 5.21: The signatures of the flower and the flower with a background using the centripetal autowave model [2].

5.3 Summary

For quite some time, cortical models have been expressed in mathematical form. Coupled oscillators or reaction-diffusion systems still apply to current models. Furthermore, in an image processing application the differences between the different models may not be that important. Therefore, speed and simplicity of implementation are more important here than replication of a biological system [2].

For image processing applications the model selected here is the ICM which is built on just three simple equations. Each neuron has two oscillators (the neuron potential and the neuron threshold) and each neuron has a non-linear operation. Thus, when stimulated, each neuron is capable of producing a spike sequence, and groups of locally connected neurons have the ability to synchronize pulsing activity. When stimulated by an image, these collectives can represent inherent segments of the stimulating image [2]. Thus, a cortical model can become a powerful first step in many image processing applications.

The traditional neural connection schemes, however, do not impede the neural communications, allowing it to continually progress away from the originating region. While this may be true in biological neurons, this property causes interference. This has been found to be deleterious to object recognition. Activity from one region can so drastically alter

the activity in another region that object recognition becomes very difficult. The solution to this problem is to alter the connections to the neurons so that they become sensitive to previous pulsing steps [2]. In the model presented, these connections are described as centripetal autowaves such that the wavefront progresses towards the local center of curvature of the pulsing regions. This eliminates the ever-expanding nature of the waves without altering their shape-describing form.

5.4 Comparing ICM with Gabor

This comparison justifies the use of ICM as opposed to Gabor filters. In the Gabor study a set of Gabor filters were created and each correlated with the input image. This created a data cube of correlation surfaces and training vectors (jets) were extracted. This study was replicated and the training vectors for honeycomb, vascular, and normal regions were extracted. These were then remapped in PCA space and the result is shown. Consider the results shown in Figure 5.22.

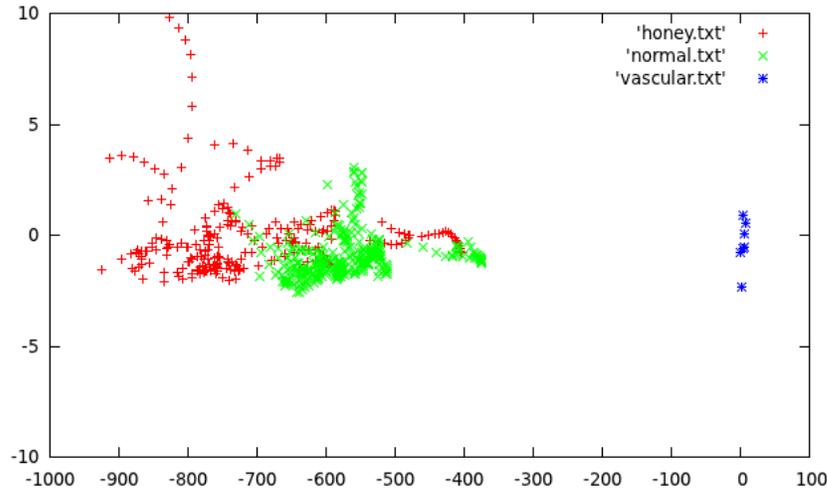


Figure 5.22: PCA on extracted vectors from Gabor filter correlations.

Clearly, the vascular training vectors are well separated from the others but honeycomb

and normal have overlapping distributions. This makes recognition difficult.

In the second study the ICM generated pulse images which were collected in a data cube and smoothed. The training vectors were extracted from this cube for several patients and again remapped by PCA. The results are shown in Figure 5.23. In a very clear fashion the

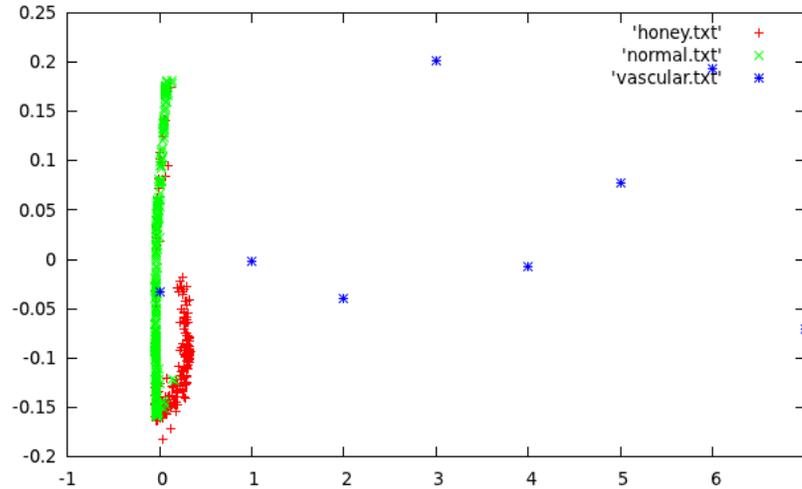


Figure 5.23: PCA on extracted vectors from pulse images.

normal and honeycomb regions showed a separation. This is precisely the reason that the ICM version is working better than the Gabor system.

5.5 Segmentation of IPF Images with Pulse Images

Each CT image is converted to a set of pulse images which through collective synchronization of pixels extract pertinent information of the diseased regions. These pulse streams are used for training and recall through an associative memory so that entire images can be segmented and analyzed. The following sections presents the algorithms and results for the analysis of patients with IPF and normal patients. Results demonstrate that segmentation of IPF images is useful in extracting quantitative information.

5.5.1 System Overview

Consider a typical image in Figure 5.24. A patchy pattern of peripheral, sub-pleural and predominantly lower lobe reticular opacities combined with honeycombing, traction bronchiectasis and absence of significant ground glass attenuation constitute the classic features of IPF.

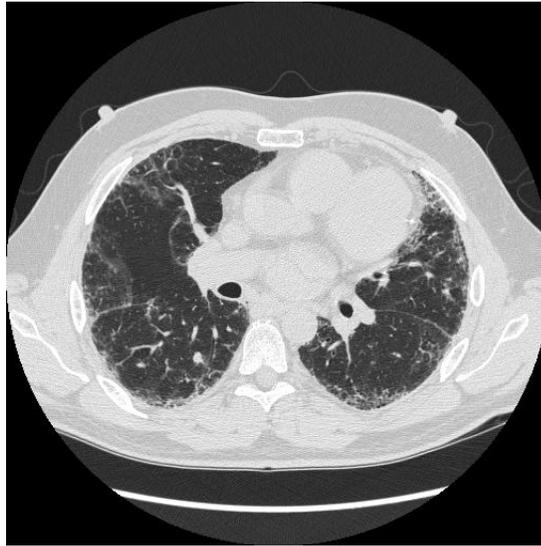


Figure 5.24: Original image of a patient with IPF.

Figure 5.25 shows an overview of the data flow for this system. The upper path depicts the training and the lower path depicts the recall system. There are two main engines which will be described in the following sections. The intersecting cortical model (ICM) is the engine that generates the pulse images and the fast analog associative memory (FAAM) is the recall engine. Each step is detailed in the next sections with examples.

5.5.2 Pulse Image Generator

The pulse stream generator has its roots in early models of the mammalian visual cortex [84]. The digitization of this model was popularized as the PCNN (pulse-coupled neural

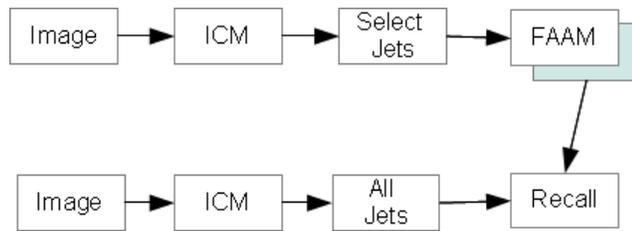


Figure 5.25: Overview of the system with a training track (upper path) and a query track (lower path).

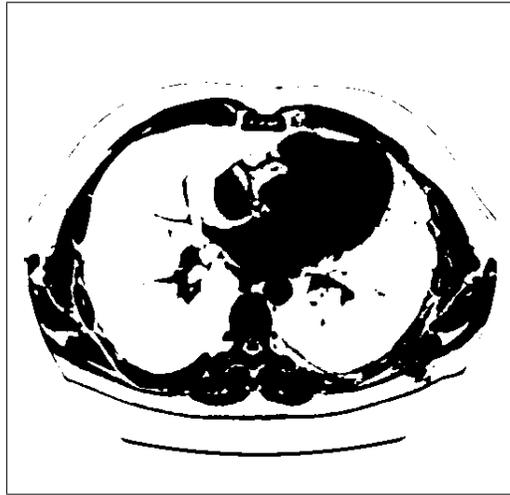
network) [86]. The PCNN contained a 2D array of neurons which were associated with the 2D array of pixels from an input image. Each neuron contained three compartments.

The first two were the feeding and linking compartments which were combined in a second order fashion to generate the neuron’s potential. This potential was then compared to a dynamic threshold (the third compartment) to produce a binary output. If the neuron fires, or produces a high output, then the threshold is raised to a significant value and then begins a decay process. The PCNN was used for many image processing applications. [2].

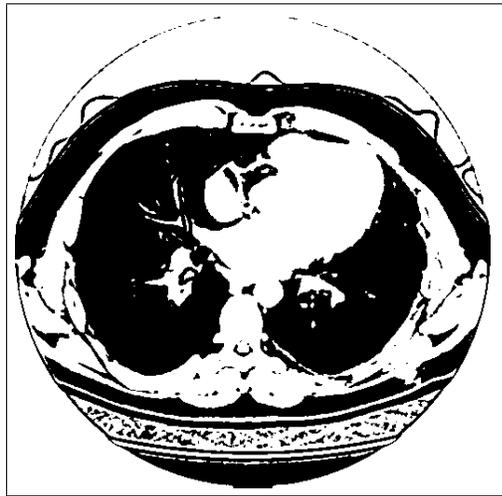
5.5.3 ICM in Texture Analysis

The ICM has been used successfully in texture analysis of medical images [97,98]. It is a modification of the PCNN which reduces the computational complexity. More importantly, it modifies the method in which the neurons communicate with each other. In the PCNN design when a neuron fires it sends a positive signal out to the neighboring neurons with the strength of the signal inversely proportional to the distance between the neurons.

In terms of image processing this type of connection caused problems and so the ICM adopted a centripetal autowave scheme [99] in which the connections to the neurons are dependent on the pulse pattern that surrounds them. The use of autowaves prevents the neural activity of one pixel from expanding too far from its location. In the original PCNN design the propagation of neural communications was not restricted and interference amongst objects within the image was significant and deleterious for image processing applications.



(a) Y at $n = 1$.



(b) Y at $n = 2$.

Figure 5.26: Two pulse images from the first cycle.

The premise of the ICM is that each pixel in an input image $\{\mathbf{S} : S_{i,j}; i = 1, \dots, V, j = 1, \dots, H\}$, where V and H are the vertical and horizontal dimensions of the image, is associated with a neuron. In the ICM the potential of the neurons are computed by,

$$F_{ij}[n+1] = fF_{ij}[n] + S_{ij} + W\{\mathbf{Y}[n]\}_{ij}, \quad (5.22)$$

where \mathbf{Y} is the output of the previous iteration, f is a decay constant less than 1, h is a strength factor, and $W\{\cdot\}$ is a centripetal autowave generator that modifies the neural connections within each generation according to level set theory [95]. The output is computed by,

$$Y_{ij}[n] = \begin{cases} 1 & F_{ij}[n] > \Theta_{ij}[n] \\ 0 & \text{Otherwise} \end{cases}, \quad (5.23)$$

and the dynamic threshold is computed by,

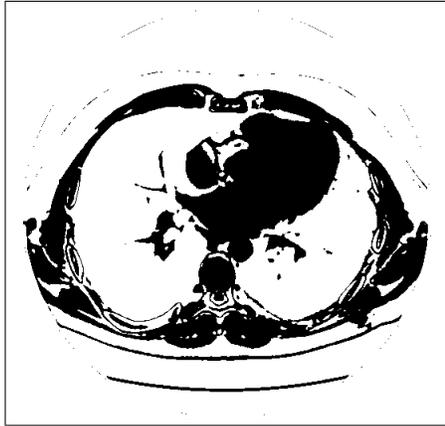
$$\Theta_{ij}[n] = g\Theta_{ij}[n-1] + hY_{ij}[n], \quad (5.24)$$

where $g < 1$ and $h > 1$.

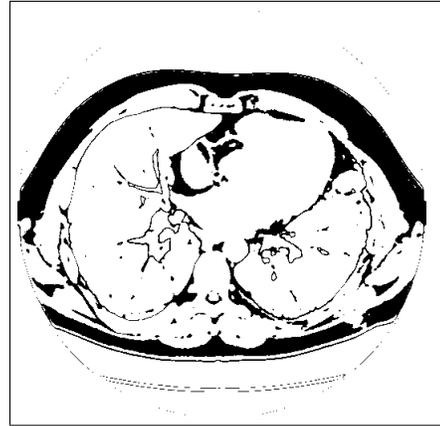
5.5.4 Pulsation in Cycles

The neurons in the ICM pulse in cycles. Figure 5.26 shows two of the pulse images from the first cycle, $Y[1]$ and $Y[2]$ where there is no distinction between the fibrosis and the healthy tissue. In the second cycle as shown in Figure 5.27, the neurons inside of the lung no longer fire in unison. As seen in Figure 5.27(d) the healthy and fibrotic portions of the image are firing at different times in a manner that is sensitive to the inherent textures.

