Multi-scale Representations for Classification of Protein Crystal Images and Multi-Modal Registration of the Lung

Ming Jack Po

Submitted in partial fulfillment of the requirements for the degree
of Doctor of Philosophy in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY
2015

ABSTRACT

Multi-scale Representations for Classification of Protein Crystal Images and Multi-Modal Registration of the Lung

Ming Jack Po

In recent years, multi-resolution techniques have become increasingly popular in the image processing community. New techniques have been developed with applications ranging from edge detection, texture recognition, image registration, multi-resolution features for image classification, and more. The central focus of this two-part thesis is the multi-resolution analysis of images. In the first part, we describe multi-resolution approaches used to help with the classification of a set of protein crystal images. In the second, we focus on similar approaches used to help register a set of 3D image volumes that would otherwise be computationally prohibitive without leveraging multi-resolution techniques.

Specifically, the first part of this work proposes a classification framework that is being developed in collaboration with NorthEast Structural Genomics Consortium (NESG) to assist in the automated screening of protein crystal images. Several groups have previously proposed automated algorithms to expedite such analysis. However, none of the classifiers described in the literature are sufficiently accurate or fast enough to be practical in a structural genomics

production pipeline. The proposed classification algorithm uses random window sampling of the regions of interest to then compute several texture and multi-resolution image descriptor features that are subsequently processed through a random forest classifier. The resulting binary classifier exceeds 90% in sensitivity and 94% in specificity. Furthermore, the classifier is able to process each image with off-the-shelf computer components at approximately 7 seconds for each image, a speed that makes this algorithm usable in high throughput settings.

The second part of this work proposes a 3D image registration algorithm to register regions of emphysema as quantified by densitometry on lung CT with MR lung volumes. The ability to register quantitatively-determined regions of emphysema with perfusion MRI will allow for further exploration of the pathophysiology of Chronic Obstructive Pulmonary Disorder (COPD). The registration method involves the registration of CT volumes at different levels of inspiration (total lung capacity to functional residual capacity [FRC]) followed by another registration between FRC-CT and FRC-MR volume pairs. We propose a registration method based on a combination of cubic b-spline registrations that is relatively quick (~4.5 minutes) and accurate (~6.3%). The methods presented in this work are being used to explore the relationships between regions of emphysema and their pulmonary microvascular blood flow during longitudinal progression of COPD.

Table of Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Professor Andrew F. Laine, who provided me the opportunity to come to Columbia and then very patiently guided me through the completion of this work. His lab provided me with the unique opportunity of working on many different collaborative research projects, as well as introducing me to colleagues whose friendships I will treasure for a lifetime. Besides my work on this thesis, Professor Laine has significantly shaped both my personal and professional development for the greater part of the past decade, and my achievements and future career path are in large part based on experiences and insights that he has imparted to me. He introduced me to the world of academia and the IEEE Engineering in Medicine and Biology Society (EMBS), which ultimately led to a journey that included participation in many conferences around the world before serving as chair of such academic conferences, and eventually even to the creation of a new academic journal at IEEE. He also introduced me to the joy of working in collaborative teams of research scientists, industry professionals, and medical professionals, which played a large part in leading me to further pursue training as a physician. I would of course also like to thank the other members of my committee, Dr. Graham Barr, Dr. Elisa Konofagou, Dr. Shunichi Homma, and Dr. Paul Sajda, whose insight and constant support were essential to me during this work.

I was given the unique opportunity of completing my thesis work while undergoing further

clinical training at the Columbia College of Physicians and Surgeons. This was an experience that was at times trying, but that has significantly shaped both the physician and the scientist that I am today, and I would not trade it for anything in the world. It would be an understatement to say that without the understanding and support of my advisor, Professor Laine, the support of the medical school administration (Dean Lisa Mellman, Dean Stephen Nicholas, Dr. Patrice Spitalnik, Dr. Katherine Nickerson, and Dr. Steve Shea), and the assistance of the many clinical faculty and residents who worked with me in the past few years, none of this would have been possible.

Thanks to my advisor, I was able to become particularly involved in the largest biomedical engineering society in the world, the IEEE Engineering in Medicine & Biology Society. Through my involvement in the society, I was given the opportunity not only to attend conferences, but to take on leadership roles in the society. This would not be possible without the tremendous support of the exceptional staff at EMBS, past and present: Laura Wolf, Dana Bernstein, Jessica Lotito, Laura Herrara, and the various colleagues that I've had the pleasure of collaborating with on the creation of a new journal and the organization of various conferences: Dr. Paolo Bonato, Dr. Martha Gray, Dr. Metin Akay, Kara McArthur, and Dr. Clifford Dacso.

My friends around the world have been an important and constant source of support both during my work on this thesis and during the boards. They have allowed me to vent ad infinitum without calling the asylum hotline, for which I am eternally grateful. Roanna Kong, Tiffany Guo,

Eric Chen, Kelly Yeh, Claire Uebbing, Anna Starikovsky, Dmitry Sagalovskiy, Parisa Rahman, Stephen Ma, Mohammed Shaik, Laura Blanar, Dan Tammuz, Joseph Brownstein, Rodica Buzescu, Justine Li, Gayan Edirisinghe, Kartik Trehan, Haytham Elhawary, Jennifer Walz, and the entire Columbia College of Physicians and Surgeons class of 2015 have made these past few years possible. During the completion of this work, one of my good friends, Stephanie Yen-mum Liem Azar, unfortunately suddenly passed away. I will always remember her fondly and treasure those many late-night conversations between us that kept me sane through trying times.

I would like to thank my colleagues and various students whom I've had the pleasure of working with in the Heffner Biomedical Imaging Lab: Dr. Matthew Niemark, Dr. Noah Lee, Dr. Ting Lee, Dr. Qi Duan, Dr. Antonio Albanese, Viktor Gamarnik, Guillaume David, Frank Provenzano, Yrjo Hame, Won Hee Lee, Arthur Mikhno, Nikolaos Karmolegkos, Angelica Chen, Alvaro Gomez, Alan Zambeli-Ljepovic, Kevin Yeroushalmi, Christel Ducroz, Murielle Hadad, Josh Aschheim, Teresa Kao, Niccola Perez, Jia Guo, Auranuch Lorsakul and the many other colleagues at Columbia and elsewhere from whom I've learned so much.

It is often said that behind every man (woman) there is a village, and it must already be clear from this long list of acknowledgements that my case is no exception. This statement is even more applicable in my case due to the incredible team of support staff that has helped me in all areas of my life these past few years. I would like to particularly thank Aileen Mistica, my chief of staff, who has stabilized my sanity for the better part of the last half-decade. The various

members of my staff past and present have contributed more to the accomplishments of my life than I can verbalize: Alex Saratan, Ravi Choudhary, Ioana Mitrea, Ruth Morallo, Hassam Sheikh, Vangjush Vellahu, Mohammad Afzaal, Enric Immus, Shirly Gerodias, and Jonathon Macfarlane. Jonathon in particular has been instrumental in editing the many papers that I've worked on in the past few years, and any and all references relating to the superiority of the Canadian healthcare system over the United States' in any of my papers are likely a result of his urging.

I am forever indebted to my family, and especially my parents. Their unconditional support and encouragement through these years has been a constant source of strength and motivation, without which I could not have completed this journey.

Finally, I would like to specifically thank my girlfriend, Lisa Huang, for her love and support. Without her infinite patience and support, I would not be where I am today.

Chapter 1: Introduction

In recent years, multi-resolution techniques have become increasingly popular in the image processing community. Techniques have been developed whose applications range from edge detection and texture recognition to image registration, multi-resolution features for image classification and more. However, the idea of analyzing images, or indeed signals, at different scales of resolutions is not itself new. Researchers in fields ranging from geology through signal analysis to physics have all considered this idea in various guises. Formalized multi-resolution architecture offers an efficient framework for extracting information from data at several different levels of resolution. Multi-resolution decompositions such as pyramid algorithms can be used to reduce the dimensions needed to represent an image, thus reducing the computational complexity involved in image processing operations. The same algorithms using the same dimension-reduction principles can also be used to efficiently represent data and images in a compact and efficient manner, allowing for applications in both data compression and image representation.

In the first major part of this work, we used multi-resolution decomposition to help with the classification of a set of protein crystallization screening images produced by a high throughput structural genomics pipeline. In the second, a multi-resolution B-spline pyramid architecture was used to help register a set of 3D image volumes that would otherwise be computationally

prohibitive.

1.1. **OUTLINE**

This dissertation begins with a description of multi-resolution image analysis methods (Chapter 2), describing in broad strokes how they are currently used in medical image registration and image classification, while providing specific details that lay the necessary groundwork for understanding the remainder of the thesis. We proceed to discuss two specific applications of multi-resolution image analysis, using two application cases. The first is the use of multi-resolution image features to classify a set of protein crystal screening images produced by a high throughput structural genomics pipeline. Chapter 2 first discusses the relevant background information on protein crystallization screening and the current pipeline that has been developed in high throughput structural genomics centers. Chapter 3 and Chapter 4 then introduce the newly-built classifier and discuss the results obtained by using its multi-resolution features. The second application case used B-spline multi-resolution registration methods in order to register two medical image volumes (CT and MRI) in order to further elucidate the pathophysiology of disease process in Chronic Obstructive Pulmonary Disorder (COPD). Chapter 5 is focused on introducing the relevant current clinical information on and pathophysiological understanding of COPD. Chapter 6 then specifically introduces the multi-resolution B-spline-based registration algorithm used to register the two different CT and MRI medical volumes.

Chapter 2: Multi-resolution Image Analysis

In partial differential equations, all current finite element numerical solvers depend on finding approximate solutions to boundary value problems by first grossly sampling the solution space and then further evaluating the neighborhood of likely candidate solutions. This allows the numerical solvers to find solutions that would otherwise be prohibitively long to derive. In image processing, similarly, the development of multi-resolution analytical frameworks has allowed for classes of algorithms that would otherwise take too long to be usable in practice. Furthermore, multi-resolution analytical methods in image processing have allowed for the development of new classes of algorithms that can leverage information that exists at high-resolution scales as well as low-resolution scales, much akin to how we humans process images. These techniques now have applications in edge detection, texture recognition, image registration, image compression, and more.

2.1. **IMAGE REGISTRATION**

Image registration is the process of determining the optimal spatial mapping that matches two images to each other [1]. The data may be from different imaging modalities, times, sensors, viewpoints, etc. In medical imaging, the most common reasons for image registration include the registration to each other of different medical imaging modalities such as CT and MRI, the registration of medical data over time, such as the tracking of cancer progression in sequential

CTs, and the registration of data between different patients so as to find commonalities within patient groups, which is commonly done in fMRI.

Though the actual algorithms used in this work were implemented in 3D, at this point we can introduce the basics of medical image registration in 2D without loss of generality. Given a pair of images $\mathbf{I}_1$ and $\mathbf{I}_2$, image registration tries to determine an optimal function that maps the pixels in $\mathbf{I}_2$ into the coordinate system of $\mathbf{I}_2$. In order to evaluate whether a particular mapping is optimal, a similarity measure (cost function) has to be defined a priori, and typically the set of possible mappings is evaluated to find the mapping that produces the best score in terms of the chosen cost function. This leaves the question of how to define the set of possible mappings. The mapping function must also be chosen a priori, and is typically constrained according to the registration problem of interest. Otherwise, the set of possible mappings would essentially be infinite and therefore computationally intractable.

2.2. LUNG VOLUME REGISTRATION

Registrations of lung data sets have been applied in establishing lung atlases [2], linking four-dimensional lung data sets [3], and tracking the motion of lung tissues [4] [5]. The registration problem tackled in this work focuses on registering lung volumes of differential breath hold statuses, as well as volumes that were acquired by different imaging modalities (CT and MRI). Due to the nature of the imaging modalities, and the desire to be consistent with

current research and clinical communities, CT lung volumes for emphysema quantification are done after patients fully inhale (Total Lung Capacity / TLC), while MRI lung volumes for perfusion, the volumes of interest in this work, are typically done after a patient exhales (Functional Residual Capacity / FRC). This is partly for historical reasons, but also partly due to the extended amount of time that MRI acquisition takes as opposed to image acquisitions in CT.

In image registration algorithms, it is necessary to define a similarity measure in order to evaluate the "fit" of candidate mappings. Typically, similarity measures include Sum of Squared Intensity Differences (SSD), Sum of Absolute Differences (SAD), Correlation Coefficients, and measures that derive from information theory such as entropy, mutual information, etc. Each of these similarity measures contains certain assumptions about the image volumes that are being registered. For example, the application of SSD as a similarity measure assumes that the two volumes being registered differ only by Gaussian noise after registration (i.e. two photos that have been rotated). Other similarity measures allow for a less strict relationship between the two image volumes, such as mutual information, which simply assumes that there is a statistical relationship between the pixel intensities of the two volumes that are being investigated. However, since lung motion, and – especially in the case of COPD patients – the relative aeration of lung tissues differ from region to region in the same lung, voxel intensities do not follow a simple global statistical relationship that can be directly applied.

In order to register these two very different types of image volumes (CT at TLC vs. MRI at

FRC), we chose to use an intermediary dataset (CT at FRC) in order to perform this image registration. Previous groups who have done extensive work on the registration of CT lung volumes found that image registration algorithms underperform for CT volumes belonging to the same patient but differing in breath-hold status [6], [7]. Thus, we did not believe that direct registration of image volumes that differ in both breath hold status and imaging modality would be appropriate. Instead, we performed two separate image registration mapping calculations: 1) we computed the mapping that will register a CT at TLC image volume with a CT at FRC image volume, and 2) we computed the mapping that will register the CT at FRC image volume with the MRI at FRC image volume. The two computed mappings were then performed serially in order to finally achieve our goal of registering the CT at TLC image volume with the MRI at FRC image volume. The two mappings were computed using the same general framework, but with different similarity measures. We adapted Choi. et al.'s sum of squared tissue difference as our first similarity measure [6], [8]. We then chose to use mutual information as the similarity measure of the second mapping, as a global statistical relationship can be assumed and has been well validated between CT and MRI image volumes.

A second major issue that must be addressed is the selection of the space of mapping functions that should be explored for these two registration steps. As mentioned earlier, lungs are differentially inflated by region during different phases of the respiratory cycle. Lamare et al. [8], [9] explored affine transformation mappings for respiratory motion correction, but reported that it

was sufficient only for a single organ and associated lesions. Effective correction of differing respiratory states thus requires non-rigid image registration, which allows for more flexible matching of local details between two images than is possible with rigid registration. It follows that we could not use rigid transformations as a basis for our mapping functions. A class of non-rigid transformations that preserve the relative computational efficiency of rigid transformations but allow for non-rigid mappings are B-splines. B-spline transformation additionally guarantees certain “nice” mathematical properties such as smoothness and compact support that makes them particularly suitable for medical imaging. However, in the absence of appropriate constraints, B-splines can still lead to non-realistic transformations such as folding of the image, which does not happen in medical image processing. Efforts have been made to regularize non-rigid image registration based on B-splines by making certain reasonable assumptions. Rueckert et al. [9], [10] enforced a smoothness constraint on the deformation of tissues by directly penalizing the bending energy of the deformation. Sorzano et al. [10], [11] proposed a regularizer based on the gradients of the divergence and the curl of the displacement field. Rohlfing et al. [11], [12] enforced a volume preserving/incompressibility constraint by requiring the Jacobian determinant of a transformation to be 1. Jacobian values represent the local volume contraction or expansion of the deformation field. A Jacobian value of less than 1 reflects local contraction, while a Jacobian value greater than 1 reflects local expansion. A Jacobian value of less than 0 would indicate that the deformation field has folded onto itself locally. Thus in cases where we are less interested in the absolute incompressibility constraint

but would like to avoid local folding deformations, we can enforce a looser constraint than Rohlfing et al.'s. In this work, our only constraint on the Jacobian was that it be strictly non-negative, which is equivalent to imposing a condition of local invertibility on the non-rigid transformation [12], [13].

2.3. **IMAGE CLASSIFICATION**

Image classification is the process of classifying an image based on its visual content. This can be as simple as determining whether an image is black or white, but can also grow in complexity into algorithms that, for instance, help military drones identify potential targets or help track the progression of cancer in patients. Image classification algorithms are generally segregated as supervised or unsupervised. In supervised algorithms, the number of classes are known ahead of time, and the problem at hand is to determine the correct class labeling for the unlabeled images. In unsupervised algorithms, the number of classes are not known a priori, and it is up to the algorithm to determine how many potential classes should be created in order to best differentiate images within the dataset. An example would be classifying fish in the ocean. A supervised algorithm may only be interested in distinguishing between a shark and a dolphin. An unsupervised algorithm might be used to determine the number of different types of fish that exist in a batch of 1000 images taken of fish. Here, we mainly focus on supervised algorithms, since they best respond to the type of problem tackled in this work.

In supervised image classification, the general framework consists of first locating the regions of interest, i.e., in our example, finding the fishes that need to be analyzed in the photos. The region of interest then needs to be described in terms of image descriptors, which are commonly termed "features" in the image processing community. After the features are extracted, the problem is similar to most data classification problems, and different types of classification schemes, including EM clustering, SVM, graph structures, min cuts, etc., can be applied to the features in order to compute the proper classifications.

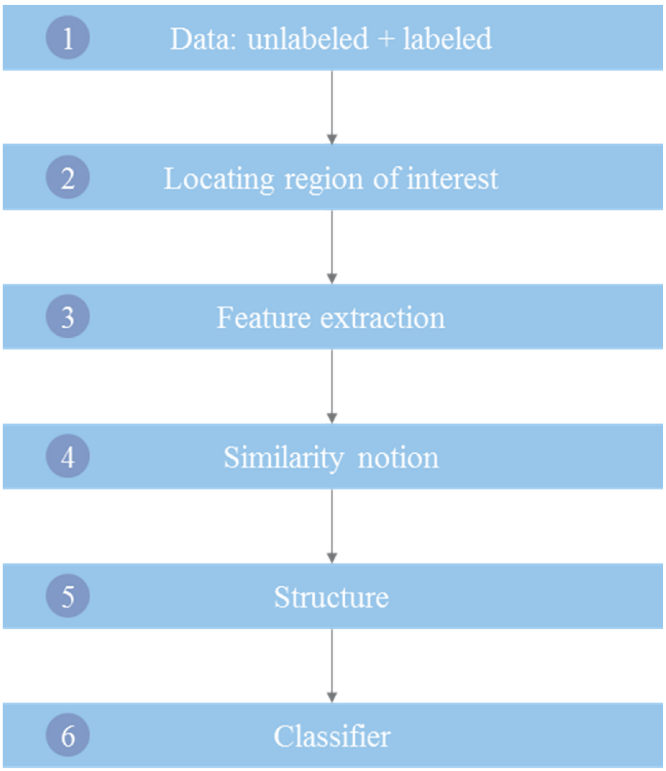

Figure 2.1: *Typical processing steps in a supervised learning algorithms.*

The performance of an image classification algorithm is dependent on two major factors: 1) the selection of features which provide the input to the classification algorithm, and 2) the

classification method used in the algorithm itself.

Feature selection is a consistent problem in machine learning. Data overfitting arises when there are a large number of features and the number of samples used in training the classifier is fairly small. Additionally, overly complex image features significantly slow down the classification process, leading to impractical algorithms for certain real-world applications.

2.4. **FEATURE SELECTION**

One of the simplest image features that can be computed is the mean of intensities of the pixels in an image. One can see how this image descriptor would be sufficient to classify images that are light vs. images that are dark. However, this simple image descriptor will not be useful for more complex classification tasks. Thus, the image processing community has generated a number of image descriptors that are also based on an image's histogram, such as using the various moments of the grey level histogram as additional image descriptors. It is important to note that all image features computed from an image's histogram will, inherently, only represent the global properties of an image, but be unable to distinguish its more complex topology, such as textures. Multi-resolution image features are therefore essential for complex classification tasks where the images of interest are globally similar, but their differences are derived from small parts of the region of interest.

Chapter 3: Crystal Image Recognition

The task of determining protein structures has been a major aim of the biology community for the past several decades [14]. With the completion of the Human Genome Project (HGP) in February 2001, it became evident to the biological community that the wealth of information available in the genomic sequence was matched by the amount of information that was inaccessible without a proper understanding of the proteins that these genomic sequences encoded. Furthermore, it became clear that many diseases cannot be attributed to errors in the genomic sequence, but rather to the interplay of abnormal proteins generated from normal genetic data. Some of these abnormalities in protein functions can be attributed to proteins with abnormal constituent components, but others can be ascribed to proteins possessing the correct constituent components but that have formed an inappropriate structure, akin to incorrectly assembling a piece of Ikea furniture and ending up with a chair rather than a bookshelf, even though none of the parts for the bookshelf are missing or poorly made.

There currently exists no method of performing a de novo prediction of a protein structure purely from its genomic sequence, and protein structure determination thus currently relies on experimentally determined structures. By the completion of the Human Genome Project, approximately 200 new protein folds were being solved per year. The spectacular success [13], [15] of the highly collaborative HGP prompted the creation of several structural genomics

consortiums in the hope that a similar systematic effort in high throughput structural genomics might be able to significantly increase the number of structures solved, and do so at a much lower cost than is currently the case [15], [16]. These consortiums focused on creating high-throughput crystallography pipelines, one component of which forms the focus of part of this thesis.

