Thesis submitted to the University of Sheffield for the degree of
Doctor of Philosophy (PhD)

Measuring Pulmonary Gas Exchange with Hyperpolarised ^{129}Xe Magnetic Resonance Imaging and Spectroscopy

Jemima Hallam Pilgrim-Morris

April 2025

Supervisors:
Professor Jim M. Wild
Dr Neil J. Stewart
Dr Guilhem J. Collier

Abstract

The purpose of this thesis was to develop hyperpolarised ^{129}Xe magnetic resonance imaging and spectroscopy (MRI/S) methodology for the assessment of pulmonary gas exchange and haemodynamics. This was accomplished through the following five sub-projects:

1. Validation and development of models to estimate the transfer factor of the lung for carbon monoxide (TL_{CO}) from ^{129}Xe metrics. Applying regression modelling on a voxel-wise level to create parametric TL_{CO} maps provides a useful tool for regional visualisation and clinical interpretation of ^{129}Xe gas exchange MRI.
2. Assessment of compressed sensing (CS)-accelerated dissolved-phase ^{129}Xe imaging. CS reconstruction enabled acquisition time to be halved, and reduced background noise, whilst preserving key gas exchange metrics.
3. Development of a keyhole reconstruction technique to map the amplitude and phase of dissolved ^{129}Xe red blood cell (RBC) signal oscillations in the lung vasculature from dissolved ^{129}Xe MRI data. This may provide a means to probe the effects of the cardiac pulse wave in the pulmonary microvasculature.
4. Evaluation of RBC oscillation amplitude and phase mapping in post-COVID-19 and chronic thromboembolic pulmonary hypertension (CTEPH) patients and initial attempts to validate against established imaging methods and a computational model. ^{129}Xe RBC oscillation amplitude and phase mapping showed potential sensitivity to lung damage at the microvascular level, and agreement with dynamic contrast enhanced MRI and numerical modelling in a small number of CTEPH patients.
5. Development of chemical shift saturation recovery (CSSR) spectroscopy, imaging and analysis techniques to assess pulmonary gas exchange and measure alveolar septal thickness (h) in healthy volunteers and patients with systemic sclerosis and/or pulmonary arterial hypertension. h was significantly greater for the patients than the healthy volunteers. Dynamic CSSR gas uptake imaging allowed for regional quantification of alveolar septal thickness, which could help identify fibrosis in heterogeneous lung disease.

Acknowledgements

First and foremost, I would like to thank my three supervisors, Jim Wild, Neil Stewart and Guilhem Collier. Jim's guidance and support throughout the course of my PhD has been invaluable. I am grateful for the many opportunities that Jim has supported me to take and for motivating and inspiring me to be the best researcher that I can be. I would not have been able to complete this PhD without the encouragement, support and patience of Neil, who has spent many hours with me at the scanner and discussing my research, becoming both a mentor and friend. I would like to thank Guilhem for all his help over the years and for always providing me with good advice - be that academic or climbing-related!

I feel incredibly lucky to have been surrounded by such a friendly and talented group of people in POLARIS, as well as the wider MR Science group and Academic Radiology team. It has been a great environment in which to start my research career. There are too many people to name - I am thankful to everyone for all their help and for being a pleasure to work with. Thank you to Ryan Munro for all the time spent polarising xenon for my experiments, as well as for teaching me to operate the polariser. Thank you to the radiography team, in particular David Capener and Jody Bray, for organising various patient studies and to Jen Rodgers and Leanne Armstrong for ensuring the smooth operation of the department.

I have been fortunate to be part of the DiMeN DTP community throughout my PhD. I am thankful for all the additional training opportunities I have had and people that I have met as a result of being a DiMeN student.

My amazing friends have been a constant source of support and joy throughout my PhD. I am so grateful to both my long-distance and Sheffield friends, especially my housemate Roberta, who has been on this PhD journey with me from the start. I am so happy to have really found my community in Sheffield.

Finally, I am enormously grateful to my family for all their love and support throughout my academic journey.

This PhD project was funded by the Medical Research Council Discovery Medicine North Doctoral Training Partnership (MR/R015902/1) and additional Collaborative Awards in Science and Engineering (CASE) funding from GE Healthcare.

Contents

List of Figures

List of Tables

List of Abbreviations

ADC	Apparent Diffusion Coefficient
AF	Acceleration Factor
AIF	Arterial Input Function
BART	Berkeley Advanced Reconstruction Toolbox
bSSFP	Balanced Steady State Free Precession
BW	Bandwidth
CA	Contrast Agent
CF	Cystic Fibrosis
CI	Confidence Interval
COPD	Chronic Obstructive Pulmonary Disease
CS	Compressed Sensing
CSI	Chemical Shift Imaging
CT	Computed Tomography
CTEPH	Chronic Thromboembolic Pulmonary Hypertension
CV	Coefficient of Variation
DAQ	Data Acquisition Window
DCE	Dynamic Contrast Enhanced
DCF	Density Compensation Function
DEV	Dose Equivalent Volume
DW	Diffusion Weighted
ECG	Electrocardiography

EPI	Echo Planar Imaging
EPSI	Echo Planar Shift Imaging
FEV_1	Forced Expiratory Volume
FID	Free Induction Decay
FOV	Field of View
FRC	Functional Residual Capacity
FWHM	Full Width at Half Maximum
GLI	Global Lung Initiative
HV	Healthy Volunteer
IDEAL	Iterative Decomposition of Water and Fat with Echo Asymmetry and Least-Squares Estimation
ILD	Interstitial Lung Disease
IPF	Idiopathic Pulmonary Fibrosis
K_{CO}	Carbon Monoxide Transfer Coefficient
MAE	(Mean Absolute Error
MESI	Multi-Echo Spectroscopic Imaging
MNS	Multi-Nuclear Spectroscopy
MOXE	Model of Xenon Exchange
MPA	Main Pulmonary Artery
mPAP	Mean Pulmonary Artery Pressure
MPV	Main Pulmonary Vein
MSE	Mean Squared Error
MTT	Mean Transit Time
NA	Natural Abundance
NMAE	Normalised Mean Absolute Error
NMR	Nuclear Magnetic Resonance
NSA	Number of Signal Averages

NSIP	Non-specific Interstitial Pneumonia
PAH	Pulmonary Arterial Hypertension
PBF	Pulmonary Blood Flow
PBV	Pulmonary Blood Volume
PE	Pulmonary Embolism
PEA	Pulmonary Endarterectomy
PFT	Pulmonary Function Testing
PH	Pulmonary Hypertension
pPTT	Pulmonary Pulse Wave Transit Time
PREFUL	Phase-Resolved Functional Lung Imaging
PVD	Pulmonary Vascular Disease
PVR	Pulmonary Vascular Resistance
RBC	Red Blood Cell
RF	Radiofrequency
RHC	Right Heart Catheterisation
RMS	Root Mean Square
RMSE	Root Mean Squared Error
RV	Residual Volume
SAR	Specific Absorbance Rate
SEOP	Spin Exchange Optical Pumping
SNR	Signal to Noise Ratio
SPGR	Spoiled Gradient Echo
SPLASH	SPectral Localization Achieved by Sensitivity Heterogeneity
SSc	Systemic Sclerosis
TE	Echo Time
TL_{CO}	Transfer Capacity of the Lungs for Carbon Monoxide
TLC	Total Lung Capacity

TLV	Total Lung Volume
TR	Repetition Time
TTP	Time to Peak
UTE	Ultra-short Echo Time
V_A	Alveolar Volume
VDP	Ventilation Defect Percentage
VV	Ventilated Volume
Xe-CTC	Xenon Clinical Trials Consortium
XTC	Xenon Polarisation Transfer Contrast
ZTE	Zero Echo Time

Chapter 1

Introduction

Respiratory disease poses an enormous burden to global health and is a leading cause of death and disability. The World Health Organisation estimates that more than 8 million deaths a year are caused by lung disease [1]. This is likely to grow in the future as a result of the rising temperatures, increased air pollution and wildfire risk which are consequences of the climate crisis. The vital importance of lung health was highlighted by the COVID-19 pandemic that emerged in 2019. More than 8 million lives were lost, mostly as a result of lung involvement [1]. In addition to deaths, the pandemic caused widespread and long-lasting social and economic effects and millions of people have been left with debilitating ongoing symptoms, known as ‘long-Covid’.

A 2016 report from the British Lung Foundation stated that, in the UK, more than 12 million people have a lung disease diagnosis and 115,000 people die from lung disease every year [2]. This is equivalent to one person every five minutes. This places an immense strain on the NHS, accounting for over 700,000 hospital admissions a year and an estimated financial burden of £11.1 billion [3]. Lung disease is strongly associated with social deprivation, and diseases such as Chronic Obstructive Pulmonary Disease (COPD) and lung cancer are far more common in underserved communities. The prevention, diagnosis and management of lung disease is therefore crucial to reduce disease prevalence, mortality and health inequalities. More research is urgently needed to increase our understanding of lung disease mechanisms and to provide better biomarkers of lung health in order to ameliorate patient management and treatment.