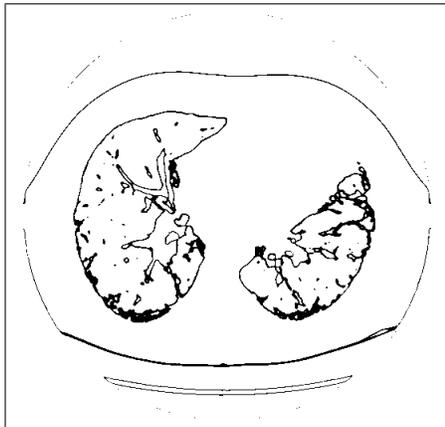
The third cycle is shown in Figure 5.28 which shows a further separation of healthy and fibrotic tissues. This stack of pulse images creates a pulse stream for each pixel. For pixel (i, j) the pulse stream is $\tau\{Y_{i,j}[k]\}$, $\forall k$ where $\tau\{\cdot\}$ is a local smoothing function. This



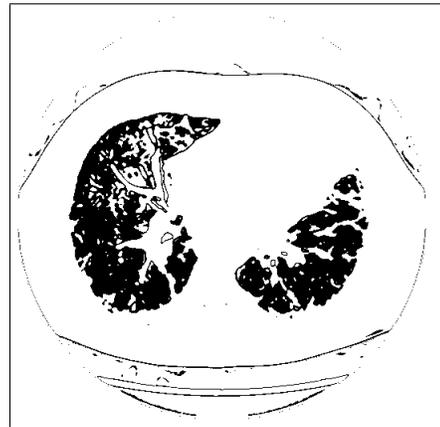
(a) Y at $n = 8$.



(b) Y at $n = 9$.



(c) Y at $n = 10$.



(d) Y at $n = 11$.

Figure 5.27: Four pulse images from the second cycle.

function is used to eliminate single pixel noise and to accommodate regions with slight differences.

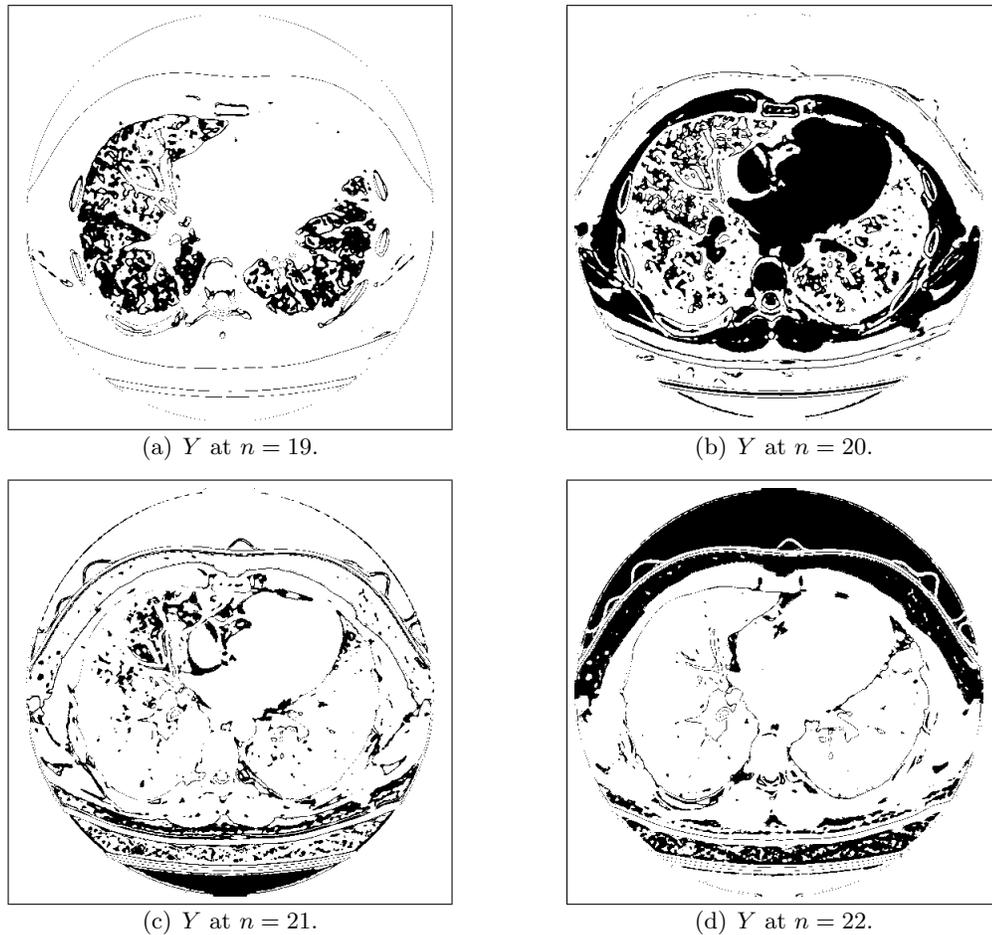


Figure 5.28: Four pulse images from the third cycle.

5.5.5 Pulse Streams for Classification

In order for the pulse streams to be useful in classification the diseased regions must behave similar to each other but different from the healthy regions. The same rule applies to the healthy region. Figure 5.29 shows two sets of error bars which represent the pulse response in the honeycomb regions and the healthy regions. Non-overlapping bars justify the use of

pulse streams for the representation of the healthy and diseased regions. The location of the error bars is the average intensity and the extent of the bars is the first standard deviation from the average.

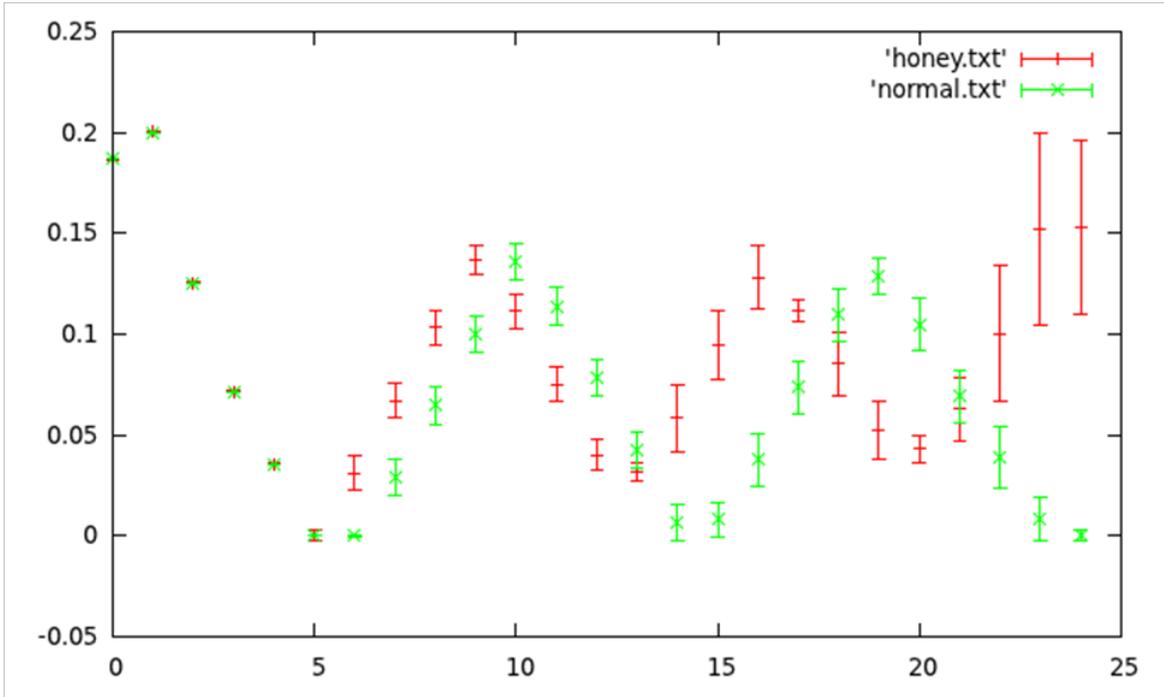


Figure 5.29: The x -axis represents the ICM iteration number, n , and the y -axis represents the pulse activity for the neurons

5.5.6 Analysis of Pulse Streams

As seen the first cycle does not separate the diseased and healthy regions. The second cycle begins at $x = 5$ and the diseased regions begin to pulse before the healthy regions. The trend continues through the third cycle. The fourth cycle for the disease regions begins at $x = 20$ and as seen the error bars increase significantly. The disintegration of the collective behavior has reached a point that is no longer useful and thus the number of iterations is terminated at $n = 20$.

Starting with the second cycle the healthy and diseased regions demonstrate unique pulse behaviors. Classification of these two regions via these pulse streams is very possible, and the use of pulse streams is justified.

5.6 Associative Memory

The next step in the training process in Figure 5.25 is the extraction of classified jets. These are pulse streams that are manually selected from fibrotic and healthy regions. The fibrotic jets were extracted from various diseased regions in the image. These are used for training in the next step. The learning algorithm FAAM (fast analog associative memory) [100] is employed to perform the training and recall operations. This is a very simple greedy training algorithm that has been successfully applied in other medical image applications. [101].

5.6.1 FAAM Demonstration

Thus, to create a system that recognizes both diseased regions and healthy regions two separate FAAMs are constructed. Figure 5.30 demonstrates the training process. Consider a case in which data vectors are associated with two classes (represented by the two markers). In Figure 5.30(a) there are only two training vectors and they have an opposing class. There does exist a boundary that crosses the line segment that connects the two points (not shown). However, the shape and location of the boundary is not known and there is no other information about this boundary. With this lack of information the boundary is defined as the bisector of the segment that joins the two data points.

As more data vectors are considered the FAAM determines if new boundaries are needed. Figure 5.30(b) shows the case in which a new vector is added and it coincidentally is on the correct side of the boundary. No training is required. Figure 5.30(c) shows the case in which a data vector is added and a new boundary is required. The FAAM sequentially considers the data vectors and creates the boundaries as necessary.

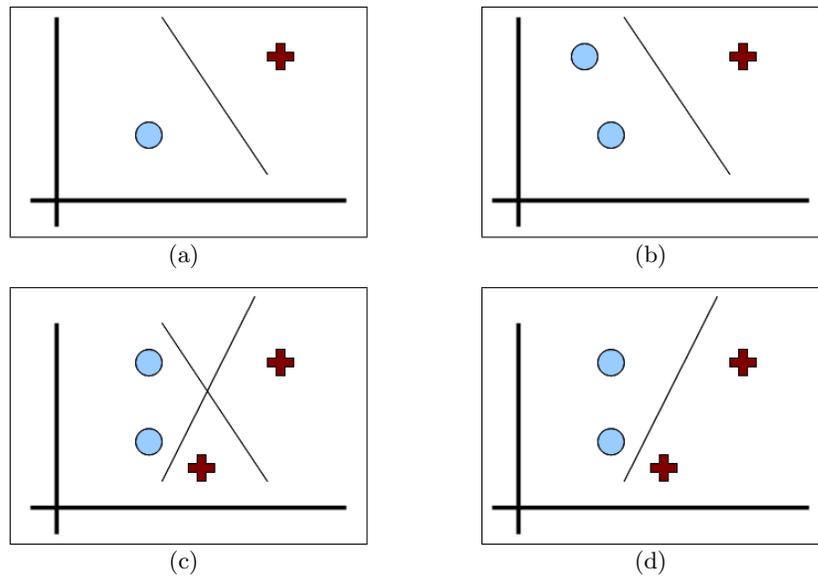


Figure 5.30: Demonstration of the FAAM.

As new boundaries come into existence older boundaries may become obsolete. This is shown in Figure 5.30(d) in which the first boundary is removed since the new boundary separates all of the known data. The FAAM optimizes the separation of the classes by creating and destroying boundaries. Unlike artificial neural networks the boundaries are not moved. The attempt to optimize boundary locations through movement is actually an expensive computation. The FAAM gains its speed by optimization via creation and destruction.

Furthermore, as the FAAM training progresses the number of boundaries is indicative of the training process. Certainly, if the number of boundaries increases as more data is considered then the older boundaries were insufficient. Training was not accomplished. If, however, the number of boundaries stabilizes then the FAAM has found boundaries that separates the two classes and confidence in the training is reached.

It is possible that as boundaries are created that pockets appear in the mapping space that are not defined. There are not examples of either class that exists on this particular area defined by the boundaries. Thus, in the recall stage there are three types of outputs. The

FAAM can classify the input as target, non-target, or unknown. There are modifications to the FAAM to address this issue but since it is so minor in this application the modifications were not employed.

The small subset of pulse streams are used to train the FAAM and each pulse stream from an image will be classified by the FAAM. The application is shown in the next section.

5.6.2 Training the FAAM

The FAAM system creates and destroys decision boundaries when necessary. Thus training is performed by the survival of the fittest boundaries. The FAAM considers training data sequentially. As the FAAM training progresses the number of boundaries is indicative of the training process.

If the number of boundaries increases as more data is considered then the older boundaries are insufficient. Training is not accomplished. If, however, the number of boundaries stabilizes then the FAAM has found boundaries that separates the two classes and confidence in the training is reached. If a network learns the problem then a leveling of the number of decisions will occur.

Therefore a network performance is demonstrated by plotting the current number of decision surfaces versus the number of training vectors considered. Three separate FAAMs: healthy, honeycomb and vascular are constructed and the randomly selected jets are trained by the network. The results are shown in Figures 5.31.

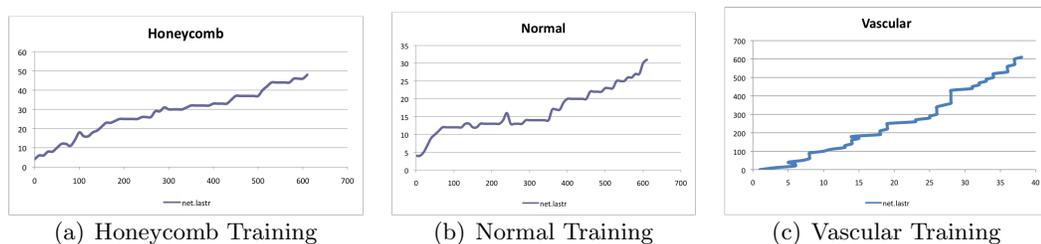


Figure 5.31: Learning curves from a typical FAAM for three different regions in the same patient.