3.1. **BIOCHEMICAL BACKGROUND**

3.1.1. **Protein Structures**

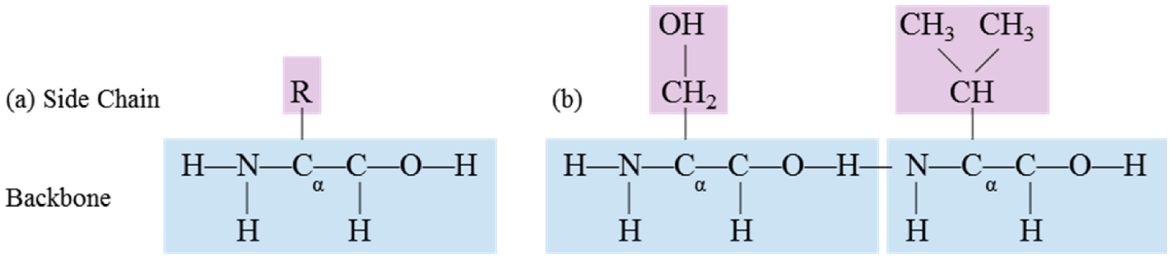

Figure 3.1: *Chemical structure of amino acids. A single, generic amino acid is shown in (a). R denotes the side-chain atoms-acid – unique to each amino-acid type. The backbone atoms are shown in the bottom (blue) box. In (b) I show two specific amino-acid residues (serine and valine) linked by a peptide bond.*

At their base, proteins are formed of amino acids, which serve as biological molecules in their own right and fulfill a variety of functions. Several hundred of these amino acids exists, but only 20 play a predominant role in the human body. A characteristic R-group determines each amino acid's chemical nature and, therefore, how it interacts with other amino acids, other molecules,

and its environment.

Amino acids link together via peptide bonds to form peptides and proteins. These peptides/proteins play an essential role in almost all cell functions in living organisms, including in catalyzing reactions, cell signaling, and structural support. The exact linear order in which amino acids are linked together is a protein's primary structure.

After amino acids are linked into this primary structure, additional processing steps take the linear structures and fold them into their final three-dimensional shape. The final three-dimensional configurations of a protein depends on a host of factors, including intrinsic bonding forces inside the protein as well as the external environment and, often, interactions with other proteins.

3.1.2. **Protein X-Ray Crystallography**

Several methods exist to determine the structure of proteins, though over 90% of the protein structures currently available in the Protein Data Bank were determined using X-ray crystallography. In recent year, NMR methods have also been used to obtain three-dimensional models of small protein molecules, but it has not yet been widely adopted and the technique is still considered immature by the crystallography community [16], [17].

The end goal of both X-ray crystallography and NMR-based protein structure determination

methods is to solve the tertiary structure of a protein of interest. We focus on X-ray crystallography, as that is the primary method used by high-throughput structural genomics consortiums.

Protein structures cannot currently be directly observed, and X-ray crystallography is thus an experimental method used to infer a protein's structure by examining its figurative shadow. This is akin to reading radiographs of the human body from different angles and building a best guess of the anatomical locations of various organs in 3D. Much like conventional X-rays, X-ray crystallography involves directing a beam of X-rays onto the crystallized version of the protein of interest. The resulting image produced is a 2D electron-density map that is then inversely solved by computationally finding the best fit among various candidate 3D conformations of the protein of interest into the electron-density map. The quality of the final tertiary structure is thus dependent on the resolution of the electron-density map, which is in turn dependent on how well the protein crystal is able to diffract the X-ray beam. The primary determinant of a protein crystal's ability to diffract the X-ray beam is how well-ordered the protein crystal is. Thus, of primary interest to X-ray crystallographers is the ability to crystallize their protein of interest, and then to further improve their results by choosing crystallization conditions that produces large and pure crystals.

Crystallization is neither easily achieved nor rapid; it can require several months for sufficiently large crystals (~0.5 mm) to grow from microcrystals. Furthermore, a number of different

parameters, including pH, temperature, protein concentration, the nature of the solvent and precipitant, and the presence of added ions or ligands to the protein have a critical effect on crystal formation. Finding the few correct combinations of these parameters that might provide crystals suitable for X-ray diffraction analysis is therefore a lengthy task requiring multiple experiments.

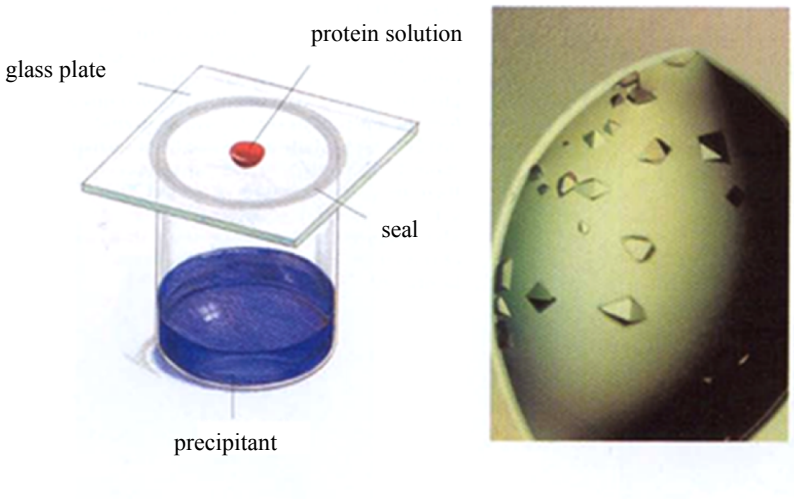

Figure 3.2: *Crystals of recombinant enzyme RuBisCo from Anacystis nidulans formed by the hanging-drop method. (Courtesy of Janet Newman, Uppsala, who produced these crystals.)*

3.2. **NORTHEAST STRUCTURAL GENOMICS CONSORTIUM (NESG)**

Due to the large number of conditions that have to be screened in order to grow crystals suitable for X-ray diffraction analysis, the structural genomics consortiums have all attempted to automate the screening of these crystallization conditions. Significant resources and efforts have

been spent on creating micro-crystallization experiments that 1) use a relatively small amount of protein that are typically difficult to produce/purify, and 2) can be standardized and reproduced at sufficient scale to simultaneously screen multiple reagent conditions and multiple proteins of interest.

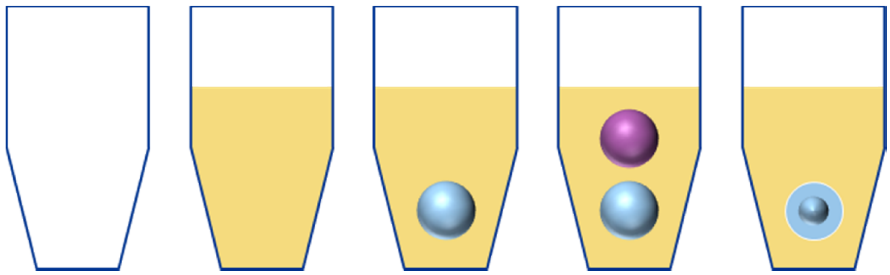

Figure 3.3: *An example of the microbatch-under-oil crystallization experiments as contrasted with the previous figure's hanging-drop experiments.*

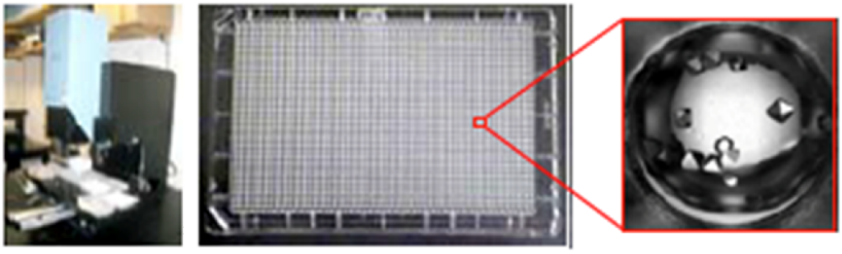

Figure 3.4: *Left: A robotic system used at NESG for the microbatch-under-oil crystallization experiments. Middle: An example of the 1536 well plates used. Right: An example of an image taken from under the well.*

Currently, in a typical high throughput setup a very small amount of protein (1.5 microliters or even less) is plated into micro-wells and a thousand combinatorial arrangements of several crystallization parameters, such as temperature, pH, and solute, are simultaneously tested to find

conditions where promising crystallization can be detected. Because time scales for the formation and dissolution of crystals differ, images from multiple time points are recorded for each experimental condition. Consortiums in structural genomics such as Northeast Structural Genomics (NESG) now perform tens of millions of such micro-experiments annually, resulting in a need to analyze an even larger number of images. At present the images are classified manually. With over 10 million images generated annually, crystallographers can only inspect one time point per experimental setup even though photos from multiple time points exist, in order to evaluate whether a specific experimental condition might have produced viable crystals. Adding to the difficulty of classification is the relatively small rate of crystal images vs. non-crystal images in the real world (1%), which leaves a particularly delicate balance for any classifier between missing images containing crystals or generating so many false positives as to make the high throughput system worthless for practical use.

Thus, one major focus of this thesis is the development of an image classification framework that can automate the evaluation of the images captured to determine whether certain micro-wells contain crystals. Several groups have previously proposed automated algorithms to help such analysis. The proposed algorithms all involve the use of supervised learning algorithms. These mainly revolve around the use of neural nets [18], [19], Linear Discriminant Analysis [20], [21], or the use of support vector machines [22] [23]. According to current literature, both classes of supervised learning algorithms have been implemented without significant success. This can be

attributed to the relatively few optimizations that were performed, as well as the lack of training and testing data available to the previous groups. Furthermore, none of the classifiers described in the literature thus far perform at fast enough speeds to be practical in a production structural biology pipeline.

The current backlog of images at the NESG alone exceeds 50 million images. To finish processing this backlog in 5 years, an algorithm must process more than 20 images per minute (i.e. one image every 3 seconds) as opposed to the current average speed of one image per 30 seconds, or 2 images a minute [18]. That is, the algorithm must perform at ten times the current speed to prevent the backlog from taking 50 years to process. This speed up can be achieved either via the use of a supercomputing grid, as proposed by Cumbaa et al., or by a classification algorithm that can be parallelized and run sufficiently fast such that two or three off-the-shelf computers can accomplish the task, as proposed in this paper.

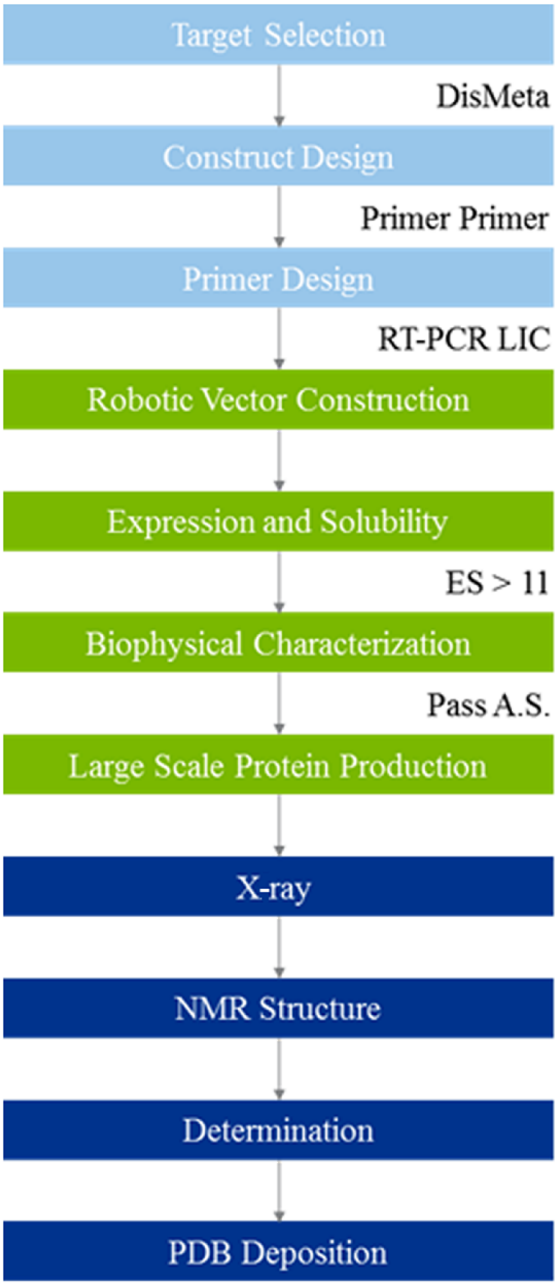

Figure 3.5: *Adapted from [17], [18], showing the NESG structural genomics high-throughput pipeline for protein production and protein crystallization screening.*

Chapter 4: Protein Crystal Image Classification

One of the primary issues in automating the classification of protein crystallization images is the varying size, shapes, and textures of the crystals present in the images. In some images, crystals with regular polyhedral shapes can occupy the majority of the region of interest. In other images that are equally promising of potential crystallization, small crystals spanning only a few pixels can be distributed throughout the image or localized in a small area of the image. Thus, any classification scheme must be able to cope both with image features that sometimes span the entire well and with image features that can be extremely localized. Textural-based classification methods can be very suitable for the detection of global features, but special attention must be paid to correctly identifying the Region of Interest (ROI). Similarly, purely shape-based techniques are extremely dependent on the correct identification of potential crystals, and are additionally sensitive to the added noise of elements in the image such as skin, precipitates, etc. To tackle both the ROI concern and the need for speed and thus parallelizable calculations, we chose to resample the images into sub-images (windows). Each window was then put into our classification algorithm and a voting system was used to evaluate the feature vectors computed from each window. The flowchart of our method is shown in Figure 4.1. First, we identified the image's ROI; the system then extracted and selected a set of windows from the ROI. For each selected window, a feature vector was then computed, based on the texture and shape analysis.

Finally the feature vectors were classified using a random forest classifier.

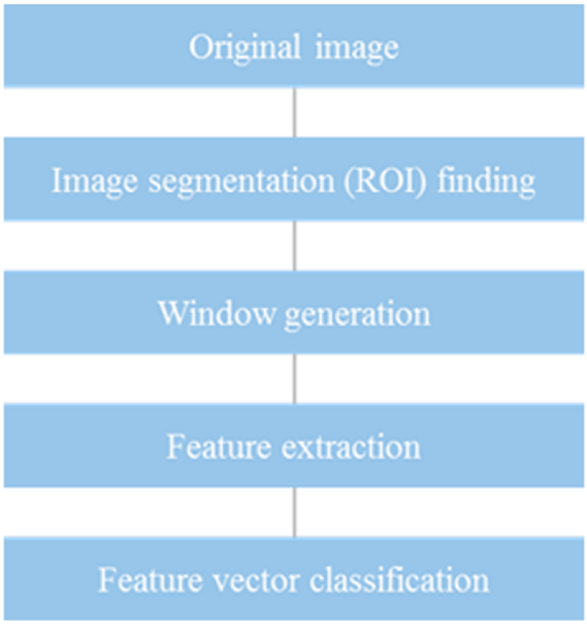

Figure 4.1: *Flowchart of the method used.*

4.1. **DATA**

Three data sets were used in the validation of the algorithm described in this work. Each data set consisted of grayscale (8-bit) protein crystal micro-well images measuring 632 by 504 pixels. Sample images can be seen in Figure 4.3. One data set consisting of 63,445 images derived from micro-well experiments on 96 different proteins was manually classified by 3 independent crystallographers at the Hauptman Woodward Medical Research Institute at the State University of New York at Buffalo [24], [25]. Each image was categorized into one or more of the

following categories: Clear, Phase Separation, Precipitate, Skin, Crystal, Garbage, Unsure. Unsure images were not used for the analysis, resulting in a total of 63,023 images from the f96 study. Examples of the 10 outcomes of the f96 study are shown in Figure 4.4 below.

Due to the relatively small number of protein crystal images that resulted in the f96 protein study, two additional datasets were obtained in order to better improve the classifier. These two datasets consisted of micro-well images that were taken by HWI in collaboration with Structural and Functional Genomics Consortium (SGPP) and Northeast Structure Genomics Consortium (NESG). Both datasets consisted of images that were classified by crystallographers as containing crystals, but that do not further delineate between images that contain only crystals, or images that may contain crystals & precipitate or crystal & phase separation, as the f96 study does. Thus these two additional datasets were used later in the work to train the binary classifier, but not used for experiments involving the 9-way classifier.

Table 4.1 *The distribution of image classes from the three datasets used in the study. Only approximately 1% of all images in a representative study contain crystals.*

Class	F96 Images	SGPP Images	NESG Images
Clear	26,136		
Phase separation & precipitate	388		
Phase separation & crystal	26		
Phase separation	5,374		
Precipitate & crystal	63		
Precipitate & skin	3,822		
Precipitate	26,457		
Skin	608		
Crystals only	149		
Crystals or crystals + another class		10,373	5,814
Total	**63,023**		

4.2. METHODS

4.2.1. Image segmentation

As can be seen in figure 4.3, each crystallization experiment image consisted of a centered circular well and surrounding margins to ensure that the full well had been captured. Due to the number of experiments that are routinely performed, high throughput protein crystallization centers have developed robotics systems to perform the setup, loading / unloading, and subsequent imaging of the micro-well experiments [26]. This allows for the images that are obtained to be fairly consistent in brightness and contrast, and their region of interest (micro-well) to be reasonably centered within the picture's field of view. Nevertheless, even with the

development of the robotics system, there remained a sizable number of images where this was not achieved (Figure 4.2). Thus, the first step in our pre-processing algorithm required the correct identification of the micro-well.

Our earlier attempts at segmenting the region of interest included simply choosing the middle 64 x 64 boxes for further processing. Though this did not fully identify the entirety of the micro-well, it had the advantage of being extremely fast. It is conceivable that this would be sufficient for a classifier if the entirety of the micro-well were indeed fairly uniform in nature. However, we quickly noticed that the segmentation results would at times miss key identifiers in pictures (Figure 4.2). This led to an image classifier unable to recognize those images as crystals.

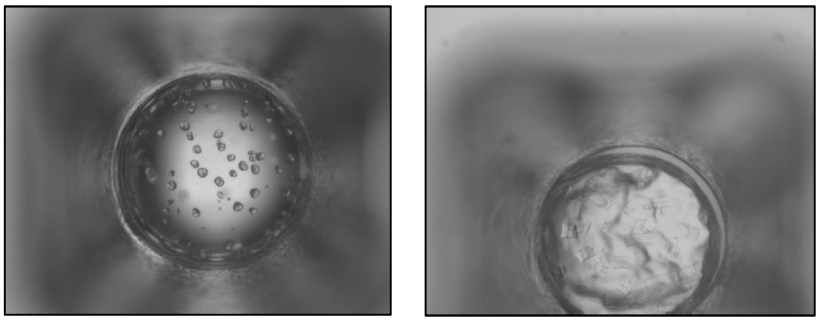

Figure 4.2: *On the left, a typical image that is relatively well centered with good contrast. The right image shows a well that is significantly off centered and is in fact not fully in the camera's field of view.*

We next tried to use an ellipsoidal multiple population genetic algorithm in order to segment the entire circular region that represented the micro-well. The choice of the ellipsoidal algorithm was made because it was noted that several micro-well regions in the training data set were not

perfectly circular. Further details about that algorithm can be found in a previous publication by Po et al. [18], [27]. Though the algorithm turned out to be extremely accurate, the processing time took over 50 seconds per image, and further optimization was unable to significantly reduce the processing time.

That segmentation algorithm was therefore abandoned and Circular Hough Transform (CHT) was evaluated. A circle with radius R and center (a, b) can be described with a set of two parametric equations [18], [27]:

$$x = a + R\,cos(\theta)$$

$$y = b + R\,sin(\theta).$$

CHT determines the location of circular shapes in an image by systematically cataloging possible (a, b, R) triplets in an image.

Specifically, for each pixel in an image, the set of possible circles as defined by their parametric triplets (a, b, R) were computed and incremented appropriately in a three-dimensional accumulator matrix. This process was iterated over all foreground pixels in an image. The maximal count in the accumulator matrix denotes the parametric triple (a, b, R) that belongs to the circle with the largest perimeter points in the image. This method is sometimes difficult to implement in practice due to its sensitivity to noise and the need to identify the approximate dimensions of the circle(s) that is/are being identified.

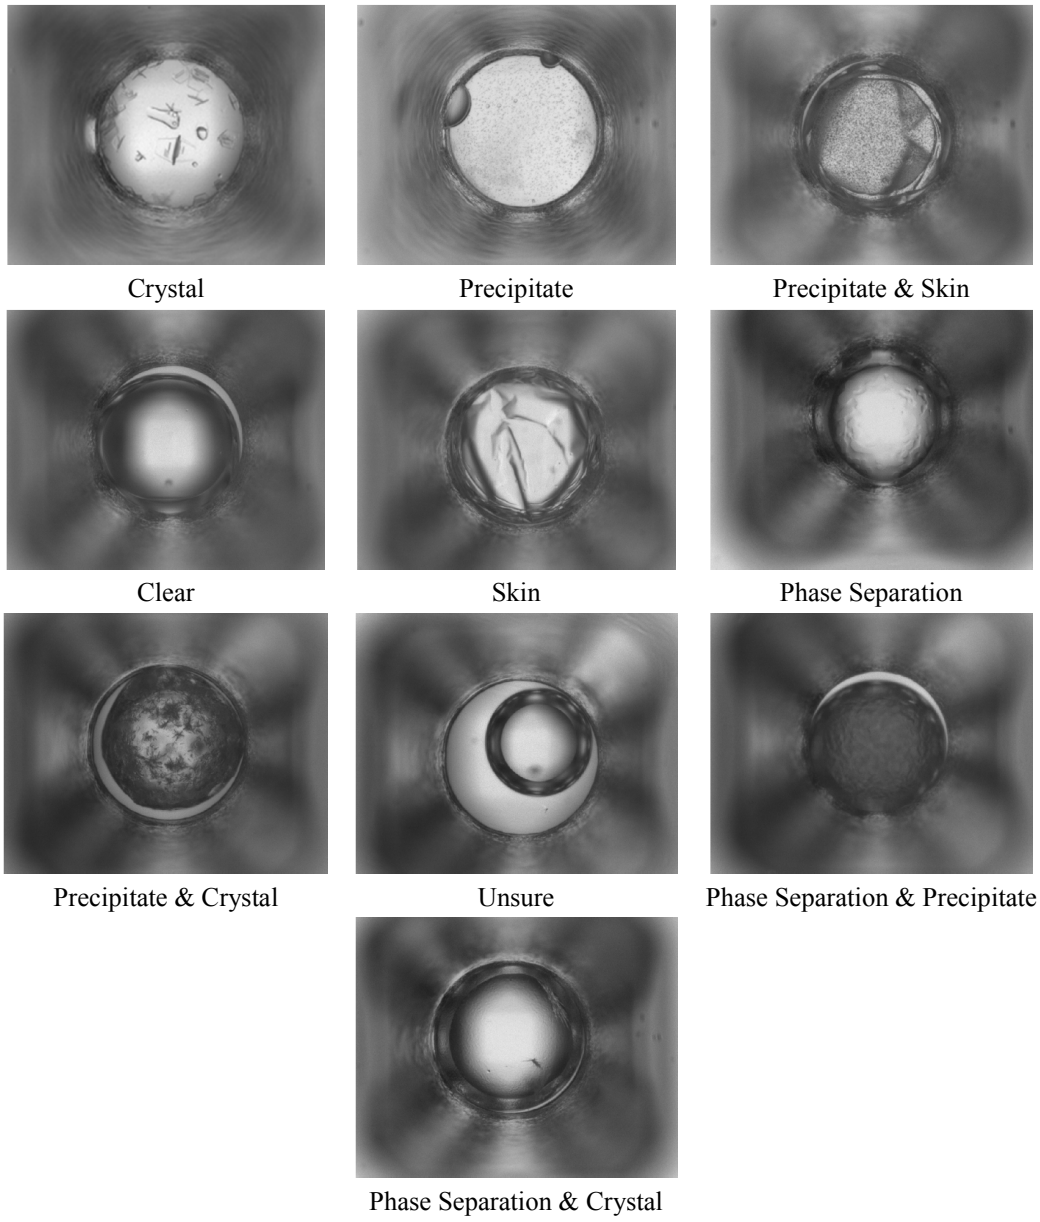

Figure 4.3: *Sample outcome images.*

However, the wells in our images were generally fairly well-centered and are similar in size. Thus, we were able to implement CHT and significantly speed up execution by limiting the accumulator matrix to only reasonable parametric triplets (a, b, R).