Respiratory diseases can be divided into several categories, depending on the affected region and on the disease aetiology. These include: obstructive lung disease, restrictive lung disease, pulmonary vascular disease, infectious diseases, inflammatory diseases, lung cancer and pleural cavity disease. This thesis is focused on obstructive, restrictive and vascular lung diseases.

Obstructive lung disease is caused by the blockage or narrowing of airways which increases resistance to air flow. The most common forms of obstructive lung disease are asthma, which affects more than 350 million people globally [4], and COPD, which is the fourth-leading cause of death worldwide [1]. Both asthma and COPD cause symptoms of cough, wheezing and breathlessness. COPD is characterised by the destruction of alveolar tissue (emphysema) and obstruction of small airways from mucus (bronchitis) and inflammation. This causes a progressive decline in lung function along with acute exacerbations [5]. There is a demand for better diagnostic and phenotyping strategies for COPD and asthma, including for patients who have both diseases. Improved methods for monitoring and staging COPD and assessing the impact of treatments and interventions such as endobronchial valves [6] are also needed.

Restrictive lung diseases cause reduced distensibility of the lungs, restricting lung expansion and leading to diminished total lung capacity (TLC). They originate from the destruction or inflammation of lung tissue (intrinsic causes) or from extra-pulmonary conditions which restrict pulmonary mechanisms, such as neuromuscular disease or obesity (extrinsic causes) [7]. Interstitial Lung Disease (ILD) is a term used to describe a group of more than 100 restrictive diseases affecting the lung parenchyma, especially the interstitial space within the alveolar walls [8, 9]. ILD is characterised by interstitial thickening, due to either inflammation or fibrosis or both, which can result in impaired oxygen transfer to the bloodstream. Prognosis can often be very poor and progression unpredictable, particularly for Idiopathic Pulmonary Fibrosis (IPF), one of the most common ILDs [10]. ILD-like pulmonary fibrosis can develop as a sequelae of COVID-19 infection [11] and is also a frequent complication and cause of morbidity in systemic sclerosis (SSc) [12]. Currently there is a shortage of robust, sensitive biomarkers of gas exchange impairment, which makes monitoring disease progression and developing novel treatments for ILD challenging [13].

Lung disease which affects the pulmonary circulation is known as pulmonary vascular disease (PVD). This is often associated with other diseases such as ILD and COPD and is linked to a worsened prognosis in these groups [14]. There are two main types of PVD: pulmonary embolism (PE) and pulmonary hypertension (PH). PE is the obstruction of pulmonary arteries by a blood clot, generally a result of deep vein thrombosis in the lower limbs. PH is a condition where the blood pressure in the pulmonary arteries is elevated, causing right heart strain. PH diagnosis can often be difficult, because the most common symptoms of dyspnoea and fatigue are similar to other respiratory and cardiac disorder symptoms [15]. It can also be challenging to differentiate between different PH subtypes, such as pre-capillary and post-capillary PH. Consequently, the time between the onset of PH symptoms and diagnosis often exceeds two years, which is detrimental to patient outcomes and delays the initiation of treatment [16]. PH evaluation includes multiple steps such as blood testing, electrocardiogram (ECG), pulmonary function testing (PFT), echocardiogram and Computed Tomography (CT). Right heart catheterisation (RHC) can confirm a PH diagnosis, however this is a highly invasive procedure with a risk of complications such as pulmonary arterial rupture [17, 14]. Hence, more sensitive, non-invasive measures of pulmonary vascular function are needed to aid PH diagnosis.

The gold standard of pulmonary imaging is CT, but this method is not well-suited for long-term disease monitoring, due to its use of ionising radiation, and can only provide limited functional information. Functional lung measurements are clinically vital and are typically acquired through PFT, such as the transfer factor of the lungs for carbon monoxide (TL_{CO}) test used to measure gas exchange. However, PFT lacks the regional specificity that CT can provide, which is especially needed in heterogeneous disease. In contrast, MRI does not require ionising radiation or invasive processes and is remarkably versatile, providing both structural and functional images of the body. Traditionally, imaging the lungs with conventional ^{1}H MRI has been challenging, primarily due to their low proton density ($\sim$0.1 g cm^{-3} in healthy lungs [18]). The MR signal is directly proportional to proton density, so this results in weak signals and low signal-to-noise ratio (SNR). Moreover, the large magnetic susceptibility differences resulting from multiple air-tissue interfaces and the effects of cardiac and respiratory motion further degrade the MR signal. The development of ultra-short and zero echo time (UTE/ZTE) MRI sequences for the lungs has overcome some of these barriers and enabled improved structural lung imaging [19, 20], but this does not satisfy the clinical need for functional, non-ionising lung imaging.

Hyperpolarised xenon-129 (^{129}Xe) MRI/S uses inhaled ^{129}Xe gas - a safe and inert contrast agent - to yield a myriad of information about lung ventilation, microstructure and gas exchange. ‘Hyperpolarised’ refers to the increased nuclear polarisation that is achieved via spin exchange optical pumping. This is necessary to increase the MR signal by 4 - 5 orders of magnitude, counteracting the low spin density and thermal polarisation of ^{129}Xe gas.

The mainstay of hyperpolarised ^{129}Xe lung MRI is ventilation imaging, whereby the distribution of gas-phase ^{129}Xe in the lungs is measured. Areas of low or no signal in the ventilation images represent a lack of ^{129}Xe signal, signifying airway obstruction or restriction. The first hyperpolarised ^{129}Xe ventilation images were obtained in 1994 in excised mouse lungs [21] and in vivo lung imaging in humans was demonstrated soon after in 1997 [22]. Since then, progress in polarisation and imaging techniques has helped establish ^{129}Xe ventilation imaging as a robust technique which can provide high sensitivity to ventilation defects in obstructive lung diseases such as COPD [23], asthma [24] and cystic fibrosis (CF) [25].

^{129}Xe MRI can also be used to make measurements of lung microstructure [26, 27]. When inhaled, the ^{129}Xe gas molecules diffuse through the alveolar airspaces with random Brownian motion. Reflections occur at the alveolar walls, restricting the diffusion and leading to a reduced apparent diffusion coefficient (ADC) which is reflective of the alveolar dimensions. By fitting the diffusion-weighted (DW) MRI signal to theoretical models of gas diffusion [28, 29], measurements of morphological parameters analogous to those from histology can be extracted. DW ^{129}Xe MRI is particularly sensitive to emphysema, where loss of alveolar tissue causes less restricted diffusion and therefore increased ADC [30].

One of the most novel and unique aspects of hyperpolarised ^{129}Xe MRI is its ability to quantify the transfer of gas from the alveoli into the bloodstream. ^{129}Xe is mildly soluble in the parenchymal tissue, pulmonary capillary blood plasma and red blood cells (RBCs). Different chemical shifts are exhibited for ^{129}Xe in the alveolar airspaces, the lung tissue and blood plasma (collectively referred to as ‘membrane’, M), and the RBCs. This means that ^{129}Xe can act as a surrogate for oxygen and be used to probe gas exchange by measuring the MR signal from each environment [31]. The ratios of these signals (RBC:M, RBC:Gas and

M:Gas) provide sensitivity to gas exchange function, tissue thickening and perfusion limitation in diseases such as ILD, asthma, and COPD [32, 33, 34, 35, 36, 37, 38, 39]. Time-resolved ^{129}Xe spectroscopy measurements can be used in conjunction with analytical diffusion models to estimate clinically relevant parameters such as the alveolar septal wall thickness [40, 41]. Recently, measurement of oscillations in the ^{129}Xe RBC signal has emerged as a way to assess microvascular function [37].

1.1 Thesis Aims and Objectives

The overarching aim of this thesis was to develop hyperpolarised ^{129}Xe MRI/S methodology for the assessment of pulmonary gas exchange and haemodynamics.

To achieve this goal, the following specific objectives were set:

1. Investigate the relationship between ^{129}Xe MRI metrics and standard clinical measures of gas exchange and to validate and expand on a model which links the two.
2. Develop a compressed sensing reconstruction for non-Cartesian ^{129}Xe gas exchange imaging data in order to reduce scan time and increase image SNR.
3. Implement a post-acquisition keyhole reconstruction to spatially resolve the amplitude and phase of cardiogenic oscillations in the ^{129}Xe RBC signal and compare with previous methodology which does not account for oscillation phase.
4. Develop MR pulse sequences and analysis protocols for time-resolved ^{129}Xe spectroscopy and imaging and use these to quantify gas exchange in healthy volunteers and patients with SSc and/or pulmonary arterial hypertension (PAH).

1.2 Thesis Overview and Author Contributions

The thesis is organised into ten chapters, including this introductory chapter.

Chapter 2 provides an overview of the essential theoretical principles which form the foundations of this thesis, covering MR physics, image reconstruction and an introduction to hyperpolarised ^{129}Xe MRI.

