

Developing clinical measures of lung function in COPD patients using medical imaging and computational modelling

DPhil Thesis

Tom Doel

UNIVERSITY OF OXFORD

Preface

This thesis is a submission for Doctor of Philosophy (DPhil or PhD) at the University of Oxford. This work was undertaken in the Computational Biology Group of the Department of Computer Science, and in the Life Sciences Interface Doctoral Training Centre.

This work was funded by EPSRC (Engineering and Physical Sciences Research Council) and the Oxford Biomedical Research Centre.

The open-source Pulmonary Toolkit software developed by the author is available here:

http://code.google.com/p/pulmonarytoolkit/

Supervisors

Doctor Vicente Grau

Institute of Biomedical Engineering, Department of Engineering Science
Oxford e-Research Centre

Professor David Gavaghan

Department of Computer Science
Life Sciences Interface Doctoral Training Centre

Abstract

Developing clinical measures of lung function in COPD patients using medical imaging and computational modelling

Thomas MacArthur Winter Doel, Wolfson College, University of Oxford.
Submission for D.Phil., Michaelmas Term 2012.

Chronic obstructive pulmonary disease (COPD) describes a range of lung conditions including emphysema, chronic bronchitis and small airways disease. While COPD is a major cause of death and debilitating illness, current clinical assessment methods are inadequate: they are a poor predictor of patient outcome and insensitive to mild disease. A new imaging technology, hyperpolarised xenon MRI, offers the hope of improved diagnostic techniques, based on regional measurements using functional imaging. There is a need for quantitative analysis techniques to assist in the interpretation of these images. The aim of this work is to develop these techniques as part of a clinical trial into hyperpolarised xenon MRI.

In this thesis we develop a fully automated pipeline for deriving regional measurements of lung function, making use of the multiple imaging modalities available from the trial. The core of our pipeline is a novel method for automatically segmenting the pulmonary lobes from CT data. This method combines a Hessian-based filter for detecting pulmonary fissures with anatomical cues from segmented lungs, airways and pulmonary vessels. The pipeline also includes methods for segmenting the lungs from CT and MRI data, and the airways from CT data. We apply this lobar map to the xenon MRI data using a multi-modal image registration technique based on automatically segmented lung boundaries, using proton MRI as an intermediate stage. We demonstrate our pipeline by deriving lobar measurements of ventilated volumes and diffusion from hyperpolarised xenon MRI data.

In future work, we will use the trial data to further validate the pipeline and investigate the potential of xenon MRI in the clinical assessment of COPD. We also demonstrate how our work can be extended to build personalised computational models of the lung, which can be used to gain insights into the mechanisms of lung disease.

Acknowledgements

Thank you to my supervisors Vicente Grau and David Gavaghan for all their work on this DPhil. Special thanks also go to Kelly Burrowes, Jon Whiteley and Sarah Waters for all their help, particularly with the modelling aspects of this project. I am grateful to the Oxford Biomedical Research Centre for helping to fund this work. Fergus Gleeson and Rob Davies were instrumental in creating and supporting this project, and Tahreema Matin provided invaluable help with the validation work. Many people have provided valuable assistance, including Michal Plotkowiak, Boguslaw Obara, Mattias Heinrich, Raf Bordas, Minsuok Kim, Ramón Casero Cañas, Felix Bragman, Annabel Nickol, Peggy Xu, Steve Turnbull, Eleanor Mishra, Jenni Lee and many others. I am grateful to them all.

Thank you also to our collaborators, who have been kind enough to provide assistance and to share their data with us. In particular I would like to thank Jim Wild at the University of Sheffield, Bastiaan Driehuys at Duke University, Xavier Maître and Pascal Hagot at Université Paris-Sud and Jan Wolber and Chris Buckley at GE Healthcare. Finally, I would like to thank everyone at the Doctoral Training Centre, the Department of Computer Science and the Oxford e-Research Centre who has provided help and support.

Contents

1 Introduction 8

1.1 Thesis overview 11
1.2 Aims 13
1.3 Contribution 14
1.4 Clinical trial 15
1.5 Software 18
1.6 Data 18

2 Background 21

2.1 Lung physiology 23
2.1.1 Alveoli 25
2.1.2 Airways 26
2.1.3 Breathing 28
2.1.4 Blood vessels 29
2.2 Chronic obstructive pulmonary disease (COPD) 30
2.3 Clinical assessment of COPD 31
2.3.1 Pulmonary function tests (PFTs) 32
2.3.2 CT imaging 34
2.3.3 Hyperpolarised gas MRI 37
2.4 Discussion 40

3 Lung and airway segmentation 41

3.1 Introduction 43
3.2 Lung segmentation from CT 44
3.2.1 Previous work 45
3.2.2 Method 48
3.2.3 Computational efficiency 52
3.2.4 Results and discussion 53
3.2.5 Automated verification of results 53
3.2.6 Performance 56
3.3 Lung segmentation from proton MRI 57
3.3.1 Previous work 57
3.3.2 Method 59

3.3.3 Limitations of method 64
3.3.4 Results 66
3.3.5 Validation 67
3.3.6 Discussion 70
3.3.7 Performance 73
3.4 Airway segmentation from CT 74
3.4.1 Previous work 75
3.4.2 Trachea location 76
3.4.3 Region-growing with explosion control 80
3.4.4 Centreline extraction 84
3.4.5 Radius calculation 86
3.4.6 Labelling of the airways by lobe 89
3.4.7 Results 93
3.4.8 Validation 96
3.4.9 Discussion 97
3.4.10 Performance 99
3.5 Conclusions 100

4 Lobe segmentation 102

4.1 Introduction 104
4.2 Previous work 106
4.2.1 Validation methods 110
4.2.2 Comparison of previous work 111
4.2.3 Limitations of existing work 115
4.3 Method 117
4.3.1 Choice of methodology 117
4.3.2 Overview of method 120
4.3.3 Lung segmentation 121
4.3.4 Blood vessel detection 121
4.3.5 Fissureness 126
4.3.6 Initial lobar approximations 128
4.3.7 Selection of fissure points 130
4.3.8 Construction of fissure surface 132
4.3.9 Lobe segmentation 134
4.4 Results 135
4.4.1 Validation 136
4.5 Discussion 138
4.5.1 Performance 139
4.6 Conclusions 140

5 Regional measurements from xenon MRI 141
5.1 Introduction 143
5.2 Previous work 145
5.2.1 Regional analysis of hyperpolarised MRI 145
5.2.2 Registration of lung images 147
5.3 Methods 149
5.3.1 Registration 149
5.3.2 Ventilated volume analysis 156
5.3.3 Apparent diffusion coefficient (ADC) 158
5.3.4 Emphysema percentage and percentile density 159
5.4 Results 160
5.4.1 Validation 173
5.5 Discussion 174
5.5.1 Data challenges 174
5.5.2 Comparison of COPD and control data 175
5.5.3 Comparison with previous results 176
5.5.4 Errors in registration 177
5.6 Conclusions 178

6 Future work 179
6.1 Regional image analysis applied to the clinical trial 181
6.2 Personalised computational modelling 182

7 Discussion and conclusions 185

Bibliography 189

List of acronyms 205

List of figures 207

List of tables 210

Index 211

1 Introduction

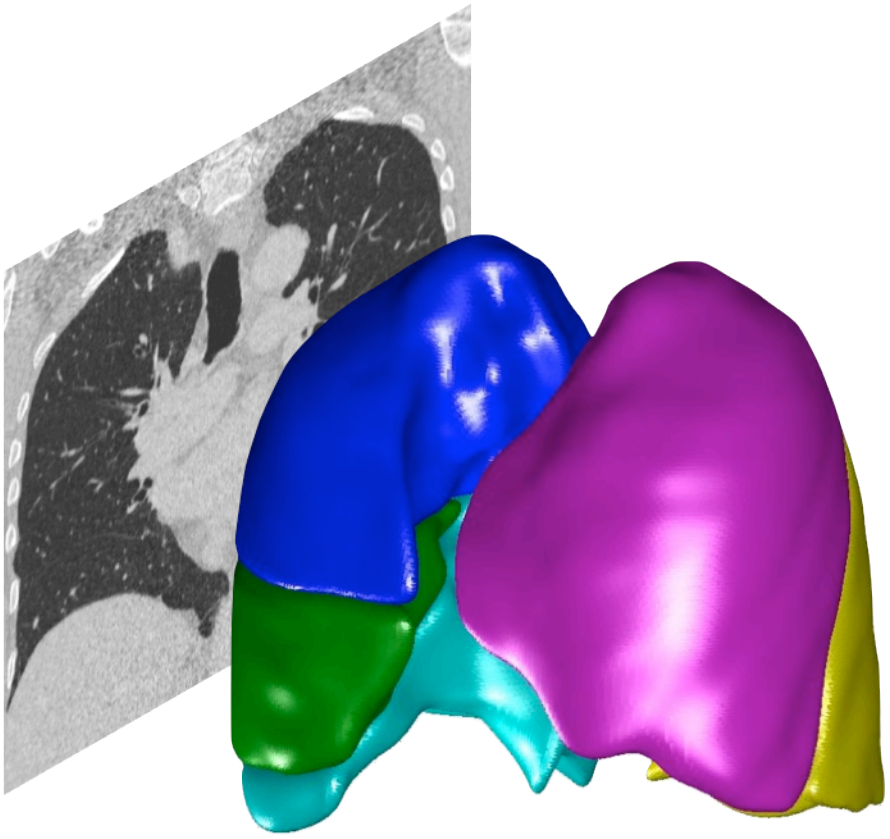

Chapter 1: Introduction



This work concerns the development of imaging techniques to underpin clinical approaches to the treatment of a subset of lung disease termed chronic obstructive pulmonary disease (COPD). COPD includes a range of conditions, such as emphysema, chronic bronchitis and small airways disease, where patients experience long-term respiratory difficulties due to restriction in their airways or destruction of their lung tissue. COPD ranks as the 4th leading cause of death worldwide, accounting for more than 3 million deaths each year with steadily increasing morbidity and mortality rates [1]. In the UK, COPD is responsible for 27,000 deaths each year [2].

Existing clinical methods are inadequate for effective treatment and a poor predictor of patient outcome [3,4]. They are based on global measurements of lung function and ignore the local changes within the lung that characterise most cases of the disease. A new approach to disease assessment, based on regional evaluation of lung function, may provide a better clinical predictor of outcome [5]. This could ultimately lead to improved treatment planning for patients.

To this end, a collaboration group has been formed including clinicians, MRI physicists, image analysis scientists and computational modellers, with the overall goal of improving assessment and treatment options in COPD using

advanced imaging and computational models. In particular, this includes a new imaging technique, hyperpolarised xenon MRI [6], which allows visualisation of the airflow (ventilation) and blood flow (perfusion) in the lung. The group is currently recruiting COPD patients and control subjects into a clinical trial. This cohort will undergo a range of imaging (including hyperpolarised xenon MRI) and function tests repeated at different time points to follow the evolution of the disease. This clinical trial will be the largest and most complete to date to use this technique.

This DPhil aims to assist in the interpretation of multimodal (CT and xenon MRI) images. CT is frequently used in pulmonary disease assessment. It allows visualisation of the main lung structures and is sensitive to regional structural changes in lung tissue resulting from disease [7]. However, CT provides no information about lung function. Conversely, xenon MRI offers little information on lung structure, but provides regional imaging of lung function, such as ventilation, diffusion and perfusion [5]. Our aim is to link the structural information from CT with the functional information from xenon MRI. By combining state-of-the-art methods in image analysis and computational modelling, we will work towards the development of a technique for deriving quantitative regional measurements of lung function from these clinical images.

The aim of this DPhil is to develop a complete, fully automated pipeline for deriving regional measurements from xenon MRI data, and for generating personalised geometries that can be used in computational models to generate results that can be compared regionally with functional images. Specifically, we demonstrate this pipeline by computing lobar measurements of ventilated volumes and apparent diffusion coefficients (ADCs) from xenon MRI data. The lobes are segmented in CT and applied to the xenon MRI data using a multi-modal registration technique. Our pipeline is illustrated in Figure 1.1.

1.1 Thesis overview

The structure of the thesis is as follows. In the remainder of this chapter we outline the main aims and contributions of this thesis, the clinical trial, the data we are using in this thesis and the software we have developed. Chapter 2 presents additional background material on lung physiology and methods for lung disease assessment, including hyperpolarised xenon MRI imaging.

The majority of work in this thesis comprises the segmentation methods of chapters 3 and 4. In chapter 3 we present fully automated methods for segmenting the lungs and airways from CT data. Our work includes calculation of airway radius and centreline information, and the anatomical labelling of the airways according to the pulmonary lobe they serve. We also present a novel algorithm for automatic segmentation of the lungs from MRI data.

In chapter 4 we describe a novel algorithm for segmenting the pulmonary lobes from CT data. This work has been validated and was presented at the ISBI 2012 conference.

The main aim of this thesis, regional analysis of xenon MRI data, is achieved in chapter 5, where we perform a lobar based analysis of xenon MRI data using the lung and lobar segmentations of the previous chapters, via a multi-modal image registration technique. Our xenon MRI measurements are based on ventilated volume and apparent diffusion coefficient (ADC) analysis.

Chapter 6 presents future work we will be carrying out as part of work funded by the Oxford Biomedical Research Centre. This includes analysis of future clinical trial data as it becomes available, and the link to personalised computational modelling of airflow, which is intended to provide more insights into lung disease. Finally, chapter 7 presents discussions and conclusions for the thesis. The thesis structure is illustrated in Figure 1.1

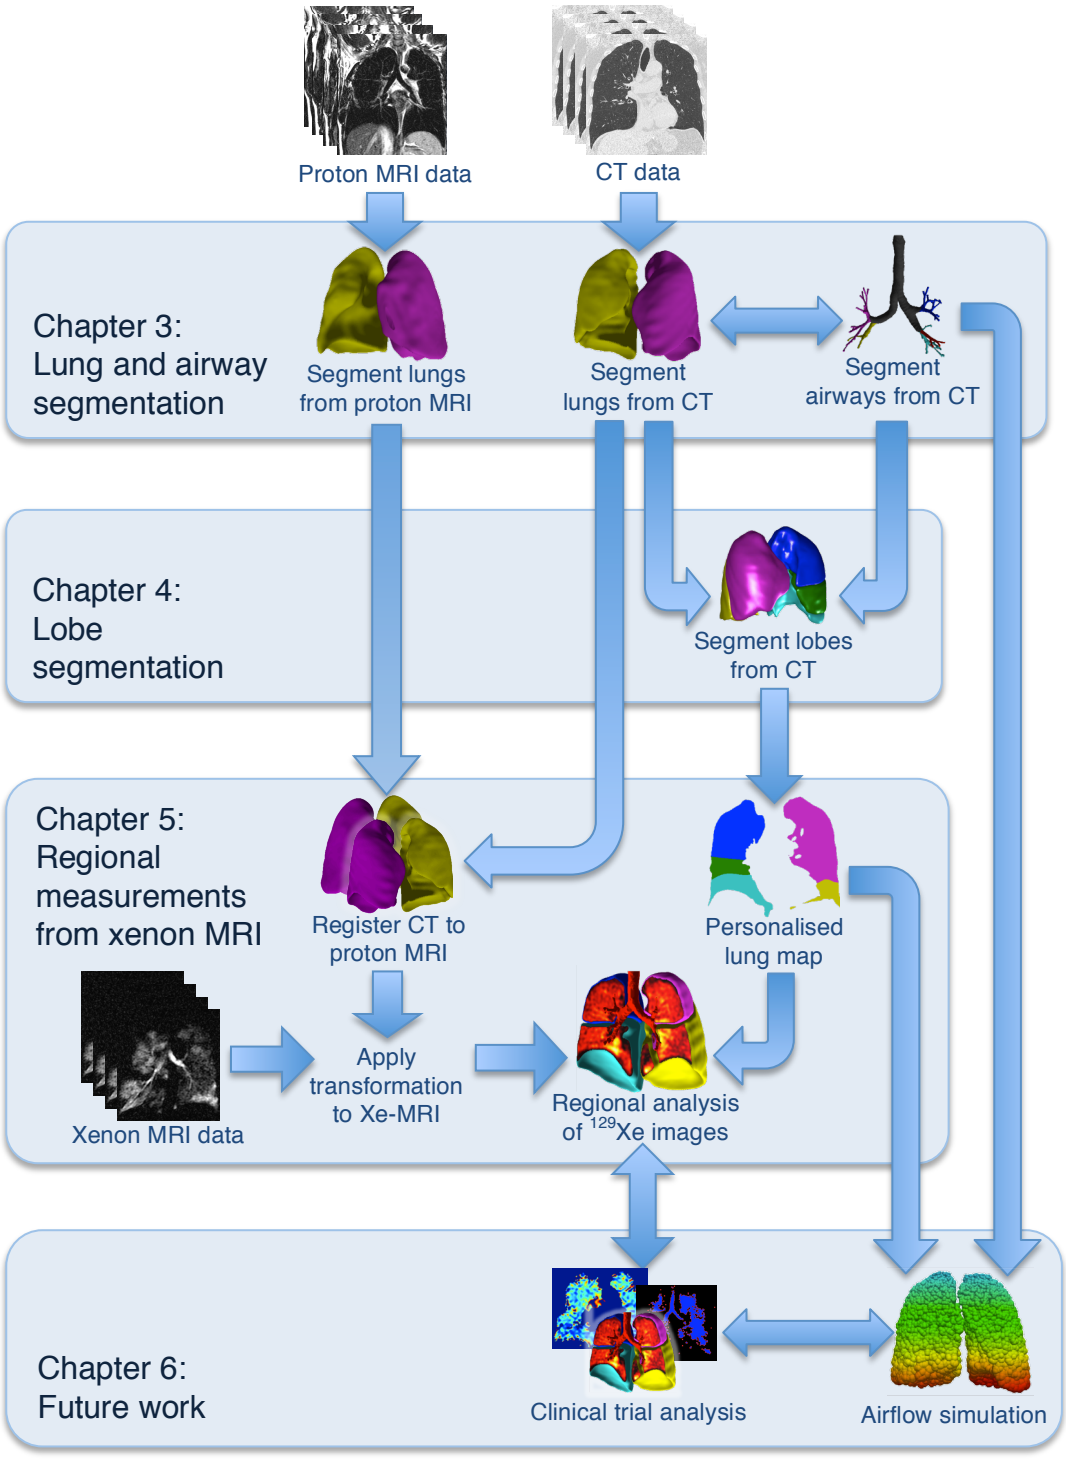

Figure 1.1: Flowchart showing the components of this thesis. We develop a pipeline for deriving regional measurements from hyperpolarised xenon MRI data. Our multi-modal method combines image segmentation and registration techniques from CT, proton MRI and xenon MRI data.

1.2 Aims

1. **To develop a fully automated pipeline for deriving quantitative regional measurements from hyperpolarised gas MRI images**. This pipeline is deigned for use with the clinical trial data, which includes CT, proton MRI and xenon MRI data.
2. **To develop the automatic segmentation techniques** and other algorithms required by the above pipeline. These methods are the main focus of the DPhil; in particular the segmentation of lungs from CT and proton MRI data, and the segmentation of lobes and airways from CT data.
3. **To illustrate the pipeline by demonstrating the lobar analysis of ventilated volume percentage and apparent diffusion coefficient (ADC) measurements for xenon MRI data**. Our data includes COPD and control subjects.

1.3 Contribution

1. Fully automated algorithms for segmenting the lungs (section 3.2) and airways (section 3.4) from clinical CT data, by adapting existing successful methods. Our airway segmentation method includes a novel wavefront definition, which allows a natural and efficient determination of branching structure and eliminates the need for the pruning and merging of branches with post-processing.
2. A novel method for automatically segmenting the lungs from clinical proton MRI data (section 3.3), based on region-growing with automatic threshold selection and lung contour refinement using level sets.
3. A novel, fully automated CT lobe segmentation method (section 4.3) offering the following improvements over existing methods:
   - Robust to small vessels crossing the lobar boundary;
   - Fast and computationally efficient, requiring no training or prior data;
   - Robust to incomplete fissures.

   This work was presented at the ISBI 2012 conference [8].
4. The first fully automated lobar analysis of hyperpolarised xenon MRI images. Our work builds on lobar analysis techniques for hyperpolarised helium-3 previously developed in our group by Plotkowiak et al. [9]. We make the following novel adaptions:
   - We replace semi-automated CT lobe segmentation with the fully automated method we introduce in chapter 4;
   - We replace manual MRI lung segmentation with the automated method we introduce in section 3.3;
   - We replace rigid registration with a nonlinear fluid registration which can better take into account the differences in lung inflation level between modalities;
   - We illustrate our pipeline by deriving preliminary results for ventilated volume percentages, mean ADCs and emphysema percentages at a lung and lobar level.

1.4 Clinical trial

The Oxford Biomedical Research Centre is developing state-of-the-art imaging techniques using hyperpolarised gas MRI (see section 2.3.3). The gas is inhaled by the subject prior to imaging and effectively acts as a contrast agent, allowing measurements of ventilation and gas diffusion in the lung. Data from a hyperpolarised ^{3}He study was previously provided to the Oxford respiratory collaboration (a multi-disciplinary group of researchers) by GE Healthcare and was analysed by Plotkowiak [9]. Although this trial showed promising results, ^{129}Xe offers a more promising clinical technology than ^{3}He because of its lower cost, wider availability and (unlike ^{3}He) its solubility in the blood stream, allowing additional measurement of gas transfer and blood flow (perfusion) [10]. This has led to the use of Xe-MRI as part of a clinical trial into COPD at the Oxford University Hospitals NHS Trust. The research in this thesis is aimed at analysing data from this trial.

In the trial, subjects first undergo baseline assessment and subsequently pulmonary function tests, CT, standard (proton) MRI and hyperpolarised xenon MRI imaging at regular intervals. Volunteers consist of COPD patients, plus a control group of non-COPD patients who require CT chest imaging for other medical reasons. ^{129}Xe gas is hyperpolarised in a machine at the hospital, and one litre is inhaled by the subject immediately prior to imaging. Xe-MRI imaging is performed using a standard clinical MRI scanner, using custom receiver coils which are worn by the subject as a flexible body vest. In the trial, Xe-MRI imaging initially consists of ventilation (airflow) and diffusion imaging. ^{129}Xe perfusion imaging, also known as the dissolved phase, will be performed later in the trial. The modalities, inflation levels and registration methods are summarised in Figure 1.2.

CT and MRI imaging are performed at different levels of lung inflation. For this reason we use non-rigid registration techniques in our analysis to align these images while taking account of the different inflation levels. CT imaging is performed at total lung capacity (TLC), which involves the subject holding their breath after maximal inhalation. The increased volume of gas in the lung improves the tissue contrast in CT images. Xe-MRI imaging is performed at breathhold at residual volume (RV) plus one litre of gas, which involves the subject breathing out normally, then inhaling one litre of hyperpolarised ^{129}Xe from a bag. Proton MRI imaging is carried out in the same scanning session and in the same manner as Xe-MRI, but using air instead of ^{129}Xe. We propose using a rigid registration to relate the proton and xenon MRI scans, based on the assumption that measurements at residual volume show good reproducibility. Brown et al. [11] measure CT lung volumes for 120 COPD subjects, computing an intraclass correlation coefficient (ICC) of 0.886. They also measure lobar volumes for a single subject, giving ICCs of 0.85–0.93.

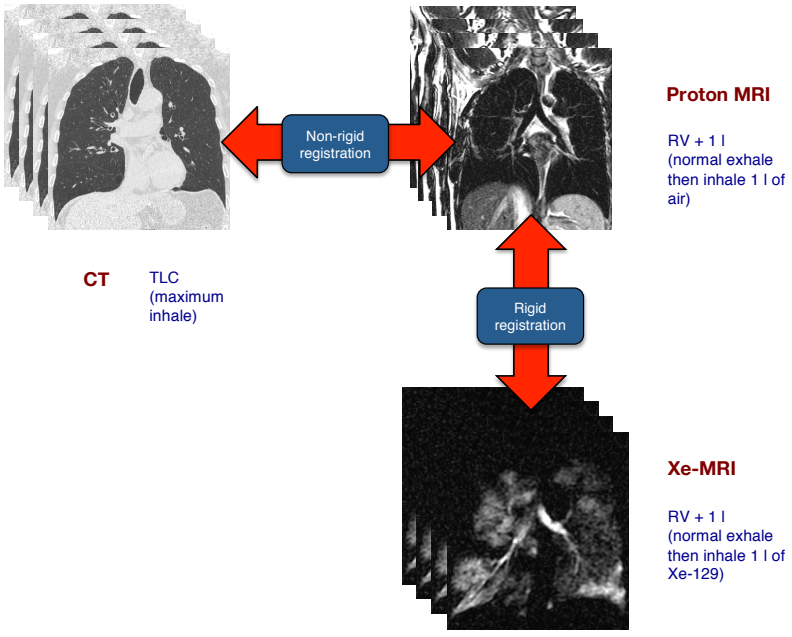

Figure 1.2: Schematic diagram of the imaging modalities available from the clinical trial. The lung inflation levels are shown for each modality. A non-rigid registration is used between the proton MRI and CT images, while a rigid registration is used between the proton and xenon-129 MRI.

Additional challenges and their mitigation

The clinical trial began imaging patients in January 2013, following a number of unexpected challenges. At the time of writing, four patients in the control group have been imaged, with further control and COPD patients to follow in 2013.

Originally the hyperpolarised gas imaging was to have used ^{3}He. However, this gas effectively became unavailable following a global shortage of ^{3}He [12]. To address this issue, the trial switched to ^{129}Xe. This gas is less well established and required some changes to hardware and protocols. However, ^{129}Xe offers a number of advantages over ^{3}He (see section 2.3.3), and its abundance makes it a viable gas for clinical use.

A second challenge was the withdrawal from the trial of the industrial partner GE Healthcare, due to global changes in the company's research priorities. Despite this, GE Healthcare continued to provide the polarising equipment. Dr. Jim Wild, an experienced researcher in hyperpolarised gas imaging from the University of Sheffield, has provided assistance in the development of the required MRI sequences and protocols. Further delays occurred following the sudden and untimely death of the PI of the project, and construction and legal issues regarding the building where the scanning equipment was installed. These issues have now been resolved.

As a result of these delays to the start of the trial, our final chapter (chapter 5) describes preliminary results, based on analysis of all the data available. To mitigate the effects of these delays on this work, we obtained a range of additional data (section 1.5) to help in the development and validation of our algorithms prior to the trial data becoming available. This data includes CT images from the Trust, and ^{129}Xe images from a previous patient trial [13] kindly provided by Dr. Bastiann Driehuys at Duke University. The protocol used to obtain this data closely matches that of our trial.

1.5 Software

The software developed during this research is part of the Pulmonary Toolkit (PTK). This is an open-source software suite, providing a library of pulmonary image processing algorithms, a framework for the rapid prototyping development of new algorithms and a graphical user interface (GUI) for simple interactive use. An application programming interface (API) allows flexible integration into existing software and supports a non-graphical mode.

The Pulmonary Toolkit is written in object-orientated Matlab, with some optimisations written as Mex files in C++. This software is available under an open source licence at the following URL:

http://code.google.com/p/pulmonarytoolkit/

1.6 Data

This thesis uses data from the Oxford University Hospitals NHS Trust, including images from the clinical trial (section 1.4), and data from a previous study into Xe-MRI at Duke University [13], which the authors have kindly given us permission to use. All data were obtained with the appropriate ethical approval for research use. Imaging parameters are summarised in Table 1.1.

For the CT lung, airway and lobe segmentation methods in chapters 3 and 4, images were retrospectively obtained for lung cancer patients at the Churchill Hospital, Oxford. The imaging protocols are the same as for the clinical trial.

For the proton MRI lung segmentation method in section 3.3, results are presented using proton MRI images from the Duke data. The protocols used

for these images are similar to those used in the Oxford trial. In particular, for both trials, the resolution and slice spacing of the proton MRI have been chosen to match that of the xenon MRI, in order to help with registration.

The regional image analysis of chapter 5 is performed using CT, proton MRI and xenon MRI data from both the Oxford and the Duke trials. Four control subjects are analysed from the Oxford trial, and one COPD subject from the Duke trial. Both trials have broadly similar protocols, although there is some variation in the xenon MRI scanning parameters.

Due to the pioneering nature of these trials, not all imaging sequences were successful, and some of the data present additional challenges. Ventilation images (showing the distribution of xenon in the airways) were produced for four subjects. Diffusion-weighted (DW) images (see section 2.3.3) were produced for three subjects (two from Oxford and one from Duke). Diffusion-weighted imaging is performed at a lower resolution than ventilation imaging, and an unweighted image is captured in addition to the diffusion-weighted image. This allows computation of apparent diffusion coefficient (ADC) maps (see section 5.3.3).

In the Oxford trial, inflation levels were not correctly matched between proton MRI and xenon MRI for one subject. The flexible body coil was too small for one subject, resulting in reduced image quality. We discuss these issues and how they were mitigated further in chapter 5.

For the Duke data, the most anterior and posterior slices were not scanned for diffusion imaging due to acquisition time constraints. Additionally, prior to being saved in DICOM format, the Duke proton and xenon MRI images were interpolated by the reconstruction algorithm of the scanner to a higher resolution of 256 × 256.

Table 1.1: Imaging parameters for the data used in this thesis. Results are grouped according to the modality of the images and the institution which provided them. MRI imaging includes standard proton (H-MRI), xenon ventilation (Xe-MRI) and xenon diffusion-weighted (Xe-MRI DW) images.

Institution		Duke				Oxford			
Modality		**CT**	**H-MRI**	**Xe-MRI**	**Xe-MRI DW**	**CT**	**H-MRI**	**Xe-MRI**	**Xe-MRI DW**
Slice orientation		Axial	Coronal	Coronal	Coronal	Axial	Coronal	Coronal	Coronal
Slice thickness	mm	0.6	15.0	15.0	15.0	0.6–2.5	8.0–20.0	15.0–20.0	30.0
Slice spacing	mm	0.6	15.0	15.0	15.0	0.6–2.5	8.0–20.0	15.0–20.0	30.0
Pixel size	mm^2	0.7 × 0.7	3.1 × 3.1	3.1 × 3.1	6.2 × 6.2	0.7–0.9 × 0.7–0.9	1.9 × 1.9	4.2 × 4.2	12.5 × 12.5
Resolution	pixels	512 × 512	128 × 128	128 × 128	64 × 64	512 × 512	256 × 256	96 × 96	32 × 32
X-ray tube current	mA	400				60–697			
Exposure time	s	0.50				0.50–0.62			
Peak kilovoltage	kV	120				120			
Flip angle	°		45	18	17		75–90	90	90
Repetition time	ms		2.8	7.9	12.0		4.2–2000.0	19.1	18.0–20.0
Echo time	ms		1.2	2.0	8.1		1.8–103.1	3.6–3.7	11.5
DW gradient maximum	$G{\cdot}cm^{-1}$				3.2				3.3
DW gradient ramp time	ms				0.5				0.3
DW gradient duration	ms				2.9				3.6
DW gradient separation	ms				5.0				5.0
DW b-value	$s{\cdot}cm^{-2}$				12.0				25.3

2 Background

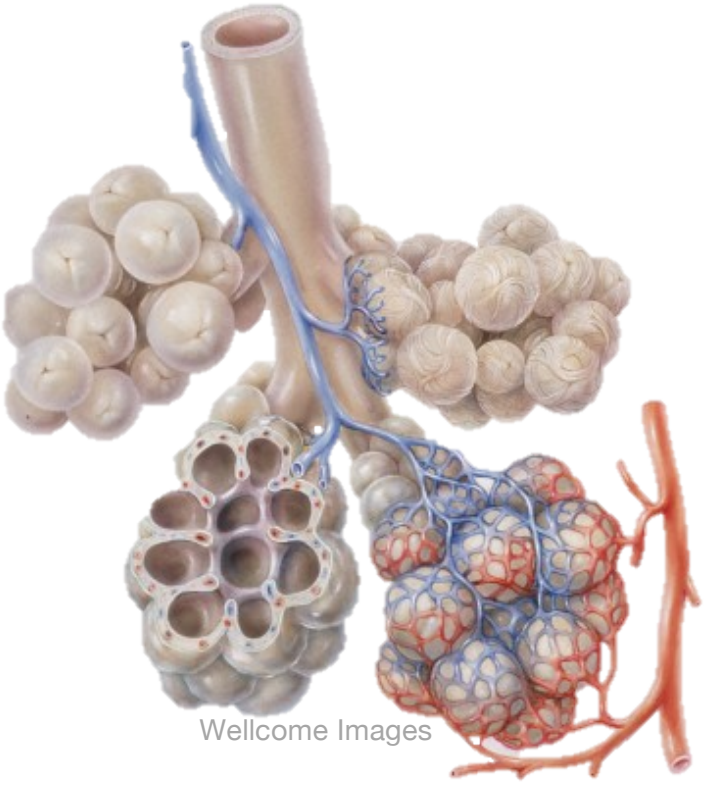

Chapter 2: Background



In this chapter we present background material concerning the lungs and lung disease. We first give an overview of lung physiology (section 2.1) and a discussion of COPD (section 2.2). We then discuss disease assessment methods (section 2.3), including hyperpolarised gas imaging (section 2.3.3), which is being used as part of the clinical trial described in section 1.4.

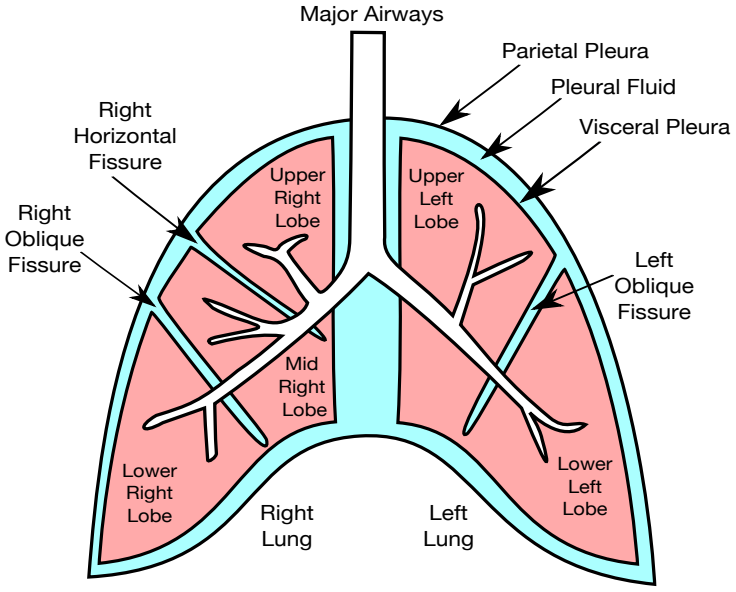

Figure 2.1: Schematic diagram of the respiratory system. The left lung is divided into two pulmonary lobes and the right lung into three. The lobes are supplied by the airways (bronchi) and separated by the pulmonary fissures. The lung is bounded by two membranes (pleura) separated by pleural fluid, which extends into the fissures. Pulmonary anatomy can vary between individuals, and incomplete fissures are common.

2.1 Lung physiology

The principal function of the lung is gas exchange: extracting oxygen from air and moving it into the blood, and removing carbon dioxide from the blood and expelling it from the body. The concentrations of oxygen and carbon dioxide in the blood are precisely controlled by the lungs, and the oxygenated blood is supplied to cells throughout the body by the circulatory system. The lung exhibits sophisticated regulatory mechanisms to match ventilation and perfusion as the flows of air and blood change in response to the demands of exercise. The lung also performs secondary functions such as metabolism, blood filtering and acting as a blood reservoir for the systemic circulation [14]. However, gas exchange is its primary function and the one most relevant to lung disease.

Gas exchange occurs in over 400 million air sacs known as *alveoli* (Figure 2.3). Air is supplied and removed through the airways (*bronchi*), while blood continually flows along the walls of the alveoli, supplied by the pulmonary arteries and removed by the pulmonary veins. The airways, arteries and veins each form a tree-like network of bifurcating elastic tubes, becoming increasingly shorter, narrower and more numerous as they approach the alveoli located in the last few generations. These networks follow similar paths and branching structures, allowing the lung to be described in terms of regions that are anatomically and functionally distinct. Initially the airways branch into the left and right lung, then into lobes, which are major subdivisions separated by fissures (Figure 2.1). The lungs are not symmetric, due to the shape and location of the heart. The left lung has a smaller volume, consisting of two lobes (left upper and lower) while the right lung consists of three lobes (right upper, middle and lower).

The lungs are bounded by two layers of membrane (pleura) separated by a thin layer of lubricating pleural fluid. The inner membrane (visceral pleura) folds into the lung between the lobes, forming the pulmonary fissures (Figure 2.1). Three main fissures separate the five lobes: the major (oblique) and minor (horizontal) fissures in the right lung; and the major (oblique) fissure in the left lung. The fissures contain pleural fluid and physically separate the lobes while permitting some relative movement. Lobar anatomy can vary significantly between individuals; fissures are often incomplete (not fully separating the lobes) and additional fissures may be found within the lobes [15].

The pulmonary lobes further divide into segments and then lobules. Segments correspond to regions perfused by the same airway (the segmental bronchi). Further down in the branching structure, the regions where gas exchange takes place are known as the *acini* (Figure 2.2). These independent respiratory regions consist of several generations of small airways with attached alveoli in formations which are traditionally described as grape-like, although are more accurately described as polyhedral (Figure 2.3) [16].

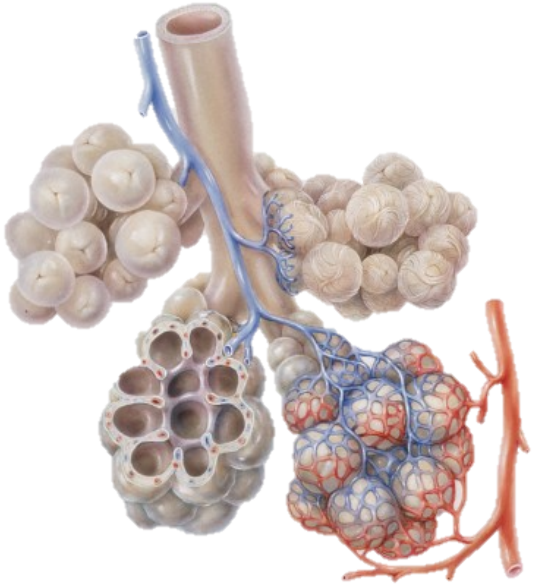

Figure 2.2: Idealised diagram of pulmonary acini, showing the small airways, vessels and alveoli. The acini are independent respiratory regions consisting of alveoli which are commonly illustrated as grape-like formations. The alveoli are air sacs surrounded by a capillary mesh. (Medical Art Service, Munich / Wellcome Images.)

2.1.1 Alveoli

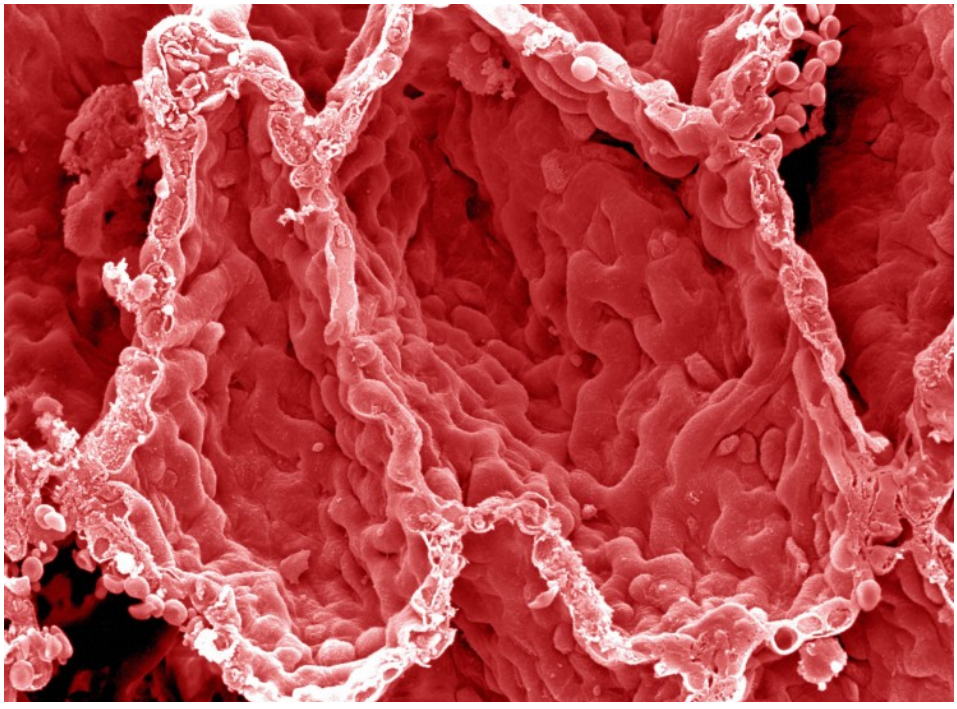

Figure 2.3: Scanning electron micrograph of alveoli in the lung. The polyhedral shape of the alveoli can be seen. The colouring is artificial. (David Gregory and Debbie Marshall / Wellcome Images.)

The alveoli are irregularly shaped air sacs around 0.3 mm in diameter mostly surrounded by a blood capillary mesh. The huge number of alveoli provides a large surface area for gas exchange (between 50 and 100 square metres) with an extremely thin membrane between the air and blood (0.1 to 0.5 microns) [17]. This permits very rapid gas diffusion. Passive diffusion is driven by the partial pressure gradients of the gases, which are maintained through the continual supply of oxygen and removal of carbon dioxide during breathing.

The capillaries have a diameter of approximately 10 μm, only wide enough for the passage of a single red blood cell. Each red blood cell traverses one or more alveoli and spends under a second in the capillary mesh. During exercise this time may decrease to 0.25 s, as the blood flow rate increases. Normally this is sufficient to achieve equilibration of oxygen and carbon dioxide, but in

lung disease patients the wall between the alveoli and capillaries may be thicker, causing oxygen diffusion to take longer. Hence patients with severe lung disease may experience symptoms during normal breathing, but for milder conditions symptoms only appear during exercise. Carbon dioxide transfer occurs much more rapidly than oxygen due to its higher solubility [14].

The volume of air in the alveoli (around three litres) far exceeds the volume of blood in the capillary mesh (around 70 ml). Nevertheless, blood flow and airflow are regulated to achieve an approximate matching of ventilation to perfusion, typically around five litres per minute. The alveoli also contain cells secreting surfactant. This lowers the surface tension of the liquid layer lining the alveoli, which increases lung compliance and may play a role in preventing alveolar collapse [14].

2.1.2 Airways

The airways consist of a tree-like system of branching tubes (*bronchi*), beginning with the trachea, which divides at the *carina* into the left and right main bronchi, serving the left and right lungs. These further divide into lobar and then segmental bronchi (Figure 2.4). This branching continues through bronchioles and respiratory bronchioles down to the *alveolar ducts*. The entire airway branching tree consists of approximately 24 generations in all, although only the smallest airways (approximately the last seven generations) form the *respiratory zone*, containing alveoli where gas exchange takes place. The larger airways form the *conducting airways*, which do not contain alveoli and therefore do not participate in gas exchange. Their purpose is to supply air to the respiratory zone, helping to warm the air as it travels to facilitate gas exchange. The majority of the lung is occupied by the respiratory zone, even though the individual acini within this zone are only millimetres in length.

The pulmonary airways are internally lined with a thin liquid film consisting of several layers. Adjacent to the airway wall is an aqueous phase, above which is a layer of surfactant. The conducting airways also have a layer of mucus, which may extend to the smaller airways for certain pathological conditions. Surfactant helps reduce the surface tension and distribute the liquid film across the airway surface. The film lining is believed to facilitate the transfer of gases and protect against infection and oxidation. It is also known to have a significant effect on the mechanical properties of the lung, particularly its elastic recoil [18].

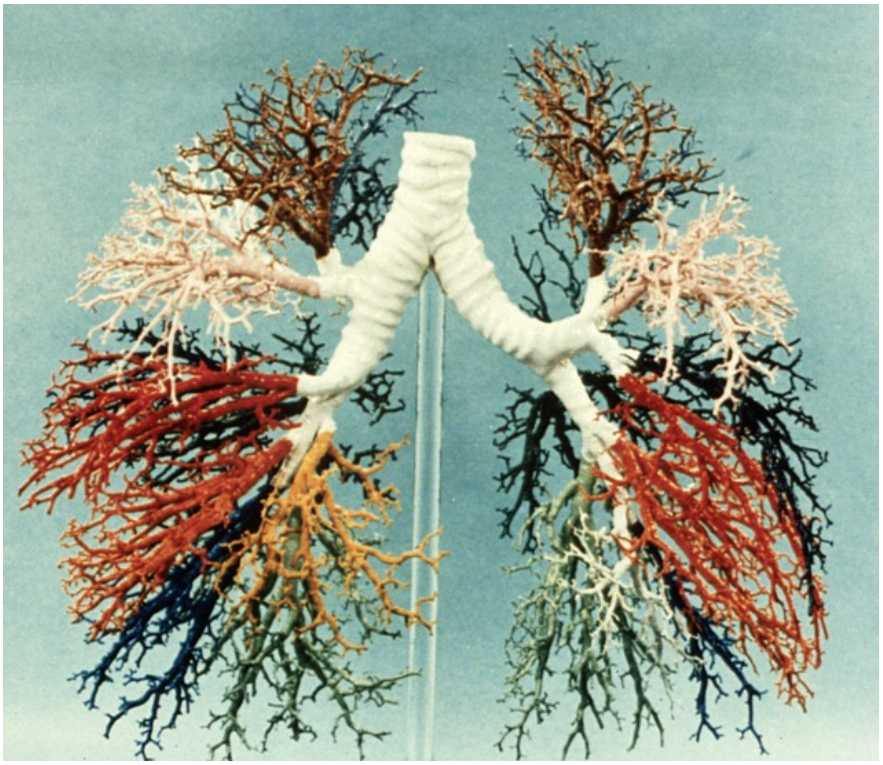

Figure 2.4: Model of the bronchial tree, with each airway coloured according to the bronchopulmonary segments it serves. (Wellcome Images.)

2.1.3 Breathing

The process of breathing involves the interaction of various muscles, which contract the diaphragm and raise the ribs. This increases the cross-sectional area of the thorax, causing a change in pressure. The structure of the lungs means that even this small change causes rapid airflow into the lungs. By contrast, expiration is a passive process for normal breathing, during which the muscles relax and the elastic lung returns to its original form. With increased breathing rates such as during exercise, expiration becomes active, recruiting additional muscles which act to raise the diaphragm and lower the ribs.

The lung tissue is elastic and its *compliance* describes the relationship between lung pressure and volume. The compliance changes with lung disease; fibrous tissue formation decreases the compliance, whereas emphysema leads to an increase. Compliance is increased by the effects of the pulmonary surfactant, which reduces the surface tension of the alveolar lining layer [14]. This is one of the factors making it difficult to measure lung tissue properties from excised lungs.

A typical lung can hold up to six litres of air, but the volume entering during normal breathing (*tidal volume*) is around 500 ml. This increases dramatically during exercise. Breathing is controlled by the brainstem, which the cortex can override to provide voluntary control. Ventilation drive is increased when chemoreceptors in the brain detect an increase in carbon dioxide levels in the bloodstream. Peripheral chemoreceptors in the vasculature can also detect a fall in oxygen concentration, although their effect on ventilation is more limited.

2.1.4 Blood vessels

The pulmonary circulation consists of two branching networks of blood vessels (Figure 2.5). The arterial tree carries deoxygenated blood from the right ventricle of the heart down to the walls of the alveoli. Here, gas exchange takes place in the dense capillary mesh of the walls. The oxygenated blood is then transported by the venous tree to four veins draining into the left atrium of the heart. From there it is circulated through the body by the systemic circulation. In the lungs, the arteries run close to the airways, with the veins following up to the lobules. The veins then separate and pass between the lobules while the arteries and bronchi travel down the centre of the lobules. The proximity of the arteries to the bronchi raises the temperature of the inspired air before it reaches the alveoli.

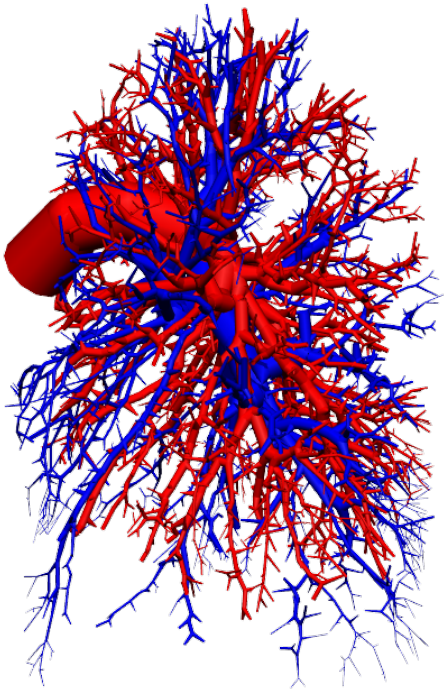

Figure 2.5: Computational model of the pulmonary blood vessels of the left lung. The vessels have been extracted from CT images and manually separated into the arterial and venous trees. (Courtesy of Kelly Burrowes.)

2.2 Chronic obstructive pulmonary disease (COPD)

A broad range of conditions are classified as lung disease, including lung cancer, pleural disease, pulmonary vascular disease and respiratory infections. One of the largest disease groups is *obstructive disease*, which includes asthma and chronic obstructive pulmonary disease (COPD). Obstructive disease occurs due to the premature closing or other blockage of airways, limiting the flow of air from the lung. This limit cannot be exceeded even with increased expiratory effort, leading to reduced lung function. While asthma leads to temporary obstructions caused by hyperresponsive airway muscles, COPD is characterised by permanent irreversibility, as the name implies.

The main conditions classified as COPD are emphysema and chronic bronchitis, and the two often occur together. Emphysema is characterised by the destruction of alveolar tissue, resulting in a loss in elasticity which causes the alveoli and small airways to collapse during expiration. The loss of tissue results in larger alveoli, with a decreased surface-to-volume ratio and a decreased total surface area for gas exchange. Chronic bronchitis is an infection-related inflammation of the airway linings with associated production of mucus, causing a restriction to airflow. Smoking is the main contributory factor in the development of COPD, but there are a number of other factors. For example, *alpha 1-antitrypsin deficiency* is a genetic disorder that increases the likelihood of developing emphysema.

Another class of lung disease is *restrictive disease*, which includes pulmonary fibrosis. Restrictive disease causes a reduction in compliance of the lung, limiting lung expansion. This results in a decrease in total lung capacity (TLC) and makes breathing more difficult. Sufferers of restrictive disease experience

difficulties breathing in enough air, while those with obstructive disease find it difficult to exhale fully before their next breath cycle. This makes it possible to distinguish obstructive and restrictive diseases using functional testing. However, while obstructive and restrictive disease are distinct classes of disease with different treatments, their causes are often related and it is common for patients with COPD to exhibit symptoms of both disease types.

2.3 Clinical assessment of COPD

Historically, COPD patients were classified as “pink puffers” or “blue bloaters” as a way of distinguishing between emphysema and chronic bronchitis as the primary cause of COPD. Pink puffers are emphysema sufferers whose breathing rate increases to compensate for the decreasing surface area for gas transfer. Thus they maintain their ventilation-perfusion ratio at the cost of hyperventilation and increased exertion – hence the name “pink puffer”. Blue bloaters suffer from chronic bronchitis and their ventilation decreases in response to the airway obstruction. This results in hypoxemia (oxygen deficiency in the blood) which, combined with the inability to fully exhale, lends the term “blue bloater”.

While these terms are still used, they refer to the extreme ends of the COPD spectrum and are of limited clinical value [19]. Current clinical assessment is based on pulmonary function tests (section 2.3.1) and CT imaging (section 2.3.2), while new techniques such as hyperpolarised gas MRI are beginning to emerge (section 2.3.3).

2.3.1 Pulmonary function tests (PFTs)

Currently, most clinical measures for lung disease are based on pulmonary function tests (PFTs). These typically measure global lung function through controlling inhaled gases and analysing those that are exhaled. A standard technique is spirometry (Figure 2.6). The most common measurements taken are forced expiratory volume in one second (FEV_1), the volume of air that can be expelled in one second, and forced vital capacity (FVC), the total volume of air that can be exhaled. Both measurements are made after maximum inspiration. Normally, FEV_1 is about 80% of FVC [14]. A reduced percentage (known as FEV_1%) is characteristic of obstructive lung disease, due to the limitation to the flow of air from the lung. A value of less than 50% may be characterised as severe COPD [19]. The total lung volume may also be abnormally large in these patients. FEV_1% is commonly reported as FEV_1% predicted, which is FEV_1% normalised by a typical value for a healthy subject of similar age, weight etc. Restrictive disease is characterised by a reduction in both FEV_1 and FVC. Lung disease patients frequently exhibit symptoms of both obstructive and restrictive disease.

Other tests measure different lung parameters. For example, D_{LCO} (diffusing capacity for carbon monoxide) is used as a measure of gas transfer between the alveoli and blood. Body plethysmography (Figure 2.6) is used to estimate functional residual capacity (FRC), the volume of gas remaining in the lung after expiration. The subject sits in an airtight box and the change in box pressure is measured as their lung expands when attempting to inhale.

Despite their widespread use in clinical practice, pulmonary function tests on their own are insufficient for diagnostic needs. FEV_1 is a poor predictor of dyspnoea (shortness of breath), survival, hospitalisation frequency and functional capacity [3]. Single measurements may not be the optimal way to assess the effects of treatments on the different components of COPD [20].

PFTs only measure global lung function and cannot detect the regional changes which are characteristic of lung diseases. For this reason medical imaging is increasingly being used in the diagnosis of lung disease.



Figure 2.6: Spirometer and plethysmograph. In the spirometer (left), the subject makes forced inspiratory and expiratory efforts while breathing through the mouthpiece, and the spirometer measures flow rate and the composition of exhaled gases. This can calculate a variety of global functional measurements, the most common of which are forced vital capacity (FVC) and forced expiratory volume in one second (FEV_1). In the plethysmograph (right), the subject sits in an airtight box and the change in box pressure is measured as their lung expands when attempting to inhale. This can be used to measure functional residual capacity (FRC).

2.3.2 CT imaging

Computed tomography (CT) (Figure 2.7) is the modality of choice for lung imaging [7]. It is widely used in patient assessment, treatment planning and follow-up, and with its high spatial resolution (typically < 0.7 mm in each dimension) and good signal-to-noise ratio, it offers a wealth of regional measurements that may relate to pulmonary disease. While the voxel size is too large to resolve individual alveoli, areas of emphysema damage are apparent from their decreased attenuation. CT values are normally expressed in Hounsfield units (HU), a radiodensity scale where, by definition, air and water have values of -1000 HU and 0 HU respectively at standard pressure and temperature (STP). Radiodensity and mass density are closely related, and to a good approximation, pulmonary tissue densities can be computed in g/l by adding 1000 to the radiodensity values in Hounsfield units.