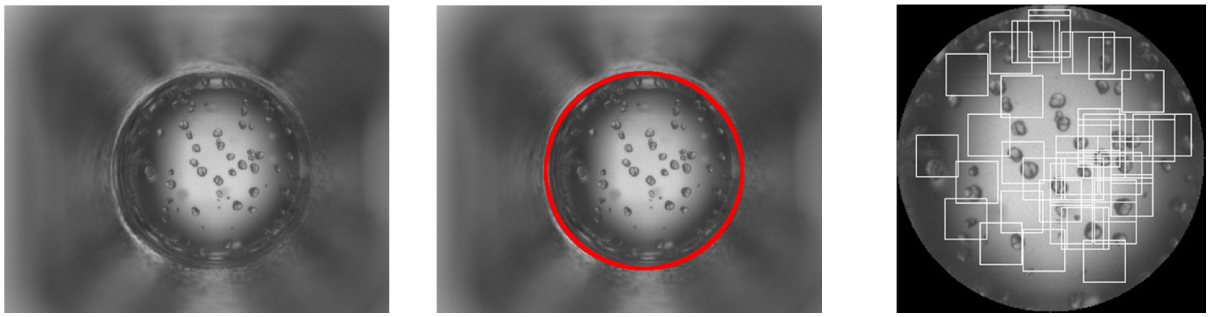

Figure 4.4: *On the left, a typical image as inputted to the system. The middle picture shows the identification of the well using CHT. The right picture shows the random window generation inside the segmented well.*

4.2.2. **Feature extraction**

Feature extraction allows us to represent high-dimensional image data with a set of vectors with lower dimensions, typically termed feature vectors. These feature vectors are not unique for each image, and are typically defined differently per application. All classification schemas are significantly dependent on these chosen feature vectors, and we therefore experimented with a large number of these before settling on our final set of features. Some of the features that were previously tested can be found in our earlier publications [18], [28]. In prior work, we primarily focused on using first-order and second-order image features derived from a Laplacian pyramid decomposition of the original full image. After significant parameter tuning of the classifier described above, we were unable to improve the positive predictive value (PPV) above 85% and we began to look towards the addition of other features to help improve the classification

performance.

We settled on four complementary image features to summarize the images: 1) Gabor wavelets transform, 2) Gray Level Co-occurrence Matrix, 3) Local Binary Pattern, and 4) Radon Transform. The Gabor wavelets transform captured the micro-texture that small crystals form in the image by analyzing a multi-scale multi-orientation representation of the images. The Gray Level Co-occurrence Matrix (GLCM) extracted second-order statistics from the image histogram by analyzing several aspects of the texture such as skewness, entropy, etc. Local Binary Pattern (LBP) was used to analyze the invariant texture properties. Radon transform was able to capture the relative orientation of any big crystals present in the image.

4.2.2.1. Gabor Wavelets

The Gabor wavelets transform was created in 1946 by Denis Gabor [28], [29]. It is one of the most popular feature extraction methods due to its biological relevance and its multi-resolution and multi-orientation properties, which are optimal for measuring local spatial frequencies and characterizing textures [29], [30]. The 2D Gabor transform is a Gaussian function modeled by a sinusoid oriented with a specific frequency W and a specific direction O. Its 2D definition in space domain is given by g(x,y) and in frequency domain by G(u,v) respectively [30], [31]:

$$g(x,y) = \left(\frac{1}{2\pi\sigma_x\sigma_y}\right)exp\left\{-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2}+\frac{y^2}{\sigma_y^2}\right)+2\pi jWx\right\}$$

$$G(u,v) = exp\left\{-\frac{1}{2}\left[\frac{(u-W)^2}{\sigma_u^2}+\frac{v^2}{\sigma_v^2}\right]\right\}.$$

The Gabor transform can be adapted into a wavelet by creating a family of pairwise similar Gabor functions that differ by dilations and rotations [30], [31]. A filter dictionary can be generated by dilatations and rotations of $g(x,y)$ through the function [30]:

$$g_{mn}(x,y) = a^{-m}g(x',y')$$

where $> 1, m = 0,1,2,\dots,M$, $n = 0,1,2,\dots,N$ and m,n specify the scale and orientation respectively. M is the total number of scales, and N is the total number of orientations.

$$x' = a^{-m}(xcos\theta + ysin\theta)$$

$$y' = a^{-m}(-xsin\theta + ycos\theta)$$

where $= \frac{nk}{N}$. The scale factor a^{-m} is necessary to make sure the energy is independent of m . The ideal sigmas can then be obtained with

$$a = \left(\frac{U_h}{U_l}\right)^{\frac{1}{M-1}}$$

$$\sigma_x = \frac{(a+1)\sqrt{2ln2}}{2\pi a^m(a-1)U_l}$$

$$\sigma_y = \frac{1}{2\pi\tan\left(\frac{\pi}{2N}\right)\sqrt{\frac{U_h^2}{2ln2}-\left(\frac{1}{2\pi\sigma_x}\right)^2}}$$

where the values σ_x and σ_y characterize the spatial domain and the filter band extension in the x, y direction. The values of U_l and U_h are the values of the lower and higher cutoff frequencies respectively [30], [32].

In order to extract information from an image using this method, it is necessary to convolve each image with a set of predefined filters with a number of scales S, rotations K, and frequency intervals $\mathrm{U_l}$, and U_h. Let $f(x, y)$ be a window image and $g_{mn}(x, y)$ the Gabor filter with scale m and orientation n, then the set of convolved images is obtained using [32]:

$$r_{m,n}(x, y) = f(x, y) * g_{mn}(x, y)$$

where r_{mn} is the convoluted image of scale m and orientation n. The descriptor E_{mn} is obtained by extracting the energy value of the convoluted image, defined by:

$$\mathrm{E_{mn}} = \sum_{\mathrm{x,y}} \left[\mathrm{r_{mn(x,y)}}\right]^2.$$

The feature vector is then constructed by concatenating the energy values of each convoluted image with the number of scales K and orientations S determinate the size of the vector [32], [33]. Figure 4.5 shows this process:

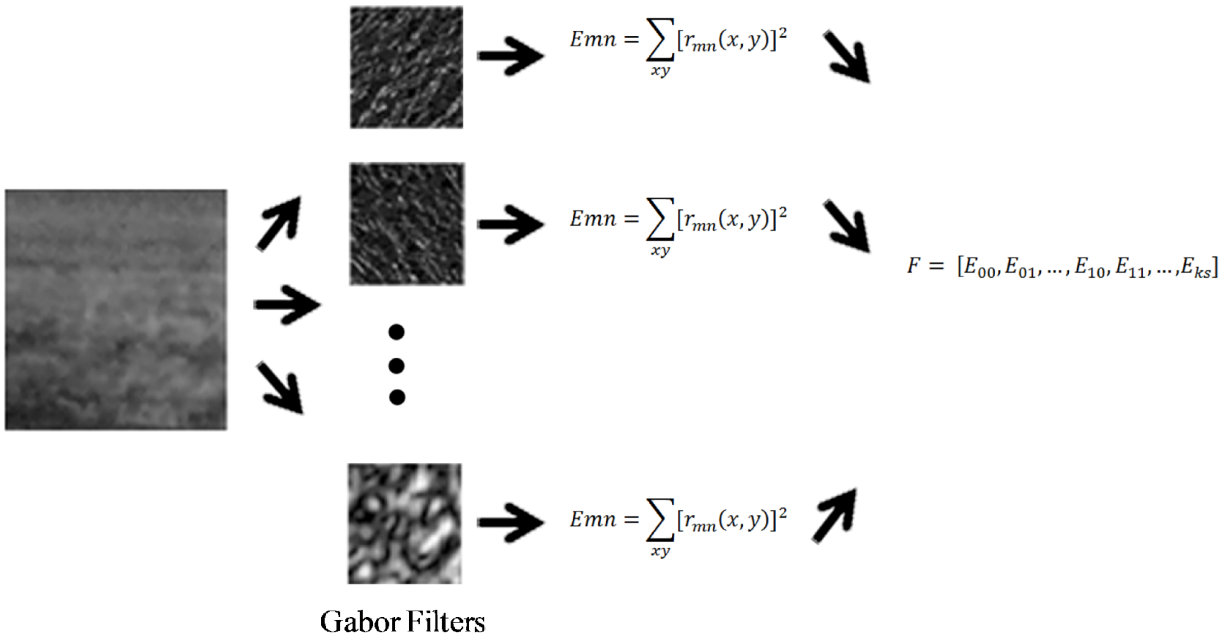

Figure 4.5: *Feature extraction using Gabor wavelets.*

4.2.2.2. Local Binary Pattern (LBP)

The multiresolution grayscale and rotation invariation Local Binary Pattern feature was first described in 2001 by Ojala et al. for the expressed purpose of fast and highly effective texture classification [34], [33]. Ojala's implementation of the classifier as described in their publications in 2001 and 2002 was used in this work. The LBP operator is defined by a convolution mask that labels pixels in a 3x3 neighborhood $f_p(p = 0,1,\ldots,7)$ of a central pixel f_c by the following transformation [33]:

$$S(f_p - f_c) = \begin{cases} 1, f_p <= f_c \\ 0, f_p = f_c \end{cases}.$$

The convolved 3x3 neighborhood is then converted into one number using the rule:

$$LBP = \sum_{p=0}^{7} S(f_p - f_c)2^p.$$

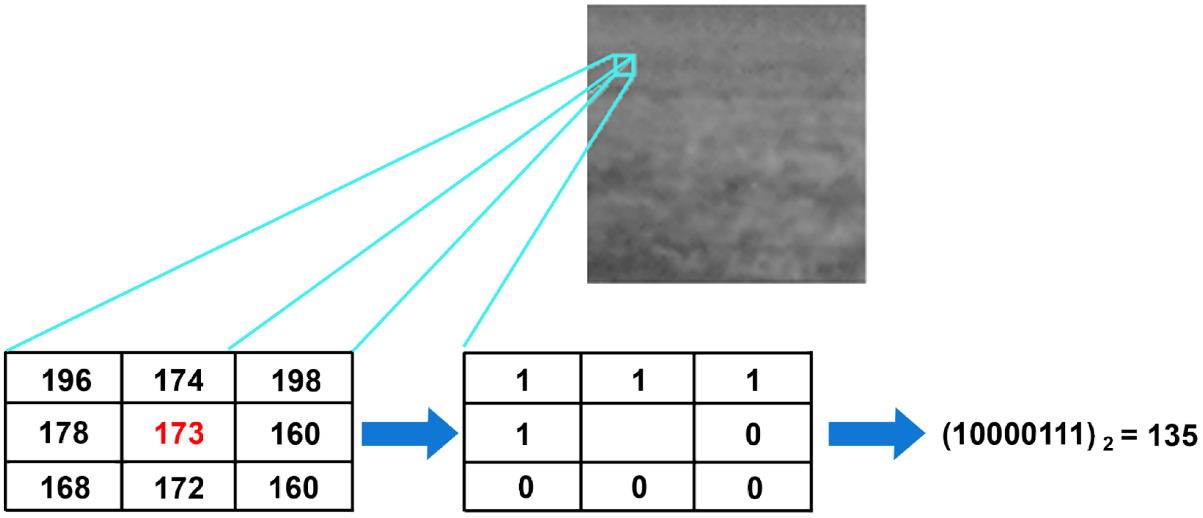

Figure 4.6: *Feature extraction using LBP method.*

The derived LBP binary codes are then bit shifted to derive the smallest possible bit code in order to preserve rotation invariance. Specifically, for binary numbers such as 01001000 and 00010010, both would be bit shifted to 00001001. This process is then repeated for every pixel in the image, resulting in a new local binary pattern coded image. The count of each unique LBP codes derived from the image histogram of the resulting LBP binary pattern coded image are then used as the actual features in the classifier.

The original LBP binary codes with 8 neighbors would result in 256 unique LBPs, but the bit shift procedure described above decreases the LBP feature space into only 36 unique LBPs. Thus this feature ultimately allows for 36 total dimensions for an input window of interest.

4.2.2.3. Gray Level Co-occurrence Matrix (GLCM)

GLCM describes second order statistics from an image histogram that are also locally invariant to gray level monotonic transformations, and are representative of the texture of an image [7]. The use of GLCM in texture classification is common in image processing, but the exact implementation varies to accommodate the need of the classification task at hand [35].

For all pairwise selected points in an image $f(x,y)$ with intensities i, j:

$$f(x_1, y_1) = i \; and \; f(x_2, y_2) = j,$$

we can define a co-occurrence matrix that describes the relationship between all such pairs by tabulating the relative frequencies $pd\theta(i,j)$ appearing in an image separated by a distance d in direction θ where:

$$(x_2, y_2) = (x_1, y_1) + (d\,cos\theta, d\,sin\theta).$$

This generates a matrix of dimensions equal to the number of intensities in the image for each distance d and orientation θ. For our implementation, we explored a radius of 4 around each point, or a total of 24 neighboring boxes. To reduce the computational cost of the algorithm, we

computed GLCM matrices quantized to 8 gray levels and four orientations $(0^0, 45^0, 90^0, 135^0)$. After the GLCM was fully calculated, 22 features were then calculated from the GLCM matrices, resulting in 176 total dimensions represented in the feature space. The 22 features included common features such as sum of average, sum of variance, sum of entropy, difference of variance, difference of entropy, sum of squared variance, maximum probability, entropy, energy, correlation, autocorrelation, and features that others have proposed for GLCM: inverse difference [36], normalized inverse difference [36], normalized first moment of the inverse difference [36], two types of correlation measures [37], dissimilarity [35], cluster shade [35] and cluster prominence [35].

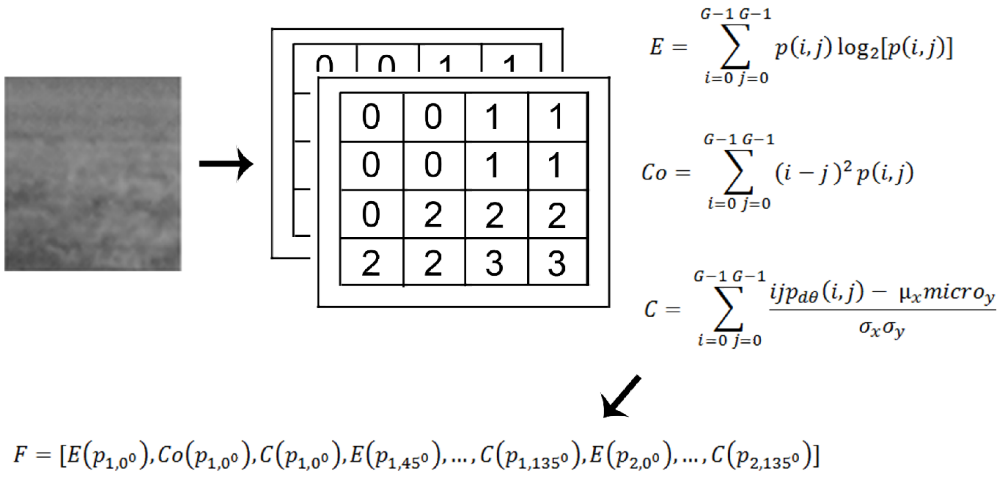

Figure 4.7: *Feature extraction using co-occurrence matrix.*

4.2.2.4. Modified Radon transform

The Radon transform has been used in image analysis to capture the directional information of

textures in an image. In our novel implementation, we used it as a shape descriptor, given its ability to detect large objects such as large crystals and accurately capture object orientation (in this case, crystal edges). A particularly useful property of this method is that when a texture has no significant orientation, the resultant feature vector from the Radon transform will be relatively empty of information and therefore non-contributory during classification. Thus, this image feature is able to pick out images with large orientations while minimizing its contribution during classification for images without such properties.

The Radon transform calculates the projections in specified directions of a texture or series of objects in an image by the equation [38]:

$$R_\theta = \int_{-\infty}^{+\infty} f(x'\cos\theta - y'\sin\theta, x'\sin\theta + y'\cos\theta)dy',$$

where

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

The operator maps the image $f(x,y)$ to the projection domain (p,θ), where θ is the angle and p is the distance to the origin. The transform generates an accumulator array of 180 points $[0\text{-}179]$ based on the spatial distribution of points in space. The maximum intensity in the accumulator array determinates the most probable angle of the texture or object in it. Figure 4.8 shows the result of applying the modified Radon Transform to a region of interest with a large crystal and thus significant directionality in the texture [38].

The transform accumulator array was then bucketed into 60 bins and, similarly to LBP, the count of each unique radon accumulator array code were then used as the actual features in the classifier.

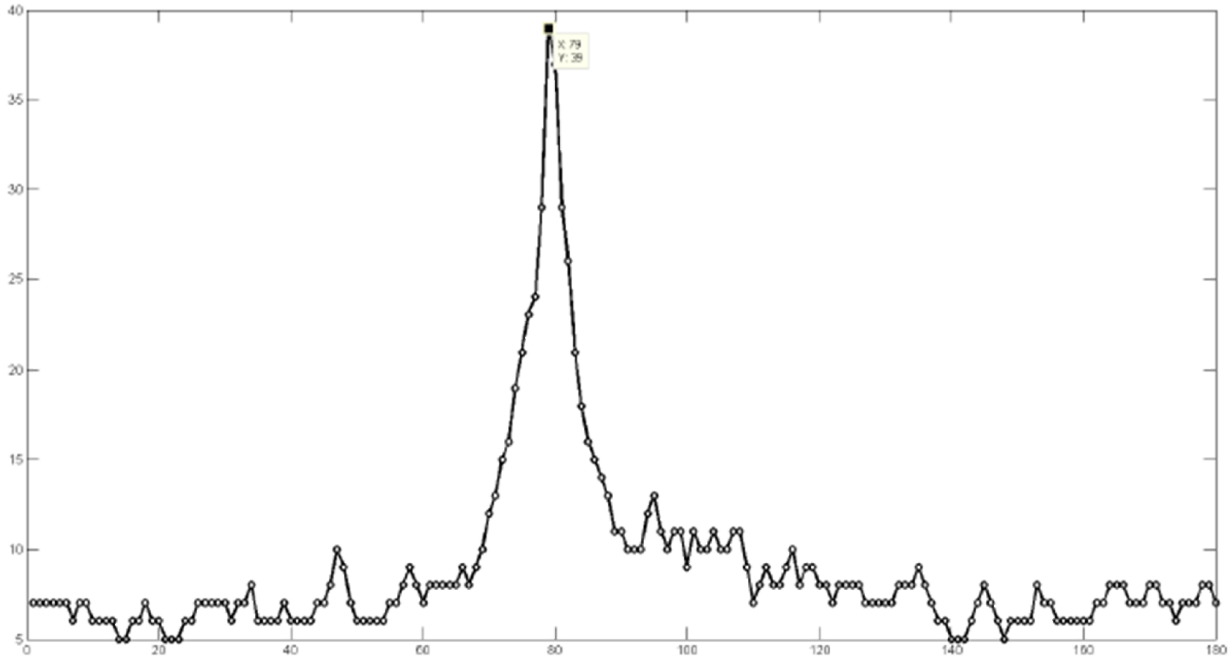

Figure 4.8: *Radon transform of a window image containing a large crystal.*

4.2.3. **Window generation**

After the well was located in the acquired image, we sought to subsample the well by identifying windows of 40 pixels x 40 pixels. 100 overlapping windows were generated by randomly generating X,Y positions to serve as the top left corner of each window (Figure 4.4). The large number of windows ensured that at least a small portion of any crystal present in the image was

captured. The windows were then clustered into 5 clusters, whereupon we sought to analyze the 25 most representative windows for further analysis. We computed the "distance" between image windows by using the variance and energy of each window as a metric. This approximation has shown success in different texture feature extraction methods [39], [40].