5.7 Application to a Single Patient

In this section a single patient is considered. There are 24 CT images that accompany this scan and they include the full height to the lungs.

5.7.1 Single Image Response

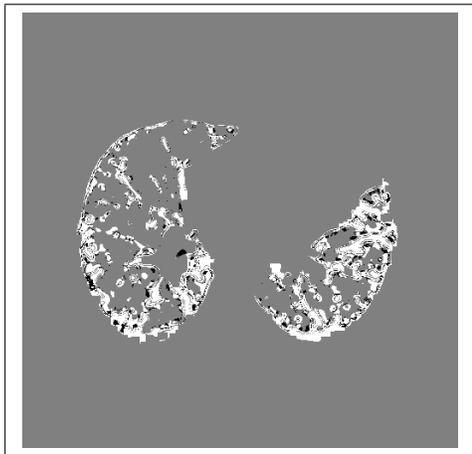
Figure 5.24 is image number 14 from the scans and it is located in the lower half of the lung. Since the honeycomb is more prevalent in the lower lung in IPF patients these images are more useful. Figures 5.26, 5.27, and 5.28 show some of the pulse images from this input. The others are not shown since they are less informative. Figure 5.29 shows the behavior of the pulses for all iterations.

About 30 pixels from diseased and healthy regions from two CT images were manually selected and the pulse streams of these pixels were used to train the fibrotic FAAM and the healthy FAAM (denoted now as fFAAM and hFAAM). The pulse stream from every pixel was then used as a query and the response from both FAAMs were collected. Figure 5.32(a) shows the response from the fFAAM. Pixels that are white are classified as diseased. Pixels that are gray (and inside the lung region) are classified as not diseased. The few black pixels are classified as unknown. The regions outside of the perimeter of the lung were not classified. Figure 5.32(b) shows the response from the hFAAM in which the white pixels are classified as healthy.

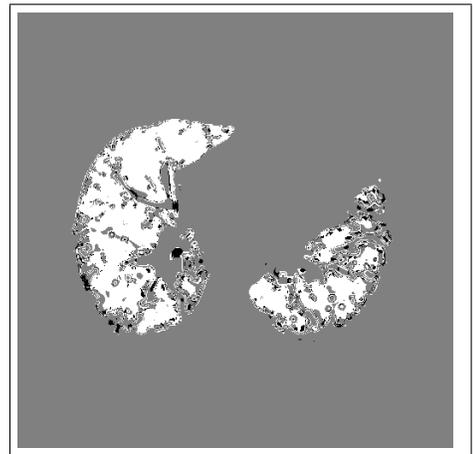
The response of these networks corresponds with analysis of the images by physicians specializing in lung disease. At this time there is no method to exactly quantify the response of the networks.

5.7.2 Non-Training Image Response

The training in this example used pulse streams from images 8 and 14 for training. Figure 5.33(a) shows image 11 from which no training pulse streams were extracted. Figure 5.33(b) shows the classification for the fFAAM and Figure 5.33(c) shows the classification of healthy regions from hFAAM.

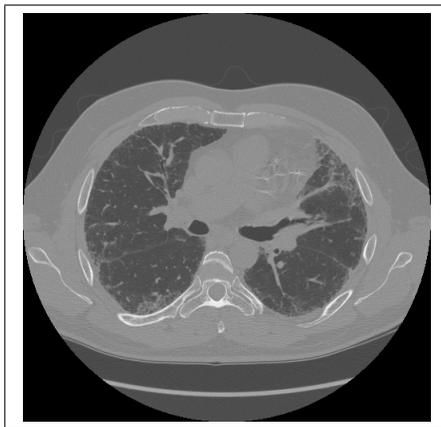


(a) Response from fFAAM.

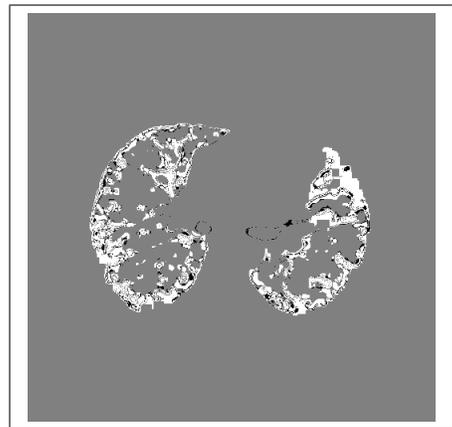


(b) Response from hFAAM.

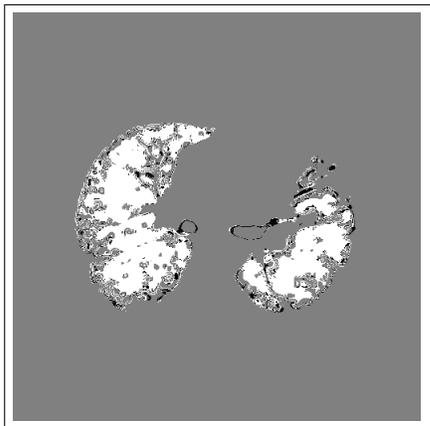
Figure 5.32: The response from image 14 from both FAAMs.



(a) Original image.



(b) Response from fFAAM.



(c) Response from hFAAM.

Figure 5.33: The response from image 11 from both FAAMs.

5.8 Analysis

Classification of the fibrotic and healthy regions allows for analysis of the patient's state and progression of the illness. Of importance is the volume of the lung that has become diseased compared to that that remains healthy. Figure 5.34 shows classification areas for this sample patient.

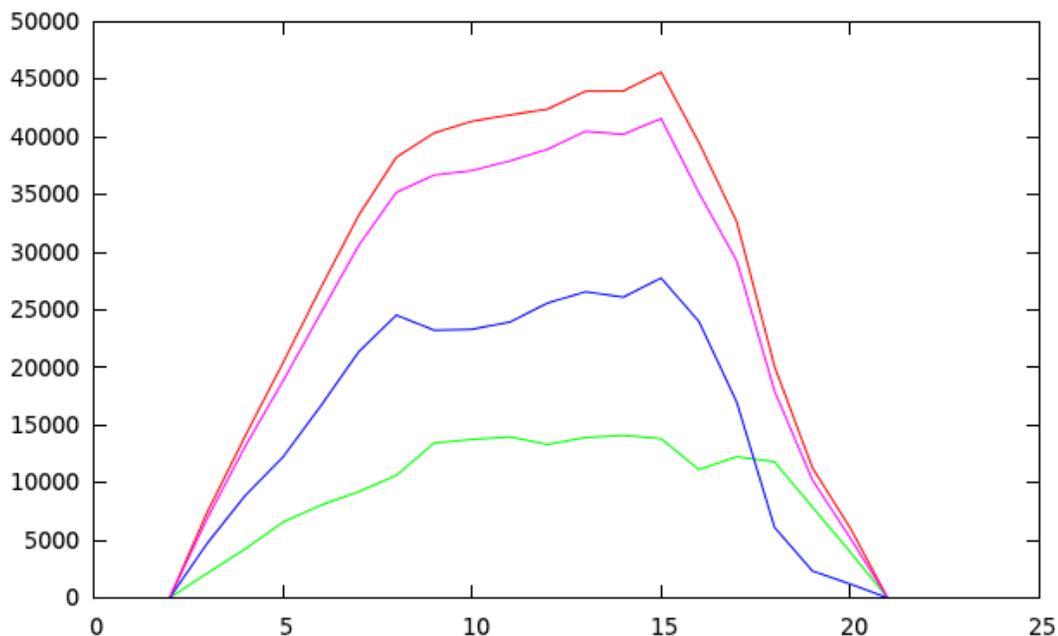


Figure 5.34: Area metrics of classification. The top most graph depicts the total area of the lung for each slice. The lower most graph depicts the area classified as fibrotic and the graph above it depicts the area classified as healthy. The second plot depicts the summation of the healthy and diseased regions.

The x -axis corresponds to the CT slide number. The range of x corresponds to scans from the top of the lung towards the bottom of the lung. There are a few images above and below the lung in the data set. The top of the lung begins at $x = 2$ and the bottom of the lung ends at $x = 21$. The difference between it and the total lung volume stems from those pixels that were classified as unknown or tissue.

The percentage of fibrotic pixels to healthy tissues increases towards the bottom of the

lung in congruence with the progression of the diseased starting at the bottom of the lungs. Of note is also the cross over point at $x = 18$ where the area of diseased lung exceeds the area of healthy lung. In this patient the total lung volume classified as diseased was 37%.

5.9 Multiple Patients

Three CT images are considered from three distinct patients at different stages of IPF. There are apparent textural and contrast differences seen in image numbers 22, 33 and 39. The response is recorded from multiple patients.

5.9.1 Normalization

To successfully and easily apply the algorithm across multiple patients, automatic histogram normalization is necessary. With this in mind the normalization algorithm is best described using the Figure 5.35(a) and its corresponding histogram 5.35(b) as shown below.

Find the highest peaks right at the beginning which correspond to the background region. See annotated Figure 5.36. The images contain low contrast contours. The histogram modes corresponding to the different regions such as bone, tissue, air and background regions (1) and (2), vascular, diseased and healthy often overlap.

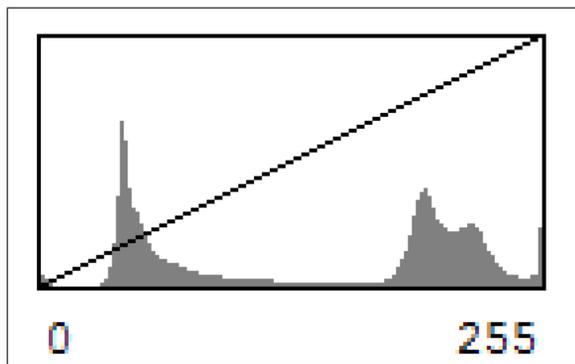
Therefore any leading edge peak corresponding to the non-diseased/non-healthy region is removed. The two remaining highest peaks located in the bimodal histogram are retained and the image is now ready for training purposes. To further reduce noise, additional pre-processing is done using a smoothing spline function. Once the image has been normalized, we apply the Gaussian fit and use the mean and standard deviation to scale the image.

5.9.2 Training and Recall

For the purpose of training we select three CT images from three different patients seen in Figure 5.37 which cover a broad spectrum of disease stages of IPF. Next the healthy and fibrotic FAAMs are trained. The pulse stream from the test images from another patient is then used as a query and response from both FAAMs are collected as seen in Figure 5.38.



(a) Original image.



(b) Corresponding Histogram

Figure 5.35: Image and Histogram before normalization.

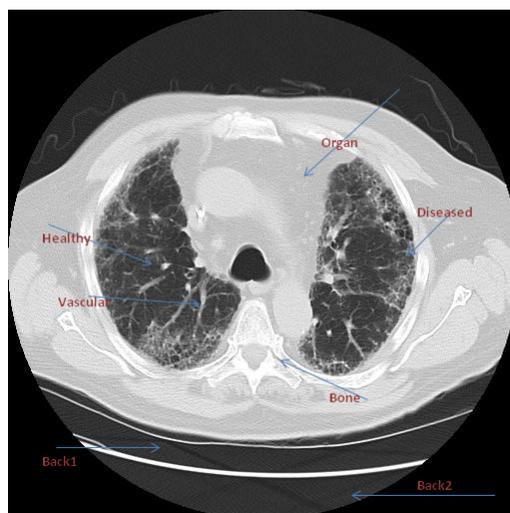


Figure 5.36: The image is divided into different regions(back1, back2, organ, tissue, bone, vascular, diseased and healthy).

5.9.3 Analysis

The test images as seen in Figures 5.39, 5.40, 5.41, 5.42 do not belong to the training series. In spite of the obvious difference in contrast, volume and texture the healthy and diseased regions are easily distinguished and classified.

5.10 Discussion

The technique used is successful in classifying the healthy and diseased portions of the lung. The goal was to train adequate and varied stages of IPF images and to be able to extract sufficiently enough information from test images. The algorithm was tested on HRCT scans procured through Inova Fairfax Hospital, Department of Radiology. Two expert radiological reviewers compared the initial results of the segmentation algorithm with the manual segmentation of the original scans. Comparison revealed agreement regarding the presence or absence of honeycombing. However for achieving a high and reliable performance for non-training cases, a large number of training cases are commonly required.