Figure 2.7: Reconstructed coronal slice from a 3D multi-detector CT (MDCT) dataset showing the lungs. Many pulmonary structural features are detectable in CT, such as airways, vessels and fissures. CT can also visualise many types of pathology.

A great deal of research has been carried out into automated assessment of COPD from CT images [21]. Two distinct COPD characteristics which can be investigated using CT are emphysema (destruction of the lung tissue) and bronchiolitis (inflammation of the small airways). Emphysema causes decreased attenuation of CT measurements due to the enlarged airspaces [22]. Voxels with decreased attenuation have been shown through histological studies to correspond well with emphysema [23], as well as showing a correlation with pulmonary function test measurements [24]. Emphysema can be quantified by computing the percentage of voxels which fall below a simple threshold such as –910 HU [25], or by identifying the value (in HU) on the frequency distribution curve below which fall 15% of the lung voxels (percentile density) [26]. Not all authors exclude the airways from analysis, which can lead to the airway lumen (interior) being classified as emphysema due to its low density.

These methods detect the presence of emphysema but provide limited information on severity. A variety of different methods for assessing emphysema severity have been proposed [27]. One method is to estimate the volume of emphysema lesions. Voxels with Hounsfield values below an emphysema threshold are grouped into connected components, with the number of voxels in each component determining the volume of that emphysema lesion. By plotting the lesion volume against the cumulative number of lesions exceeding that volume, a power law relationship is obtained [28]. However, studies have mixed success when validating against morphometric data and patient outcome [23].

Authors are increasingly looking at regional heterogeneity, for example as a way of phenotyping COPD to help with surgical decisions. Revel et al. [29] calculate the proportion of emphysema in each lobe for patients selected for lung volume reduction surgery (LVRS) and find the majority to have significant

heterogeneity between lobes. Automated quantification shows good agreement with visual scoring estimates performed by a radiologist. Uppaluri et al. [30] carry out a texture-based feature analysis to characterise emphysema. They find an improvement over histogram analysis but poor correlation with pulmonary function tests. Chabat et al. [31] use a supervised 2D Bayesian classifier combining texture and shape features to differentiate types of obstructive disease. Xu et al. [32] extend this work into 3D, using it to classify different stages of emphysema. Charemza et al. [33] and Boehm et al. [34] independently developed methods for characterising emphysema and fibrotic tissue using Minkowski functionals, which describe the shape and topology of binary regions segmented by applying an image threshold. Further methods for the supervised classification of fibrotic lung tissue have been developed by Hoffman et al. [35] and Zavaletta et al. [36].

Emphysema is typically accompanied by airway obstruction. While chronic bronchitis is associated with the larger airways, most obstructive effects occur in the smaller airways not visible on a CT scan [37] and are therefore not directly measureable. This is one of the reasons why CT-derived measurements do not always correlate with functional measurements. However, Nakano et al. [38] find a relationship between large and small airway dimension changes, suggesting that large airway wall thickening may be a surrogate marker for pathology in the small airways. Hasegawa et al. [39] find correlations between pulmonary function tests and airway luminal area and wall thickening. Diaz et al. [40] find a correlation between the number of airways detectable in CT (counted manually) and emphysema measurements and indicators of clinical outcome in smokers. Aziz et al. [41] measure a number of airway features and find a correlation between FEV_1 and wall thickness, but not with bronchial dilation. However, their airway measurements are a manual grading based on comparison with adjacent pulmonary arteries. The same study measures emphysema heterogeneity by dividing the lung into

core and rind regions, and finds a higher proportion of emphysema in the core region correlated with a reduction in D_{LCO} (diffusing capacity for carbon monoxide).

Unlike pulmonary function tests, medical imaging offers a regional approach to disease detection. Standard CT is insensitive to mild emphysema and is limited as a predictor of lung function [42], although using xenon-enhanced CT it is possible to measure ventilation and perfusion [43]. Techniques such as Spect and scintigraphy [44] provide functional information but suffer from poor signal-to-noise ratio and, like CT, have the drawback of exposing the subject to ionising radiation. This has motivated further research into non-ionising imaging techniques that can provide functional information such as hyperpolarised gas MRI.

2.3.3 Hyperpolarised gas MRI

In recent years there has been increasing interest in hyperpolarised gas imaging [5]. Prior to imaging, the subject inhales a gas such as helium-3 (^{3}He) or xenon-129 (^{129}Xe) that effectively acts as a contrast agent. In its normal state the gas would not be detectable on MRI scans, but optical pumping [45] techniques are used to artificially increase the polarisation of the gas such that a detectable signal is produced even with low densities of gas [6]. Spin-density measurements allow for imaging of static ventilation in 2D and 3D, highlighting regions of poor ventilation, which may be an indication of obstructive diseases such as asthma [46]. Dynamic ventilation imaging uses various techniques such as interleaved spiral k-space acquisition to allow for imaging of inhalation and forced exhalation, which can, for example, highlight gas trapping in asthma [47].

Using different MRI pulse sequences, an alternative imaging method known as diffusion-weighted (DW) imaging can be used, which is a particularly useful measure for assessing alveolar structure. Imaging is performed during breath hold, and in a healthy lung the gas in the lung parenchyma is highly restricted within the alveoli, resulting in a low diffusion measurement. Increased gas diffusion in a region indicates less containment of air by the alveoli, which may indicate damage to the alveolar tissue due, for example, to emphysema. Although the resolution of the diffusion imaging (typical pixel size around $6 \times 6\ \text{mm}^2$) cannot resolve the individual alveoli, an averaged diffusion measurement over a voxel can be determined, known as the *apparent diffusion coefficient* (ADC) (Figure 2.8), also referred to as the *effective diffusion coefficient* [48]. ADC measurements appear to correlate with regional structural changes caused by disease [49] and show a significantly increased regional inhomogeneity in COPD and asthma sufferers [50]. For COPD this can be explained by the non-uniform destruction of alveolar tissue, but the precise mechanisms linking asthma and ADC changes are uncertain [50].

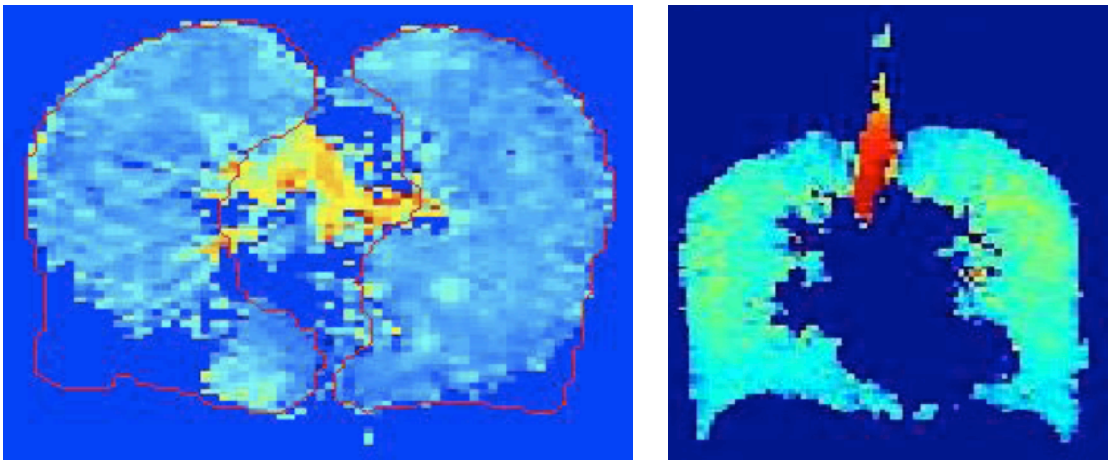

Figure 2.8: Ventilation and diffusion images from hyperpolarised helium MRI. Left: Ventilation helium MRI image with the lung contour superimposed (red line). A region of impaired ventilation is clearly observed in the lower left corner. Right: Diffusion imaging showing ADC values of a healthy subject. Higher diffusion can be observed at the large airways. (Courtesy of Michal Plotkowiak.)

The reproducibility of some regional measurements is uncertain. Mathew et al. [51] measure changes in ADC values and ventilation defects over a seven-day period. Ventilated defect volumes are manually segmented from the central coronal slice of each dataset, defined as the slice containing the carina. Mean ADC values are computed for this central slice and for the whole lung. The ADC measurements are found to be reproducible but ventilation defect measurements are significantly altered. Quantitative validation of gas MRI-based disease assessment is ongoing, for example with histology [52], spirometry [53] and CT [54]. Some authors create computational models of the diffusion imaging process and simulate emphysema to predict its effects on imaging [55,56,57].

Until recently, the gas of choice for hyperpolarised imaging research was ^{3}He. However, ^{3}He is a limited resource, and in recent years a number of factors have led to a worldwide shortage and soaring prices [12]. With recent technological advances there is growing interest in xenon-129 (^{129}Xe) as a viable hyperpolarised imaging gas [58]. ^{129}Xe has some disadvantages compared to ^{3}He, in particular the reduced magnetic moment resulting in a lower spatial resolution and signal-to-noise ratio [5]. However, the relative abundance and lower cost of ^{129}Xe make it a more viable option for clinical use [10]. For this reason, ^{129}Xe was adopted in the clinical trial by the Oxford Biomedical Research Centre (section 1.4). Hyperpolarised ^{129}Xe makes use of existing clinical MRI scanners, with the additional equipment required being a polarising machine, a supply of ^{129}Xe gas and a removable body coil worn by the subject. Several groups have previously conducted ^{129}Xe trials on human subjects [13,59,60,61]. Safety studies have found ^{129}Xe imaging to be a safe technique with only minor transient anaesthetic side effects [13,59].

^{129}Xe has other advantages over ^{3}He. Kirby et al. [62] report that static ventilation defect measurements made using ^{129}Xe images are more sensitive

to COPD detection than those from ^{3}He, due to the different properties of xenon. A major advantage is that following inhalation, the hyperpolarised ^{129}Xe dissolves in the blood, and the dissolved phase can also be visualised using MRI [63]. This allows for measurements of perfusion. A healthy lung is characterised by a good matching between freshly inspired air and blood flow, known as the *alveolar ventilation-to-perfusion ratio* [5]. Respiratory diseases commonly result in a localised change of this ratio, for example a decrease in the case of asthma, or an increase due to damage to the alveolar wall tissue and capillaries caused by emphysema. Simultaneous measurement of ventilation and perfusion could allow computation of this ratio.

2.4 Discussion

While pulmonary function tests are currently the gold standard for assessing lung disease, they are poor predictors of patient outcome and limited in their ability to distinguish between disease phenotypes. Crucially, they provide global measurements of lung function that cannot reflect the regional characteristics of COPD. CT imaging is increasingly used in pulmonary assessment as it can provide regional structural measurements that correlate with COPD, in particular the change in tissue density resulting from emphysema. However, CT exposes the subject to ionising radiation and its diagnostic abilities are limited because it measures the structure rather than the function of the lung. Hyperpolarised xenon MRI is potentially a powerful diagnostic tool that offers safe, non-ionising imaging of regional lung function which complements pulmonary function tests and could be offered in hospitals at relatively low cost. To realise this potential, new quantitative methods are required to analyse these images. This thesis works towards this goal.

3 Lung and airway segmentation

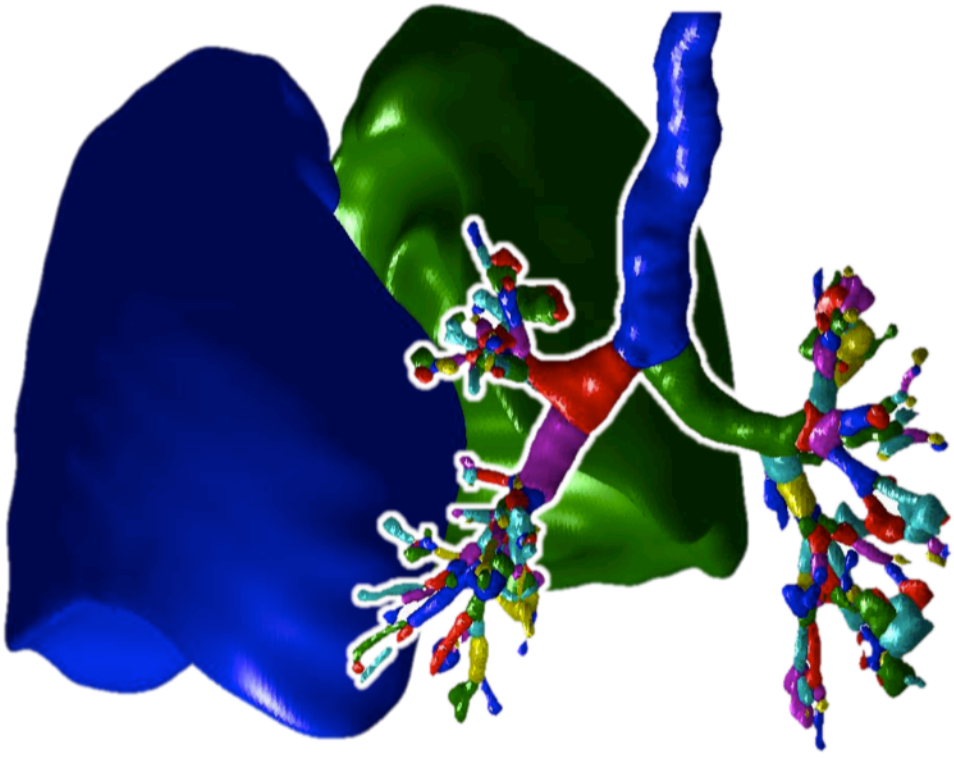

Chapter 3: Lung and airway segmentation



In this chapter we develop image segmentation algorithms that we require in later chapters. We present methods for automatically extracting three-dimensional representations of the lungs and airways from clinical CT data, and for extracting the lungs from proton MRI data. These lung segmentations, along with the lobe segmentations of chapter 4, will be used to build our regional image analysis methods in chapter 5. The airway segmentations are required by both the lung and lobe segmentations, and are also used to build personalised computational airflow models in our proposed future work (chapter 6).

In section 3.1 we introduce lung and airway segmentation. No standard method has emerged for lung or airway segmentation; there are still significant improvements to be made in existing algorithms, and they are not generally available for us to use in open source packages. For these reasons we have implemented our own routines. Our CT lung segmentation method is described in section 3.2. In section 3.3 we present a novel method for automatically segmenting the lungs from MRI data. Our airway segmentation algorithm is introduced in section 3.4. In addition to segmentation, our algorithm computes the radius and centreline for each bronchus, and determines which pulmonary lobe it serves. In section 3.5 we present conclusions for this chapter.

3.1 Introduction

Segmentation is the process of dividing images into regions. The aim of this chapter is to generate binary masks representing the lungs and the larger airways. Masks for the left and right lungs are generated from two imaging modalities: CT and proton MRI. Airway masks are generated from CT data.

The gold standard for clinical segmentation is the manual tracing of features by a pulmonary radiologist on each two-dimensional slice of a three-dimensional volume. This process has two major drawbacks. Firstly, it is an extremely laborious process, since datasets typically comprise hundreds of slices. A full lung and lobe segmentation may require several hours of work using this method. Secondly, the process has some degree of subjectivity, leading to variation in results between different radiologists (inter-observer variability) and even between results by the same radiologist (intra-observer variability). For these reasons there is considerable interest in developing segmentation methods that are entirely reproducible and require either no (fully automatic) or minimal (semi-automatic) interaction from the clinician.

Computational segmentation is a challenging problem and there is no single method that works in all situations. Methods are highly dependent on the modality, source data and the required application of the results. It is rare that published methods can be applied without modification, and it is difficult to compare existing methods since most groups do not publish either the source code or data used to generate their results. Challenge workshops such as Lola and Exact [64] can provide useful direct comparison of contributing methods, but the reference segmentations and source code are not generally made available. In developing our own fully automated segmentation algorithms, we have evaluated existing methods and adapted them wherever possible, introducing novel adaptations to overcome specific limitations.

3.2 Lung segmentation from CT

Lung segmentation is a fundamental part of pulmonary image analysis with important clinical applications in disease assessment [65]. For example, radiodensity analysis of the lung region in CT images can be used to assess emphysema [25], even though the alveolar structures are not visible at this resolution. Lung segmentation has two important applications in this work. The lung masks are used in our multi-modal image registration, which forms part of our regional analysis method (chapter 5). The CT lung boundaries are a fundamental part of our lobar segmentation (chapter 4), which is also used in the regional analysis.

No precise definition exists for the lung region [65], but in the context of clinical image analysis it is usually considered to be the volume enclosed within the pleura, including the lung parenchyma, airways and vessels, but excluding those parts of the major airways and vessels which extend beyond the pleural boundary (Figure 3.1).

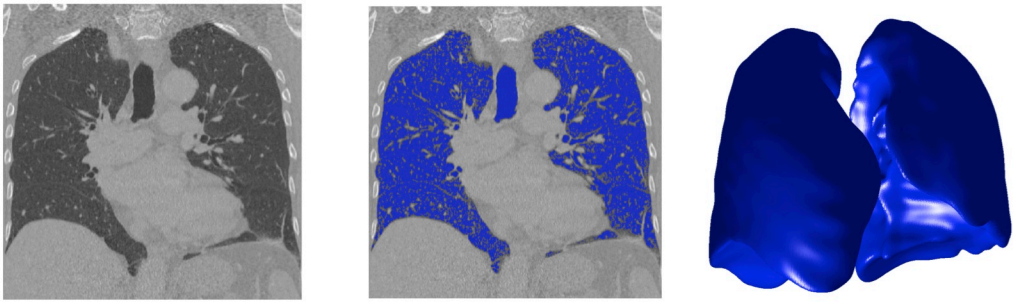

Figure 3.1: Illustration of automatic lung segmentation from CT. Lung segmentation depends on the definition of the lung region. Our definition includes the lung parenchyma inside the pleura and those parts of the vessels and airways which lie inside the lung volume. Left: A coronal slice of the CT lung image. Centre: The same image with a radiodensity threshold applied, showing areas of the lung parenchyma and airways highlighted in blue. Right: A surface rendering of the segmented 3D lung volume after removal of the airways. The 3D rendering has been smoothed in order to present a more realistic visualisation of the lung.

Clinical multi-detector CT scanners offer 3D resolutions below 0.5 mm in each dimension. In pulmonary CT scans, the lungs are readily visible, appearing as two large regions of low-density parenchymal tissue and air, bounded by high-density tissue at the pleura (see Figure 3.1). Most lung tissue appears in the range −910 to −500 Hounsfield units (HU), while the chest wall has much higher values [66]. High-density structures also appear within the lungs, including blood vessels, airways, fissures and pathological features. Nevertheless, the predominance of low-density "air-like" voxel values within the lung has led to threshold-based approaches being popular for CT lung segmentation, particularly since CT calibration ensures consistency in values within and between scans.

3.2.1 Previous work

For CT lung segmentation, there is a substantial body of previous work. A widely used lung extraction method is to search for a large connected region of air-like values within the image [66,67,68,69,70,71]. A threshold is chosen using analysis of the grey level histogram, and an initial lung region is then extracted either by region-growing from the airways [71] or by finding the largest connected component in the image [66]. When used with a fixed threshold, both methods produce the same result, a single region comprising the left and right lungs connected by the airway tree (see Figure 3.2). Removal of the airway tree is discussed below.

In a recent Miccai lung segmentation grand challenge workshop, these threshold-based methods gain the best overall scores [70,72,73,74], compared to other methods such as active shape methods [75] and supervised voxel classification [76]. However, while these methods are fast, robust and accurate for healthy lungs, they may perform less well in the presence of pathology [73,77,78]. In particular, pathological tissue near the lung borders may be

interpreted as being exterior to the lung. Results may be improved using knowledge-based analysis of the segmented lung border. In their lung segmentation algorithm, Pu et al. [78] extend the threshold-based method to include lung nodules adjacent to the lung border. Their method refines the lung border to remove convex regions which may indicate oversegmentation. Prasad et al. [79] adjust the threshold used for the region-growing method based on how close the resulting lung border matches the shape of the segmented ribs. Sun et al. [75] also use the ribs, this time to guide an active shape model. El-Ba et al. [80] use stochastic methods to segment the lungs from low-dose CT data.

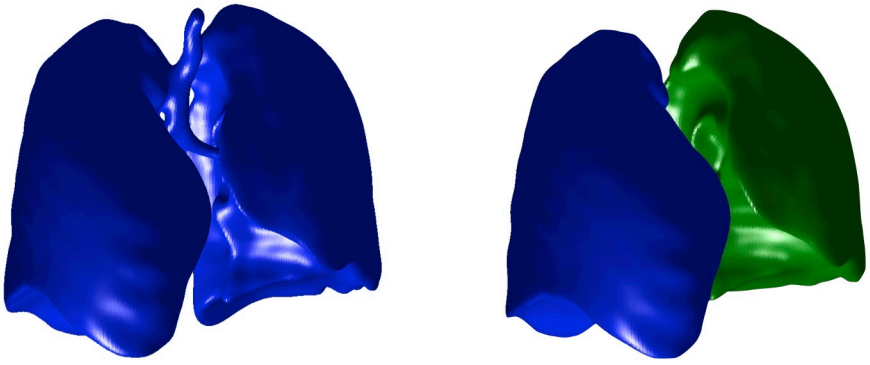

Figure 3.2: 3D surface renderings showing example results from automatic segmentation and labelling of the lungs. Left: The first stage in the lung segmentation process, showing the lungs and airways generated by applying an air threshold to the CT dataset, and then extracting the largest connected component. Right: The final lung segmentation results, after removal of the major airways, separation of the left and right lungs, and morphological smoothing.

Multi-atlas registration methods have also been applied to lung segmentation [77]. Van Rikxoort et al. [81] use a hybrid approach. A conventional threshold-based lung segmentation method is applied. Then, various shape and volume measurements are used to detect if a segmentation error occurred, such as in the presence of pathology. If these error checks fail, a multi-atlas lung registration is performed. This hybrid approach has two benefits: the atlas method is more accurate in the error cases, but the conventional approach is both more accurate and much faster in the non-error cases.

Separation of left and right lungs, and boundary refinement

Threshold-based approaches generally produce a segmentation where the left and right lungs are connected by the airways. In order to separate the lungs, the main bronchi must be identified and then removed from the segmentation. This can be done using the airway segmentation methods described in section 3.4 below. Frequently, removing the airways completely separates the segmentations of the left and right lungs. However, in some datasets they remain connected in regions where the physical boundary separating the lungs is particularly thin, such as near the apex (the top of the lung). This results from partial volume effects; the boundary is so thin compared to the voxel size that the CT values of the voxels around the border lie within the air-like threshold and can connect across the boundary. However, the border is still generally visible as a higher-density surface separating the lungs. Authors employ slightly different methods to find this surface and use it to separate the left and right lungs. Kuhnigk et al. [71] and Lassen et al. [72] use a watershed transform [82] seeded from automatically chosen points in each lung; this separates the original segmentation into two basins (lungs) at the higher-density ridge between them. Other authors [66,67,69,70,73,74] find the high-density boundary between the lungs using dynamic programming.

Morphological operations are typically applied to fill holes within the segmented lung volumes and to smooth the lung surfaces (see Figure 3.2). This helps to correct small errors at the boundaries.

3.2.2 Method

Our algorithm is summarised in Figure 3.3. We use a region-growing method with a global threshold to extract the initial lung region, similar to the method of Hu et al. [66]. The methods for left and right lung separation and boundary refinement are an adaptation of work by Kuhnigk et al. [71].

1. **Smoothing and noise reduction using a Gaussian filter**

   We initially apply a Gaussian filter with a standard deviation of 0.5 mm. This smoothing filter reduces image noise and suppresses small structures within the lung such as vessels which are not important to our segmentation.

2. **Threshold to identify low-density voxels**

   Next, we generate a mask of voxels with low-density values using the density range [−1024 HU, −400 HU]. This detects both air and low-density parenchymal tissue. We used a fixed threshold value of −400 HU because it lies between typical density values for tissue interior and exterior to the lung, and provides good results for the datasets we are analysing. A variable threshold approach such as used by Hu et al. [66] is potentially more robust to factors which affect mean lung density, such as lung pathology, imaging noise and scanner characteristics. However, we found a fixed threshold was adequate for our purposes, and the detection of the correct threshold is a challenge on its own and if it fails can lead to incorrect results. We also note that variable threshold approaches such as those used in [66] find the optimal threshold for separating voxels into two classifications (lung and non-lung). The optimal threshold for classification is not necessarily the optimal threshold for lung segmentation, because the threshold could be affected by high-density structures within the lung such as vessels.

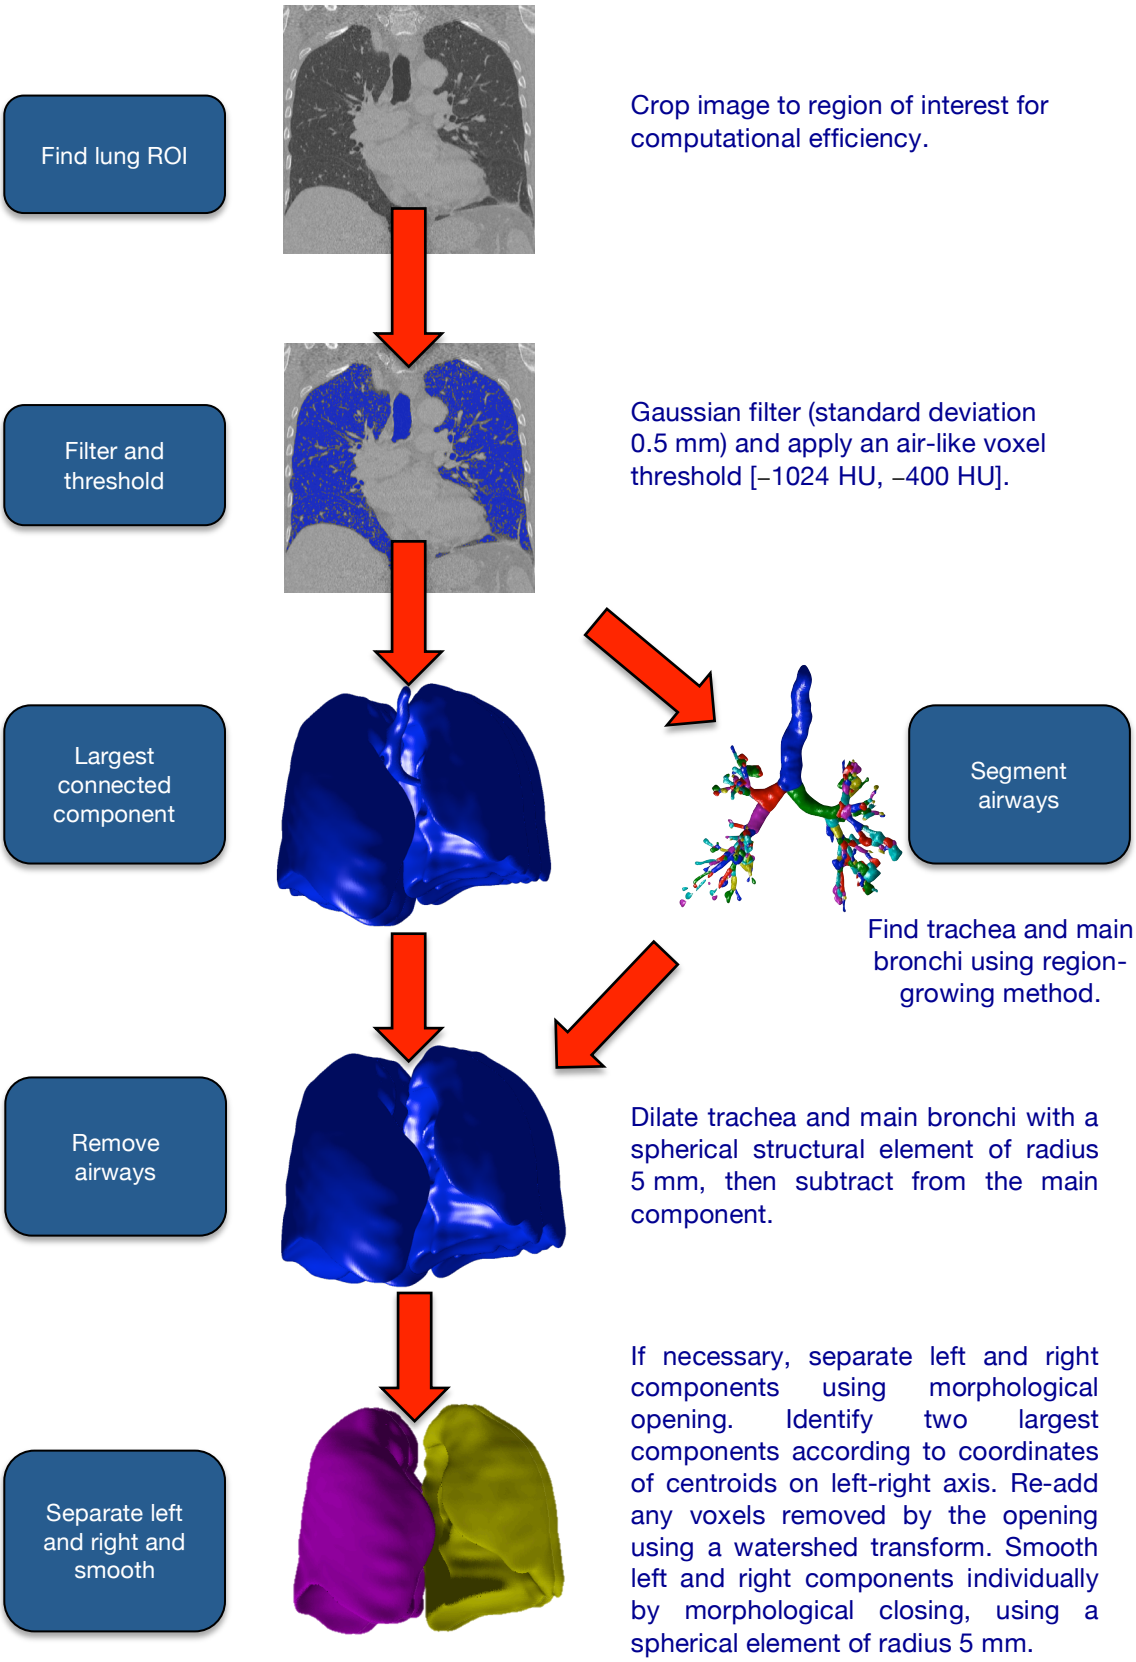

Figure 3.3: Schematic of the lung segmentation algorithm.

3. Identification of lung and airways

The threshold operation produces a mask which includes the parenchyma and airway interiors, but also air spaces exterior to the body as well as other low-density areas within the body which are not part of the lung. This is difficult to avoid with thresholding as the low tissue density within the lungs cannot easily be distinguished from air density. We can however identify and remove regions that are not connected to the lungs. We apply a 26-way connected-component analysis to the threshold mask to identify all the separate regions within the mask. We remove components outside the body, which we identify as being those touching the left, right, anterior or posterior borders of the image. Unlike other work [66], we permit connectivity to the superior or inferior borders. We have found that clinical scans occasionally do not completely extend to the full axial length of the lung, omitting small parts of the superior or inferior lung. This leaves lung mask components touching the superior or inferior borders, which would be incorrectly removed by other methods [66], resulting in a segmentation failure in these cases. Our approach avoids this problem.

4. Removal of airways

The resulting component generally includes the left and right lungs connected by the airway tree. To separate the lungs we first segment the airways using the methods discussed in section 3.4 below. We apply a morphological dilation operator to the trachea and main bronchi segmentation (including the interior lumen) using a spherical element of radius 5 mm, and remove this from the lung component. The dilation ensures the complete removal of these airways from the lung segmentation. This is necessary as the lung segmentation may include voxels within the airways which are not included in the airway segmentation, due to the different filtering and morphological operations employed in airway and lung segmentation.

5. Separation of left and right lungs

Removal of the airways will often separate the left and right lung masks, but in many cases they may still appear to be connected near the lung apex. This occurs when the membrane separating the lungs is close to or thinner than the voxel size, and a combination of partial volume effects and the image smoothing causes voxels on the separating membrane to fall within the lung density threshold and hence be included in the lung mask. To completely separate the lungs, we employ a method similar to that of Kuhnigk et al. [71]. We check if the lung mask consists of two separate large components. If this is not the case we apply a morphological image opening on the lung mask using spherical structural elements of increasing radius (1 mm, 2 mm, …, 10 mm) until the lungs are separated. The left and right lung masks are then selected as the two largest connected components and are labelled by examining the separation of the centroids of the components in the left and right direction. Finally, the voxels removed from the segmentation by the opening operation are re-added by applying a two-region watershed transform [82] on the CT density image with the left and right components as seed areas. This works because even with partial volume effects, the membrane tissue separating the lungs has higher intensity than the lung interior.

6. Smoothing of lung contour

Finally, we apply a morphological closing operation separately to the left and right lung segmentations with a structural element of radius 5 mm. The large size of this element fills any segmentation holes resulting from the removal of the airways and from vessels, airway walls and other high-density tissue within the lung. The closing operation also smooths the surface of the lung segmentation. The output from the connected-component analysis can be uneven at the lung surface due to partial volume effects, and the smoothing operation gives it a more realistic shape.

3.2.3 Computational efficiency

To improve the execution time of our algorithms without affecting the quality of the results, we introduce an additional step at the start of the segmentation algorithm which crops the lung image (and auxiliary images) to a lung region of interest. The resolution is not changed, but the reduction in the number of voxels significantly decreases the execution time of subsequent algorithms.

Lung region of interest (ROI)

We define the lung ROI as the minimal parallelepiped containing the lungs and bronchial tree. The use of an ROI provides a consistent frame of reference in which our algorithms can operate. By cropping our images to the region of interest, we also reduce processing and memory usage without affecting image quality. After determination of the ROI, a border of surrounding voxels is added, so that all points within the lung still retain access to a local neighbourhood of voxels after the cropping. This is important for those subsequent processing operations which rely on voxel neighbours.

The ROI is computed by first performing the CT lung segmentation (excluding the airway removal and smoothing) on a downsampled image. This is executed very quickly due to the small image size. The resulting low-resolution segmentation is used to define the ROI. The ROI is then used to crop the original image before performing the full segmentation. The cropping operation reduces the image size, but there is no loss of quality since there is no rescaling of the original image. This method works because the CT segmentation algorithm can operate at multiple scales, and the low-resolution mask is sufficiently accurate to define the ROI, which is only an approximation. Airway removal and lung contour smoothing is not required when computing the ROI.

3.2.4 Results and discussion

Figure 3.5 shows the segmentation results for an isotropic CT dataset. The methods we are using have been extensively validated in the literature (for example Hu et al. [66] or Kuhnigk et al. [71]; see section 3.2.1). We have tested our CT lung segmentation method on more than 50 datasets including those from the clinical trial, and verified by visual inspection that the segmentations are accurate. We discuss possible failure cases below.

3.2.5 Automated verification of results

Our automated segmentation algorithm is designed to be robust to common variations in lung images, such as the varying lung volume between subjects and the presence of emphysema. However, as with any automated method, it could potentially fail or prove inaccurate under certain conditions. If large imaging artefacts were introduced by subject motion or metal implants, this might lead to an apparent change of tissue density in parts of the image, disrupting the intensity-based segmentation algorithm. Another example is the change in the lung shape or appearance due to severe pathology. For example, Figure 3.4 illustrates a lung condition known as pulmonary empyema, where fluid collects in the lungs.

Our software allows the user to review the automated lung segmentation, so that if an error is found the segmentation mask can be manually corrected. To alert the user to possible segmentation errors, the volume of each lung is computed and automated checks are performed. The user is warned if the volume of either lung is below one litre, or if one lung has a volume more than 50 % larger than the other.

Additional checks could be used to alert the user to possible segmentation errors. The lung volume can be analysed with comparison to a shape model. A

simple example is a principal component analysis of the lung segmentation, to analyse the distribution of voxels relative to the main axes. The algorithm would alert the user if the lung shape does not match an expected lung shape subject to scaling and an error tolerance. Additionally, the two lung segmentations should be, to a coarse approximation, a mirror image of each other. This can also be verified by the shape model algorithm.

For the pulmonary empyema example in Figure 3.4, both the volume and shape model checks would alert the user to this unusual case. The automatic segmentation shown in Figure 3.4 has correctly detected the ventilated regions of the lung, but has excluded the pathological fluid regions. This may be considered correct depending on the intended application of the segmentation.

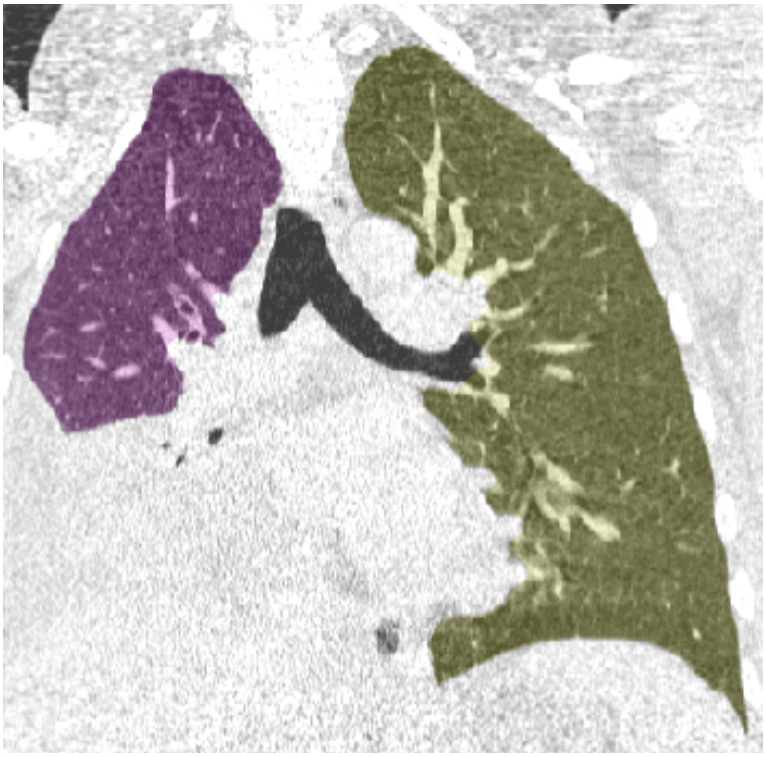

Figure 3.4: Example of incomplete lung segmentation due to pathology. The entire lung volume is not segmented, due to pulmonary empyema in the right lung. However, the ventilated regions of the lung are correctly segmented.

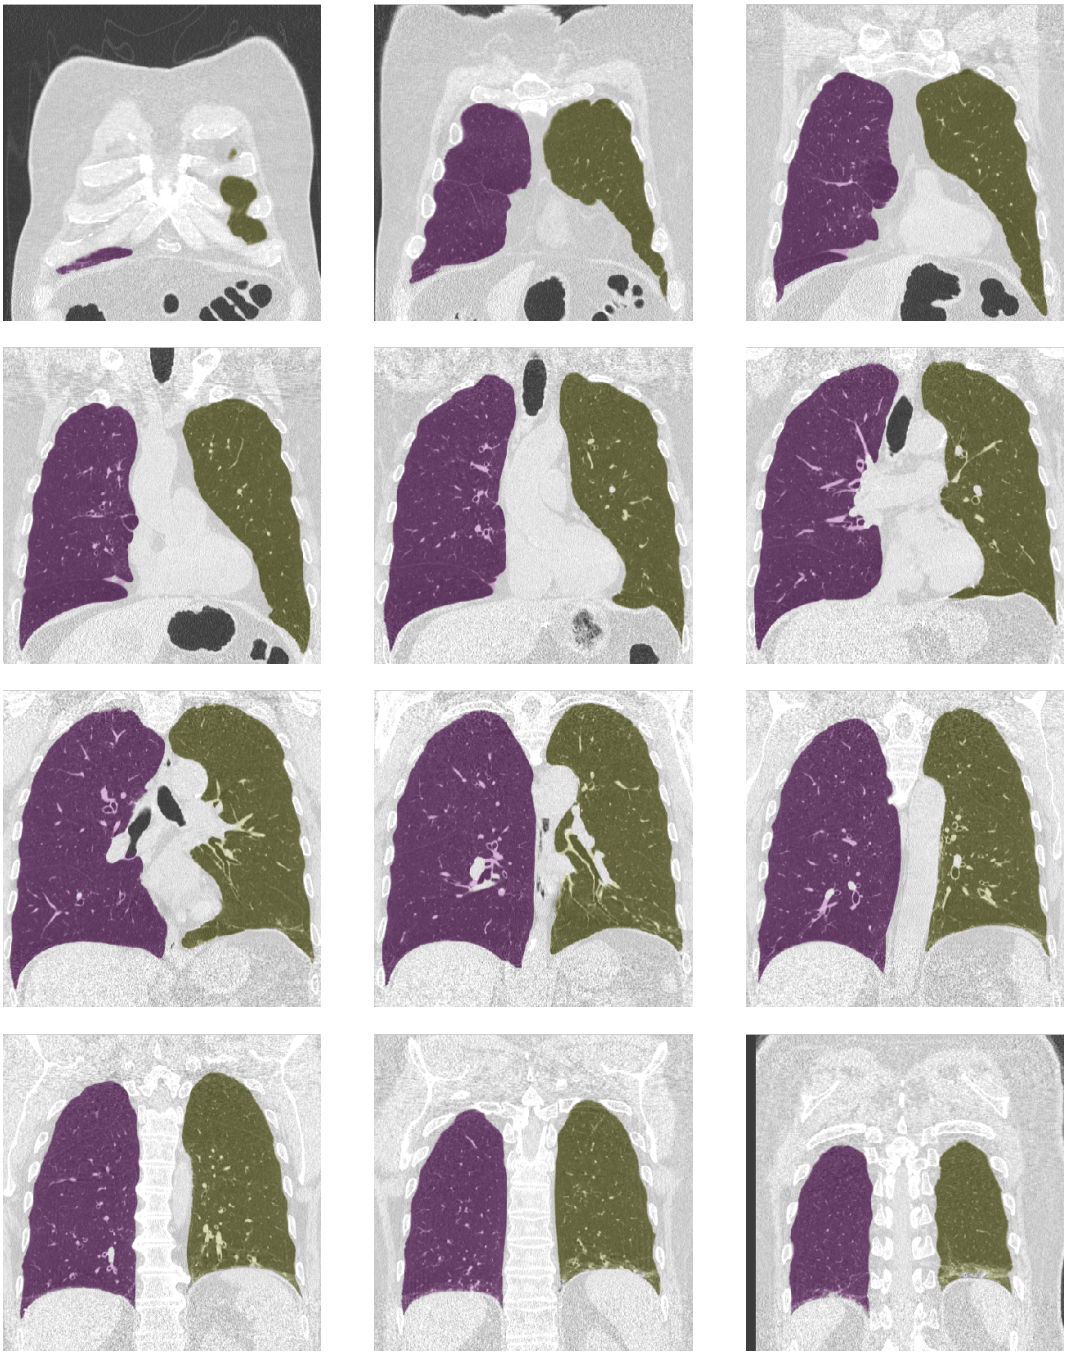

Figure 3.5: Results of automatic lung segmentation from a CT dataset. The dataset is isotropic and has been resliced in a coronal orientation. Only a regularly spaced subset of coronal slices is shown; intermediate slices are not shown but are part of the segmentation. A transparent overlay shows the masks for the right (magenta) and left (yellow) lungs. In addition to the lungs, dark air-like regions can be seen in the colon (base of the top images) and the trachea (centre of the top images). These are correctly excluded from the lung segmentation mask.

3.2.6 Performance

Table 3.1 shows the algorithm execution time for 10 datasets on a laptop computer (2.66 GHz MacBook Pro running OS 10.8 and Matlab 8). The algorithm is broken down into three stages: loading of the original data; segmentation of the airways; and the final segmentation of the lungs. The total algorithm execution time is also shown, and times are compared with the volume size of the lungs. The mean time for lung segmentation is 94.4 s, excluding time for data loading and airway segmentation.

Table 3.1: Algorithm execution time and lung volume for CT lung segmentation.

Subject	Lung Volume (cm^3)	Algorithm execution time (s)			
		Data loading	**Airway segmentation**	**Lung segmentation**	**Total**
1	3167	46.5	27.4	76.3	150.1
2	4890	43.1	28.5	103.3	174.9
3	6208	69.5	28.3	103.6	201.4
4	4343	46.4	31.4	110.0	187.7
5	3293	47.4	26.7	84.0	158.0
6	5436	50.5	27.7	81.1	159.3
7	4894	52.1	20.4	105.3	177.7
8	7247	53.2	37.1	125.9	216.2
9	3374	50.4	19.0	74.1	143.5
10	3467	49.9	20.0	80.8	150.7
Mean		**50.9**	**26.7**	**94.4**	**172.0**

3.3 Lung segmentation from proton MRI

Segmenting the lung contours from proton MRI data is an important part of our regional analysis method (chapter 5). The lung mask is used in our multi-modal image registration technique which relates CT-derived lobar maps to hyperpolarised xenon MRI images. Direct registration between CT and xenon MRI is difficult because the xenon MRI images do not differentiate between unventilated lung regions and regions exterior to the lung. Proton MRI provides an ideal intermediate stage, since images can be rigidly registered to xenon data using image metadata, and the visible lung boundaries can be used to perform a non-rigid registration with CT data.

However, lung segmentation remains a challenging problem in MRI, due to its poor contrast, noise and lack of calibration. We have been motivated to develop our own automated algorithm by the very limited number of existing publications and the particular difficulties presented by our proton MRI data, which consists of thick coronal slices. Our method is adaptable to higher resolution proton MRI data.

3.3.1 Previous work

MRI lung segmentation is more challenging than for CT due to the non-uniform nature of the MRI signal, noise and poor contrast, resulting in a much less well-defined lung boundary [83]. The literature on MRI lung segmentation is very limited. For standard proton MRI, Ray et al. [84] develop an active contour approach as a prerequisite for segmentation of hyperpolarised gas MRI. Their method allows the merging of multiple contours to segment the lung region. Results are presented for ten subjects. Middleton and Damper [83] use an active contour model initialised by a trained neural network classifier. They

show results for 13 subjects. Osareh and Shadgar [85] use region-aided geometric snakes to find the lung contours on each axial slice of the MRI image. An initial slice is chosen based on template matching to choose the best-defined lung contours. The snake solution then initialises segmentation on neighbouring slices. Results are presented for 30 subjects.

Hyperpolarised gas MRI

Segmentation from hyperpolarised gas MRI is related to that of proton MRI, but presents a different set of challenges since only ventilated regions are visible. Lung segmentation is used to measure unventilated regions (ventilation defects), but it is difficult to distinguish unventilated lung regions near the lung boundary from areas exterior to the lung. Furthermore, the slice thickness is typically much larger for gas imaging than for CT or standard proton MRI. Many authors perform the segmentation using manual or semi-automatic techniques [54,86,87,88]. Tustison et al. [89] develop an automated segmentation method based on statistical shape analysis. Other authors segment the lung boundary in proton MRI or CT and use a rigid registration to apply it to the gas MRI data [54,90].

3.3.2 Method

Our algorithm is summarised in Figure 3.6. Our novel method comprises an initialisation based on 3D region-growing with an adaptive threshold. Contour refinement is then performed for each coronal slice using a 2D level-set algorithm with an image-gradient force, region-based speed term and curvature regulariser. We describe these steps in more detail below.

1. **Smoothing and noise reduction using a Gaussian filter**

   We initially apply a Gaussian filter with a standard deviation of 1 mm. This smoothing filter reduces image noise and reduces the inhomogeneity of image values around the lung border which can make the lung segmentation difficult.

2. **Identify initial seed points for left and right lung**

   We next find seed points to initialise our segmentation of the left and right lungs. The lung image is divided into two parts along the left-right axis, and a search region is created around the centre of each image. The search region is a parallelepiped with size equal to one tenth of the image size in each respective dimension. The seed points are chosen as the voxels with the lowest intensity value in each search region.

3. **Adaptive region-growing**

   For each lung, a threshold value is applied to the image and a region-growing performed on the resulting mask, starting from the seed points using six-way connectivity in 3D. This operation is repeated with different threshold values, in order to find the highest threshold value which produces a segmentation that does not leak out of the lung volume. Leakage is defined as being where the segmented region connects with the image boundaries. For each lung, the largest segmentation which does not leak (the segmentation from the highest threshold value) is our initial guess for the lung segmentation.

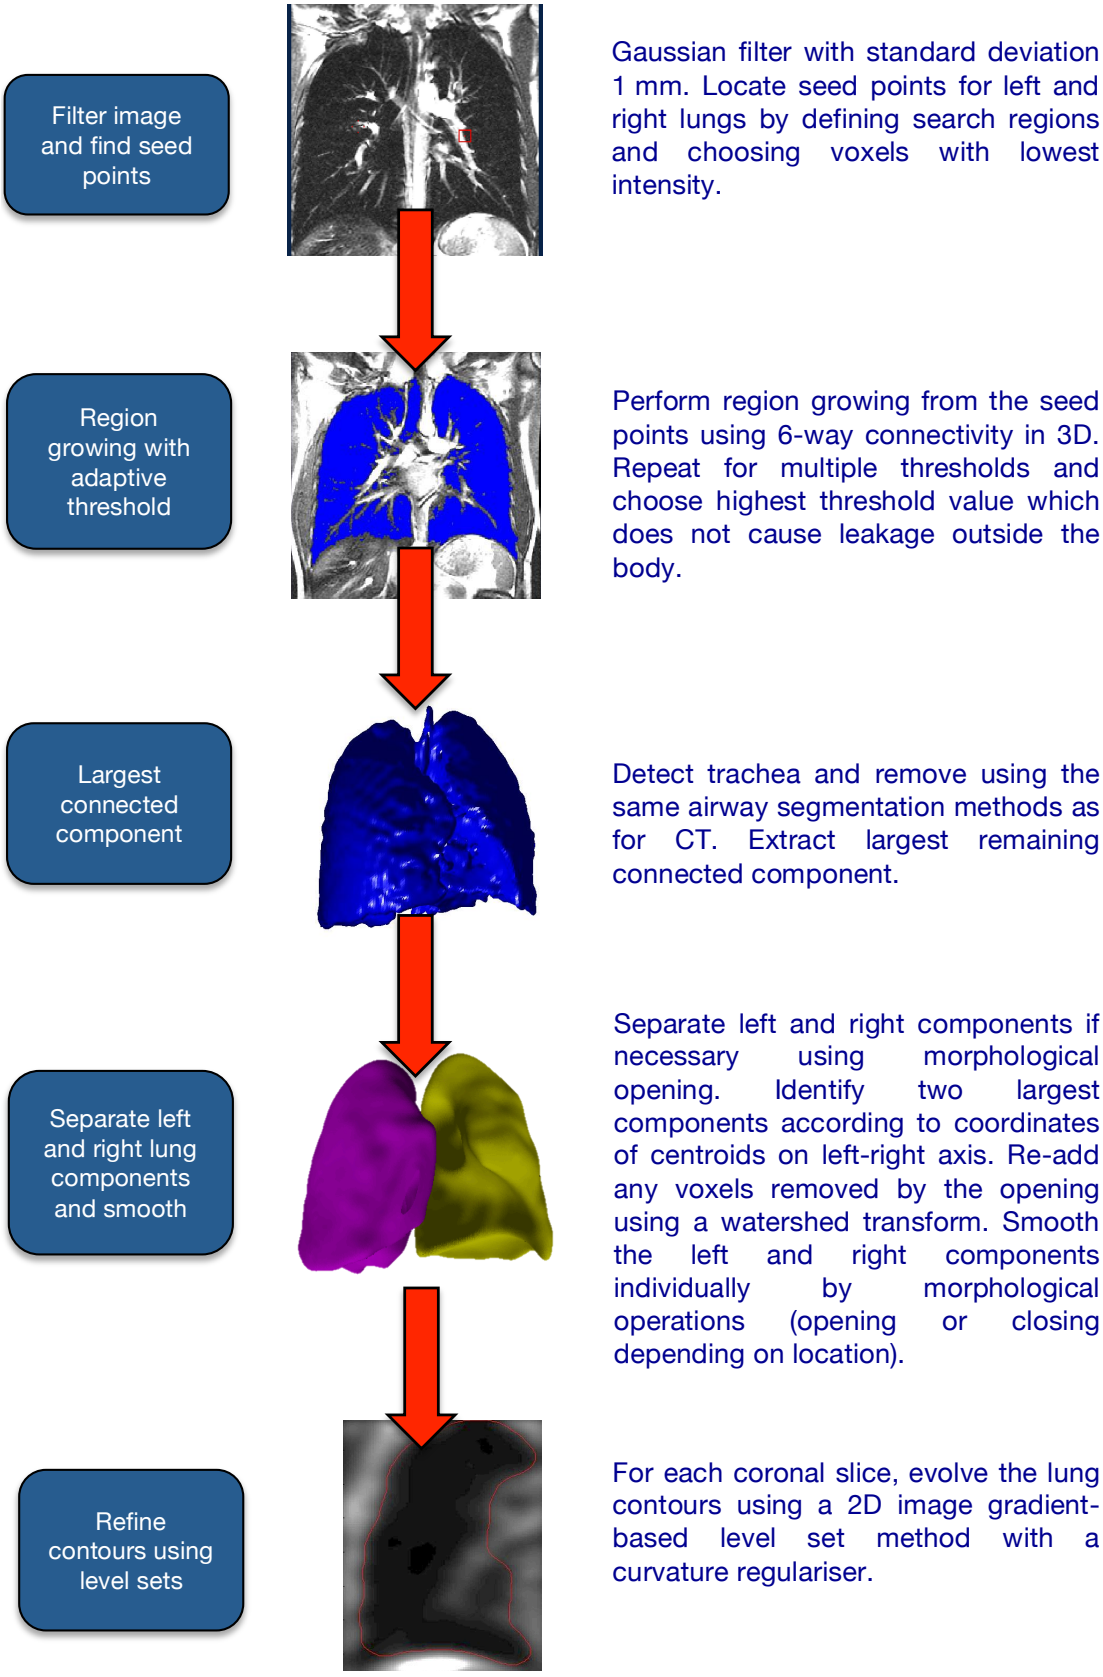

Figure 3.6: Schematic diagram of the proton MRI lung segmentation algorithm.

4. Removal of trachea

We detect the trachea by applying the airway segmentation methods described later (section 3.4). Airway segmentation from proton MRI data is much less reliable than for CT data, but is sufficient to identity and remove the trachea. We do not attempt to remove the main bronchi, as bronchial labelling is not sufficiently reliable in the thick slice proton MRI data we are using.

5. Left and right lung segmentation

The initial lung approximation is separated into left and right lungs using the same method as for CT data (section 3.2). A sample result is shown in Figure 3.7.