Variance and energy are first order statistics derived from the image histogram. Let $h(i)$h(i) be the histogram of a grayscale image and $p(i)$ a probability density function, then:

$$p(i) = \frac{h(i)}{NM}.$$

Where $i = 0,1,2,\dots,G-1$ and G is number of grayscale tones used in the image. The variance (ϑ^2) and energy (ee) can then be defined from the image histogram and p(i) as:

$$\vartheta^2 = \sum_{i=0}^{G-1}(i-u)^2 p(i)$$

$$e = \sum_{i=0}^{G-1}[p(i)]^2.$$

The above statistics were calculated for each window and the windows were then clustered using the Mahalanobis distance. Specifically, the distance between image p and image q can be defined as:

$$R(p;q) = \sqrt{(p-q)^T \sum\nolimits_{\cdot}^{-1}(p-q)},$$

where $\sum_{\cdot}^{-1}(p-q)\,\Sigma^{-1}$(p-q) represents the covariance matrix. The 25 windows that most represent

the different clusters were then chosen for further processing. This allowed the algorithm to process multiple windows simultaneously in parallel cores while also ensuring that all interesting regions of the image of interest would be examined by the classifier.

4.2.4. **Random Forest Classification**

The Random Forest (RF) classification model is an ensemble learning method that constructs numerous decision trees at training time and classifies results using the mode of the classes output by each individual tree. The basic premise is that even a group of "weak" classifiers can together outperform any individual "strong" classifier. Thus, instead of relying on any single classifier, a large number of separate decision trees are used as an ensemble. The use of a large number of separate decision trees also helps prevent high variance or high bias issues that are common in single decision tree classifiers. RF attempts to mitigate these two problems by creating a set of random decision trees (ensemble), and then averaging the results of each tree, thereby finding a natural balance between the two extremes. The strength of a RF classifier is dependent on the strength and the correlation between individual tree classifiers in the ensemble. If all trees in the ensemble were identical, the RF classifier's performance would be identical to the one decision tree, which defeats the point of an ensemble.

In order to ensure that individual tree classifiers in the ensemble are different, the RF implementation as suggested by Breiman [41] and adapted in this work trains the individual tree

classifiers using different portions of the training dataset. Specifically, each individual tree classifier trained on only 2/3 of the total training sample provided to the RF classifier. The 2/3 portion of the training dataset shown to each individual tree classifier was different; thus, no portion of the training set was wasted, while simultaneously this helped ensure that no individual tree was identical to another. In addition to only showing each tree classifier a portion of the training dataset, each tree classifier was also only shown a portion of the feature vector. This has the same effect as that described previously, and further served to differentiate each individual tree classifier. For each individual tree where M is the total number of dimensions in the feature vector, we only presented $\lfloor \log_2 M + 1 \rfloor$ of the dimensions. This parameter was chosen according to work by Breimen that showed it to be optimal for a number of different applications [41].

Other models that were attempted for this classifier include Naive Bayes, Support Vectors, Multilayer Perceptrons, and K*. In our and other groups' unpublished works [18], [42], we have found RF to be the strongest and most consistent classification model for this work.

4.3. **EXPERIMENTS AND RESULTS**

All experiments related to this study were run on a CentOS 5.6 server with dual Intel Xeon E5-2690 2.9Ghz CPU processors on available cores and 128 gigs of RAM.

4.3.1. **Parameter Optimizations**

The goal of this work was to identify micro-well conditions that generated promising crystallization conditions that could then be manually replicated in a scaled-up environment by bench crystallographers. Thus, not only was the detection of crystals in micro-well experiments important, it was also vital to minimize the number of false positives that had to be evaluated. Since our algorithm was tuned to optimize against false negatives, it necessarily resulted in a non-trivial number of false positive images. Thus our crystallographer collaborators agreed to manually evaluate all machine-classified crystal images as a second screen instead of simply replicating all the micro-well conditions suggested by this algorithm. However, if this algorithm was too liberal in classifying images as containing crystals, then the manual crystallographers performing the second screening of machine-classified crystal images would again be dealing with an overwhelming number of images, the very problem that we began with and were attempting to fix. Furthermore, if the algorithm were too slow, then an insufficient number of images would have been processed, thus creating a new bottleneck in the high throughput pipeline. Therefore, significant efforts in the development of this work were put towards balancing the need for speed and high positive predictive value with avoiding an un-necessarily high negative predictive value.

Thus, different parameters in the implementation of the aforementioned algorithm were varied in order to tune this algorithm for the best possible performance. This process included experimentation with various features, including but not limited to additional first order image histogram features such as described in early work by our group [18], other types of wavelet expansions [43] and fractal dimension texture methods [44]. Ultimately, the first order image histogram features were found to be significantly correlated with features proposed in the current implementation of the algorithm, and thus contributed very little to the accuracy of the classifier other than to serve as potential sources of over-fitting. The Gabor wavelet feature expansions were the fastest in execution compared with other wavelet texture classification techniques, and classification performance did not differ significantly between all the tested wavelet-based texture features. The fractal dimensions calculated from the images ultimately decreased classification accuracy; it is likely that this is due to the proximity of some of our precipitate images with our crystal images.

Aside from feature selection, our algorithm is particularly sensitive to the window selection method. As, ultimately, what the classification algorithm examined were the windows, window size, window number and the clustering technique used in the window grouping were all tested in order to further optimize the algorithm. Table 4.2 describes results from experiments done during cluster technique optimization using a reduced sample of 1000 images. The number of clusters was held at 5, the cluster distance metric was defined using the Mahalanobis distance, as

described earlier in this work, and the number of representative windows from each cluster passed on to the qualifier was varied from 1 to 7 (resulting in a total of 5 to 35 windows). We can see that the calculation time per image increases linearly to the number of extracted images, as one would expect. However, we can also observe that the PPV and NPV peak at 25 windows (5 per cluster), beyond which the performances no longer improve. Thus, the final algorithm generated 5 images from each of the 5 clusters identified for further processing.

Table 4.2 *Results during window selection optimization*

# of Windows	Calculation Time per Image	PPV	NPV
10	2.756541121	0.875	0.9794
15	3.356735079	0.875	0.9794
20	3.84880265	0.878	0.9896
25	4.888238579	0.881	1
30	5.981462071	0.881	1
35	6.736029357	0.881	1

The f96 protein study is representative of the typical class distribution of micro-well images that might be obtained from one study. Even though there are a large number of images processed on

a daily basis, only approximately 1% of the images in any particular batch might contain crystals (Table 4.1).

4.3.2. **9-way Classifier**

We first began to explore the possibility of building a 9-way classifier that could correctly distinguish the 9 classes defined by the crystallographers in the f96 dataset. Since the NESG and SGPP datasets have not been fully sorted into these 9 classes, they were not included in this experiment. Approximately 80% of the 63,023 images were used for training, and the remaining 20% were used for testing. The classification technique and feature sets were identified above, except that 4 images were chosen from each of the 5 clusters of window images (instead of the 5 images chosen for the binary classifier as discussed above) and the random forest classifier was asked to sort the resulting images into 9 classes. Table 4.3 shows the number of images used in each category, the number of windows sampled from the images during classification, and finally the number of windows used during training and testing. It is important to note that each of the windows that were used to train for the class "crystal" was individually manually optimized. Specifically, we examined each of the 20 generated windows from crystal class images and manually discarded any windows that were non-indicative of crystals and thus likely to confuse the classifier.

Table 4.3 *Number of images, and windows used for training and testing*

Class	N° of images	N° of windows	N° windows for training	N° windows for testing
Clear	26,136	522,720	418,176	104,544
Phase separation & precipitate	388	7,760	6,208	1,552
Phase separation & crystal	26	520	416	104
Phase separation	5,374	107,480	85,984	21,496
Precipitate & crystal	63	1,260	1,008	252
Precipitate & skin	3,822	76,440	61,152	15,288
Precipitate	26,457	529,140	423,312	105,828
Skin	608	12,160	9,728	2,432
Crystal	149	2,980	2,384	596
Total	**63,023**	**1,260,460**	**1,008,368**	**252,092**

Table 4.4 *Percentage of ground truth labeling on first column and the result of the 9-way classifier in columns*

Truth	Clear	Phase separation & precipitate	Phase separation & crystal	Phase separation	Precipitate & crystal	Precipitate & skin	Precipitate	Skin	Crystal
Clear	**72.22**	6.11	7.26	8.42	2.23	2.41	0	0	1.35
Phase separation & precipitate	1	**32**	19.28	10.27	22.4	1.04	0.98	1.24	10.8
Phase separation & crystal	2.14	18.43	**49.37**	12.29	1.41	0.02	0.08	3.82	12.44
Phase separation	2.44	4.80	7.43	**47.9**	4	6.47	13.31	9	3.83
Precipitate & crystal	0	0	7.74	1.22	**59.9**	4.72	12.07	1.58	12.77
Precipitate & skin	1.33	7.18	5.92	2.31	14.27	**32.1**	15.34	17	3.99
Precipitate	2	9.84	9.32	4	8	14.22	**42.3**	6.77	2.78
Skin	6.37	14.32	9	12.07	19	9	8.84	**18.12**	2.44
Crystal	0.32	0	10.38	1.72	8.24	3.41	12.4	1.43	**62.1**

Analyzing Table 4.4, we can see that our classification method had difficulty distinguishing between all 10 classes together. However, given that certain classes had relatively few images available, and that several of the classes overlapped, this result was expected, albeit still disappointing.

4.3.3. **Binary Classifier**

Since the primary objective of our image classifier was to differentiate promising crystallization conditions from non-promising ones, we focused most of our efforts on the development of a binary classifier. Looking at the f96 database of images, we batched any image that contained crystals into a crystal group, and batched all other crystals into a non-crystal group. One can see that there is a relative scarcity of crystal images, accounting of only ~1% of the f96 proteins.

Table 4.5 *Mapping of the 9-way ground truth classes to the binary ground truth classes*

Class	Binary Class
Clear	Non-crystal
Phase separation & precipitate	Non-crystal
Phase separation & crystal	Crystal
Phase separation	Non-crystal
Precipitate & crystal	Crystal
Precipitate & skin	Non-crystal
Precipitate	Non-crystal
Skin	Non-crystal
Crystal	Crystal

In the previously described 9-way classifier, we see that the low number of images per class caused significant problems in the evaluation of the classification algorithm. To remedy this, we were able to obtain a large number of crystal images from the NESG and SGPP studies, as described in Table 4.1, to further supplement our datasets. Similar to the 9-way classifier, the crystal training windows were manually examined and any generated windows not displaying clear crystal characteristics were removed from the training dataset prior to the actual training of the algorithm.

Table 4.6 *Number of images, and windows used for training and testing in the binary classifier*

Class	N° of images	N° of training images	N° of testing images	N° windows for training
Non-Crystal	62,785	50,228	12,557	**1,004,560**
Crystal	16,425	13,140	3,285	109,523
Total	79,210	63,368	15,842	1,114,083

The results of the binary classifier are shown in Table 4.6 and Table 4.7. It is important to note that the training dataset for the binary classifier is no longer representative of the general distribution of crystal images in typical experiments (~1%).

Table 4.7 *Percentage of ground truth crystals and non-crystal images that were classified into the two labels by the two-class classifier*

Classifier Result	Ground Truth	
	Crystal	Non-Crystal
Crystal	**2,977**	682
Non-crystal	308	**11,875**

Table 4.8 *Sensitivity, specificity and other relevant statistics of the final two-way classifier*

Statistic	Value
Positive Predictive Value (PPV)	97.47%
Negative Predictive Value (NPV)	81.36%
Sensitivity	90.62%
Specificity	94.57%

Figure 4.9: *Examples of true positives, false positives and false negatives in the two-way classifier.*

True Positive	False Positive	False Negative
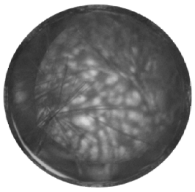		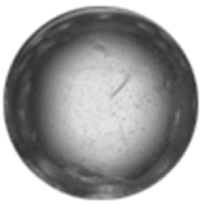
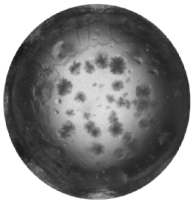	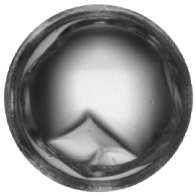	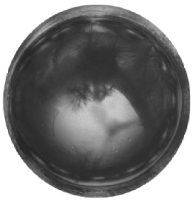

True Positive	False Positive	False Negative
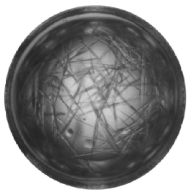	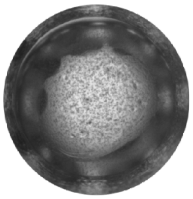	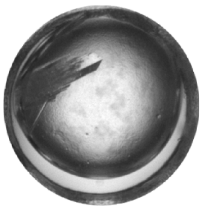
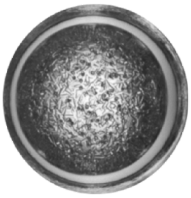	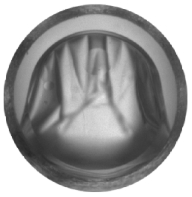	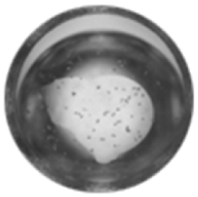
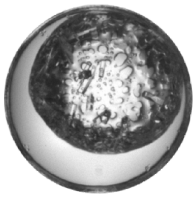	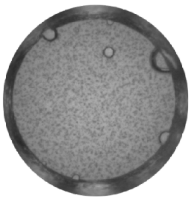	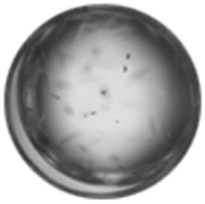
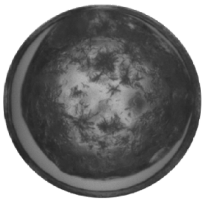		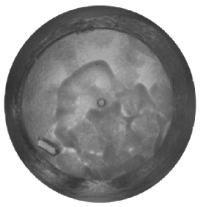

4.4. **CONCLUSIONS**

Studying the results of the nine-way classifier reveals that our classification method had significant difficulty distinguishing directly between all ten categories. However, we can see that the leading categories into which true crystal images were sorted included crystal (62.1%),

precipitate and crystal (12.77%), and phase separation and crystal (12.44%). As our method depended on using windows of each image, it is perhaps not surprising that images that were categorized by the 3 crystallographers as having multiple attributes were sometimes mis-categorized by the classifier as possessing only one of these attributes.

Previous work in this area by Cumbaa et al. using a similar but different training and testing set demonstrated much better nine-way classification results. As mentioned above, this may be due to the windowing method used here, compared to the global image features used by Cumbaa. Furthermore, their group did not choose to optimize processing time due to their use of the world community grid. This allowed them to use a larger dimension in the feature vector space for their classifier.

Ultimately, the two most important outcomes to crystallographers and to the project are a strong performance in the binary classification of crystal vs. non-crystal images, and the classifier's ability to process images at sufficient speed to be practical in a high-throughput protein crystallization pipeline. As can be seen in table 4.7, our two-way classifier was able to classify images with extremely high sensitivity and specificity. Our results for the binary classification of crystal vs. non-crystal demonstrate significantly better performance than previously published classifier results. Furthermore, our classifier was particularly tweaked for bias against missing crystal images (i.e. against false negatives), leading to extremely high sensitivity, though at the cost of introducing a higher false positive rate. It is important to note that ~1% of the images in

the production environment are expected to be true crystals. Thus, in a situation of one million images of which only 4,000 images are expected to contain crystals, a classifier with a false positive rate of 5% will generate, on average, 50,000 false positives. Our algorithm was therefore refined to specifically produce a false positive rate that is less than 5%, or producing no more than roughly 10 false positive classifications for every true classification.

Secondly, our algorithm was also significantly optimized so that the training process for 1,000 images took approximately 4 hours; once a model has been created, the classification of images takes approximately 7.4 seconds per image on a Centos 5 server running with one core i7 processor and 8 gigs of RAM (and only approximately 2.4 seconds per image on our computational cluster described earlier in this work). Thus, our algorithm would be able to fully process the batch of backlog images at the NESG using only three desktop servers over five years, and additional speedups can be realized with additional servers or the use of computational class servers such as the quadcore i7 Xeon processor used in the development of this algorithm. This is in contrast to Cumbaa et al.'s algorithm, which required the use of the world community grid, a scarce resource that cannot be used by crystallographers for any extended period of time.

There are of course limitations to the algorithm we present. While we were able to achieve the speedups of the algorithm with the use of a random forest classifier, there is currently no ability for online training. There are now groups working on RF implementation that allows for

real-time training [45], but those recent developments were not implemented in this work. It is also important to note that even if the online RF algorithm were implemented, our current pipeline would still not allow for complete automation of the process. Due to our implementation's window creation strategy, a significant number of windows in images containing crystals do not cover regions where crystals are evident. Thus, simply feeding all windows generated from an image containing crystals back into the training dataset would not only be computationally wasteful, it would in fact serve to decrease the accuracy of the algorithm.

As discussed earlier, the training dataset ultimately used for the current classifier was manually curated to ensure that only windows of interest existed in the training dataset. This was a painstaking ordeal involving the examination of over 300,000 images, a process that might make sense perhaps once a year, but would be highly unreasonable to recreate with any type of frequency. Thus, a better window selection method would have to be derived in order to allow the incorporation of any type of online training into the current classification framework.

Chapter 5: Chronic Obstructive Pulmonary Disease

Chronic Obstructive Pulmonary Disease (COPD) is an extremely common respiratory condition involving the airways, characterized by a limitation in airflow that is not fully reversible [46]. It is now evident that several different subtypes of COPD exist, and that these different subtypes have different prognoses and/or responses to treatment. It is however poorly understood how these subtypes arise and why they follow different clinical courses, and thus unsurprisingly COPD is currently the subject of intense research by the clinical community.

In the previous chapters, we introduced a multi-scale classification algorithm to help classify images in a high throughput environment. In the next 3 chapters, we will leverage multi-scale techniques to analyze radiographical images from COPD patients to help elucidate the pathophysiology of one subtype of COPD – emphysema. First, however, we will briefly present the illness in context.

5.1. EPIDEMIOLOGY

The currently accepted diagnostic criterion for COPD in the United States (GOLD guidelines) is spirometry showing airflow limitation (FEV_1/FVC ratio less than 0.70 or less than the lower limit of normal PLUS an FEV_1 less than 80 percent of predicted) that is incompletely reversible

with inhaled bronchodilator and cannot be otherwise explained. This last part of the diagnostic criteria is particularly important, as it excludes a large group of patients suffering from diseases such as asthma that are more acute in presentation, but ultimately not chronic in nature. In the general literature, COPD is sometimes confused with emphysema, one particular subtype of COPD that accounts for only around 1/3 of all COPD prevalence but attracts significant amounts of attention due to its strong association with smokers.

COPD is currently the third leading cause of death in the United States [47], and its mortality rate has doubled in the last two decades [48]. Since 1980, COPD deaths have increased in the US by more than 70%, up to a high of 141,000 persons in 2008, and the number continues to increase [47], [49]. Worldwide, COPD is the sixth leading cause of death [50] and is projected to rank fifth by 2020 [51]. Between 1980 and 2004, the US saw diagnoses of COPD by physicians increase from 7 million to 12 million [52], with an estimated further 12 million cases remaining undiagnosed [49]. In 2006 there were 672,000 hospitalizations linked to COPD [47], with total and direct costs of $50 billion and $30 billion respectively in 2010 [49].

The prevalence of COPD increases with age and smoking status, with a five-fold increased risk for those aged over 65 years old compared to patients aged less than 40 years, and a similar five-fold increased risk for those who smoke vs. those who don't smoke [53]. COPD currently affects twice as many males as females, but the difference is decreasing as the gender gap between smokers and non-smokers closes in the developing world.

Due to its high prevalence and rates of morbidity, mortality, and readmission, COPD has garnered increasing interest from both healthcare professionals and policy makers. However, COPD continues to present numerous ongoing, contentious issues, including even its definition and staging.

5.2. **BREATHHOLD TERMINOLOGY**

A number of terms have been standardized in the clinical community to describe various inspiratory / expiratory statuses of a person's breath hold. This terminology is not only used by the research community, but the measurements of these volumes is used in the diagnosis and maintenance of patients with respiratory disorders.

Table 5.1 *Explanation of terminology typically used by the pulmonary research community to describe various states during the respiratory cycle*

Total Lung Capacity	**TLC**	The maximum volume of the lungs. This is the respiratory state in which CT scans for COPD patients are normally done.
Tidal Volume	**TV**	The volume of air exchanged from the lungs during normal breathing.
Residual Volume	**RV**	The volume of air remaining in the lungs after maximal exhalation.
Expiratory Reserve Volume	**ERV**	The additional volume that can be exhaled if a person tries to blow out all possible air from his / her lungs.
Inspiratory Reserve Volume	**IRV**	The additional volume that can be inhaled if a

		person tries to take in all possible air from his / her lungs.
Inspiratory Capacity	**IC**	The sum of IRV and TV.
Functional Residual Capacity	**FRC**	The volume of air left in lungs after normal exhalation. This is the respiratory state in which cardiac MRI / blood flow related MRI is done.
	FEV_1	Forced expiratory volume in 1 second
	PEF	Peak Expiratory Flow, or the highest flow rate measured by flow meter during exhalation.