Chapter 3 begins with a description of lung anatomy and physiology and some conventional methods for measuring pulmonary gas exchange. This is followed by a review of alternative techniques to quantify gas exchange using hyperpolarised ^{129}Xe MRI/S.

Chapter 4 recounts the methodology used throughout this thesis, including an overview of the MRI hardware and some of the pulse sequences used.

Chapter 5 is the first chapter of original research and compares three models for predicting TL_{CO} from ^{129}Xe MRI metrics. This includes validation of a previously published physiology-based model and the development of two machine learning models, using data from patients with asthma and/or COPD. A regional model is implemented, which allows for the creation of TL_{CO} maps. This chapter is based on the following publication: ‘A framework for modelling whole-lung and regional TL_{CO} using hyperpolarised ^{129}Xe lung MRI’, **JH Pilgrim-Morris**, LJ Smith, H Marshall, BA Tahir, GJ Collier, NJ Stewart, JM Wild, *European Respiratory Journal Open Research*, 2025; 11 (1): 00442-2024. doi: 10.1183/23120541.00442-2024 Author contributions:

- Study design: JHP-M, LJS, BAT, NJS, JMW
- Data acquisition: LJS, HM, GJC
- Data analysis/model development: JHP-M
- Manuscript preparation: JHP-M, LJS, HM, BAT, GJC, NJS, JMW

Chapter 6 describes the development and assessment of a compressed sensing ^{129}Xe gas exchange imaging reconstruction which permits a 50% reduction in scan time. This chapter is based on the following publication: ‘Compressed sensing reconstruction for high-SNR, rapid dissolved ^{129}Xe gas exchange MRI’, **JH Pilgrim-Morris**, GJ Collier, RS Munro, G Norquay, NJ Stewart, JM Wild. *Magnetic Resonance in Medicine*, 2025; 93: 741-750. doi: 10.1002/mrm.30312. Author contributions:

- Study design: JHP-M, GJC, NJS, JMW
- Data acquisition: JHP-M, GJC, RSM, GN, NJS
- Data analysis: JHP-M
- Manuscript preparation: JHP-M, GJC, NJS, JMW

Chapter 7 reports a novel reconstruction method for dissolved-phase ^{129}Xe lung spectroscopic images, which is capable of spatially resolving the amplitude and phase of oscillations in the signal from ^{129}Xe dissolved in the pulmonary capillary RBCs.

Chapter 8 builds on Chapter 7, by applying the RBC oscillation mapping method to data from two patient groups: COVID-19 and chronic thromboembolic pulmonary hypertension (CTEPH). Also included in this chapter is preliminary work to validate the oscillation mapping method against other imaging modalities and numerical simulations.

Both Chapter 7 and 8 are partially based on the following publication: ‘Mapping the amplitude and phase of dissolved ^{129}Xe red blood cell signal oscillations with keyhole spectroscopic lung imaging’, **JH Pilgrim-Morris**, GJ Collier, M Takigawa, S Strickland, R Thompson, G Norquay, NJ Stewart, JM Wild. *Magnetic Resonance in Medicine*, 2025; 93: 584-596. doi: 10.1002/mrm.30296. Author contributions:

- Study design: JHP-M, GJC, NJS, JMW

- Data acquisition: JHP-M, GJC, SS, RT, GN, NJS
- Data analysis: JHP-M, MT
- Manuscript preparation: JHP-M, GJC, NJS, JMW

Additionally, the dynamic-contrast enhanced MRI data included in Chapter 8 were acquired and analysed using techniques developed by Dr Paul Hughes and Dr Laura Saunders from POLARIS, University of Sheffield. The patient-specific pulmonary blood flow simulations were performed by Dr Behdad Shaarbaf Ebrahimi from the University of Auckland Bioengineering Institute.

Chapter 9 presents methods for measuring temporal gas exchange using CSSR spectroscopy and an analysis method which uses bootstrapping to estimate the uncertainty on model-derived parameters. This technique is applied to investigate gas exchange in patients with SSc and/or PAH. Initial work towards a regional CSSR imaging sequence is reported, including evaluation in a small number of healthy volunteers. The CSSR pulse sequences were developed by JHP-M and NJS and all data analyses were performed by JHP-M.

Chapter 10 summarises the conclusions of the previous five chapters of research and the future research opportunities stemming from this PhD project.

Chapter 2

Theoretical Background

This chapter covers the theoretical background relevant to the experimental work in this thesis, covering MR Physics, image reconstruction and an introduction to hyperpolarised ^{129}Xe MRI. The following textbooks were used as a reference: [42, 43].

2.1 Principles of Nuclear Magnetic Resonance

MRI has been used clinically since the early 1980s, following a series of important discoveries by pioneers in the fields of Physics, Chemistry and Medicine, many of whom were awarded a Nobel Prize for their work. The nuclear magnetic resonance (NMR) phenomenon was first observed by Isidor Rabi [44] in 1938, who measured the spin of a proton and its interaction with a magnetic field, earning him the 1944 Nobel Prize in Physics. Felix Bloch [45, 46] and Edward Purcell [47] were, independently, the first to measure magnetic resonance in solids and liquids. For this discovery, they shared the 1952 Nobel Physics Prize. Another Nobel Prize, this time for Chemistry, was won by Richard Ernst in 1991 for his development of Fourier transform NMR spectroscopy [48] in 1966.

In 1971, NMR was first used for biomedical applications by Raymond Damadian, who found increased relaxation times in mouse tumours in comparison to healthy tissues [49]. Paul Lauterbur and Peter Mansfield were awarded the 2003 Nobel Prize in Physiology/Medicine for independently demonstrating the use of magnetic field gradients to localise NMR signals in 1973 [50].

In this section, the phenomenon of NMR is reviewed, using both classical and quantum mechanics descriptions.

2.1.1 Spin

Spin angular momentum, or ‘spin’, is the intrinsic angular momentum that is a property of fundamental particles. In the presence of a magnetic field, the spin of a proton precesses about the field direction and the foundation of MRI is based on manipulating and detecting the bulk precession of spins within organic molecules in the body.

2.1.1.1 Classical Mechanics Description

In the classical picture of a proton in a magnetic field, we treat the proton as an electric charge rotating with motion which is analogous to that of a spinning top or gyroscope, as shown in Figure 2.1.1.

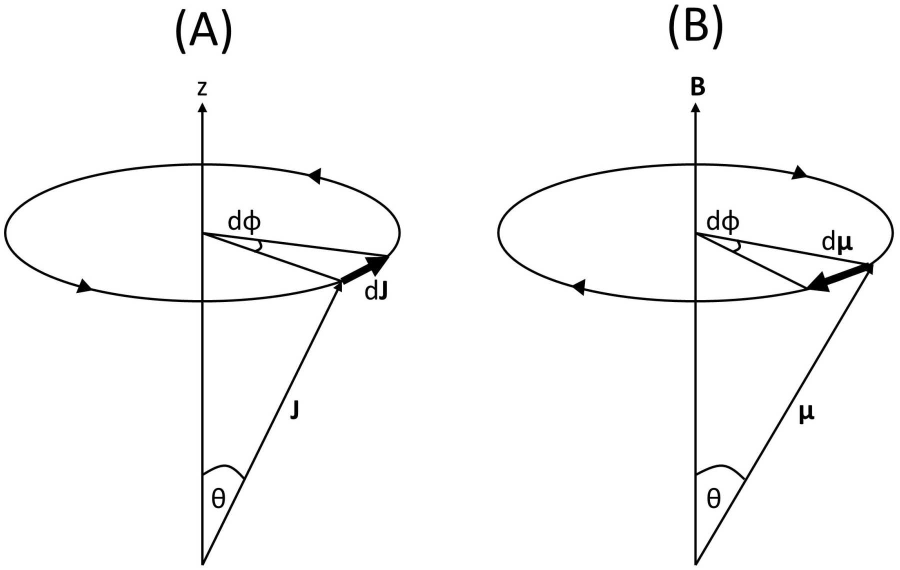

Figure 2.1.1: Angular momentum diagrams illustrating the precession of (A) a gyroscope with angular momentum **J** in a gravitational field and (B) a proton with magnetic moment μ in a magnetic field, where the torque, and therefore the precession, are in the opposite direction.