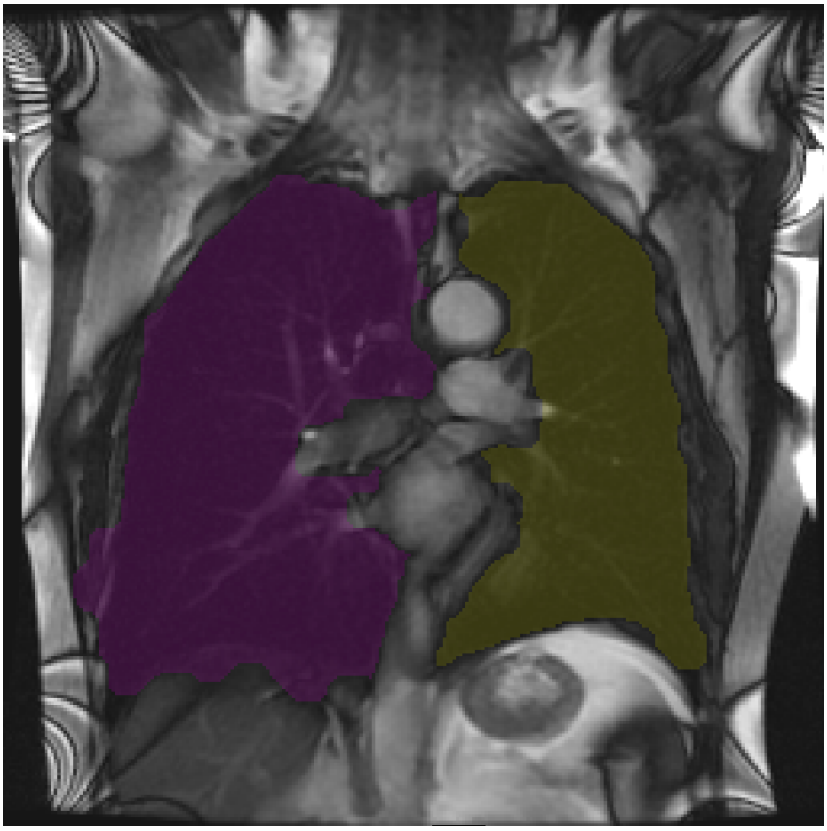

Figure 3.7: Illustration of the MRI lung segmentation initialisation. Regions of the image close to the heart tend to show undersegmentation of the lung due to the high density of the blood vessels. Regions closest to the image boundaries tend to show oversegmentation of the lung where the region-growing leaks past the lung boundary. We perform morphological closing in the undersegmented region and morphological opening in the oversegmented region.

6. Smoothing of lung contour

The initial lung approximations contain gaps and leakage, which do not form an optimal initialisation for the level set refinement. We apply morphological operations separately to the left and right lung segmentations with a structural element of radius 10 mm, in order to fill any segmentation holes and smooth the lung contours. We have found that the initial lung segmentation tends to undersegment in the central region of the image (i.e. close to the heart), due to the high-density blood vessels, while away from the heart, it tends to oversegment, i.e. the segmentation leaks outside of the thoracic cavity. This is illustrated in Figure 3.7. To deal with this, we perform morphological closing on voxels close to the image centre, and morphological opening on voxels further away. Additionally, we fill any image holes separately for each lung, on each 2D coronal slice. The result is a smooth initialisation for the left and right lungs.

7. Level set contour refinement

The final lung boundary contours are created by applying a 2D level set-based refinement method separately to the left and right lungs for each coronal slice of the image. The starting contours are derived from the smoothed initial approximations generated in the previous step. The purpose of this refinement step is to produce smoothed image boundaries which correct for segmentation errors caused by noise and other limitations of MRI lung imaging.

Level sets [91] are widely used in image processing. Here we present only a sketch of our method. A detailed overview of level set methods is given by Angelini et al. [92].

We represent the lung contour parametrically as the zero level set $\Psi = 0$ of a higher dimensional function Ψ. We choose Ψ to be a signed distance function from the contour, and we define Ψ on grid points that correspond to the lung

image I. The function Ψ evolves over time, according to a speed function F which acts normal to the contour

$$\Psi_t + F(x,y)|\nabla\Psi| = 0\,. \tag{3.1}$$

We choose a speed function with a region-based term and a stopping function based on the contour curvature and the image gradient, given by

$$F(x,y) = \alpha F_{\text{region}} + (1-\alpha)\kappa g(\nabla I)\,. \tag{3.2}$$

Here, κ is the image curvature, which regularises the curve to produce a smooth contour

$$\kappa = \nabla \cdot \left(\frac{\nabla\Psi}{|\nabla\Psi|}\right). \tag{3.3}$$

The function $g(\nabla I)$ is a stopping term, which seeks to halt the curve evolution over a boundary with a large image gradient

$$g(\nabla I) = \frac{1}{1+|\nabla I * G_\sigma|}\,, \tag{3.4}$$

where G_σ is a Gaussian filter with standard deviation σ. We choose a value of 3 mm for σ. The region term attempts to guide the contour between high- and low-valued regions of the image

$$F_{\text{region}}(x,y) = \begin{cases} I(x,y) - L & \text{if} \quad I(x,y) < L + \dfrac{U-L}{2} \\ U - I(x,y) & \text{otherwise} \end{cases}, \tag{3.5}$$

where L and U are lower and upper image thresholds. We set L to be 0 and U to be the threshold found by the adaptive region-growing initialisation method above. We introduce the region term as it helps to correct for leakage by pulling the lung contour into the main lung region. It also helps to force the segmentation into sharp corner regions at the base of the lungs, which can be removed by the morphological smoothing operations and regularisation terms. The parameter α defines the weighting between the region and stopping terms.

We have found 0.01 to be a good value for our images, which have intensity values scaled between 0 and 255.

We discretise equation (3.1) in time using a Forward Euler method, with the time step set to $\Delta t = 0.5$ which we have found to be numerically stable. Every ten iterations we re-compute the signed distance function Ψ to prevent sharp gradients from developing which could lead to numerical instability. We solve for maximum number of 1000 iterations, although convergence generally happens before this limit is reached. The resulting contours are converted into a binary segmentation mask for each lung.

3.3.3 Limitations of method

Our lung segmentation consists of two main stages: lung approximation and contour refinement. For the lung approximation, we adopt a region-growing method because it is fast, simple to implement in 3D and requires no training or model information. MRI data are not normalised, but we account for this by using an adaptive threshold. However, we encounter a significant issue when using region-growing with MRI data. Parts of the lung boundary are not well defined in the image, for reasons which include partial volume effects and variations in the MRI signal. While some uncertainties in the lung boundary can be quantified using partial volume correction algorithms (see section 3.3.6), a more serious problem can occur when the region-growing leaks out of the lung parenchyma and segments a low-intensity region exterior to the lung. This effect is visible in several places in Figure 3.7. We correct for these effects using morphological operations in step 6 above; however, this heuristic approach may not be robust to all leakage. Improved results for the initialisation might be obtained by replacing the region-growing with a random walker algorithm [93]. This method assigns voxels to regions according to their connectivity to marker points, which could be automatically placed inside and

exterior to the lung. Importantly, this algorithm is often robust to incomplete or poorly defined boundaries and much less prone to the leakage of region-growing methods. Another approach which could be adopted with the clinical trial data is to use the high quality CT lung segmentations we have for each subject. While the lung inflations (and therefore lung shape) are different than for MRI, the segmentations could be adaptively scaled using an affine transformation. This approach is robust in the sense that the initial approximation should be within quantifiable limits of the actual lung contour.

The second stage of our lung segmentation is contour refinement, which aims to correct for errors in the segmentation and smooth the lung boundaries. We have chosen a level set approach because it can often remove areas of leakage, and it naturally deals with the concavity at the base of the lung and holes in the segmentation due to airways or vessels. Level set methods also generalise straightforwardly to 3D, although we have found our algorithm does not perform well in 3D for thick-slice coronal datasets, due to the large curvature of the lung at the anterior and posterior slices. In 2D, the lung may appear as more than one region in these slices, but the level set approach naturally deals with this situation. However, level sets are computationally expensive and can suffer from stability issues. Furthermore, our method, which includes a region term to help maintain the lung shape, is quite sensitive to the initial approximation. As discussed above, the initial approximation is not always accurate in its current form due to leakage. The level set algorithm can sometimes remove this leakage, but not in all cases. Switching from the implicit level set method to an explicit active contour approach [94] for the contour refinement, using a similar gradient-based approach, may therefore offer advantages. These include being more robust and computationally faster than the level set approach, and the ability to introduce explicit constraints on the lung shape.

3.3.4 Results

We performed MRI lung segmentation on five of the proton MRI datasets described in section 1.5. An example result is shown in Figure 3.8.

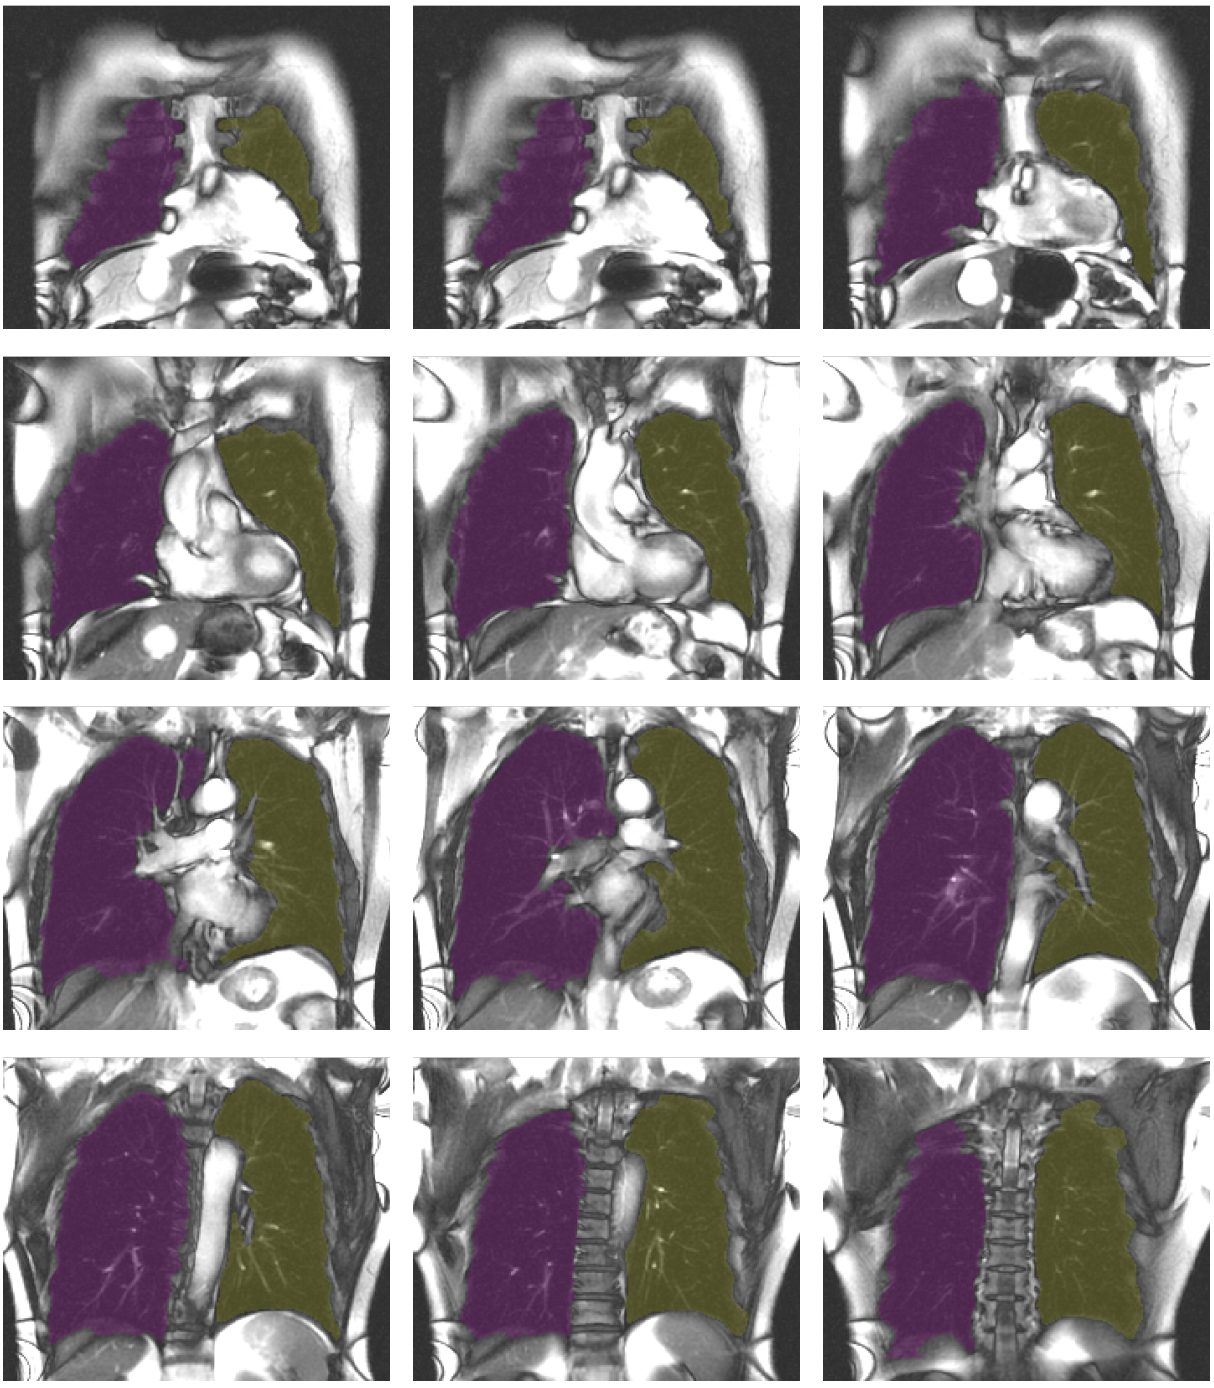

Figure 3.8: Results of automatic lung segmentation from a proton MRI dataset. Each coronal slice is shown with a transparent overlay showing the masks for the right (magenta) and left (yellow) lungs.

3.3.5 Validation

To validate the results for this dataset, we performed a manual segmentation of the lung contours for each coronal image slice using the ITK-Snap software package [95]. We use this manual segmentation as ground truth. The automatic and manual segmentations for the regional analysis dataset are shown overlaid in Figure 3.9. Red voxels represent true positives (voxels present in both the ground truth and automated segmentation masks). Blue voxels are false positives (voxels in the automated segmentation mask that do not appear in the ground truth). Green voxels are false negatives (voxels that appear in the ground truth but not the automated segmentation mask).

We can quantify the performance of the automated segmentation method by computing the Dice coefficient [96]. If TP is the number of true positives, FN is the number of false negatives and FP is the number of false positives, then the Dice coefficient is given by

$$\mathrm{Dice} = \frac{2\mathrm{TP}}{2\mathrm{TP} + \mathrm{FP} + \mathrm{FN}}\,. \tag{3.6}$$

The closer the value of the Dice coefficient to 1, the better the segmentation. The Dice coefficient is equivalent to the segmentation performance metric F described in Middleton and Damper [83]. For comparison with previous work [83,85], we also compute measurements of precision and recall. Precision measures the proportion of voxels from the automatic segmentation that are correct, while recall measures the proportion of true lung voxels that have been correctly identified. These quantities are computed by

$$\mathrm{precision} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}\,, \tag{3.7}$$

$$\mathrm{recall} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}\,. \tag{3.8}$$

None of these quantities provides a direct measure of how the segmented lung boundary deviates from the true boundary. To measure this, we find the absolute distance in 3D from every voxel on the surface of the automatic segmentation mask to the nearest voxel on the surface of the ground truth mask. We also perform the inverse calculation, finding the distance from every voxel on the surface of the ground truth mask to the nearest voxel on the surface of the automatic segmentation. We compute the mean and maximum distance across all these values. The mean gives an overall measure of the accuracy of the surface segmentation, while the maximum value quantifies large segmentation errors that can occur in some regions where the automatic segmentation performs poorly. By computing the distances in both directions, our measurement takes into account both false positives and false negatives, and is symmetric between the two segmentation masks.

Table 3.2 shows the Dice, precision, recall and distance measurements for the automated MRI segmentations compared to the manual ground truth.

Table 3.2: Validation results for proton MRI lung segmentation. For each of the five datasets, precision and recall measurements and Dice coefficients are computed. We also show the mean and maximum distances in 3D between the segmented lung mask and the ground truth.

	Mask measures			Distance between surfaces (mm)	
Subject	**Precision**	**Recall**	**Dice**	**Mean**	**Maximum**
1	0.94	0.98	0.96	2.20	22.4
2	0.96	0.91	0.93	3.13	26.7
3	0.96	0.95	0.96	2.06	37.8
4	0.91	0.98	0.94	2.65	27.9
5	0.87	0.96	0.92	3.77	43.6
Mean	**0.93**	**0.96**	**0.94**	**2.76**	**31.7**

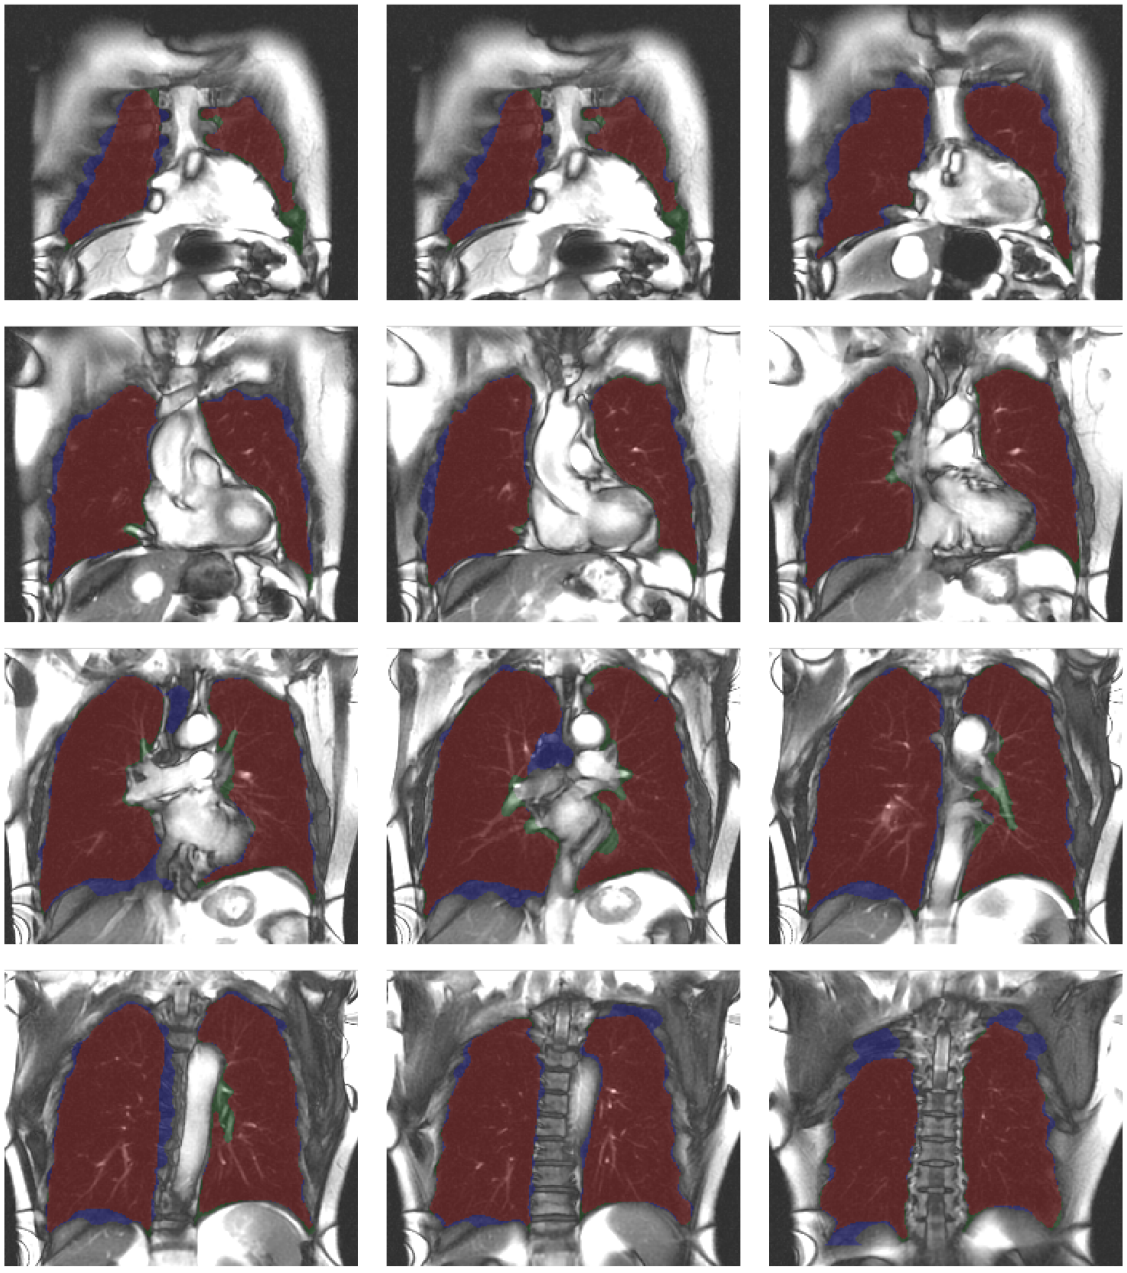

Figure 3.9: Slice by slice comparison of automatic and manual MRI lung segmentations. Each coronal slice is shown with a transparent overlay. Taking the manual segmentations as ground truth, red shows true positives (voxels in the manual and automatic segmentations); blue shows false positives (voxels in the automatic segmentation but not in the manual segmentation); and green shows false negatives (voxels in the manual segmentation but not in the automatic segmentation).

3.3.6 Discussion

From Table 3.2 (p. 68), the algorithm shows good overall performance. All the Dice coefficients are above 0.9, and the mean distances between the segmented surfaces are comparable to the scanning voxel width (3.12 mm). The main segmentation errors are due to the leakage problem described above, which is not always corrected for by the refinement stage.

In Table 3.3 we compare our validation results with previous work [83,85], showing minimum, maximum and mean values for precision, recall and the Dice coefficient. Although our sample size is small, these results compare well with existing work, particularly considering the lower resolution of our images, which were scanned at 12 × 128 × 128 voxels, compared to 256 × 256 × 80 (Osareh and Shadgar [85]) and 256 × 256 × 30 (Middleton and Damper [83]). As described in section 1.5, the images we are using were upsampled to 12 × 256 × 256 voxels by the scanner.

Table 3.3: Comparison of MRI lung segmentation results against previously published methods. Precision, recall and Dice coefficients are computed for each dataset. The minimum, maximum and mean values across all datasets are shown.

	Mask measures								
Method	**Precision**			**Recall**			**Dice**		
	Min	**Max**	**Mean**	**Min**	**Max**	**Mean**	**Min**	**Max**	**Mean**
Our method	0.87	0.96	**0.93**	0.91	0.98	**0.96**	0.92	0.96	**0.94**
Middleton and Damper [83]	0.82	0.99	**0.92**	0.60	0.91	**0.80**	0.69	0.92	**0.86**
Osareh and Shadgar [85]	0.88	0.96	**0.93**	0.84	0.98	**0.91**	0.87	0.97	**0.92**

Comparing the mean Dice coefficients, which provide an overall measure of algorithm performance, our algorithm offers a marginal improvement over Osareh and Shadgar [85] and a significant improvement over Middleton and Damper [83]. The Dice coefficient is a combination of the complementary

metrics precision and recall, and the low Dice coefficients for Middleton and Damper are explained by their low recall values. Our mean recall values are significantly higher than the other two methods, while our precision values are similar. This suggests our algorithm is segmenting more of the true lung volume than the other methods, but a greater proportion of our detected points are incorrect. This oversegmentation is consistent with the leakage issues described above.

Partial volume effects

Comparing the segmentations against the ground truth such as in Figure 3.9, we have identified the main regions showing segmentation errors. These are in the anterior (front) and posterior (rear) slices of the lung where, due to the thick coronal slices and the high curvature of the lung boundary in these regions, the lung boundaries are indistinct and ambiguous due to partial volume effects. In Figure 3.9 there appears to be a prevalence of false negatives (green) towards the anterior of the lung and false positives (blue) towards the posterior. This may be suggestive of a systematic MRI artefact such as a non-stationary bias field across the image [97]. However we have analysed the locations of the false positives and false negatives for each dataset, and found no strong indication of bias in the false positives or false negatives along the principal directions.

For lung segmentation, partial volume effects occur in voxels crossing the lung surface. To assess their importance, we compute the number of surface voxels in the ground truth segmentation, and present this as a percentage of the total number of voxels in Table 3.4. We define a surface voxel as any voxel in the segmentation having a neighbour in any of the six principal directions which is not part of the segmentation. An accurate calculation of the effect on volume measurements requires a partial volume model such as those used in analysis

of brain MRI [98,99]. A partial volume model for the lungs would need to take into account the properties of the vasculature and airways inside the lungs, as well as bone, pleural fluid and other tissue exterior to the chest cavity. For the purposes of estimation, we can make a reasonable assumption that the upper error bound in our volume computation due to partial volume effects is 50 % of the surface voxel volume. From Table 3.4, this corresponds to a 16.7 % error in our mean volume measurements, which may warrant the use of partial volume correction. This effect is large due to the very thick image slices, and is significantly greater than the 5.4 % mean absolute difference in volume between the automatic and ground truth segmentations (expressed as a percentage of the ground truth volume), which incorporates the effect of segmentation errors. Partial volumes affect manual ground truth as well as the automated segmentations, since they are a cause for uncertainty in both cases.

Table 3.4: Metrics for assessing partial volume effects. For each dataset, the lung volume and total number of voxels in the ground truth segmentation are shown. Other metrics are given as a percentage of this total number of voxels. The number of surface voxels is computed for the ground truth segmentation, and approximately quantifies the number of voxels subject to partial volume effects. The absolute difference in lung volume is computed between the automatic and ground truth segmentations, and is due to partial volume effects and other segmentation issues. The other metrics show the number of false negatives (voxels not segmented by the automatic method) which lie on the segmentation of the surface, and the number of false positives (voxels incorrectly segmented) which lie adjacent to the surface. These represent voxels whose misclassification may be due to partial volume effects.

Subject	Lung volume	Total number of voxels in lung	Number of surface voxels	Segmented lung volume difference	Number of FNs on surface	Number of FPs adjacent to surface
	(cm^3)		**(% of the total number of voxels in the lung)**			
1	6376	174101	30.2	3.9	5.5	1.9
2	5639	153989	32.7	6.2	2.9	7.3
3	4364	119161	36.6	0.7	3.8	3.0
4	5015	136953	31.4	7.0	7.6	1.8
5	4917	134276	36.3	9.2	11.3	1.3
Mean			**33.4**	**5.4**	**6.2**	**3.1**

Table 3.4 includes two additional metrics, which aim to quantify the effect of partial volumes on the automated segmentation results. As described above, segmentation errors consist of false negatives and false positives, as defined relative to the ground truth. Many of these error points lie close to the lung surface, and may be due to the partial volume. To place a bound on this effect on the segmentation, we compute the number of false negative voxels which lie on the surface of the ground truth segmentation. We similarly compute the number of false-positive voxels which are adjacent to this surface (by definition, false positives cannot lie on the surface itself). We express these as percentages of the number of voxels in the ground truth. Our results suggest that on average, partial volume effects may cause up to 9 % of the lung voxels to be classified differently between automated and manual segmentations.

3.3.7 Performance

Table 3.5 shows the algorithm execution time on a laptop computer (2.66 GHz MacBook Pro running OS 10.8 and Matlab 8). Times are shown for three stages of the algorithm: data loading; initial lung segmentation and separation; and level set contour refinement. The total execution time varies between approximately 1.5 and 2 minutes according to the size of the lung.

Table 3.5: Execution time for proton MRI lung segmentation.

Subject	Lung volume (cm^3)	Algorithm execution time (s)			
		Data loading	**Initial segmentation**	**Level set refinement**	**Total**
1	6376	1.3	8.0	99.5	108.9
2	5639	1.5	8.7	98.1	108.5
3	4364	1.5	8.8	73.5	83.9
4	5015	1.4	8.0	78.2	88.1
5	4917	2.2	10.4	75.4	88.2
Mean		**1.6**	**8.8**	**84.9**	**95.5**

3.4 Airway segmentation from CT

We segment the airways (bronchi) in order to separate the left and right lung segmentations (section 3.2) and segment the lobes (chapter 4), since the branching structure of the bronchial tree has a close relationship with the lobes. Our airway segmentation is used as the starting point for building a computational model of airflow simulation (chapter 6). Airway segmentation also has potential clinical applications, since regional airway counts and radius measurements correlate with disease [38,39].

In CT data, airways appear as irregular hollow tubes, with high-density airway walls containing low-density air. The airway tree is visible down to around the seventh generation [65]. In the larger airways, the airway walls and air interior can be clearly distinguished, but as the airways become smaller, partial volume effects become more significant. As a result, smaller airways are hard to detect, as the airway wall becomes increasingly hard to separate from the lumen and surrounding parenchymal tissue.

We develop a novel airway segmentation method based on region-growing with an explosion control heuristic. Our method is similar to that used by Wood et al. for canine lungs [100] and van Ginneken et al. [101] for human lungs, but we have introduced a novel adaptation of a variable-thickness wavefront to detect bifurcations. Our method has a number of advantages including robustness to noise and requires no heuristic pruning steps to remove spurious branches.

After segmenting the airways, we compute the centreline and radius for each bronchus, and label the bronchi according to the pulmonary lobes they supply. We explain each stage of the method in the following sections.

3.4.1 Previous work

Segmentation of the airways is an ongoing research area and was the subject of a recent Miccai workshop challenge [64]. The most widely used methods are based on region-growing and wave propagation [101,102,103,104]. Other approaches include detecting tube-like airway structures and connecting them to form a tree [105,106], fuzzy connectivity [107] and morphological methods [108]. Model-based approaches make use of additional information such as vessel locations [109]. Other authors use a hybrid approach, typically combining region-growing with other methods [110,111,112,113].

The threshold-based region-growing algorithm starts from the trachea, and extracts the air space inside the bronchial tree by adding neighbouring points with intensities similar to that of air (–1000 Hounsfield units) to the segmentation. The trachea location can be manually selected or automatically determined through analysis of the lung segmentation, due to the trachea's distinctive shape and situation. An explosion control heuristic prevents the segmentation leaking into the lung parenchyma at points where airway walls are not detectable due to partial volume effects [104,110,113].

Region-growing is commonly used because it is a simple, fast and robust method for segmenting the larger airways. In these airways there is a high air/tissue contrast at the inner bronchial wall, and the thick walls prevent algorithmic failures resulting from partial volume effects. However, for smaller airways the algorithm may fail where partial volume effects cause apparent blockages. Partial volume effects also cause leakage, which is corrected for by the explosion control, but this may also terminate the branch early. Variable threshold adaptations of region-growing do not fix these problems because fundamentally region-growing is based on local threshold analysis, and small airways are locally hard to distinguish from other features. We discuss region-growing further in section 3.4.3, where we introduce our method.

3.4.2 Trachea location

We first identify the location of the trachea, required as a starting point for airway segmentation. Published airway segmentation algorithms typically locate the trachea by searching for approximately circular airspace regions over consecutive axial slices (e.g. Tschirren et al. [102]). Other methods rely on manual trachea identification [110] or the assumption that the trachea can always be found in the centre of the first axial slice [103]. For the purposes of our airway segmentation, we require only a single point to be identified within the top region of the trachea, so for this we have developed an alternative approach based on connected-component analysis, creating an algorithm which is fast, effective and robust. Essentially, our algorithm searches for large superior-inferior aligned airspaces, removing those too wide or too far away from the image centre to be the trachea. Once we have identified a section of the trachea, we choose the highest voxel in the superior direction to be the seed point for our airway segmentation.

The three stages of our trachea segmentation algorithm are shown in Figure 3.10, which illustrates a 2D axial cross-section and a 3D surface rendering of the detected trachea mask at each stage. In the first stage we apply an air threshold [–1024 HU, –775 HU] to the region of interest (ROI) image determined in section 3.2.2 above. This will identify voxels interior to the trachea and airways, in addition to low-density parenchymal tissue and regions outside the body (Figure 3.10 A). We need to separate the trachea from these other regions, but these regions can appear connected; the trachea connects to the exterior of the body via the upper airways (if they are part of the scan), and the lower airways may appear to connect to the parenchyma in the small airways, where sections of the thin airway walls are undetectable due to partial volume effects.

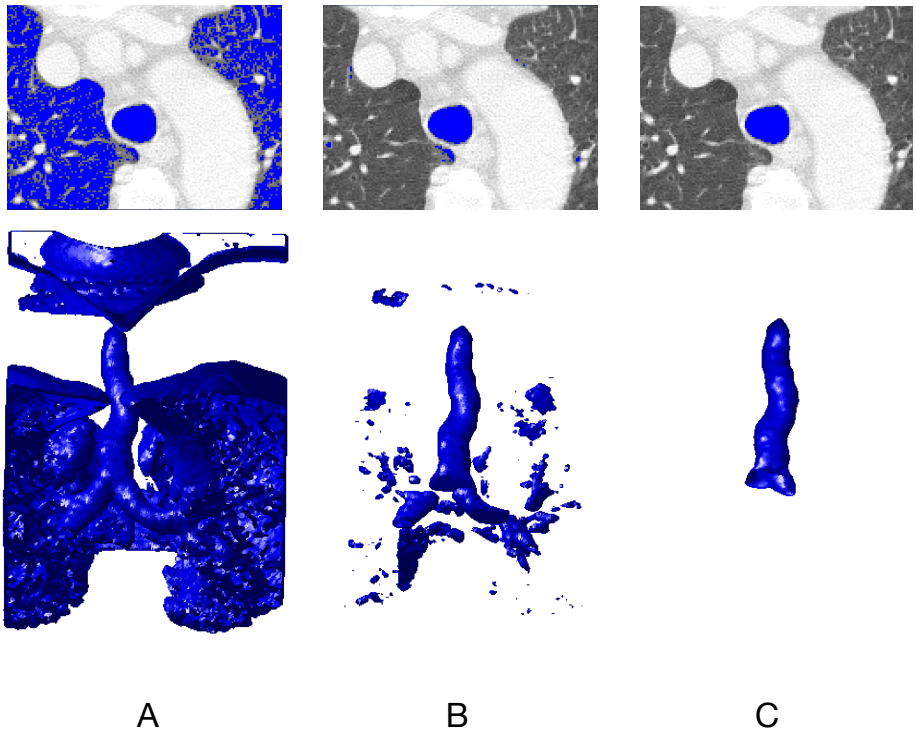

Figure 3.10: Illustration of the trachea detection algorithm. The three stages A, B and C are each illustrated with 2 images. The top image shows a 2D axial cross-section superimposed with a blue overlay showing the trachea mask. The bottom image shows a 3D surface rendering of the trachea mask. In stage A, an air threshold has been applied to the region of interest. The resulting mask includes the airways, the lung parenchyma and air regions exterior to the body. In stage B, the image has been divided into axial slices of thickness 1.3 mm and a 3D connected-component analysis applied to each slice. Components greater than 80mm wide or more than a certain distance from the image centre are removed. This stage separates the trachea from other parts of the segmentation and removes most of the lung parenchymal components. In the final stage C, the slices are reassembled in 3D and the largest remaining component is selected. The trachea location is chosen as the most superior point in this mask.

To separate the trachea from other air-like regions we exploit the knowledge that the trachea is aligned approximately vertically, has a characteristic width of ~40 mm and has thick walls which provide a clear separation from other lung regions. By analysing the threshold mask in the axial plane, we can remove components which do not have these characteristics. We divide the threshold mask into axial slices of 1.3 mm thickness (rounded upwards to the nearest voxel; for our images the slice thickness is 0.63 mm so we divide the image into 2-voxel thick slices). For each slice we find all 26-way connected

components in 3D. Components with a cross-sectional size greater than 80 mm in the axial plane are removed since they are too big to be part of the trachea. We also remove any component which includes voxels further than 100 mm from the anterior-posterior axis, or further than 75 mm along the left-right axis. This is because the subject is approximately centred in the axial plane of a CT scanner, so the trachea should also be approximately centred. The results are shown in Figure 3.10 B. For our purposes we are defining the axes to lie along the centre of the original image before it is cropped to the region of interest.

This second step naturally separates the trachea from the lower airways (Figure 3.10 B), since the bifurcation at this point causes the cross sectional area to exceed 80 mm. If the upper airways have been included in the scan, this step similarly separates the trachea from where the upper airways connect to the outside of the body through the mouth. This solves the problem of separating the trachea from other regions; however, the slice thickness cannot be chosen too large, otherwise removing these regions would also remove large parts of the trachea to which they are connected in 3D. The value of 1.3 mm is chosen so that most of the trachea remains intact, while permitting a limited amount of 3D connectivity in the superior-inferior direction. The 3D connectivity is necessary to help identify and remove regions of the lung parenchyma. The parenchymal regions are not always well connected in 2D due to large holes where the parenchymal tissue has density values outside of the air threshold. The 3D connectivity ensures the connected regions are large enough to fall outside of the 80 mm threshold and be removed. The parenchymal regions could be better detected (and therefore more easily removed) through use of a higher threshold, morphological closing or image smoothing; however all of these methods could lead to false connections appearing between the trachea and the lung parenchyma, which could cause a failure in the trachea detection.

The final step of the algorithm reassembles the mask slices into a 3D volume and identifies the largest single 3D connected component (Figure 3.10 C). Since the previous step will have removed most parts of the lung parenchyma, the single largest remaining component is a section of the trachea. We choose the most superior point in this segmentation as the trachea location.

We have verified that the trachea location algorithm performed correctly on the CT datasets used in this thesis as well as a number of other CT datasets, including low-dose CT.

Our algorithm makes a number of assumptions. The trachea should be located in the central region of the left-right axis and be aligned approximately vertically. A substantial number of voxels inside the trachea lumen must fall into the air threshold. The trachea lumen must be well separated from air regions in the lung parenchyma by high-density tissue, and there must be a large section of the trachea for which the apparent diameter of the lumen in the axial plane does not exceed 80 mm. These assumptions are anatomically reasonable, since the trachea has thick walls, contains a large volume of air throughout the entre breathing cycle, and its location and shape does not change significantly. However, image artefacts, such as those caused by metallic implants or subject motion, could affect the apparent tissue density in the image. A severe artefact might lead to the appearance of a blockage in the trachea, or a false gap in the trachea wall. This could potentially disrupt the algorithm, although it will still succeed if a large enough component of the trachea exists which does obey the above assumptions. Similar problems might occur in low-resolution images due to partial volume effects.

A further assumption is that the lung parenchyma does not contain large air-like regions with similar properties to the trachea. It is possible that small sections of the parenchyma may exhibit this appearance due to pathology, unusual airway anatomy or surgical interventions. Normally the trachea

detection algorithm will still preferentially select the correct region for the trachea due to its greater size, but if the trachea segmentation were disrupted due to the reasons explained above, other regions of the lung might be incorrectly selected instead.

To verify the trachea detection, the user is shown the selected point highlighted on the original image. In the case of an algorithm failure, this allows the user to manually select an alternative point.

3.4.3 Region-growing with explosion control

We segment the airways by again applying a threshold [−1024 HU, −775 HU] to construct a mask showing low-density air voxels within the airway lumen. Region-growing methods start at the trachea and iteratively add neighbouring voxels from the mask to create a segmentation of the airways. A problem with this approach is that large parts of the lung parenchyma fall into the low-density threshold range, and the algorithm has to be able to separate these from the lumen. This is straightforward in the large airways where the thick high-density airway walls provide a clear barrier between lumen and parenchyma, but in the smaller airways, partial volume effects can reduce the voxel density values of the walls, leading to apparent connections between the lumen and the parenchyma. This leads to the problem of *leakage*, where the computational airway segmentation escapes from the lumen and grows into the surrounding tissue. Leakage is a standard problem in region-growing methods for airway segmentation, and must be addressed by detecting and removing leakages, a process also known as *explosion control*.

A number of region-growing methods employ global explosion detection [102,104,113]. This is linked to a variable air threshold, which is set at a low value and then increased until leakage occurs. This is a very conservative

approach, since a leakage anywhere in the airway tree will terminate the entire segmentation. As a result, these methods may omit smaller airways, but can be suitable for initialising more sensitive airway detection methods [103,110].

To obtain more generations with region-growing, a wavefront-based method can be introduced which allows for localised explosion control [100,101,114]. The wavefront mechanism we have developed has similarities to that used by van Ginneken et al. [101], but where their method involves a wavefront of single voxel thickness, our method uses a wavefront with variable thickness (Figure 3.11).

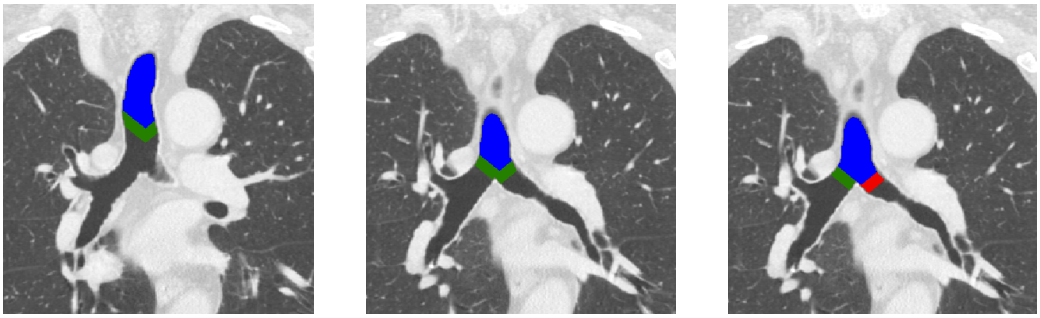

Figure 3.11: Illustration of the airway segmentation in progress. The images show a 2D coronal slice cross-section of the 3D image with the airway segmentation superimposed as a colour overlay. A: The wavefront (green) is moving downwards through the trachea, connecting voxels which lie within the air threshold. Voxels at the rear of the wavefront are added to the current bronchus segmentation (blue). B and C: Bifurcations are detected by performing 3D connected-component analysis on the wavefront. In B the wavefront (green) still forms a single component, while in C it has divided into 2 components (green and red). When this occurs, the previous bronchus (blue) is terminated and two new child bronchi are created, each with their own new wavefront (green and red).

The wavefront propagates through each bronchus using 26-way connectivity between voxels in the air mask. As new points are added to the wavefront, voxels at the rear of the wavefront are added to the bronchus segmentation. When the wavefront splits into separate components (see Figure 3.11 B), this is determined to be a bifurcation point. The current bronchus segmentation is

terminated and two new child bronchi added, each with its own new wavefront (see Figure 3.11 C). In this way, the airway segmentation is naturally separated into bronchi as part of the growing process.

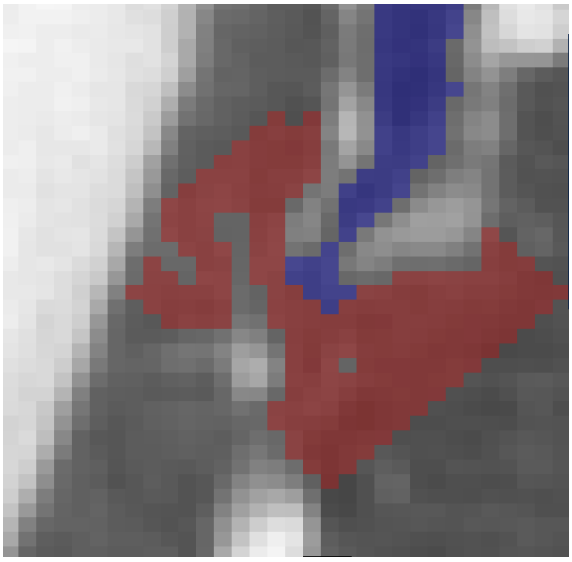

Figure 3.12: Illustration of segmentation leakage from the airway lumen into the surrounding parenchyma during region-growing. The correctly segmented lumen is shown in blue. The leakage is shown in red; this is automatically detected and removed. Leakage is due to partial volume effects near the airway wall, where the mean density within a voxel is low due to the thinness of the airway wall.

The advantage of this bronchus-by-bronchus approach is that leakage can be detected and removed without terminating the whole segmentation. We detect leakages by monitoring the wavefront size at each iteration, which should decrease as the airways get smaller, except for temporary increases around bifurcation points. When the wavefront size exceeds a particular threshold, we assume a leakage has occurred (Figure 3.12). Our algorithm tracks the point at which the wavefront started to increase in size, and we remove all voxels after that point. Our method is an improvement on that of van Ginneken et al. [101] where the whole bronchus is deleted in the event of an explosion, which may in some cases erase true sections of the airway lumen. We choose the threshold experimentally as seven times the minimum wavefront size attained along the airway path from the trachea to the current point. This was found to provide a

good balance between detecting explosions and permitting the temporary increase in bronchial size that occurs near to bifurcations.

We have introduced a variable thickness wavefront in order to make the bifurcation detection more robust to noise. Noise can result in holes appearing in the segmentation mask. With a wavefront of single voxel thickness, this noise could divide the wavefront into separate components, leading to an apparent bifurcation where none exists. To get round this problem, van Ginneken et al. [101] employ a large (30-way) connectivity matrix to determine whether their wavefront has split. Our thick wavefront approach naturally deals with this noise in a more robust way, since the thick wavefront permits connectivity in 3 dimensions. We use a wavefront thickness of 10 mm in the trachea, followed by 5 mm in smaller airways (rounded to the nearest voxel). Occasionally, the airway segmentation shows small pockets of air branching off the trachea, particularly around the larynx and upper airways. The thicker wavefront in the trachea helps to prevent spurious bifurcations from appearing at these points. For the same reasons we impose an additional constraint, which is that the trachea segmentation cannot bifurcate until the wavefront has moved by at least a quarter of the ROI height. This prevents premature bifurcations in the upper airways, where small air pockets may appear to the algorithm as bifurcations.

After the region-growing segmentation is complete (there are no more neighbouring voxels within the threshold, except those removed as leakage), any noise-related holes in the segmentation are removed using a morphological closing operation with a spherical element of radius 5 mm. The closing is applied to each bronchus individually to avoid creating spurious connections between bronchi.

Figure 3.13 shows a segmented airway tree visualised as a surface rendering.

3.4.4 Centreline extraction

We compute a centreline from the extracted airways. This will be used to measure quantities such as airway radius and wall thickness for each airway segment. The centreline is also used in the algorithm to label bronchi according to the lobe they serve, in section 3.4.6 below.

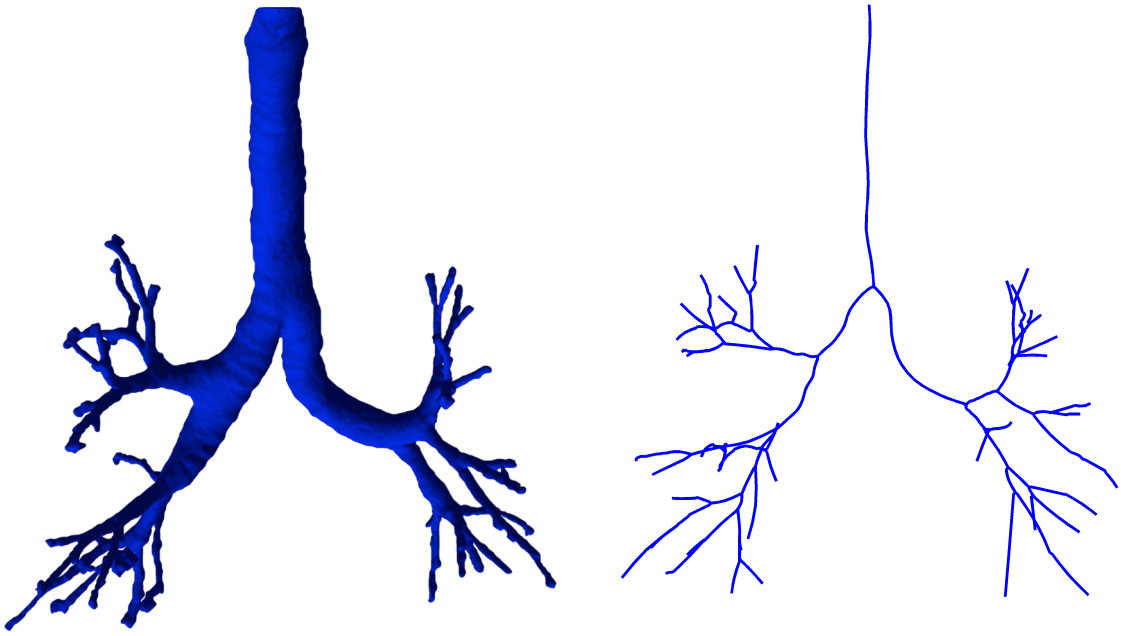

Figure 3.13: 3D visualisation of the airways segmented from CT data. The left image shows the segmented airway lumen as a surface rendering, which has been smoothed to improve visualisation. The right image shows a cubic B-spline fitting to voxels in the airway centreline.

Methods for detecting airway centrelines fall into two categories. The first consists of those which detect centrelines directly from the local properties of the image, for example using the Hessian matrix [106]. These methods may be optimised towards tube detection and hence may perform poorly near to bifurcations where the airway wall structure is more complicated. The second category extracts a centreline from an existing airway segmentation using skeletonisation algorithms, for example the method of Palágyi et al. [115]. We

choose to use a skeletonisation algorithm in order to avoid the bifurcation problems of direct-based approaches, and because we already have a segmented airway tree.

Our method is similar to the generalised thinning method described by Palágyi et al. [115]. The thinning process works by iteratively removing boundary points from the image in each of the 6-connected neighbour directions, but preserving fixed points and topologically non-simple points. A simple point is one that does not cause a change in topology if it is removed from the skeleton. We determine whether a point is simple using an implementation of the fast algorithm described by Malandain and Bertrand [116]. Our thinning method differs from that of Palágyi et al. [115] as we make use of anatomical knowledge to define known fixed points: one at the trachea, and one at each endpoint of the airway tree. This ensures that our resulting skeleton has the correct number of branches and bifurcations. By contrast, Palágyi et al. [115] use no a priori information about the airway tree, which can result in the loss of some airway branches in addition to the creation of spurious branches which must be removed using heuristic methods.

Due to the discrete voxel nature of the skeletonisation procedure, the output is not a smooth curve. To produce a more anatomically realistic centreline, a final step is to smooth the centreline. When computing the airway radius in section 3.4.5 below, we do this by fitting a cubic B-spline to the centreline voxels. We show a smoothed airway centreline result in Figure 3.13. An additional final step would be to prune the small terminating branches.

3.4.5 Radius calculation

We compute a radius for each segmented bronchus. We use these radii in our lobar labelling method (section 3.4.6), which is required by our lobe segmentation algorithm (chapter 4). Airway measurements have been shown to correlate with disease [37,38,39,40,41], and could therefore contribute to regional measures of lung disease, although we do not explore this possibility here.

We adopt the widely used full-width-at-half-maximum (FWHM) method (described below) for estimating the airway radius. There are known issues with the FWHM method: it tends to overestimate the outer airway boundary and underestimate the inner boundary [117]. It also assumes a circular geometry for the airway cross-section, which is not true in general [118]. However, for our purposes (the lobar labelling of section 3.4.6) we require only relative measurements of radius. The FWHM algorithm is an efficient and accurate method for this purpose. We choose to measure the inner radius, i.e. the lumen size excluding the airway walls, as the outer radius is more difficult to detect due to adjacent high-density structures, which may result in inaccuracies in outer radius measurements.

We initially construct a radius approximation for each bronchus. To do this, we take the segmented airway lumen and compute a 3D distance transform to the exterior of the lumen. The radius approximation for each bronchus is found as the maximum value of the distance transform within that bronchus. This radius approximation is used to limit the search space used for the computing the true radius using the methods below.

For more accurate computation, we return to the CT image and use the FWHM method (described below) to compute radius values at points along the centreline of each bronchus. Radius measurements can be unreliable close to

bifurcations, so we ignore points near the ends of each bronchus by considering only the central 50 % of the centreline.

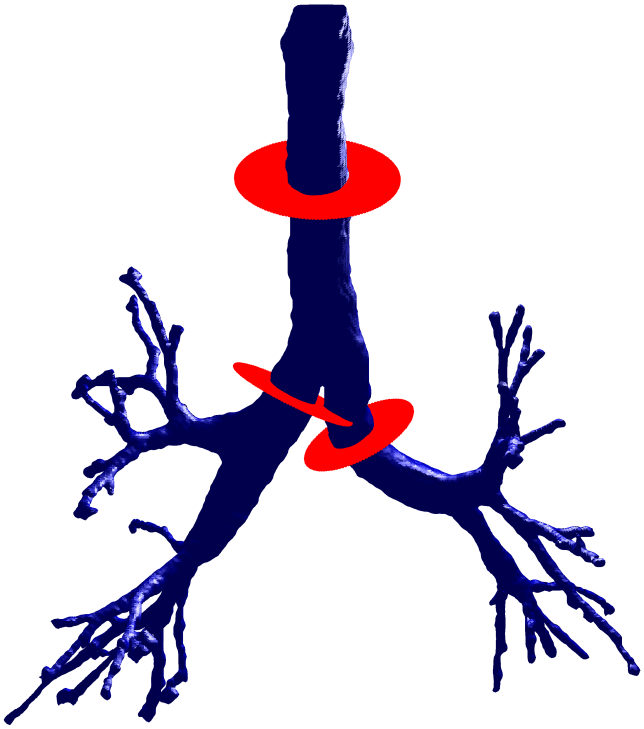

Figure 3.14: 3D surface rendering of the segmented airway lumen, showing the radial grids used to compute the airway radius. This diagram shows the grids (red) for the trachea and main bronchi, but the method is applied to all bronchi in the airway tree. The grids are constructed in planes perpendicular to the bronchi, with the origin on the bronchial centreline. The CT volume is interpolated onto the grid points. This is used to compute the inner radius of each bronchus by detecting the inner airway wall using the full-width-at-half-maximum (FWHM) method.

For each centreline point being considered, we compute the bronchus orientation by approximating the centreline as a straight line at that point. We then construct a plane perpendicular to the bronchus. Within the plane, we construct a radial grid in polar coordinates r, θ, with the centreline point as the origin (see Figure 3.14). The maximum value of r is twice the radius approximation computed above – this limits the search space to the immediate neighbourhood of the bronchus. The grid spacings in r, θ are chosen so that the Cartesian distance between any two grid points is less than half the minimum voxel size. In this way, the grid size is adjusted to each bronchus,

ensuring that the airway radius calculation makes use of all available voxels without oversampling, which would increase execution time.