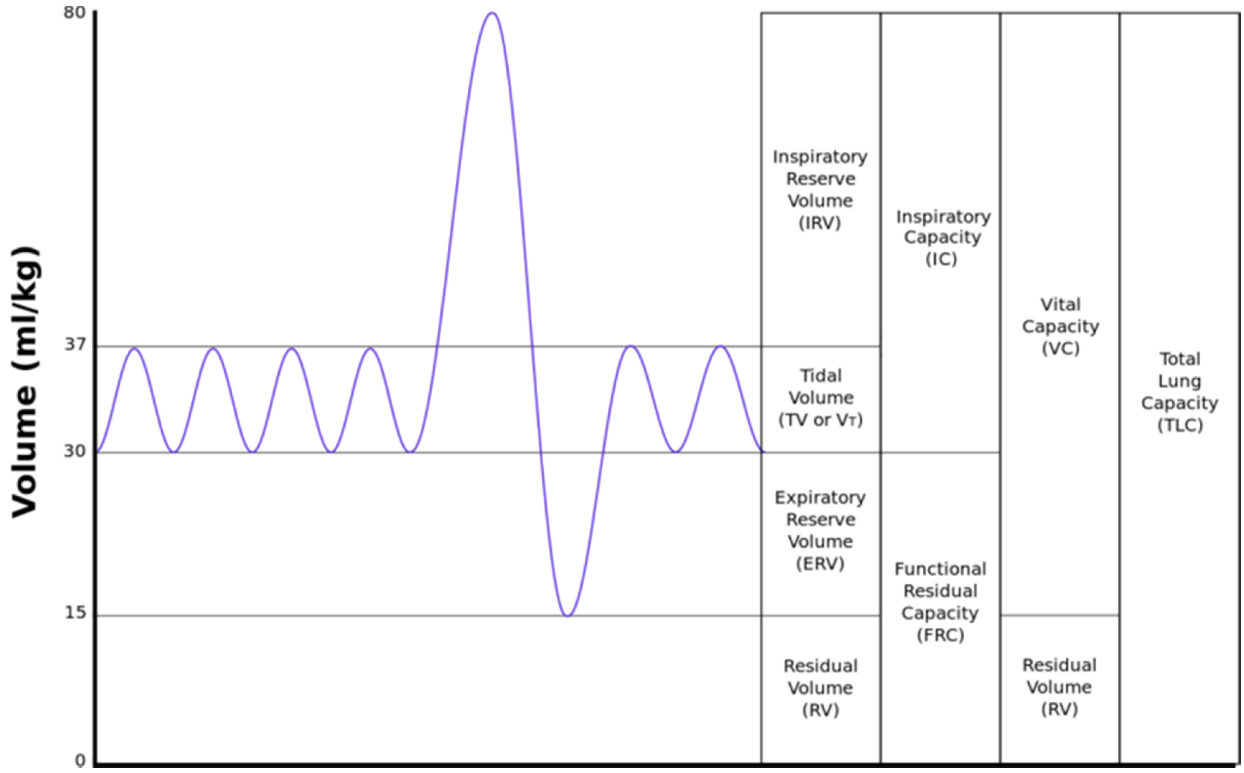

Figure 5.1: *Terminology used when discussing Lung Volumes.*

5.3. **PATHOPHYSIOLOGY OF COPD**

Airflow limitation, specifically expiratory airflow limitation, is the pathophysiological hallmark of COPD. The exact etiology of the expiratory flow limitation is still unknown, but there are two general mechanisms that are known to contribute to the pathology. First, any type of physical obstruction in the airways can contribute to airflow limitations. The obstructions of large airways such as the trachea can produce airflow limitations, but similarly, the obstructions of many smaller airways can produce an equal if not greater airflow limitation, as will be discussed later. Secondly, alterations of the intrinsic mechanical properties of the lung can also produce expiratory airflow limitations. While the process of normal respiratory exchange requires the active use of the respiratory muscles, normal expiration is largely a passive process that depends on the elastic recoil of the lung tissue itself to create the necessary positive pressure in the airways to expel the de-oxygenated air. If this elastic recoil is compromised for any reason, then it becomes necessary to use the respiratory muscles in order to generate the force necessary to create that positive pressure.

Regardless of the reason for the increased expiratory resistance, the body initially compensates by using respiratory muscles to help increase the expiratory pressures and thus maintain a positive airway pressure. Analysis of flow-volume loops of COPD patients by Mergoni et al. showed that as COPD progresses, the flow generated during expiration of a tidal volume becomes very close to the flow generated during forced maximal expiration [54]. When the body

no longer has any reserve of respiratory flow, higher tidal volumes are needed in order to reach the same expiratory flow, and hyperinflation of the lung then occurs. Thus COPD patients with expiratory airflow resistance will then have increased functional residual capacity (FRC).

Hyperinflation of the lungs causes several consequences: 1) The increased expiratory resistance causes the body to start using respiratory muscles for both inspiration and expiration, resulting in much easier fatigue of the respiratory muscles, and in particular, of the diaphragm. This, together with an increased lung volume that pushes on the diaphragm, is thought to contribute to the flat diaphragm frequently seen in CT images of COPD patients. 2) The increased expiratory resistance and increased FRC cause hyperinflation of the alveolar and thus constrict the proximal alveolar vessels [54]. This is known to lead to an increase in pulmonary vascular resistance and also contribute to hypoxic vasoconstriction. This increase in pulmonary vascular resistance leads to an increased right ventricular preload; eventually, right-sided hypertrophy ensues. 3) At lung volumes that are close to TLC, the lung compliance is physiologically reduced, and thus requires an even larger elastic force to generate the same inspiratory volume. This is particularly problematic in COPD patients who already have elastic recoil deficits. One can observe that each of the 3 factors above worsen with further progression of COPD, creating a positive feedback system that makes it difficult for the body to compensate for these physiological changes. In severe cases, cor pulmonale, or right ventricular heart failure secondary to pulmonary artery hypertension, can occur [55], [56].

Given the above, one can conclude that the development of increased expiratory airway resistance that cannot be compensated by the body will eventually lead to the development and progression of COPD. However, the etiology of COPD, or the cause(s) of the initial insult(s) that cause chronic increased expiratory airway resistance, is less well understood.

5.4. **DIAGNOSIS OF SUBTYPES OF COPD**

Several different morphological "phenotypes" have been observed in patients currently considered to have COPD, and it is now clear that these different morphologies have different prognoses and/or responses to treatment.

Clinically, four major subtypes of COPD have been recognized: 1) Chronic Bronchitis, 2) Emphysema, 3) Asthma, and 4) Alpha-1 Antitrypsin Deficiency. Chronic Bronchitis is defined by constant inflammation of the lining of the airways, causing increased secretion, scarring and resulting restriction of airflow. Emphysema is defined by abnormal and permanent enlargement of the airspaces themselves, leading to air trapping that effectively becomes a restriction of airflow. Patients with chronic asthma can eventually develop COPD as a result of processes similar to those of chronic bronchitis as described above. Alpha-1-Antitrypsin Deficiency is a genetic disorder causing defective production of an anti-protease named alpha-1-antitrypsin (AAT), ultimately resulting in destruction of the alveolar walls causing COPD, following an etiology similar to pulmonary emphysema.

Though the destruction of pulmonary tissue due to AAT deficiency is well understood, the etiologies of chronic bronchitis, emphysema, and asthma remain obscure. Indeed, these subtypes of COPD frequently coexist in the same patient, making these categorizations artificial and clinically ineffectual.

Aside from AAT deficiency, the different sub-types of COPD are distinguished based on history and clinical symptoms alone and no reliable imaging or biomarker exists to distinguish between them. Thus, clinical symptoms are used to attempt to differentiate the subtypes. Chronic bronchitis is defined by a cough productive of sputum for over three months' duration during two consecutive years [57]. Dyspnea and airway obstruction, often with an element of reversibility, are intermittently to continuously present. Pulmonary emphysema, on the other hand, is primarily caused by a pathology of the respiratory alveoli and not the airways, and thus is rarely accompanied by significant productive cough. Imaging modalities hold an important role in the monitoring of progression of pulmonary emphysema. Additionally, recent research has suggested that high resolution computed tomography (CT) is an especially reliable tool for demonstrating the pathology of emphysema, even for subtle changes in secondary pulmonary lobules [58].

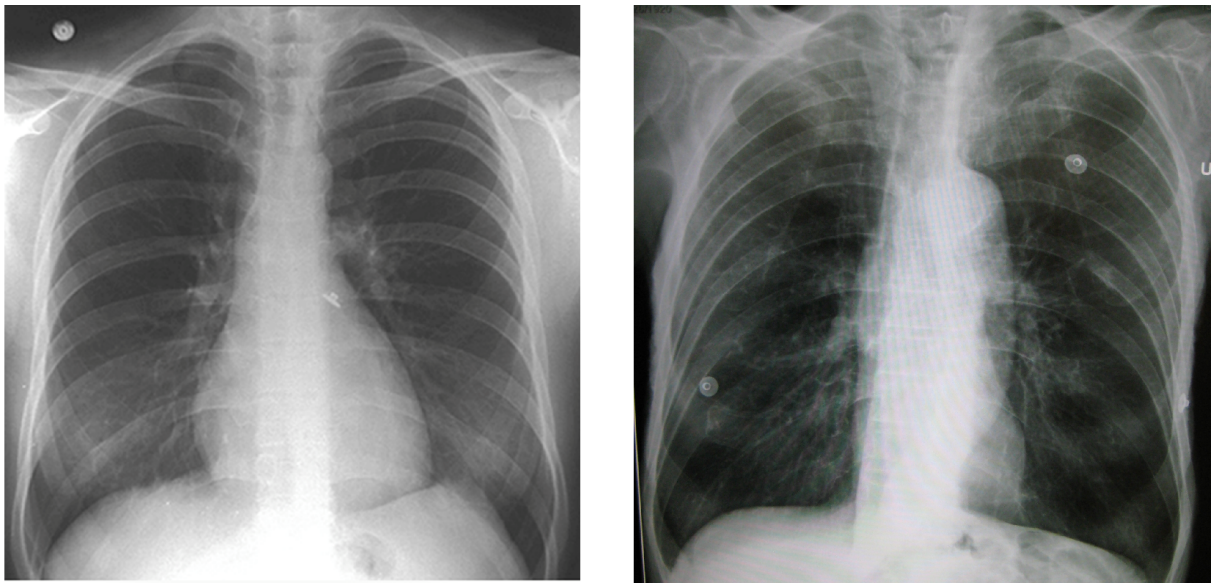

Figure 5.2: *Left: Normal Chest Radiograph from Yale Radiology. Right: Typical Late stage COPD chest radiograph with lung hyperexpansion and narrowed heart, flat diaphragm, and prominent hilar vascular shadows due to pulmonary hypertension and cor pulmonale.*

5.4.1. **Pathophysiology of Chronic Bronchitis**

The principal pathologic features in Chronic Bronchitis are inflammation of the airways and hypertrophy and hyperplasia of airway mucous glands. The airway mucosa is variably infiltrated with inflammatory cells, including polymorphonuclear leukocytes and lymphocytes. Due to the chronic inflammation, the normal ciliated pseudostratified columnar epithelium is frequently replaced by patchy squamous metaplasia. This results in severely diminished mucociliary clearance function. This combined with the mucus hypersecretion secondary from the hypertrophy and hyperplasia of large airway mucous glands, further contributes to luminal narrowing. Mucus impaction and luminal obstruction of smaller airways are often seen. The

result of these combined changes is chronic airway obstruction and impaired clearance of airway secretions.

The non-uniform airway obstruction of chronic bronchitis leads to areas of both high and low V/Q ratios. The regions of low V/Q ratios are largely responsible for the more significant resting hypoxemia seen in chronic bronchitis compared with emphysema. True shunting, that is, perfusion with no ventilation, or V/Q of 0, is unusual in chronic bronchitis.

5.4.2. **Pathophysiology of Pulmonary Emphysema**

In contrast to chronic bronchitis, which is a disease of the airways, emphysema is a disease of the surrounding lung parenchyma. The principal pathology in emphysema is the progressive destruction of alveoli tissue and permanent enlargement of the airspaces themselves, thought to be due to an inability of proteolytic activity by protease inhibitors (e.g., AAT) to compensate for local oxidant injury, whether endogenous or exogenous (e.g. cigarette smoke). If present, airway inflammatory changes are minimal. Alveolar capillaries are also lost, which can result in decreased diffusing capacity and progressive hypoxemia, particularly when exercising. However, alveolar destruction is not uniform in all cases of emphysema. Moreover, while emphysema is one of four subsets of COPD, it has itself been classified into three major subtypes (centriacinar, panacinar, and paraseptal), based on the disease distribution within secondary pulmonary lobules [59], [60]. Centriacinar emphysema, most commonly associated with prolonged smoking,

presents a pattern of destruction that is focused in the center of the terminal respiratory unit; the respiratory bronchioles and alveolar ducts are comparatively rarely affected. Panacinar emphysema involves destruction of the terminal respiratory unit globally, with diffuse airspace distention. This pattern is typically, although not uniquely, seen in cases of α1-protease inhibitor deficiency. Paraseptal emphysema is characterized by an enlarged airspace at the periphery of acini, and unlike panlobular emphysema, the lesions are usually limited in extent, occurring most commonly along the dorsal surfaces of the upper lung.

All three types of emphysema cause fairly similar physiological consequences which can be traced to three important pathologies: 1) destruction of terminal respiratory units, 2) destruction of the alveolar capillary bed, and 3) loss of the supporting structures of the lung, including elastic connective tissue. The loss of elastic connective tissue further results in a lung with diminished elastic recoil and increased compliance, leading to a loss of the normal support of noncartilaginous airways. Thus the noncartilaginous airways prematurely collapse during expiration, resulting in the characteristic obstructive symptoms and physiologic findings.

As discussed in the earlier section on the pathophysiology of COPD, severe and uncompensated increases in expiratory pressures can eventually lead to compression of the pulmonary vasculature and thus leading to pulmonary HTN and eventually right heart failure. However, evidence currently suggests that there are changes in hemodynamics even for mild and moderate COPD patients that cannot be explained purely by pressure changes from hyperinflation and

hypoxic shunting [61]. Indeed, Barr et al. found a significant association between percent emphysema and left ventricular measures in the entire spectrum of mild to severe COPD patients. Thus, one possible mechanism that could account for this result is that there are separate and independent direct injuries to the pulmonary capillary bed beyond simple compressive effects. If this is the case, then such an injury to the pulmonary capillary bed could also be an independent contributing factor to the progression of increased expiratory pressure and thus further progression of COPD. In order to further examine this potential mechanism, chapter 7 is dedicated to the development of registration techniques for CT and perfusion MRI images that would allow us to more directly observe and correlate regions of decreased pulmonary vascular perfusion with regions of potential COPD progression.

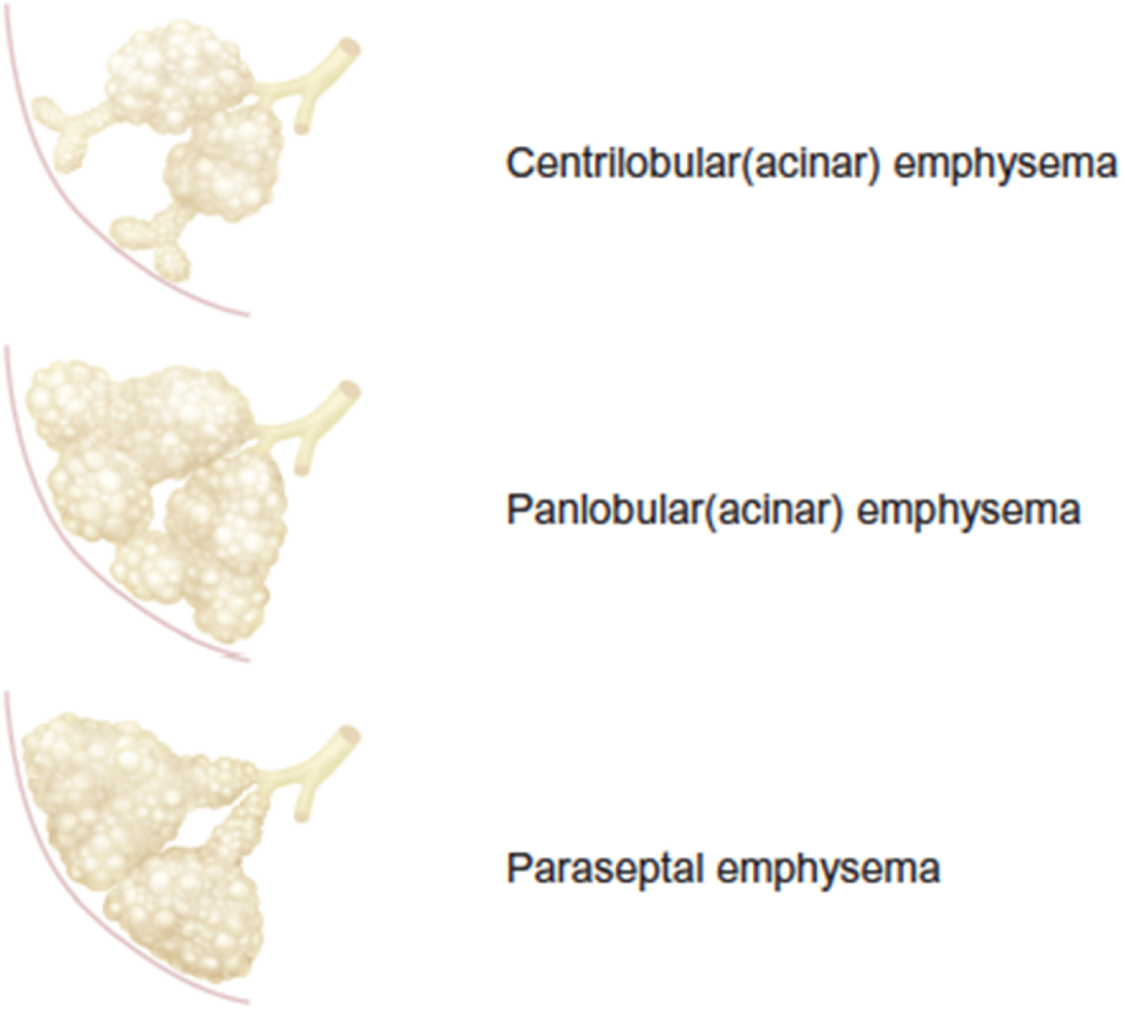

Figure 5.3: *Different subtypes of emphysema – Centrilobular, Panlobular, and Paraseptal emphysema are shown here. The above figure is adapted from [58].*

5.5. **CLINICAL PRESENTATION**

Most patients with chronic obstructive pulmonary disease (COPD) seek medical attention late in the course of their disease due to the slowly progressive nature of the disease. When they do present, they typically present with a combination of cough productive of sputum and progressively worsening exercise tolerance. Dyspnea is the most significant symptom and the

most common symptom leading to ER visits/admissions, but it usually does not occur until late in the course of the disease. Systemic manifestations such as pulmonary hypertension, cor pulmonale, and left-sided heart failure is also common in the later stages of the disease.

5.6. **TREATMENT**

Currently, no treatment (pharmacological or otherwise) for COPD other than lung transplantation and lung reduction surgery has been shown to significantly improve lung function or decrease mortality [62], [63]. The development of effective treatments has been hindered, in part, by a lack of clearly defined, patho-physiologically distinct sub-phenotypes. Thus, the goal of COPD management is primarily to improve a patient's quality of life by alleviating symptoms and decreasing the number of exacerbations and hospitalizations. Since there are no treatments for COPD, most of the medication-based therapies are directed towards any reversible causes of airflow limitation in order to remove any reversible component of dyspnea. This is done in the hopes that this will at least improve the patient's functional status. The 4 primary causes for reversible airflow limitation in COPD patients are:

- Bronchial smooth muscle contraction;
- Bronchial mucosal congestion and edema;
- Airway inflammation;
- Increased airway secretions.

Thus it is not surprising that medical therapies are only somewhat effective in stages 1 and 2 of COPD, where exacerbations are more likely a result of reversible airflow exacerbations on top of mild but chronic irreversible airflow obstruction. As the disease progresses, the chronic airflow obstruction worsens and thus contributes in a larger way to the airflow limitations, making medical therapies much less effective. Indeed, medical therapies have not shown to be effective in slowing the progression in patients with Stage 3 and 4 COPD either.

Surgery may be used to remove parts of the diseased lung in some patients with emphysema, but this only benefits a minority of emphysema patients and is ineffective in cases of chronic bronchitis. While in severe cases, a lung transplant is an option, donors are limited and the procedure is very costly.

5.7. QUANTIFICATION OF PULMONARY EMPHYSEMA

As mentioned previously, radiographs are frequently obtained for patients suspected of having COPD. However, most of the signs of COPD observable on radiographs, such as focal absence of pulmonary vessels, lung hyper-expansion, and flattening of the diaphragm, are only present in late stage COPD, and are therefore impractical for the early detection of COPD. What's more, the sensitivity of radiographic COPD findings is as low as 40% [64]. The low sensitivity and specificity of radiographic images have led researchers to explore using CT to further

characterize COPD, and in particular the pulmonary emphysema subtypes. Early studies demonstrated that emphysematous pathology can be visible in CT sections with a thickness of 5-8 mm, and is even more evident in the high-resolution CT scans now available with reconstructed thicknesses of 1 to 2 mm.

Chapter 6: Computed Tomography (CT) of Pulmonary Emphysema

CT images are generated by the differential attenuation of various materials as X-rays pass through the object of interest. Thus, areas where there is least material (i.e. air) typically have the strongest signal and the lowest X-ray attenuation. The unit of measurement in CT is the Hounsfield unit, a unit of measurement that is linearly related to the X-ray attenuation. Thus, an area consisting of only air will have the lowest possible Hounsfield unit, and an area consisting of extremely dense material (for example, bone) will have very high Hounsfield units.

It is therefore possible to use Hounsfield units to distinguish normal lung tissue from lung regions with lower tissue density. The accuracy of high-resolution CT in assessing the presence and extent of emphysema has been documented in numerous studies [64], [65], [66], [67]. In one study based on necropsy specimens, investigators demonstrated that CT sensitivity was sufficiently high to identify even mild centrilobular emphysematous case [66]. The correlation between the in-vitro CT emphysema score done via manual visual scoring and the pathological grade of emphysema was excellent ($r = 0.91$).