(a) HRCT Scan from Patient 22

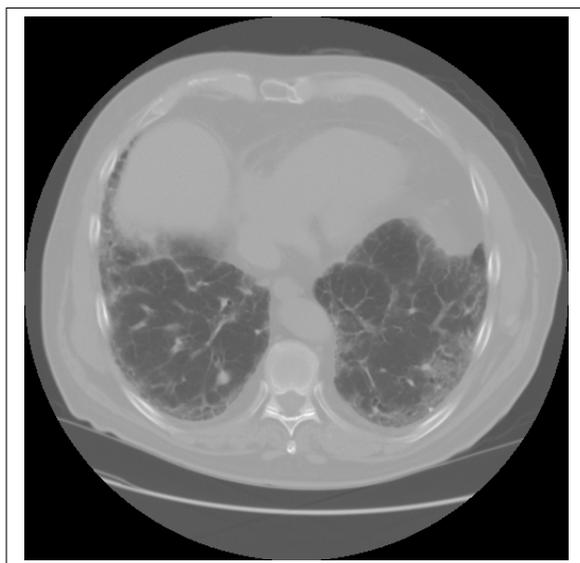


(b) HRCT Scan from Patient 39

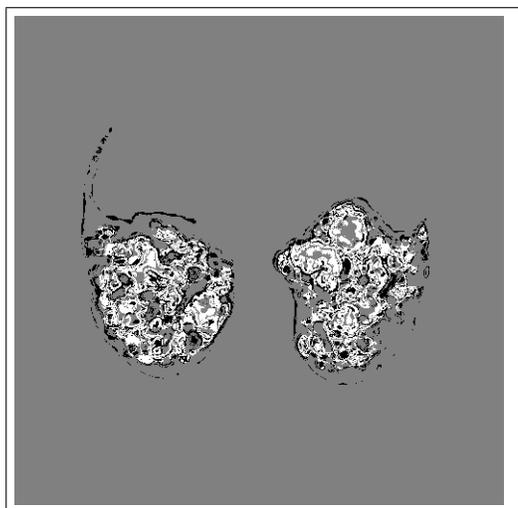


(c) HRCT Scan from Patient 17

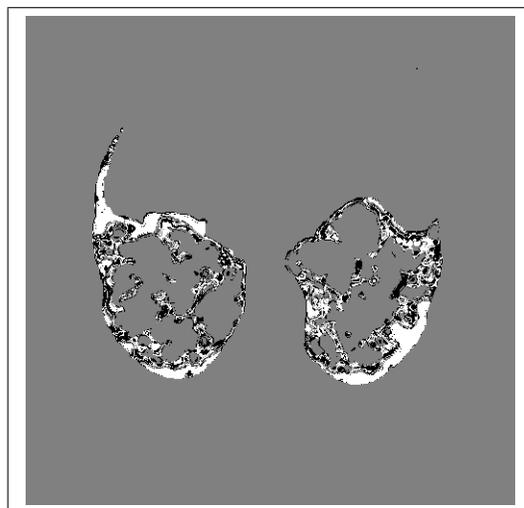
Figure 5.37: Training scans from multiple patients.



(a) HRCT Scan from Patient 33



(b) Response from hFAAM from Patient 33



(c) Response from fFAAM from Patient 33

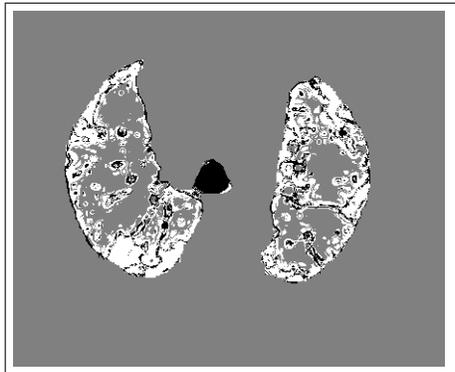
Figure 5.38: Demonstration of the FAAM from multiple patients.

5.11 Enhancements and Conclusions

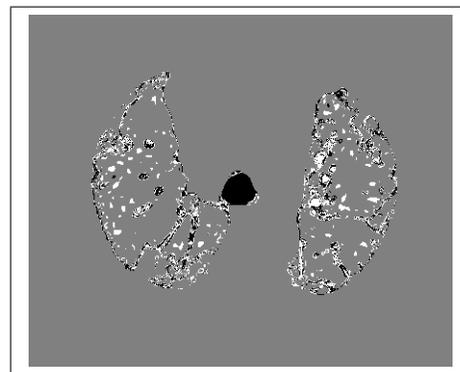
Past experiences with the FAAM indicate that in these circumstances the FAAM performs better if the data pairs are presented in random order. Given a data set $\{\mathbf{X} : \mathbf{Y}; \vec{x}_i : y_i; i = R(1, \dots, N)\}$ where \vec{x}_i is the i -th training vector which is associated with y_i and y_i is a binary classification ($y_i \in [0, 1]$). The value of N is the total number of training pairs. It is possible to present $\vec{x}_i : y_i$ in sequential order but in this case it is better to randomize the indices which is represented by the function $R(\cdot)$. The reason for this is the data contains subgroups ($i \in i_1, \dots, i_2$) in which the data vectors are extracted from the same region of a single image. Thus, these vectors are self-similar. The number of training cases were also increased and this aided the quality of segmentation as seen in Figures 5.39, 5.40, 5.41, 5.42.



(a) HRCT Scan



(b) Honeycombing

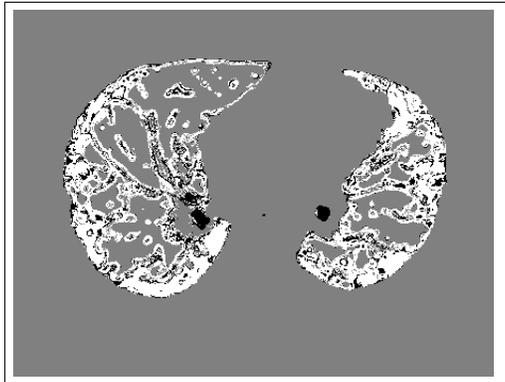


(c) Vascularization

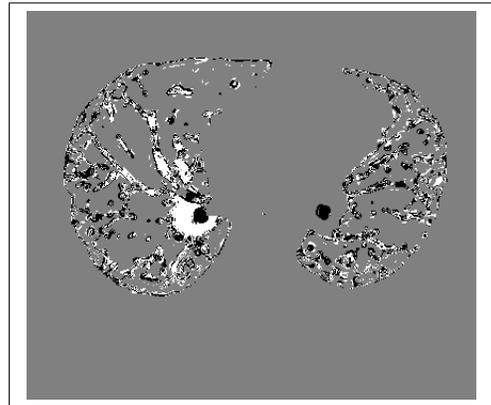
Figure 5.39: Honeycombing and Vascular regions from Patient a0017.



(a) HRCT Scan

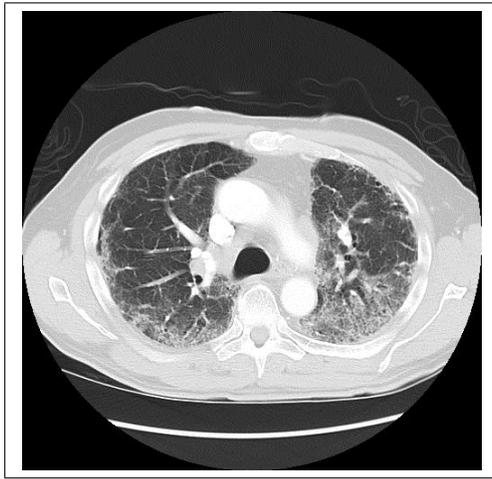


(b) Honeycombing

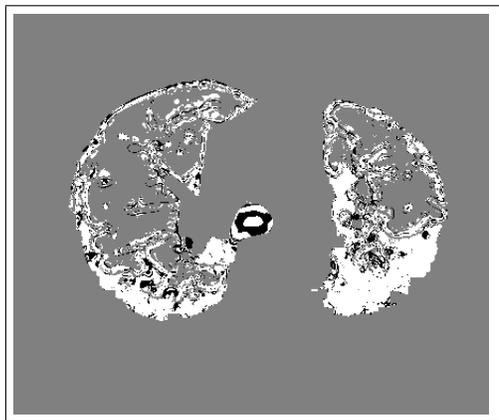


(c) Vascularization

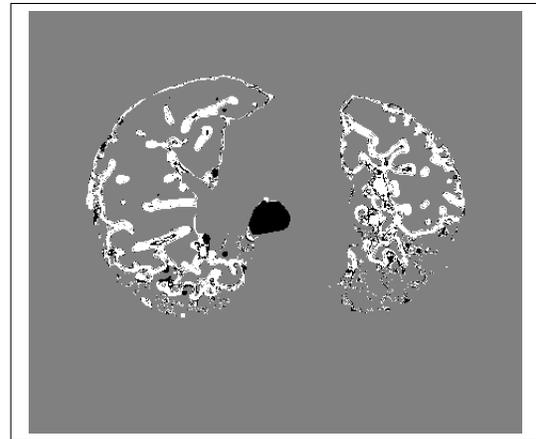
Figure 5.40: Honeycombing and Vascular regions from Patient d0034.



(a) HRCT Scan from Patient p0020

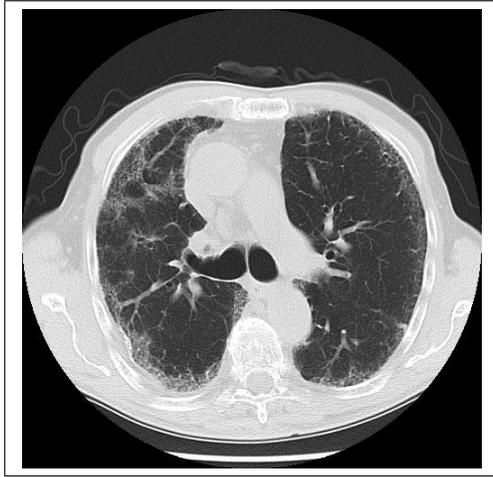


(b) Honeycombing

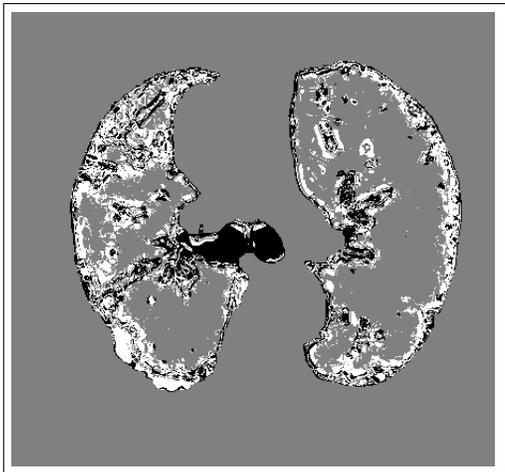


(c) Vascularization

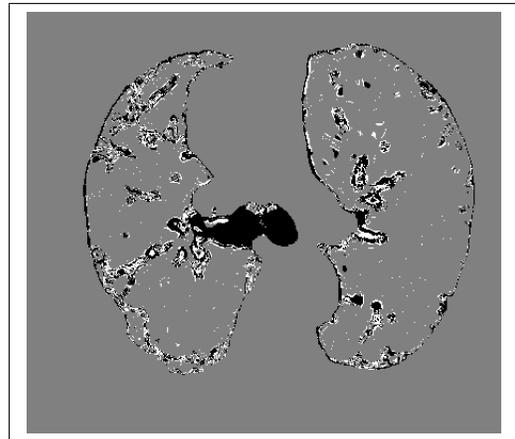
Figure 5.41: Honeycombing and Vascular regions.



(a) HRCT Scan



(b) Honeycombing



(c) Vascularization

Figure 5.42: Honeycombing and Vascular regions from Patient p0023 .

Chapter 6: Comparisons and Conclusions

The diagnostic approach to IPF is multidisciplinary, involving primary care physicians, pulmonologists, radiologists, and pathologists as seen in Figure 6.1. It begins with a complete clinical evaluation, including history, physical examination, chest radiograph, laboratory studies, and pulmonary physiologic testing [102]. After this thorough assessment, patients suspected of having an IIP should undergo a HRCT scan of their lungs. In some instances, an experienced chest radiologist can make a confident diagnosis of IPF or other diffuse lung disease without further diagnostic intervention. When the HRCT results are unclear, many patients will proceed directly to surgical lung biopsy. Ultimately, surgical lung biopsy is the gold standard in diagnosing IPF.

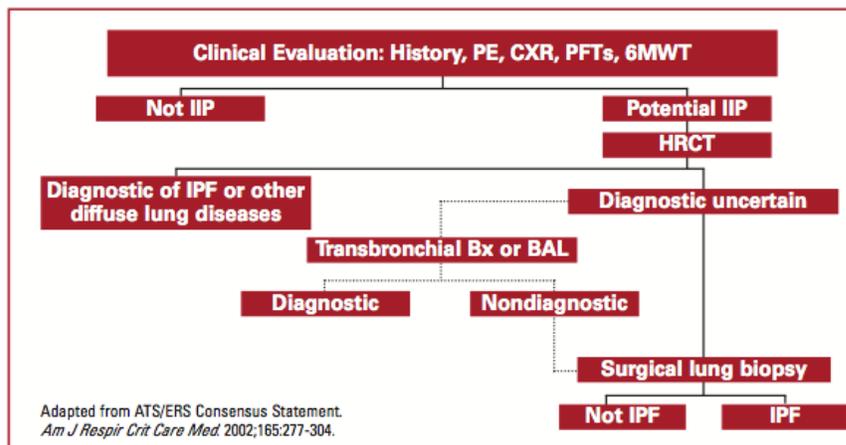


Figure 6.1: Diagnostic approach to IPF.

6.1 Gold Standards and Image Processing

The absence of gold standards in image processing makes quantification challenging for early stage images of IPF and blinded images. Gene expression profiling offers the tantalizing prospect of providing a new gold standard for the diagnosis of IPF. However, given the current absence of a definitive gold standard test, classification should incorporate clinical, radiological and pathological aspects of disease manifestation and behavior. By relying on this combined approach to classification, it is to be hoped that all points along the phenotypic continuum of disease representing IPF can be recognized, something that is crucial to the ongoing success of both clinical and basic scientific research into this devastating disease [103].

Even though there may be some misclassification or decreased discrimination between similar abnormalities for the classifiers, it is more important that the methods provide a consistent result that agrees with the experts as frequently as the experts agree among themselves, particularly in a setting where an objectively defined gold standard outside of the expert opinion does not exist [104]. With this in mind lung volumes derived from PFT results served as established clinical parameters and were used as "gold standard" [5].