The image volume is resampled to the radial grid using cubic interpolation (Figure 3.15). Intensity values are constrained to a maximum of 0 HU, since values above this are unlikely to correspond to airway wall tissue and may affect the FWHM calculation. For each value of θ we compute the value of r for which the intensity attains half its maximum value for that value of θ. The final radius value is found by averaging all values of r for all centreline points.

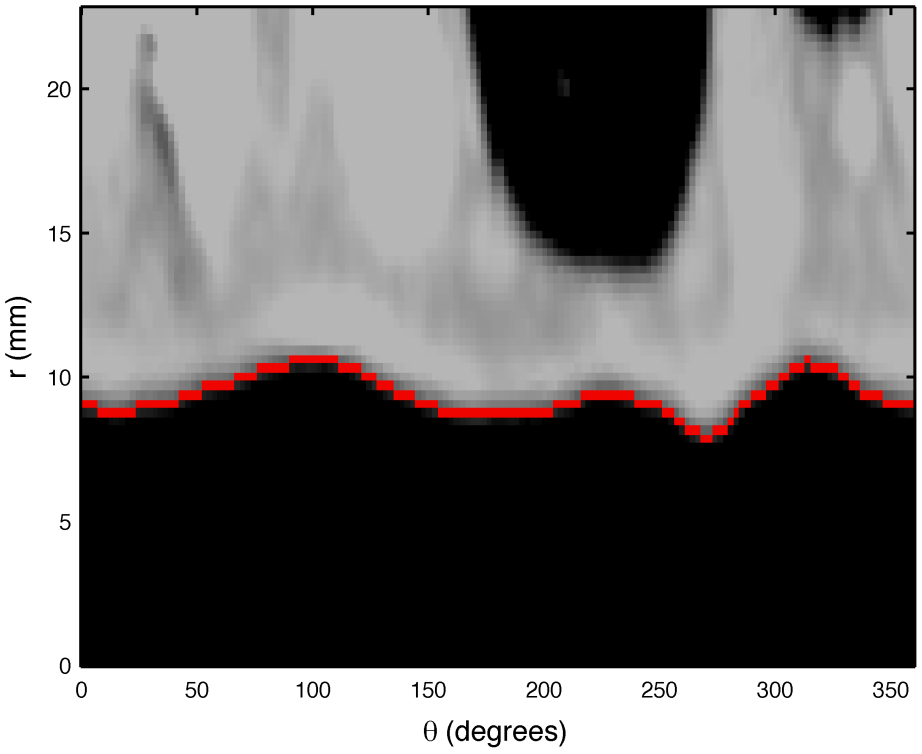

Figure 3.15: Stretched-out airway cross-section illustrating computation of the inner airway wall radius using the full-width-at-half-maximum (FWHM) method. The 3D CT volume is interpolated onto a radial grid (r, θ) constructed on a plane perpendicular to the bronchus with the centreline point as the origin. The image shown here is for the trachea radial grid illustrated in Figure 3.14. Each value of θ represents a radial line starting at a point on the airway centreline. Along each radial line (ray), the inner wall location (red) is estimated by finding the inner location at which the image intensity achieves half of its maximum value along that ray. The bronchus radius is estimated by computing the mean of these values, after repeating for all the points in the central 50 % of the bronchus centreline.

3.4.6 Labelling of the airways by lobe

Our lobe segmentation algorithm in chapter 4 requires us to determine which bronchi serve each of the five pulmonary lobes. In existing work this lobar labelling is often performed manually, but a few authors have developed automatic methods by using a knowledge-based approach where the branching structure of the bronchial tree is compared to a standard textbook model, with heuristics to deal with segmentation errors and minor variations in anatomy [101,119,120]. Here we describe a simple algorithm we have developed for automated lobar labelling which makes use of the airway centreline and radius information. Our algorithm is based on finding the bifurcation points in the airway tree which separate the airways into subtrees serving each lobe (see Figure 3.16).

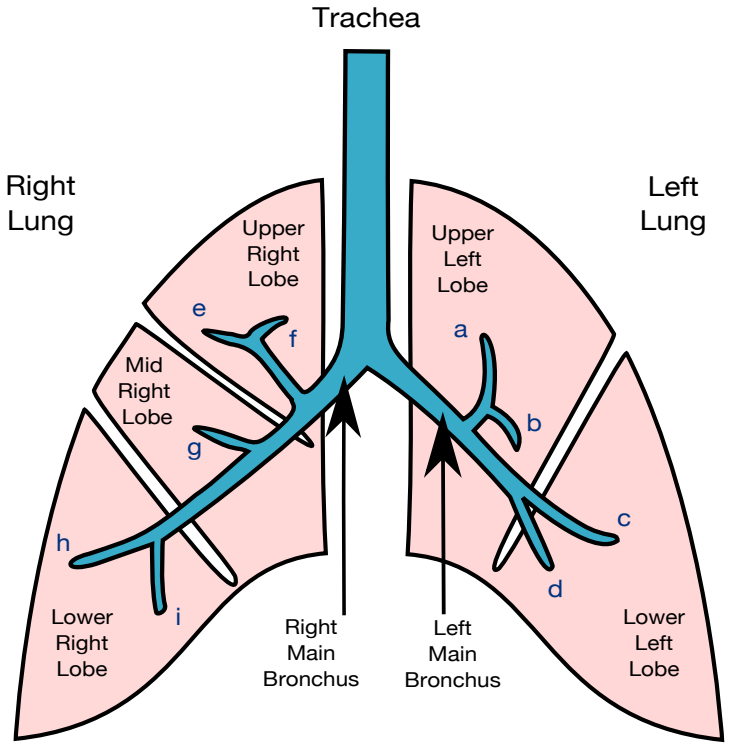

Figure 3.16: Diagram showing the branching structure of the main airways in relation to the pulmonary lobes. The trachea divides into the left and right main bronchi. In the left lung, two subtrees serve the upper left lobe (a, b) and two serve the lower left lobe (c, d). In the right lung, two subtrees serve the upper right lobe (e, f), one subtree serves the middle right lobe (g) and two subtrees serve the lower right lobe (h, i). Our lobar labelling algorithm aims to determine the location of these subtrees.

Separation of left and right lungs

We first label the main left and right bronchi. This is simple since they always bifurcate from the trachea at the first bifurcation point in the airway tree, serving the left and right lungs respectively. To determine which is left and right, we consider the two centreline trees arising from the first airway bifurcation and label them according to the relative position of their centroids along the left-right axis.

Determining the branching structure beyond the main bronchi is less straightforward, due to anatomical variation between subjects and the proximity of bifurcations in certain parts of the tree, which leads to ambiguity about their anatomical ordering. The separation of the bronchi into two lobar trees may occur at more than one bifurcation point (see Figure 3.17).

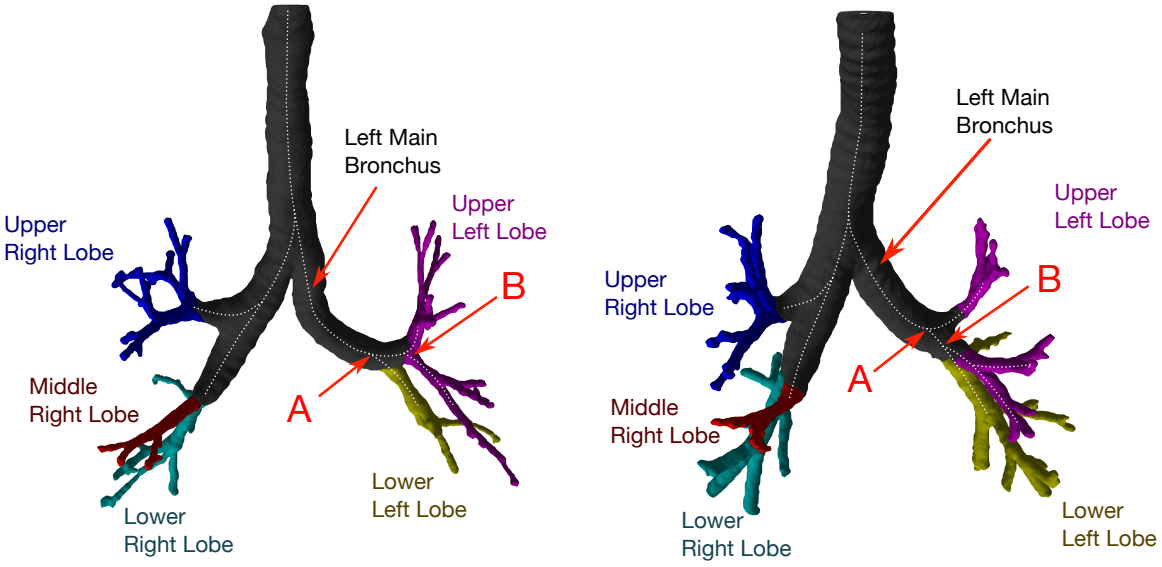

Figure 3.17: 3D surface renderings of two airway trees automatically labelled according to the lobes served by each bronchus. The airway colours are: magenta: upper left lobe; yellow: lower left lobe; blue: upper right lobe; cyan: lower right lobe; red: middle right lobe. Grey bronchi serve more than one lobe. A white dotted line illustrates the centrelines of the main airways. In these two airway trees, the separation between upper left lobe and lower left lobe occurs at different topological points. In the left image, this separation occurs at the first bifurcation point of the left main bronchus (point A), while in the right image, the separation occurs at the second bifurcation point (point B). This illustrates why lobar labelling cannot be performed based on the topology alone.

A simple textbook model which searches for a single bifurcation point to divide the lobes will fail in these cases. To deal with this, our algorithm permits lobar separation to occur at one or more bifurcation points. This results in a number of possible permutations; to choose the correct one we exploit anatomical knowledge about the number of large branches expected at each lobar bifurcation.

Left lung

In the left lung, the main left bronchus divides (via one or more generations) into four bronchi, two serving each of the lobes. These bronchi have a much larger radius than any of their children. Of the two bronchi serving the upper left lobe, one serves the upper segments and is directed towards the apex of the lung, while the other serves the lower segments (the lingula) and is directed towards the base of the lung. Using this knowledge, we use the following heuristic to label the bronchi in the left lung:

- Starting at the left main bronchus we find all the possible divisions of the bronchial tree into exactly four subtrees (labelled a–d in Figure 3.16).
- For each of these permutations, two subtrees are selected for the upper lobe (a, b) and the remaining two for the lower lobe (c, d). This selection is based on finding the centreline of the subtree and computing its centroid. We then compute the distance of each centroid to a plane which passes diagonally from the superior-posterior edge of the image to the inferior-anterior edge of the image. This plane roughly approximates the direction of the left oblique fissure and therefore the distance to this plane can be used to separate the two bronchi serving the upper lobe from the two serving the lower lobe.
- One of the two subtrees selected for the upper lobe should be directed towards the base of the lung, since it is serving the lower segments of the upper lobe (lingula). We verify this by comparing the distance along

the superior-inferior axis from the start of the subtree to the most apical and most basal points of the subtree. If the basal distance is greater, we assume the subtree is directed towards the base of the lung. If neither subtree is directed towards the base of the lung, we reject this permutation.

- Of the remaining permutations, we choose the one for which *all* of the supplying bronchi have a large radius value. To do this, we consider the smallest radius value of the four bronchi in each permutation. We then choose the permutation which has the maximum value of this smallest radius. The aim of this step is to eliminate anatomically unrealistic combinations of large bronchi serving one of the lobes with small bronchi serving the other.

Right lung

For the right lung, the process is slightly different due to the differing anatomy. In particular, there is a significant separation in the bronchial tree between the upper lobe and the other two lobes, which allows us to perform the labelling in two stages, separating the upper lobe first and then the two other lobes.

- We first find the division of the bronchial tree into that supplying the upper lobe and that supplying the other lobes. Starting at the right main bronchus we find all the possible divisions of the bronchial tree into exactly three subtrees.
- For each of these permutations, two subtrees are selected for the upper lobe (labelled e and f in Figure 3.16) and the remaining one for the other lobes. As above, this selection is based on computing the distance of each subtree centroid to the plane which passes diagonally from the superior-posterior edge of the image to the inferior-anterior edge of the image. As with the left fissure, this plane roughly approximates the direction of the right oblique fissure.

- The permutation is chosen based on radius information in the same way as for the left lung.
- We now separate the bronchi for the middle and lower lobes. We use the same algorithm as above, but in this we select one subtree for the middle lobe (g) and two for the lower lobe (h, i). To choose between them we use the centroid separation to a coronal plane. This provides a rough approximation of the part of the right oblique fissure which separates the middle right and lower right lobes.

Two example results from the lobar labelling process are shown in Figure 3.17.

3.4.7 Results

We computed airway tree segmentations for ten of the CT datasets described in section 1.5. Table 3.6 shows the total number of bronchi detected for each subject, and the number of bronchi detected in each lobe. The lobar counts exclude bronchi that supply more than one lobe (for example, the trachea and left and right main bronchi). Hence the total is greater than the sum of lobar counts.

Table 3.7 shows a count of bronchi in each generation of the segmented tree, starting from the trachea (generation zero). We compare these results with the number of bronchi that would exist in a perfectly bifurcating bronchial tree (i.e. powers of two). A number lower than this may indicate that some branches in this generation have not been detected. A generation count higher than the ideal may occur as a result of a trifurcation (as in subject 3). This can also be interpreted as two bifurcations occurring in close proximity. The generation numbers are derived from the segmented tree and do not take into account any possible mismatch with the true tree resulting from a segmentation error. For example, if the algorithm completely fails to detect a bronchus, then it will

not insert a corresponding bifurcation point in the segmentation. The generation numbers of branches distal to that point will then be out by one.

Table 3.8 shows the computed lengths and radius values from the segmented airways. The results are shown for each generation, as a mean value across all of the segmented bronchi in the ten datasets. We also show the standard deviation to give a measure of the variability. For comparison, typical reference values are quoted from Hlastala and Berger [17]. We discuss these results in the following sections.

Table 3.6: Airway segmentation results for CT datasets showing the total number of bronchi segmented and the number found in each lobe. Bronchi beyond generation nine are excluded from the counts.

Subject	Number of bronchi segmented per lobe					Total
	LU	**LL**	**RU**	**RM**	**RL**	
1	52	26	47	28	49	209
2	59	18	45	21	36	186
3	37	48	39	17	34	182
4	46	32	57	25	42	209
5	30	46	29	7	29	148
6	43	42	47	15	36	190
7	8	6	27	3	6	57
8	44	76	52	23	52	254
9	28	22	15	7	18	97
10	28	37	19	3	16	110

Table 3.7: Airway segmentation results for CT datasets showing the number of airways in each generation of the segmented tree. The generation number refers to the segmented tree, starting at the trachea with generation zero. The ideal case is shown for comparison; it shows the number of actual branches that would exist in a perfectly bifurcating bronchial tree.

Subject	Generation									
	0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
1	1	2	4	8	16	28	40	44	39	27
2	1	2	4	8	16	28	47	34	32	14
3	1	2	4	8	16	33	38	30	30	20
4	1	2	4	8	16	32	40	44	44	18
5	1	2	4	8	16	29	32	30	16	10
6	1	2	4	8	16	30	45	38	32	14
7	1	2	4	8	10	14	10	6	2	0
8	1	2	4	9	18	34	50	56	56	24
9	1	2	4	8	16	26	20	14	4	2
10	1	2	4	8	16	24	20	17	10	8
Ideal	**1**	**2**	**4**	**8**	**16**	**32**	**64**	**128**	**256**	**512**

Table 3.8: Airway segmentation results for CT datasets showing the mean radius and mean length of the airways in each generation. The generation number refers to the segmented tree, starting at the trachea with generation zero. The reference values are based on the values given in Hlastala and Berger (Table 2.1, p.22) [17]. The standard deviations of the radius calculations are also shown for each generation.

Airway generation	Radius of bronchus (cm)			Length of bronchus (cm)		
	Reference	**Mean**	**st. dev.**	**Reference**	**Mean**	**st. dev.**
0	**0.90**	0.88	0.11	**12.00**	11.42	1.03
1	**0.61**	0.67	0.11	**4.76**	3.97	1.50
2	**0.42**	0.53	0.11	**1.90**	1.60	0.73
3	**0.28**	0.36	0.11	**0.76**	1.23	0.65
4	**0.23**	0.25	0.10	**1.27**	1.05	0.58
5	**0.18**	0.19	0.07	**1.07**	0.93	0.69
6	**0.14**	0.17	0.06	**0.90**	0.76	0.59
7	**0.12**	0.16	0.06	**0.76**	0.64	0.49
8	**0.09**	0.16	0.07	**0.64**	0.65	0.66
9	**0.08**	0.17	0.09	**0.54**	0.65	0.67

3.4.8 Validation

In the above tables we compare our bronchus count, radius and length results with typical values. We have also performed the following manual validation of our methods:

- Verification of the trachea location for all the ten CT datasets;
- Verification of the labelling by lobe for all ten CT datasets;
- Verification of the first four generations of the segmented airway tree for each of the ten datasets, including ensuring there is no leakage outside the airways and no missing bifurcations in these four generations.

The manual verification of the first four generations (0–3) of airways was performed by visual inspection using our image visualisation software by displaying the segmented airways over the image data (Figure 3.18).

The purpose of our validation is to provide confidence in the detection and lobar labelling of the major airways, which is a required step in the lobar segmentation methods of chapter 4. Ideally, a complete validation of airway segmentation methods requires the manual creation of a gold-standard airway segmentation. However, this would be extremely difficult and time consuming due to the number, small size and three-dimensional geometry of the airways. For this reason, full gold-standard validation is rarely performed in the literature. Even the 2009 Miccai workshop airway challenge [64] did not attempt to manually create a gold standard; a consensus segmentation was created by combining all the submissions with manual correction by clinicians to identify and remove false-positive branches.

For the purposes of this work, we believe the validation methods we have described are sufficient. The primary use of our airway segmentation is as an initialisation stage in the lobe segmentation method (chapter 4). Provided the

lobar labelling is correct (which we have verified) then the segmentation accuracy does not have a critical impact on the results.

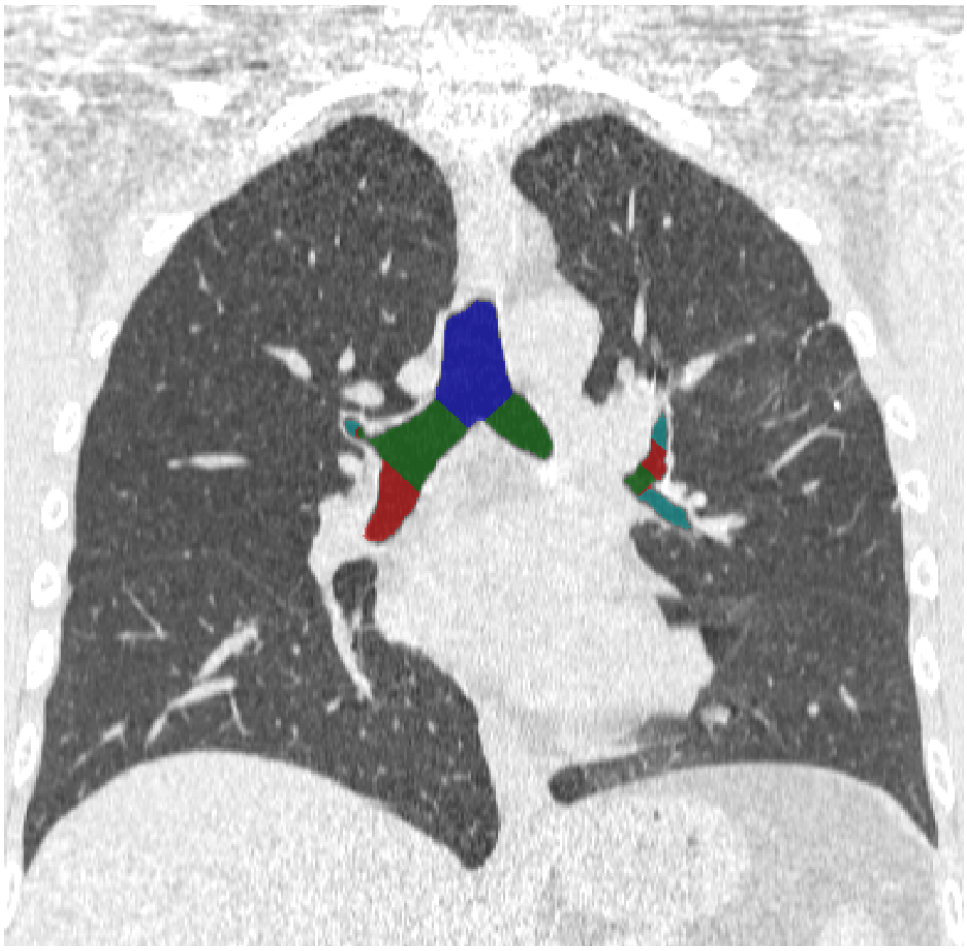

Figure 3.18: Illustration of manual validation of airway segmentation. Manual validation was performed on the first four generations. These were coloured by generation number and displayed as a transparent overlay on the CT image on our image viewer which includes cine and orientation tools. We verified by visual inspection that each bronchus was complete, did not leak into the parenchyma, and that no bronchi in the first four generations had been missed from the segmentation.

3.4.9 Discussion

Using manual validation (section 3.4.8) we verified that airway segmentation in the first four generations was correct, and the airway counts in Table 3.7 suggest that our segmentation detects all branches up to generation four. Beyond this, the results are good in most cases, with some bronchi not being

detected for a number of reasons, including partial volume effects making the airway walls or lumen unclear, the airways becoming too small to be visible or region-growing leakage resulting in removal of a segmented bronchus. Nevertheless, results suggest that the majority of bronchi are detectable down to generation six, consistent with expected results for region-growing [65], and some bronchi are visible to generation nine.

The generation numbers in Table 3.7 are derived from the segmentation and therefore do not necessarily correspond exactly with the true generation number. In particular, if the algorithm fails to detect a child branch at a bifurcation point, then the remaining branch will simply form a continuation of the previous branch and the generation numbering of distal points in the segmentation will not match that of the real tree. Therefore, these results give us some indication of the sensitivity of the airway segmentation, but not the specificity. The manual validation of the first four generations gives us some confidence in the specificity of the algorithm in the larger branches of the tree.

Visual verification of the trachea location confirmed its correct findings on all ten datasets. The trachea location algorithm is designed to be robust and we expect it to perform well for most scans. Although we did not explicitly validate the radius calculation (which uses the standard FWHM algorithm), Table 3.8 shows that the results are consistent with expected values. We expect radius and length values in the smaller airways to become unreliable as the airway dimensions approach the voxel resolution (a typical resolution for these datasets is 0.68 mm × 0.68 mm × 0.63 mm). The radius algorithm could be extended to measure the airway wall thickness, if this was required. This is done by detecting the outer airway wall in addition to the current inner wall value. However, this calculation is more difficult because the outer wall boundary is often unclear due to adjacent high-density tissue.

We have further confirmed by visual inspection that the automatic allocation of bronchi to lobes (lobar labelling) performed successfully for all airways in all datasets. This algorithm is robust to some of the common variations in bronchial anatomy, such as the location, ordering and orientation of the bifurcations. However, it does require correct segmentation of the principal bronchi serving each lobe, and the identification of the bronchus serving the lingula of the upper left lobe. Lobar labelling could still fail for poor-quality datasets or subjects with unusual anatomy or pathology that affects the airway segmentation. For these reasons we will consider introducing a manual verification and correction stage. This would be implemented in a simple manner through the user interface, where the user can change the lobar assignment of a bronchus (and all its descendent bronchi) with a mouse click. Another possibility is to compute the lobe segmentations (chapter 4) for all possible permutations of the lobar labelling, and develop a comparative metric to evaluate the resulting lobe segmentations and select the best one. Such a metric could for example be based on the output of the *fissureness* filter (which we introduce in section 4.3.5) near the lobar boundaries.

3.4.10 Performance

Table 3.1 shows the execution time for the airway segmentation algorithms. The algorithms are: detecting the location of the trachea; segmenting the airways; computing the centreline and radius of each airway; and labelling the airways by the lobe they serve. Times are recorded for ten datasets on a laptop computer (2.66 GHz MacBook Pro running OS 10.8 and Matlab 8). For comparison the number of airways detected and volume of the segmented lung are also shown for each dataset.

Table 3.9: Execution time for the airway segmentation algorithms.

Subject	Lung volume (cm^3)	Airways found	Algorithm execution time (s)			
			Trachea	**Airways**	**Radius and centreline**	**Labelling by lobe**
1	3167	217	4.4	22.6	55.6	3.9
2	4890	183	5.1	22.9	65.6	3.3
3	6208	181	4.5	23.2	91.2	3.9
4	4343	193	5.8	25.0	70.1	3.7
5	3293	138	4.5	21.7	71.4	3.1
6	5436	189	4.2	22.8	73.7	3.7
7	4894	53	5.6	14.1	66.1	2.0
8	7247	251	5.5	28.8	101.6	4.7
9	3374	97	4.0	14.5	48.2	2.1
10	3467	100	4.8	14.7	54.3	2.4
Mean			**4.9**	**21.0**	**69.8**	**3.3**

3.5 Conclusions

In this chapter we have presented algorithms for segmenting the lungs from CT (section 3.2) and proton MRI (section 3.3) datasets. We have also segmented the airways from CT data and computed the bronchial centrelines, radii and their allocation to lobes (section 3.4). These segmentations are prerequisites for the remainder of work in this thesis; the CT lung and airway segmentations are used in the lobar segmentation of chapter 4, while the proton MRI lung segmentations are used along with the CT lung segmentations as part of the registration methods for regional xenon MRI analysis (chapter 5). Airway segmentations will also contribute towards future work in airflow simulation (chapter 6).

Our CT lung segmentation method is based on a region-growing method with removal of the airways and morphological operations to smooth the lung contours and fill holes. This method is well established and has been widely validated in the literature. For proton MRI segmentation, prior work is more limited and we have developed a novel two-stage algorithm. An initial segmentation is generated using a 3D region-growing method with an adaptive threshold. After morphological smoothing and hole filling, boundary refinement is performed on each coronal slice using a 2D level set method with a region term, an image gradient stopping criterion and a curvature-based regulariser. Validated results for five datasets showed a performance comparable to existing work in the field, despite the larger slice thickness of our images causing high errors at the edges of the images. We designed this method for the thick slice MRI datasets available to us as part of matched proton and xenon MRI data. However, this method can be adapted to thinner slice MRI data from the clinical trial, which could allow us to improve our results by allowing us to detect sharper boundaries and implement the level set refinement in 3D.

Our CT airway segmentation is based on a region-growing method with an explosion control heuristic. While this method is widely used, we have introduced a novel adaptation of a variable thickness wavefront which makes the segmentation more robust to noise and eliminates the need for pruning of spurious branches. We implemented standard centreline and radius computation based on skeletonisation and full-width-at-half-maximum analysis respectively. Finally, we implemented a novel method for labelling bronchi according to the pulmonary lobes they serve, which is robust to common variations in the anatomy of the bronchial tree. We performed a validation of the airway results on ten CT datasets, based on number of bronchi detected, assignment to the correct lobe and manual verification of the trachea detection and segmentation of the first five generations.

4 Lobe segmentation

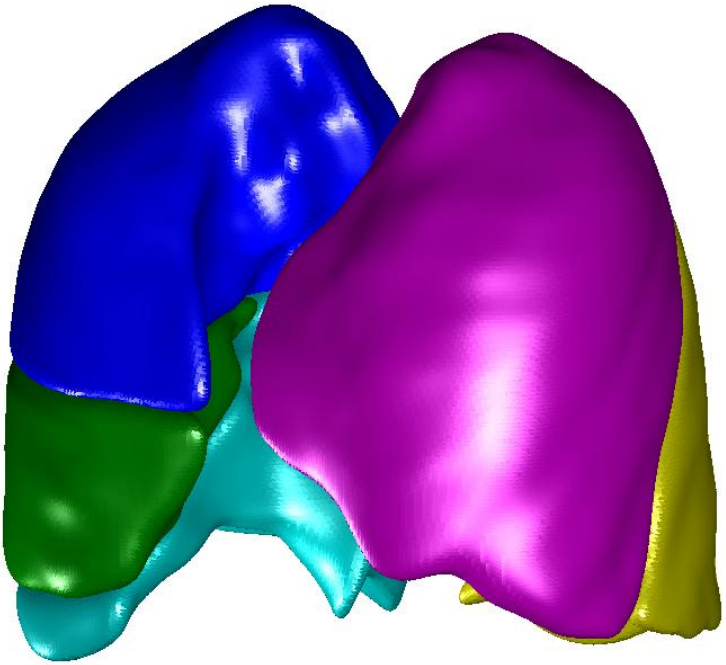

Chapter 4: Lobe segmentation

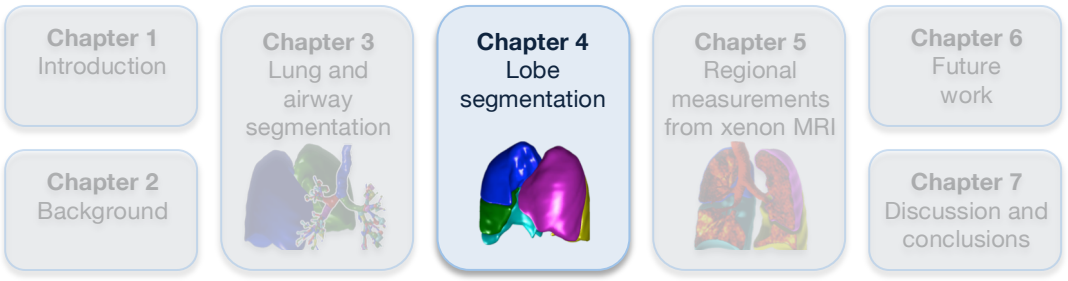

This chapter describes the automatic segmentation of pulmonary lobes from clinical CT data. Our aim is to create a lobar map of the lungs, which we use to perform lobar-based regional analysis of hyperpolarised xenon-129 images in chapter 5. As we discuss in section 4.1 below, the lobes are natural candidates for regional analysis, since they are largely anatomically and functionally independent, and many lung diseases such as emphysema act at a lobar level. We also plan to use these lobar segmentations in future work, to help build personalised airflow models (chapter 6).

In section 4.1 we introduce lobe segmentation and its potential for disease assessment and treatment planning. Lobe segmentation is an active area of research, and in section 4.2 we review published work. We develop our own fully automated lobar segmentation method in section 4.3. Our method includes novel adaptations to address several limitations in previous methods. We presented this work at the ISBI 2012 conference as a poster and conference paper [8]. In section 4.4 we apply this method to clinical CT data and present results validated by a pulmonary clinician. In section 4.5 we discuss the results, and in section 4.6 we present conclusions for this chapter.

4.1 Introduction

The human lung normally divides into five lobes (see section 2.1). Lobe detection is clinically important because many lung diseases act at a lobar level. For example, emphysema [121], postprimary tuberculosis [122] and silicosis [123] usually affect the upper lobes, while idiopathic pulmonary fibrosis is commonly associated with the lower lobes [124]. Regional measurements could help with phenotyping and assessing the severity of these diseases using clinical imaging [28]. Lobe detection has further applications in surgical treatment planning. For example, lung volume reduction surgery (LVRS) offers improved outcomes for patients with upper lobe emphysema, but may not be appropriate for those with lower lobe emphysema [125]. A key stage in lobe segmentation is detection of the pulmonary fissures that separate them (see Figure 4.1). In isotropic CT images (where the axial slice thickness is similar to the pixel dimensions), large parts of the fissure are visible as a locally bright sheet crossing the 3D volume of the lung. In MRI, pulmonary fissures are much more difficult to detect. This is due to the low proton density inside the lung and susceptibility artefacts (magnetic field distortions) caused by the many air-tissue interfaces within the lung [28]. For this reason we perform lobe segmentation only on CT data. For lobar analysis of MRI data, we segment the lobe mask from CT and then use an image registration technique to apply this mask to MRI data (see chapter 5).

The current gold standard for CT lobar segmentation is the manual tracing of lobar boundaries by an experienced pulmonary radiologist. However, this is an extremely laborious task, which could take several hours per subject. As a result, manual lobe segmentation is rarely performed in clinical practice and clinicians instead rely on subjective visual inspection. It is not surprising therefore that there are high hopes and considerable interest in computational

techniques for extracting the lobes automatically (no user interaction) or semi-automatically (minimal user interaction).

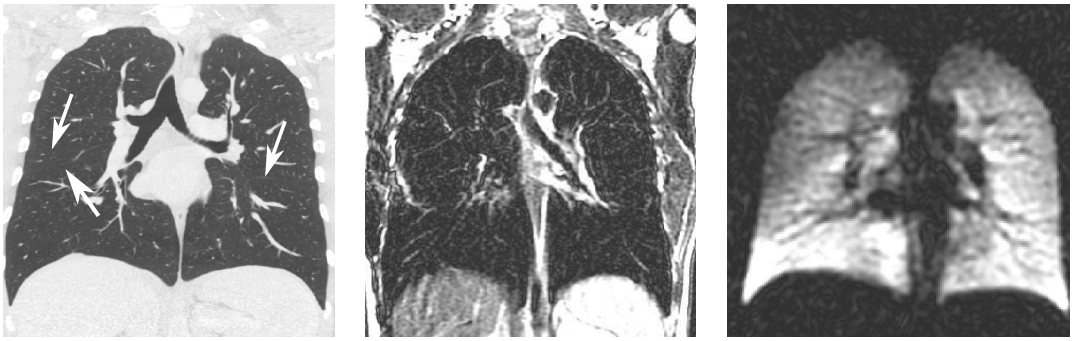

A B C

Figure 4.1: Comparison of CT, proton MRI and hyperpolarised xenon MRI data for detecting pulmonary fissures. A: 2D reconstructed coronal image from an isotropic CT dataset. The fissures are clearly visible as shown by the white arrows. B: Coronal slice from a proton MRI dataset. Fissures are hard to detect due to the low proton density, susceptibility artefacts and partial volume effects of the thick (15 mm) slice. C: Coronal slice from a hyperpolarised xenon MRI dataset. Fissures are not visible at all, although their locations can sometimes be inferred as lying approximately between ventilated regions.

Computational lobe segmentation is usually based on segmenting the lungs and the three main pulmonary fissures. Lung segmentation methods are well established and were discussed in chapter 3. The segmentation of fissures is more challenging, despite their visibility in CT. Fissures can exhibit similar density values to other lung features, and this is also true for higher-order metrics such as the structure tensor and Hessian matrix. Fissures frequently appear blurred due to partial volume effects and subject motion. Furthermore, fissures vary between subjects, they can be incomplete, and lobe-separating fissures can be confused with other minor and accessory fissures which are often visible on CT scans [15,126]. These challenges make fissure and lobe segmentation an ongoing area of research.

4.2 Previous work

Several automated and semi-automated published methods aim to segment the pulmonary lobes from CT data [71,127,128,129,130,131,132,133,134]. We summarise each method here and then compare them in section 4.2.2. In section 4.2.3 we discuss the limitations of existing work which have motivated the development of our own method in section 4.3.

Van Rikxoort et al. [127] use an atlas-based registration to segment the pulmonary fissures. A template is selected from an atlas of manually segmented images, and registered to a cost image based on lung boundaries and fissure points detected using a supervised filter. They also construct an approximate lobar segmentation from the airways, which they use to guide the final registration in regions where no fissure points are found. Quantitative validation is performed for 20 datasets, by manually placing a number of marker points on the true fissure curves and computing the distances to the nearest points on the fissure segmentation. A subjective assessment is also performed for 120 datasets, by visually assessing the quality of each lobar segmentation and awarding a five-point score.

Zhou et al. [131] initialise their fissure search using the assumption that blood vessels are connected within lobes but do not cross the lobar boundaries. Vessels are segmented using a region-growing method, and then allocated to lobes using information from the airway tree. The gaps between lobar groups then define regions of interest for the fissures. Edge-detection filtering is used to find fissure points, which are connected using morphological operations. Lobe separation is only performed where complete fissures have been found. Quantitative validation is performed by manually tracing lobar boundaries on seven datasets, chosen because the fissures are clearly visible. They compute

coincidence degree percentage measures for the volumes, and average distance separation for the boundaries.

Kuhnigk et al. [71] use the vasculature in a similar way. However, instead of assigning vessels directly to lobes using the airway tree, they perform a watershed transform [82] on a search image combining intensity with a distance transform to the segmented vessels. The lobar allocation is based on marker points derived from the bronchial tree. The resulting fissure lines broadly follow the gaps between vessels but track high-intensity fissure lines if they are close enough. To improve fissure accuracy, Lassen et al. [132] add fissure detection based on eigen analysis of the second-order Hessian matrix for each voxel. To evaluate the results, reference segmentations are generated using a semi-automated method, and Dice coefficients [96] are computed for 42 subjects.

Zhang et al. [128] use an atlas to initialise a fuzzy reasoning search for regions containing the oblique fissures. A graph search then finds optimal fissure curves by searching for ridge-like structures on 2D slices using the MLSEC-ST approach [135], based on eigen analysis of the first-order structure tensor. Manually placed marker points are required to segment the right horizontal fissure. Validation is performed by computing the distance from the segmented oblique fissures to manually defined gold standards on a selected subset of image slices for 22 datasets. In later work, Ukil and Reinhardt [130] replace the atlas initialisation with a vasculature-based watershed transform similar to that used by Kuhnigk et al. [71], but they use a Hessian approach to segment the vessels [136]. The resulting watershed boundary is used to define the region of interest for the final graph searching stage. This is different to Lassen et al. [132], who incorporate fissure detection directly into the watershed transform. Ukil and Reinhardt [130] present results for 29 subjects, although manual correction was performed in some cases to obtain correctly the right horizontal

fissure. Fissure segmentations are evaluated using manually placed marker points, similar to the method used by Van Rikxoort et al. [127].

Wei et al. [133] use estimated fissure directions to search for a planar region of interest free of vessels and airways. Fissures are then detected using a 2D wavelet transformation and 2D graph searching, with fissure lines extrapolated between slices. Manually placed marker points are again required to segment the right horizontal fissure. Results are validated against a manually segmented gold standard for nine subjects. Ohkawa et al. [137] use a similar idea, fitting fissure planes between possible vessel and bronchial structures detected using local maxima filtering. However, the method is only applied to two subjects and validation results are not presented for the right horizontal fissure. Further work is presented by Kobashi et al. [138].

Some methods do not depend on auxiliary pulmonary structures such as vessels and airways. However, some form of a priori information is generally still needed to correctly identify and separate the lobes.

Ross et al. [129] use a particle sampling approach to identify candidate fissure points. Fissure evaluation is based on second-order Hessian analysis, similar to that used by other methods [132]. However, rather than directly extracting fissure points, the Hessian is used to generate a vector field which directs the motion of a randomly created system of interacting particles. After a number of iterations, particles tend to concentrate along the fissures. False candidates are removed using morphological analysis. A fissure surface is then estimated by a statistical approach, using a simple manually defined atlas of marker points as a prior. The final fissure curves are computed by fitting a smooth thin plate spline to the remaining points. Results are presented for six datasets, computing distances from the automatic fissure curves to gold-standard segmentations by two radiologists.

Pu et al. [134] apply a radiodensity threshold to detect candidate fissure points which they use to build a 3D geometric mesh. Successive smoothing operations reduce the mesh to a series of approximate fissure patches, which are separated and labelled according to their orientation. This approach is very sensitive to the fissures having the expected orientations, and large patches of genuine fissure points are removed by the process. However, this is corrected for using radial basis functions to extend the detected fissure patches to form a complete fissure surface. Qualitative results are presented for 65 subjects, using a subjective three-point score covering all five lobes; they do not report results for individual lobes. For quantitative validation, they quote a previous evaluation of their fissure segmentation [139] where they compute mean distances from radiologist-placed markers. However, the latter results cannot be compared directly with other methods [127,130] because the fissures detected in that paper do not represent lobe-separating curves; they are incomplete fissure sections and the lobes are not separated. A related method with quadratic B-spline surface fitting is introduced by Gu et al. [140].

Lobe segmentation has been used in applications such as bronchial labelling [141] and model-building [142]. Some authors extend lobar segmentation ideas to extract the pulmonary segments [71,143,144]. However, validation is difficult as the segment boundaries are not generally identifiable on CT scans.

Other methods exist which detect fissures but do not extend this to lobe segmentation. Saita et al. [145] use a vasculature-based method similar to Zhou et al. [131], but they do not perform lobe segmentation and their fissure surfaces can also be incomplete. Wang et al. [146] describe a semi-automated method for extraction of the oblique fissures only. The user segments an initial slice, which is used to iteratively initialise a 2D curve-searching algorithm in neighbouring slices.

4.2.1 Validation methods

As discussed in section 4.1, the gold standard for pulmonary lobe segmentation is the manual delineation of the lobar boundaries on each CT slice by an experienced pulmonary radiologist. Even this is not considered completely accurate, as many authors report small but significant inter- and intra-observer variability [127,129,132]. The main difficulty with manual segmentation, however, is the time required, which makes it impractical for a large number of datasets. Zhou et al. [131] validate only seven of their 44 subjects using this method. Lassen et al. [132] speed up this process by replacing the manual segmentation with a semi-automatic method. Other authors adopt different validation approaches.

Several authors use a subjective visual scoring of the lobar segmentations [127,130,134]. This method provides a fast assessment of the predicted lobar boundary, including extrapolation regions where the fissure is incomplete. However, evaluation is subjective and observer-dependent, and must be repeated for every algorithm to be compared, whereas a gold standard need only be created once for each dataset.

Other approaches concentrate on validation only of the fissure surfaces, since lung segmentation methods have been well validated elsewhere (see section 3.2). Distances are computed between the fissure segmentation and a manually segmented gold standard, either for the whole image [129] or selected 2D slices [128,133,146]. A faster approach is to place marker points on the fissure, rather than segmenting the whole fissure [127,130]. However, this has the drawback that a two-way validation is not possible: a distance cannot be computed from the automatic fissure lines to the manually placed points, since this would be dependent on which slices were chosen and where the markers were placed. Some authors mitigate these issues by performing more than one type of validation [127,130].

4.2.2 Comparison of previous work

Table 4.1 summarises the validation results, execution time and methodology for each published work in lobe segmentation. All the works adopt a similar overall strategy, which can be broken down into four modular stages:

1. Segmentation of the lung using standard methods (see section 3.2.1);
2. Detection of candidate fissure voxels using local image filtering;
3. Use of prior anatomical information (such as airway and vessel segmentations, or anatomical atlases) to reduce the set of candidate fissure points to a more likely subset;
4. Construction of a continuous fissure surfaces from the remaining candidate points, with extrapolation to span the whole lung volume.

The methodologies for stages 2–4 are summarised in Table 4.1. The first stage (lung segmentation) is not included in the table, since all authors use broadly the same approach, similar to that of Hu et al. [66] (see section 3.2.1).

In terms of quantitative results, the best comparative measure we have for these publications is the root mean square (RMS) distance from manually placed fissure points to the automatically segmented lobar boundary. We note that all of these methods perform similarly, giving typical RMS measures below 3 mm. This is comparable to the inter-observer variability.

However, direct quantitative comparison is difficult and of limited value, because these results use a variety of different validation metrics, have different levels of automation (some allow for manual correction) and are based on different sets of data, each of which may present their own unique challenges. True quantitative comparison requires algorithms to be run on the same data, which is not possible because few authors openly publish their code or validation data. However, a methodological comparison is useful for informing our own lobe segmentation method (section 4.3).

Table 4.1: Comparison of published work in lobe segmentation.

Authors	Notes	Number of validated datasets	Results	Time for complete algorithm execution	Methodology: 1. Identifying fissure points; 2. Use of anatomical information; 3. Fitting fissure surface.
Saita et al. 2005 [145]	Do not segment lobes.	20	Fissures: extraction rate measured; varies according to parameters	not recorded	1. 3D Sheet filter; 2. Voronoi division based on distance transform to blood vessels; 3. Morphological operations.
Zhang et al. 2006 [128]	Semi-automatic for horizontal right fissure. Only oblique fissures validated.	22	Fissures: 1.96 ± 0.71 mm RMS error; 0.998 ± 0.011 similarity index (oblique fissure only)	12–16 min (excluding lung segmentation)	1. 2D MLSEC-ST; 2. Atlas of fissures and lobe boundaries; 3. 2D graph searching and fuzzy reasoning with smoothness constraint.
Zhou et al. 2006 [131]	Do not segment lobes if fissures are incomplete	7	Fissures: 0.30–0.68 mm ASED (mean difference); 98% coincidence degree	105 min	1. Edge detection; 2. Voronoi division based on distance transform to vessels; 3. Morphological operations and connected-component analysis.
Wang et al. 2006 [146]	Oblique fissures only. Semi-automatic.	10	Fissures: 1.01 mm RMS difference (oblique fissures only)	< 5 min	1. 2D ridge map based on Sobel operators; 2. Manual initialisation and directed extrapolation to adjacent slices; 3. Shape-based Bayesian curve growing.
Ukil and Reinhardt 2009 [130]	Semi-automatic for horizontal right fissure	29	< 2.7 mm RMS across all parameters and datasets	10–12 min	1. 2D MLSEC-ST; 2. Watershed based on intensity and distance transform to vessels; 3. 3D Graph searching.
Wei et al. 2009 [133]	Semi-automatic for horizontal right fissure	9	Fissures: 1.13 mm mean difference in 2D	< 5 min	1. 2D wavelet transformation; 2. 2D searching for areas free of vessels and airways, or manual marker points, with extrapolation to adjacent slices; 3. 2D graph searching.
Pu et al. 2009 [134]		65	Lobes: 50.8% of examinations rated good or better, 1.5% rated poor or worse	25 min	1. Threshold; 2. Orientation of fissure patches; 3. Radial basis functions.
Ross et al. 2010 [129]		6	Fissures: 2.22 mm mean difference	30–45 min	1. Hessian-based particle sampling; 2. Simple atlas of control points; 3. Thin plate splines.
Van Rikxoort et al. 2010 [127]		120 qualitative; 20 quantitative	Fissures: 1.71 mm RMS error (20 datasets); Lobes: 85/120 datasets rated good or excellent	119 min	1. Hessian analysis; 2. Atlas of fissures and lobe boundaries, plus airway-based Voronoi division; 3. Smoothed registration.
Lassen et al. 2010 [132]		42	Lobes: 0.97 mean Dice coefficient	3–31 min (average 10min)	1. Hessian analysis; 2. Watershed transform based on vessels and airways; 3. Watershed boundary.

The key parts of a lobe segmentation algorithm lie in stages 2 and 3; i.e. the methods for detecting candidate fissure points and differentiating true fissure points from false positives. Points on the fissure surface are identified using local filtering. This is a trade-off between finding as many true positives as possible (sensitivity) and reducing the number of false positives (specificity). A simple intensity threshold [134] gives high sensitivity but poor specificity. Other authors find improved specificity using various edge detectors [131,145,146], 2D wavelet transforms [133], supervised filtering [127] or eigen analysis of the first-order structure tensor [130] or second-order Hessian matrix [129,132] at each voxel.

There is a limit to the specificity that can be achieved with these methods. This is due to inconsistency in fissure appearance caused by anatomical variability, subject breathing and partial volume effects. The filters have to be sensitive enough to capture this variation, but this makes is hard to distinguish fissures from other lung structures which have a similar local appearance, such as airway and vessel walls or fibrotic tissue. As a result, many false positives are generated. The distribution of these false positives is not random, which can make removing them difficult. Furthermore, filtering may detect minor fissures between segments, which must also be removed. Morphological operations such as direction-based connected-component analysis [129,134] can be used to reduce false positives, but these are generally insufficient on their own.

Anatomical knowledge is often used to help locate and delineate the lobar boundaries and remove the false positives. Atlas-based approaches [127,128] guide fissure detection based on similarity to existing human anatomy, but these have not proven sufficiently accurate on their own and so have been combined other approaches. Segmenting the airways and vasculature provides a powerful approach for assisting with fissure detection [71,127,128,130,131,132,133,145]. This is based on the idea that small airways

and blood vessels do not cross lobar boundaries, so fissures should be located in the gaps between airway and vessel trees. However, inaccuracies or mislabelling in the segmentations may require manual correction [130].

Authors who avoid the use of atlas or vasculature-based methods employ simplified anatomical and morphological assumptions applied to statistical [129] or graph searching [134] methods. A major advantage of these approaches is that they avoid some of single points of failure inherent in the "pipeline" approaches. However, a wider range of validation is required to verify that these simplified assumptions will hold for a wide range of data, especially where pathology is present. For example, it may be difficult to geometrically separate the main fissures in the right lung, if the right horizontal fissure is curved so that it is partially parallel to the right oblique fissure. A vasculature-based segmentation is less likely to have difficulty in this situation.

The final stage in lobar segmentation is the use of the detected fissure points to construct continuous, smooth (in general) surfaces to separate the lobes. Watershed and atlas registration methods do this naturally [71,127,132]. Other authors use curve-fitting [129,134], graph searching [128,130] or morphological techniques to complete the fissure surface [131,145]. Of these, morphological techniques are the least reliable, since they are highly dependent on the points initially found and do not guarantee finding a continuous surface in regions where no candidate fissure points have been detected. Curve fitting methods are likely to perform well, provided outlying false positives have been removed; these methods are in general only robust to outliers that follow a normal distribution, which is not the case for the output of a fissure filter. Another form of graph searching involves extrapolating a 2D fissure curve found on one slice to initialise graph searching on adjacent slices [133,146]. These approaches are prone to error propagation if the fissure curve becomes poorly segmented on one slice.

4.2.3 Limitations of existing work

We have examined a number of lobar segmentation methods. Some require manual initialisation [128,133,146] or manual correction [130], which may not lend them to being part of an automated clinical workflow. Certain methods do not guarantee to deal with incomplete fissures [131,145], which means they cannot be directly used in automated lobe segmentation procedures. Other methods are fully automated and naturally provide interpolation and extrapolation to completely separate the lobes [127,129,132,134]. However, most methods have typical execution times in excess of ten minutes, and some much longer [127,129].

Atlas-based approaches can be computationally expensive and require training on manually segmented datasets [127,128]. Furthermore, they may perform poorly on datasets which are not well represented by their training data, such as when the shape and position of the middle right lobe changes in response to disease.

Methods which use the vasculature to detect fissure regions are sensitive to the reliability of the vessel segmentation [130,132,133]. The use of the distance transform means that the predicted fissure locations will tend to track the line equidistant to lobar vessel groups. There is no anatomical reason this should be the case, and it means results are very sensitive to vessels near the fissure boundaries. High-density features such as pulmonary nodules can be incorrectly labelled as vessels due to their similar local appearance. Any vessel-like structures which appear close to the fissure boundaries are likely to disrupt the fissure detection algorithm, since the distance transform pushes the predicted path away (see Figure 4.2). A similar problem occurs with vessels labelled with the wrong lobe. Missing vessels can also cause problems such as deviations in the predicted fissure path or a widening of the extracted regions of interest. The inaccuracies caused by vessel segmentation are responsible

for Ukil and Reinhardt requiring manual correction in some cases [130]. They are also responsible for the unsatisfactory results by Kuhnigk et al. [71], causing Lassen et al. to introduce fissure detection directly into their watershed transform algorithm to help correct for this [132].

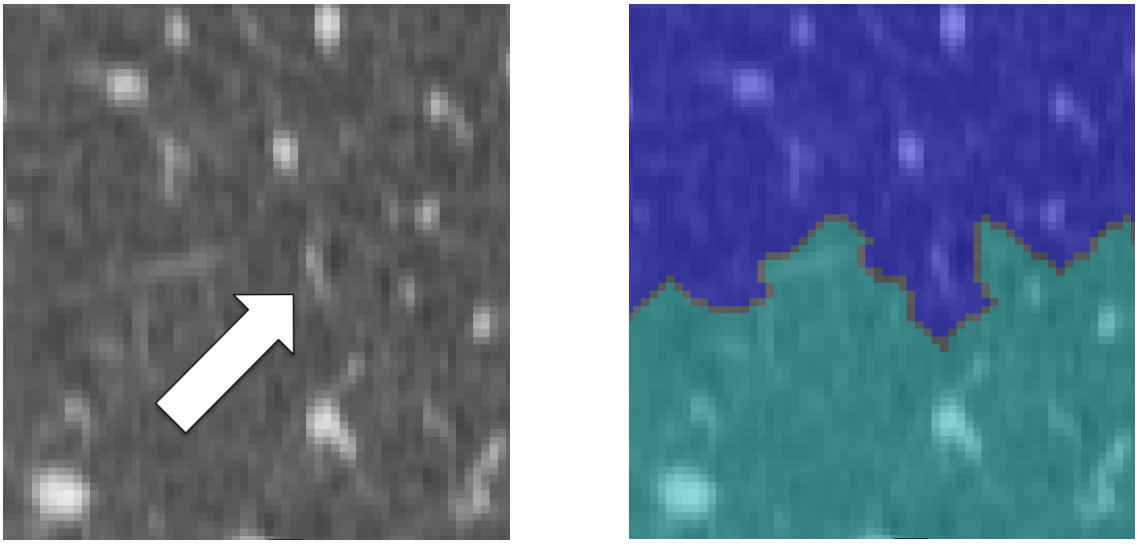

Figure 4.2: Close-up of a coronal CT lung slice showing an example of a small vessel-like structure close to a fissure boundary. Left: These structures (white arrow) can disrupt algorithms which search for fissures in the regions between blood vessels. Right: Result of a lobar segmentation using a watershed transform [82] based on distance transform to segmented blood vessels. The vessel-like structure near the boundary pushes the detected curve away from the true fissure.

It is clear that no lobe segmentation algorithm has yet emerged as a standard method suitable for all situations. The relative modularity of the approaches outlined above suggests an improved method can be constructed by combining features of different methods, with speed, accuracy and automation balanced according to the particular application. In the following section we propose a new method that combines the more successful ideas of previous algorithms with novel adaptations to overcome a number of the current limitations.

4.3 Method

In this section we present a new method for lobe segmentation. Our method addresses some of the limitations of previous work (section 4.2.3): it is fully automated; computationally efficient; requires no training data and is robust to incomplete fissures and small vessels crossing the lobar boundaries. Before describing the method in detail (section 4.3.2), we explain the reasoning for the approach we have taken, which is informed by our analysis of previous work (see section 4.2).

4.3.1 Choice of methodology

Our lobe segmentation strategy is to detect the fissure curves and use these to divide our CT lung segmentations (see section 3.2) into lobar regions. The choice of fissure detection method is therefore a crucial consideration.