Substance	HU
Air	-1000
Lung	-700
Soft Tissue	-300 to -100
Fat	-84
Water	0
CSF	15
Blood	+30 to +45
Muscle	+40
Bone	+700 (cancellous bone) to +3000 dense bone)

Figure 6.1: *Typical Hounsfield units used for a generic CT scanner.*

The inherent limitations of subjective visual scoring and the consistent ranges of Hounsfield units in CT datasets between different patients have fostered significant interest in the use of computer algorithms to quantify pulmonary emphysema [68]. In 1988, Müller et al. compared a single 1cm-thick CT section of lung tissue acquired during full inspiration (after injection of contrast material) with the corresponding macroscopic section of the fixed lung cut in the same plane as the CT section [69]. The goal was to correlate emphysematous regions on pathology with Hounsfield unit ranges on CT that might uniquely identify those regions. The highest correlation was observed with attenuation values lower than -910 HU, and consequently, this threshold was recommended for the identification of emphysema. Though the emphysema research community generally accepted this threshold in the 1990s, for a number of reasons significant uncertainty remained regarding the accuracy of the resulting analysis. Firstly, the

baseline of -910HU was derived under the influence of contrast material. Secondly, though the pathology scores and -910HU were statistically significant, they did not guarantee that the percentage areas obtained by CT quantifications were equal to the percentage area occupied by emphysema on the pathological specimen.

In an attempt to provide a more precise attenuation threshold for the recognition of emphysema, Gevenois et al. compared CT data acquired during full inspiration from 63 patients about to undergo lung resections with the macroscopic extent of emphysema measured on horizontal paper-mounted whole lung sections post resection [70]. On thin CT sections obtained from the lung apices to the bases with 1-cm intervals, the authors calculated the relative area of lung tissue occupied by HU values lower than various thresholds ranging from -900 HU to -970 HU. They showed that the only threshold for which there was no statistically significant difference between the distribution of the CT measurements and the distribution of macroscopic measurements was -950 HU. Thresholds lower than -950 HU underestimated the amount of emphysema, while thresholds above -950 HU overestimated the amount of emphysema [70].

More recent work using airspace wall surface area per unit volume of lung tissue (AWVU) to quantify microscopic pulmonary emphysema also showed -950HU on CT as having the largest correlation with AWVU indices. Thus, both the macroscopic and the microscopic study suggest that -950 HU is a valuable parameter for quantifying emphysema on CT during full inspiration [69].

However, not all CTs are taken at full inspiration, and the possible role of CT obtained after deep expiration in the assessment of emphysema was first suggested by Knudson et al. [71]. Gevenois et al. found two different thresholds, respectively validated by comparisons against macroscopy (-910 HU) and microscopy (-820 HU) that were quite different from the threshold found valid for CT images obtained at full inspiration (-950 HU) [69].

6.1. **PULMONARY MICROVASCULAR PERFUSION**

Emphysema is characterized by progressive airflow obstruction secondary to the destruction of alveolar tissue, but in the late stages of disease is also associated with systemic vascular symptoms such as cor pulmonale and pulmonary hypertension. However, less is known about the interplay of pulmonary disease with its associated vascular response in the early stages of COPD. Recent evidence suggests that endothelial dysfunction and alterations in pulmonary vascular response occur early in COPD and may represent an important vascular pathway in the development of smoking-associated emphysema [72]-[74]. In response to chronic oxidate stress, an overwhelming inflammatory response is triggered, leading to obstruction of alveoli and regional hypoxia. The lung's normal response to alveolar hypoxia is hypoxic pulmonary vasoconstriction (HPV), which results in the shunting of blood toward better-ventilated lung regions for improved oxygenation. In mice, sheep [75], and humans [76], [77], imaging demonstrates that HPV is locally blocked in the presence of inflammation [78]. Thus, various

groups have suggested that one possible mechanism of emphysema progression is an impaired HPV response that perpetuates an otherwise self-limited inflammatory response, resulting in chronic inflammation, tissue destruction, and worsening hypoxemia and emphysema.

In order to test this hypothesis, Alford et al. used a dynamic 4D ECG-gated MDCT perfusion sequence to measure the differences in quantitative regional perfusion parameters such as pulmonary blood flow and mean transit time between people who have never smoked and smokers with radiographic evidence of centrilobular emphysema [79]. The study found no significant differences in PBF and MTT between the two groups, but did note differences in the heterogeneity of the PBF between the two groups [79]. This study was the first we know of that attempted to investigate pulmonary microvascular perfusion using any imaging modality in emphysema patients. Unfortunately, CT perfusion imaging is not well validated and no significant work was done after this study until the maturation of dynamic contrast enhanced MRI.

Several approaches for quantifying pulmonary microvascular perfusion by MRI have been proposed [80], [81]. They are based on the indicator dilution theory [82], which applies to blood pool contrast agents. In 2013, Hueper et al. used dynamic contrast-enhanced MRI to evaluate the differences in pulmonary microvascular perfusion that were calculated from the signal intensity-time curves in the lung parenchyma of COPD vs. non-COPD patients [83]. They noted substantially reduced perfusions in cases of mild, moderate and severe COPD as compared to

patients without COPD. Furthermore, the reduction in perfusion was linearly related to the percentage of emphysema on CT scans. The findings suggest marked microvascular damage in early COPD.

6.2. **CONCLUSION**

Despite the many substantial research efforts on the role of CT and MRI in pulmonary emphysema, an important series of issues remains to be investigated. The proper selection of an HU threshold for emphysema quantification outside of 1 mm CT slice thickness and total breath hold continues to be unexplored. Furthermore, integrating MRI perfusion findings with regions of emphysema as identified in CT has yet to be achieved. This would allow for proper investigation of the pathogenesis of vascular abnormalities in the pathophysiology of emphysema progression.

Chapter 7: Registration of CT and MRI Lung Volumes

Changes in pulmonary vasculature and perfusion are thought to occur in severe COPD secondary to hypoxemia and destruction of lung parenchyma [55], [56]. Various ex vivo studies have confirmed the existence of inflammatory reactions against both pulmonary vasculatures [84] and lung alveolar epithelial and endothelial cells [85] in mild and severe COPD. However, it is not well understood whether the destruction of lung parenchyma and the changes in pulmonary vascular perfusion are two different pathological processes in the progression of COPD or whether one is the predisposing mechanism of the other. This difficulty in correlating the two processes results from the inability prior to recent imaging developments in both perfusion MRI [80] and CT [86] to quantitatively evaluate what degree of parenchyma damage or changes in pulmonary vascular perfusion existed in humans before autopsy. With both modalities now clinically validated, we propose a registration algorithm between quantitative CT and microvascular perfusion MRI that will allow for the correlation of emphysematous regions as identified on CT, with pixel resolution perfusion changes as identified on microvascular pulmonary perfusion MRI.

Several approaches to quantify pulmonary microvascular perfusion by MRI exist, though all are based on indicator dilution theory as applied to the pooling of contrast agents [83]. Due to

concerns related to reduced pulmonary perfusion during breath holds at Total Lung Capacity (TLC) versus Functional Residual Capacity (FRC) [87], perfusion MRIs are typically done at FRC, or the end of passive expiration.

The identification of emphysematous regions on CT is a more straightforward task and is made through a hard threshold of Hounsfield units. There is now widespread consensus in the COPD research community that a value of -950 Hounsfield units (HU) has strong correlations with other clinical indicators of COPD in severe COPD patients [88]. However, in studies of early/mild COPD, the threshold of -910 Hounsfield is sometimes used in order to more sensitively capture potentially emphysematous regions on CT. Regardless of the threshold used, emphysema thresholding is optimally performed on CT images captured and analyzed during TLC breath holds due to both higher replicability by patients across multiple visits and decreased effects of confounding factors such as obesity and physiological gas trapping [89]. As COPD is a chronic disease, many patients have had CTs done at different lung inflations over many years and even decades. Thus, work has been done to register CT images at different time points [5], [90]-[92], and with different patient breath hold statuses [6], [7], [93], [93].

To register TLC-CT images with FRC-MRI perfusion images, one could imagine carrying out a registration between TLC-CT and FRC-CT images, and then performing a simple rigid transformation between FRC-CT images and FRC-MRI perfusion images. Unfortunately, due to the patients' chronic lung disease and the breath hold time intervals needed for FRC-MRI scans,

the FRC-MRI lung volumes are often differentially inflated when compared with FRC-CT volumes.

We propose a method of registering lung CT volumes with MRI volumes, taking into account the different breath hold statuses that could be acquired during CT and MRI, as well as the inherent differences in the imaging characteristics of the two modalities. The registration is ultimately being performed to localize functional findings from MRI imaging (pulmonary microvascular blood flow) to changes in lung structure on CT lung imaging (low attenuation areas).

Our registration strategy involves 1) segmentation of CT volumes, 2) segmentation of MRI volumes, 3) registration of TLC-CT volume to FRC-CT volume for each individual patient and creation of pixel density volume mapping, 4) registration of FRC-CT volume to FRC-MRI volume for each individual patient and creation of pixel density volume mapping, 5) mapping of emphysematous regions from TLC-CT to FRC-MRI volume through the composite transformation functions described above, and 6) morphological processing of resulting masks from step 5 to improve results.

Validation of the algorithm is then performed by selecting the mid-coronal slice of corresponding MRI and CT volumes and comparing the relative areas of the bronchial tree.

7.1. **METHODS**

7.1.1. **Registration Transformation Model**

Our processed algorithm requires two separate registration steps: 1) the registration of CT images of different breath hold statuses to each other and 2) the registration of CT and MRI images of nominally identical breath hold statuses. Due to the complexity of both registration steps, a simple rigid or affine transformation is unlikely to produce satisfactory results for either registration. Thus, we adapted a B-spline free-form deformation (FFD) -based registration algorithm for both steps.

To define a spline-based FFD, the image volume domain is defined as $\Omega = \{(x, y, z) | 0 \leq x < X, 0 \leq y < Y, 0 \leq z < Z\}$. We can then denote a $n_x \times n_y \times n_z$ mesh of control points $\emptyset_{i,j,k}$ with spacings between the control grids in the x, y, and z directions denoted by δ_x, δ_y, and δ_z respectively. Then, the FFD can be written as the 3D tensor product of the 1-D cubic B-splines [94]:

$$T_{local}(x, y, z) = \sum_{i=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_1(u) B_m(v) B_n(w) \emptyset_{i+l, j+m, k+n}$$

where $i = \left[{}^{x}/_{\delta_x} \right] - 1, j = \left[{}^{y}/_{\delta_y} \right] - 1, k = \left[{}^{z}/_{\delta_y} \right] - 1, u = {}^{x}/_{\delta_x} - (i + 1), v = {}^{y}/_{\delta_\gamma} - j + 1,$ $w = {}^{z}/_{\delta_z} - (k + 1)$ and where B_l *(u)* represents the *l*th basis function of the B-spline [95], [9].

$$B_0(u) = {(1-u)^3}/{6} \quad B_1(u) = {(3u^3 - 6u^2 + 4)}/{6};$$

$$B_2(u) = {(-3u^{3+}3u^2 + 3u + 1)}/{6} \quad B_2(u) = {u^3}/{6}$$

where $0 \leq u < 1$. B-splines are locally controlled, which allows them to be computationally efficient even for large numbers of control points. Cubic B-splines in particular have limited support, which allows the transformation to affect only the local neighborhood of a control point.

In registration processes where both larger displacements, as during breath hold changes, and smaller displacements, as those due to noise and minor displacements of the bronchial tree, can be expected, the registration algorithm must be able to cope with both changes. Since cubic B-splines have limited support, neither an exclusively fine nor an exclusively coarse spacing between control grids will capture both scales of displacements. Thus, a multilevel B-spline adaption was adopted [6], [7], [95]. In the multilevel B-spline algorithm, a hierarchy of control grids with progressively finer spacing between control grids is used, $\emptyset_0, \emptyset_{1,\dots,}\emptyset_n$. The registration process is then iteratively performed using each control grid, resulting in a sequence of transformations $T_0, T_{1,\dots,}T_n$. The final transformation is then defined as the composite of all the operations $T_{final} = T_n{\circ}T_{n-1}{\circ} \dots {\circ}T_0$.

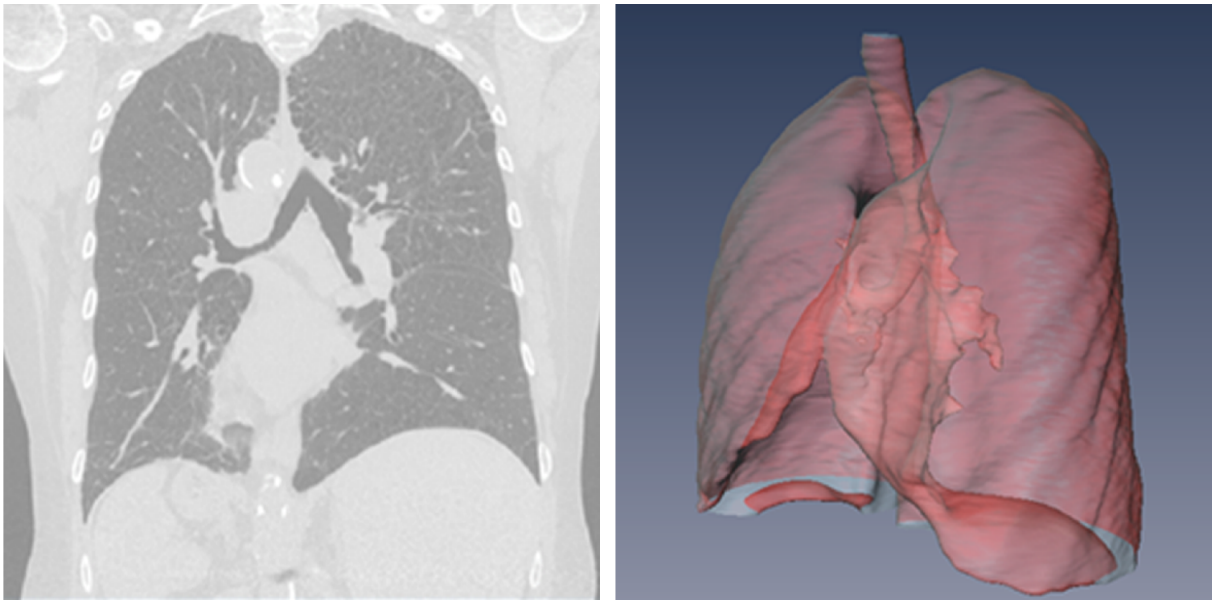

Figure 7.1: *On the left is a sample coronal slice of the TLC-CT volume analyzed. On the right is the segmented lung and brunchial airway.*

7.1.2. **CT Registration Cost Function**

A typical cost function for image registration is to use the sum of squared differences (SSD),

$$SSD = \frac{1}{n}\sqrt{\sum \Big(I(t_o) - T\big(I(t)\big)\Big)}.$$

The implicit assumption with SSD as a cost function for registration is that images differ only by Gaussian noise after registration. When lung parenchyma are differentially inflated during different breath hold statuses, the voxel intensities for the same region of the lung will change accordingly. Specifically, we can separate the Hounsfield units of the lung CT as a linear combination of parenchyma plus blood and air. We can assume that the lung parenchyma plus blood is not significantly affected during changes in breath hold, but the volume of air is

naturally affected [6], [7]. Thus, we wish to adapt SSD to register only the tissue and perfusion component of the voxel intensities between CT images of different breath hold statuses. Using 55 HU as the voxel intensity of parenchyma plus blood, and -1000 HU as the voxel intensity of air, the volume of tissue and air can be estimated as:

$$V_{tissue} = v(x)\frac{I(x)+1000}{1000+55} = v(x)\tilde{I}(x)$$

$$V_{air} = v(x)\frac{55-I(x)}{1000+55} \ ,$$

with $v(x)$ as the volume of a voxel. Using this technique, we can modify the CT registration cost function as the sum of squared differences of only the local tissue volume difference:

$$E(\emptyset) = \sum_{x\in\Omega}^{3}[V_r(x) - V_f\big(T(\phi)\big)]^2 \ ,$$

with $V_r(x)$ and $V_f(x)$ denoting the tissue volumes in a voxel in the reference image and the floating images respectively.

7.1.3. **MRI Registration Cost Function**

As described above, SSD as a cost function for registration makes the implicit assumption that images differ only by Gaussian noise. It is therefore particularly unsuitable for intra-modality registration such as FRC-CT to FRC-MRI. We thus similarly adapted our B-spline registration described in section A with a cost function that does not depend on a linear intensity relationship

between images intensities: Mutual Information.

$$MI(E, I_F, I_{m)} = \sum_{m \in L_m} \sum_{f \in L_f} p(f, m; E) \times log_2 \left(\frac{p(f,m;E)}{p_f(f) p_m(m;E)} \right),$$

where L_F and L_m are sets of regularly spaced intensity bin centers, p is the discrete joint probability, and P^F and P^M are the marginal discrete probabilities of the fixed and the moving image, obtained by summing p over the indices m and f, respectively [96]. The joint probabilities can then be estimated from the B-spline Parzen windows

$$p(f, m; E) = \frac{1}{|I_F|} \sum_{x_i \in I_F} w_f \left({}^{f}\!/_{\delta_F} I_F(x_i) \right)$$

$$x w_M \left({}^{m}\!/_{\delta_M} - {}^{I_m(x_i + x_\mu(x_i))}\!/_{\delta_m} \right),$$

where x_i denotes the spatial coordinates of voxel I in the fixed image I_F, u_μ is the B-spline deformation field, and w_m and w_F represent the moving and fixed Parzen windows respectively. Further details of the implementation can be found in [97].

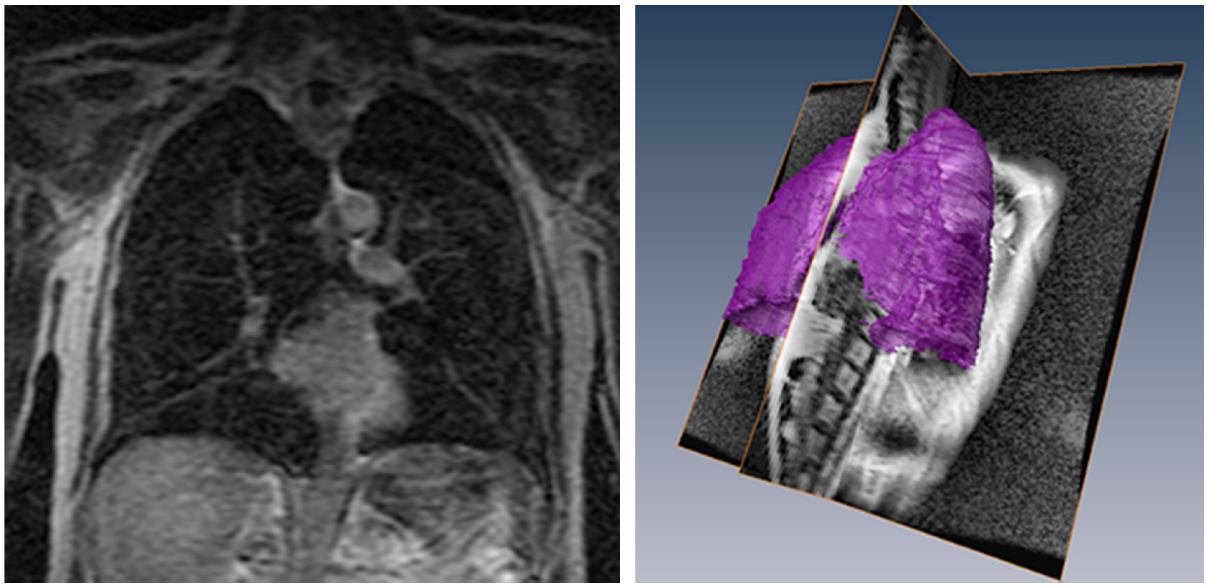

Figure 7.2: *On the left is a sample sagittal slice of the MRI image analyzed. On the right is a surface reconstruction of the manual segmentation of the MRI lung volume super imposed on the original MRI volume [98].*

7.2. **IMAGE ACQUISITION**

The TLC and FRC-CT data sets were acquired under an approved IRB protocol at the Columbia University Medical Center, using a GE VCT 64 multi-detector row CT scanner during breath holds at TLC and FRC from patients diagnosed with COPD. Every volumetric dataset contains 638-671 image sections, each 0.625 mm thick and spaced 0.5 mm apart, and with a reconstruction matrix of 512 by 512 pixels using a B31F kernel. In-plane pixel spatial resolution was 0.863 mm by 0.863 mm.

The MR angiography images were acquired using the TRICKS sequence with a GE scanner during breath holds at lung volume of 15% of vital capacity. Each volumetric dataset contains

32-52 image sections, each 256 by 256 pixels, measuring 1.875 mm by 1.875 mm by 10 mm (thickness) with a slice spacing of 5 mm. Each volumetric data set was accompanied by a time-lapse subtracted perfusion sequence that was not used for the purpose of this paper. Only the non-subtracted, non-enhanced, MR images from the perfusion sequence at time 0 before the injection of contrast agent were used for registration.

7.3. CT IMAGE SEGMENTATION AND REGISTRATION

To prepare the images for TLC to FRC registration, segmentations of both the TLC and FRC-CT lung volumes were performed, using Hu et al.'s [99] automated, gray-level-thresholding algorithm and Pulmonary Workstation 2 (PW2: VIDA Diagnostics, Coralville, Iowa). Non-rigid 3D registration was then performed for each TLC – FRC lung volume pair using the registration process described in Sections A and B. The result of this registration is a transformation map from the TLC lung volume onto the FRC lung volume.

Using the original TLC voxel intensities, two emphysema masks differing in their Hounsfield unit (HU) thresholds were created on the TLC volume. One was created using an upper threshold of -910 HU to mark mild emphysema-like regions, and the other was created using a higher-specificity threshold of -950 HU. Both masks were then mapped into the FRC-CT lung volumes using the aforementioned TLC-to-FRC transformation map.

7.4. **MRI IMAGE SEGMENTATION AND REGISTRATION**

Thorax MRI images are notoriously difficult to segment, and completely automated algorithms for the segmentation of diseased lungs from MRI images have been particularly difficult to develop. Additionally, accurate and fully quantitative MRI perfusion imaging requires a high temporal resolution. Thus, the spatial resolution of MRI perfusion sequences is lower than usual, in favor of a high temporal resolution. We therefore opted to use a manual process for the segmentation of the MRI lung images.