The gyroscopic motion of a spinning top is a consequence of a gravitational torque, which is perpendicular to the spin axis. For a proton in a static magnetic field, B_0, the interaction of the spin with the field results in a torque, which causes it to precess about the field direction with a frequency proportional to the magnetic field strength, called the Larmor frequency. The rotation of charge produces an effective electric current loop around the spin axis, which both interacts with the static magnetic field and produces its own magnetic field, with magnetic moment $\boldsymbol{\mu}$. The change in μ in time dt is given by:

$$|d\boldsymbol{\mu}| = \gamma|\boldsymbol{\mu} \times \mathbf{B}|dt = \gamma\mu B\sin\theta dt, \tag{2.1.1}$$

where γ is the gyromagnetic ratio, which is a constant that is unique to each nuclei. Using the geometric representation in Figure 2.1.1, $d\mu$ can also be written as:

$$|d\boldsymbol{\mu}| = \mu\sin\theta|d\phi| \tag{2.1.2}$$

Combining Equations 2.1.1 and 2.1.2 and using $\omega \equiv \left|\frac{d\phi}{dt}\right|$, we arrive at the Larmor precession frequency:

$$\omega_0 = \gamma B_0 \tag{2.1.3}$$

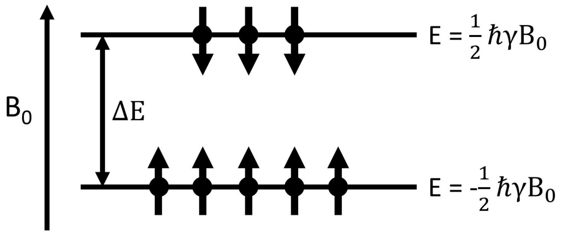

Figure 2.1.2: The energy levels of spin-$\frac{1}{2}$ nuclei (with a positive gyromagnetic ratio) in a magnetic field B_0. Spins which are aligned with the direction of B_0 have a lower energy and therefore there is a larger population of spins in this state.

2.1.1.2 Quantum Mechanics Description

Spin is a quantum property and therefore can only take discrete values. Electrons, protons and neutrons each have a spin of $\frac{1}{2}$ and the overall spin of a nucleus is determined by its number of protons and neutrons. For MRI, nuclei must have a non-zero total spin. The magnetic moment of nuclei with non-zero spin is given by:

$$\boldsymbol{\mu} = \gamma\hat{\mathbf{I}}, \tag{2.1.4}$$

where $\hat{\mathbf{I}}$ is the spin angular momentum quantum operator. To 'observe' the nuclear spin, $\hat{\mathbf{I}}$ is projected along one an axis (chosen here to be the z-axis) onto the eigenstates $|m\rangle$:

$$\hat{\mathbf{I}}_z|m\rangle = m|m\rangle \tag{2.1.5}$$

m is the spin angular momentum quantum number, which can take $2I$+1 possible values, in the range $-I$ to $+I$, where I is the absolute value of the nuclear spin. In the absence of a magnetic field, the eigenstates are degenerate and exist in a linear superposition, but this degeneracy is lifted when the nuclei are placed in an external magnetic field (B_0). The energy levels are separated according to the Zeeman splitting Hamiltonian:

$$\hat{\mathbf{H}} = -\gamma\hbar\mathbf{B}_0\hat{\mathbf{I}}_z \tag{2.1.6}$$

where $\hbar$ is the reduced Plank's constant ($h/2\pi$). The eigenenergies of this Hamiltonian are:

$$E = -\boldsymbol{\mu}\cdot\mathbf{B} = -\gamma\hbar B_0 m \tag{2.1.7}$$

For a spin-$\frac{1}{2}$ nucleus, $m = \pm\frac{1}{2}$ and so this corresponds to two energy levels, shown in Figure 2.1.2. The $|\frac{1}{2}\rangle$ state is referred to as 'spin-up' and the $|-\frac{1}{2}\rangle$ state is referred to as 'spin-down' and these correspond to the spin vectors being aligned and anti-aligned with the magnetic field, respectively.

The energy difference between the spin-up and spin-down states is given by:

$$\Delta E = \left(\frac{1}{2} - \left(-\frac{1}{2}\right)\right)\gamma\hbar B_0 = \gamma\hbar B_0 \tag{2.1.8}$$

which has an associated frequency:

$$\Delta E = \hbar\omega \tag{2.1.9}$$

Equating Equations 2.1.8 and 2.1.9:

$$\hbar\omega_0 = \gamma\hbar B_0 \tag{2.1.10}$$

Therefore, the Larmor frequency is the frequency of the transition between the two energy levels.

For an ensemble of nuclei in a magnetic field, the number of nuclei in each energy level ($N_{\uparrow/\downarrow}$) is given by the thermal Boltzmann distribution:

$$\frac{N_\uparrow}{N_\downarrow} = \exp\left(\frac{\Delta E}{k_B T}\right), \tag{2.1.11}$$

where k_B is the Boltzmann constant (1.38 x 10^{-23} J K^{-1}) and T is temperature. At body temperature and clinical field strengths, $\Delta E = \gamma\hbar B_0 \lll k_B T$, and Equation 2.1.11 can be simplified by taking the Taylor series expansion:

$$\frac{N_\uparrow}{N_\downarrow} = 1 + \frac{\gamma\hbar B_0}{k_B T} \tag{2.1.12}$$

Assuming a body temperature of 37° and a 1.5 T field strength, this fraction is approximately equal to 1.000004, or in other words, for every million protons in the spin-down state, there are a million and four protons in the spin-up state. The difference between the number of nuclei in each state is defined as the spin excess:

$$N_\uparrow - N_\downarrow = \frac{N}{2}\frac{\gamma\hbar B_0}{k_B T} \tag{2.1.13}$$

where N is the total number of spins per unit volume. The polarisation P is the ratio of the spin excess to the total number of nuclei:

$$P = \frac{|N_\uparrow - N_\downarrow|}{N} \tag{2.1.14}$$

This creates the net magnetisation, M_0, which is the vector sum of all spins. Assuming that B_0 is acting along the z-axis:

$$M_0 = N\mu_z P = N\frac{\gamma^2\hbar^2 B_0}{4k_B T}, \tag{2.1.15}$$

where μ_z is the projection of $\boldsymbol{\mu}$ along the z-axis, equal to $\frac{\gamma\hbar}{2}$ for a spin-$\frac{1}{2}$ system. Although at thermal equilibrium, proton polarisation is small (P $\sim$ 5 x 10^{-6}), the large spin density of protons within the body compensates for this and results in a measurable net magnetisation.

2.1.2 Excitation and Relaxation

The net magnetisation is described by a vector **M**, the sum of all magnetic moments per unit volume. To understand the effects of excitation and relaxation, it is useful to consider a rotating frame of reference, which rotates around the z-axis (the direction of B_0) at frequency ω_0. In this frame, spins which precess with frequency ω_0 are stationary. The magnetisation in the rotating frame, denoted (x', y', z') can be expressed in terms of the magnetisation in the fixed frame (x, y, z):

$$\left(\frac{d\mathbf{M}}{dt}\right)_{rot} = \left(\frac{d\mathbf{M}}{dt}\right)_{fixed} - \boldsymbol{\omega} \times \mathbf{M} \tag{2.1.16}$$

where $\boldsymbol{\omega}$ is the frequency of the rotating frame.

$$\left(\frac{d\mathbf{M}}{dt}\right)_{fixed} = \gamma\mathbf{M} \times \mathbf{B} \tag{2.1.17}$$

and so:

$$\left(\frac{d\mathbf{M}}{dt}\right)_{rot} = \gamma\mathbf{M} \times \mathbf{B} + \frac{\omega}{\gamma} \tag{2.1.18}$$

At the Larmor frequency, $\omega = \omega_0 = \gamma/B_0$. Therefore, for a static magnetic field in the z direction, $\left(\frac{d\mathbf{M}}{dt}\right)_{rot} = 0$ and the magnetisation is constant. **M** precesses at the Larmor frequency with respect to the fixed frame.

The net magnetisation vector can be perturbed from its thermal equilibrium state by inputting energy in the form of a secondary magnetic field, B_1, which rotates at the Larmor frequency. This frequency corresponds to the radiofrequency (RF) energy on the electromagnetic spectrum and so this excitation is referred to as the RF excitation or RF pulse. Applying an oscillating RF excitation of frequency ω for time t displaces **M** from the z-axis by a flip angle θ:

$$\theta = \omega t = \gamma\mathbf{B}_1 t \tag{2.1.19}$$

$\mathbf{B}_1$ is the strength of the RF magnetic field, which is applied in the direction perpendicular to B_0. In the rotating frame, Equation 2.1.18 becomes:

$$\left(\frac{d\mathbf{M}}{dt}\right)_{rot} = \gamma\mathbf{M} \times \mathbf{B}_{eff} = \gamma\mathbf{M} \times \left[\left(B_0 - \frac{\omega}{\gamma}\right)\hat{z}' + B_1\hat{x}'\right] \tag{2.1.20}$$

In the fixed frame, the magnetisation precesses about both the static B_0 field and the rotating B_1 field. At $\omega = \omega_0$, $\mathbf{B}_{eff} = B_1\hat{x}'$ and **M** precesses about B_1 with frequency ω_1; the RF pulse is described as 'on-resonance'. When $\omega \neq \omega_0$, **M** precesses around B_{eff} with a frequency ω_{eff}; this is called 'off-resonance' behaviour.