In section 4.2.2 we discussed the existing fissure detection methods used in the literature. We compare three of these methods in Figure 4.3. Most methods are based on local filtering of image intensity and its derivatives. Fissures show variability in their appearance, and may exhibit a similar local appearance to other lung structures such as vessel walls. Fissure detection is therefore a trade-off between sensitivity and specificity. Simple intensity thresholds [134] or edge detectors [131,145,146] do not give very high specificity, so we look for higher-order methods involving analysis of the structure tensor [130] or Hessian matrix [129,132]. Our final choice is to develop a Hessian-based *fissureness* filter (see section 4.3.5). Fissureness is analogous to the *vesselness* filter [147] used for blood vessel detection (see section 4.3.4), which is also built from the Hessian matrix. This property allows us to build vessel suppression directly into our fissureness filter. In addition to these intensity-

based methods, phase space approaches such as phase congruency [148] have recently been applied to curvilinear feature detection [149], although not specifically to fissure segmentation.

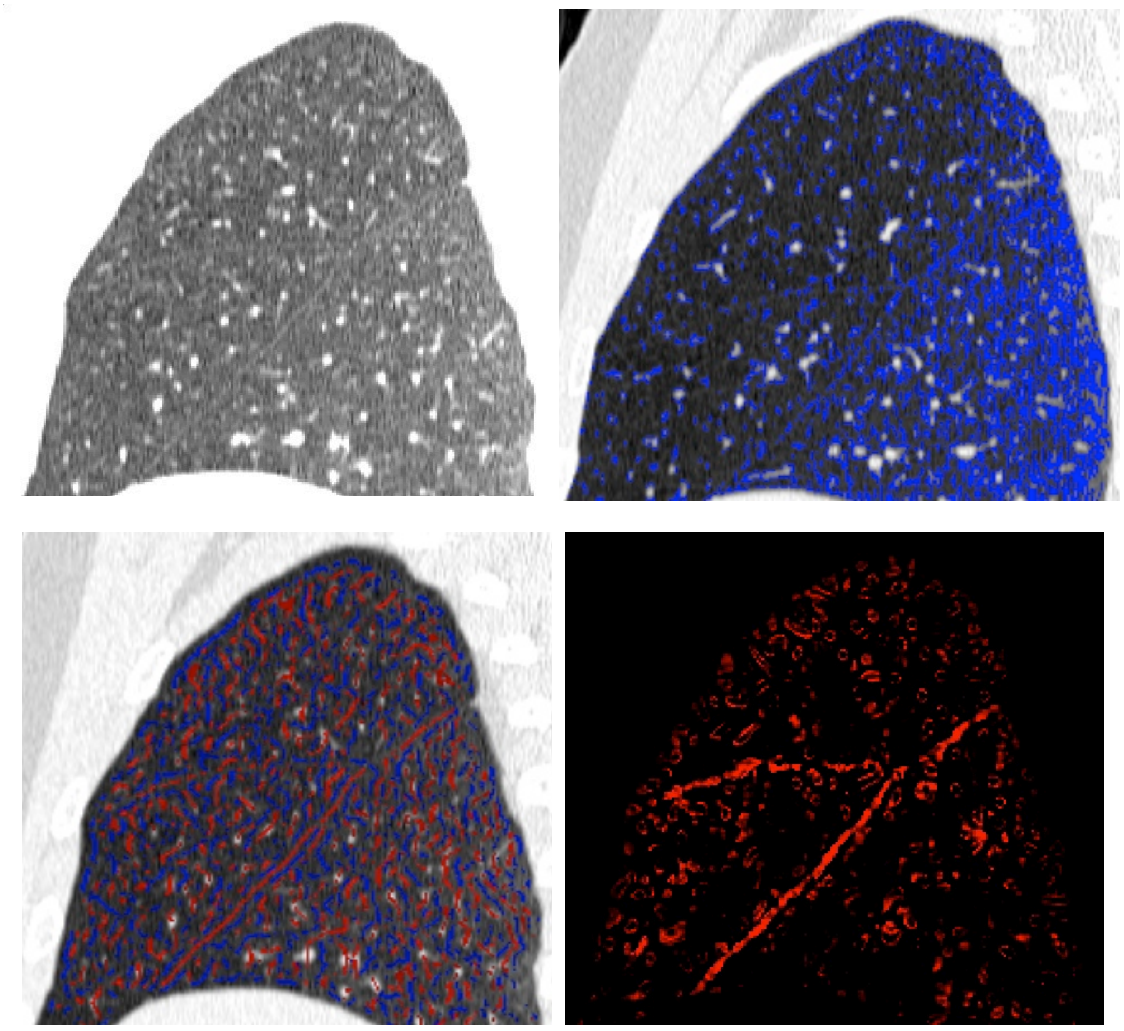

Figure 4.3: Examples of fissure detection methods. Top-left: Sample sagittal slice from a CT density image. Top-right: Corresponding output of a simple fissure threshold operation, where blue shows voxels inside the lung with intensity values in the range –700 HU to –800 HU. This simple thresholding is sensitive but non-specific, which is why higher-order methods are used in an attempt to balance specificity and sensitivity. Bottom-left: Output of a 3D creaseness filter based on an extrinsic curvature level set method (MLSEC-ST) [135]. Red voxels show ridge-like structures, whereas blue show valley-like structures. Bottom-right: Output of a 3D Hessian-based fissureness filter which emphasises plane-like structures (this is part of the fissureness filter we develop in section 4.3.5). The brightness of the image represents the output of the filter, representing the probability of there being a fissure at each point.

Vesselness is an example of multiscale analysis [150]; the filter can be applied at multiple scales to detect vessels of different sizes, with the size of the detected features governed by the chosen scale, not the image resolution. The fissureness is analogous to this, but we only need to apply the fissureness filter at one scale, since the characteristic size of a fissure does not vary significantly. However, the multiscale property is useful because it allows the fissureness filter to work for different image resolutions, and features of the wrong scale are automatically suppressed.

In section 4.2.2 we found that fissure detection is not generally sufficient on its own and needs to be supplemented by prior information in order to correctly select the true fissure points. Atlas methods, which have large training and computational overheads, do not show any significant improvement over anatomy-based approaches. We therefore develop an initialisation method based on airway and vessel segmentations, which is similar to the approach used in other methods [130], but overcomes some of their limitations. The initialisation is used to generate a fissure reference plane, which we use to automatically select and reject the candidate fissure points that will be used to model the fissure surface.

The final stage in lobe segmentation is to use the detected fissure points to create a smooth, continuous fissure surface with which we can segment the lobes. We have not adopted the methods use in previous work, as they are not very robust to possible errors in the choice of fissure points. Instead we apply an approximation method known as *multilevel B-splines* method [151]. This robustly fits an approximation surface to a set of scattered points. If necessary, the final fissure surface is extrapolated to the lung boundaries using thin plate splines. This surface is then used to segment the lobes.

4.3.2 Overview of method

In summary, the method performs the following pipeline:

- Segment and separate the left and right lungs using the methods described previously in section 3.2;
- Segment the airways, find the centreline and label the bronchi by lobe using the methods described previously in section 3.4;
- Compute vesselness and vesselness density filters for detecting blood vessels (section 4.3.4);
- Calculate a fissureness measure based on a combination of the eigenvalues of the Hessian matrix and suppression of points close to blood vessels (section 4.3.5);
- Calculate an approximation to the lobar boundaries using a watershed transform applied to the vesselness density, seeded from the airways labelled by lobes (section 4.3.6);
- Detect fissureness maxima in regions of interest near to the approximate lobar boundaries computed above (section 4.3.7); and
- Compute smooth, final fissure surfaces using multilevel B-spline approximations (section 4.3.8).

This pipeline is illustrated in Figure 4.4, and the steps are discussed in detail below.

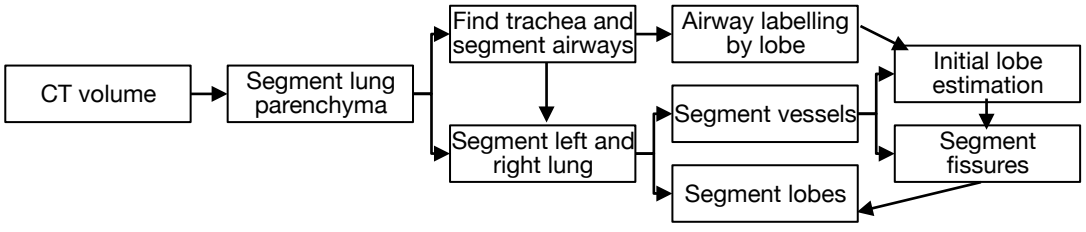

Figure 4.4: Simplified illustration of the automated lobe segmentation pipeline.

4.3.3 Lung segmentation

Our lobe segmentation is guided by anatomical information from the lung boundaries, airways and blood vessels. We described our lung and airway segmentation methods in chapter 3. The next stage is to introduce the vessel detection filters we will use in later parts of the algorithm.

4.3.4 Blood vessel detection

We first introduce a *vesselness* filter, which produces a high response in the presence of the pulmonary blood vessels. This filter serves two purposes:

1. We use the vesselness to derive a *vesselness density* image. This is used to guide our initial lobar approximations, using the prior knowledge that few vessels cross the lobar boundaries;
2. Voxels in the vessel walls can exhibit a locally fissure-like appearance, leading to false positives in fissure detection. Our fissureness filter (section 4.3.5) uses the vesselness to suppress the response from these voxels. Such voxels do not usually exhibit a high vesselness value, but are located close to voxels which do (those in the centre of the vessels). Hence the filter uses a distance transform from voxels of high-valued vesselness to suppress voxels in their vicinity.

A large body of work exists for vessel segmentation [152,153]. In the context of lung vessels, published work is more limited. A recent segmentation challenge, Vessel ’12, was part of a workshop at the ISBI 2012 conference. Most submitted methods used Hessian-based approaches with differing methods to produce connected trees [154,155,156,157,158,159]. Other approaches included variational region-growing [160,161,162] and hysteresis thresholding [163]. No approach emerged significantly better than any others. In general, most methods work well for larger vessels which are well resolved on the

images, and work poorly for small vessels where partial volume effects make detection difficult and the vessels appear to connect with other lung structures with similar local density values.

For our purposes, we do not need to segment a connected vessel tree and we are not interested in the large vessels. We seek to detect the likely presence of small vessels. For these reasons, we turn to a filter-based approach, which can successfully identify a large proportion of the visible blood vessels in the lung but does not attempt to connect them into a hierarchical tree. We adopt a widely used multiscale vesselness filter, which was introduced by Frangi et al. [147] and has been successfully applied to CT and other types of imaging data. This filter is based on second-order voxel analysis of the Hessian matrix, and is outlined in the next section. For each voxel in the image, the Hessian matrix is constructed at a number of different scales and its eigenvalues computed. The ratios of the eigenvalues are used to enhance tubular structures.

The resulting vesselness filter is not a segmentation but a probabilistic image enhancing vessel-like structures. Vessels can be segmented by applying a simple threshold [147] or using hysteresis methods to trace vessel paths [136]. The former approach will produce disconnected vessel segments, while the latter can segment complete sections of the vessel trees. In our fissureness filter (section 4.3.5), we use a simple threshold of 0.1 on the vesselness filter output and compute a distance transform to the resulting segmentation. For our vesselness density filter below, we apply a Gaussian filter directly to the vesselness filter output.

Vesselness

Our vesselness filter is an implementation of the method used by Frangi et al. [147]. The vesselness $\mathcal{V}(s)$ depends on a characteristic length scale s, which relates to the diameter of the vessel structures we aim to detect. The final measure of vesselness $\mathcal{V}_{\max}$ chooses, for each voxel in the image, the maximum value of $\mathcal{V}(s)$ across the range of chosen scale values $\{\text{scales}\}$

$$\mathcal{V}_{\max} = \max_{s\in\{\text{scales}\}} \mathcal{V}(s) \ . \tag{4.1}$$

To compute $\mathcal{V}(s)$, the vesselness at scale s, we first need to compute the Hessian matrix $\mathcal{H}_s$ for each voxel of the image at this scale. We define our image as I and our coordinates $\boldsymbol{x}$ in the basis $(\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3)$. Following [147], we define the Hessian at scale s as the convolution of the original image with the second derivative of the three-dimensional Gaussian filter G

$$\mathcal{H}_s\big(I(\boldsymbol{x})\big)_{i,j} = I(\boldsymbol{x}) * \frac{\partial^2}{\partial x_i \partial x_j} G(\boldsymbol{x}, s) \ , \tag{4.2}$$

where the three-dimensional Gaussian function at scale s is defined by

$$G(\boldsymbol{x}, s) = \frac{1}{(2\pi s^2)^{\frac{3}{2}}} e^{-\frac{\|\boldsymbol{x}\|^2}{2s^2}} \ , \tag{4.3}$$

where $\|\boldsymbol{x}\|$ is the Euclidean norm. We note that scale s is the standard deviation of the Gaussian filter. Next, we compute the eigenvalues λ_1, λ_2, λ_3 of the Hessian matrix for each voxel, and order them such that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. The geometric interpretation of these eigenvalues is defining an ellipsoid describing the local principal directions of curvature. The ratios between eigenvalues therefore give insight into the local geometry. Following this geometric interpretation, the vesselness at each voxel is defined as a product of components

$$\mathcal{V}(s) = \Gamma \mathcal{V}_{\mathcal{A}} \mathcal{V}_{\mathcal{B}} \mathcal{V}_{\mathcal{C}} \,. \tag{4.4}$$

The first factor suppresses points whose largest eigenvalues λ_2, λ_3 are positive. Points in the vessels should have negative values, since vessels are locally bright

$$\Gamma = \begin{cases} 1, & \lambda_2 < 0 \ \text{ and } \ \lambda_3 < 0 \,, \\ 0, & \text{otherwise.} \end{cases} \tag{4.5}$$

The second factor distinguishes between plate-like and line-like structures:

$$\mathcal{V}_{\mathcal{A}} = \exp\left(1 - \frac{\mathcal{R}_{\mathcal{A}}{}^2}{2a^2}\right), \qquad \mathcal{R}_{\mathcal{A}} = \frac{|\lambda_2|}{|\lambda_3|} \,. \tag{4.6}$$

The third factor separates blob-like structures from plane-like or line-like structures

$$\mathcal{V}_{\mathcal{B}} = \exp\left(-\frac{\mathcal{R}_{\mathcal{B}}{}^2}{2b^2}\right), \qquad \mathcal{R}_{\mathcal{B}} = \frac{|\lambda_1|}{\sqrt{|\lambda_2 \lambda_3|}} \,. \tag{4.7}$$

The final factor removes voxels with low eigenvalues, which are likely to correspond to background noise. This is done by computing the Frobenius norm of the Hessian $\|\mathcal{H}\|_F$ which can be expressed in terms of the eigenvalues

$$\mathcal{V}_{\mathcal{C}} = \exp\left(1 - \frac{\|\mathcal{H}\|_F{}^2}{2c^2}\right) \,, \qquad \|\mathcal{H}\|_F = \sqrt{\textstyle\sum_{j \le 3} \lambda_j^2} \,. \tag{4.8}$$

The threshold values a, b, c are respectively set to 0.5, 0.5 and half the maximum value of the Hessian norm, in line with [147]. Since vessel sizes vary, we choose a range of length scales $\{\text{scales}\} = \{0.5 \text{ mm}, 1 \text{ mm}, 1.5 \text{ mm}, 2 \text{ mm}\}$ appropriate for the size of the vessels we seek.

Vesselness density

We use our vesselness filter to initialise our fissure searching algorithm, since fissures are located in regions with few vessels. Authors using this approach [130,131,132,145] typically threshold their vessel detection output and apply a distance transform to identify regions away from the vessels. One problem with this approach is that occasionally small vessels or vessel-like structures cross or appear in close proximity to the fissures. An example is shown in Figure 4.2 (p.116). This can significantly disrupt the distance transform, leading to the estimated fissure moving away from the true fissure.

To address this issue, we seek an alternative quantity to the distance transform which is robust to small vessels near the fissures. We introduce the vesselness density, which is a simple application of a wide Gaussian filter of standard deviation 10 mm to the multiscale vesselness output. This provides a relatively smooth density function which is high-valued in the centre of the lobes and low-valued between the lobes. Figure 4.5 shows an example output.

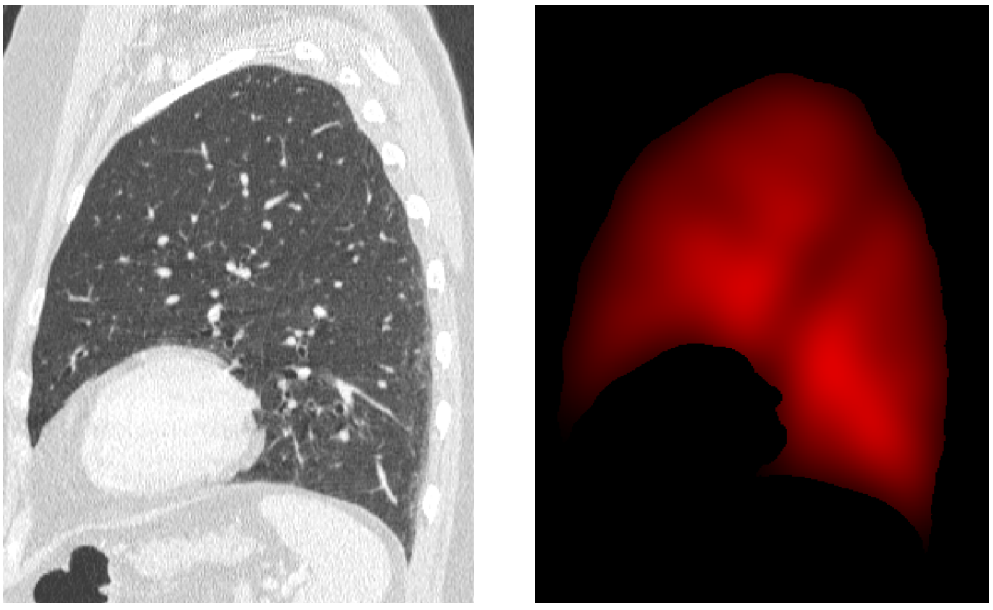

Figure 4.5: Example output of the vesselness density filter. Left: Sagittal slice reconstruction of a CT image of the left lung. Right: Corresponding output of a 3D vesselness density filter. The filter result is scaled between 0 (black) representing low vesselness density and 1 (red) representing high vesselness density. The left oblique fissure is clearly visible as a band of low vesselness density.

4.3.5 Fissureness

We develop a fissureness filter analogous to vesselness, which provides a strong response in the vicinity of pulmonary fissures. We apply a Gaussian filter to the original CT data with a kernel size of 1.5 mm, which is an appropriate scale to capture the features of the fissure plane. At each image point the Hessian matrix is constructed. Its eigenvalues λ_1 , λ_2 , λ_3 are computed and ordered such that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. At each point we define the fissureness

$$\mathcal{F} = \Gamma \, \mathcal{F}_{\text{plane}} \, \mathcal{F}_{\text{wall}} \, \mathcal{F}_{\text{vessels}} \, . \tag{4.9}$$

The first factor suppresses points whose largest eigenvalue λ_3 is positive, since fissures are locally bright:

$$\Gamma = \begin{cases} 1, & \lambda_3 < 0 \\ 0, & \lambda_3 \geq 0 \end{cases} . \tag{4.10}$$

The second factor detects plane or curve-like structures by searching for differences in the two largest eigenvalues:

$$\mathcal{F}_{\text{plane}} = \exp\left(-\frac{\mathcal{R}_{\text{plane}}{}^2}{2p^2}\right), \quad \mathcal{R}_{\text{plane}} = \frac{|\lambda_2|}{|\lambda_3|} , \tag{4.11}$$

with p (set to 0.5 in our case) acting as a soft threshold on $\mathcal{R}_{\text{plane}}$. The third factor suppresses signals from vessel walls, which, in contrast to plane-like fissures, have relatively large second, and possibly also third, eigenvalues. A soft threshold parameter *w* is used, in this case set to 3:

$$\mathcal{F}_{\text{wall}} = \exp\left(-\frac{\mathcal{R}_{\text{wall}}{}^2}{2w^2}\right), \quad \mathcal{R}_{\text{wall}} = \sqrt{\lambda_1^2 + \lambda_2^2} \, . \tag{4.12}$$

The final term decreases the fissureness value at points in the vicinity of blood vessels:

$$\mathcal{F}_{\text{vessels}} = \exp\left(1 - \frac{DT\{\text{vessels}\}^2}{2v^2}\right) . \tag{4.13}$$

Here $DT\{\text{vessels}\}$ is the distance transform to segmented blood vessels, found by thresholding the vesselness filter described above at a threshold value of 0.1. The idea is to suppress points in the vessel walls, which locally appear fissure-like and are close to high-vesselness voxels in the vessel interior. v is a characteristic distance over which fissureness should be suppressed close to vessels. We set v = 5 mm. A sample fissureness output is shown in Figure 4.6.

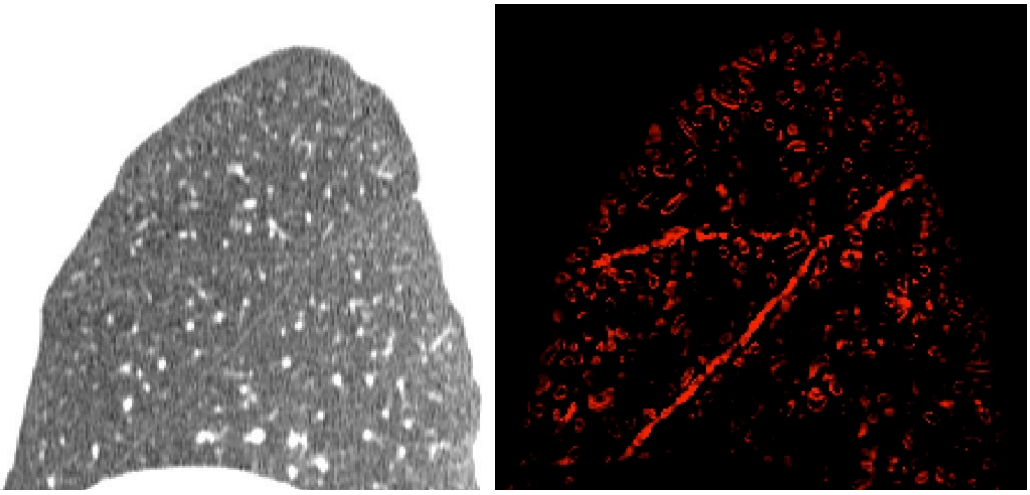

Figure 4.6: Example output of the fissureness filter. Left: A reconstructed sagittal CT slice from the right lung. Right: the corresponding output of the fissureness filter. The filter result is scaled between 0 (black) and 1 (red).

Properties of the fissureness filter

We have observed the following characteristics of the fissureness filter, which we have taken into account while designing our lobe segmentation method:

- Fissureness values are highly heterogeneous across the fissure. A global threshold is therefore not a sufficiently specific segmentation method as this would remove large sections of the fissures or produce too many false positives. Discontinuities in the fissure values also make curve tracing approaches such as Canny methods unreliable. For this reason our method is based on local maxima detection combined with a high-specificity threshold.

- Values of fissureness are normally at a local maximum on the fissure, *in the direction perpendicular to the fissure*. There are two consequences to this. Firstly, fissure detection by local maxima filtering is only effective if it is directional. Secondly, methods which can alter the local maxima of the fissureness may bias fissure detection away from the true points. These include algorithms such as watershed [82] and graph searching methods, which (directly or indirectly) apply a weighting to the fissureness values. To take these effects into consideration, our method introduces directionality through the use of a reference plane. Maxima are detected perpendicular to this plane using unweighted values.
- Outlying points of high fissureness may occur systematically. For example, high fissureness values can appear on minor fissures between segments. The systematic nature of such outliers is important as it means many standard curve-fitting methods such as least-squares and Ransac cannot be used; these statistical methods are robust to random outliers, but not systematic outliers. For this reason, we use morphological operations rather than statistical methods to remove outliers, and we perform the final curve fitting using multilevel B-splines, which is based on local curve-fitting.

4.3.6 Initial lobar approximations

To generate a first approximation for the lobar boundaries, we first compute the Hessian matrix-based vesselness filter above using four scales (0.5, 1.0, 1.5, and 2.0 mm). We then create the vesselness density field described above by applying a strong (10 mm) 3D Gaussian filter to the vesselness filter. A watershed transform [82] is applied to this vesselness density to separate the image into lobar approximations. The watershed algorithm only operates within the volume defined by the lung segmentation (see section 3.2). Each lobar region is seeded using points from the airway tree that have been assigned to

that lobe using the airway labelling described in section 3.4.6, excluding the major airway generations. The resulting watershed regions define the lobar approximations (see Figure 4.7). The boundaries between the resulting watershed regions provide our estimation for each fissure.

Our watershed approach differs from that used in existing work such as by Ukil and Reinhardt [130], where the watershed is applied to a distance transform from the vessels, and is thus sensitive to the correct labelling of vessels by lobe. Vessels close to the fissure boundary are easy to mislabel, and this leads to a requirement of manual correction in some cases. Furthermore, while lobes are largely independent, small vessels occasionally cross the lobar boundaries, which introduces significant errors if the distance transform is used. In contrast, the vesselness density is a much smoother function with minima close to the fissure boundaries, and does not require any vessel labelling.

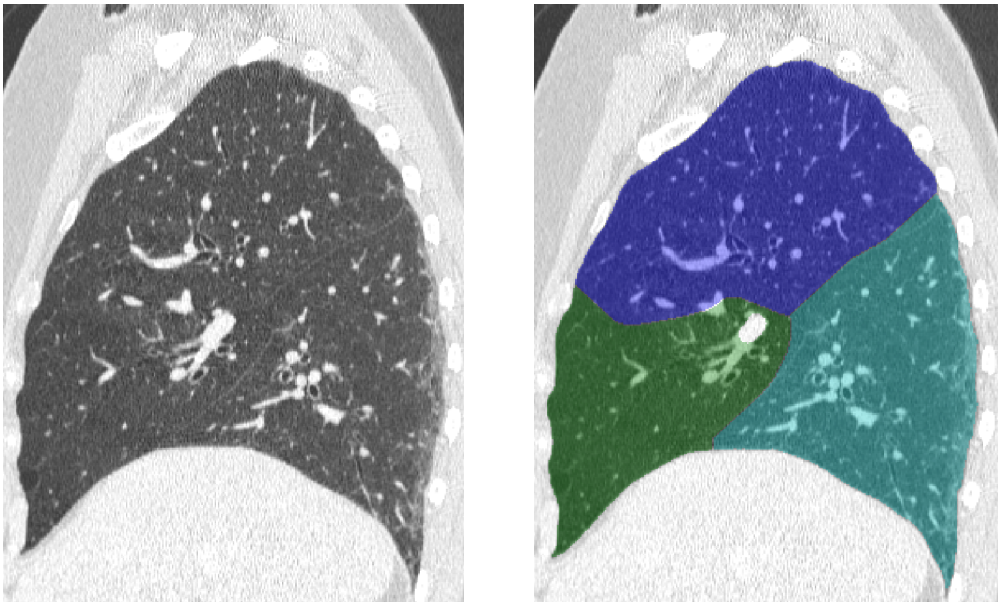

Figure 4.7: Results of the initial lobar approximation method. Left: Sagittal slice reconstruction of a CT image of the right lung. Right: Corresponding output of the lobar initial approximation based on a 3D watershed algorithm [82] on the vesselness density seeded from airway bronchi labelled by lobe. The transparent overlay is coloured according to the lobe mask. Dark blue shows the upper right lobe approximation; green shows the middle right lobe approximation; and cyan shows the lower right lobe approximation.

4.3.7 Selection of fissure points

To construct the final fissure surfaces, we combine information from the initial lobar approximation (section 4.3.6) with the output of the fissureness filter (section 4.3.5). We first select a set of candidate fissure points $\{p\}$ for each fissure, consisting of voxels in the lung volume that lie within 20 mm of the initial approximation for that fissure and which have a fissureness value greater than a minimum threshold (set to half the maximum fissureness value in the image). Next, we construct a *fissure reference plane*. We do this by using principal component analysis to fit a plane through the points on the initial fissure approximation. We define coordinate axes x' and y' in the reference plane and z' perpendicular to the plane (see Figure 4.8).

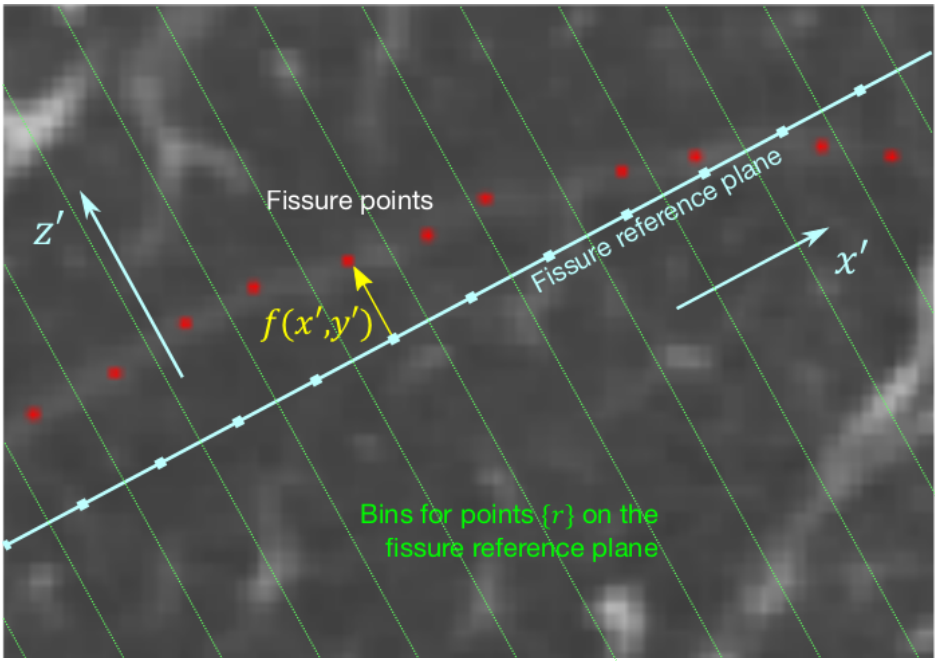

Figure 4.8: Schematic illustration of the process for selecting points on the fissure surface. The algorithm is three-dimensional but has been illustrated in 2D for clarity. A fissure reference plane (blue) is constructed based on the initial lobar boundary approximation. The reference plane is divided into a 2D grid of points (blue) with coordinates x' and y' in the plane. In this figure, y' is perpendicular to the image. The candidate voxels in the 3D image are grouped into bins, according to the nearest point on the plane. In this diagram, the green lines show the boundaries of the bins. Candidate voxels whose centres lie between the green lines will fall into the bin for that grid point. Fissure points (red) are found by examining the candidate voxels in each bin and choosing the voxel with the highest fissureness value.

We can represent the true fissure surface relative to the plane by the function

$$z' = f(x', y') \,. \tag{4.14}$$

Voxels on the fissure surface in general exhibit a local maximum of fissureness in the z' direction. Our aim is to find a subset of the candidate points $\{q\} \subseteq \{p\}$ with this property.

To do this we create a 2D grid of points $\{r\}$ on the reference plane using 1 mm intervals in x' and y'. We assign each candidate fissure point p to the nearest grid point r, which groups the candidate fissure points into 3D spatial bins (see Figure 4.8). We select one point from each bin, which is the voxel with the highest fissureness value. This forms the new subset of candidate fissure points $\{q\} \subseteq \{p\}$.

This method is designed to select a large number of points across the fissure, each of which is at a local fissureness maximum. The binning approach ensures the whole of the fissure surface is covered, without biasing the number of points selected towards areas of high fissureness. Points outside the lung are naturally excluded, since they do not appear in the set of initial candidate points $\{p\}$. Candidate points never change their location, so they always map directly to the original image. This allows us to use fissureness values directly, avoiding resampling and interpolation errors.

The resulting candidate fissure points $\{q\}$ lie mostly on the fissure, but may include some outlying points of high fissureness. Outliers are removed by examining the connectedness of candidate points after a morphological dilation using a spherical structural element of radius 2.5 mm. For any connected components comprising a volume less than 300 mm^3, the associated candidate points are removed.

4.3.8 Construction of fissure surface

The final stage in our lobe segmentation method is to construct a smooth, continuous fissure surface using the set of candidate fissure points $\{q\}$ for each fissure. We use these fissure approximations to separate the lung segmentations into lobes. We first create a new reference plane based on the final set of candidate fissure points $\{q\}$. As before, this is done by fitting a plane through the principal axes. Each point in $\{q\}$ has coordinates (x, y, z) relative to the new plane. The final smooth fissure surface $z = f(x, y)$ is computed using a multilevel B-spline method [151]. This method offers a number of advantages. It is a fast, multiscale method for computing a surface approximation from a set of scattered data points. It produces a smooth curve based on local point fitting, but does not force the curve through every point, making it more robust to overfitting in the presence of noise. The multiscale nature of the method ensures that each part of the curve is primarily controlled by the nearest data points, even though the proximity of those data points varies throughout the image.

Multilevel B-spline method

Here we present a sketch of the multilevel B-spline method; a full explanation is given by Li et al. [151]. The method computes an initial approximation surface f_0 to the data points $\{q\}$ on a coarse grid Φ_0. Then, a new set of points $\{q_1\}$ is created, containing the coordinates of the original points relative to f_0. A new approximation surface f_1 is computed for the points $\{q_1\}$. This is computed on a finer grid Φ_1 with twice the spacing of the previous grid. A new fissure surface is computed by summing the approximation surfaces $f_0 + f_1$, and this is used to compute a new set of points $\{q_2\}$. This process is repeated at k scales, to obtain the final fissure surface f as a sum of the approximation surfaces

$$f = f_0 + f_1 + \cdots + f_k \,. \tag{4.15}$$

Each surface f_k is found by computing a cubic B-spline surface approximation to the point set $\{q_k\}$. Without loss of generality, we can take the grid points of the lattice Φ_k (the lattice at scale k) to lie on the integer boundaries of x and y. Then f_k is given by

$$f_k(x,y) = \sum_{l=0}^{3}\sum_{m=0}^{3} B_l(s)B_m(t)\phi_{(i+l)(j+m)'} \,, \tag{4.16}$$

where ϕ_{ij} are weightings at each grid point of Φ_k, whose coordinates are given by $i = \lfloor x \rfloor - 1$ and $j = \lfloor y \rfloor - 1$. The uniform cubic B-spline basis functions are

$$\begin{aligned} B_0(t) &= \frac{1}{6}(1-t)^3 \,, \\ B_1(t) &= \frac{1}{6}(3t^3 - 6t + 4) \,, \\ B_2(t) &= \frac{1}{6}(-3t^3 + 3t^2 + 3t + 1) \,, \\ B_3(t) &= \frac{1}{6}t^3 \,, \end{aligned} \tag{4.17}$$

where $s = x - \lfloor x \rfloor$ and $t = y - \lfloor y \rfloor$.

Each grid point weighting ϕ_{ij} is found using a least squares minimisation based on a subset of points $\{c\} = \{q_k\}_{ij}$, consisting of all the points in $\{q_k\}$ which lie in a 4 × 4 grid point-size neighbourhood of the grid point (i,j)

$$\phi_{ij} = \frac{\sum_c w_c^2 \phi_c}{\sum_c w_c^2}. \tag{4.18}$$

For each point c, we denote its coordinate as (x_c, y_c, z_c), and its weighting is given by

$$\phi_c = \frac{w_c z_c}{\sum_{a=0}^{3}\sum_{b=0}^{3} w_{ab}^2} \,, \tag{4.19}$$

where w_c is defined by the basis functions

$$w_c = w_{lm} = B_l(s)B_m(t) \tag{4.20}$$

and $l = (i+1) - \lfloor x_c \rfloor$, $m = (j+1) - \lfloor y_c \rfloor$, $s = x_c - \lfloor x_c \rfloor$ and $t = y_c - \lfloor y_c \rfloor$.

In this way, the multilevel B-splines method minimises the error between the currently computed surface and a set of control points, which changes at each scale.

4.3.9 Lobe segmentation

After computing each fissure surface using the multilevel B-spline method, the surface is extrapolated to the lung boundaries using thin plate splines. The final lobe segmentation is computed by performing a dual-region flood fill operation, seeded from both sides of the fissure surface. This ensures the lobes are fully separated even if there are holes in the fissure surface resulting from the discretisation. In the right lung, the lower lobe is first separated from the upper and middle lobes by finding the oblique fissure. Once these are separated, the horizontal fissure is found and the upper and middle lobes are separated.

4.4 Results

Figure 4.9 shows one of the lobe segmentation results, which is illustrated as a coloured lobe mask overlaid on the corresponding sagittal CT image slice. A 3D visualisation of the resulting lobe masks is shown in Figure 4.10.

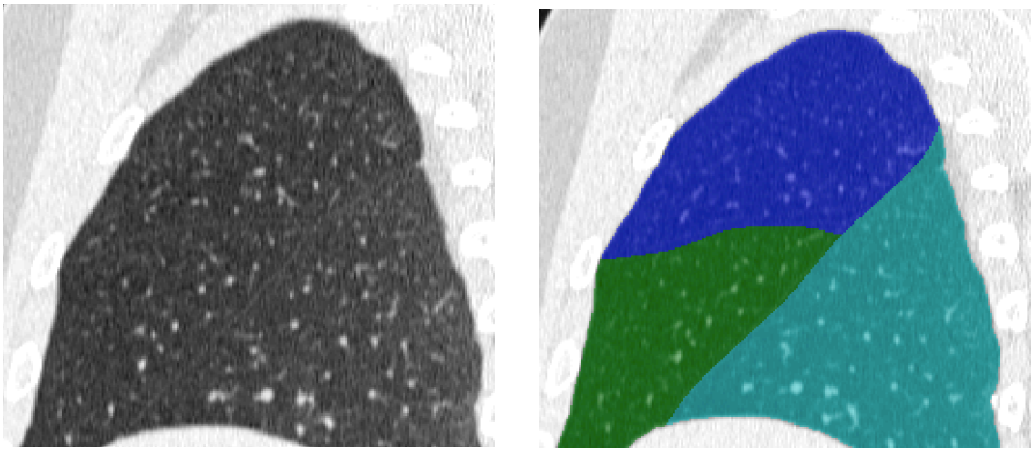

Figure 4.9: Example results from the lobe segmentation. Left: A reconstructed sagittal slice of the right lung from CT data. Right: The same slice with transparent lobe segmentation masks overlaid, showing the upper (dark blue), middle (green) and lower (cyan) lobes.

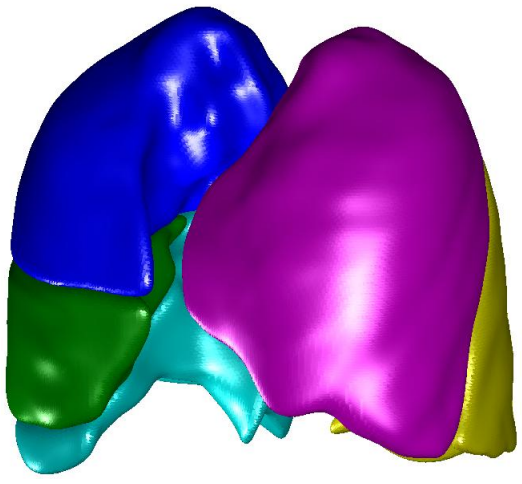

Figure 4.10: Smoothed 3D surface rendering of the segmented lobe masks.

4.4.1 Validation

Our lobe segmentation method was validated with the assistance of an experienced pulmonary clinician. We performed segmentation on ten CT lung cancer datasets described in section 1.5. For each dataset, the clinician created validation points by manually placing fissure markers on every 10th axial slice of the CT data. This generated approximately 300 fissure points per dataset. Markers were placed across the whole span of the fissure present in each slice, including points close to the lung boundaries.

Markers were placed using the graphical interface of our Pulmonary Toolkit software (Figure 4.11). This presents an interface similar to clinical visualisation software, with the ability to zoom, pan, cine through adjacent slices and change window and level settings. The image can be switched to coronal, sagittal or axial view, and markers are visible and can be placed in any orientation. Markers are added by clicking on the image, moved by dragging, and can be deleted or assigned to different fissures using a right-click menu.

Table 4.2 shows the results of the validation for each of the ten datasets. For comparison with other works (e.g. Ukil and Reinhardt [130]), we quote the mean, maximum and RMS distance in mm of the manually placed marker points to the nearest point on the automatic fissure segmentation (see discussion below).

Table 4.2: Validation results for the automated fissure detection algorithm. Results are grouped by fissure. We show mean, max and root mean square of the distance in mm from a set of manually placed fissure points to the nearest point on the automatically segmented fissure curve.

	Right oblique fissure			Right transverse fissure			Left oblique fissure		
	Mean	**Max**	**RMS**	**Mean**	**Max**	**RMS**	**Mean**	**Max**	**RMS**
Subject 1	1.53	8.94	2.54	0.62	1.37	0.77	2.72	31.61	5.98
Subject 2	2.26	20.80	4.22	1.78	8.99	2.78	1.26	14.57	2.68
Subject 3	4.31	17.94	5.57	6.67	13.21	7.78	7.59	46.62	11.90
Subject 4	2.43	15.64	3.53	1.13	3.64	1.46	1.55	14.58	2.65
Subject 5	2.79	27.37	6.00	2.54	10.25	3.73	1.61	8.99	2.34
Subject 6	2.61	11.64	3.51	4.79	26.19	8.53	2.76	15.63	4.18
Subject 7	3.68	15.01	4.81	3.66	16.31	5.27	4.09	19.10	5.90
Subject 8	1.52	10.22	2.22	4.57	11.22	5.72	2.42	19.71	4.38
Subject 9	3.26	21.03	4.76	7.27	25.03	10.79	1.40	15.38	2.55
Subject 10	3.85	29.08	6.70	4.90	21.94	8.09	2.98	27.90	5.72
Mean	2.82	17.77	4.38	3.79	13.82	5.49	2.84	21.41	4.83

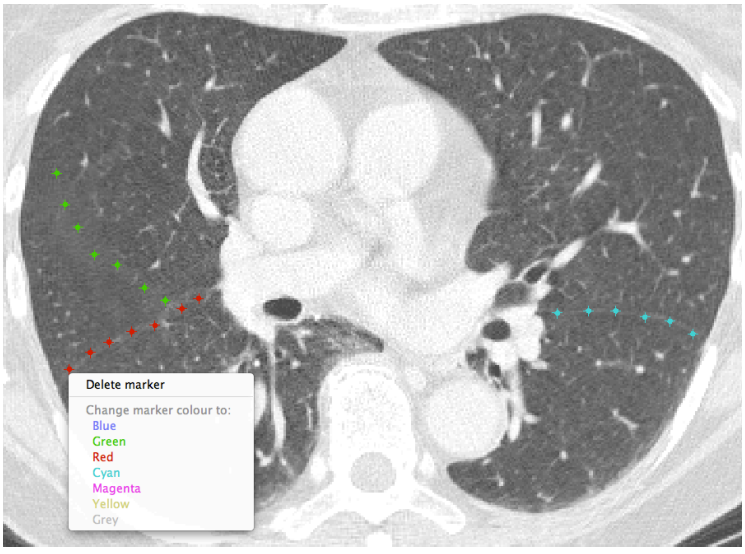

Figure 4.11: Screenshot of the Pulmonary Toolkit graphical interface showing the interactive placing and editing of marker points. Markers are shown for the right oblique (red), left oblique (cyan) and right horizontal (green) fissures. The markers shown here were placed by a pulmonary clinician for the validation of the lobe segmentation method.

4.5 Discussion

The oblique fissures show typical mean distances below 3 mm, comparable to other published results [127,128,129,130,133]. We expect values for the right transverse fissure to be higher, since this fissure is much harder to detect and many existing methods use manual correction [128,130,133]. Our results are broadly within the range of intra-observer variability reported in [127,129,132]. However, some of our results show relatively high maximum and RMS values, resulting from inaccurate curve extrapolation to the lung boundaries. This is one area we seek to improve in future work. Another is our method for removing outlying points, which sometimes removes true positives in regions with fissures that are difficult to detect, such as the horizontal fissure. In general the method does not depend on the quality of the vessel and airway segmentations, since these are only used as an initialisation step for the fissure detection. However, as with most other methods, it does require the airway segmentation algorithm to correctly find and label the main airway serving each lobe. In future work, we aim to test the effectiveness of this on a larger database.

4.5.1 Performance

Table 4.3 shows the execution time for the each of the five stages of the lobe segmentation method. The stages are: computing the vesselness filter; generating an initial lobar approximation by calculating the vesselness density filter and performing a watershed transform; computing the fissureness filter; selection of fissure points using an approximation plane; and the final lobe segmentation generated by fitting smooth multilevel B-spline surfaces to the selected points and using these to divide the lung regions into lobes. Times for the lung and airway segmentation algorithms are not included; these are shown in Table 3.1 (p.56) and Table 3.9 (p.100) respectively. Times are recorded for ten datasets on a laptop computer (2.66 GHz MacBook Pro running OS 10.8 and Matlab 8). For comparison the segmented lung volumes are also shown.

Table 4.3: Execution time for the lobe segmentation algorithms.

Subject	Volume (cm^3)	Algorithm execution time (s)				
		Vesselness	**Initial lobar approximation**	**Fissureness**	**Fissure point selection**	**Final lobar segmentation**
1	3167	178.6	57.1	38.9	20.5	53.2
2	4890	236.1	88.8	51.1	27.4	81.6
3	6208	284.1	106.9	69.9	36.0	87.5
4	4343	232.1	95.4	52.2	28.1	67.1
5	3293	206.5	70.7	49.1	26.2	68.7
6	5436	229.8	80.6	56.0	28.5	69.8
7	4894	260.5	91.6	59.7	30.8	92.6
8	7247	303.3	126.1	71.1	37.2	118.8
9	3374	185.9	61.1	41.4	23.1	63.9
10	3467	201.1	67.2	46.4	24.5	65.0
Mean		**231.8**	**84.6**	**53.6**	**28.2**	**76.8**

4.6 Conclusions

In this chapter we developed a novel, fully automated lobe segmentation method which addresses some of the limitations of existing work. We have found that published methods broadly follow a standard approach, which involves lung segmentation using standard methods; filter-based detection of candidate fissure points; use of prior information to reduce false positives from the candidate points; and using the final set of fissure points to generate a smooth, continuous fissure surface, with which the lobes can be segmented. We found a number of limitations in existing works, including slow execution times, manual correction requirements and lack of robustness to incomplete fissures or unusual anatomical features.

Our lobe segmentation algorithm is a fully automated, computationally efficient method based on a fissureness filter combining image appearance and anatomical information from the pulmonary vasculature and airways. This anatomical information allows the algorithm to produce a continuous plane even in the absence of clearly defined features in the image. We achieve this using an adaptive minimisation plane and a multilevel B-spline fitting. Furthermore, we have found the combination of several information sources adds robustness to the detection, e.g. in the case of small vessels crossing the lobar boundaries which are compensated for by the presence of other anatomical cues. Our results are broadly within the error range of existing published works, but we also aim to make further improvements and increase our test database.

5 Regional measurements from xenon MRI

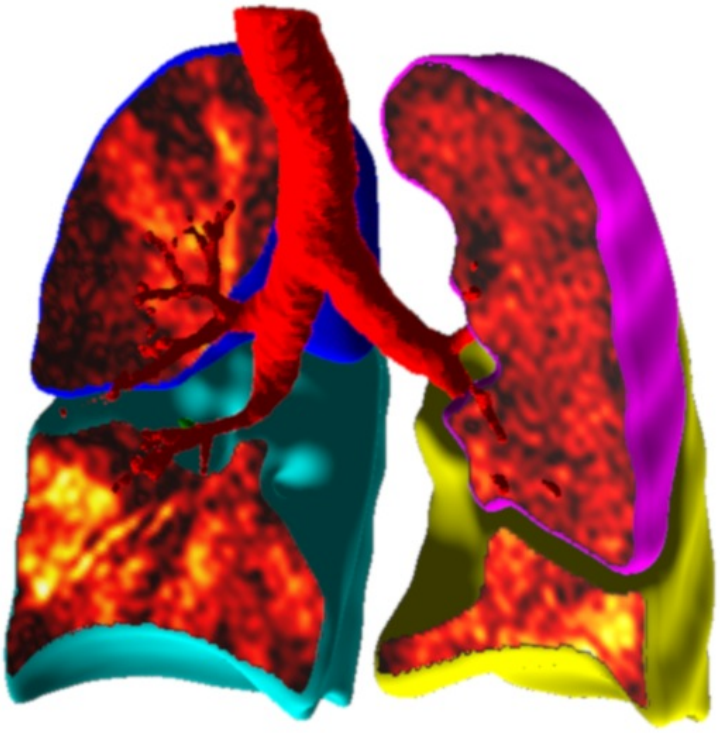

Chapter 5: Regional measurements from xenon MRI



In this chapter we develop an automated method for obtaining lobar measurements from hyperpolarised xenon MRI. Our method forms part of our investigation into correlating regional xenon image analysis with disease severity using data from the clinical trial (section 1.4). We demonstrate the complete pipeline using the first data from the clinical trial, comprising four control subjects. For comparison we perform analysis for a COPD subject using data from a previous study [13].

In section 5.1 we introduce regional analysis and in section 5.2 we describe previously published work. In section 5.3 we describe our methods. Lobar analysis is achieved by using image registration to apply the CT lobar masks from chapter 4 to the xenon MRI images. The registration method is described in section 5.3.1, and uses the automatically segmented lung boundaries for CT and proton MRI data generated using the methods developed in chapter 3.

We illustrate the pipeline with several examples of regional measurements. Ventilated volume analysis is described in section 5.3.2, apparent diffusion coefficient measurements in section 5.3.3 and CT-based emphysema evaluation in section 5.3.4. Results are presented in section 5.4 and discussed in section 5.5. Finally, we present conclusions for this chapter in section 5.6.

5.1 Introduction

This research aims towards the development of new methods for the clinical assessment of lung disease. These methods combine two recent advances. The first is hyperpolarised xenon MRI, which produces three-dimensional images of ventilation and diffusion in the lungs (see section 2.3.3), providing a wealth of functional information not available from standard anatomical imaging or from global function tests. The second advance is the use of automated segmentation techniques, enabling quantitative regional image analysis. Many lung diseases act at a lobar level, including emphysema [121], postprimary tuberculosis [122], silicosis [123] and pulmonary fibrosis [124] (see chapter 4). Lobar-based measurements could help with disease localisation and may be more sensitive to the detection and phenotyping of disease than current methods (see section 2.3).

A particular challenge arises when combining these two advances. Lobar analysis of xenon MRI images is difficult because the lobar boundaries are not directly visible on xenon MRI (see Figure 4.1, p.105). However, lobar boundaries are visible on CT data, which are available for all subjects in the clinical trial. We segment the lobes using the methods we developed in chapter 4, and then map them to the xenon images using image registration. Previous work has made use of rigid registration techniques [9,90]; however, these are inaccurate for this purpose, because CT and xenon MRI are typically performed at different levels of lung inflation (total lung capacity and residual volume plus one litre, respectively). The positions of the lobes can move by several centimetres between these two inflation levels. We aim to account for this motion by using a non-rigid registration method.

Direct registration between CT and xenon MRI is difficult due to the lack of anatomical features present in the xenon MRI data. For this reason we perform

the registration in two steps, using proton MRI as an intermediate stage. We use a simple registration method based on a distance transform to the lung contours. This method does not make use of all the information available in the images, but it has the advantage of being robust to thick-slice MRI data and sliding issues at the lung boundaries. The combined registration computes a deformation field which maps CT images to the xenon images. In this chapter we use this deformation field to derive and compare lobar measurements from CT, ventilation and diffusion images.

We illustrate our pipeline with image-based measurements which have shown correlations with COPD in previous work, but for which there has been little or no investigation of their properties at a lobar level. The first two measurements are derived from functional xenon MRI data. The *ventilated volume* measures the volume of lung receiving gas during the breathhold ventilation scan [88]. Since we have automatically segmented lung masks from proton MRI and CT data, we can compute this as a percentage of total volume, which is a more meaningful inter-patient measure. Secondly, the *apparent diffusion coefficient* (ADC) is derived from xenon MRI diffusion data and is sensitive to disease-related microstructural changes in the alveoli which increase the apparent diffusion of the xenon gas [164]. Finally, *emphysema percentage* (EP) and the related *percentile density* (PD) are derived from CT data and detect emphysema from regional decreases in radiodensity [26]. We compute these measurements at a whole lung, individual lung and lobar levels.

The same deformation field can be used to register other CT-derived data to the xenon images. For example, a CT-derived segmental map could be used to further refine the regional analysis, while segmented airways can be used to build computational airflow models. Our future work will investigate these possibilities, and provide further assessment, validation and correlation of regional measurements with disease (see chapter 6).

5.2 Previous work

5.2.1 Regional analysis of hyperpolarised MRI

Hyperpolarised xenon MRI is an emerging technology and only a few groups have presented results for human imaging [13,59,60,61]. Quantitative regional analysis has been very limited: generally, measurements are computed across the whole of each coronal image slice. While this can show gravitational dependence, it does not reflect the regional nature of many lung diseases which may act at a lung, lobar or segmental level. Furthermore, without bringing in additional information about lung structure, it is difficult to accurately detect regions of the lung which are not being ventilated. This is due to unventilated regions being indistinguishable in xenon MRI from areas exterior to the lung.

Regional analysis methods are similarly limited in hyperpolarised helium MRI, despite this being a longer established modality [47]. Investigating cystic fibrosis, Woodhouse et al. [165] divide each coronal helium-3 ventilation slice into six regional zones based on identifiable landmark points such as bifurcations in the airway tree. Tzeng et al. [87] compare ventilated volumes in asthmatic subjects before and after a bronchial challenge test by manually segmenting the ventilated volumes and rigidly registering corresponding slices before and after the test. Other methods use registration with proton MRI or CT images to provide anatomical information to assist segmentation and analysis. Woodhouse et al. [88] perform manual segmentation of the lung boundaries in proton MRI images and ventilated boundaries in helium-3 images, and use these to quantify ventilation defects per coronal slice. Fain et al. [54] quantify unventilated lobar regions in asthma patients by performing manual lobe segmentation of the helium-3 images guided by overlaid proton MRI images.

Recent work by Wild et al. [166] shows concurrent acquisition of helium and proton MRI, allowing improved registration between the two modalities. Ireland et al. [90] demonstrate a method for registering CT to helium-3 images using a rigid registration based on manually selected landmarks.

Plotkowiak [9] demonstrates a semi-automated technique for performing lobar-based analysis of helium-3 diffusion images. The lobar boundaries are semi-automatically detected in CT images and registered to the helium-3 data using a rigid registration method, using proton MRI data as an intermediate step. The registration is based on a distance transform image computed from the lung boundaries, which are manually segmented.