Mid-coronal views of the FRC-MRI lung volumes were manually segmented for validation using Adobe Photoshop CS5. A binary mask was then made for each mid-coronal slice and a preliminary 3D volume mask was generated. The preliminary 3D volume mask and the FRC-MRI lung volume were then both imported into AMIRA (FEI Scientific Visualization Group) and overlaid. The preliminary 3D volume mask was further manually refined into the finalized 3D volume mask.

Non-rigid 3D registration was then performed for each FRC-CT – FRC-MRI lung volume pair, using the registration process described in Sections A and C, resulting in a transformation map from the FRC-CT lung volume onto the FRC-MRI lung volume. The Hounsfield unit thresholded masks were also mapped from the FRC-CT lung volumes onto the FRC-MRI lung volumes.

7.5. **MAPPING OF EMPHYSEMATOUS REGIONS**

Having segmented the emphysematous regions from the TLC volumes, and now possessing a registration pipeline from TLC-CT to FRC-CT to FRC-MRI, we were able to map the masks of emphysematous regions onto the MRI volumes. However, CT resolution is significantly higher than that of MRI, and thus emphysematous regions mapped in CT must be further re-sampled prior to comparisons with microvascular perfusion MRI.

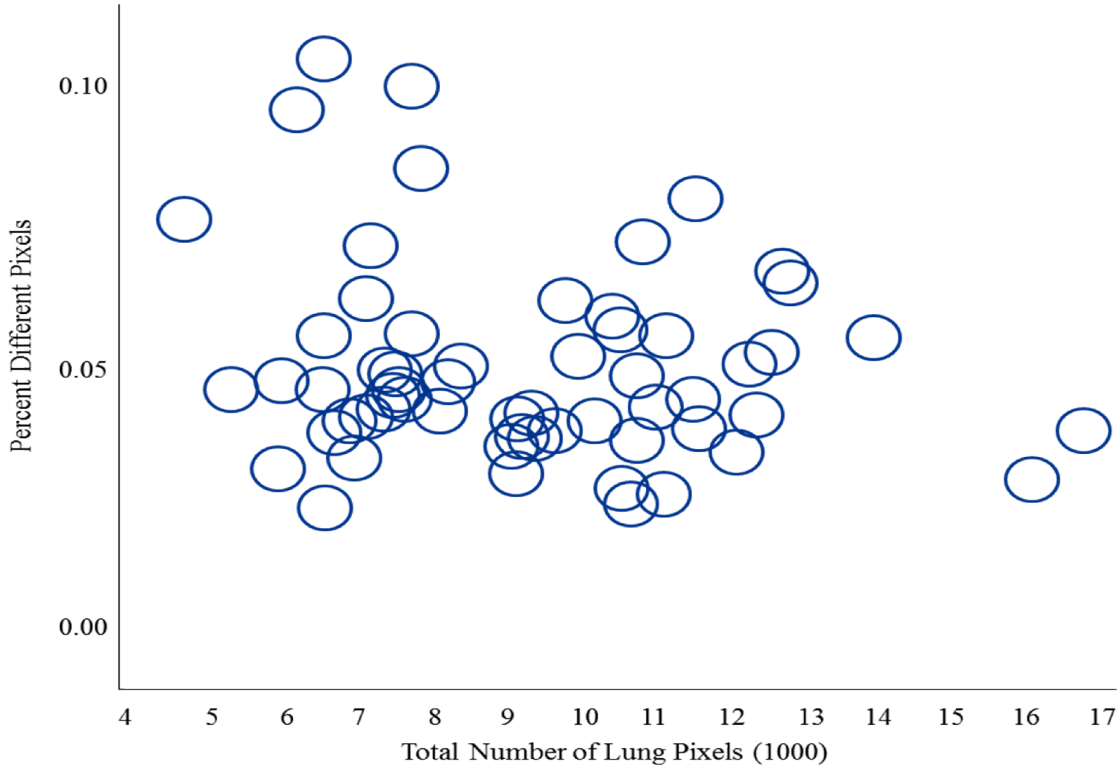

Figure 7.3: *The y-axis shows the percentage of pixels labeled as bronchial tree that differ between the TLC-CT to FRC-CT to FRC-MRI post registered mapping and the pixels manually labeled as bronchial tree by our expert tracers. Total numbers of lung pixels in segmented TLC-CT lung volumes are shown on the x-axis.*

The MR coronal slice thickness of our data was 10 mm, and the coronal thickness of each CT slice was 0.863 mm. Thus, approximately 11 CT slices and mapped emphysematous projections have to be fused in order to approximate one MRI slice thickness. There are three methods for thickening CT slices: maximum intensity projection (MIP), mean intensity projection, and minimum intensity projection (minIP).

The most common methods are maximum and minimum intensity projections, and both methods were explored to determine which would be the clinically valid approach. The minIP applies a "minimum" operation on each set of voxels that forms a ray normal to the slices to be joined. In contrast, MIP transfers only the intensity of the voxel with the greatest level of attenuation onto the thickened slice and is thus a less sensitive measure of emphysema.

The transformed emphysema masks were further post-processed with erosion and dilation morphological operators to remove any stray pixels and create a mask that is more relevant for later correlation with MR angiography results.

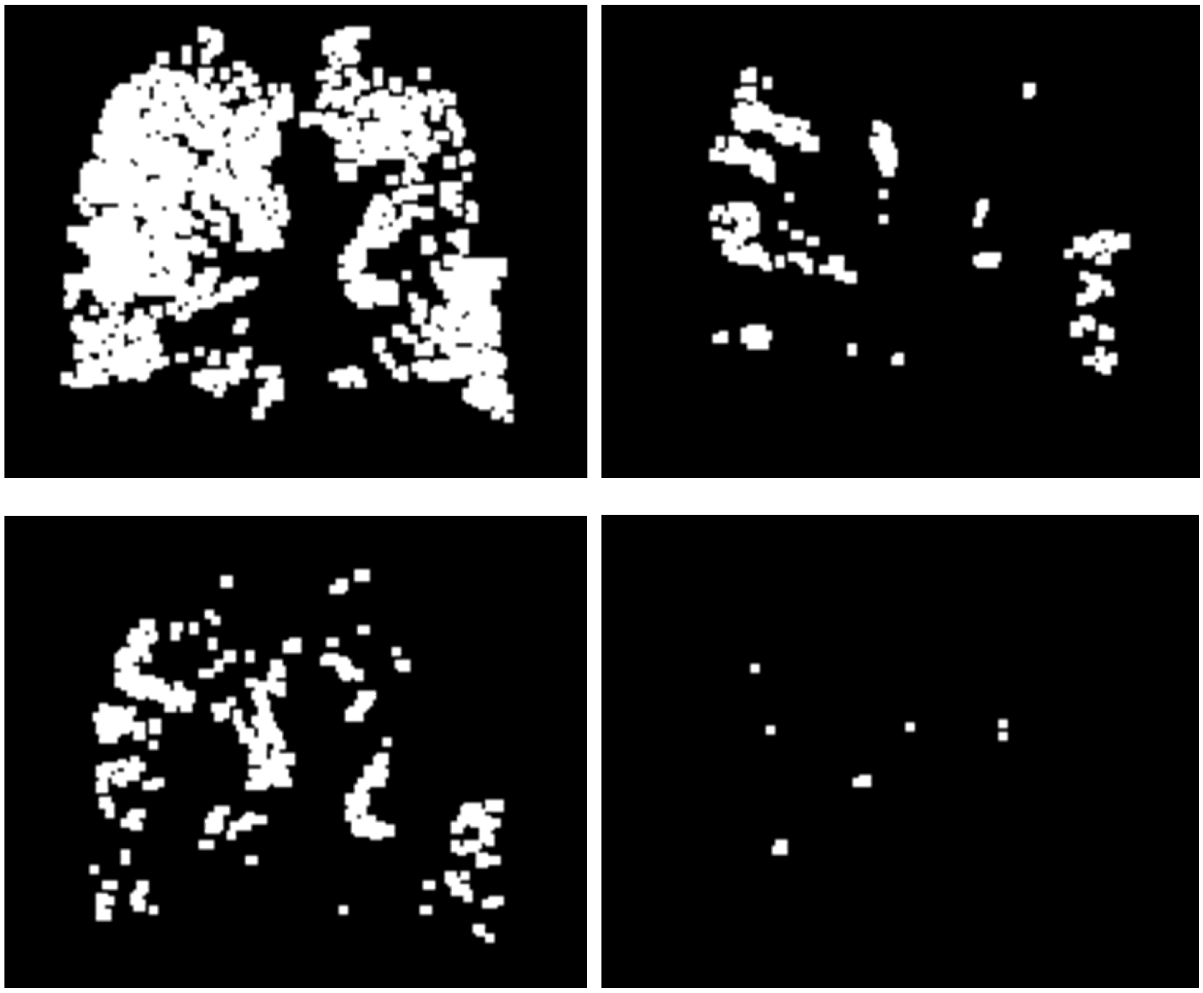

Figure 7.4: *A) MinIP emphysema mask (-910) registered to MR. B) MIP emphysema mask (-910) registered to MR. C) MinIP emphysema mask (-950) registered to MR. D) MIP emphysema mask (-950) registered to MR.*

7.6. **RESULTS**

Our methods successfully registered 59 image sets, in each case creating an emphysema mask to overlay the target MR volume. The resulting volumes allows simple, superimposed correlations of emphysematous regions as identified on CT with areas of decreased perfusion on MRI. At approximately 4.5 minutes per registration on one computing node, the process is

computationally efficient to a significant degree and allows for this process to be performed on a routine clinical basis without significant impact to clinical workflow.

To determine the correct parameters for registration, slice thickening, and emphysematous thresholding, mid-coronal projections of the MRI-registered emphysema masks were created and manually examined for clinical suitability by our collaborating radiologists as described in methods.

To locate the mid-coronal slice from the MR volumes, the most anterior and posterior coronal slices containing lung parenchyma were identified and the mid-coronal slice was assumed to be halfway between those two image slices (Figure 7.4 A-D).

Since the mid-coronal MR slice corresponds to the eleven central coronal CT slices, the latter set had to be fused by projection to create comparable image slabs. After fine-tuning the registration grid resolutions, mid-coronal CT projections from 59 participants were generated in four different ways: combinations of two CT thickening techniques – maximum and minimum intensity projections – and two emphysema quantification techniques – thresholds at -910 and -950 HU.

As figure 7.4 demonstrates, regions defined as emphysema-like are highly sensitive to the slice thickening algorithm and the Hounsfield unit threshold. Applying MIP thickening resulted in the detection of significantly fewer emphysema-like regions on the 10 mm constructed slabs

compared to minIP thickening. As expected, a -950 HU threshold yielded fewer lung regions classified as emphysematous on the slabs compared to a -910 HU threshold. The combination of MinIP and -950 HU thresholding conveyed emphysema-like regions even smaller than 10 mm (along the anteroposterior axis) to the final thickened slice, while retaining the requirement of significantly low CT attenuation. The application of well-established morphological operators allowed for stray pixels to be eliminated without affecting the sensitivity of the COPD masks.

Having established matching image slabs, we validated our serial registration algorithm by comparing the post-registration anatomical landmark from the original CT volume to the final MR volume with the same anatomical landmark that was directly manually segmented on the MR volume (Figure 7.4). Specifically, we used a validated algorithm to automatically segment voxels of the main bronchial tree from the mid-coronal slices of each TLC-CT volume. We then applied the transformation matrices derived earlier in the registration process to the segmented bronchial trees, thus allowing for the bronchial tree to be mapped from TLC-CT volume, to the FRC-CT volume, and ultimately to the FRC-MR volume. We then also manually segmented the main bronchial tree in the mid-coronal slice on the FRC-MRI. Thus, we were able to use the manual segmentation as ground truth to evaluate the transformed segmentation from the TLC-CT volume. We quantified the error by computing the number of pixels that differed between the FRC-CT transformed bronchial segmentation and the directly manually segmented bronchial tree in the FRC-MR volume. We then normalized this value by dividing the absolute difference in

number of pixels by the number of segmented lung pixels in the FRC-MR mid-coronal slice. This normalization was performed in order to compensate for the differences in field of view between different patients' MRI scans. The graph of the normalized difference ratio plotted against the total number of segmented lung volume pixels on TLC-CT is displayed in Figure 7.3. It is important to note that the total number of segmented lung volume pixels cannot be directly used as a proxy for lung volume due to similar concerns regarding the field of view changes between CT scans. We were able to obtain a difference ratio with a mean of approximately 6.3%.

Except in a few outlier cases, registration values tended to be very stable. After visual inspection of these instances, it was determined that their FRC-CT volumes and FRC-MRI volumes differed significantly and the registration process detailed above was unable to sufficiently compensate for the significant changes in volume. Large registration errors were therefore created. It is important to note that those cases should most likely not be registered or used for further perfusion/tissue damage comparison, as functional comparisons on those volumes would most likely yield clinically misleading results.

7.7. **CONCLUSION**

This chapter has presented a fast and clinically relevant algorithm to combine the detection of emphysematous regions from TLC-CT images with perfusion data from microvascular perfusion FRC-MRI images. A multi-scale B-spine registration algorithm was used with two different

similarity measures in order to serially register the imaged volumes from TLC-CT to FRC-CT, and then from FRC-CT to FRC-MRI.

As part of the MESA COPD study, our collaborators have collected an extensive number of TLC-CT, FRC-CT, and FRC-MRI volumes for hundreds of patients who additionally have had several scans of each modality over the past decade. The application of this algorithm will aid in the clarification of the pathophysiology behind the interplay of vascular changes and parenchymal destruction in the evolution of COPD.

Chapter 8: Conclusions and Future Research

This thesis rests on the application of multi-resolution algorithms in the analysis of images and 3D image volumes. The current research focused on two specific applications of multi-resolution techniques to 1) provide a framework for automated screening of high throughput protein crystallization images and 2) to provide imaging tools necessary for the validation of a hypothesis on the pathophysiology of pulmonary emphysema.

8.1. PROTEIN CRYSTAL IMAGE CLASSIFICATION

This work proposed a classification algorithm that exceeds 99% in sensitivity and 96% in specificity. Furthermore, the classifier is able to process each image with off-the-shelf computer components at approximately 7 seconds per image, a speed that makes this algorithm usable in high throughput settings. The classification algorithm has been tested on a previous published and available dataset from NESG and is now well validated. Further work can be done to convert the algorithm into C to further improve the algorithm's processing speed. The current implementation of the classifier does not allow for online learning. That is to say that any updates to the model will require the entire classifier to be retrained using a new set of labeled images. Recent advances in online random forests can be implemented into the algorithm to allow it to perform online learning. [45]. The ultimate goal of this classifier is for it to be used in

a high throughput structural genomics setting. Further work should be done in order to implement this now finalized algorithm into the structural genomics pipeline at NESG to hopefully detect previously unsolvable protein structures that did include protein crystals but were missed either by crystallographers or because the wrong time-based image was examined by the crystallographers.

8.2. **PATHOPHYSIOLOGY OF PULMONARY EMPHYSEMA**

The second part of this work proposed a 3D image registration algorithm to register regions of emphysema as quantified by densitometry on lung CT with MR lung volumes. The ability to register quantitatively determined regions of emphysema with perfusion MRI will allow for further exploration of the pathogenesis and cardiopulmonary implications of Chronic Obstructive Pulmonary Disorder (COPD). As discussed earlier, recent evidence suggests that endothelial dysfunction and alterations in pulmonary vascular response occur early in COPD and may represent an important vascular pathway in the development of smoking-associated emphysema [54], [85], [100]. Several groups present global perfusion differences between patients with COPD and patients without COPD. However, no work has yet quantified whether this global perfusion decrease is explained by a uniform decrease in global perfusion due to uniform microvascular damage, or whether, as in the case of emphysema damage, impairment of

microvascular perfusion is also non-uniform and, indeed, the precursor to further emphysematous damage. Unpublished preliminary work by our collaborators and us currently suggest that the latter is true, but more investigations are required in order to properly verify this claim.

8.3. **MULTI-RESOLUTION IMAGE PROCESSING**

This thesis additionally introduced several novel applications of multi-resolution image processing techniques, including new image descriptor features and the introduction of a new multi-resolution image registration architecture that can be applied to other areas of biomedical imaging and computer vision. The registration framework proposed here will be increasingly important as high resolution 3D and 4D volumes become increasingly prevalent in clinical practice and the need for near real-time registration becomes not only useful, but often a necessary condition for busy clinicians. The new image features proposed in this thesis should be tested in other computer vision applications that privilege texture in the classification of those images. It is likely that these image features will help produce a more accurate classifier, as is this case in this thesis.

References

[1] D. Paquin, D. Levy, E. Schreibmann, and L. Xing, “Multiscale image registration,” *Math Biosci Eng*, vol. 3, no. 2, pp. 389–418, Apr. 2006.

[2] B. Li, G. E. Christensen, E. A. Hoffman, G. McLennan, and J. M. Reinhardt, “Establishing a normative atlas of the human lung: computing the average transformation and atlas construction.,” *Academic radiology*, vol. 19, no. 11, pp. 1368–1381, Nov. 2012.

[3] W. Lu, J. H. Song, G. E. Christensen, P. J. Parikh, J. D. Bradley, and D. A. Low, “Modeling lung motion using consistent image registration in 4D computed tomography for radiation therapy,” presented at the Medical Imaging, 2006, vol. 6144, pp. 61442L–61442L–8.

[4] Y. Yin, “MDCT-based dynamic, subject-specific lung models via image registration for CFD-based interrogation of regional lung function,” 2011.

[5] Y. Yin, E. A. Hoffman, and C.-L. Lin, “Lung lobar slippage assessed with the aid of image registration.,” *Med Image Comput Comput Assist Interv*, vol. 13, no. 2, pp. 578–585, 2010.

[6] Y. Yin, E. A. Hoffman, and C.-L. Lin, “Mass preserving nonrigid registration of CT lung images using cubic B-spline.,” *Med Phys*, vol. 36, no. 9, pp. 4213–4222, Sep. 2009.

[7] Y. Yin, E. A. Hoffman, K. Ding, J. M. Reinhardt, and C.-L. Lin, “A cubic B-spline-based hybrid registration of lung CT images for a dynamic airway geometric model with large deformation.,” *Phys Med Biol*, vol. 56, no. 1, pp. 203–218, Jan. 2011.

[8] F. Lamare, T. Cresson, J. Savean, C. Cheze Le Rest, A. J. Reader, and D. Visvikis, “Respiratory motion correction for PET oncology applications using affine transformation of list mode data.,” *Phys Med Biol*, vol. 52, no. 1, pp. 121–140, Jan. 2007.

[9] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes, “Nonrigid registration using free-form deformations: application to breast MR images.,” *IEEE Trans Med Imaging*, vol. 18, no. 8, pp. 712–721, Jul. 1999.

[10] C. O. S. Sorzano, P. Thevenaz, and M. Unser, “Elastic registration of biological images using vector-spline regularization.,” *IEEE Trans Biomed Eng*, vol. 52, no. 4, pp. 652–663, Apr. 2005.

[11] T. Rohlfing, C. R. J. Maurer, D. A. Bluemke, and M. A. Jacobs, “Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint.,” *IEEE Trans Med Imaging*, vol. 22, no. 6, pp. 730–741, Jun. 2003.

[12] S. Y. Chun and J. A. Fessler, “A Simple Regularizer for B-spline Nonrigid Image Registration That Encourages Local Invertibility,” *IEEE J. Sel. Top. Signal Process.*, vol. 3, no. 1, pp. 159–169.

[13] S. K. Burley, S. C. Almo, J. B. Bonanno, M. Capel, M. R. Chance, T. Gaasterland, D. Lin, A. Sali, F. W. Studier, and S. Swaminathan, “Structural genomics: beyond the human genome project.,” *Nat Genet*, vol. 23, no. 2, pp. 151–157, Oct. 1999.

[14] P. Bertone, Y. Kluger, N. Lan, D. Zheng, D. Christendat, A. Yee, A. M. Edwards, C. H. Arrowsmith, G. T. Montelione, and M. Gerstein, “SPINE: an integrated tracking database and data mining approach for identifying feasible targets in high-throughput structural proteomics.,” *Nucleic Acids Res.*, vol. 29, no. 13, pp. 2884–2898, Jul. 2001.

[15] T. C. Terwilliger, “Structural genomics in North America - Nature Structural & Molecular Biology,” *Nat. Struct Biol.*, vol. 7, pp. 935–939, Nov. 2000.

[16] J. M. Aramini, P. Rossi, C. Anklin, R. Xiao, and G. T. Montelione, “Microgram-scale protein structure determination by NMR,” *Nat Meth*, vol. 4, no. 6, pp. 491–493, May 2007.

[17] R. Xiao, S. Anderson, J. Aramini, R. Belote, W. A. Buchwald, C. Ciccosanti, K. Conover, J. K. Everett, K. Hamilton, Y. J. Huang, H. Janjua, M. Jiang, G. J. Kornhaber, D. Y. Lee, J. Y. Locke, L.-C. Ma, M. Maglaqui, L. Mao, S. Mitra, D. Patel, P. Rossi, S. Sahdev, S. Sharma, R. Shastry, G. V. T. Swapna, S. N. Tong, D. Wang, H. Wang, L. Zhao, G. T. Montelione, and T. B. Acton, “The high-throughput protein sample production platform of the Northeast Structural Genomics Consortium,” *J. Struct. Biol.*, vol. 172, no. 1, pp. 21–33, Oct. 2010.

[18] M. J. Po and A. F. Laine, “Leveraging genetic algorithm and neural network in automated protein crystal recognition.,” *Conf Proc IEEE Eng Med Biol Soc*, vol. 2008, pp. 1926–1929, 2008.

[19] I. J. Christian A Cumbaa, “Protein crystallization analysis on the World Community Grid,” *Journal of structural and functional genomics*, vol. 11, no. 1, pp. 61–69, Mar. 2010.

[20] C. A. Cumbaa, A. Lauricella, N. Fehrman, C. Veatch, R. Collins, J. R. Luft, G. DeTitta, and I. Jurisica, “Automatic classification of sub-microlitre protein-crystallization trials in 1536-well plates,” vol. 59, no. 9, pp. 1619–1627, Sep. 2003.