Medical image processing validation often can not rely on availability of true gold standards. Gold standards based on experts interpretation or correlation with other imaging modalities and/or with pathology have been developed [105]. The comparison of a manual segmentation to a computer algorithm's classification is another issue.

6.2 Volumetrics for IPF images

Using pulse images pertinent information are extracted from defined regions with a distinctly superior separation of classes. From these regional classifications volumetric measures of honeycomb, vascular and normal regions are computed. Figure 6.2 shows classification areas for sample patients. The x -axis corresponds to the CT slide number. The range of x corresponds to scans from the top of the lung towards the bottom of the lung. There are a

few images above and below the lung in the data set. The top of the lung begins at $x = 1$.

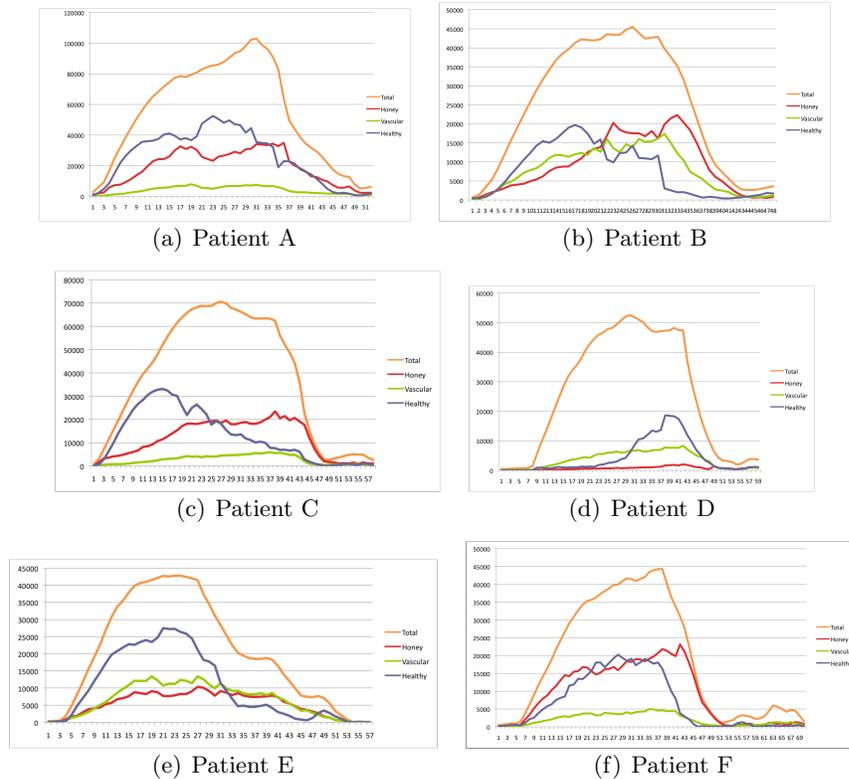


Figure 6.2: Area metrics of classification. The top most graph depicts the total area of the lung for each slice. The graph in red depicts the area classified as fibrotic, green depicts vascular and the graph in blue depicts the area classified as healthy.

The percentage of fibrotic pixels to healthy tissues increases towards the bottom of the lung in congruence with the progression of the diseased starting at the bottom of the lungs as seen in Figures 6.2(b), 6.2(c) and 6.2(f). The volumetric percentage by healthy, fibrotic and vascular regions for patients with IPF is shown in Table 6.1. The top most graph depicts the total area of the lung for each slice. The graph in red depicts the area classified as fibrotic and the graph in blue depicts the area classified as healthy.

In 6.2(c) the percentage of fibrotic pixels to healthy pixels increases towards the bottom of the lung. Of particular note is the cross over point at $x=25$ where the area of the diseased lung far exceeds the area of the healthy lung. The corresponding volumetric percentage

Table 6.1: Volumetric Percentage by Patient.

Patient ID	Percent Vacular	Percent Honey	Percent Healthy
A	7.4	34.99	48
B	32.55	37.85	30
C	7.02	29.75	33.25
D	14	2.67	17
E	31.56	25.96	50.34
F	11	49.13	36.06

across all the slices are seen in Table 6.1. In this patient the total lung volume classified as diseased is 29.75 %.

In 6.2(d) the percentage of fibrotic to healthy pixels is almost equal at the beginning of the lung. However the percentage of healthy tissue increases towards the bottom of the lung. The total lung volume classified as diseased is only 2.67 % as compared to the total lung volume classified as healthy which is 17% as seen in Table 6.1.

Finally consider Figure 6.2(b). It is seen that towards the top of the lung, the percentage of fibrotic pixels is the least compared to the healthy and vascular regions. This trend changes towards the bottom of the lung wherein the percentage fibrotic regions far exceeds that of the healthy region. The total volumes of the fibrotic, healthy and vascular regions remain on par at 38%, 30% and 33% respectively as seen in Table 6.1.

6.3 PFT Measurements

Pulmonary function testing is often used and recommended in the assessment and management of diffuse parenchymal lung diseases (DPLD). The potential clinical application of physiologic testing includes to aid in diagnosis, although its value in differential diagnosis is limited. Pulmonary function testing also aids in establishing disease severity and in defining prognosis. Potential clinical applications include: (1) aiding in diagnosis; (2) establishing disease severity; (3) defining prognosis; and (4) monitoring response to therapy and disease progression.

It is evident that the physiologic presentation of DPLDs, while typical, is not specific. As such, pulmonary function tests (PFTs) should be used in conjunction with clinical, radiographic, and histologic information in assuring an accurate diagnosis. On the other hand, in patients with appropriate symptoms, PFTs can serve as early diagnostic tools [106].

Much of the PFT data have come from an external provider/laboratory. Some are from Inova's pulmonary laboratory. There can be variances between equipment and references used for completing the PFTs within the facilities. The PFT measurements collected for analysis included but were not limited to:

- FEV1 forced expiratory volume in the first 1 second of expiration; volume of air forcibly expired from a maximum inspiratory effort in the first one second.
- FVC forced vital capacity; the total volume that can be forcefully expired from a maximum inspiratory effort.
- TLC total lung capacity; the total volume of air in the lungs at full inhalation; the sum of all volume compartments.
- DLco diffusing capacity of the lung; the capacity of the lungs to transfer carbon monoxide.

The components of the respiratory cycle are labeled as lung volumes and lung capacities (a capacity is the sum of one or more volumes) as seen in Figure 6.3 [107].

These PFT measurements were chosen for correlation with volumetric measures computed from regional classifications. A sample of the pulmonary function measurements in patients from INOVA with IPF are shown in Table 6.2.

6.4 Materials and Methods

Sixteen patients (mean age, 66.2 ± 5.28 years [\pm SD]; range, 55 to 74 years) with IPF were analyzed. PFTs were used to measure forced vital capacity (FVC), forced expiratory

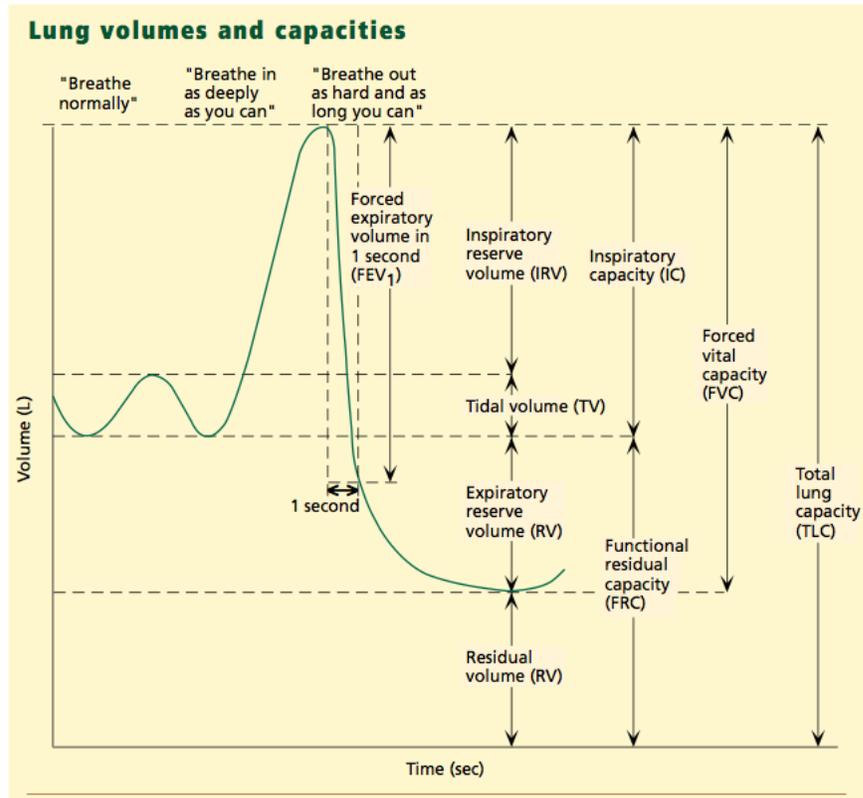


Figure 6.3: Lung volumes and capacities are depicted on a volume-time graph.

Table 6.2: Pulmonary function measurements in patients with IPF.

Patient ID	Age	CT Date	Date of PFT's	FVC	FVC %	FEV1	FEV1%	TLC	TLC %	DLco	DLco %
A	68	5/1/07	5/1/2007	2.71	68	2.3	72	4.43	65	11.6	36
C	64	6/29/09	1/8/09	3.08	77	2.77	89			6.9	29
F	55	12/28/09	9/30/09	2.48	43	2.18	49	3.81	48	9.99	26
G	58	8/27/07	12/10/07	1.93	61	1.66	66	2.92	57	8.9	36
H	67	7/3/08	7/3/08	1.5	64	1.35	72	2.56	67	8	52
I	64	10/10/07	08/21/2007	3.05	75	2.39	73	4.12	65	6.8	27

volume in 1 second (FEV1), total lung capacity (TLC), and diffusing capacity (DLco). The mean time between PFT administration and thin-section CT scanning was 64 ± 130 days.

6.4.1 Statistical Analysis

Statistical analyses were performed by using SAS 9.2 (SAS Institute, Cary, NC). Pearson correlation coefficients (*r-values*) were used to quantify the relationships between volume measurements of segmentation and each of the four PFT measurements (DLco, TLC, FEV1, and FVC). A *p-value* of less than 0.05 was considered to indicate a significant difference.

6.5 Results

The correlation coefficients (*r-values*) between volumetric measures computed from regional classification (honeycomb, vascular and normal regions) and PFTs and given in Table 6.3 and Table 6.4. Table 6.4 provides a direct comparison of segmentation volumes with PFT measures.

The mean TLC% in patients with IPF was 63.33 ± 11.48 , and mean VC% was 67.5 ± 11.22 . Mean FEV was 74.38 ± 13.96 . Mean DLco was 42.06 ± 14.71 .

There is no well-defined threshold for acceptability of strength of correlation. A high degree of correlation may indicate that volumetric measures from regional classification can provide a useful indicator of disease progression [5]. The greatest correlation was between honeycomb regions and FVC ($r=0.51$). The healthy and honeycomb regions correlated negatively with PFT measure DLco ($-0.3,-0.42$). Of all the PFTs, FEV showed the least correlation with volumetric measures, while FVC showed the greatest correlation. Of the volumetric measures, normal regions showed the least magnitude of correlation with PFTs, while honeycomb regions showed the greatest correlation.

As expected, correlation coefficients for results between PFT values were high within this group, correlation between FVC and TLC was the greatest ($r=0.72$). This is primarily because these lung function values determine how much air the lungs can hold, how quickly

Table 6.3: Correlation Coefficients (r) and p -value in patients with IPF. Note: pctvas, pcthealthy, pcthon: Volumetric measures computed from regional classification for vascular, healthy and honeycomb regions respectively, tot = vascular+healthy measure of the lung.

Pearson Correlation Coefficients								
Prob > r under H0: Rho=0								
Number of Observations								
	tot	pctvas	pcthealthy	pcthon	FVC	TLC	FEV	Dlco
tot	1	0.68149	0.77747	-0.34847	0.32818	0.23155	0.14554	-0.11653
		0.0036	0.0004	0.1859	0.2146	0.4063	0.5907	0.6674
	16	16	16	16	16	15	16	16
pctvas	0.68149	1	0.06958	-0.22939	0.13265	0.04703	0.21997	0.16652
			0.7979	0.3928	0.6243	0.8678	0.413	0.5376
	16	16	16	16	16	15	16	16
pcthealthy	0.77747	0.06958	1	-0.27788	0.33335	0.25852	0.00935	-0.30195
				0.2974	0.2071	0.3522	0.9726	0.2557
	16	16	16	16	16	15	16	16
pcthon	-0.34847	-0.22939	-0.27788	1	-0.51478	-0.13913	-0.40166	-0.42321
					0.0413	0.6209	0.123	0.1024
	16	16	16	16	16	15	16	16
FVC	0.32818	0.13265	0.33335	-0.51478	1	0.72741	0.71013	0.05392
						0.0021	0.0021	0.8428
	16	16	16	16	16	15	16	16
TLC	0.23155	0.04703	0.25852	-0.13913	0.72741	1	0.48137	-0.13357
							0.0693	0.6351
	15	15	15	15	15	15	15	15
FEV	0.14554	0.21997	0.00935	-0.40166	0.71013	0.48137	1	0.41387
								0.111
	16	16	16	16	16	15	16	16
Dlco	-0.11653	0.16652	-0.30195	-0.42321	0.05392	-0.13357	0.41387	1
	0.6674	0.5376	0.2557	0.1024	0.8428	0.6351	0.111	
	16	16	16	16	16	15	16	16

Table 6.4: Coefficients of correlation (r -values) between volumetric measures and PFT measures. Note: pctvas, pcthealthy, pcthon: Volumetric measures computed from regional classification for vascular, healthy and honeycomb regions respectively, tot = vascular+healthy measure of the lung.