The precession of **M** around the B_1 field as a result of RF pulse application is finite, and eventually **M** will return to its equilibrium state of alignment with $\mathbf{B}_0$ through longitudinal relaxation. This is analogous to damping and is caused by energy lost through interactions of the spin system with the surrounding nuclei. The change in magnetisation in the z direction is given by:

$$\frac{dM_z}{dt} = \frac{1}{T_1}(M_0 - M_z), \tag{2.1.21}$$

where T_1 is the characteristic time-scale of longitudinal relaxation. The solution to Equation 2.1.21 describes the exponential regrowth of the longitudinal magnetisation from an initial value, $M_z(0)$, to the equilibrium value, M_0:

$$M_z(t) = M_z(0)\mathrm{e}^{-t/T_1} + M_0(1 - \mathrm{e}^{-t/T_1}) \tag{2.1.22}$$

T_1 is the exponential time constant, representing the time after which the initial longitudinal magnetisation has recovered ~63% of the equilibrium value. Figure 2.1.3A shows the relaxation of the longitudinal magnetisation back to its initial value after excitation with 90° and 180° RF pulses.

In additional to longitudinal magnetisation relaxation, spin dephasing processes result in a reduction in net transverse magnetisation. Each spin experiences a slightly different magnetic field due to the effect of local magnetic fields created by neighbouring spins. This leads to different precessional frequencies and thus, in the rotating frame, the spin vectors disperse over time, as shown in Figure 2.1.3B.

The rate equation for transverse relaxation in the rotating frame is:

$$\frac{dM_{xy}}{dt} = -\frac{M_{xy}}{T_2}, \tag{2.1.23}$$

which is solved by:

$$M_{xy}(t) = M_{xy}(0)\mathrm{e}^{-t/T_2}, \tag{2.1.24}$$

where T_2 is the transverse relaxation time constant, equivalent to the time after which M_{xy} has decayed to ~37% of its initial value. B_0 field inhomogeneities and susceptibility differences also affect the effective magnetic field experienced by each spin and cause additional transverse relaxation with relaxation time T_2'. The effective relaxation time T_2^* is a combination of these two time constants:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{2.1.25}$$

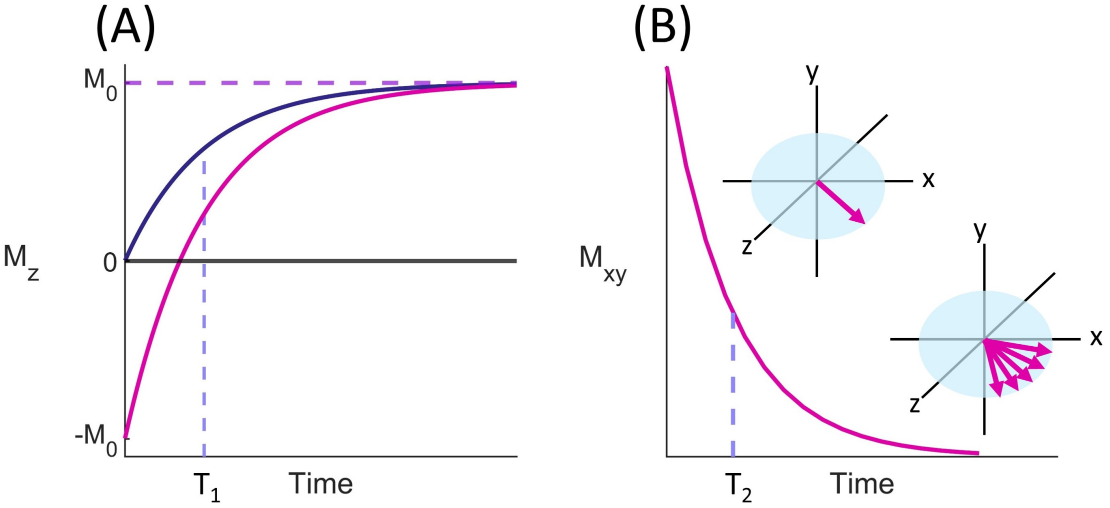

Figure 2.1.3: (A) The regrowth of the longitudinal magnetisation, M_z, from initial values of 0 (after the application of a 90° RF pulse) and -M_0 (after the application of a 180° RF pulse) back to the equilibrium value, M_0. T_1 is the time taken for M_z to reach 63% of M_0. (B) The decay of the transverse magnetisation, M_{xy}, from an initial value. T_2 is the time taken for M_{xy} to decay to 37% of its initial value. The transverse decay is due to the spin dephasing: over time, the spin vectors fan out, due to their different precessional frequencies.

Magnetic field induced dephasing is reversible through rephasing RF pulses, however the loss of coherence from spin-spin interactions is irreversible, because this is related to random, time-dependent field variations.

The magnetisation evolution can be described by a set of differential equations, known as the Bloch equations, which have the following form in the rotating frame:

$$\frac{dM_{x'}}{dt} = (\omega_0 - \omega)M_{y'} - \frac{1}{T_2}M_{x'} \tag{2.1.26}$$

$$\frac{dM_{y'}}{dt} = (-\omega_0 + \omega)M_{x'} + \omega_1 M_{z'} - \frac{1}{T_2}M_{y'} \tag{2.1.27}$$

$$\frac{dM_{z'}}{dt} = \omega_1 M_{y'} + \frac{1}{T_1}(M_0 - M_{z'}) \tag{2.1.28}$$

where ω_0 is the Larmor frequency, ω_1 is the RF field spin frequency and ω is the RF oscillation frequency. Assuming that the system was originally in equilibrium and was excited by a 90° pulse applied along the x-axis, such that $M_z(0) = M_x(0) = 0$ and $M_y(0) = M_0$, the Bloch equations can be solved by the following functions:

$$M_{x'}(t) = M_0 \mathrm{e}^{-t/T_2} \sin((\omega_0 - \omega_1)t) \tag{2.1.29}$$

$$M_{y'}(t) = M_0 \mathrm{e}^{-t/T_2} \cos((\omega_0 - \omega_1)t) \tag{2.1.30}$$

$$M_{z'}(t) = M_0\left(1 - \mathrm{e}^{-t/T_1}\right) \tag{2.1.31}$$

2.1.3 Signal Detection

Signal detection in MRI can be understood using the principles of electromagnetism and reciprocity [51]. The precession of the transverse magnetisation after the application of RF excitation creates a time-varying magnetic flux Φ through the receive coil of the MR machine:

$$\Phi = \int_{area} \boldsymbol{B} \cdot d\boldsymbol{S} \tag{2.1.32}$$

where $\boldsymbol{B}$ is the field generated by the magnetisation $\boldsymbol{M}$ of the sample and $d\boldsymbol{S}$ is a vector which is perpendicular to the coil area. According to Faraday's law, this will induce an electromotive force ε which is proportional to the rate of change of Φ but, in accordance with Lenz's law, in the opposite direction:

$$\varepsilon = -\frac{d\Phi}{dt} \tag{2.1.33}$$

Equation 2.1.33 can also be used to describe the magnetic field produced by current I flowing through a coil. Using this principle of reciprocity, Φ can be derived from the magnetic field per unit current that would be produced by the receive coil. This current normalised receive field is defined as:

$$\mathcal{B}^{receive}(\boldsymbol{r}) \equiv \frac{\boldsymbol{B}^{receive}(\boldsymbol{r})}{I} \tag{2.1.34}$$

The electromotive force induced in the MRI receive coil is given by:

$$\varepsilon = -\frac{d}{dt}\int_{sample} \boldsymbol{M}(\boldsymbol{r}, t) \cdot \mathcal{B}^{receive}(\boldsymbol{r}) d^3 r \tag{2.1.35}$$

This force induces motion in the electrons in the coil, creating a current known as the free induction decay (FID). This current oscillates at frequency ω_0 and decays at a rate controlled by T_2^*.

The FID signal is demodulated to remove the oscillations at the Larmor frequency, because these are typically too high for accurate conversion to digital form. This is done by multiplying the signal by a sinusoid or cosinusoid function with a frequency at or near the Larmor frequency, such that the resulting FID oscillates at $\Delta\omega = \omega_0 - \omega_{ref}$. The FID signal after excitation by an RF pulse is given by:

$$S(t) = S_0 \mathrm{e}^{i\Delta\omega t} \mathrm{e}^{-1/T_2^*} \tag{2.1.36}$$

where S_0 is the initial signal amplitude at time $t = 0$. The real and imaginary components of the signal are detected separately, by two different receiver channels.