[21] K. Saitoh, K. Kawabata, H. Asama, T. Mishima, M. Sugahara, and M. Miyano, “Evaluation of protein crystallization states based on texture information derived from greyscale images,” *Acta Crystallogr D Biol Crystallogr*, vol. 61, no. 7, pp. 873–880, Jul. 2005.

[22] X. Zhu, S. Sun, and M. Bern, “Classification of protein crystallization imagery.,” *Conf Proc IEEE Eng Med Biol Soc*, vol. 3, pp. 1628–1631, 2004.

[23] S. Pan, G. Shavit, M. Penas-Centeno, D. H. Xu, L. Shapiro, R. Ladner, E. Riskin, W.

Hol, and D. Meldrum, “Automated classification of protein crystallization images using support vector machines with scale-invariant texture and Gabor features,” *Acta Crystallogr D Biol Crystallogr*, vol. 62, no. 3, pp. 271–279, Mar. 2006.

[24] E. H. Snell, J. R. Luft, S. A. Potter, A. M. Lauricella, S. M. Gulde, M. G. Malkowski, M. Koszelak-Rosenblum, M. I. Said, J. L. Smith, C. K. Veatch, R. J. Collins, G. Franks, M. Thayer, C. Cumbaa, I. Jurisica, and G. T. DeTitta, “Establishing a training set through the visual analysis of crystallization trials. Part I: approximately 150,000 images.,” *Acta Crystallogr D Biol Crystallogr*, vol. 64, no. 11, pp. 1123–1130, Nov. 2008.

[25] E. H. Snell, A. M. Lauricella, S. A. Potter, J. R. Luft, S. M. Gulde, R. J. Collins, G. Franks, M. G. Malkowski, C. Cumbaa, I. Jurisica, and G. T. DeTitta, “Establishing a training set through the visual analysis of crystallization trials. Part II: crystal examples.,” *Acta Crystallogr D Biol Crystallogr*, vol. 64, no. 11, pp. 1131–1137, Nov. 2008.

[26] A. Georgiev and P. K. Allen, “Two-stage robotic crystal mounting of protein crystals for X-ray data collection,” presented at the Automation Science and Engineering, 2008. CASE 2008. IEEE International Conference on, 2008, pp. 1019–1024.

[27] M. Nixon, M. S. Nixon, and A. S. Aguado, *Feature Extraction & Image Processing for Computer Vision*. Academic Press, 2012.

[28] D. Gabor, “Theory of communication. Part 1: The analysis of information,” *Electrical Engineers - Part III: Radio and Communication Engineering, Journal of the Institution of*, vol. 93, no. 26, pp. 429–441, Nov. 1946.

[29] J. G. Daugman, “Two-dimensional spectral analysis of cortical receptive field profiles.,” *Vision Research*, vol. 20, no. 10, pp. 847–856, Jan. 1980.

[30] B. S. Manjunath and W. Y. Ma, “Texture features for browsing and retrieval of image data,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on DOI - 10.1109/34.295913*, vol. 18, no. 8, pp. 837–842, 1996.

[31] S. Arivazhagan, R. A. Priyadharshini, and S. Seedhanadevi, “Object recognition based on gabor wavelet features,” presented at the Devices, Circuits and Systems (ICDCS), 2012 International Conference on, 2012, pp. 340–344.

[32] D. Casanova, J. J. de Mesquita Sá Junior, and O. M. Bruno, “Plant leaf identification using Gabor wavelets,” *International Journal of Imaging Systems and Technology*, vol. 19, no. 3, pp. 236–243, Sep. 2009.

[33] T. Ojala, M. Pietikainen, and T. Maenpaa, “Multiresolution gray-scale and rotation invariant texture classification with local binary patterns,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on DOI - 10.1109/34.295913*, vol. 24, no. 7, pp. 971–987, Jul. 2002.

[34] T. Ojala, M. Pietikäinen, and T. Mäenpää, “A Generalized Local Binary Pattern Operator for Multiresolution Gray Scale and Rotation Invariant Texture

Classification," in *Lecture Notes in Computer Science*, vol. 2013, no. 41, Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 399–408.

[35] L.-K. Soh and C. Tsatsoulis, "Texture analysis of SAR sea ice imagery using gray level co-occurrence matrices," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 37, no. 2, pp. 780–795, 1999.

[36] D. A. Clausi, "An analysis of co-occurrence texture statistics as a function of grey level quantization," *Canadian Journal of Remote Sensing*, vol. 28, no. 1, pp. 45–62, Jun. 2002.

[37] R. M. Haralick, "Statistical and structural approaches to texture," *Proceedings of the IEEE*, vol. 67, no. 5, pp. 786–804, May 1979.

[38] Y. Hai-peng, L. Yi-xing, and L. Zhen-bo, "Auto detection of wood texture orientation by Radon transform," *Journal of Forestry Research*, vol. 16, no. 1, pp. 1–4, Mar. 2005.

[39] A. G. Zuñiga, J. B. Florindo, and O. M. Bruno, "Gabor wavelets combined with volumetric fractal dimension applied to texture analysis," *Pattern Recogn Lett*, 2014.

[40] A. G. Zuniga and O. M. Bruno, "Enhancing Gabor Wavelets Using Volumetric Fractal Dimension," in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, vol. 6419, no. 49, I. Bloch and R. M. Cesar, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 362–369.

[41] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[42] G. Xu, C. Chiu, E. D. Angelini, and A. F. Laine, "An Incremental and Optimized Learning Method for the Automatic Classification of Protein Crystal Images," presented at the Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE IS -, 2006, vol. Supplement, pp. 6526–6529.

[43] N. Sebe and M. S. Lew, "Wavelet based texture classification," presented at the Pattern Recognition, 2000. Proceedings. 15th International Conference on, 2000, vol. 3, pp. 947–950.

[44] S. W. Myint, "Fractal approaches in texture analysis and classification of remotely sensed data: Comparisons with spatial autocorrelation techniques and simple descriptive statistics," *International Journal of Remote Sensing*, vol. 24, no. 9, pp. 1925–1947, Jan. 2003.

[45] A. Saffari, C. Leistner, J. Santner, M. Godec, and H. Bischof, "On-line Random Forests," presented at the Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on IS - SN - VO -, 2009, pp. 1393–1400.

[46] R. A. Pauwels, A. S. Buist, P. Ma, C. R. Jenkins, S. S. Hurd, GOLD Scientific Committee, "Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease: National Heart, Lung, and Blood Institute and World Health Organization Global Initiative for Chronic Obstructive Lung Disease (GOLD): executive summary.," presented at the Respiratory care, 2001, vol. 46, no. 8, pp. 798–

825.
[47] A. M. Miniño, S. L. Murphy, J. Xu, and K. D. Kochanek, “Deaths: final data for 2008.,” *Natl Vital Stat Rep*, vol. 59, no. 10, pp. 1–126, Dec. 2011.
[48] T. L. Petty, “Definition, epidemiology, course, and prognosis of COPD,” *Clinical cornerstone*, 2003.
[49] National Heart Lung and Blood Institute, *Morbidity and mortality: 2009 chart book on cardiovascular, lung, and blood diseases*. Rockville (MD): US Department of Health, 2004.
[50] A. D. Lopez, C. D. Mathers, M. Ezzati, D. T. Jamison, and C. J. L. Murray, “Global and regional burden of disease and risk factors, 2001: systematic analysis of population health data.,” *The Lancet*, vol. 367, no. 9524, pp. 1747–1757, May 2006.
[51] A. D. Lopez and C. D. Mathers, “Measuring the global burden of disease and epidemiological transitions: 2002-2030.,” *Ann Trop Med Parasitol*, vol. 100, no. 5, pp. 481–499, Jul. 2006.
[52] J. S. Schiller, *Summary Health Statistics for the U.S. Population (April 2005): National Health Interview Survey, 2003 (Vital and Health Statistics)*. National Center for Health Statistics, 2005.
[53] C. Raherison and P.-O. Girodet, “Epidemiology of COPD,” *European Respiratory Review*, vol. 18, no. 114, pp. 213–221, Jan. 2009.
[54] M. Mergoni and A. Rossi, “[Physiopathology of acute respiratory failure in COPD and asthma].,” *Minerva Anestesiol*, vol. 67, no. 4, pp. 198–205, Apr. 2001.
[55] W. MacNee, “Pathophysiology of cor pulmonale in chronic obstructive pulmonary disease. Part One.,” *Am J Respir Crit Care Med*, vol. 150, no. 3, pp. 833–852, Sep. 1994.
[56] “Pathophysiology of cor pulmonale in chronic obstructive pulmonary disease. Part two.,” vol. 150, no. 4, pp. 1158–1168, Oct. 1994.
[57] J. M. Heath and R. Mongia, “Chronic bronchitis: Primary care management,” *Am Fam Physician*, vol. 57, no. 10, pp. 2365–2372, 1998.
[58] M. Takahashi, J. Fukuoka, N. Nitta, R. Takazakura, Y. Nagatani, Y. Murakami, H. Otani, and K. Murata, “Imaging of pulmonary emphysema: a pictorial review.,” *Int J Chron Obstruct Pulmon Dis*, vol. 3, no. 2, pp. 193–204, 2008.
[59] W. M. Thurlbeck and N. L. Müller, “Emphysema: definition, imaging, and quantification.,” *AJR Am J Roentgenol*, vol. 163, no. 5, pp. 1017–1025, Nov. 1994.
[60] E. J. Stern, M. S. Frank, J. F. Schmutz, R. W. Glenny, R. A. Schmidt, and J. D. Godwin, “Panlobular pulmonary emphysema caused by i.v. injection of methylphenidate (Ritalin): findings on chest radiographs and CT scans.,” *Am J Roentgenol*, vol. 162, no. 3, pp. 555–560, Mar. 1994.
[61] R. G. Barr, D. A. Bluemke, F. S. Ahmed, J. J. Carr, P. L. Enright, E. A. Hoffman, R. Jiang, S. M. Kawut, R. A. Kronmal, J. O. A. C. Lima, E. Shahar, L. J. Smith, and K. E.

Watson, “Percent emphysema, airflow obstruction, and impaired left ventricular filling.,” *New Engl J Med*, vol. 362, no. 3, pp. 217–227, Jan. 2010.

[62] I. Peytremann-Bridevaux, P. Taffe, B. Burnand, P. O. Bridevaux, and M. A. Puhan, “Mortality of patients with COPD participating in chronic disease management programmes: a happy end?,” *Thorax*, vol. 69, no. 9, pp. 865–866, Sep. 2014.

[63] A. Manoharan, P. M. Short, W. J. Anderson, and B. J. Lipworth, “Impact of Long-Acting Bronchodilators and Exposure to Inhaled Corticosteroids on Mortality in COPD: A Real-Life Retrospective Cohort Study,” *Lung*, vol. 192, no. 5, pp. 649–652, Oct. 2014.

[64] D. Litmanovich, P. M. Boiselle, and A. A. Bankier, “CT of pulmonary emphysema - current status, challenges, and future directions,” *Eur Radiol*, vol. 19, no. 3, pp. 537–551.

[65] H. O. Coxson, R. M. Rogers, and K. P. Whittall, “A quantification of the lung surface area in emphysema using computed tomography,” *American journal of …*, 1999.

[66] K. T. Bae, R. M. Slone, D. S. Gierada, R. D. Yusen, and J. D. Cooper, “Patients with emphysema: Quantitative CT analysis before and after lung volume reduction surgery - Work in progress,” *Radiology*, vol. 203, no. 3, pp. 705–714, Jun. 1997.

[67] C. J. Galbán, M. K. Han, J. L. Boes, K. A. Chughtai, C. R. Meyer, T. D. Johnson, S. Galbán, A. Rehemtulla, E. A. Kazerooni, F. J. Martinez, and B. D. Ross, “Computed tomography-based biomarker provides unique signature for diagnosis of COPD phenotypes and disease progression,” *Nature Medicine*, vol. 18, no. 11, pp. 1711–1715, Nov. 2012.

[68] R. Uppaluri, T. Mitsa, M. Sonka, and E. A. Hoffman, “Quantification of pulmonary emphysema from lung computed tomography images,” *American journal of …*, 1997.

[69] P. A. Gevenois, P. De Vuyst, V. de Maertelaer, J. Zanen, D. Jacobovitz, M. G. Cosio, and J. C. Yernault, “Comparison of computed density and microscopic morphometry in pulmonary emphysema.,” *American journal of …*, 1996.

[70] P. A. Gevenois, M. C. Koob, D. Jacobovitz, P. De Vuyst, J. C. Yernault, and J. Struyven, “Whole lung sections for computed tomographic-pathologic correlations. Modified Gough-Wentworth technique.,” *Investigative Radiology*, vol. 28, no. 3, pp. 242–246, Mar. 1993.

[71] R. J. Knudson, J. R. Standen, W. T. Kaltenborn, D. E. Knudson, K. Rehm, M. P. Habib, and J. D. Newell, “Expiratory computed tomography for assessment of suspected pulmonary emphysema.,” *Chest*, vol. 99, no. 6, pp. 1357–1366, Jun. 1991.

[72] R. G. Barr, S. Mesia-Vela, J. H. M. Austin, R. C. Basner, B. M. Keller, A. P. Reeves, D. Shimbo, and L. Stevenson, “Impaired flow-mediated dilation is associated with low pulmonary function and emphysema in ex-smokers: the Emphysema and Cancer Action Project (EMCAP) Study.,” *Am J Respir Crit Care Med*, vol. 176, no. 12, pp. 1200–1207, Dec. 2007.

[73] H. Kanazawa, K. Asai, K. Hirata, and J. Yoshikawa, “Possible effects of vascular endothelial growth factor in the pathogenesis of chronic obstructive pulmonary disease.,” *Am J Med*, vol. 114, no. 5, pp. 354–358, Apr. 2003.

[74] D. A. McAllister, J. D. Maclay, N. L. Mills, G. Mair, J. Miller, D. Anderson, D. E. Newby, J. T. Murchison, and W. Macnee, “Arterial stiffness is independently associated with emphysema severity in patients with chronic obstructive pulmonary disease.,” *Am J Respir Crit Care Med*, vol. 176, no. 12, pp. 1208–1214, Dec. 2007.

[75] R. B. Easley, M. K. Fuld, A. Fernandez-Bustamante, E. A. Hoffman, and B. A. Simon, “Mechanism of hypoxemia in acute lung injury evaluated by multidetector-row CT.,” *Academic radiology*, vol. 13, no. 7, pp. 916–921, Jul. 2006.

[76] R. Gust, J. Kozlowski, A. H. Stephenson, and D. P. Schuster, “Synergistic hemodynamic effects of low-dose endotoxin and acute lung injury.,” *Am J Respir Crit Care Med*, vol. 157, no. 6, pp. 1919–1926, Jun. 1998.

[77] D. P. Schuster and G. F. Marklin, “The effect of regional lung injury or alveolar hypoxia on pulmonary blood flow and lung water measured by positron emission tomography.,” *Am Rev Respir Dis*, vol. 133, no. 6, pp. 1037–1042, Jun. 1986.

[78] E. A. Hoffman, B. A. Simon, and G. McLennan, “State of the Art. A structural and functional assessment of the lung via multidetector-row computed tomography: phenotyping chronic obstructive pulmonary disease.,” *Proc Am Thorac Soc*, vol. 3, no. 6, pp. 519–532, Aug. 2006.

[79] S. K. Alford, E. J. R. van Beek, G. McLennan, and E. A. Hoffman, “Heterogeneity of pulmonary perfusion as a mechanistic image-based phenotype in emphysema susceptible smokers.,” *P Natl Acad Sci Usa*, vol. 107, no. 16, pp. 7485–7490, Apr. 2010.

[80] Y. Ohno, H. Hatabu, K. Murase, T. Higashino, H. Kawamitsu, H. Watanabe, D. Takenaka, M. Fujii, and K. Sugimura, “Quantitative assessment of regional pulmonary perfusion in the entire lung using three-dimensional ultrafast dynamic contrast-enhanced magnetic resonance imaging: Preliminary experience in 40 subjects.,” *J Magn Reson Imaging*, vol. 20, no. 3, pp. 353–365, Sep. 2004.

[81] F. Risse, W. Semmler, H.-U. Kauczor, and C. Fink, “Dual-bolus approach to quantitative measurement of pulmonary perfusion by contrast-enhanced MRI.,” *J Magn Reson Imaging*, vol. 24, no. 6, pp. 1284–1290, Dec. 2006.

[82] P. Meier and K. L. A. Zierler, “On the Theory of the Indicator-Dilution Method for Measurement of Blood Flow and Volume,” *J Appl Physiol*, vol. 6, no. 12, pp. 731–744, Jun. 1954.

[83] K. Hueper, M. A. Parikh, M. R. Prince, C. Schoenfeld, C. Liu, D. A. Bluemke, S. M. Dashnaw, T. A. Goldstein, E. A. Hoffman, J. A. Lima, J. Skrok, J. Zheng, R. G. Barr, and J. Vogel-Claussen, “Quantitative and Semiquantitative Measures of Regional Pulmonary Microvascular Perfusion by Magnetic Resonance Imaging and Their

Relationships to Global Lung Perfusion and Lung Diffusing Capacity," *Investigative Radiology*, vol. 48, no. 4, p. 1, Feb. 2013.

[84] S. Santos, V. I. Peinado, J. Ramirez, T. Melgosa, J. Roca, R. Rodriguez-Roisin, and J. A. Barbera, "Characterization of pulmonary vascular remodelling in smokers and patients with mild COPD," *European Respiratory Journal*, vol. 19, no. 4, pp. 632–638, Apr. 2002.

[85] G. G. Brusselle, G. F. Joos, and K. R. Bracke, "New insights into the immunology of chronic obstructive pulmonary disease," *The Lancet*, vol. 378, no. 9795, pp. 1015–1026, Sep. 2011.

[86] M. Kinsella, N. L. Müller, R. T. Abboud, N. J. Morrison, and A. DyBuncio, "Quantitation of emphysema by computed tomography using a 'density mask' program and correlation with pulmonary function tests.," *Chest*, vol. 97, no. 2, pp. 315–321, Feb. 1990.

[87] S. R. Hopkins, T. J. Arai, A. C. Henderson, D. L. Levin, R. B. Buxton, and G. K. Prisk, "Lung volume does not alter the distribution of pulmonary perfusion in dependent lung in supine humans," *J Physiol*, vol. 588, no. 23, pp. jphysiol.2010.196063–4768, Oct. 2010.

[88] G. Camiciottoli, F. Bigazzi, M. Bartolucci, L. Cestelli, M. Paoletti, S. Diciotti, E. Cavigli, C. Magni, L. Buonasera, M. Mascalchi, and M. Pistolesi, "BODE-index, modified BODE-index and ADO-score in Chronic Obstructive Pulmonary Disease: Relationship with COPD phenotypes and CT lung density changes," *COPD*, vol. 9, no. 3, pp. 297–304, May 2012.

[89] C. E. Come, A. A. Diaz, D. Curran-Everett, N. Muralidhar, C. P. Hersh, J. A. Zach, J. Schroeder, D. A. Lynch, B. Celli, and G. R. Washko, "Characterizing functional lung heterogeneity in COPD using reference equations for CT scan-measured lobar volumes.," *Chest*, vol. 143, no. 6, pp. 1607–1617, Jun. 2013.

[90] B. Li, G. E. Christensen, E. A. Hoffman, G. McLennan, and J. M. Reinhardt, "Pulmonary CT image registration and warping for tracking tissue deformation during the respiratory cycle through 3D consistent image registration.," *Med Phys*, vol. 35, no. 12, pp. 5575–5583, Dec. 2008.

[91] G. Li, H. Xie, H. Ning, D. Citrin, J. Capala, R. Maass-Moreno, P. Guion, B. Arora, N. Coleman, K. Camphausen, and R. W. Miller, "Accuracy of 3D volumetric image registration based on CT, MR and PET/CT phantom experiments.," *J Appl Clin Med Phys*, vol. 9, no. 4, pp. 2781–2781, Jan. 2008.

[92] T. Guerrero, G. Zhang, T. C. Huang, and K. P. Lin, "Intrathoracic tumour motion estimation from CT imaging using the 3D optical flow method," *Phys Med Biol*, vol. 49, no. 17, pp. 4147–4161, 2004.

[93] K. Ding, K. Cao, M. K. Fuld, K. Du, G. E. Christensen, E. A. Hoffman, and J. M. Reinhardt, "Comparison of image registration based measures of regional lung

ventilation from dynamic spiral CT with Xe-CT.," *Med Phys*, vol. 39, no. 8, pp. 5084–5098, Aug. 2012.

[94] D. Rueckert, P. Aljabar, R. A. Heckemann, J. V. Hajnal, and A. Hammers, "Diffeomorphic Registration Using B-Splines," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, vol. 4191, no. 86, Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 702–709.

[95] S. Lee, G. Wolberg, and S. Y. Shin, "Scattered data interpolation with multilevel B-splines," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 3, no. 3, pp. 228–244, 1997.

[96] S. Klein, M. Staring, and J. P. W. Pluim, "Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines," *IEEE Trans Image Process*, vol. 16, no. 12, pp. 2879–2890, 2007.

[97] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT image registration in the chest using free-form deformations.," *IEEE Trans Med Imaging*, vol. 22, no. 1, pp. 120–128, Dec. 2002.

[98] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," *Image Processing*, 2000.

[99] S. Hu, E. A. Hoffman, and J. M. Reinhardt, "Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images.," *IEEE Trans Med Imaging*, vol. 20, no. 6, pp. 490–498, Jun. 2001.

[100] D. R. Fraidenburg, G. Voiriot, N. M. Pohl, K. A. Smith, and J. X. Yuan, "Lung Inflammation Downregulates TRPC Expression And Impairs Hypoxic Pulmonary Vasoconstriction By Activating The Calcineurin/NFAT Signaling Pathway," in *American Thoracic Society International Conference Abstracts*, American Thoracic Society, 2014, pp. A4144–A4144.