	Dlco	FVC	FEV	TLC	pctvas	pcthealthy	pcthon	tot
Dlco	1							
FVC	0.05	1						
FEV	0.41	0.71	1					
TLC	-0.13	0.73	0.48	1				
pctvas	0.17	0.13	0.22	0.05	1			
pcthealthy	-0.3	0.33	0.009	0.26	0.07	1		
pcthon	-0.42	-0.51	-0.4	-0.14	-0.23	-0.28	1	
tot	-0.12	0.33	0.15	0.23	0.68	0.78	-0.35	1

one can move air in and out of the lungs, and how well the lungs put oxygen into and remove carbon dioxide from the blood and therefore correlate well amongst themselves. There was no significant correlation within the volumetric measures themselves [106].

The corresponding graphs for correlation between volumetric measures and PFT measures are provided in Figures 6.4 and 6.5. These graphs are scatterplots, which is the graphical version of the correlation. This graph shows the strength and direction of the relationship between the two variables just like the correlation coefficient. The Figure 6.4 gives the correlation between the sum of vascular and healthy against the PFT measures. The Figure 6.5 gives the correlation between the honeycomb regions against PFT measures.

In Figure 6.4(a), PFT values are plotted on the y -axis and correlated against volumetric measures comprising of the sum of vascular and healthy regions (tot). There is a negative relationship (correlation) seen between DLco and tot. As DLco increases, the sum of vascular and healthy regions decreases. Unlike Figure 6.2(a), Figures 6.2(b), 6.2(c) and 6.2(d) show a positive relationship between the two parameters. The strongest correlation ($r = 0.33$) is seen between tot and FVC values and the least is seen between tot and DLco ($r = 0.11$).

In Figure 6.5(a), PFT values are plotted on the x -axis and correlated against volumetric measure from the honeycomb regions (pcthon) on the y -axis. Overall a negative relationship (correlation) is seen between pcthon and the PFT values. As the percentage volume of the honeycomb region decreases, the PFT value increases. The strongest relationship is seen between pcthon and FVC ($r = 0.51$) closely followed by DLco ($r = 0.42$) and FEV ($r = 0.4$). The weakest correlation is seen between pcthon and TLC values ($r = 0.14$). It is however worth re-iterating that there is no well-defined threshold for acceptability of strength of correlation [5].

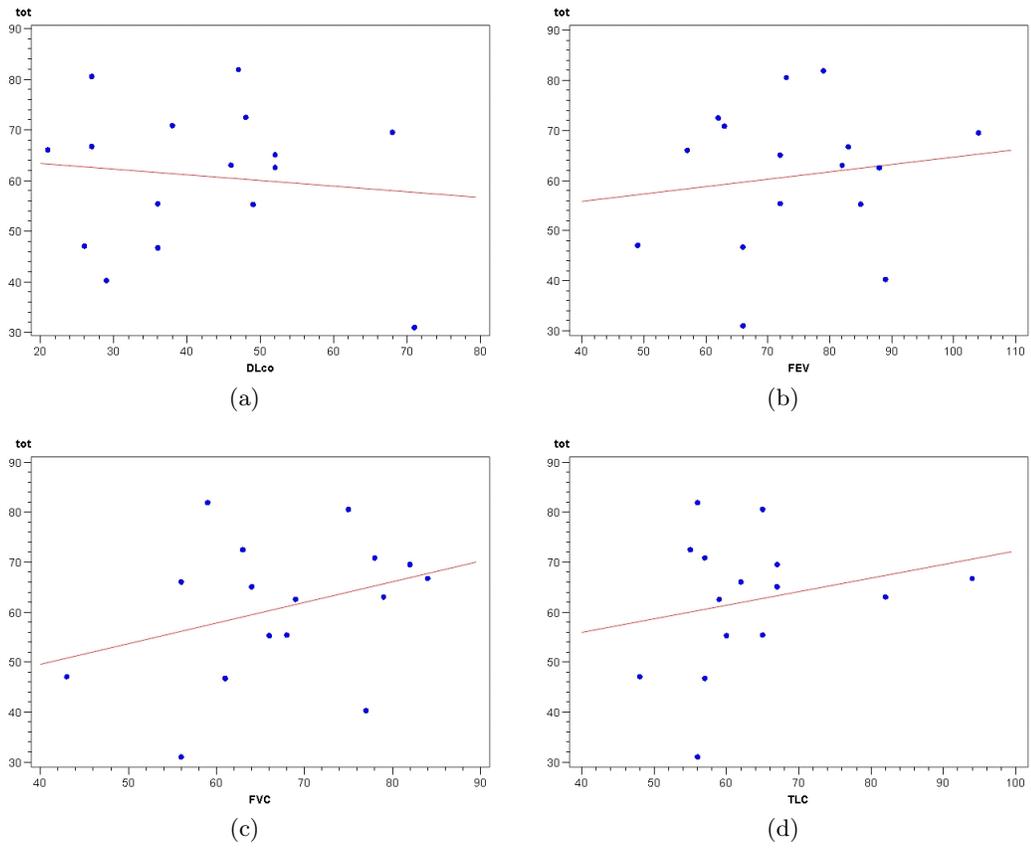


Figure 6.4: Correlation graphs between sum of vascular and healthy (tot) against PFT measures.

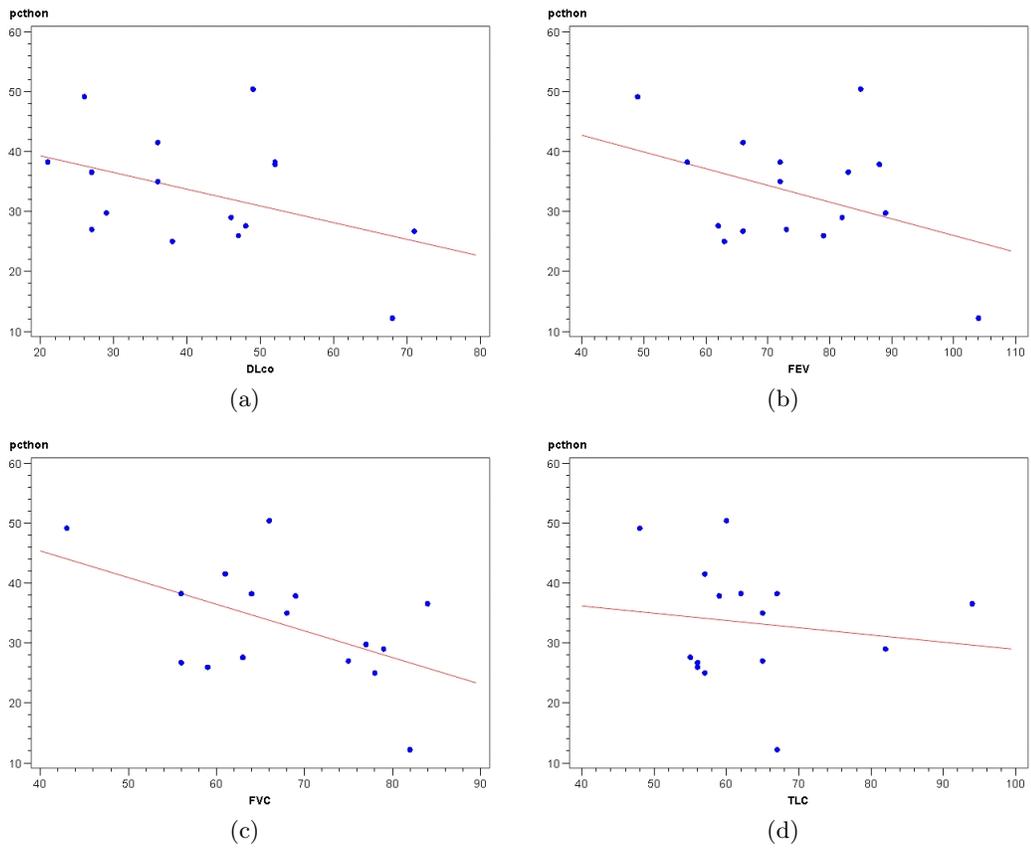


Figure 6.5: Correlation graphs between honeycomb regions (pcthon) and PFT measures.

6.6 Discussion

There is currently no widely accepted standardization system for scoring the extent of ILD. It is especially important to define disease extent in patients with IPF for the purposes of selection of treatment, determination of disease progression and evaluation of effectiveness of investigational new treatments. The absence of gold standards in image processing makes quantification challenging for early stage images of IPF and blinded images. Thus medical image processing validation often cannot rely on availability of true gold standards.

Hence lung volumes derived from Pulmonary Function Tests (PFT) results served as established clinical parameters and were used as gold standards. Volumetric measurements of honeycomb, vascular and normal regions are found to correlate with results of PFTs in patients with IPF. The greatest correlation was between honeycomb regions and forced vital capacity (FVC). The healthy and honeycomb regions correlated negatively with PFT measure diffusing capacity (DLco). Thus, in spite of the obvious difference in contrast, volume and texture, the healthy and diseased regions are distinguished and classified using pulse images.

Chapter 7: Summary

High-Resolution Computed Tomography (HRCT) has changed the diagnostic evaluation of patients with ILD and is particularly useful in the diagnosis of idiopathic pulmonary fibrosis (IPF) [3]. The quantification of disease by CT is important to indicate prognosis and to evaluate progression of the disease or response to treatment. In this thesis automated tools are presented which extract information from the CT images and isolate visual evidence of the disease from healthy lung tissue.

7.1 Previous Approaches

All existing methods of CT quantification of ILD are less than ideal and are dependent on the depth of the patient's inspiration. Visual estimation techniques are subjective and show only moderate reproducibility. Quantitative methods based on CT density histogram are particularly susceptible to variation with differing depths of inspiration. A significant problem with quantitative CT technology is the lack of user-friendly automated method for segmenting the lung and deriving quantitative indices. It seems particularly surprising that none of the major CT manufacturers provide a user-friendly package for performing quantitative analysis of lung images [4].

Texture based methods examines the relationship between the attenuation of pixels within a certain distance of each other. However the methods failed to compare the results of the segmentation with measurements such as pulmonary function tests, which is important for the treatment option and the patient prognosis [46].

7.2 Proposed Algorithms

Three different approaches were undertaken which resulted in some drawbacks and lessons learned. It effectively led to an important conclusion when working with IPF images: Quantization of IPF images is challenging because it is difficult and potentially misleading to use a single metric to convey the complex textural information offered by the CT scan.

7.2.1 Adaptive multiple Feature Method

An adaptive thresholding algorithm has been presented which utilized quantitative CT indexes to correlate IPF with pulmonary abnormality. Simulation results demonstrated that the algorithm performed well in identified IPF images. However the absence of gold standards made quantification challenging for early stage images of IPF and blinded images.

7.2.2 Neural Networks

Image processing requires complex computations which explains the need for artificial neural networks [108,109]. However there were several disadvantages which came to the forefront during the task of training the network. The lack of any improvement in segmentation quality was accompanied by a considerable increment in human intervention: manual processing time for interactive labeling of the training data was approximately 20 to 30 minutes per image. The second problem related to the amount of input data. An unblinded number of 16 IPF patients was insufficient to train and evaluate the quality of the segmentation results.

7.2.3 Gabor Filter

Gabor filter is a very useful tool for texture analysis in both spatial and frequency domains and hence combines the advantages of both filters. However the disadvantage of image dependence unexpectedly resulted in image enhancements. Certain configuration settings for wavelength, orientation, phase offset or aspect ratio worked well with one set of images

but not with another. It was challenging to identify identical texture with less or more resolution than the original image presented to the Gabor filter as the same texture.

7.3 PCNN and IPF

The ICM has been used successfully in texture analysis of medical images [97,98] is a modification of the PCNN which reduces the computational complexity. The technique used is successful in classifying the healthy and diseased portions of the lung. The goal was to train adequate and varied stages of IPF images and to be able to extract sufficiently enough information from test images. The algorithm was tested on HRCT scans procured through INOVA Fairfax Hospital, Department of Radiology. Two expert radiological reviewers compared the initial results of the segmentation algorithm with the manual segmentation of the original scans. Comparison revealed agreement regarding the presence or absence of honeycombing.

7.4 Volumetric Analysis and Gold standards

It is especially important to define disease extent in patients with IPF for the purposes of selection of treatment, determination of disease progression and evaluation of effectiveness of investigational new treatments. The absence of gold standards in image processing makes quantification challenging for early stage images of IPF and blinded images. Hence lung volumes derived from Pulmonary Function Tests (PFT) results served as established clinical parameters and were used as gold standards. Volumetric measurements of honeycomb, vascular and normal regions are found to correlate with results of PFTs in patients with IPF. Thus, in spite of the obvious difference in contrast, volume and texture, the healthy and diseased regions are distinguished and classified using pulse images.

7.5 Conclusion

In conclusion, the diagnostic approach to IPF remains multidisciplinary, involving primary care physicians, pulmonologists, radiologists, pathologists and image analysts. Quantitation of radiological images of IPF remains challenging. But despite its limitation, visual estimation remains the most widely used technique for the estimation of the extent of the diffused lung disease. Quantitative evaluation is a critical component of lung disease, and when validated against an appropriate gold standard, it may one day replace routine physiologic assessment. Until then, the need for the ideal quantitative technique that would be simple to perform, reproducible, observer independent and offer a valid measurement of disease extent would remain. In that perspective, clinicians and image analysts will continue to team up to build the best information retrieval methodologies.