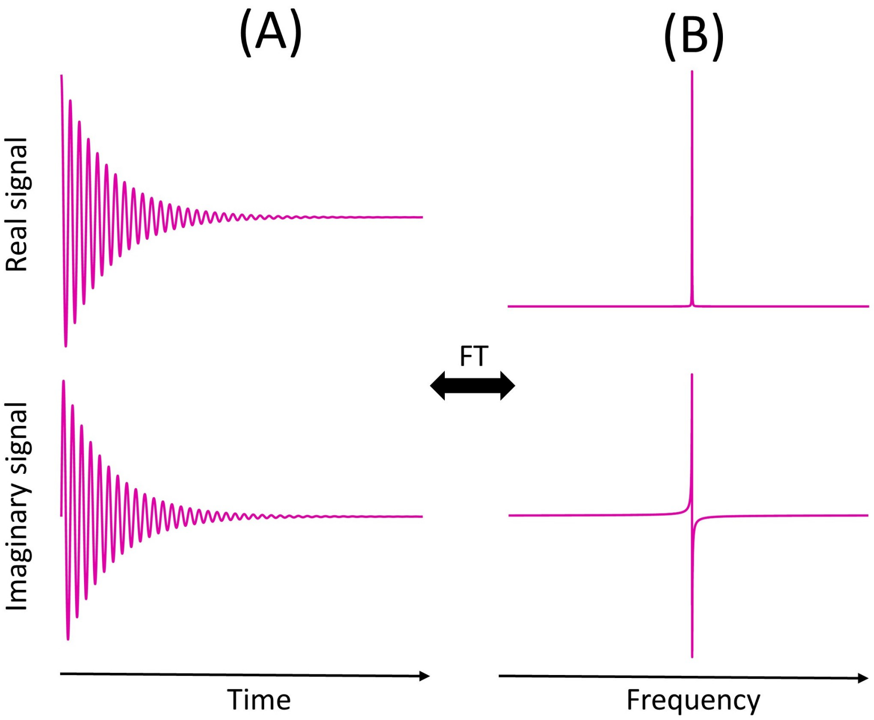

Figure 2.1.4: (A) An example time-domain free induction decay signal and (B) the Fourier transform (FT) of the signal, which gives the frequency domain spectrum, shown for both the real and imaginary parts.

Figure 2.1.4 shows the real and imaginary components of an example FID and its corresponding frequency-domain spectrum. These are related by the Fourier transform.

2.1.4 Chemical Shift

When a nucleus is placed in an external magnetic field, an electric current is induced in its surrounding electron cloud. This generates a magnetic field within the electron cloud which, in accordance with Lenz's law, opposes the direction of the external magnetic field. Therefore, the nucleus is 'shielded' from the full strength of the external field and experiences a reduced field, $\mathbf{B}_{loc}$:

$$\mathbf{B}_{loc} = \mathbf{B}_0 + \mathbf{B}_{induced} = \mathbf{B}_0(1-\sigma) \tag{2.1.37}$$

where σ is the shielding tensor. Shielding depends on the electronic configuration and therefore, $\mathbf{B}_{loc}$ is dependent on the local molecular environment of the nucleus. This leads to different resonant frequencies for the same nucleus in different chemical environments, which allows identification of nuclei and chemical configurations through MRS.

2.2 Spatial Encoding and Image Reconstruction

Applying magnetic field gradients along each axis introduces a spatially-dependent frequency and phase of the precessing nuclear spins. This allows spatial localisation by encoding this information into ‘k-space’, an empty matrix which is populated with the acquired data according to the gradient trajectory. The k-space data are then reconstructed using the inverse Fourier transform to obtain the MR image. In this section, we discuss the concepts of k-space, spatial encoding and image reconstruction.

2.2.1 Gradients and Echoes

In addition to the static B_0 field, magnetic field gradients are used in MRI for spatial localisation. The application of these gradients results in a position-dependent magnetic field strength and Larmor frequency. For example, for a gradient applied in the x direction:

$$B_z(x) = B_0 + G_x x \tag{2.2.1}$$

$$\omega_0(x) = \gamma B_0 + \gamma G_x x \tag{2.2.2}$$

where $G_x = \frac{\partial B_z}{\partial x}$ is the gradient strength in the x direction. This is illustrated in Figure 2.2.1.

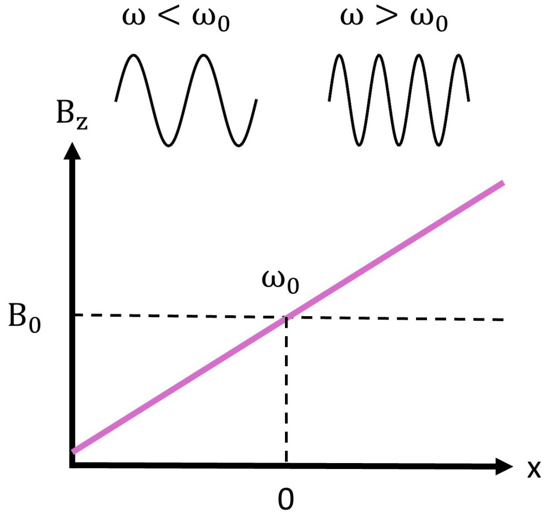

Figure 2.2.1: The application of a magnetic field gradient means that the spins experience a position-dependent magnetic field, which in turn induces a spatially varying resonant frequency.

When a gradient is applied, phase is accumulated by the transverse magnetisation:

$$\phi(x,t) = \gamma \boldsymbol{r} \int_0^t G(\boldsymbol{r},t')dt' \tag{2.2.3}$$

The phase accumulation remains encoded when the gradient is switched off. If the gradient polarity is reversed, this reverses the magnetic field gradient and spins which were previously precessing at a frequency $\omega < \omega_0$ will start to precess at a higher frequency $\omega > \omega_0$ and vice versa. This causes the spins to rephase and after a certain time called the echo time (TE), they are refocused, creating a signal 'echo' known as a gradient echo:

$$S_{GE} = S_0 \exp\left(-\frac{TE}{T_2^*}\right) \tag{2.2.4}$$

Echoes can also be formulated by RF pulses. After the application of a 90° RF pulse, the spins undergo transverse relaxation. Applying a 180° pulse on the $+y$ axis flips the spins by 180°, reversing their phases. The spins continue to dephase, but the direction of dephasing has been reversed. After a time TE, the spins refocus, creating a spin echo, with a signal given by:

$$S_{SE} = S_0 \exp\left(-\frac{TE}{T_2}\right) \tag{2.2.5}$$

The spin echo signal depends on the inherent transverse relaxation time, T_2, instead of T_2^*, because the 180° rephasing pulse compensates for magnetic field inhomogeneities.

2.2.2 k-Space

k-Space refers to the spatial frequency domain of a Fourier transform, and is the reciprocal of the spatial domain. Each location in k-space corresponds to a spatial frequency; a periodic variation in signal spatial distribution. The magnitude of each k-space point determines how much that spatial frequency is represented in the image. Figure 2.2.2A shows an MRI image of the brain and its corresponding 2D k-space representation. In MR images, low spatial frequencies contribute most to the signal and relate to the bulk structures, whereas high spatial frequencies relate to fine details and edges. The centre of k-space, $k = 0$, represents the average image intensity. These concepts are illustrated in Figure 2.2.2B and C.

To show how the MR signals are represented in k-space, we start by modifying the Bloch equations to include the effect of a magnetic field gradient $\boldsymbol{G}(t) = G_x(t)\hat{\boldsymbol{\imath}} + G_y\hat{\boldsymbol{\jmath}} + G_z\hat{\boldsymbol{z}}$. In matrix form:

$$\begin{bmatrix} \frac{dM_x}{dt} \\ \frac{dM_y}{dt} \\ \frac{dM_z}{dt} \end{bmatrix} = \begin{bmatrix} 0 & \gamma \boldsymbol{G}\cdot\boldsymbol{r} & -\gamma B_{1,y} \\ -\gamma \boldsymbol{G}\cdot\boldsymbol{r} & 0 & \gamma B_{1,x} \\ \gamma B_{1,y} & -\gamma B_{1,x} & 0 \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} + \begin{bmatrix} -\frac{1}{T_2} & 0 & 0 \\ 0 & -\frac{1}{T_2} & 0 \\ 0 & 0 & -\frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{1}{T_1} \end{bmatrix} M_0 \tag{2.2.6}$$