Our work in this chapter extends the methods of Plotkowiak [9], while introducing a substantial number of improvements and additions. The previous work computed ADC measurements for helium-3 images; we have extended this ADC computation to new xenon-129 images and added regional analysis of ventilation defects. To perform the regional analysis we have developed a novel, fully automated CT lobar segmentation method (chapter 4), replacing the semi-automated method by Plotkowiak [9]. We replace the rigid registration with a non-rigid method, and replace manual segmentation of the MRI lung boundaries with a method which is fully automated (section 3.3) while still permitting manual correcton if necessary. Finally, while the modelling work by Plotkowiak [9] focussed on diffusion processes, we have developed airway segmentation and measurement methods (section 3.4) which allow us to build personalised models of airflow as part of our planned future work (chapter 6).

5.2.2 Registration of lung images

Lung image registration is challenging due to the multiple imaging modalities used and the large, non-uniform deformations of the lung during breathing. While registration is a major area in medical image analysis [167], work on CT-MRI registration of lung images is very limited. Previous analysis of hyperpolarised gas MRI relies on rigid techniques based on distance transform maps to the lung boundaries [9] or manually selected landmarks [90]. The changes in lung volume can be substantially reduced if all imaging is performed at similar levels of inflation. However, this is not generally possible in a clinical setting, where inflation levels are chosen to achieve the optimal imaging for each modality. CT is typically performed at total lung capacity (TLC), while hyperpolarised gas MRI is performed at residual volume plus one litre (RV + 1 l) or functional residual capacity plus one litre (FRC + 1 l). A non-rigid registration is necessary to correctly account for these variations.

A recent Miccai workshop challenge was held into the registration of lung images [168]. However, this only covered registration within a single modality, CT. Many intensity-based methods are not applicable to CT-MRI registration due to the different imaging properties of the two modalities [97]. Approaches based on mutual information [169] have been applied to multi-modal image analysis, including CT-MRI lung registration [170]. Heinrich et al. [97] demonstrate a modality-independent method based on self-similarity of small image patches. These methods can provide accurate CT-proton MRI registration, but generally require high-resolution images for both modalities.

An additional challenge is the sliding at the lung boundaries. The pleural fluid permits the lungs to slide within the thoracic cavity during breathing, so that adjacent structures such as the ribs may have a very different motion pattern from the lung surface itself. Many registration methods include isotropic spatial regularisation and may fail in this situation. In the recent challenge workshop

[168], the most successful methods were those that mitigated the sliding effects by registering the lungs after extraction from the image background [171]. Risser et al. [171] develop a piecewise-diffeomorphic registration technique which allows for sliding between regions as well as the movement of the boundary between them. They demonstrate their method for sliding at the lung boundaries for CT-CT [171] and CT-proton MRI registration [170].

Sliding can also occur between lobes along the separating fissures, although the effect of this is likely to be smaller than for the lung boundaries. The piecewise-diffeomorphic method of Risser et al. [171] could in principle be applied to registration with lobar sliding. However, the method in its current form could not be used for CT-proton MRI registration; accurate lobe segmentations would be required in both modalities, and lobe-separating fissures are not generally visible in proton MRI. Piecewise-affine methods [172,173] divide the image into regions which are each registered using their own local affine transformation. They could potentially model sliding at the lung and lobar boundaries while capturing the different rigidity properties of lung features such as the large airways. However, they have not been applied to CT-proton MRI lung registration, and may be very sensitive to the MRI quality.

Our method adapts the lung boundary approach of [9] to use a non-rigid registration. By removing the background we mitigate the effects of sliding at the lung boundaries, although as with existing methods, we do not model sliding between the lobes. This method is robust to the thick slice MRI data we are using, since the lung boundaries are more visible in thick slice datasets than other features which may be used by a mutual information approach. Two disadvantages of this method are the reduced accuracy caused by ignoring the image information interior to the lung, and the potential sensitivity of the distance transform to image noise. For these reasons, alternative registration methods remain an important consideration in our future work (chapter 6).

5.3 Methods

To perform regional analysis at the lung and lobar levels, we use the lung and lobe maps segmented from CT data using the methods from chapters 3 and 4. We apply these maps to the xenon MRI data using the non-rigid registration method described in section 5.3.1. We use the registered masks to derive three regional measures of COPD at the lung and lobar level: ventilated volume percentage (section 5.3.2); apparent diffusion coefficient (section 5.3.3); and emphysema percentage (section 5.3.4).

5.3.1 Registration

CT segmentation masks cannot be applied directly to the xenon MRI data due to variations in the subject position and inflation level between the scans, arising from the different scanners and protocols used for these imaging modalities. We correct for these differences using a non-rigid registration method, which generates a deformation field relating the images.

Direct registration between the CT and xenon MRI data is difficult due to the lack of anatomical features present in the xenon MRI data. For this reason we perform the registration in two steps, using proton MRI as an intermediate stage. Xenon and proton MRI data are registered using a rigid transformation based on the image metadata. This is reasonable because these images are acquired in the same imaging session by the same scanner, and the lung inflation levels are matched (see section 1.4). To relate the CT and proton MRI data, we use a non-rigid fluid registration [174]. Modality-independent registration images are built from the lung contours segmented using the methods of chapter 3. This method is robust to the thick-slice proton MRI data we are using.

Ideally we aim to achieve accuracy to within the resolution of the xenon images (typical resolution is 15.0 × 4.2 × 4.2 mm). This resolution is much lower than for CT (typical resolution 0.7 × 0.7 × 0.7 mm) and hence we do not require the same level of accuracy as state-of-the-art CT-MRI registration methods.

The sliding of the pleura presents a difficulty for most registration algorithms because they include regularisation of the deformation field. The regularisation causes problems across the boundary where tissues on either side have different movements. To avoid this problem, we remove all influence from tissue outside of the lung by basing our registration on the lung volume masks segmented from the CT and MRI data. This approach is similar to that used by Plotkowiak [9]. Reducing the images to lung masks removes useful information but has the advantage of being robust; segmentation errors in the lung masks can easily be detected and corrected by an observer if necessary.

The lung masks are converted into distance transforms from the lung exterior, which are each normalised by their maximum value. This defines modality-independent images which are used to compute the deformation field. The registration aims to minimise a sum of squared difference similarity measure between the normalised distance transforms. The non-rigid fluid registration is performed with a curvature-based regulariser (see below).

Once the two registrations have been performed, a combined deformation field is computed which can be used to directly register images derived from the CT data (such as lobe masks) with the xenon data, without any intermediate registration step.

Our method is summarised in Figure 5.1 and the individual stages are described in more detail below.

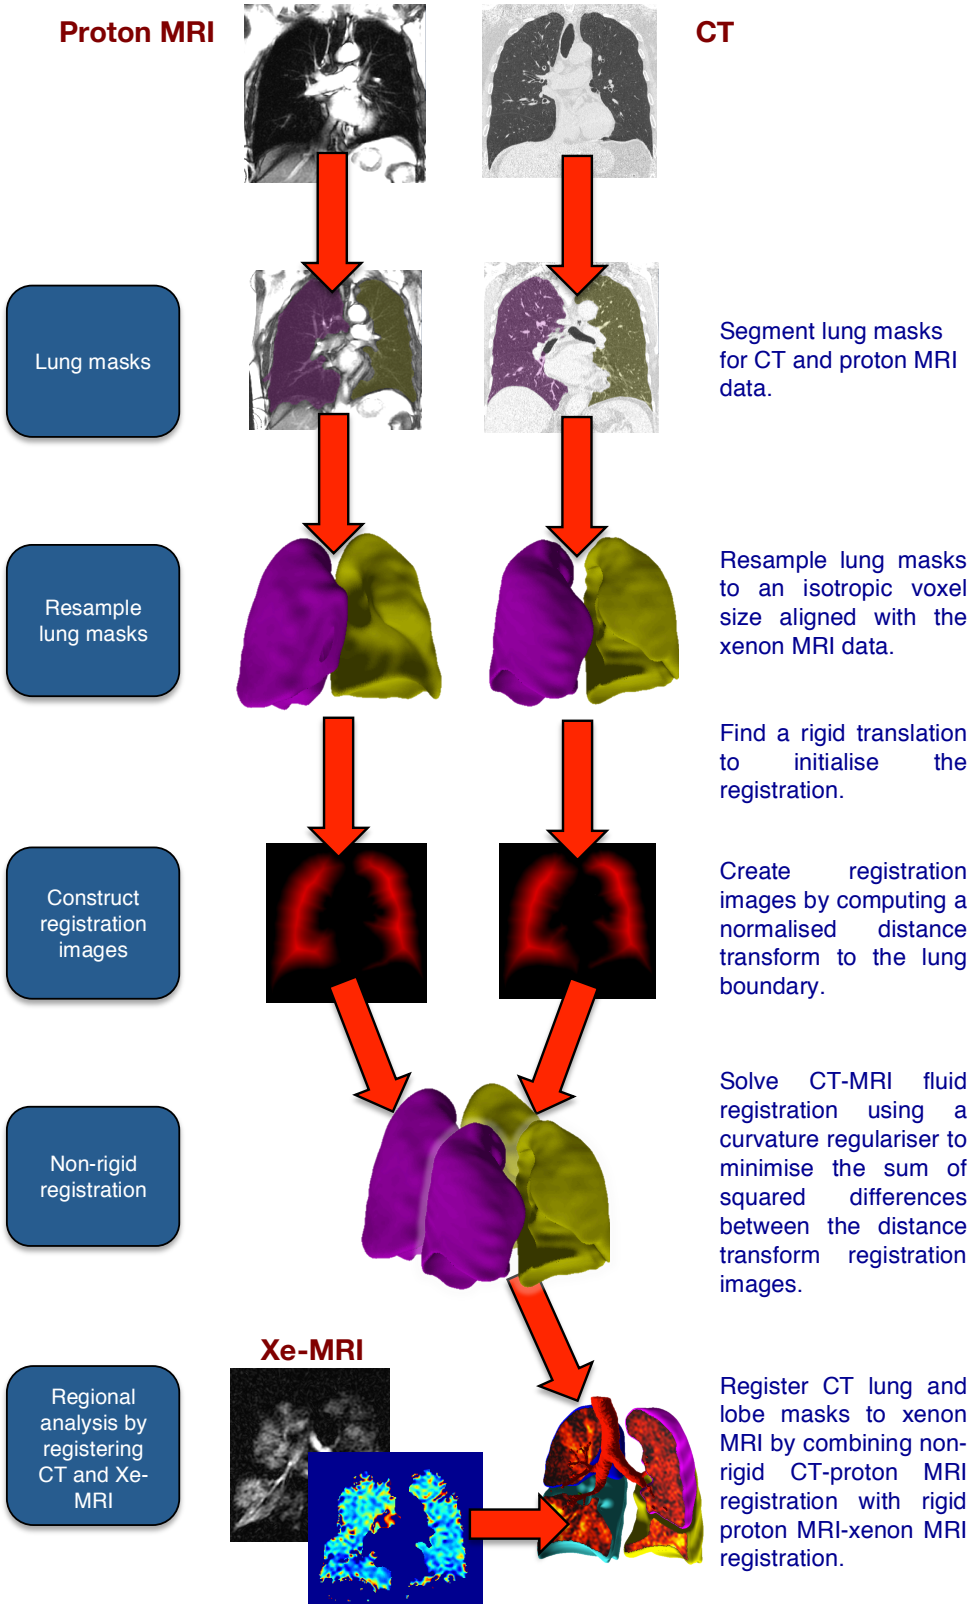

Figure 5.1: Schematic of the pipeline for regional image analysis using multi-modal image registration. Regional analysis is performed by registering the lobe boundaries from CT to proton MRI by combining a non-rigid method from CT to proton MRI with a rigid translation from proton MRI to xenon MRI.

CT to proton MRI registration

There are three stages in the non-rigid registration between CT and proton MRI. Firstly, an initial translation provides an approximate alignment of the images. Secondly, modality-independent images are created by computing distance transforms to the segmented lung masks. The lung masks are used since they are some of the few features visible in both modalities. Finally, a fluid registration is performed separately for the left and right lung images, which computes a deformation field to take into account the different lung inflation levels. These stages are described in more detail in the following sections.

1. Initial CT-proton MRI translation

We initialise the CT-proton MRI registration using a translation which approximately aligns the CT and proton images based on the centroids of the segmentation masks. We first align the left and right lung masks separately, computing a translation for each. We then compute the mean of the two translations and use this as the initial registration for both lungs. In this way we initially align both lungs to the same fixed point, while giving each equal prominence in the determination of this point. Using thin-slice proton MRI data, the translation could be based on an automatically detected landmark such as the carina, which is not well defined on thick-slice data.

2. Creation of registration images

To perform the CT-MRI registration we need to create modality-independent images. We first resample the CT and MRI lung masks to a new coordinate system based on the xenon MRI data, but we increase the number of coronal slices in order to create an approximately isotropic voxel size. For the COPD dataset, this results in a 10-times increase in the number of coronal slices giving a voxel size of $1.50 \times 1.56 \times 1.56\ \text{mm}^3$ (taking into account the upsampling of the MRI images by the scanner). We need to do this additional

resampling because the registration implementation we are using requires the two input images to be isotropic with the same voxel size. By choosing this almost isotropic voxel size, our interpolation grid is aligned exactly with the xenon MRI image coordinates, which will allow us to use the resulting deformation field directly in our alignments of CT to xenon MRI data.

The modality-independent registration image is created from the CT and MRI data by performing a distance transform to the exterior of the lung. This distance transform is normalised by its maximum value so that the range of values is the same for different lung volumes. However, this distance transform approach treats the lung as a homogeneous entity and does not take into account the regional nature of deformation.

3. Non-rigid CT-proton MRI registration

The initialisation above adjusts for variation in subject position between the MRI and CT scans. To compensate for changes in lung inflation level, we use a fluid registration to deform the lung masks. This is based on minimising the mean squared difference between the registration images (the normalised 3D distance transforms from the lung exterior).

We use the viscous fluid registration code by Cahill et al. [174] which generalises previous work [175] to use a curvature-based regulariser. Full details are given by Cahill et al. [174]; here we present a summary.

This registration seeks to deform a floating image I_{float} onto a reference image I_{ref} using equations of viscous fluid flow. In our case these images are the normalised distance transforms to the lung boundaries described above. The reference image is derived from the proton MRI lung mask, while the floating image is derived from the CT lung mask.

The deformed floating image can be expressed as

$$I_{\mathrm{def}}(\boldsymbol{x}) = I_{\mathrm{float}}\big(\boldsymbol{x} - \boldsymbol{u}(\boldsymbol{x})\big)\,, \tag{5.1}$$

where $\boldsymbol{u}(\boldsymbol{x})$ is a displacement field describing the movement of a point at $\boldsymbol{x}$ from its original position (an *Eulerian reference frame*). If the displacement $\boldsymbol{u}$ evolves in time, then we can define the velocity of the displacement field (*Eulerian velocity field*) as

$$\boldsymbol{v}(\boldsymbol{x}, t) = \frac{d\boldsymbol{u}}{dt} \tag{5.2}$$

$$= \frac{\partial \boldsymbol{u}}{\partial t} + \nabla \boldsymbol{u}(\boldsymbol{x}, t)\boldsymbol{v}(\boldsymbol{x}, t)\,. \tag{5.3}$$

In fluid registration, the displacement field is the steady-state solution of the Navier-Lamé partial differential equations for viscous fluid flow

$$\mu \Delta \boldsymbol{v}(\boldsymbol{x}, t) + (\mu + \lambda)\nabla(\nabla \cdot \boldsymbol{v}(\boldsymbol{x}, t)) = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}(\boldsymbol{x}, t))\,, \tag{5.4}$$

where t is time, $\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u})$ is a driving force, and μ and λ are the Lamé constants, typically set to 1 and 0 respectively [174]. We impose Neumann boundary conditions on the PDE. The driving force is equivalent to minimising a cost function equal to the sum of squared differences between the reference image and the deformed image

$$\boldsymbol{f}\big(\boldsymbol{x}, \boldsymbol{u}(\boldsymbol{x}, t)\big) = -[I_{\mathrm{def}}(\boldsymbol{x}, t) - I_{\mathrm{ref}}(\boldsymbol{x})]\nabla I_{\mathrm{def}}(\boldsymbol{x}, t)\,. \tag{5.5}$$

We solve equations (5.2)–(5.5) numerically using a forward Euler method applied to equation (5.3) in time, giving the displacement at iteration $k + 1$ as

$$\boldsymbol{u}^{k+1}(\boldsymbol{x}) = \boldsymbol{u}^{k}(\boldsymbol{x}) + \Big(I - \nabla \boldsymbol{u}^{k}(\boldsymbol{x})\Big)\boldsymbol{v}^{k}(\boldsymbol{x}, t)\Delta \mathrm{t}^{k}\,, \tag{5.6}$$

where I is the identity and Δt^k is the time step at iteration k. Equation (5.4) now becomes a system of linear PDEs, which implicitly define the velocity field $\boldsymbol{v}^k$ at each iteration $k = 0, 1, \ldots$

$$\mu\Delta\boldsymbol{v}^k(\boldsymbol{x}) + (\mu+\lambda)\nabla\Big(\nabla\cdot\boldsymbol{v}^k(\boldsymbol{x})\Big) = \boldsymbol{f}\Big(\boldsymbol{x},\boldsymbol{u}^k(\boldsymbol{x})\Big). \tag{5.7}$$

As the reference and floating images have been initially aligned according to the translation described above, the initial displacement field is set to zero for $k=0$. The velocity field is computed by solving equation (5.7), and the new displacement field is found using equation (5.6). This process is repeated until the velocity field converges.

The time step Δt^k at each iteration is chosen [175] such that

$$\Big(I-\nabla\boldsymbol{u}^k(\boldsymbol{x},t)\Big)\boldsymbol{v}^k(\boldsymbol{x},t) < \Delta\boldsymbol{u}_{\max}\ , \tag{5.8}$$

where $\Delta\boldsymbol{u}_{\max}$ is the maximum permitted displacement in one iteration, which we set to 1.5 mm. To maintain the conditioning of the system, the Jacobian of the displacement field is monitored. If it falls below 0.5, a new floating image is generated by applying the current deformation, and the deformation field is set back to zero. At the end of the algorithm, the final deformation is then found by summing all deformation fields generated during the calculation [175].

Proton MRI to xenon MRI registration

The registration between proton MRI and xenon MRI is straightforward, because these images are acquired on the same scanner during the same session, so the image coordinates have matching origins. This does not necessary mean that the images overlap, but they can be registered using a simple translation computed from the image metadata. The inflation levels are similar due to the same protocols being followed for breathhold imaging.

CT to xenon MRI registration

The final deformation field is found by combining the fields for the left and right lungs, and then applying the proton MRI-xenon MRI translation. The resulting deformation field directly relates the CT to the xenon MRI images, but the

coordinate grid is an oversampling of the xenon image coordinate grid. By extracting an appropriate subset of coronal slices, we obtain a deformation field from every voxel in the xenon MRI data to the corresponding point in CT data. Lung and lobe are then registered to the xenon image by nearest neighbour interpolation.

5.3.2 Ventilated volume analysis

Figure 5.2 shows coronal slices from two hyperpolarised ^{129}Xe ventilation images. In the right image, airway obstruction leads to regions of low signal described as *ventilation defects*. Quantification of these defects could help in the localisation, severity assessment and phenotyping of lung disease.

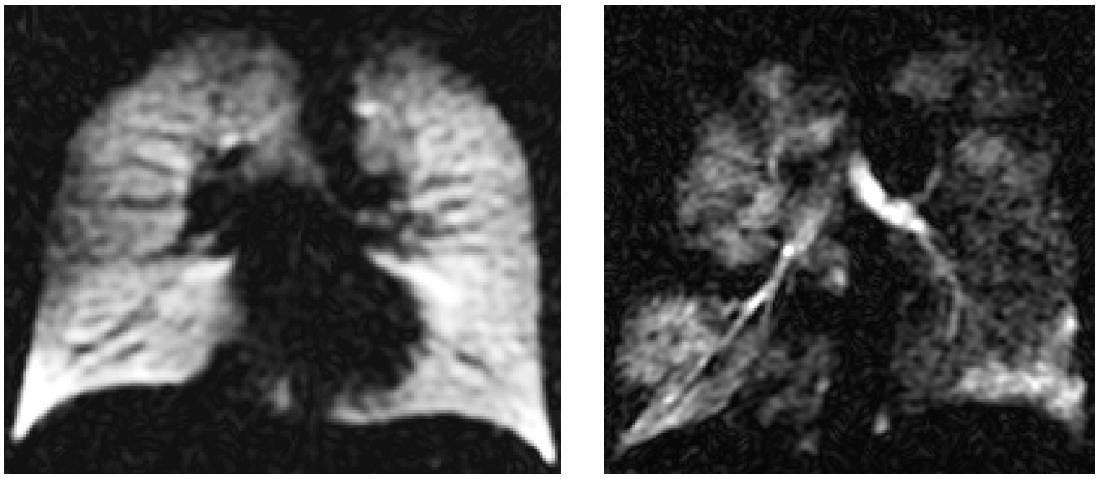

Figure 5.2: Coronal slices from hyperpolarised xenon MRI ventilation datasets. Left: control subject. Right: COPD subject. In the COPD subject, ventilation defects cause a lack of signal in some areas of the lung.

The change in the absolute volume of ventilated lung regions has been used to monitor the effects of exercise on healthy and asthmatic subjects [87], while the percentage of ventilated voxels in the lung is sensitive to the effects of smoking even in subjects who are asymptomatic of lung disease [88].

We measure the ventilated volume percentage by applying a threshold to the ventilation signal in a method similar to previous ^{3}He studies [88]. However, the signal-to-noise characteristics of ^{129}Xe are different to ^{3}He and for this reason we use the ventilation threshold used in a previous ^{129}Xe study [164] which included the COPD dataset we are using in this chapter. The threshold is set as the mean of the background noise plus twice its standard deviation. The background noise is determined from a 50 × 100 pixel region outside the lungs. The relatively low threshold results in some areas of noise exterior to the lung remaining in the ventilation mask (see left image in Figure 5.3). To remove these we perform a morphological opening operation using a spherical structural element of radius 5 mm (see right image in Figure 5.3).

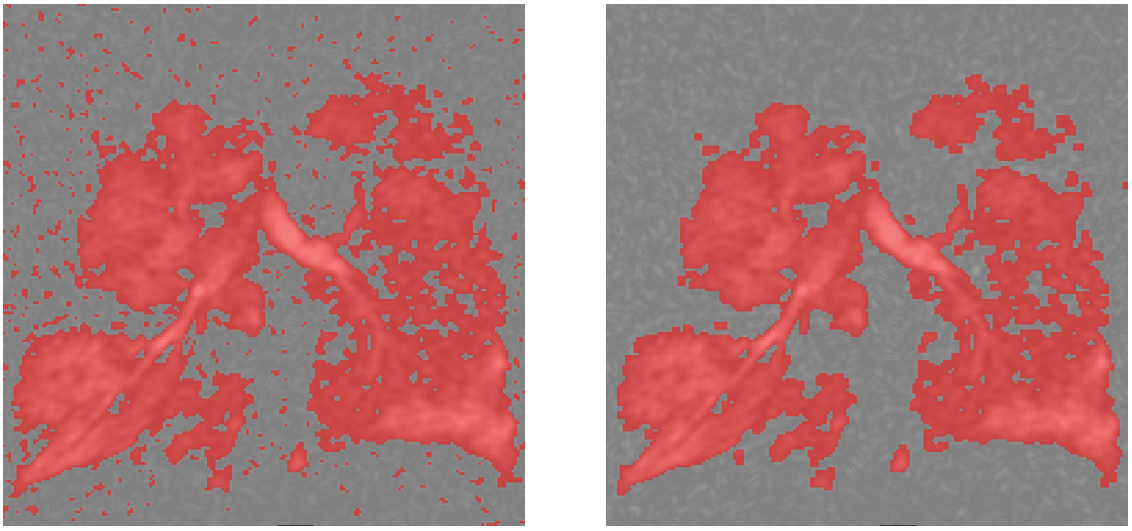

Figure 5.3: Ventilation masks for a hyperpolarised xenon MRI ventilation dataset. The ventilation mask is shown as a transparent red overlay. Left: Ventilation mask constructed by applying a threshold based on the background image noise. Ventilation defects cause a lack of signal in some areas of the lung. Right: Removal of noise outliers by applying a morphological opening using a spherical structural element of radius 5 mm.

The ventilated volume percentage is computed by registering CT lung and lobar masks to the ventilation data and computing the percentage of voxels inside the masks which are ventilated. Figure 5.9 (p.169) and Figure 5.10 (p.170) show the ventilated region masks for control and COPD subjects.

5.3.3 Apparent diffusion coefficient (ADC)

Apparent diffusion coefficients (ADCs) can be measured using diffusion-weighted imaging [56]. ADC values are sensitive to the lung microstructure and can be used to detect the destruction of alveolar tissue that is characteristic of emphysema [164].

ADC values are measured by comparing a diffusion-weighted (DW) image S against an unweighted reference image S_0. For hyperpolarised gas MRI, diffusion-weighted images are generated using a spoiled gradient-recalled-echo (SPGR) sequence with a bipolar diffusion-weighted gradient [56]. The degree of diffusion weighting depends on the shape of the gradient waveform, but can be conveniently characterised as a *b-value* [56]. A standard waveform is shown in Figure 5.4. The b-value for this can be computed by [56]

$$\mathrm{b} = (\gamma G_{\max})^2 \left[\delta^2 \left(\Delta - \frac{\delta}{3}\right) + \tau \left(\delta^2 - 2\Delta + \Delta\tau - \frac{7}{6}\delta\tau + \frac{8}{15}\tau^2\right)\right], \tag{5.9}$$

where γ is the gyromagnetic ratio of ^{129}Xe (7.3997 G·ms), $G_{\max}$ is the maximum value of the bipolar gradient, τ is the gradient ramp time, δ is the gradient duration and Δ is the gradient separation, as shown in Figure 5.4.

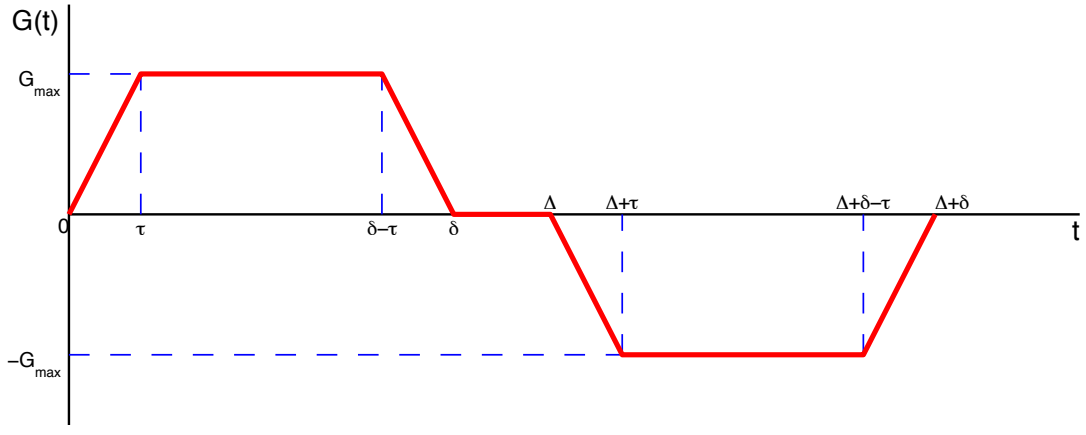

Figure 5.4: Standard bipolar gradient waveform used in hyperpolarised gas MRI diffusion-weighted imaging. The bipolar gradient strength G varies over time, with maximum amplitude $G_{\max}$. τ is the gradient ramp time, δ is the gradient duration and Δ is the gradient separation.

For the Duke images used here, the b-value is 12 s·cm^{-2} [164], while for the Oxford data, we computed the b-value as 25.3 s·cm^{-2}. Using the b-values, the ADC values at each voxel can be then approximated by [164]

$$\mathrm{ADC} = \frac{1}{b}\ln\left(\frac{S_0}{S}\right). \tag{5.10}$$

Following Kaushik et al. [164], we assume any ADC values below 0 cm^2 s^{-1} or above 0.14 cm^2 s^{-1} are due to noise, 0.14 cm^2 s^{-1} being the free diffusion coefficient of xenon in nitrogen at infinite dilution at 37°C. These voxels are excluded from the ADC map.

We compute the ADC values over lungs and lobes, using the CT masks registered to the xenon data. Voxels falling outside the masked region are also excluded. This also removes high ADC values in the trachea and main bronchi, which have large diameters. For each region we calculate the mean and standard deviation of ADC values.

The computed ADC map for the COPD patient is shown in Figure 5.11 (p.171).

5.3.4 Emphysema percentage and percentile density

Emphysema can be detected from CT images by applying a radiodensity threshold such as –950 HU [28]. Figure 5.12 (p.172) shows an example output mask for the COPD patient. The emphysema percentage (EP) is the percentage of voxels within a particular region below this threshold. An alternative measure is computing the percentile density (PD) [26] which is the value (in HU) on the frequency distribution curve below which fall 15% of the lung voxel values. We compute EP and PD values for lung and lobar regions defined by our segmentation masks. Results are shown in Table 5.5 and Table 5.6 (p.163).

5.4 Results

The following tables present the results of applying the above regional analysis methods to CT, proton MRI and hyperpolarised xenon MRI data for five subjects, comprising the four normal subjects scanned to date in the Oxford trial, and one COPD patient from the Duke study for which CT data was available. For each subject, measurements are computed for the entire lung; for the left and right lungs separately; and for each of the five pulmonary lobes.

Table 5.1 shows regional measurements for the ventilated volume percentages. These measurements are computed from hyperpolarised xenon MRI ventilation images, with the exception of subject 4. For this subject, the unweighted image in the diffusion-weighting sequence was used, as a standard ventilation image was not acquired (see section 5.5.1).

Table 5.1: Regional measurements of ventilated volume percentages. Results are shown for five patients, comprising one COPD patient from the Duke study and four control subjects from the Oxford study. The percentage of ventilated volume is computed for the whole lung, for the left and right lungs individually, and for the five pulmonary lobes. Ventilation measurements are derived from hyperpolarised xenon MRI data using a CT lobar map applied to the xenon data using image registration.

	Ventilated volume (%)				
Subject	**1**	**2**	**3**	**4**	**5**
COPD	no	no	no	no	yes
Both lungs	95.5	82.9	89.2	97.9	60.0
Left lung	94.4	78.3	89.7	97.3	65.7
Right lung	96.6	86.3	88.7	98.5	55.9
Upper right lobe	97.9	93.5	93.1	96.0	48.3
Middle right lobe	95.5	91.6	94.3	95.4	25.9
Lower right lobe	94.7	85.1	87.6	93.9	73.7
Upper left lobe	95.9	93.4	88.1	95.2	50.8
Lower left lobe	91.6	67.5	86.6	89.6	90.8

Table 5.2 shows the means and standard deviations of the apparent diffusion coefficients (ADC) computed for the image regions in each subject. These values were calculated using the diffusion-weighted hyperpolarised xenon MRI images. No diffusion-weighted imaging was performed for subject 2.

Table 5.2: Regional measurements of the apparent diffusion coefficient (ADC). Results are shown for three control subjects from the Oxford study and one COPD patient from the Duke study. ADC measurements are derived from diffusion-weighted hyperpolarised xenon MRI data using a CT lobar map applied using image registration. Mean ADC values are computed for the whole lung, for the left and right lungs individually, and for the five pulmonary lobes. Standard deviations are shown in parentheses. No diffusion-weighted images were acquired for Subject 2.

	ADC mean values (standard deviations) /cm^2s^{-1}			
Subject	**1**	**3**	**4**	**5**
COPD	no	no	no	yes
Both lungs	**0.039** (0.013)	**0.046** (0.012)	**0.041** (0.014)	**0.053** (0.027)
Left lung	**0.040** (0.010)	**0.047** (0.010)	**0.044** (0.012)	**0.049** (0.023)
Right lung	**0.039** (0.014)	**0.046** (0.012)	**0.040** (0.014)	**0.055** (0.029)
Upper right lobe	**0.038** (0.007)	**0.052** (0.009)	**0.045** (0.013)	**0.056** (0.023)
Middle right lobe	**0.037** (0.006)	**0.049** (0.012)	**0.045** (0.015)	**0.049** (0.025)
Lower right lobe	**0.037** (0.010)	**0.044** (0.010)	**0.036** (0.010)	**0.047** (0.025)
Upper left lobe	**0.043** (0.010)	**0.050** (0.010)	**0.051** (0.012)	**0.053** (0.023)
Lower left lobe	**0.035** (0.009)	**0.044** (0.010)	**0.038** (0.010)	**0.042** (0.021)

For comparison, we show values for previous ^{129}Xe research in Table 5.3, and values for previous ^{3}He research in Table 5.4. We discuss these works in section 5.5.3 below.

Table 5.3: Regional measurement results for previously published work in ^{129}Xe imaging. Results are shown for percentage ventilated volume and apparent diffusion coefficients (ADC). The ADC standard deviations are shown in parentheses.

^{129}Xe	Regional measure			
Published work	**Ventilated volume /%**		**ADC /cm^{2}s^{-1} (standard deviation)**	
	Healthy	**COPD**	**Healthy**	**COPD**
Kirby et al. [62]	96	66	0.053 (0.002)	0.079 (0.015)
Kaushik et al. [164]			0.036 (0.003)	0.056 (0.008)

Table 5.4: Regional measurement results for previously published work in ^{3}He imaging. Results are shown for percentage ventilated volume and apparent diffusion coefficients (ADC). The ADC standard deviations are shown in parentheses. *Ventilated volume percentages for Stavngaard et al. [44] are visual estimates.

^{3}He	Regional measure			
Published work	**Ventilated volume /%**		**ADC /cm^{2}s^{-1} (standard deviation)**	
	Healthy	**COPD**	**Healthy**	**COPD**
Plotkowiak et al. [9]			0.214 (0.042)	0.358 (0.115)
Kirby et al. [62]	96	80	0.246 (0.021)	0.481 (0.121)
Stavngaard et al. [44]	93*	71*	0.19	0.30
Woodhouse et al. [88]	90	68		

Table 5.5 shows regional measurements for the emphysema percentage (EP). The corresponding percentile density (PD) measurements are shown in Table 5.6. These measurements are derived from the CT data.

Table 5.5: Regional measurements of emphysema percentage (EP). Results are shown for five patients, including one COPD patient from the Duke study and four control subjects from the Oxford study. The emphysema percentage (EP) is computed from density values in the CT image falling below an emphysema threshold (–950 HU). This percentage is computed for the whole lung, for the left and right lungs individually, and for the five pulmonary lobes.

	Emphysema percentage (%)				
Subject	**1**	**2**	**3**	**4**	**5**
COPD	no	no	no	no	yes
Both lungs	0.4	0.0	0.1	1.3	13.3
Left lung	0.5	0.0	0.2	1.5	14.1
Right lung	0.3	0.0	0.1	1.2	12.7
Upper right lobe	0.1	0.0	0.0	0.6	11.7
Middle right lobe	0.5	0.0	0.1	3.7	14.5
Lower right lobe	0.3	0.1	0.1	0.3	12.9
Upper left lobe	0.5	0.0	0.1	1.8	17.0
Lower left lobe	0.4	0.0	0.2	1.3	8.9

Table 5.6: Regional measurements of percentile density (PD). Results are shown for five patients, including one COPD patient from the Duke study and four control subjects from the Oxford study. The percentile density (PD) is the value (in HU) on the frequency distribution curve below which fall 15% of the lung voxels. This measure is computed for the whole lung, for the left and right lungs individually, and for the five pulmonary lobes.

	Percentile density (HU)				
Subject	**1**	**2**	**3**	**4**	**5**
COPD	no	no	no	no	yes
Both lungs	–904	–861	–900	–916	–946
Left lung	–904	–863	–899	–917	–948
Right lung	–904	–859	–900	–914	–945
Upper right lobe	–896	–840	–895	–912	–942
Middle right lobe	–915	–817	–914	–929	–949
Lower right lobe	–902	–870	–898	–907	–945
Upper left lobe	–907	–853	–901	–920	–953
Lower left lobe	–899	–871	–896	–914	–934

Figure 5.5 shows the rigid proton MRI to xenon MRI registration by illustrating the xenon ventilation image superimposed on the MRI data for the COPD patient. To compute lobar results for the xenon images, the CT lobe map is registered to the xenon data using the non-rigid registration technique described above. The registered lobar map is shown overlaid on the ten coronal slices of the xenon ventilation images in Figure 5.6 for the COPD subject, and Figure 5.7 for a control subject. To illustrate ventilation images in a 3D context, the lobar segmentations are visualised in Figure 5.8 as a surface rendering with the xenon ventilation images displayed as cut-away coronal slices.

Ventilated regions are shown in Figure 5.9 for a control subject, and Figure 5.10 for the COPD patient. The xenon MRI ventilation image is shown with a transparent blue overlay illustrating the ventilation masks produced by applying a threshold and morphological operations. The boundary of the ventilated region is shown in red. The percentage of ventilated voxels is computed by comparing the ventilated mask to the registered lung and lobar masks.

Figure 5.11 shows the apparent diffusion coefficient (ADC) values computed for the xenon MRI diffusion images for the COPD patient. Only voxels contained within the lung and lobe masks are considered, and voxels whose ADC value falls outside the range [0, 0.14] $cm^2 s^{-1}$ are also excluded.

Finally, Figure 5.12 illustrates the computation of emphysema percentage (EP) from CT images from the COPD patient. Voxels with density values below the emphysema threshold (–950 HU) are highlighted in red and counted to determine the percentage of emphysematous voxels in each region.

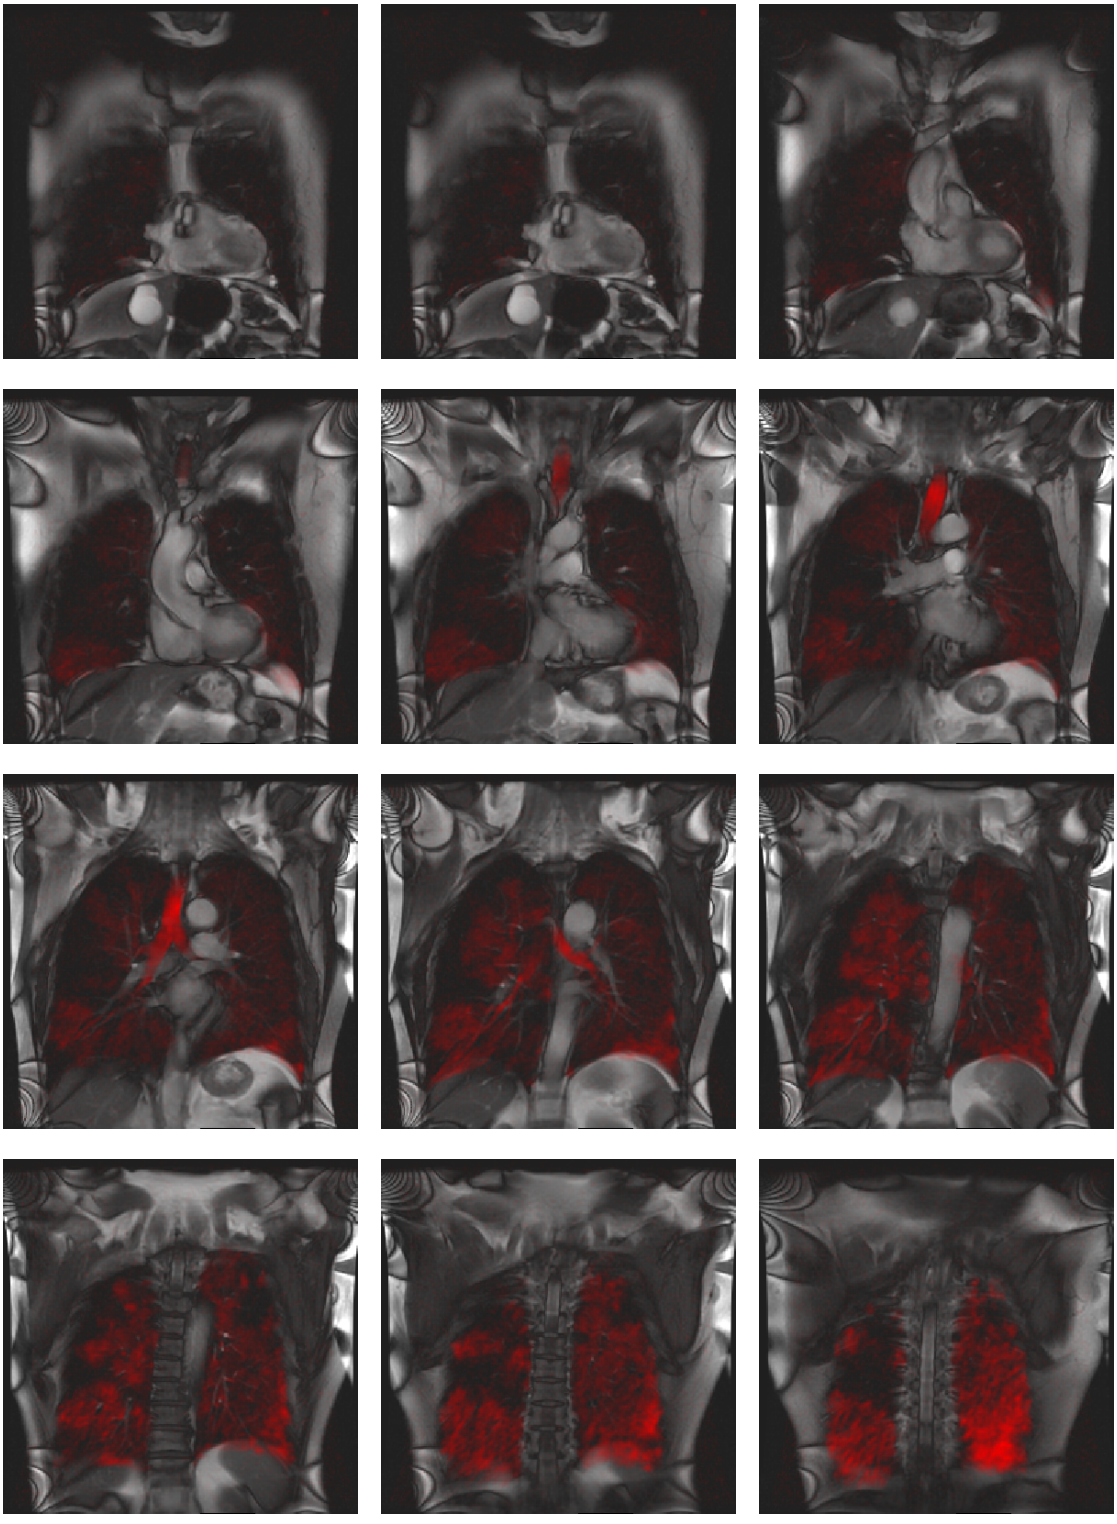

Figure 5.5: Fusion of anatomical proton MRI and ventilation xenon MRI for a COPD patient. The background is the proton MRI data, and the xenon ventilation data are shown overlaid in transparent red, with the opacity proportional to the signal.

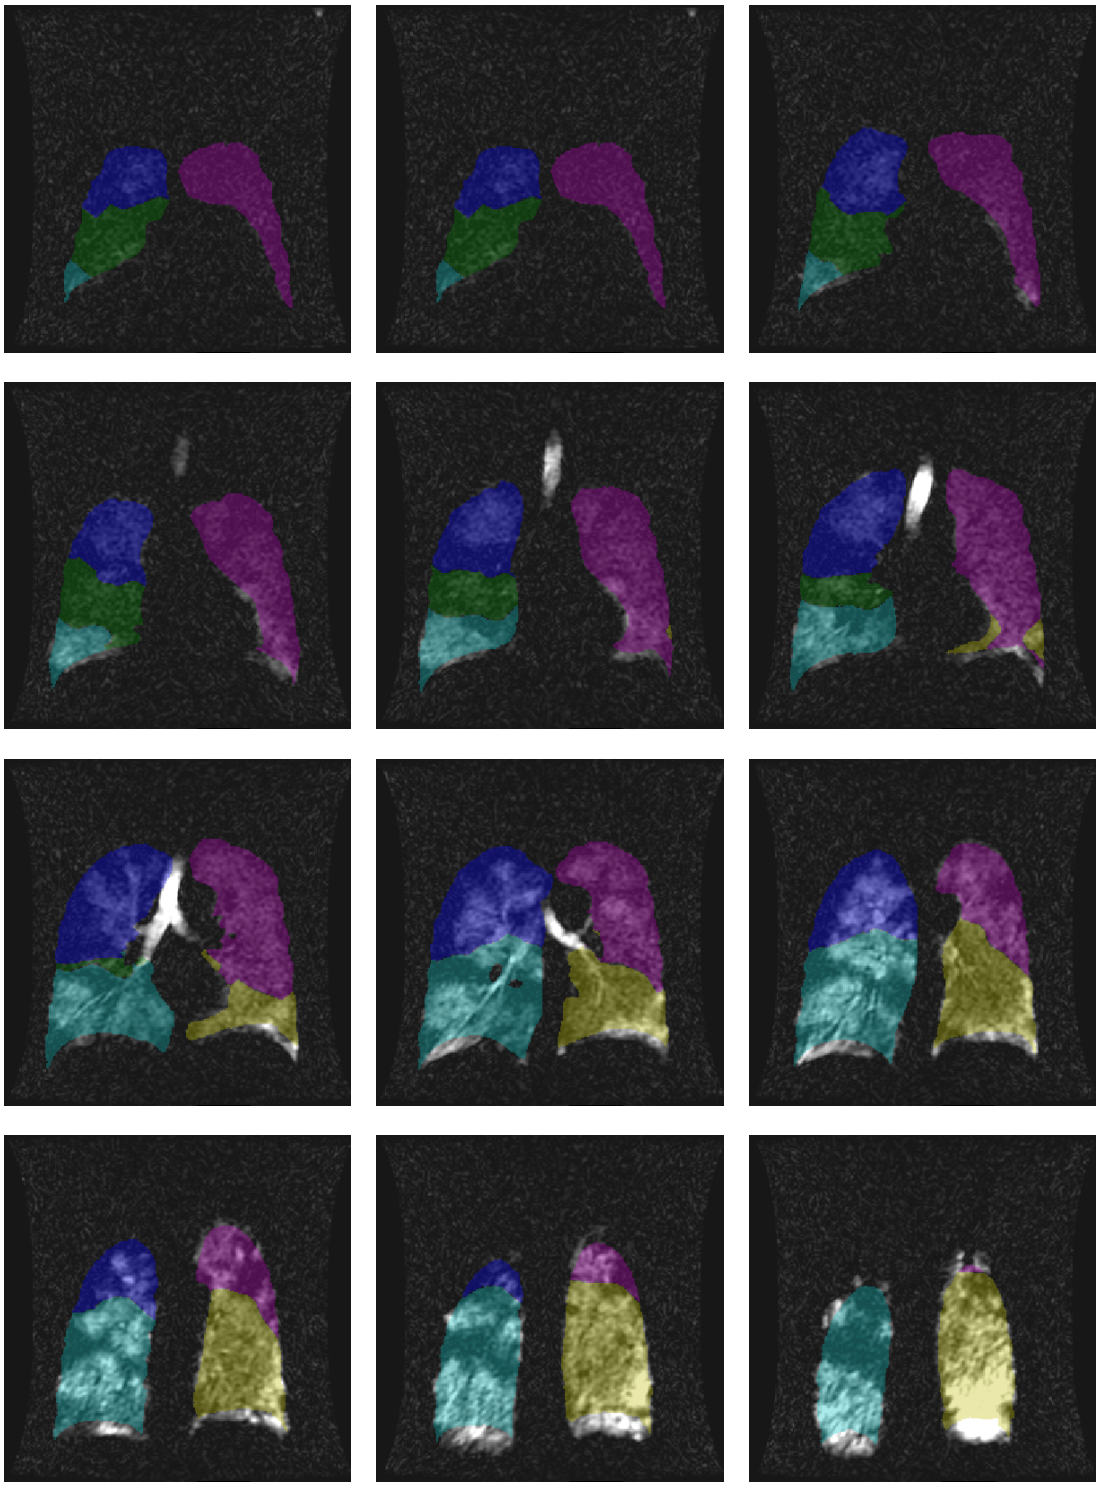

Figure 5.6: Fusion of lobe masks and xenon MRI ventilation images for a COPD patient. Each image shows one coronal slice of the xenon ventilation image in greyscale, with a colour lobe map shown as a transparent overlay. The lobe map is generated by registration of the CT lobe map to the xenon images.

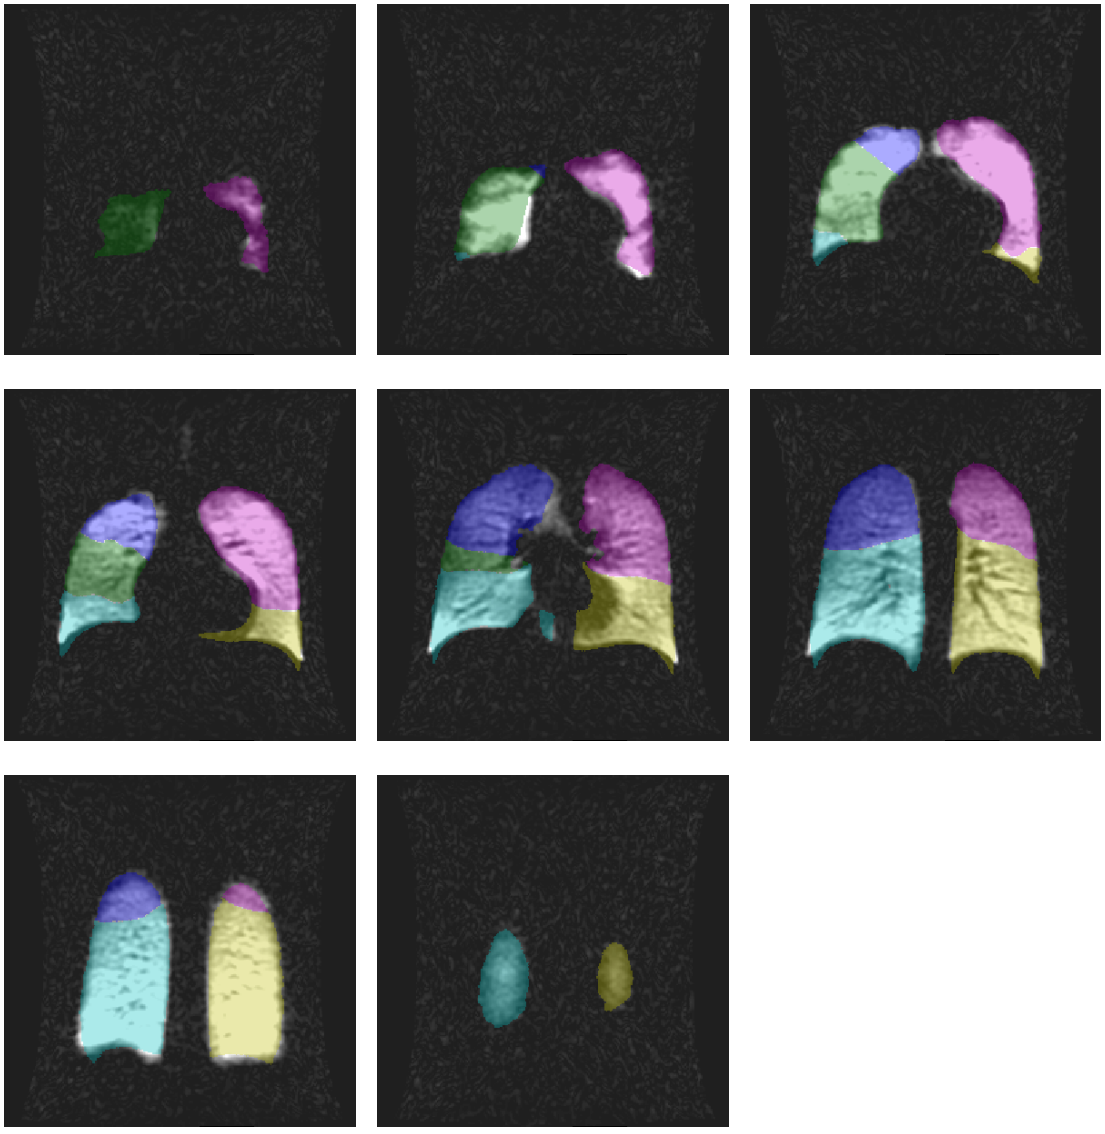

Figure 5.7: Fusion of lobe masks and xenon MRI ventilation images for a control subject. Each image shows one coronal slice of the xenon ventilation image in greyscale, with a colour lobe map shown as a transparent overlay. The lobe map is generated by registration of the CT lobe map to the xenon images.

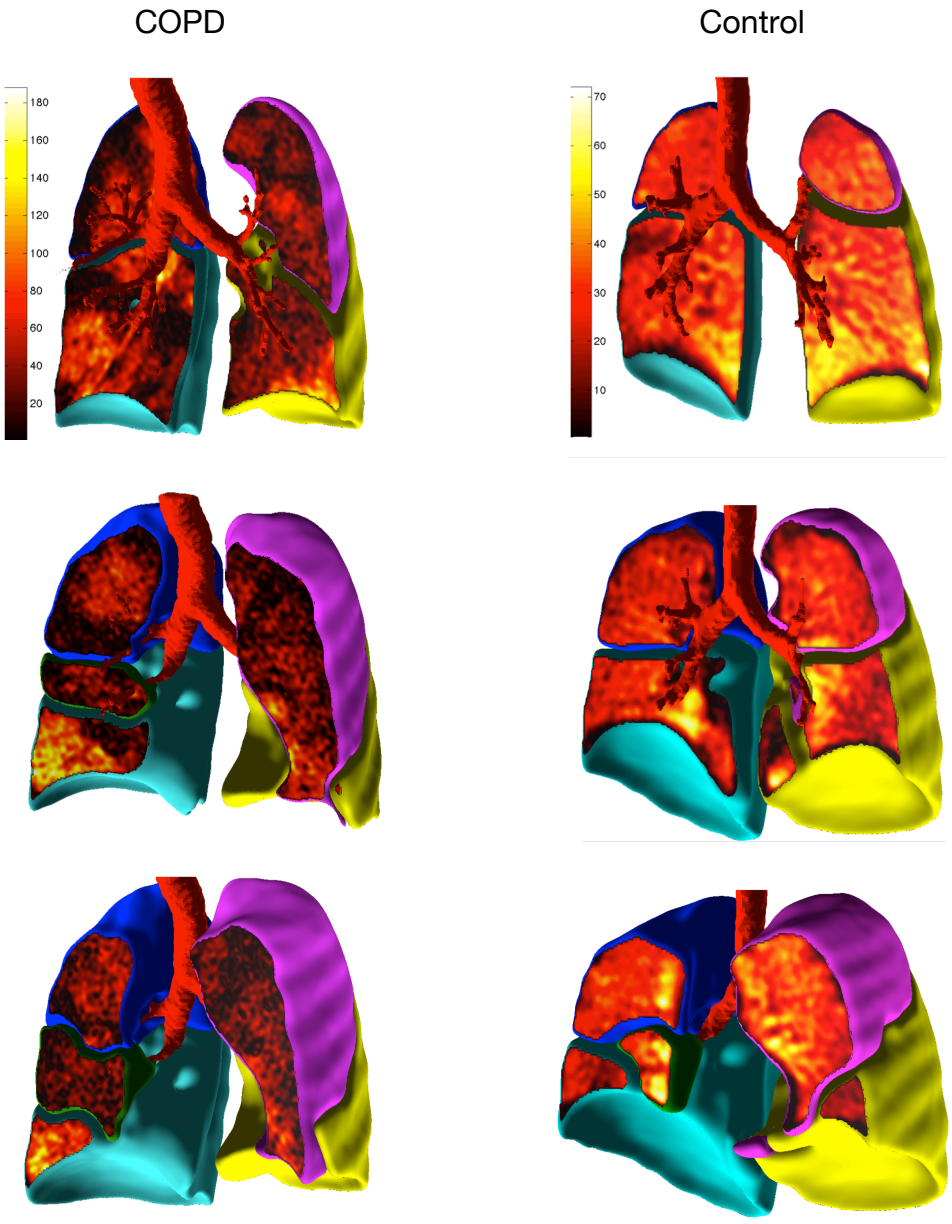

Figure 5.8: Xenon MRI ventilation images illustrated using a 3D rendering of the lobe masks. Images are shown for a patient with COPD (left) and a control subject (right). These images are intended to assist with the qualitative interpretation of xenon MRI images by showing ventilation data in the context of the lobes. Similar images can be generated for other image measurements such as ADC maps. The lobes are segmented from CT and registered to the xenon ventilation images. The lobes are visualised as smoothed surface renderings cut away at coronal slice positions (three of which are visualised here). The ventilation images are shown for these positions, using a colour scale of dark red for low ventilation and yellow for high ventilation.