Bibliography

Bibliography

- [1] “Idiopathic Pulmonary Fibrosis: Diagnosis and Treatment . International Consensus Statement,” *Am. J. Respir. Crit. Care Med.*, vol. 161, pp. 646–664, 2000.
- [2] T. Lindblad and J. M. Kinser, *Image Processing Using Pulse-Coupled Neural Networks*, 2nd ed. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2005.
- [3] M. B. Gotway, M. M. Freemer, and T. E. King, “Challenges in pulmonary fibrosis 1: Use of high resolution CT scanning of the lung for the evaluation of patients with idiopathic interstitial pneumonias,” *Thorax*, vol. 62, pp. 546–553, 2007.
- [4] D. A. L. MB and J. D. N. MD, *Imaging of Diffuse Lung Disease*. pmph USA: pmph USA, 2000.
- [5] J. Zaporozhan, S. Ley, K. K. Gast, J. Schmiedeskamp, A. Biedermann, B. Eberle, and H.-U. Kauczor, “Functional analysis in single-lung transplant recipients*a comparative study of high-resolution ct, 3he-mri, and pulmonary function tests,” *CHEST Journal*, vol. 125, no. 1, pp. 173–181, 2004. [Online]. Available: + <http://dx.doi.org/10.1378/chest.125.1.173>
- [6] I. Johnston, J. Britton, W. Kinnear, and R. Logan, “Rising mortality from cryptogenic fibrosing alveolitis.” [Online]. Available: <http://www.pubmedcentral.gov/articlerender.fcgi?artid=1664054>
- [7] J. A. BJORAKER, J. H. RYU, M. K. EDWIN, J. L. MYERS, H. D. TAZELAAR, D. R. SCHROEDER, and K. P. OFFORD, “Prognostic Significance of Histopathologic Subsets in Idiopathic Pulmonary Fibrosis,” *Am. J. Respir. Crit. Care Med.*, vol. 157, pp. 199–203, 1998.
- [8] S. Harar and A. Caminat, *IPF: new insight on pathogenesis and treatment*, 201, vol. 6. [Online]. Available: <http://dx.doi.org/10.1111/j.1398-9995.2009.02305>.
- [9] J. Chapman and C. Farver, “Idiopathic interstitial lung disease: a review of recent classifications.” *Clin Pulm Med.*, vol. 11, pp. 17–24, 2004.
- [10] A. C. Society, “Cancer facts and figures 2006.” 2006. [Online]. Available: www.cancer.org
- [11] G. Raghu, D. Weycker, J. Edelsberg, W. Z. Bradford, and G. Oster, “Incidence and prevalence of idiopathic pulmonary fibrosis.” *Am. J. Respir. Crit. Care Med.*, vol. 174, pp. 810–816, 2006.

- [12] C. F. Foundation, "What is cf?" 2006. [Online]. Available: www.cff.org
- [13] T. K. Jr, M. Schwarz, K. Brown, J. Tooze, T. Colby, J. W. Jr, A. Flint, W. Thurlbeck, and R. Cherniack, "Idiopathic pulmonary fibrosis: relationship between histopathologic features and mortality." *Am. J. Respir. Crit. Care Med.*, vol. 164, pp. 1025–1032, 2001.
- [14] F. Martinez, S. Safrin, D. Weycker, K. Starko, W. Bradford, T. K. Jr, K. Flaherty, D. Schwartz, P. Noble, G. Raghu, and K. Brown, "The clinical course of patients with idiopathic pulmonary fibrosis." *Ann Intern Med.*, vol. 142, pp. 963–967, 2005.
- [15] J. R. Mathieson, J. R. Mayo, C. A. Staples, and N. L. Mller, "Chronic diffuse infiltrative lung disease: comparison of diagnostic accuracy of CT and chest radiography." *Radiology*, vol. 171, pp. 111–116, 1989.
- [16] K. T. Tung, A. U. Wells, M. B. Rubens, J. M. Kirk, R. M. du Bois, and D. M. Hansell, "Accuracy of the typical computed tomographic appearances of fibrosing alveolitis." *Thorax*, vol. 48, pp. 334–338, 1993. [Online]. Available: <http://thorax.bmj.com/content/48/4/334.abstract>
- [17] P. Grenier, D. Valeyre, P. Cluzel, M. W. Brauner, S. Lenoir, and C. Chastang, "Chronic diffuse interstitial lung disease: diagnostic value of chest radiography and high-resolution CT." *Radiology*, vol. 179, pp. 123–132, 1991.
- [18] D. Lynch, J. Newell, P. Logan, T. King, and N. Muller, "Can CT distinguish hypersensitivity pneumonitis from idiopathic pulmonary fibrosis?" *American Journal of Roentgenology*, vol. 165, no. 4, pp. 807–811, Oct. 1995. [Online]. Available: <http://www.ajronline.org/cgi/content/abstract/165/4/807>
- [19] S. J. Swensen, G. L. Aughenbaugh, and J. L. Myers, "Diffuse lung disease: diagnostic accuracy of CT in patients undergoing surgical biopsy of the lung." *Radiology*, vol. 205, no. 1, pp. 229–234, Oct. 1997. [Online]. Available: <http://radiology.rsna.org/content/205/1/229.abstract>
- [20] C. Collins, A. Wells, D. Hansell, R. Morgan, J. Macsweeney, R. Dubois, and M. Rubens, "Observer variation in pattern type and extent of disease in fibrosing alveolitis on thin section computed tomography and chest radiography," *Clinical Radiology*, vol. 49, no. 4, pp. 236–240, 1994. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0009926005818471>
- [21] T. Beinert, J. Behr, F. Mehnert, P. Kohz, M. Seemann, and M. Reiser, "Spirometrically controlled quantitative CT for assessing diffuse parenchymal lung disease," *Journal of Computer Assisted Tomography*, vol. 19, no. 6, 1995.
- [22] J. D. Pham, Chenyang Xu, "Current methods in medical image segmentation," *Annual Review of Biomedical Engineering*, vol. 2, pp. 315–337, 2000.
- [23] K. A.L. Baert, M. Knauth, *MRI of the Lung*. Springer, 2009.
- [24] A. Reijg and C. Kroegel, "Transthoracic sonography of diffuse parenchymal lung disease," *Journal of Ultrasound in Medicine*, vol. 22, no. 2, pp. 173–180, Feb. 2003. [Online]. Available: <http://radiology.rsna.org/content/175/1/265.abstract>

- [25] C. J. Bergin, D. Y. Bell, C. L. Coblenz, C. Chiles, G. Gamsu, N. R. MacIntyre, R. E. Coleman, and C. E. Putman, "Sarcoidosis: correlation of pulmonary parenchymal pattern at CT with results of pulmonary function tests." *Radiology*, vol. 171, no. 3, pp. 619–624, Jun. 1989. [Online]. Available: <http://radiology.rsna.org/content/171/3/619.abstract>
- [26] N. L. Mller, J. B. Mawson, J. R. Mathieson, R. Abboud, D. N. Ostrow, and P. Champion, "Sarcoidosis: correlation of extent of disease at ct with clinical, functional, and radiographic findings." *Radiology*, vol. 171, no. 3, pp. 613–618, 1989. [Online]. Available: <http://radiology.rsna.org/content/171/3/613.abstract>
- [27] W. A. Kalender, R. Rienmller, W. Seissler, J. Behr, M. Welke, and H. Fichte, "Measurement of pulmonary parenchymal attenuation: use of spirometric gating with quantitative CT." *Radiology*, vol. 175, no. 1, pp. 265–268, Apr. 1990. [Online]. Available: <http://radiology.rsna.org/content/175/1/265.abstract>
- [28] H. Coxson, J. Hogg, J. Mayo, H. Behzad, K. Whittall, D. Schwartz, P. Hartley, J. Galvin, J. Wilson, and G. Hunninghake, "Quantification of idiopathic pulmonary fibrosis using computed tomography and histology," *American Journal of Respiratory and Critical Care Medicine*, vol. 155, no. 5, pp. 1649–1656, May 1997. [Online]. Available: <http://ajrcm.atsjournals.org/cgi/content/abstract/155/5/1649>
- [29] P. Wollmer, K. Jakobsson, M. Albin, U. Albrechtsson, K. Brauer, L. Eriksson, B. Jonson, S. Skerfving, and U. Tylene, "Measurement of lung density by x-ray computed tomography. relation to lung mechanics in workers exposed to asbestos cement," *Chest*, vol. 91, no. 6, pp. 865–869, 1987.
- [30] L. Hedlund, P. Vock, and E. Effmann, "Evaluating lung biopsy computed tomography," *Semin Respir Med*, vol. 5, pp. 76–87, 1983.
- [31] R. Mull, "Mass estimates by computed tomography: physical density from ct numbers." *Am J Radiol*, vol. 143, pp. 1101–1104, 1984.
- [32] R. M, H. I, E. L, B. G, B. E, S. D, and H. M., "Quantitative computed tomography of the lung in the early diagnosis of asbestosis: initial results in asbestos-exposed navy personnel," *Rofo*, pp. 288–295, 1995.
- [33] A. C. Best, A. M. Lynch, C. M. Bozic, D. Miller, G. K. Grunwald, and D. A. Lynch, "Quantitative CT indexes in idiopathic pulmonary fibrosis: Relationship with physiologic impairment1," *Radiology*, vol. 228, no. 2, pp. 407–414, 2003. [Online]. Available: <http://radiology.rsna.org/content/228/2/407.abstract>
- [34] T. Beinert, J. Behr, P. Kohz, M. Seemann, T. Egge, and M. Reiser, "Spirometrically controlled high resolution computed tomography - quantitative assessment of density distribution in patients with diffuse fibrosing alveolitis." *Eur J Med Res.*, vol. 1, pp. 269–272, 1996.
- [35] J. A. Verschakelen, L. V. fraeyenhoven, G. Laureys, M. Demedts, and A. L. Baert, "Differences in ct density between dependent and nondependent portions of the lung: influence of lung volume." *Am J Roentgenology*, vol. 161, pp. 713–717, 1993.

- [36] T. Beinert, J. Behr, F. Mehnert, P. Kohz, M. Seemann, and M. Reiser, "Quantitative computerized tomography of the lung–respiration controlled diagnosis of diffuse lung diseases." *Pneumologie.*, vol. 49, pp. 678–683, 1995.
- [37] N. Muller, C. Staples, R. Miller, and R. Abboud, "Density mask: an objective method to quantitate emphysema using computed tomography," *Chest*, vol. 95, pp. 782–787, 1988.
- [38] G. Kemerink, H. Kruize, R. Lamers, and J. van Engelshoven, "Ct lung densitometry: dependence of ct number histograms on sample volume and consequences for scan protocol comparability," *J. Comput. Assist. Tomogr*, vol. 21, pp. 948–954, 1997.
- [39] H. Zagers, H. Vrooman, N. Aarts, J. Stolk, L. Schultze, Kool, J. Dijkman, A. V. Voorthuisen, and J. Reiber, "Assessment of the progression of emphysema by quantitative analysis of spirometrically gated computed tomography images," *Invest. Radiol.*, vol. 31, pp. 761–767, 1996.
- [40] M. Gilman, R. Laurens, J. Somogyi, and E. Honig, "Ct attenuation values of lung density in sarcoidosis," *J. Comput. Assist. Tomogr*, vol. 7, pp. 407–410, 1983.
- [41] R. Rienmuller, J. Behr, W. Kalender, M. Schatzl, I. Altmann, M. Merin, and T. Beinert, "Standardized quantitative high resolution ct in lung diseases," *J. Comput. Assist. Tomogr*, vol. 17, pp. 742–749, 1991.
- [42] T. Beinert, J. Behr, F. Mehnert, P. Kohz, M. Seemann, R. Rienmuller, and M. Reiser, "Spirometrically controlled quantitative ct for assessing diffuse parenchymal lung disease," *J. Comput. Assist. Tomogr*, vol. 19, pp. 924–931, 1995.
- [43] R. UPPALURI, E. HOFFMAN, M. SONKA, G. HUNNINGHAKE, and G. MCLENNAN, "Interstitial lung disease . a quantitative study using the adaptive multiple feature method," *American Journal of Respiratory and Critical Care Medicine*, vol. 159, no. 2, pp. 519–525, Feb. 1999. [Online]. Available: <http://ajrcm.atsjournals.org/cgi/content/abstract/159/2/519>
- [44] Y. Uchiyama, S. Katsuragawa, H. Abe, J. Shiraishi, F. Li, Q. Li, C. Zhang, K. Suzuki, and K. Doi, "Quantitative computerized analysis of diffuse lung disease in high-resolution computed tomography," *Medical Physics*, vol. 30, no. 9, p. 2440, 2003.
- [45] S. DELORME, M. KELLER-REICHENBECHER, I. ZUNA, W. SCHLEGEL, and G. V. KAICK, "Usual interstitial pneumonia: Quantitative assessment of High-Resolution computed tomography findings by Computer-Assisted Texture-Based image analysis," *Investigative Radiology*, vol. 32, no. 9, 1997.
- [46] L. H. Rodriguez, P. F. Vargas, U. Raff, D. A. Lynch, G. M. Rojas, D. M. Moxley, and J. D. Newell, "Automated discrimination and quantification of idiopathic pulmonary fibrosis from normal lung parenchyma using generalized fractal dimensions in high-resolution computed tomography images," *Academic Radiology*, vol. 2, no. 1, pp. 10–18, Jan. 1995. [Online]. Available: <http://www.sciencedirect.com/science/article/B75BK-4HNF68G-3/2/dbfa72f6c2cd299f246944af151b1290>