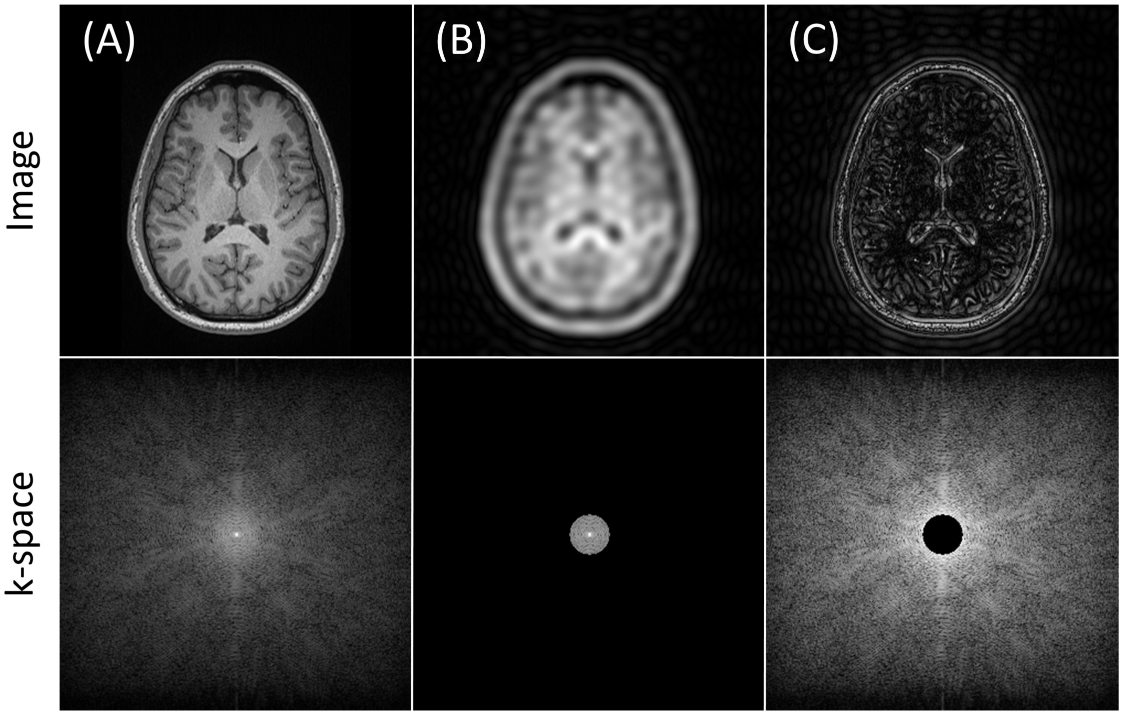

Figure 2.2.2: Example MR images of the brain (top row) and the corresponding k-space (bottom row). (A) Image reconstructed from all k-space data. (B) Reconstructing only the low-frequency k-space data results in an image which contains the gross structure of the image but lacks detail. (C) The high-frequency k-space data contains the fine structure of the image. This figure was created using the ‘k-Space Explorer’ app [52].

The Bloch equation for the transverse magnetisation, $M_{xy} = M_x + iM_y$, becomes:

$$\frac{dM_{xy}}{dt} = \left(-\frac{1}{T_2} - i\gamma \boldsymbol{G} \cdot \boldsymbol{r}\right) M_{xy} \tag{2.2.7}$$

which has the solution:

$$M_{xy}(\boldsymbol{r}, t) = M_{xy}(\boldsymbol{r}, 0) \exp\left(-\frac{t}{T_2}\right) \exp\left(-i\gamma \boldsymbol{r} \cdot \int_0^t \boldsymbol{G}(t')dt'\right) \tag{2.2.8}$$

Defining the k-space vector as:

$$\boldsymbol{k}(t) = \frac{\gamma}{2\pi} \int_0^t \boldsymbol{G}(t')dt', \tag{2.2.9}$$

Equation 2.2.8 can be rewritten:

$$M_{xy}(\boldsymbol{r}, t) = M_{xy}(\boldsymbol{r}, 0) \exp\left(-\frac{t}{T_2}\right) \exp(-i2\pi \boldsymbol{k}(t) \cdot \boldsymbol{r}) \tag{2.2.10}$$

Integrating the transverse magnetisation over the sample volume gives the MR signal:

$$S(t) = \int_{sample} M_{xy}(\boldsymbol{r}, 0) \exp\left(-\frac{t}{T_2}\right) \exp(-i2\pi \boldsymbol{k}(t) \cdot \boldsymbol{r}) d^3\boldsymbol{r} \tag{2.2.11}$$

Equation 2.2.11 is the Fourier transform of $M_{x,y}(\mathbf{r}, 0)$ with an additional relaxation term $\mathrm{e}^{-\frac{t}{T_2}}$, which acts like a damping term and imposes a convolutional k-space filter on the MR signal. This term does not affect the phase evolution or the frequency dispersion imparted to the spins with phase and frequency encoding. Hence, the inverse Fourier transform can be used to reconstruct the magnetisation from the MR signal that is nevertheless T_2-weighted. Upon Fourier transformation to the spatial frequency domain, the T_2 k-space filter is manifested as a broadening of the point spread function in the frequency direction (and also in the phase direction for single-shot sequences), following the convolution theorem.

2.2.2.1 Discrete Sampling

Although the Fourier transform of the k-space vector is a continuous function, in practise it can only be sampled discretely, with a sampling interval Δk between measurements. In the time domain, the spacing between samples is called the dwell time, t_d. The rate at which the samples are acquired is the receiver bandwidth (BW):

$$BW = \frac{1}{t_d} \tag{2.2.12}$$

The bandwidth is the full range of frequencies across an image (on GE Healthcare systems; other MRI machine manufacturers use the bandwidth per pixel) in the frequency encoding direction. The acquired signal is the product of the continuous signal ($S(k)$) and a comb function:

$$S_{acq}(k) = S(k) \cdot \mathrm{comb}(k), \tag{2.2.13}$$

where the comb function is equal to the sum of Dirac delta functions (δ) multiplied by the sampling interval:

$$\mathrm{comb}(k) = \Delta k \sum_n \delta(k - n\Delta k) \tag{2.2.14}$$

An example of discrete sampling of a sinc function is shown in Figure 2.2.3. For a continuous function to be accurately represented by the discrete sampling function, there exists a minimum sampling rate, called the Nyquist sampling rate. This is equal to twice the highest frequency of the continuous function. Sampling below this rate results in signal misrepresentation and 'aliasing' in image space: the parts of the image corresponding to the pixels near the edge of the object are 'folded back' inside of the field of view (FOV).

To reconstruct the discretised k-space data, the inverse fast Fourier transform is used, which is an algorithm for computing the inverse discrete Fourier transform.

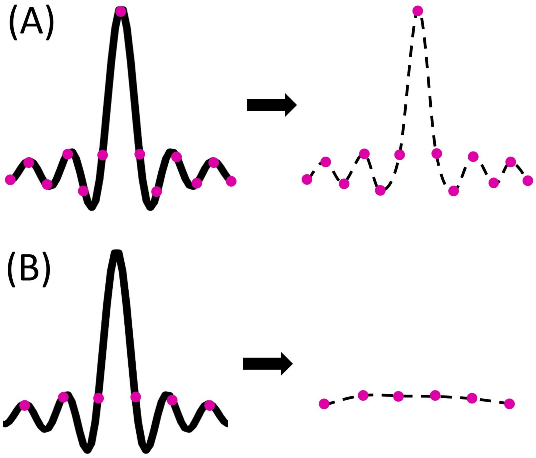

Figure 2.2.3: Discrete sampling of a sinc function: (A) with a sampling rate which is sufficient to reconstruct the continuous function and (B) with a sampling rate that is too low and leads to misrepresentation of the continuous function.

2.2.3 Cartesian k-Space Encoding

The most common way to acquire k-space data in MRI is with a Cartesian trajectory, which acquires the data row-by-row to fill a regularly spaced grid.

2.2.3.1 Frequency Encoding

Spatial localisation along one dimension (conventionally the x-axis) is achieved by 'frequency encoding'. A constant gradient is applied such that the Larmor frequency of the nuclei is linearly proportional to the location along the gradient direction. Figure 2.2.4 shows a typical pulse sequence diagram for frequency encoding in a gradient echo sequence. The gradient waveform has two lobes; a prephasing lobe and a readout lobe. The prephasing lobe initially dephases the spins such that an echo will form when the readout gradient is played out. In a gradient echo sequence, the lobes have opposite polarity, whereas in a spin echo sequence the lobes have the same polarity but are separated by a 180° pulse as was described in Section 2.2.1.

Data are acquired during an acquisition time T_{acq}, which is determined by the number of sampling points N_x and the dwell time:

$$T_{acq} = N_x t_d = \frac{N_x}{BW} \tag{2.2.15}$$

The k-space spacing is given by:

$$\Delta k_x = \frac{\gamma G_x t_d}{2\pi} \tag{2.2.16}$$

for a constant readout gradient.

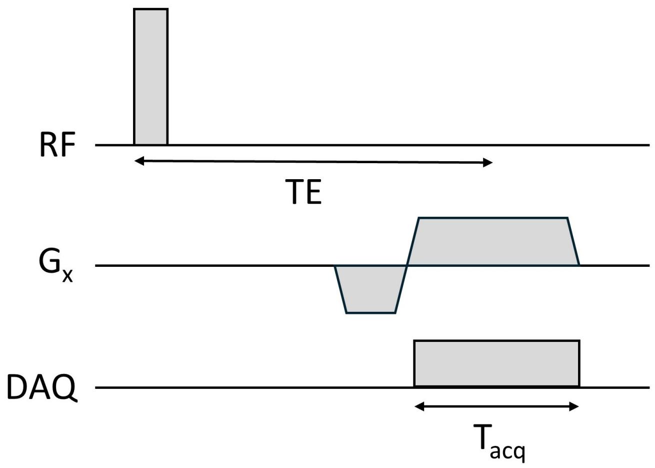

Figure 2.2.4: Pulse sequence diagram for frequency encoding. The RF pulse is used to excite the magnetisation and is followed by a frequency encoding gradient applied along the x axis (G_x). DAQ represents the data acquisition window, which has a duration T_{acq}.