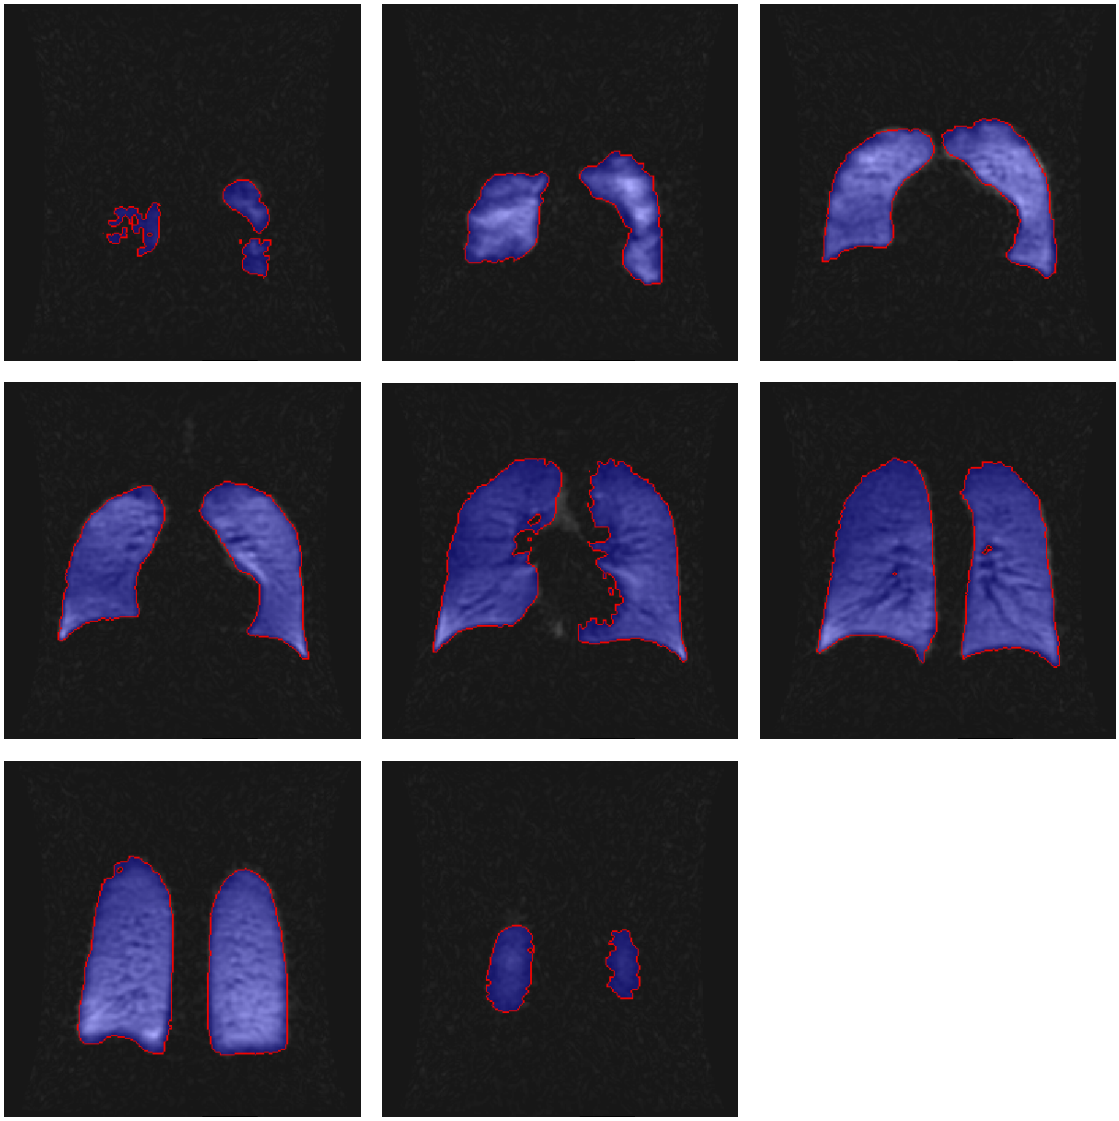

Figure 5.9: Ventilation masks generated for a control subject using hyperpolarised xenon MRI. Each image shows one slice of the hyperpolarised xenon MRI ventilation image. The transparent blue overlay highlights voxels within the ventilation threshold, after removal of outliers using morphological opening. The boundaries of the ventilated regions are outlined in red.

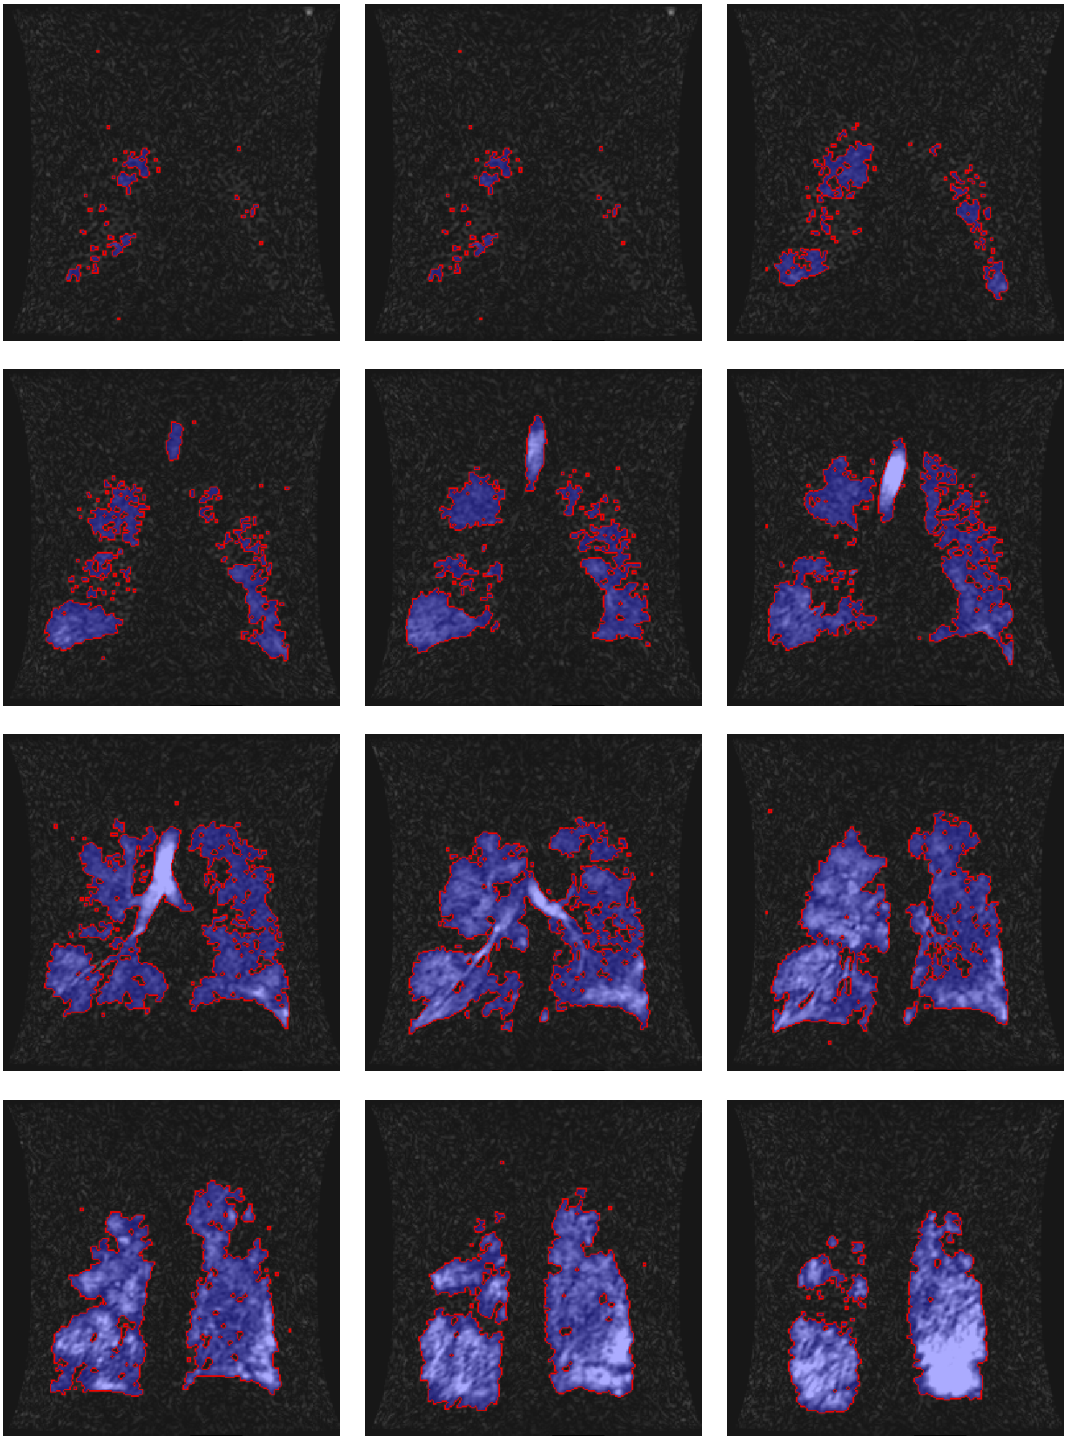

Figure 5.10: Ventilation masks generated for a COPD subject using hyperpolarised xenon MRI. Each image shows one slice of the hyperpolarised xenon MRI ventilation image. A transparent blue overlay highlights voxels within the ventilation threshold, after removal of outliers using morphological opening. The boundaries of the ventilated regions are outlined in red.

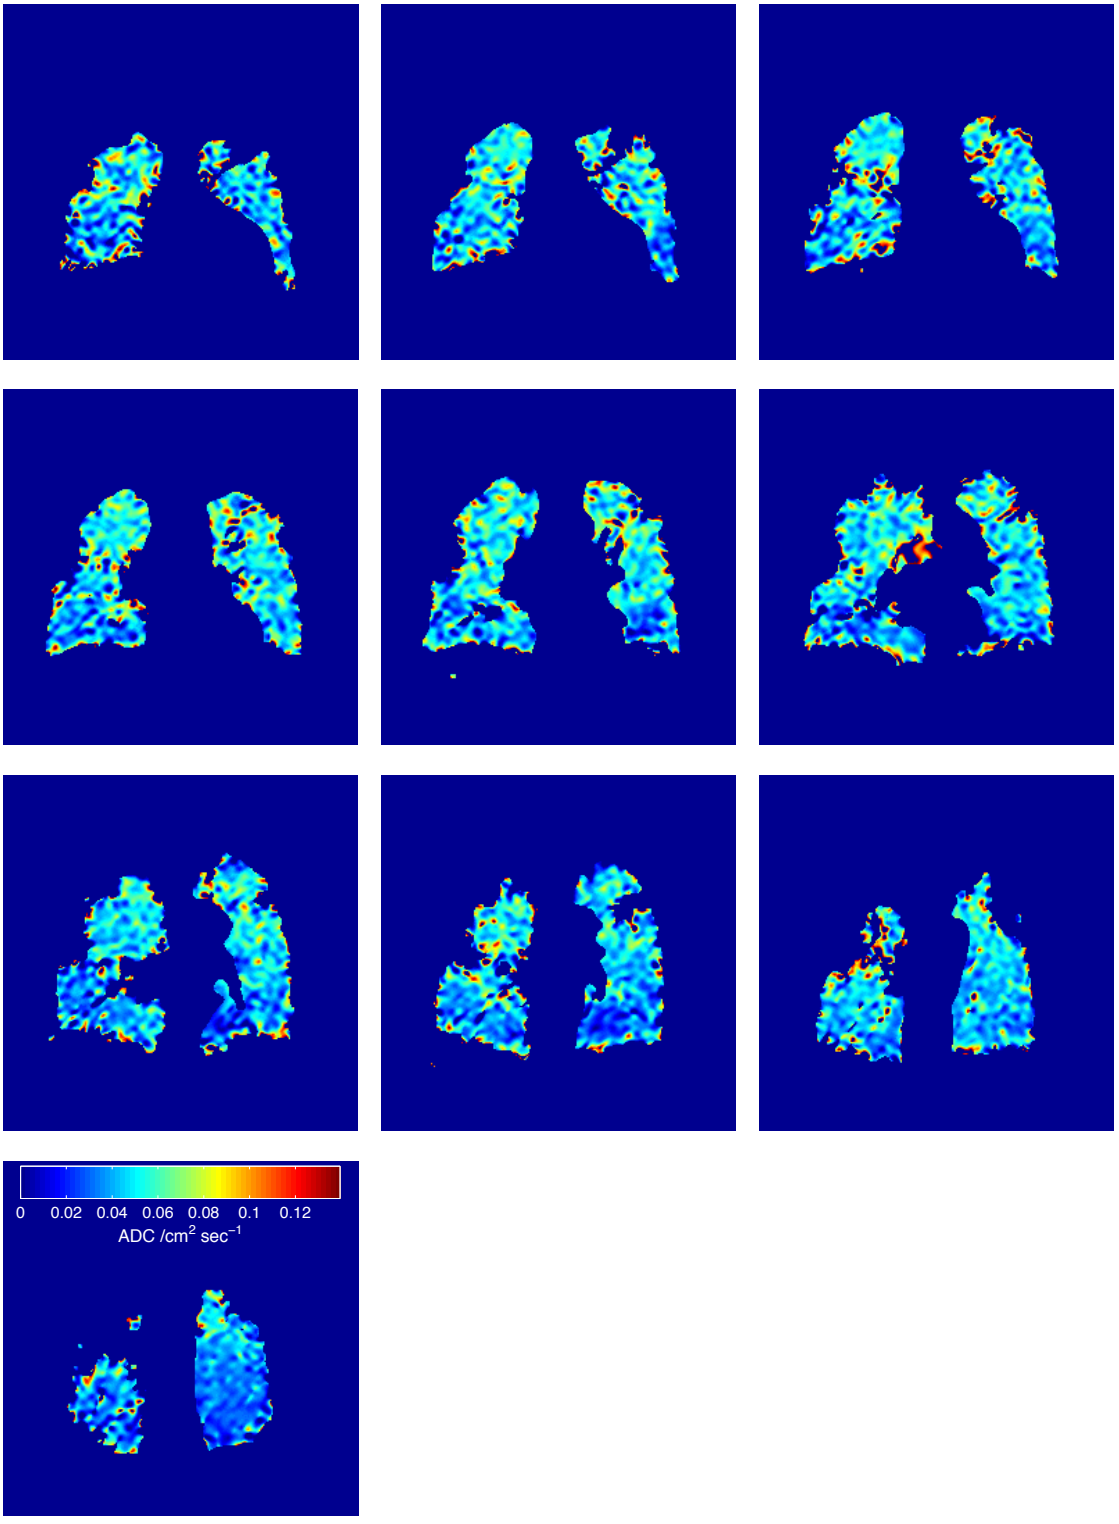

Figure 5.11: Apparent diffusion coefficient maps for a COPD patient using hyperpolarised xenon MRI diffusion-weighted imaging. The images show the ADC map for each coronal slice of the xenon MRI diffusion data. Red colouring denotes a higher ADC value. Voxels outside of the lung mask are omitted. ADC values below zero or above 0.14 cm^2s^{-1} are assumed to be noise and the corresponding voxels are removed from the ADC maps.

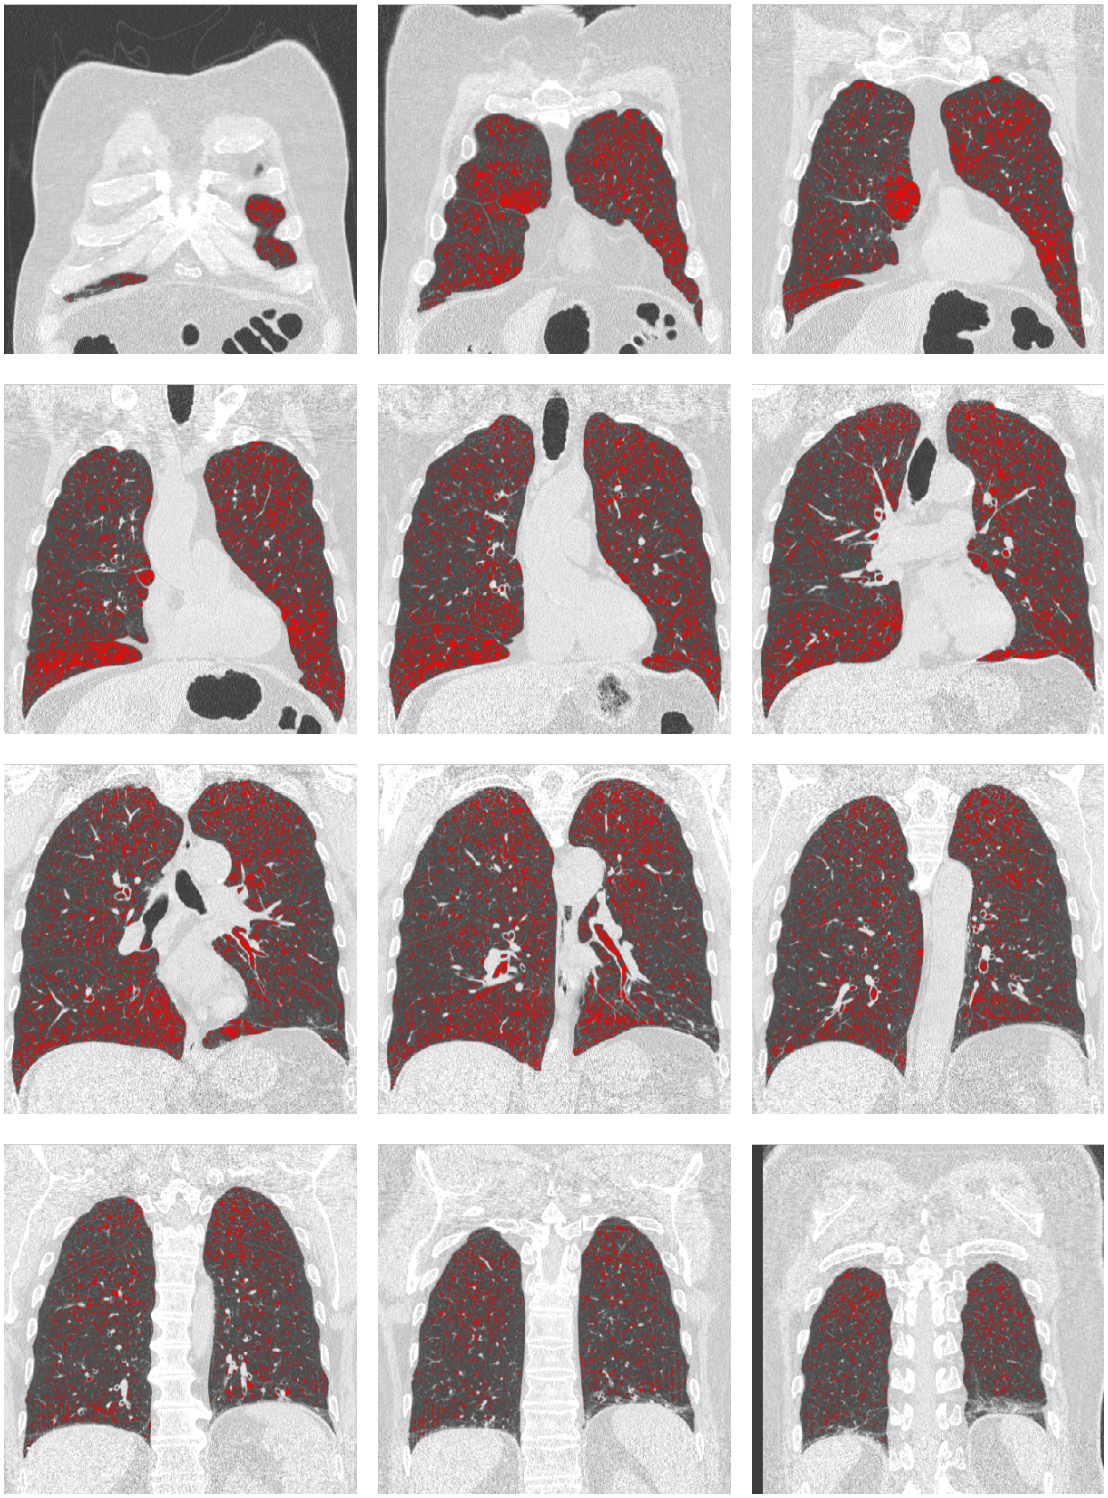

Figure 5.12: Emphysema detection results for a COPD patient. The dataset has been resliced in a coronal orientation and a regularly spaced subset of coronal slices is shown. Voxels below the emphysema threshold (–950 HU) are highlighted in red. The percentages of these voxels within a given region are used to quantity emphysema.

5.4.1 Validation

To give an indication of the registration performance, we compute the Dice, precision and recall measures previously introduced in equations (3.6)–(3.8) (see section 3.3.5, p. 67), along with the mean and maximum displacements of the lung boundaries. Our reference images are created by performing a rigid registration of a manually segmented lung volume from proton MRI onto the coordinates of the xenon images. We compare these with the fluid registration of the CT lung segmentation masks onto the xenon coordinates. These figures are not an independent measure of registration accuracy, but they give some indication of the expected performance of the registration method, which balances matching of the images with maintaining a smooth deformation field.

Table 5.7: Validation of automatic registration methods between CT and MRI. Lung masks from CT are registered to the xenon MRI data using the initial rigid and the final fluid methods. To assess these registrations, we create approximate ground truth masks by registering manually segmented lung masks from proton MRI data to the xenon MRI using a translation between coordinate origins defined by image metadata. We compute Dice, precision and recall measurements between the masks, and the mean and maximum distances between the mask boundaries.

		Mask measures			Distance between surfaces (mm)	
Subject	**Registration**	**Dice**	**Precision**	**Recall**	**Mean**	**Maximum**
1	Rigid (initial)	0.81	0.72	0.93	8.53	43.8
	Fluid (final)	0.92	0.94	0.90	4.68	42.4
2	Rigid (initial)	0.88	0.86	0.90	5.91	42.8
	Fluid (final)	0.93	0.97	0.90	5.44	48.8
3	Rigid (initial)	0.79	0.72	0.87	9.76	86.2
	Fluid (final)	0.93	0.95	0.91	5.56	90.1
4	Rigid (initial)	0.83	0.95	0.74	10.59	45.1
	Fluid (final)	0.84	0.81	0.89	10.53	63.6
5	Rigid (initial)	0.85	0.80	0.89	7.04	42.1
	Fluid (final)	0.92	0.92	0.91	4.02	32.2
Mean	**Rigid (initial)**	**0.83**	**0.81**	**0.87**	**8.37**	**52.0**
	Fluid (final)	**0.91**	**0.92**	**0.90**	**6.06**	**55.4**

5.5 Discussion

5.5.1 Data challenges

We analysed the images from the first four patients of the clinical trial. Being the initial data from a novel trial, these images presented several challenges which we discuss here.

For two patients (subjects 1 and 3), the ventilation levels were not correctly matched for the proton MRI and the xenon MRI imaging due to the protocols not being correctly followed. To adjust for this we manually corrected the lung boundaries that were computed from the proton MRI segmentation. To do this, the automatic segmentations were registered to the xenon data and then manually edited prior to being used in the CT-MRI registration. This is possible for the control subjects since they do not show ventilation defects. However, manual correction might be more subjective in the case of a COPD patient where severe ventilation defects may make the lung boundaries unclear. This highlights the importance of the inflation level in our analysis. Improvements to the protocol should reduce the likelihood of a large ventilation mismatch.

One patient (subject 2) was unusually large, which necessitated an increase in the slice spacing to 20 mm. Ventilation images were successful but the diffusion-weighted sequences did not work for this subject, so ADC maps cannot be computed. For a different patient (subject 4), no ventilation sequence was performed, as the xenon supply was exhausted performing diffusion-weighted imaging and then multiple unsuccessful attempts at dissolved-phase imaging. However, we were able to generate a ventilation map using the unweighted image from the diffusion-weighted sequence, which is equivalent to a lower-resolution ventilation image.

Diffusion-weighted sequences were obtained for three subjects. At the time of writing, these sequences and the reconstruction methods are still under development. The ADC results are thus shown here for the purposes of validation of the registration and segmentation methods proposed, and the ADC values themselves should not be taken to provide a reliable characterisation of the subjects' lungs. This does not affect the ADC results for the COPD patient, which was part of a previous study.

5.5.2 Comparison of COPD and control data

Comparing the CT analysis of emphysema percentage in Table 5.5 and percentile density in Table 5.6 (p.163) with the ventilated volume percentages of Table 5.1 (p.160), the COPD subject shows clear effects of disease with increased emphysema percentages and decreased ventilated volume percentages. The effect is particularly pronounced in the upper lobes, which is consistent with the expected location of emphysema. The relationship between emphysema and ventilation defects is interesting, as it suggests mechanisms linking the two such as airway remodelling or premature airway closing which may be related to the change in compliance. We can explore this relationship further using clinical trial data for subjects with COPD.

Interestingly, for the COPD patient, regional variations in the ADC are small and within the standard deviations. We might expect a stronger correlation between the ADC and the emphysema percentage, since the two may be related. The upper left lobe shows a slightly higher ADC (0.052 cm^2s^{-1}) than the others, but the other lobes do not show the same correlations. The lobar ADC values we compute for the COPD patient are consistent with emphysema across the whole lung, in contrast to the variability in emphysema percentage and ventilated volume percentage. A wider range of ADC data will allow us to

assess the expected variation of values across the lobes for healthy and COPD subjects.

As we would expect, the control subjects show high ventilated volume percentages and low emphysema percentages. Ventilated volume percentages in the COPD patient show more variability between lobes, which is not surprising due to the clearly visible ventilation detects in Figure 5.10 (p.170). One aim of our future work is to derive a regional "map of normality" for the control subjects, against which we can assess the COPD patients.

5.5.3 Comparison with previous results

Previous results are shown in Table 5.3 for ^{129}Xe and Table 5.4 for ^{3}He (p.162). Aziz et al. [41] define a "high emphysema percentage" as above 25% and a "low emphysema percentage" as below 25%. Over the whole study from which our data were derived [164], the mean parenchymal ADC is 0.036 $cm^2\,s^{-1}$ for healthy volunteers and 0.056 $cm^2\,s^{-1}$ for COPD subjects with emphysema. Kirby et al. [62] find higher ADC values for healthy and COPD subjects, although their protocol is slightly different, using 0.5 l of ^{129}Xe with 0.5 l of ^{4}He compared to 1 l of ^{129}Xe for Kaushik et al. [164]. Their b-value is 1.6 $s{\cdot}cm^{-2}$ compared to 12 $s{\cdot}cm^{-2}$ for Kaushik et al. The change in b-value can have an effect on the results because equation (5.10) on p.159 is only an approximation.

Across the whole lung, Kirby et al. [62] report ^{129}Xe ventilated volume percentages of 96% for healthy subjects, compared to 66% for subjects with COPD, which is consistent with our measurements. Comparing with ^{3}He measurements, Kirby et al. find the healthy ventilated volume percentages to be unchanged at 96%, while the COPD patients show an higher average of 80% for ^{3}He, suggesting that for static ventilation images, ^{129}Xe is more sensitive to disease than ^{3}He, due to the different diffusing properties of the

gas [62]. For ^{3}He data, Woodhouse et al. [88] report ventilated volume percentages of 90% for healthy never-smokers, compared to 68% for smokers with COPD. Due to the differing gas properties, ADC measurements for ^{3}He data are an order of magnitude higher than for ^{129}Xe.

5.5.4 Errors in registration

Assessment of the registration method is important for validation of the pipeline. The metrics in Table 5.7 (p.173) show broadly good results. There is a clear difference between the values for the rigid and non-rigid registrations, suggesting the importance of using a non-rigid approach in registering these images from different inflation levels. Values could be improved by reducing the weighting of the regularisation term in the registration algorithm, but this would lead to a less smooth (and less anatomically realistic) deformation field.

The maximum boundary displacement is very high and this suggests some significant errors in parts of the registration. There are three major sources of error in the segmented lung boundaries used for registration: unclear boundaries caused by partial volume effects in the thick-slice MRI data; inconsistent boundary segmentation around the main vessels in the centre of the lung; and leakage of the automatic proton MRI segmentation into regions outside the lung, again arising from partial volume effects. Improved proton MRI data with thinner slices should improve all of these issues, and allow additional validation of the registration methods using anatomical landmarks. Such validation can be used to assess the effects of registration inaccuracy and determine whether alternative registration methods could be used to improve results.

5.6 Conclusions

In this chapter we have developed a complete, automated, general-purpose pipeline for deriving lobar measurements from hyperpolarised xenon MRI data. This pipeline provides a key component in the analysis of data from the clinical trial. We illustrated our methods with preliminary results, using four datasets from the clinical trial and one COPD patient from a previous study. We plan to extend our pipeline to derive additional measurements, and to extend the regional analysis, for example to a segmental level.

In section 5.3.1 we developed a simple non-rigid registration method for mapping CT-derived data (in this case the pulmonary lobe segmentations) onto xenon MRI data to facilitate regional analysis. Our method is based on the distance transform to the segmented CT and proton MRI lung boundaries, and is robust to thick slice data and the sliding effects at the lung boundaries. Thin-slice proton MRI data may allow us to improve the accuracy of our registration method by introducing mutual information, and to perform further validation using visible landmarks.

In sections 5.3.2–5.3.4 we illustrated the pipeline in use by deriving four regional measurements at a lung and lobar level. These comprised of two related CT-based emphysema measures (emphysema percentage and percentile density) and two functional measures (ventilated volume percentage and apparent diffusion coefficient). We found that our measurements fall within the expected values for the COPD and control subjects. Further data will be used to analyse the lobar distribution and determine “deviation from normality”.

6 Future work

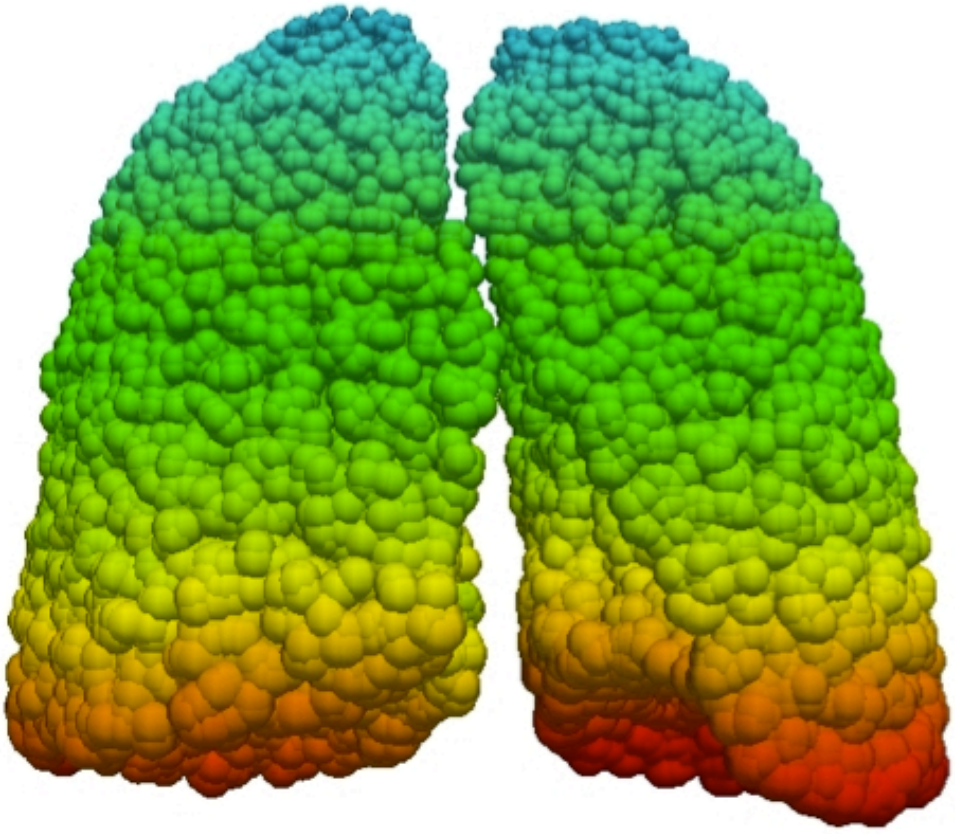

Chapter 6: Future work



In this chapter we briefly describe our future plans, in the form of two research proposals. We summarise the two proposals here and present more detail in the following sections.

1. We will apply the regional analysis methods we have developed in this thesis to new data from the clinical trial. We will compute regional measurements of lung function (ventilated volume percentage and ADC values) at a lobar level, and investigate extending this image analysis to a pulmonary segmental level. We will develop further validation methods as required. The aim of this work is to validate our methods and demonstrate their potential effectiveness in guiding diagnosis and choice of treatment.
2. We will investigate characteristics of lung disease using personalised computational modelling. We will extend our segmentation methods to build personalise geometric models based on the clinical trial data and integrate these with the ventilation models already developed in our group. We will simulate ventilation images based on the modelling results and compare these with the functional imaging. Using clinical knowledge we will investigate whether disease-specific modification of model parameters (for example, introducing airway restrictions to model small airways disease) can qualitatively induce the pathology seen in the functional images.

6.1 Regional image analysis applied to the clinical trial

In chapter 5 we demonstrated the complete pipeline for lobar-based analysis of ventilated volume and ADC measurements in hyperpolarised xenon MRI data. We plan to extend this pipeline to develop further regional measurements and apply it to more data from the clinical trial. Our aim is to validate the pipeline and investigate the potential of regional measurements as a future aid to disease assessment and treatment planning in clinical environments.

Our work plan is as follows:

- From the clinical trial, obtain CT, proton MRI, xenon MRI ventilation and xenon ADC images for at least 20 subjects, including COPD and control subjects;
- Compute lung- and lobe-based ventilated volume percentages using the methods described in chapter 5;
- Compute lung- and lobe-based apparent diffusion coefficient (ADC) measurements from the diffusion images;
- Compare these regional measurements (ventilated volume percentage and ADC) to traditional quantifiers of COPD severity, i.e. pulmonary function tests and CT emphysema percentage (EP) and percentile density (PD);
- By comparing data across the normal and COPD subjects, investigate whether it is meaningful to construct a “normal” ventilation and ADC distribution and use this to define a quantity describing “deviation from normality”. Such a measure has the potential to be a clinical indicator of disease;

- Develop and validate a method for pulmonary segment extraction. This method will be based on labelling of the bronchial tree by segment and subdivision into segments by either performing a watershed transform or by allocating voxels to the nearest segmental bronchus, according to which performs better in validation tests. Validation will involve a clinician assigning a visual score based on shape and location of the segments, relation to the bronchial tree and correspondence with visible features such as minor fissures;
- If the pulmonary segment extraction is shown to be meaningful, extend lobar analysis of ventilated volumes, ADC measurements and density masks to a segmental level;
- Validate registration techniques and investigate whether an improved registration method is required.

6.2 Personalised computational modelling

In this work we will develop personalised airway models and couple them to the computational airflow simulations currently being developed by our group [176]. These simulations are being used to simulate particular disease characteristics by modifying model parameters guided by clinical knowledge. Our aim is to link our work to these models at the input and output stages. We will construct personalised airway models and drive simulations directly from our software and simulate ventilation images based on the results. By comparing the simulated images with functional imaging we have a tool for validating the personalised airflow simulations and investigating the disease mechanisms which influence functional imaging results.

The simulation we will use has been developed in our group by Kim et al. [176] and is based on a mathematical model coupling acinar dynamics to airflow in the conducting airways. Airflow through the conducting airway tree is modelled using a pressure distribution model with conservation equations governing flow at the bifurcation points (Figure 6.1 B). At the end of the conducting airways, the acini are modelled using a lumped parameter model [177], computing acinar pressure based on the fluctuation in pleural pressure during breathing (Figure 6.1 A).

A sample output of the airflow model is shown in Figure 6.2. Two models were generated. The left image shows airflow results for a model generated with no constrictions. The right image shows results for the model where some of the small airways were constricted to simulate the effects of disease.

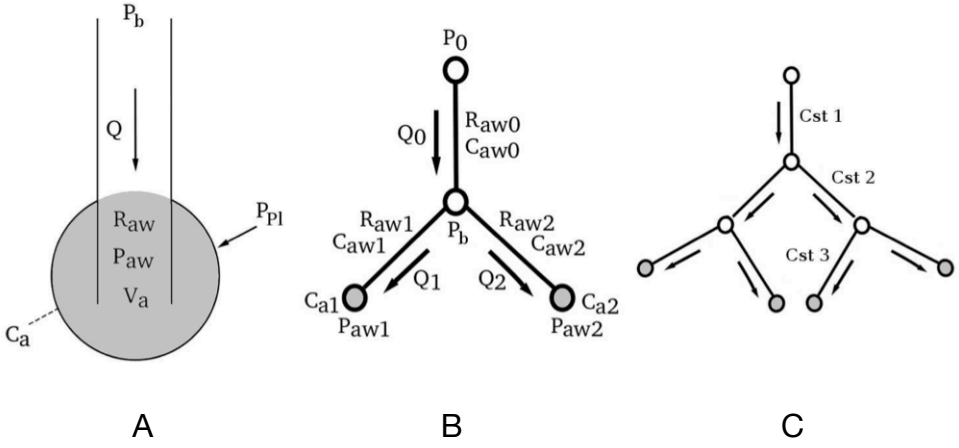

Figure 6.1: Schematic diagram of the coupled ventilation model used in simulation studies. A: Acinar units are represented by a lumped parameter model. B: Airflow in the conducting airways is modelled using a pressure distribution model. C: An ideal airway model with three bifurcations, indicating the location of an applied bronchoconstriction (the sites labelled Cst 1–3). (Courtesy of Minsuok Kim.)

Our work plan is as follows:

- Identify data from the clinical trial representing the disease phenotypes we wish to investigate, plus control subjects;
- Integrate existing work on personalised airway model development to grow the segmented airway tree to a full bronchial tree down to the acinar level using a volume-filling method based on the lobe segmentations;
- Develop a method for simulating ventilation images based on the results of the airflow modelling;
- Using clinical knowledge, identify model parameters in the simulation or the airway growing model which can be modified in order to simulate known pathological effects of disease, for example airway density and constriction, or mechanical properties such as compliance;
- Perform a range of airflow simulations for a sweep of parameter values and compare the results with the simulated images derived from the functional images obtained from the clinical trial.

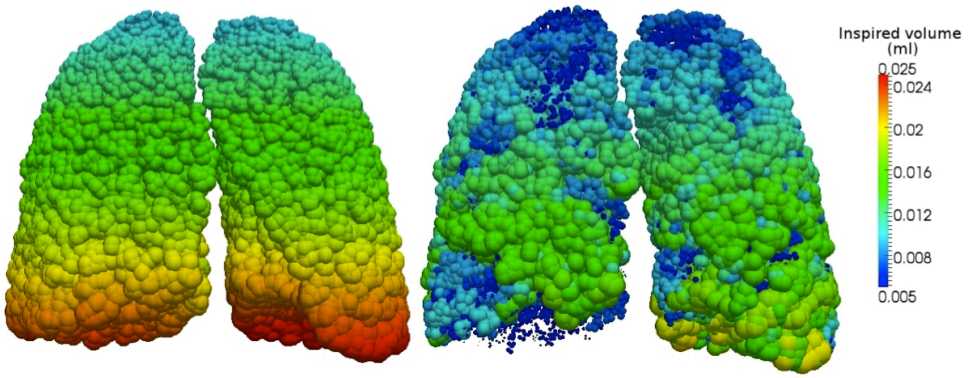

Figure 6.2: Flow rate results for a computational airflow simulation with and without constriction in the airways. The model is derived from airways and lobes segmented from CT data, with the airways extended down to the acinar level by an artificial volume-filling tree-growing algorithm. The left image has been generated by simulating airflow with unrestricted airways. For the right image, the radii of some of the small airways have been adjusted to simulate the effects of construction. The acinar units at the endpoints of the conducting airway tree are represented by spheres for visualisation purposes. Colours represent the volume of inspired air reaching each acinar unit during the simulation. (Courtesy of Minsuok Kim.)

7 Discussion and conclusions

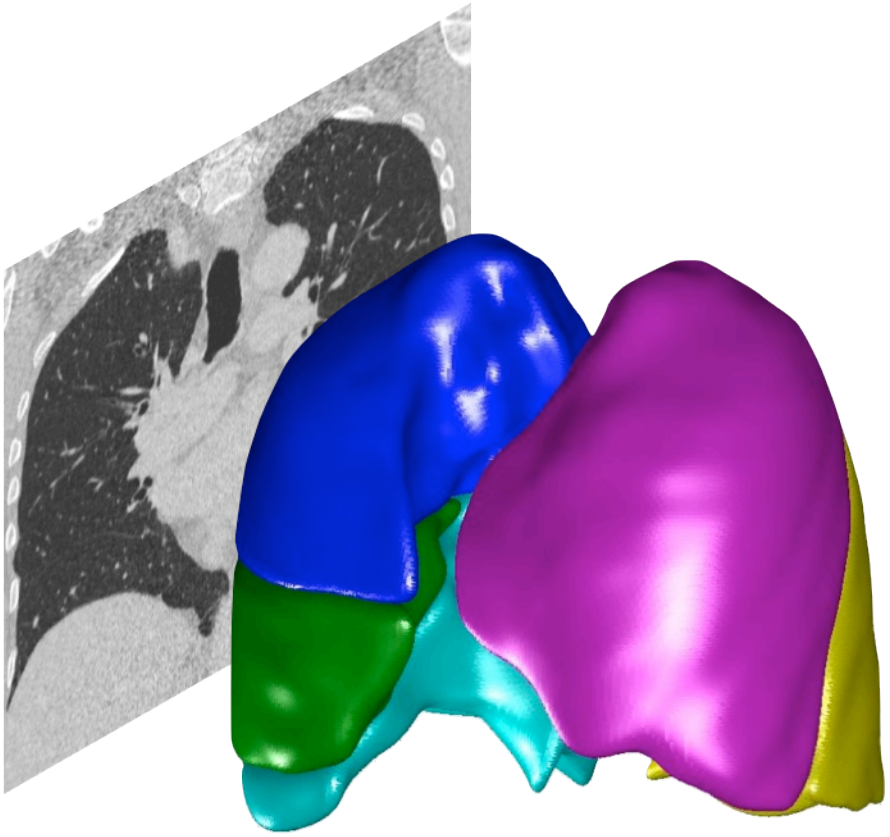

Chapter 7: Discussion and conclusions



In this thesis we have developed an automated pipeline for deriving regional measurements from hyperpolarised xenon MRI images. The development of this pipeline is the first stage in our plan to analyse data from a clinical trial currently underway at the Oxford University Hospitals. As more imaging data becomes available, we will further validate our pipeline and extend our regional analysis to derive ventilation and ADC measurements at a lobar and segmental level. The eventual aim of this work is to develop new measures of lung function that could offer better predictions of patient outcome than current pulmonary function tests. Ultimately, this could lead to improved diagnosis and treatment planning.

Our pipeline makes use of the multi-modal data available from the clinical trial, combining the structural information present in CT and proton MRI data with the functional measurements from the hyperpolarised xenon images. The core elements of this thesis are 1) the development of automatic segmentation algorithms for extracting the lungs, airways and pulmonary lobes; and 2) the development of a software pipeline for automatically performing regional analysis across multiple modalities.

In chapter 3 we developed algorithms for segmenting the lung boundaries and airways. These segmentations were required for the lobe segmentation method

in chapter 4 and the image registration method in chapter 5. In section 3.2 we developed a CT lung segmentation method based on a region-growing technique which has been widely used and well validated by previous authors. In section 3.3 we developed a novel method for automatically segmenting lung boundaries from thick coronal slice proton MRI data. Our method uses a 3D region-growing method with an adaptive threshold for initialisation, followed by contour refinement using 2D level sets. Very little work has been published in automatic MRI lung segmentation, and our validation metrics compared well with previous works. Our method could provide more accurate results with thinner slice MRI data by extending the level set refinement into 3D.

In section 3.4 we developed a method for segmenting airways from CT data. Airway segmentation is a required part of our lobe segmentation algorithm in chapter 4 and will be used in our future computer simulation work described in section 6.2. Our method uses a novel wavefront approach, which allows explosion control to be applied at a bronchial level and eliminates the need for pruning of the airway tree in post-processing steps. We validated the algorithm by examining the number of branches segmented across a number of datasets and by manual verification of the first four segmented generations. We computed the centreline and radius for each bronchus using standard methods, and we developed a novel algorithm for labelling the airways by the lobe they serve, which is robust to common variability in bronchial anatomy between subjects.

In chapter 4 we reviewed methods for automatically segmenting the pulmonary lobes from CT data and developed our own novel algorithm. The generated lobe masks form the basis of our regional analysis of chapter 5. Our method uses a novel fissureness filter to fit multilevel B-spline curves to the pulmonary fissures, guided by anatomical cues from the lungs, airways and pulmonary

vessels. We validated the results using fissure markers placed by a pulmonary clinician.

In chapter 5 we described several regional analysis measures: ventilated volume analysis; ADC measurements; and CT density-based emphysema analysis. We applied these measures to the first four datasets available from the clinical trial, comprising CT, proton MRI and xenon MRI data. For comparison, we performed the same analysis for a COPD patient from a previous study with similar imaging protocols. We developed a simple multi-modal registration method to demonstrate the computation of lobar measurements through the application of the CT lobar segmentation methods developed in chapter 4. In future work we aim to extend this methodology to pulmonary segment analysis We propose additional validation using the clinical trial data, and investigating the use of improved registration methods to achieve increased accuracy from the regional analysis.

Finally, in chapter 6 we described two proposals for future work. Firstly, we will apply our methods to the clinical trial data as it becomes available and extend the range of measurements we obtain. Secondly, we will link our methods to existing airflow simulation work being developed in our group and use this to investigate the mechanisms of lung disease which act on scales not visible through clinical imaging.

Bibliography

1. Chan, M.: The world health report 2008: primary health care now more than ever. World Health Organisation (2008).

2. British Thoracic Society: The burden of lung disease. Second edition. British Thoracic Society (2006).

3. Pinto-Plata, V.M., Cote, C., Cabral, H., Taylor, J., Celli, B.R.: The 6-min walk distance: change over time and value as a predictor of survival in severe COPD. European Respiratory Journal 23(1), 28-33 (2004).

4. Wilt, T.J., Niewoehner, D., Kim, C.-B., Kane, R.L., Linabery, A., Tacklind, J., MacDonald, R., Rutks, I.: Use of spirometry for case finding, diagnosis, and management of chronic obstructive pulmonary disease (COPD). Agency for Healthcare Research and Quality. Evidence Report/Technology Assessment(121) (2005).

5. Fain, S.B., Korosec, F.R., Holmes, J.H., O'Halloran, R.L., Sorkness, R.L., Grist, T.M.: Functional lung imaging using hyperpolarized gas MRI. Journal of Magnetic Resonance Imaging 25(5), 910-923 (2007).

6. Albert, M.S., Cates, G.D., Driehuys, B., Happer, W., Saam, B., Springer, C.S., Wishnia, A.: Biological magnetic resonance imaging using laser-polarized 129Xe. Nature 370(6486), 199-201 (1994).

7. Coxson, H.O.: Computed tomography and monitoring of emphysema. European Respiratory Journal 29(6), 1075-7 (2007).

8. Doel, T., Matin, T.N., Gleeson, F.V., Gavaghan, D.J., Grau, V.: Pulmonary lobe segmentation from CT images using fissureness, airways, vessels and multilevel B-splines. IEEE International Symposium on Biomedical Imaging, 1491-4 (2012).

9. Plotkowiak, M.: Hyperpolarised gas MRI as a promising technique for regional chronic lung diseases assessment. D.Phil. Thesis, University of Oxford (2010).

10. Patz, S., Hersman, F.W., Muradian, I., Hrovat, M.I., Ruset, I.C., Ketel, S., Jacobson, F., Topulos, G.P., Hatabu, H., Butler, J.P.: Hyperpolarized 129Xe MRI: a viable functional lung imaging modality? European Journal of Radiology 64(3), 335-44 (2007).

11. Brown, M.S., Kim, H.J., Abtin, F., Da Costa, I., Pais, R., Ahmad, S., Angel, E., Ni, C., Kleerup, E.C., Gjertson, D.W., McNitt-Gray, M.F., Goldin, J.G.: Reproducibility of lung and lobar volume measurements using computed tomography. Academic Radiology 17(3), 316-22 (2010).

12. Shea, D., Morgan, D.: The helium-3 shortage: supply, demand, and options for Congress. Congressional Research Service (2010).

13. Driehuys, B., Martinez-Jimenez, S., Cleveland, Z.I., Metz, G.M., Beaver, D.M., Nouls, J.C., Kaushik, S.S., Firszt, R., Willis, C., Kelly, K.T., Wolber, J., Kraft, M., McAdams, H.P.: Chronic obstructive pulmonary disease: safety and tolerability of hyperpolarized 129Xe MR imaging in healthy volunteers and patients. Radiology 262(1), 279-89 (2012).

14. West, J.B.: Respiratory physiology: the essentials. Seventh edition. Lippincott Williams & Wilkins (2005).

15. Cronin, P., Gross, B.H., Kelly, A.M., Patel, S., Kazerooni, E.A., Carlos, R.C.: Normal and accessory fissures of the lung: evaluation with contiguous volumetric thin-section multidetector CT. European Journal of Radiology 75(2), e1-8 (2010).

16. Prange, H.: Laplace's law and the alveolus: A misconception of anatomy and a misapplication of physics. Advances in Physiology Education 27(1), 34-40 (2003).

17. Hlastala, M., Berger, A.: Physiology of respiration. Second edition. Oxford University Press (2001).

18. Heil, M., Hazel, A.L., Smith, J.A.: The mechanics of airway closure. Respiratory Physiology & Neurobiology 163(1-3), 214-21 (2008).

19. Siafakas, N.M., Vermeire, P., Pride, N.B., Paoletti, P., Gibson, J., Howard, P., Yernault, J.C., Decramer, M., Higenbottam, T., Postma, D.S.: Optimal assessment and management of chronic obstructive pulmonary disease (COPD). European Respiratory Journal 8(8), 1398-420 (1995).

20. Stolk, J., Putter, H., Bakker, E.M., Shaker, S.B., Parr, D.G., Piitulainen, E., Russi, E.W., Grebski, E., Dirksen, A., Stockley, R.A., Reiber, J.H.C., Stoel, B.C.: Progression parameters for emphysema: a clinical investigation. Respiratory Medicine 101(9), 1924-30 (2007).

21. Sverzellati, N., Molinari, F., Pirronti, T., Bonomo, L., Spagnolo, P., Zompatori, M.: New insights on COPD imaging via CT and MRI. International Journal of Chronic Obstructive Pulmonary Disease 2(3), 301-12 (2007).

22. Hoffman, E.A., Simon, B.A., McLennan, G.: State of the art. A structural and functional assessment of the lung via multidetector-row computed tomography: phenotyping chronic obstructive pulmonary disease. Proceedings of the American Thoracic Society 3(6), 519-32 (2006).

23. Madani, A., Van Muylem, A., de Maertelaer, V., Zanen, J., Gevenois, P.A.: Pulmonary emphysema: size distribution of emphysematous spaces on multidetector CT images – comparison with macroscopic and microscopic morphometry. Radiology 248(3), 1036-1041 (2008).

24. Kinsella, M., Müller, N.L., Abboud, R.T., Morrison, N.J., DyBuncio, A.: Quantitation of emphysema by computed tomography using a "density mask" program and correlation with pulmonary function tests. Chest 97(2), 315-21 (1990).

25. Müller, N.L., Staples, C.A., Miller, R.R., Abboud, R.T.: "Density mask". An objective method to quantitate emphysema using computed tomography. Chest 94(4), 782-7 (1988).

26. Dirksen, A.: Monitoring the progress of emphysema by repeat computed tomography scans with focus on noise reduction. Proceedings of the American Thoracic Society 5(9), 925-8 (2008).

27. Litmanovich, D., Boiselle, P.M., Bankier, A.A.: CT of pulmonary emphysema-current status, challenges, and future directions. European Radiology 19(3), 537-51 (2009).

28. Coxson, H.O., Mayo, J., Lam, S., Santyr, G., Parraga, G., Sin, D.D.: New and current clinical imaging techniques to study chronic obstructive pulmonary disease. American Journal of Respiratory and Critical Care Medicine 180(7), 588-97 (2009).

29. Revel, M.-P., Faivre, J.-B., Remy-Jardin, M., Deken, V., Duhamel, A., Marquette, C.-H., Tacelli, N., Bakai, A.-M., Remy, J.: Automated lobar quantification of emphysema in patients with severe COPD. European Radiology 18(12), 2723-30 (2008).

30. Uppaluri, R., Mitsa, T., Sonka, M., Eric A. Hoffman, E.A., Mclennan, G.: Quantification of pulmonary emphysema from lung computed tomography images. American Journal of Respiratory and Critical Care Medicine 156(1), 248-54 (1997).