- [47] S. Katsuragawa, K. Doi, H. MacMahon, N. Nakamori, Y. Sasaki, and J. J. Fennessy, "Quantitative computer-aided analysis of lung texture in chest radiographs." *Radiographics*, vol. 10, no. 2, pp. 257–269, Mar. 1990. [Online]. Available: <http://radiographics.rsna.org/content/10/2/257.abstract>
- [48] S. Katsuragawa, K. Doi, and H. MacMahon, "Image feature analysis and computer-aided diagnosis in digital radiography: classification of normal and abnormal lungs with interstitial diseases in chest images." *Medical Physics*, vol. 16, pp. 38–44, 1989.
- [49] T. Zrimec and J. Wong, "Methods for automatic honeycombing detection in hrct images of the lung," in *11th Mediterranean Conference on Medical and Biomedical Engineering and Computing 2007*, ser. IFMBE Proceedings, R. Magjarevic, J. H. Nagel, T. Jarm, P. Kramar, and A. Zupanic, Eds. Springer Berlin Heidelberg, 2007, vol. 16, pp. 830–833, hC.
- [50] R. Uppaluri, E. Hoffman, M. Sonka, G. Hunninghake, and G. Mclellan, "Interstitial lung disease . a quantitative study using the adaptive multiple feature method," *American Journal of Respiratory and Critical Care Medicine*, vol. 159, pp. 519–525, 1999.
- [51] G. W. Davis and S. T. Wallenslager, "Improvement of chest region CT images through automated Gray-Level remapping," *IEEE Transactions on Medical Imaging*, vol. 5, no. 1, pp. 30–34, 1986.
- [52] R. protection of patients. (2012, Jun.) Image quality and quality control. [Online]. Available: <https://rpop.iaea.org/RPOP/RPoP/Content/index.htm>
- [53] A.P.Dhawan, *Medical Image Analysis*. John Wiley and Sons Ltd, 2003.
- [54] K. Zuiderveld, *Contrast limited adaptive histogram equalization*. San Diego, CA, USA: Academic Press Professional, Inc., 1994, pp. 474–485. [Online]. Available: <http://portal.acm.org/citation.cfm?id=180895.180940>
- [55] W. Qian, L. Clarke, B. Zheng, M. Kallergi, and R. Clark, "Computer assisted diagnosis for digital mammography," *IEEE Engineering in Medicine and Biology Magazine*, vol. 14, no. 5, pp. 561–569, Oct. 1995.
- [56] I.N.Bankman, *Handbook of medical imaging: processing and analysis*, ser. Academic Press series in biomedical engineering. Academic Press, 2000.
- [57] C. H. Chen, L. F. Pau, and P. S. P. Wang, *The Handbook of Pattern Recognition and Computer Vision*. World Scientific Publishing Co., 1998.
- [58] R. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Trans. on Systems, Man and Cybernetics*, vol. 3, pp. 610–621, 1973.
- [59] F. Albreghsen, *Statistical Texture Measures Computed from Gray Level Run Length Matrices*, Image Processing Laboratory Department of Informatics University of Oslo, 1995.

- [60] E. M. Ferdeghini, B. Pinamonti, E. Picano, F. Lattanzi, R. Bussani, G. Slavich, A. Benassi, F. Camerini, L. Landini, and A. L'Abbate, "Quantitative texture analysis in echocardiography: Application to the diagnosis of myocarditis," *Journal of Clinical Ultrasound*, vol. 19, no. 5, pp. 263–270, 1991. [Online]. Available: <http://dx.doi.org/10.1002/jcu.1870190503>
- [61] D.Gabor, "Theory of communication." *J. Inst. Elect. Eng London.*, vol. 93, pp. 429–457, 1946.
- [62] J. P. Havlicek, A. Bovik, and D. Chen, "Am-pm image modeling and gabor analysis." *Visual Information Representation, Communication, and Image Processing.*, vol. 99, pp. 343–386, 1999.
- [63] D. J. Field, "Relation between the statistics of natural images and the response properties of cortical cell." *J. Opt. Soc. Am. A.*, vol. 4, pp. 2379–2394, 1987.
- [64] A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequency by gabor wavelets." *IEEE Trans. on Inf. Theory.*, vol. 38, pp. 691–712, 1992.
- [65] O. Nestares, R. Navarro, T. Portilla, and A. Taberero, "Efficient spatial-domain implementation of multiscale image representation based on gabor functions." *J. Electronic Imaging.*, vol. 7, pp. 166–173, 1998.
- [66] J.Daugman, "Uncertainty relation for resolution in space spatial frequency and orientation optimised by two dimensional visual cortical filter." *J. Opt. Soc. Am.*, vol. 2, pp. 1160–1169, 1985.
- [67] J. G. 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, pp. 1160–1169, Jul 1985. [Online]. Available: <http://josaa.osa.org/abstract.cfm?URI=josaa-2-7-1160>
- [68] N. Petkov, "Biologically motivated computationally intensive approaches to image pattern recognition," *Future Generation Computer Systems*, vol. 11, no. 4-5, pp. 451–465, 1995.
- [69] N. Petkov and P. Kruizinga, "Computational models of visual neurons specialised in the detection of periodic and aperiodic oriented visual stimuli: bar and grating cells," *Biological Cybernetics*, vol. 76, no. 2, pp. 83–96, 1997.
- [70] P. Kruizinga and N. Petkov, "Non-linear operator for oriented texture," *IEEE Trans. on Image Processing*, vol. 8, no. 10, pp. 1395–1407, 1999.
- [71] N. P. S.E. Grigorescu and P. Kruizinga, "Comparison of texture features based on gabor filters," *IEEE Trans. on Image Processing*, vol. 11, no. 10, pp. 1160–1167, 2002.
- [72] N. Petkov and M. A. Westenberg, "Suppression of contour perception by band-limited noise and its relation to non-classical receptive field inhibition," *Biological Cybernetics*, vol. 88, no. 10, pp. 236–246, 2003.

- [73] C. Grigorescu, N. Petkov, and M. A. Westenberg, “Contour detection based on non-classical receptive field inhibition,” *IEEE Transactions on Image Processing*, vol. 12, no. 7, pp. 729–739, 2003.
- [74] A. C. Bovik, M. Clark, and W. S. Geisler, “Multichannel texture analysis using localized spatial filters,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 1, pp. 55–73, Jan. 1990. [Online]. Available: <http://dx.doi.org/10.1109/34.41384>
- [75] T. S. Lee, “Image representation using 2d gabor wavelets,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, pp. 959–971, 1996.
- [76] B. S. Manjunath and W. Y. Ma, “Texture features for browsing and retrieval of image data,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 8, pp. 837–842, Aug. 1996. [Online]. Available: <http://dx.doi.org/10.1109/34.531803>
- [77] *GRANITE TEXTURE CLASSIFICATION WITH GABOR FILTERS*, 2006.
- [78] A. K. Jain and F. Farrokhnia, “Unsupervised texture segmentation using gabor filters,” *Pattern Recogn.*, vol. 24, no. 12, pp. 1167–1186, 1991.
- [79] *Design of Gabor wavelets for analysis of texture features in cervical images*, 2003.
- [80] R. Manthalkar, P. K. Biswas, and B. N. Chatterji, “Rotation invariant texture classification using even symmetric gabor filters,” *Pattern Recognition Letters*, vol. 24, no. 12, pp. 2061–2068, 2003.
- [81] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, “Shiftable multi-scale transforms,” *IEEE Trans Information Theory*, vol. 38, no. 2, pp. 587–607, March 1992, special Issue on Wavelets.
- [82] S. K. Rogers and D. W. Ruck, Eds., *Simplified pulse-coupled neural network*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 2760, Mar. 1996.
- [83] R. C. Muresan, “Pattern recognition using pulse-coupled neural networks and discrete fourier transforms,” *Neurocomputing*, vol. 51, no. 0, pp. 487 – 493, 2003. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0925231202007270>
- [84] R. Eckhorn, H. J. Reitboeck, M. Arndt, and P. Dicke, “Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex,” *Neural Comp.*, vol. 2, pp. 293–307, 1990.
- [85] E. Ulf, J. M. Kinser, J. Atmer, and N. Zetterlund, “The intersecting cortical model in image processing,” *Nuclear Instruments and Methods in Physics Research*, vol. 525, pp. 392–396, 2004.
- [86] J. L. Johnson, “Pulse-coupled neural nets: Translation, scale, distortion, and intensity signal invariance for images,” *Applied Optics*, vol. 33, pp. 6239–6253, 1994.
- [87] Y. I. Balkarey, M. G. Evtikhov, and M. I. Elinson, “Autowave media and neural networks,” in *Proc. of SPIE*, vol. 1621, 1991, pp. 238–249.

- [88] O. A. Mornev, *Elements of the Optics of Autowaves in Self-Organization Autowaves and Structures far from Equilibrium*. Springer-Verlag, 1984, pp. 111–118.
- [89] W. Gernster, “Time structure of the activity in neural network models,” *Phys. Rev. E*, vol. 51, no. 1, pp. 738–758, 1995.
- [90] R. E. Mirollo and S. H. Strogatz, “Synchronization of pulse-coupled biological oscillators,” *SIAM J. of Appl. Math.*, vol. 50, no. 6, pp. 1645–1662, 1990.
- [91] O. Parodi, P. Combe, and J.-C. Ducom, “Temporal encoding in vision: Coding by spiking arrival times leads to oscillations in the case of moving targets,” *Neurocomputing*, vol. 4, pp. 93–102, 1992.
- [92] Y. Yamaguchi, K. Ishimura, and M. Wada, “Synchronized oscillation and dynamical clustering in chaotic pcnn,” in *Proc. of SICE*, August 2002, pp. 730–735.
- [93] E. Neibur and F. Wörgötter, “Circular inhibition: A new concept in long-range interaction in the mammalian visual cortex,” in *Proc. IJCNN*, vol. II, 1990, pp. 367–372.
- [94] J. M. Kinser, “Image signatures: Classification and ontology,” in *Proc. of the 4th IASTED Intl. Conf. on Computer Graphics and Imaging*, 2001.
- [95] R. Malladi and J. A. Sethian, “Level set methods for curvature flow, image enhancement, and shape recovery in medical images,” in *Proc. of Conf. on Visualization and Mathematics*. Springer-Verlag, June 1995, pp. 329–345.
- [96] M. A. Grayson, “The heat equation shrinks embedded plane curves to round points,” *J. Differential Geometry*, vol. 26, pp. 285–314, 1987.
- [97] J. M. Kinser and G. Wang, “Texture recognition of medical images with the icm method,” *Nuclear Inst. and Methods in Physics Research, A*, vol. 525, no. 1-2, pp. 387–391, 2004.
- [98] J. M. Kinser, “Biological models for image processing,” pp. 151–165, 2003.
- [99] J. M. Kinser and C. Nguyen, “Image object signatures from centripetal autowaves,” vol. 21, no. 3, pp. 221–225, 2000.
- [100] J. M. Kinser and F. Selzer, “Multiordered mapping technique for target prioritization,” in *SPIE Proc.*, vol. 1699, 1992, pp. 2–7.
- [101] C. L. Wyman, M. Schreeder, W. Grundy, and J. M. Kinser, “Cancer diagnostics using neural network sorting of images,” in *Proc. of the SPIE*, vol. 2760, 1996.
- [102] “American thoracic society/european respiratory society international multidisciplinary consensus classification of the idiopathic interstitial pneumonias,” *American Journal of Respiratory and Critical Care Medicine*, vol. 165, no. 2, pp. 277–304, 2002. [Online]. Available: <http://ajrccm.atsjournals.org/content/165/2/277.short>
- [103] T. M. Maher, A. U. Wells, and G. J. Laurent, “Idiopathic pulmonary fibrosis: multiple causes and multiple mechanisms?” *European Respiratory Journal*, vol. 30, no. 5, pp. 835–839, 2007. [Online]. Available: <http://erj.ersjournals.com/content/30/5/835.abstract>

- [104] V. Zavaletta, B. Bartholmai, and R. Robb, "High resolution multidetector CT-aided tissue analysis and quantification of lung fibrosis." *Acad Radiol*, vol. 14, no. 7, pp. 772–787, Jul. 2007. [Online]. Available: <http://ukpmc.ac.uk/abstract/MED/17574128>
- [105] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (staple): An algorithm for the validation of image segmentation," *IEEE TRANS. MED. IMAG*, vol. 23, pp. 903–921, 2004.
- [106] F. J. Martinez and K. Flaherty, "Pulmonary function testing in idiopathic interstitial pneumonias," *Proceedings of the American Thoracic Society*, vol. 3, no. 4, pp. 315–321, 2006. [Online]. Available: <http://pats.atsjournals.org/content/3/4/315.abstract>
- [107] F. Al-Ashkar, R. Mehra, and P. J. Mazzone, "Interpreting pulmonary function tests: recognize the pattern, and the diagnosis will follow." *Cleveland Clinic journal of medicine*, vol. 70, no. 10, p. 866866, 2003. [Online]. Available: <http://www.ccjm.org/content/70/10/866.short>
- [108] K. Doi, "Computer-aided diagnosis in medical imaging: Historical review, current status and future potential," *Comp. Med. Imag. and Graph.*, vol. 31, no. 4-5, pp. 198–211, 2007.
- [109] A. Miller, B. Blott, and T. James, "Review of neural network applications in medical imaging and signal processing," *Medical and Biological Engineering and Computing*, vol. 30, pp. 449–464, 1992, 10.1007/BF02457822. [Online]. Available: <http://dx.doi.org/10.1007/BF02457822>

Curriculum Vitae

Mekhala Acharya grew up in Mumbai, India. She attended University of Mangalore, where she received her Bachelor of Engineering in Computer Science in 2000. She went on to receive her Master of Science in Bioinformatics and Computational Biology from George Mason University in 2006. She is working in the area of Statistics in Clinical Trials and looks forward to a teaching career in a few years.