31. Chabat, F., Yang, G.-Z., Hansell, D.M.: Obstructive lung diseases: texture classification for differentiation at CT. Radiology 228(3), 871-7 (2003).

32. Xu, Y., Sonka, M., McLennan, G., Guo, J., Hoffman, E.A.: MDCT-based 3-D texture classification of emphysema and early smoking related lung pathologies. IEEE Transactions on Medical Imaging 25(4), 464-475 (2006).

33. Charemza, M., Thönnes, E., Bhalerao, A., Parr, D.: Integral geometry descriptors for characterizing emphysema and lung fibrosis in HRCT images. First International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 1-10 (2008).

34. Boehm, H.F., Fink, C., Attenberger, U., Becker, C., Behr, J., Reiser, M.: Automated classification of normal and pathologic pulmonary tissue by topological texture features extracted from multi-detector CT in 3D. European Radiology 18(12), 2745-55 (2008).

35. Hoffman, E.A., Reinhardt, J.M., Sonka, M., Simon, B.A., Guo, J., Saba, O.I., Chon, D., Samrah, S., Shikata, H., Tschirren, J., Palagyi, K., Beck, K.C., McLennan, G.: Characterization of the interstitial lung diseases via density-based and texture-based analysis of computed tomography images of lung structure and function. Academic Radiology 10(10), 1104-1118 (2003).

36. Zavaletta, V.A., Bartholmai, B.J., Robb, R.A.: High resolution multidetector CT-aided tissue analysis and quantification of lung fibrosis. Academic Radiology 14(7), 772-87 (2007).

37. Hogg, J.C.: Pathophysiology of airflow limitation in chronic obstructive pulmonary disease. Lancet 364(9435), 709-21 (2004).

38. Nakano, Y., Wong, J.C., de Jong, P.A., Buzatu, L., Nagao, T., Coxson, H.O., Elliott, W.M., Hogg, J.C., Paré, P.D.: The prediction of small airway dimensions using computed tomography. American Journal of Respiratory and Critical Care Medicine 171(2), 142-6 (2005).

39. Hasegawa, M., Nasuhara, Y., Onodera, Y., Makita, H., Nagai, K., Fuke, S., Ito, Y., Betsuyaku, T., Nishimura, M.: Airflow limitation and airway dimensions in chronic obstructive pulmonary disease. American Journal of Respiratory and Critical Care Medicine 173(12), 1309-15 (2006).

40. Diaz, A.A., Valim, C., Yamashiro, T., Estépar, R.S.J., Ross, J.C., Matsuoka, S., Bartholmai, B.J., Hatabu, H., Silverman, E.K., Washko, G.R.: Airway count and emphysema assessed by chest CT imaging predicts clinical outcome in smokers. Chest 138(4), 880-7 (2010).

41. Aziz, Z.A., Wells, A.U., Desai, S.R., Ellis, S.M., Walker, A.E., MacDonald, S., Hansell, D.M.: Functional impairment in emphysema: contribution of airway abnormalities and distribution of parenchymal disease. American Journal of Roentgenology 185(6), 1509-15 (2005).

42. Washko, G.R., Criner, G.J., Mohsenifar, Z., Sciurba, F.C., Sharafkhaneh, A., Make, B.J., Hoffman, E.A., Reilly, J.J.: Computed tomographic-based quantification of emphysema and correlation to pulmonary function and mechanics. COPD 5(3), 177-86 (2008).

43. van Beek, E.J.R., Hoffman, E.A.: Functional imaging: CT and MRI. Clinics in Chest Medicine 29(1), 195-216, vii (2008).

44. Stavngaard, T., Søgaard, L.V., Mortensen, J., Hanson, L.G., Schmiedeskamp, J., Berthelsen, A.K., Dirksen, A.: Hyperpolarized 3He MRI and 81mKr SPECT in chronic obstructive pulmonary disease. European Journal of Nuclear Medicine and Molecular Imaging 32(4), 448-57 (2005).

45. Walker, T., Happer, W.: Spin-exchange optical pumping of noble-gas nuclei. Reviews of Modern Physics 69(2), 629-642 (1997).

46. Altes, T.A., Powers, P.L., Knight-Scott, J., Rakes, G., Platts-Mills, T.A.E., de Lange, E.E., Alford, B.A., Mugler, J.P., Brookeman, J.R.: Hyperpolarized 3He MR lung ventilation imaging in asthmatics: preliminary findings. Journal of Magnetic Resonance Imaging 13(3), 378-84 (2001).

47. Fain, S.B., Schiebler, M.L., McCormack, D.G., Parraga, G.: Imaging of lung function using hyperpolarized helium-3 magnetic resonance imaging: review of current and emerging translational methods and applications. Journal of Magnetic Resonance Imaging 32(6), 1398-408 (2010).

48. Schreiber, W.G., Morbach, A.E., Stavngaard, T., Gast, K.K., Herweling, A., Søgaard, L.V., Windirsch, M., Schmiedeskamp, J., Heussel, C.P., Kauczor, H.-U.: Assessment of lung microstructure with magnetic resonance imaging of hyperpolarized helium-3. Respiratory Physiology and Neurobiology 148(1-2), 23-42 (2005).

49. Saam, B.T., Yablonskiy, D.A., Kodibagkar, V.D., Leawoods, J.C., Gierada, D.S., Cooper, J.D., Lefrak, S.S., Conradi, M.S.: MR imaging of diffusion of 3He gas in healthy and diseased lungs. Magnetic Resonance in Medicine 44(2), 174-9 (2000).

50. Wang, C., Altes, T.A., Mugler, J.P., Miller, G.W., Ruppert, K., Mata, J.F., Cates, G.D., Borish, L., de Lange, E.E.: Assessment of the lung microstructure in patients with asthma using hyperpolarized 3He diffusion MRI at two time scales: comparison with healthy subjects and patients with COPD. Journal of Magnetic Resonance Imaging 28(1), 80-8 (2008).

51. Mathew, L., Evans, A., Ouriadov, A.V., Etemad-Rezai, R., Fogel, R., Santyr, G., McCormack, D.G., Parraga, G.: Hyperpolarized 3He magnetic resonance imaging of chronic obstructive pulmonary disease: reproducibility at 3.0 Tesla. Academic Radiology 15(10), 1298-311 (2008).

52. Woods, J.C., Choong, C.K., Yablonskiy, D.A., Bentley, J., Wong, J., Pierce, J.A., Cooper, J.D., Macklem, P.T., Conradi, M.S., Hogg, J.C.: Hyperpolarized 3He diffusion MRI and histology in pulmonary emphysema. Magnetic Resonance in Medicine 56(6), 1293-300 (2006).

53. Salerno, M., de Lange, E.E., Altes, T.A., Truwit, J.D., Brookeman, J.R., Mugler, J.P.: Emphysema: hyperpolarized helium 3 diffusion MR imaging of the lungs compared with spirometric indexes-initial experience. Radiology 222(1), 252-60 (2002).

54. Fain, S.B., Gonzalez-Fernandez, G., Peterson, E.T., Evans, M.D., Sorkness, R.L., Jarjour, N.N., Busse, W.W., Kuhlman, J.E.: Evaluation of structure-function relationships in asthma using multidetector CT and hyperpolarized He-3 MRI. Academic Radiology 15(6), 753-62 (2008).

55. Sukstanskii, A.L., Yablonskiy, D.A.: In vivo lung morphometry with hyperpolarized 3He diffusion MRI: theoretical background. Journal of Magnetic Resonance 190(2), 200-10 (2008).

56. Sukstanskii, A.L., Yablonskiy, D.A.: Lung morphometry with hyperpolarized 129Xe: theoretical background. Magnetic Resonance in Medicine 67(3), 856-66 (2012).

57. Plotkowiak, M., Burrowes, K.S., Wolber, J., Buckley, C., Davies, R., Gleeson, F.V., Gavaghan, D.J., Grau, V.: Relationship between structural changes and hyperpolarized gas magnetic resonance imaging in chronic obstructive pulmonary disease using computational simulations with realistic alveolar geometry. Philosophical Transactions Of The Royal Society A—Mathematical Physical And Engineering Sciences 367(1896), 2347-69 (2009).

58. Hersman, F.W., Ruset, I.C., Ketel, S., Muradian, I., Covrig, S.D., Distelbrink, J., Porter, W., Watt, D., Ketel, J., Brackett, J., Hope, A., Patz, S.: Large production system for hyperpolarized 129Xe for human lung imaging studies. Academic Radiology 15(6), 683-92 (2008).

59. Shukla, Y., Wheatley, A., Kirby, M., Svenningsen, S., Farag, A., Santyr, G.E., Paterson, N.A.M., McCormack, D.G., Parraga, G.: Hyperpolarized 129Xe magnetic resonance imaging: tolerability in healthy volunteers and subjects with pulmonary disease. Academic Radiology 19(8), 941-51 (2012).

60. Dregely, I.M., Mugler, J.P., Ruset, I.C., Altes, T.A., Mata, J.F., Miller, G.W., Ketel, J., Ketel, S., Distelbrink, J., Hersman, F.W., Ruppert, K.: Hyperpolarized xenon-129 gas-exchange imaging of lung microstructure: first case studies in subjects with obstructive lung disease. Journal of Magnetic Resonance Imaging 33(5), 1052-62 (2011).

61. Mugler, J.P., Altes, T.A., Ruset, I.C., Dregely, I.M., Mata, J.F., Miller, G.W., Ketel, S., Ketel, J., Hersman, F.W., Ruppert, K.: Simultaneous magnetic resonance imaging of ventilation distribution and gas uptake in the human lung using hyperpolarized xenon-129. Proceedings of the National Academy of Sciences 107(50), 21707-12 (2010).

62. Kirby, M., Svenningsen, S., Owrangi, A.M., Wheatley, A., Farag, A., Ouriadov, A.V., Santyr, G.E., Etemad-Rezai, R., Coxson, H.O., McCormack, D.G., Parraga,

G.: Hyperpolarized 3He and 129Xe MR imaging in healthy volunteers and patients with chronic obstructive pulmonary disease. Radiology 265(2), 600-610 (2012).

63. Cleveland, Z.I., Cofer, G.P., Metz, G.M., Beaver, D.M., Nouls, J.C., Kaushik, S.S., Kraft, M., Wolber, J., Kelly, K.T., McAdams, H.P., Driehuys, B.: Hyperpolarized Xe MR imaging of alveolar gas uptake in humans. Plos One 5(8), e12192 (2010).

64. Lo, P., van Ginneken, B., Reinhardt, J.M., de Bruijne, M.: Extraction of airways from CT (EXACT'09). Second International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 175-190 (2009).

65. Sluimer, I., Schilham, A.M.R., Prokop, M., van Ginneken, B.: Computer analysis of computed tomography scans of the lung: a survey. IEEE Transactions on Medical Imaging 25(4), 385-405 (2006).

66. Hu, S., Hoffman, E.A., Reinhardt, J.M.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. IEEE Transactions on Medical Imaging 20(6), 490-8 (2001).

67. Brown, M.S., McNitt-Gray, M.F., Mankovich, N.J., Goldin, J.G., Hiller, J., Wilson, L.S., Aberle, D.R.: Method for segmenting chest CT image data using an anatomical model: preliminary results. IEEE Transactions on Medical Imaging 16(6), 828-39 (1997).

68. Sun, X., Zhang, H., Duan, H.: 3D computerized segmentation of lung volume with computed tomography. Academic Radiology 13(6), 670-7 (2006).

69. Armato, S.G., Sensakovic, W.F.: Automated lung segmentation for thoracic CT impact on computer-aided diagnosis. Academic Radiology 11(9), 1011-21 (2004).

70. Weinheimer, O., Achenbach, T., Heussel, C.P., Düber, C.: Automatic lung segmentation in MDCT images. Fourth International Workshop on Pulmonary Image Analaysis, Medical Image Computing and Computer-Assisted Intervention, 241-256 (2011).

71. Kuhnigk, J.-M., Dicken, V., Zidowitz, S., Bornemann, L., Kuemmerlen, B., Krass, S., Peitgen, H.-O., Yuval, S., Jend, H.-H., Rau, W.S., Achenbach, T.: Informatics in radiology (infoRAD): new tools for computer assistance in thoracic CT. Part 1. Functional analysis of lungs, lung lobes, and bronchopulmonary segments. Radiographics 25(2), 525-36 (2005).

72. Lassen, B., Kuhnigk, J.-M., Schmidt, M., Krass, S., Peitgen, H.-O.: Lung and lung lobe segmentation methods at Fraunhofer MEVIS. Lobe and Lung Analysis 2011. Fourth International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 185-199 (2011).

73. van Rikxoort, E.M., van Ginneken, B.: Automatic segmentation of the lungs and lobes from thoracic CT scans. Fourth International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 261-268 (2011).

74. Lo, P., Goldin, J.G., Oria, D., Banola, A., Brown, M.S.: Historic automated lung segmentation method: performance on LOLA11 data set. Fourth International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 257-260 (2011).

75. Sun, S., Bauer, C., Beichel, R.: Robust active shape model based lung segmentation in CT scans. Fourth International Workshop on Pulmonary Image Analaysis, Medical Image Computing and Computer-Assisted Intervention, 213-224 (2011).

76. Montillo, A.: Context selective decision forest with applications to lung segmentation in CT. Fourth International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 201-212 (2011).

77. Sluimer, I.C., Prokop, M., van Ginneken, B.: Toward automated segmentation of the pathological lung in CT. IEEE Transactions on Medical Imaging 24(8), 1025-38 (2005).

78. Pu, J., Roos, J., Yi, C.A., Napel, S., Rubin, G.D., Paik, D.S.: Adaptive border marching algorithm: automatic lung segmentation on chest CT images. Computerized Medical Imaging and Graphics 32(6), 452-62 (2008).

79. Prasad, M.N., Brown, M.S., Ahmad, S., Abtin, F., Allen, J., da Costa, I., Kim, H.J., McNitt-Gray, M.F., Goldin, J.G.: Automatic segmentation of lung parenchyma in the presence of diseases based on curvature of ribs. Academic Radiology 15(9), 1173-80 (2008).

80. El-Ba, A., Gimel'farb, G., Falk, R., Holland, T., Shaffer, T.: A new stochastic framework for accurate lung segmentation. Medical Image Computing and Computer-Assisted Intervention 5241(Pt 1), 322-30 (2008).

81. van Rikxoort, E.M., de Hoop, B., Viergever, M.A., Prokop, M., van Ginneken, B.: Automatic lung segmentation from thoracic computed tomography scans using a hybrid approach with error detection. Medical Physics 36(7), 2934-47 (2009).

82. Vincent, L., Soille, P.: Watersheds in digital spaces: an efficient algorithm based on immersion simulations. IEEE Transactions on Pattern Analysis and Machine Intelligence 13(6), 583-598 (1991).

83. Middleton, I., Damper, R.I.: Segmentation of magnetic resonance images using a combination of neural networks and active contour models. Medical Engineering and Physics 26(1), 71-86 (2004).

84. Ray, N., Acton, S.T., Altes, T.A., de Lange, E.E., Brookeman, J.R.: Merging parametric active contours within homogeneous image regions for MRI-based lung segmentation. IEEE Transactions on Medical Imaging 22(2), 189-99 (2003).

85. Osareh, A., Shadgar, B.: A segmentation method of lung cavities using region aided geometric snakes. Journal of Medical Systems 34(4), 419-33 (2010).

86. Kirby, M., Mathew, L., Wheatley, A., Santyr, G.E., McCormack, D.G., Parraga, G.: Chronic obstructive pulmonary disease: longitudinal hyperpolarized 3He MR imaging. Radiology 256(1), 280-9 (2010).

87. Tzeng, Y.-S., Lutchen, K., Albert, M.S.: The difference in ventilation heterogeneity between asthmatic and healthy subjects quantified using hyperpolarized 3He MRI. Journal of Applied Physiology 106(3), 813-22 (2009).

88. Woodhouse, N., Wild, J.M., Paley, M.N.J., Fichele, S., Said, Z., Swift, A.J., van Beek, E.J.R.: Combined helium-3/proton magnetic resonance imaging measurement of ventilated lung volumes in smokers compared to never-smokers. Journal of Magnetic Resonance Imaging 21(4), 365-9 (2005).

89. Tustison, N.J., Avants, B.B., Flors, L., Altes, T.A., de Lange, E.E., Mugler, J.P., Gee, J.C.: Ventilation-based segmentation of the lungs using hyperpolarized 3He MRI. Journal of Magnetic Resonance Imaging 34(4), 831-41 (2011).

90. Ireland, R.H., Woodhouse, N., Hoggard, N., Swinscoe, J.A., Foran, B.H., Hatton, M.Q.F., Wild, J.M.: An image acquisition and registration strategy for the fusion of hyperpolarized helium-3 MRI and x-ray CT images of the lung. Physics in Medicine and Biology 53(21), 6055-63 (2008).

91. Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. Journal of Computational Physics 79(1), 12-49 (1988).

92. Angelini, E.D., Jin, Y., Laine, A.: State of the art of level set methods in segmentation and registration of medical imaging modalities. Handbook of Biomedical Image Analysis: Volume 3: Registration Models, 47-101 (2005).

93. Grady, L.: Random walks for image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 28(11), 1768-83 (2006).

94. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: active contour models. International Journal of Computer Vision (1988).

95. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. NeuroImage 31(3), 1116-28 (2006).

96. Dice, L.: Measures of the amount of ecologic association between species. Ecology 26(3), 297–302 (1945).

97. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, M., Schnabel, J.A.: MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration. Medical Image Analysis 16(7), 1423-1435 (2012).

98. González Ballester, M.A., Zisserman, A.P., Brady, M.: Estimation of the partial volume effect in MRI. Medical Image Analysis 6(4), 389-405 (2002).

99. Bello, F., Colchester, A.C., Röll, S.A.: A geometry-and intensity-based partial volume correction for MRI volumetric studies. Computerized Medical Imaging and Graphics 22(2), 123-32 (1998).

100. Wood, S.A., Zerhouni, E.A., Hoford, J.D., Hoffman, E.A., Mitzner, W.: Measurement of three-dimensional lung tree structures by using computed tomography. Journal of Applied Physiology 79(5), 1687-97 (1995).

101. van Ginneken, B., Baggerman, W., van Rikxoort, E.M.: Robust segmentation and anatomical labeling of the airway tree from thoracic CT scans. Medical Image Computing and Computer-Assisted Intervention 5241(1), 219-26 (2008).

102. Tschirren, J., Yavarna, T., Reinhardt, J.M.: Airway segmentation framework for clinical environments. Second International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 227-238 (2009).

103. Fabijańska, A.: Two-pass region growing algorithm for segmenting airway tree from MDCT chest scans. Computerized Medical Imaging and Graphics 33(7), 537-546 (2009).

104. Mori, K., Hasegawa, J., Toriwaki, J., Anno, H., Katada, K.: Recognition of bronchus in three-dimensional X-ray CT images with applications to virtualized bronchoscopy system. IEEE International Conference on Pattern Recognition 3, 528-532 (1996).

105. Pu, J., Fuhrman, C.R., Good, W.F., Sciurba, F.C., Gur, D.: A differential geometric approach to automated segmentation of human airway tree. IEEE Transactions on Medical Imaging 30(2), 266-278 (2011).

106. Bauer, C., Bischof, H., Beichel, R.: Segmentation of airways based on gradient vector flow. Second International Workshop on Pulmonary Image Analysis, Medical Image Computing and Computer-Assisted Intervention, 191-201 (2009).

107. Tschirren, J., Eric A. Hoffman, E.A., McLennan, G., Sonka, M.: Segmentation and quantitative analysis of intrathoracic airway trees from computed tomography images. Proceedings of the American Thoracic Society 2(6), 484-7 (2005).

108. Aykac, D., Hoffman, E.A., McLennan, G., Reinhardt, J.M.: Segmentation and analysis of the human airway tree from three-dimensional X-ray CT images. IEEE Transactions on Medical Imaging 22(8), 940-950 (2003).

109. Lo, P., Sporring, J., Ashraf, H., Pedersen, J.J.H., de Bruijne, M.: Vessel-guided airway tree segmentation: A voxel classification approach. Medical Image Analysis 14(4), 527-38 (2010).

110. Graham, M.W., Gibbs, J.D., Cornish, D.C., Higgins, W.E.: Robust 3-D airway tree segmentation for image-guided peripheral bronchoscopy. IEEE Transactions on Medical Imaging 29(4), 982-997 (2010).

111. Wang, X., Fang, C., Xia, Y., Feng, D.: Airway segmentation for low-contrast CT images from combined PET/CT scanners based on airway modelling and seed prediction. Biomedical Signal Processing and Control 6(1), 48-56 (2011).

112. Mayer, D., Bartz, D., Fischer, J., Ley, S., del Río, A., Thust, S., Kauczor, H.-U., Heussel, C.P.: Hybrid segmentation and virtual bronchoscopy based on CT images. Academic Radiology 11(5), 551-65 (2004).

113. Kiraly, A.P., Higgins, W.E., McLennan, G., Hoffman, E.A., Reinhardt, J.M.: Three-dimensional human airway segmentation methods for clinical virtual bronchoscopy. Academic Radiology 9(10), 1153-1168 (2002).

114. Schlathölter, T., Lorenz, C., Carlsen, I., Renisch, S., Deschamps, T.: Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy. Proceedings of the Society of Photo-Optical Instrumentation Engineers 4684, 103-113 (2002).

115. Palágyi, K., Tschirren, J., Hoffman, E.A., Sonka, M.: Quantitative analysis of pulmonary airway tree structures. Computers in Biology and Medicine 36, 974-996 (2006).

116. Malandain, G., Bertrand, G.: Fast characterization of 3D simple points. IEEE International Conference on Pattern Recognition 3, 232-235 (1992).

117. Estépar, R.S.J., Washko, G.R., Silverman, E.K., Reilly, J.J., Kikinis, R., Westin, C.-F.: Accurate airway wall estimation using phase congruency. Medical Image Computing and Computer-Assisted Intervention 4191(2), 125-34 (2006).

118. Reinhardt, J.M., D'Souza, N.D., Hoffman, E.A.: Accurate measurement of intrathoracic airways. IEEE Transactions on Medical Imaging 16(6), 820-7 (1997).

119. Tschirren, J., McLennan, G., Palágyi, K., Hoffman, E.A., Sonka, M.: Matching and anatomical labeling of human airway tree. IEEE Transactions on Medical Imaging 24(12), 1540-1547 (2005).

120. Mori, K., Hasegawa, J., Suenaga, Y., Toriwaki, J.: Automated anatomical labeling of the bronchial branch and its application to the virtual bronchoscopy system.

IEEE Transactions on Medical Imaging 19(2), 103-114 (2000).

121. Jeffery, P.K.: Structural and inflammatory changes in COPD: a comparison with asthma. Thorax 53(2), 129-36 (1998).

122. Leung, A.N.: Pulmonary tuberculosis: the essentials. Radiology 210(2), 307-22 (1999).

123. Rees, D., Murray, J.: Silica, silicosis and tuberculosis. International Journal of Tuberculosis and Lung Disease 11(5), 474-84 (2007).

124. American Thoracic Society, European Respiratory Society: International multidisciplinary consensus classification of the idiopathic interstitial pneumonias. American Journal of Respiratory and Critical Care Medicine 165(2), 277-304 (2002).

125. Fishman, A., Martinez, F.J., Naunheim, K., Piantadosi, S., Wise, R., Ries, A., Weinmann, G., Wood, D.E.: A randomized trial comparing lung-volume-reduction surgery with medical therapy for severe emphysema. New England Journal of Medicine 348(21), 2059-73 (2003).

126. Meenakshi, S., Manjunath, K., Balasubramanyam, V.: Morphological variations of the lung fissures and lobes. Indian Journal of Chest Diseases and Allied Sciences 46(3), 179-182 (2004).

127. van Rikxoort, E.M., Prokop, M., de Hoop, B., Viergever, M.A., Pluim, J.P.W., van Ginneken, B.: Automatic segmentation of pulmonary lobes robust against incomplete fissures. IEEE Transactions on Medical Imaging 29(6), 1286-1296 (2010).

128. Zhang, L., Hoffman, E.A., Reinhardt, J.M.: Atlas-driven lung lobe segmentation in volumetric X-ray CT images. IEEE Transactions on Medical Imaging 25(1), 1-16 (2006).

129. Ross, J.C., Estépar, R.S.J., Kindlmann, G.L., Díaz, A., Westin, C.-F., Silverman, E.K., Washko, G.R.: Automatic lung lobe segmentation using particles, thin plate splines, and maximum a posteriori estimation. Medical Image Computing and Computer-Assisted Intervention 6363, 163-171 (2010).

130. Ukil, S., Reinhardt, J.M.: Anatomy-guided lung lobe segmentation in X-ray CT images. IEEE Transactions on Medical Imaging 28(2), 202-14 (2009).

131. Zhou, X., Hayashi, T., Hara, T., Fujita, H., Yokoyama, R., Kiryu, T., Hoshi, H.: Automatic segmentation and recognition of anatomical lung structures from high-resolution chest CT images. Computerized Medical Imaging and Graphics 30, 299-313 (2006).

132. Lassen, B., Kuhnigk, J.-M., Friman, O., Krass, S., Peitgen, H.-O.: Automatic segmentation of lung lobes in CT images based on fissures, vessels, and

bronchi. IEEE International Symposium on Biomedical Imaging, 560-563 (2010).

133. Wei, Q., Hu, Y., Gelfand, G., MacGregor, J.H.: Segmentation of lung lobes in high-resolution isotropic CT images. IEEE Transactions on Biomedical Engineering 56(5), 1383-93 (2009).

134. Pu, J., Zheng, B., Leader, J.K., Fuhrman, C.R., Knollmann, F., Klym, A., Gur, D.: Pulmonary lobe segmentation in CT examinations using implicit surface fitting. IEEE Transactions on Medical Imaging 28(12), 1986-1996 (2009).

135. Lopez, A., Lloret, D., Serrat, J., Villanueva, J.: Multilocal creaseness based on the level-set extrinsic curvature. Computer Vision and Image Understanding 77(2), 111-144 (2000).

136. Shikata, H., McLennan, G., Hoffman, E.A., Sonka, M.: Segmentation of pulmonary vascular trees from thoracic 3D CT images. International Journal of Biomedical Imaging, 636240 (2009).

137. Ohkawa, T., Kobashi, S., Kondo, K., Hata, Y., Nakano, T.: Tubular tissue-based segmentation of lung lobes from chest MDCT images. IEEE EMBS Asian-Pacific Conference on Biomedical Engineering, 126-127 (2003).

138. Kobashi, S., Hata, Y.: Lung lobar segmentation using tubular tissue density from multidetector-row CT images. International Journal of Innovative Computing, Information and Control 6(3), 829-842 (2010).

139. Pu, J., Leader, J.K., Zheng, B., Knollmann, F., Fuhrman, C.R., Sciurba, F.C., Gur, D.: A computational geometry approach to automated pulmonary fissure segmentation in CT examinations. IEEE Transactions on Medical Imaging 28(5), 710-719 (2009).

140. Gu, S., Pu, J., Zheng, Q., Siegfried, J.: Robust pulmonary lobe segmentation against incomplete fissures. Proceedings of the Society of Photo-Optical Instrumentation Engineers 8315, 831535 (2012).

141. Kitasaka, T., Nakada, Y., Mori, K., Suenaga, Y., Mori, M., Takabatake, H., Natori, H.: Recognition of lung lobes and its application to the bronchial structure analysis. IEEE International Conference on Pattern Recognition 3, 288-291 (2006).

142. Blaffert, T., Barschdorf, H., von Berg, J., Dries, S., Franz, A., Klinder, T., Lorenz, C., Renisch, S., Wiemker, R.: Lung lobe modeling and segmentation with individualized surface meshes. Proceedings of the Society of Photo-Optical Instrumentation Engineers 6914, I9141 (2008).

143. van Rikxoort, E.M., de Hoop, B., van de Vorst, S., Prokop, M., van Ginneken, B.: Automatic segmentation of pulmonary segments from volumetric chest CT scans. IEEE Transactions on Medical Imaging 28(4), 621-30 (2009).

144. Busayarat, S., Zrimec, T.: Detection of bronchopulmonary segments on high-resolution CT – preliminary results. IEEE International Symposium on Computer-Based Medical Systems, 199-204 (2007).

145. Saita, S., Kubo, M., Kawata, Y., Niki, N., Ohmatsu, H., Moriyama, N.: An algorithm for the extraction of pulmonary fissures from low-dose multislice CT image. Systems and Computers in Japan 37(9), 63-76 (2005).

146. Wang, J., Betke, M., Ko, J.P.: Pulmonary fissure segmentation on CT. Medical Image Analysis 10(4), 530-47 (2006).

147. Frangi, A., Niessen, W.J., Vincken, K., Viergever, M.A.: Multiscale vessel enhancement filtering. Medical Image Computing and Computer-Assisted Intervention 1496, 130-137 (1998).

148. Kovesi, P.: Image features from phase congruency. Videre 1(3), 1-26 (1999).

149. Obara, B., Fricker, M., Gavaghan, D., Grau, V.: Contrast-independent curvilinear structure detection in biomedical images. IEEE Transactions on Image Processing 21(5), 2572-81 (2012).

150. Lindeberg, T.: Edge detection and ridge detection with automatic scale selection. International Journal of Computer Vision 30(2), 117-154 (1998).

151. Lee, S., Wolberg, G., Shin, S.Y.: Scattered data interpolation with multilevel B-splines. IEEE Transactions on Visualization And Computer Graphics 3(3), 228-244 (1997).

152. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. ACM Computing Surveys 36(2), 81-121 (2004).

153. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. Medical Image Analysis 13(6), 819-845 (2009).

154. Staring, M., Xiao, C., Shamonin, D.P., Stoel, B.C.: Pulmonary vessel segmentation using vessel enhancement filters. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

155. Estépar, R.S.J., Ross, J.C., Washko, G.R., Kindlmann, G.L.: Pulmonary vasculature segmentation based on scale-space particles: BWH-ACIL team. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

156. Öksüz, İ., Ünay, D., Kadipaşaoğlu, K.: Multi-scale Hessian based approach for segmentation of lung vessel tree in 3-D CTA data. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

157. Xue, W., Zhu, X., Liang, J.: An automatic pulmonary vessel segmentation approach. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

158. Smistad, E., Elster, A.C., Lindseth, F.: GPU-based pulmonary vessel segmentation using a tubular spline model. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

159. Foruzan: Vessel12 challenge algorithm description. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

160. Fabijańska, A.: Vessel segmentation in lungs. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

161. Prieto, J.-C., Revol-Muller, C., Odet, C., Orkisz, M., Hernández Hoyos, M.: Description of the algorithm. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

162. Meine, H., Hüllebrand, M., Stöcker, C.: Vessel12 challenge algorithm description. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

163. Fetita, C.I., Brillet, P.-Y., Lefevre, C.: Intrapulmonary vessels segmentation from CT based on topographical connectivity and anatomical knowledge. Vessel Segmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (2012).

164. Kaushik, S.S., Cleveland, Z.I., Cofer, G.P., Metz, G., Beaver, D., Nouls, J.C., Kraft, M., Auffermann, W., Wolber, J., McAdams, H.P., Driehuys, B.: Diffusion-weighted hyperpolarized 129Xe MRI in healthy volunteers and subjects with chronic obstructive pulmonary disease. Magnetic Resonance in Medicine 65(4), 1155-1165 (2011).

165. Woodhouse, N., Wild, J.M., van Beek, E.J.R., Hoggard, N., Barker, N., Taylor, C.J.: Assessment of hyperpolarized 3He lung MRI for regional evaluation of interventional therapy: a pilot study in pediatric cystic fibrosis. Journal of Magnetic Resonance Imaging 30(5), 981-8 (2009).

166. Wild, J.M., Ajraoui, S., Deppe, M.H., Parnell, S.R., Marshall, H., Parra-Robles, J., Ireland, R.H.: Synchronous acquisition of hyperpolarised 3He and 1H MR images of the lungs - maximising mutual anatomical and functional information. NMR in Biomedicine 24(2), 130-4 (2011).

167. Maintz, J.B.A., Viergever, M.A.: A survey of medical image registration. Medical Image Analysis 2(1), 1-36 (1998).

168. Murphy, K., van Ginneken, B., Reinhardt, J.M., Kabus, S., Ding, K., Deng, X., Cao, K., Du, K., Christensen, G.E., Garcia, V., Vercauteren, T., Ayache, N., Commowick, O., Malandain, G., Glocker, B., Paragios, N., Navab, N., Gorbunova,

V., Sporring, J., de Bruijne, M., Han, X., Heinrich, M.P., Schnabel, J.A., Jenkinson, M., Lorenz, C., Modat, M., McClelland, J.R., Ourselin, S., Muenzing, S.E.A., Viergever, M.A., De Nigris, D., Collins, D.L., Arbel, T., Peroni, M., Li, R., Sharp, G.C., Schmidt-Richberg, A., Ehrhardt, J., Werner, R., Smeets, D., Loeckx, D., Song, G., Tustison, N., Avants, B., Gee, J.C., Staring, M., Klein, S., Stoel, B.C., Urschler, M., Werlberger, M., Vandemeulebroucke, J., Rit, S., Sarrut, D., Pluim, J.P.W.: Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. IEEE Transactions on Medical Imaging 30(11), 1901-20 (2011).

169. Pluim, J.P.W., Maintz, J.B.A., Viergever, M.A.: Mutual-information-based registration of medical images: a survey. IEEE Transactions on Medical Imaging 22(8), 986-1004 (2003).

170. Risser, L., Heinrich, M.P., Matin, T., Schnabel, J.A.: Piecewise-diffeomorphic registration of 3D CT/MR pulmonary images with sliding conditions. IEEE International Symposium on Biomedical Imaging, 1351-1354 (2012).

171. Risser, L., Vialard, F.-X., Baluwala, H.Y., Schnabel, J.A.: Piecewise-diffeomorphic image registration: application to the motion estimation between 3D CT lung images with sliding conditions. Medical Image Analysis 17(2), 182-93 (2013).

172. Arsigny, V., Pennec, X., Ayache, N.: Polyrigid and polyaffine transformations: a novel geometrical tool to deal with non-rigid deformations – application to the registration of histological slices. Medical Image Analysis 9(6), 507-23 (2005).

173. Ourselin, S., Roche, A., Subsol, G., Pennec, X., Ayache, N.: Reconstructing a 3D structure from serial histological sections. Image and Vision Computing 19, 25-31 (2000).

174. Cahill, N.D., Noble, J.A., Hawkes, D.J., Ray, L.A.: Fast fluid registration with Dirichlet boundary conditions: a transform-based approach. IEEE International Symposium on Biomedical Imaging, 712-715 (2007).

175. Bro-Nielsen, M., Gramkow, C.: Fast fluid registration of medical images. Lecture Notes in Computer Science 1131, 267-276 (1996).

176. Kim, M., Grau, V., Kay, D., Burrowes, K.S.: Predicting ventilation distribution: a model of coupled acinar dynamics and pressure distribution in the conducting airways. American Journal of Respiratory and Critical Care Medicine 185, A2068 (2012).

177. Swan, A.J., Clark, A.R., Tawhai, M.H.: A computational model of the topographic distribution of ventilation in healthy human lungs. Journal of Theoretical Biology 300, 222-31 (2012).

List of acronyms

2D	2-dimensional
3D	3-dimensional
ADC	Apparent diffusion coefficient (equivalent to EDC)
COPD	Chronic obstructive pulmonary disease
CT	Computed tomography
D_{Lco}	Diffusing capacity for carbon monoxide
DPhil	Doctor of philosophy (PhD)
DW	Diffusion-weighted
EDC	Effective diffusion coefficient (equivalent to ADC)
Empire	Evaluation of Methods for Pulmonary Image Registration 2010
EP	Emphysema percentage
EPSRC	Engineering and Physical Sciences Research Council
Exact	Extraction of Airways from CT 2009
FEV_1	Forced expiratory volume in one second
FEV_1 %	Ratio of FEV_1 to FVC
FEV_1 % predicted	FEV_1 % normalised by a population average
FN	False negative
FP	False positive
FRC	Functional residual capacity
FVC	Forced vital capacity
FWHM	Full-width-at-half-maximum
GE	General Electric
^{3}He	Helium-3

He-MRI	Hyperpolarised helium-3 magnetic resonance imaging
HRCT	High-resolution computed tomography
HU	Hounsfield unit
ICC	Intraclass correlation coefficient
ISBI	International Symposium on Biomedical Imaging
Lola	Lobe and Lung Analysis 2011
LVRS	Lung volume reduction surgery
MDCT	Multi-detector row x-ray computed tomography
Miccai	Medical Image Computing and Computer Assisted Intervention
MLSEC-ST	Multilocal level-set extrinsic curvature with structure tensor
MRI	Magnetic resonance imaging
PD	Percentile density
PDE	Partial differential equation
PFT	Pulmonary function test
PTK	Pulmonary Toolkit
Ransac	Random sample consensus
RMS	Root mean square
ROI	Region of interest
RV	Residual volume
Spect	Single-photon emission computed tomography
SPGR	Spoiled gradient-recalled-echo
STP	Standard temperature and pressure
TN	True negative
TP	True positive
TLC	Total lung capacity
Xe-MRI	Hyperpolarised xenon-129 magnetic resonance imaging
^{129}Xe	Xenon-129

List of figures

Figure 1.1: Flowchart showing the components of this thesis 12

Figure 1.2: Schematic diagram of the imaging modalities available from the clinical trial 16

Figure 2.1: Schematic diagram of the respiratory system 22

Figure 2.2: Idealised diagram of pulmonary acini, showing the small airways, vessels and alveoli 24

Figure 2.3: Scanning electron micrograph of alveoli in the lung 25

Figure 2.4: Model of the bronchial tree, with each airway coloured according to the bronchopulmonary segments it serves 27

Figure 2.5: Computational model of the pulmonary blood vessels of the left lung 29

Figure 2.6: Spirometer and plethysmograph 33

Figure 2.7: Reconstructed coronal slice from a 3D multi-detector CT (MDCT) dataset showing the lungs 34

Figure 2.8: Ventilation and diffusion images from hyperpolarised helium MRI 38

Figure 3.1: Illustration of automatic lung segmentation from CT 44

Figure 3.2: 3D surface renderings showing example results from automatic segmentation and labelling of the lungs 46

Figure 3.3: Schematic of the lung segmentation algorithm 49

Figure 3.4: Example of incomplete lung segmentation due to pathology 54

Figure 3.5: Results of automatic lung segmentation from a CT dataset 55

Figure 3.6: Schematic diagram of the proton MRI lung segmentation algorithm 60

Figure 3.7: Illustration of the MRI lung segmentation initialisation 61

Figure 3.8: Results of automatic lung segmentation from a proton MRI dataset 66

Figure 3.9: Slice by slice comparison of automatic and manual MRI lung segmentations 69

Figure 3.10: Illustration of the trachea detection algorithm 77

Figure 3.11: Illustration of the airway segmentation in progress 81

Figure 3.12: Illustration of segmentation leakage from the airway lumen into the surrounding parenchyma during region-growing 82

Figure 3.13: 3D visualisation of the airways segmented from CT data 84

Figure 3.14: 3D surface rendering of the segmented airway lumen, showing the radial grids used to compute the airway radius 87

Figure 3.15: Stretched-out airway cross-section illustrating computation of the inner airway wall radius using the full-width-at-half-maximum (FWHM) method.....88

Figure 3.16: Diagram showing the branching structure of the main airways in relation to the pulmonary lobes 89

Figure 3.17: 3D surface renderings of two airway trees automatically labelled according to the lobes served by each bronchus 90

Figure 3.18: Illustration of manual validation of airway segmentation.....97

Figure 4.1: Comparison of CT, proton MRI and hyperpolarised xenon MRI data for detecting pulmonary fissures.....105

Figure 4.2: Close-up of a coronal CT lung slice showing an example of a small vessel-like structure close to a fissure boundary.....116

Figure 4.3: Examples of fissure detection methods 118

Figure 4.4: Simplified illustration of the automated lobe segmentation pipeline.....120

Figure 4.5: Example output of the vesselness density filter 125

Figure 4.6: Example output of the fissureness filter 127

Figure 4.7: Results of the initial lobar approximation method 129

Figure 4.8: Schematic illustration of the process for selecting points on the fissure surface 130

Figure 4.9: Example results from the lobe segmentation 135

Figure 4.10: Smoothed 3D surface rendering of the segmented lobe masks.....135

Figure 4.11: Screenshot of the Pulmonary Toolkit graphical interface showing the interactive placing and editing of marker points.....137

Figure 5.1: Schematic of the pipeline for regional image analysis using multi-modal image registration 151

Figure 5.2: Coronal slices from hyperpolarised xenon MRI ventilation datasets 156

Figure 5.3: Ventilation masks for a hyperpolarised xenon MRI ventilation dataset.157

Figure 5.4: Standard bipolar gradient waveform used in hyperpolarised gas MRI diffusion-weighted imaging.....158

Figure 5.5: Fusion of anatomical proton MRI and ventilation xenon MRI for a COPD patient 165

Figure 5.6: Fusion of lobe masks and xenon MRI ventilation images for a COPD patient 166

Figure 5.7: Fusion of lobe masks and xenon MRI ventilation images for a control subject 167

Figure 5.8: Xenon MRI ventilation images illustrated using a 3D rendering of the lobe masks....168

Figure 5.9: Ventilation masks generated for a control subject using hyperpolarised xenon MRI....169

Figure 5.10: Ventilation masks generated for a COPD subject using hyperpolarised xenon MRI....170

Figure 5.11: Apparent diffusion coefficient maps for a COPD patient using hyperpolarised xenon MRI diffusion-weighted imaging 171

Figure 5.12: Emphysema detection results for a COPD patient....172

Figure 6.1: Schematic diagram of the coupled ventilation model used in simulation studies....183

Figure 6.2: Flow rate results for a computational airflow simulation with and without constriction in the airways 184

List of tables

Table 1.1: Imaging parameters for the data used in this thesis.....20

Table 3.1: Algorithm execution time and lung volume for CT lung segmentation.....56

Table 3.2: Validation results for proton MRI lung segmentation.....68

Table 3.3: Comparison of MRI lung segmentation results against previously published methods.....70

Table 3.4: Metrics for assessing partial volume effects.....72

Table 3.5: Execution time for proton MRI lung segmentation.....73

Table 3.6: Airway segmentation results for CT datasets showing the total number of bronchi segmented and the number found in each lobe.....94

Table 3.7: Airway segmentation results for CT datasets showing the number of airways in each generation of the segmented tree.....95

Table 3.8: Airway segmentation results for CT datasets showing the mean radius and mean length of the airways in each generation.....95

Table 3.9: Execution time for the airway segmentation algorithms.....100

Table 4.1: Comparison of published work in lobe segmentation.....112

Table 4.2: Validation results for the automated fissure detection algorithm.....137

Table 4.3: Execution time for the lobe segmentation algorithms.....139

Table 5.1: Regional measurements of ventilated volume percentages.....160

Table 5.2: Regional measurements of the apparent diffusion coefficient (ADC).....161

Table 5.3: Regional measurement results for previously published work in ^{129}Xe imaging.....162

Table 5.4: Regional measurement results for previously published work in ^{3}He imaging.....162

Table 5.5: Regional measurements of emphysema percentage (EP).....163

Table 5.6: Regional measurements of percentile density (PD).....163

Table 5.7: Validation of automatic registration methods between CT and MRI.....173

Index

acini, 24, 26, 183

airflow simulation. *see* computational modelling

airways, 9, 23, 24, **26**–**27**, 29, 36, 50, 74–99, 89, 90, 94, 95, 182–84

- centreline, 84–85
- labelling by pulmonary lobe, **89**–**93**, 96, 99
- labelling by pulmonary segment, 182
- physiology, 26–27
- radius, **86**–**88**, 98
- removal from lung segmentations, 50, 61
- segmentation from CT, 74–99
- trachea, 26, 61, 75, **76**–**80**, 96, 98
- wall thickness, 26, 36, 98

alpha 1-antitrypsin deficiency, 30

alveolar ducts, 26

alveoli, 23, 24, **25**–**26**, 29, 30, 32, 38, 144

apex of lung, 47

apparent diffusion coefficient (ADC), 19, 38, 39, **158**–**59**, 161, 162, 171, 175, 176

arteries, 23, 29

asthma, 30, 37, 38, 40, 145, 156

atlas

- segmentation of lungs from CT, 46
- segmentation of pulmonary lobes from CT, 107, 108, 113, 115, 119

blood vessels. *see* vessels

blue bloaters, 31

breathing, 25, 26, **28**, 30, 31, 147

bronchi. *see* airways

bronchioles, 26

bronchiolitis, 35

bronchopulmonary segments. *see* segments

b-value, 20, **158**, 159, 176

capillaries, 24, 25, 26, 29, 40

carbon dioxide, 23, 25, 28

carina, 26

centreline of airways, 84–85

chronic bronchitis, 30, 31, 36

chronic obstructive pulmonary disease (COPD), 9, **30**–**40**

- regional analysis, 34–40, 144, 156, 160, 161, 162, 163, 165, 166, 168, 170, 171, 172, 175, 176, 177

clinical trial into hyperpolarised xenon-129 MRI, 9, 10, **15**–**17**, 18, 39, 174, 181–82

compliance, 26, 28, 30, 175

computational modelling, 39, 182–84

computational performance, 56, 73, 99, 139

computed tomography (CT), 10, 16, 18–20, **34**–**37**, 39, 40, 147

- airway segmentation, 74–99
- assessment of COPD, 10, **34**–**37**, 40, 144, 163, 175
- imaging parameters, 20
- lobe segmentation, **102**–**40**, 166, 167, 168
- lung segmentation, 44–56

pulmonary segment extraction, 109, 182

registration with MRI, 16, 143, 144, 145, 146, 147–48, 173

conducting airways, 26, 27

cystic fibrosis, 145

data used in this thesis, 18–20, 160, 174, 176

Dice coefficient, **67**, 68, 70, 71, 173

diffusing capacity for carbon monoxide (D_{LCO}), 32, 37

diffusion-weighted (DW) imaging, 15, 19, **38**, 158–59

imaging parameters, 20

dynamic ventilation imaging, 37

dyspnoea, 32

effective diffusion coefficient. *see* apparent diffusion coefficient (ADC)

emphysema, 28, 30–40, 104, 143, 144, 158, 163, 164, 172, 175, 176

percentage (EP), 35, 144, **159**, 163, 164, 175, 176

percentile density (PD), 35, 144, **159**, 163

Empire ’10 lung image registration challenge, 147

empyema, 53, 54

Eulerian reference frame, 154

Eulerian velocity field, 154

Exact ’09 airway segmentation challenge, 43, 75, 96

explosion control in region-growing, 75, 80–83

FEV_1 predicted, 32

fibrosis, 28, 30, 104, 143, 145

fissureness, 117, 118, 119, 121, **126**–**28**

fissures, 22, **23**, **24**, 105

sliding, 147–48, 150

use in lobe segmentation, 102–40

forced expiratory volume in 1 second (FEV_1), 32, 33, 36

forced vital capacity (FVC), 32, 33

full-width-at-half-maximum (FWHM), 86–88

fully automatic segmentation, 43

functional residual capacity (FRC), 32, 33

gas exchange, 23, 24, 25, 26, 29, 30

gold standard

airway segmentation from CT, 96

assessment of lung disease, 40

clinical segmentation, 43

lobe segmentation from CT, 104, 110, 136

grand challenges, 43

Empire ’10 registration of lung images, 147

Exact ’09 airway segmentation, 43, 75, 96

Lola ’11 lung and lobe segmentation, 43, 45

Vessel ’12 vessel segmentation, 121

helium-3, 15, 17, 37, 38, 39, 145, 146, 162, 176, 177

Hessian matrix, 107, 108, 113, 117, 118, 121, 122, 123, 124, 126, 128

histology, 39

horizontal fissure, 22, 24, 114, 134, 138

Hounsfield units (HU), 34, 45, 75

hyperpolarised gas MRI, **37**–**40**, 145–46

hyperpolarised helium-3, 15, 17, 37, 38, 39, 145, 146, 162, 176, 177

hyperpolarised xenon-129 MRI, 10, 15, 18–20, 37, 39, 40, 143

clinical trial, 9, 10, **15**–**17**, 18, 39, 174, 181–82

diffusion-weighted (DW) imaging, 15, 19, 20, **158**–**59**

perfusion imaging, 15, 40

regional measurements, 141–78

ventilation imaging, 15, 19, 20, 156–57, 165, 166, 167, 168, 169, 170

hypoxemia, 31

idiopathic pulmonary fibrosis, 104

ITK-Snap, 67

labelling of airways by pulmonary lobe, **89**–**93**, 96, 99

labelling of airways by pulmonary segment, 182

leakage

airway segmentation from CT, 75, 80–83, 98

lung segmentation from proton MRI, 59, 62, 63, 64, 65, 70, 177

level sets, **62**–**64**, 65

lingula, 91

lobes, 22, **23**, 36, 89–93, 102–40, 143, 148, 160, 161, 163, 166, 167, 168, 175, 176

labelling of airways, **89**–**93**, 96, 99

segmentation from CT, **102**–**40**, 166, 167, 168

sliding, 147–48

lobules, 24, 29

Lola ’11 lung and lobe segmentation challenge, 43, 45

lung volume reduction surgery, 35, 104

lungs

lung disease, 26, 30–40

physiology, 23–29

region of interest (ROI), 52

segmentation from CT, 44–56

segmentation from hyperpolarised gas MRI, 58, 145

segmentation from proton MRI, **57**–**73**, 177

separation of left and right lung segmentations, 47, 51, 61

sliding, 147–48, 150

magnetic resonance imaging (MRI)

hyperpolarised xenon-129 MRI. *see* hyperpolarised xenon-129 MRI

proton MRI. *see* proton MRI

Miccai

Empire ’10 lung image registration challenge, 147

Exact ’09 airway segmentation challenge, 43, 75, 96

Lola ’11 lung and lobe segmentation challenge, 43, 45

MLSEC-ST, 107, 118

multilevel B-spline, 119, 132–**34**

mutual information, 147, 148

oblique fissure, 22, 24, 91, 92, 93, 107, 125, 134, 137, 138

obstructive lung disease, 30, 31, 32, 36, 37

oxygen, 23, 25, 26, 28, 29, 31

parenchyma, 38, 44, 45, 77, 78, 79, 176

parietal pleura, 22

partial volume effects

airway segmentation from CT, 74, 75, 76, 79, 80, 82, 98

lobe segmentation from CT, 105, 113

lung segmentation from CT, 47, 51, 64

lung segmentation from proton MRI, 71–73, 177

vessel segmentation from CT, 122

percentile density (PD), 144, 159, 163

performance. *see* computational performance

perfusion, 23, 26, 40

imaging, 15, 37, 40

personalised computational modelling. *see* computational modelling

pink puffers, 31

plethysmography, 32, 33

pleura, 22, 24

pleural fluid, 22, 24, 147

precision, 67, 68, 70, 71, 173

proton MRI, 16, 18–20, 57–73, 105, 144, 147–48, 149, 165, 177

- imaging parameters, 20
- lung segmentation, **57**–**73**, 177
- registration with CT, 145, **147**–**48**

pulmonary arteries, 23, 29

pulmonary circulation. see vessels

pulmonary empyema, 53, 54

pulmonary fibrosis, 28, 30, 104, 143, 145

pulmonary fissures. *see* fissures

pulmonary function tests (PFTs), 31, **32**–**33**, 35, 36

pulmonary lobes. *see* lobes

pulmonary segments, 24, 27, 109, 182

pulmonary surfactant, 26, 27, 28

Pulmonary Toolkit, 2, 18, 136, 137

pulmonary veins, 23, 29

radiodensity, 34

radius of airways, **86**–**88**, 98

recall, 67, 68, 70, 71, 173

regional image analysis, 10, 34, 35, 37, 38, 39, 40, 86, 104, **141**–**78**, 181–82

region-growing

- airway segmentation from CT, 75, 80–83
- lung segmentation from CT, 45, 46
- lung segmentation from proton MRI, 59, 61, 64, 65
- vessel segmentation from CT, 121

registration, 16, 143, 144, 145–46, 147–48, **149**–**56**, 173, 177

- CT to MRI, 16, **152**–**55**, 166, 167, 168, 173
- Empire ’10 challenge, 147
- proton MRI to xenon MRI, 16, 155
- use in regional analysis, 145–46
- validation, 173, 177

residual volume, 16

respiratory bronchioles, 26

respiratory system, 22, 23–29

respiratory zone, 26

restrictive lung disease, 30, 32

ribs, 28, 46, 147

scintigraphy, 37

segmentation, 43

- airways from CT, 74–99
- lobes from CT, **102**–**40**, 166, 167, 168
- lungs from CT, 44–56
- lungs from proton MRI, **57**–**73**, 177
- pulmonary segments from CT, 109, 182
- vessels from CT, 29, 115, **121**–**25**

segments, 24, 27, 109, 182

semi-automatic segmentation, 43

silicosis, 104, 143

skeletonisation, 84–85

sliding of the pleura and fissures, 147–48, 150

small airways disease, 9

software, 2, 18, 136, 137

Spect, 37

spirometry, 32, 33, 39

structure tensor, 107, 113, 117

surfactant, 26, 27, 28

thinning, 85

tidal volume, 28

total lung capacity (TLC), 16, 30

trachea, 26, 61, 75, **76**–**80**, 96, 98

tuberculosis, 104, 143

validation

- airway segmentation from CT, 96–97
- lobe segmentation from CT, 110, 111, 112, 136–37
- lung disease assessment using hyperpolarised gas MRI, 39
- lung segmentation from CT, 53
- lung segmentation from proton MRI, 67–68, 70
- pulmonary segment extraction, 182
- registration, 173, 177

vasculature. *see* vessels

veins, 23, 29

ventilated volume, 144, 145, **156**–**57**, 160, 162, 169, 170, 175, 176

ventilation, 23, 26, 28, 31, 37, 40, 145, 174

- computational model, 182–84
- defects, 39, 58, 145, 156, 170, 174, 175
- imaging, 15, 19, 37, 38, 105, 144, 156, 165, 166, 167, 168, 169, 170, 176
- imaging parameters, 20

ventilation-to-perfusion ratio, 26, 31, 40

Vessel ’12 segmentation challenge, 121

vesselness, 119, 121, 122, **123**–**24**, 125, 128

vesselness density, 121, 122, **125**, 128, 129

vessels, 24, 28, **29**, 48, 61, 75, 107, 116, 117, 119, 121–25, 129

- segmentation, 115, **121**–**25**

visceral pleura, 22, 24

wall thickness of airways, 26, 36, 98

watershed transform

- lobe segmentation from CT, 107, 116, 128, 129
- lung segmentation from CT, 47, 51
- pulmonary segment subdivision, 182

xenon-129 MRI. *see* hyperpolarised xenon-129 MRI