This is equivalent to the maximum spatial extent in the spatial domain, otherwise known as the FOV, L_x:

$$\Delta k_x = \frac{1}{N_x \Delta x} = \frac{1}{L_x} \tag{2.2.17}$$

where Δx is the image resolution in the x direction. Therefore, by combining Equations 2.2.16 and 2.2.17, the amplitude of the readout gradient is found to be:

$$G_x = \frac{2\pi BW}{\gamma L_x} \tag{2.2.18}$$

2.2.3.2 Phase Encoding

Spatial encoding along the direction orthogonal to the frequency direction (typically the y direction) is achieved using phase encoding, whereby position-dependent phase differences are induced by a gradient. Unlike frequency encoding, where the spins return to their original precession frequency after the frequency encoding gradient is switched off, the phase differences accumulated over the duration of the phase encoding gradient are permanent. Additionally, whereas frequency encoding is able to sample the entire k_x k-space with one readout gradient, phase encoding must be repeated for each k_y data point. The Fourier transform cannot distinguish components with different phase but the same frequency and so to resolve N_y phase components, the same number of phase encoding steps are required. The amplitude of the phase encoding gradient is varied between each sequence repetition time (TR) and so the total time required to acquire 2D k-space data is N_yTR. Figure 2.2.5 shows a pulse sequence diagram for phase encoding in a gradient echo sequence.

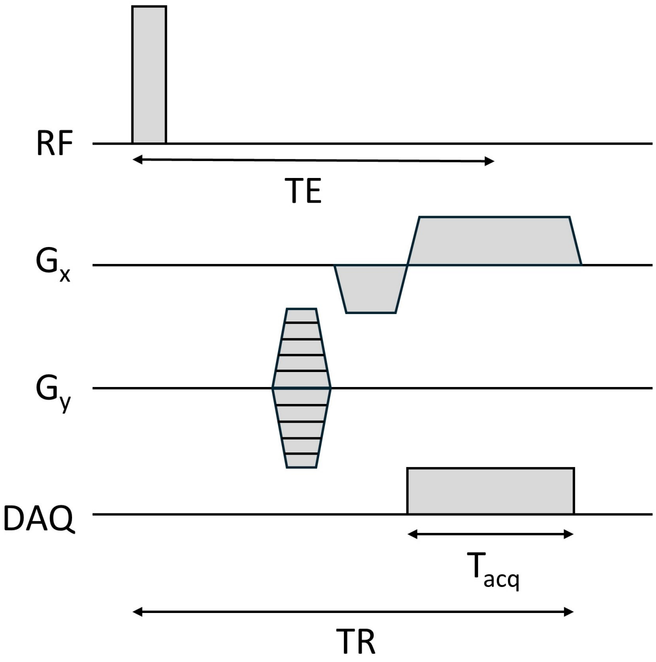

Figure 2.2.5: Pulse sequence diagram for frequency and phase encoding, which is achieved by repeating the k-space acquisition for several different phase encoding gradient amplitudes. Each repetition takes place over a time TR.

Phase encoding is symmetric about $k_y = 0$ and can either be performed sequentially, from $-k_{y,max}$ to $k_{y,max}$, or in a centric pattern, starting from $k_y = 0$ and moving outwards in each direction by alternating the polarity of the gradients. The k-space spacing is given by

$$\Delta k_y = \frac{1}{N_y \Delta y} = \frac{1}{L_y} \tag{2.2.19}$$

The maximum area of the phase encoding gradient is:

$$A_{y,max} = \int_0^t G_y(t')dt' = \frac{2\pi}{\gamma} k_{y,max} = \frac{\pi(N_y - 1)}{\gamma L_y} \tag{2.2.20}$$

To reduce TR and therefore the total sequence time, the time period of the phase encoding gradient should be minimised, which is achieved by using the maximum gradient amplitude and slew rate for $A_{y,max}$. For 3D acquisitions, phase encoding can be applied in both the y and z axes, however this increases the required time to $N_y N_z$TR. Alternatively, spatial localisation in the z direction can be achieved using slice selection.

2.2.3.3 Slice Selection

Instead of exciting the entire sample, a specific section or 'slice' can be excited by using a frequency-selective RF pulse in addition to a slice selection gradient. The simplest frequency-selective RF pulse is the sinc pulse ($\mathrm{sinc}(t) = \frac{\sin(t)}{t}$), which has a rectangular profile after Fourier transforming to the frequency domain. By applying this pulse at the same time as the plateau of a slice selection gradient, the range of frequencies excited (bandwidth) Δf, corresponds to a range of spatial locations Δz, which form the image slice (Figure 2.2.6). The plane of the slice is orthogonal to the gradient direction.

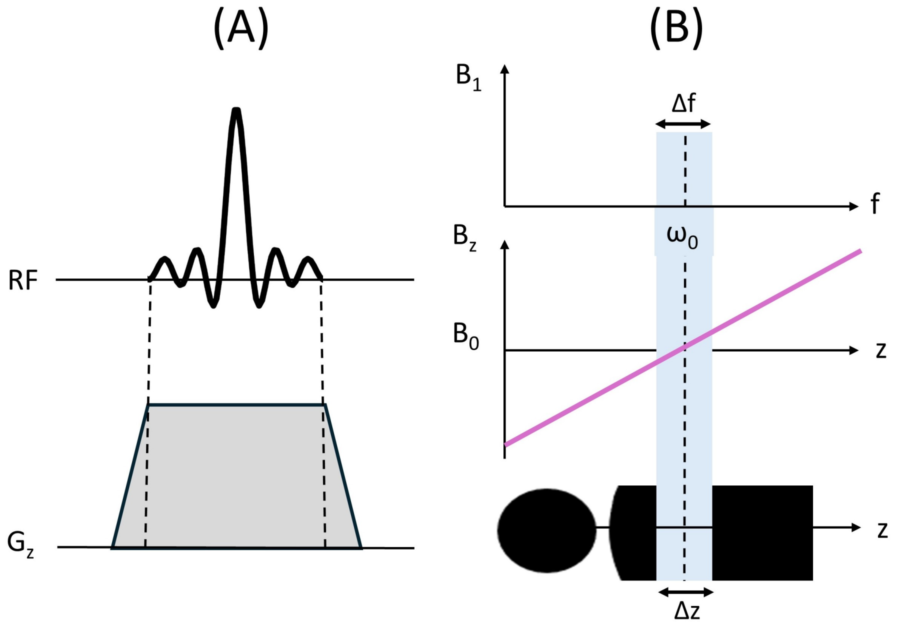

Figure 2.2.6: Slice selection: applying an RF pulse at the same time as a magnetic field gradient, as depicted in (A) results in the excitation of a slice of the body (B). The frequency bandwidth of the RF pulse (Δf) and the strength of the slice selective gradient determine the width of the image slice (Δz).

The position-dependent Larmor frequency can be expressed as:

$$f(z) = \frac{\gamma}{2\pi} G_z z \tag{2.2.21}$$

and the spatial range excited by the RF pulse is given by:

$$\Delta z = \frac{2\pi\Delta f}{\gamma G_z} \tag{2.2.22}$$

for a slice located at the gradient iso-centre. If the slice centre is offset by a distance δz from the iso-centre, the RF pulse frequency must also be offset from the Larmor frequency, by a factor of δf:

$$\delta z = \frac{2\pi\delta f}{\gamma G_z} \tag{2.2.23}$$

Therefore, the slice thickness is controlled by the strength of G_z and the bandwidth of the RF pulse.

The slice selection gradient is followed by a rephasing gradient lobe with an opposite polarity. This is needed to compensate for the dephasing that occurs across the image slice during the slice select gradient and causes signal losses. The area of the rephasing lobe is equal to half the area of the slice selecting lobe.

Full spatial coverage can be achieved using a combination of frequency, phase and slice selection encoding. This is referred to as a 2D sequence; 3D imaging involves phase encoding in two directions, with or without slice selection. The slice thickness is limited by the maximum gradient amplitude and RF bandwidth (Equation 2.2.22) in a 2D imaging acquisition. 3D acquisitions allow for thinner slices and better slice resolution because they avoid problems caused by imperfect frequency-selective RF pulse excitation profiles. However, 3D sequences tend to have increased acquisition time.

2.2.4 Signal-to-Noise Ratio

In MRI, noise arises from random variations in electrical currents originating from the MR hardware and conducting tissues of the body. The SNR is related to the image resolution ($\Delta x, \Delta y$ and Δz), number of k-space samples (N_x, N_y and N_z), number of signal averages (NSA; the number of times each k-space point is sampled) and the bandwidth. In 2D, the SNR per imaging volume (‘voxel’) is:

$$SNR_{2D} \propto \frac{\Delta x \Delta y \Delta z \sqrt{N_x N_y NSA}}{\sqrt{BW}} \tag{2.2.24}$$

In 3D, the SNR is increased by a factor of $\sqrt{N_z}$:

$$SNR_{3D} \propto \frac{\Delta x \Delta y \Delta z \sqrt{N_x N_y N_z NSA}}{\sqrt{BW}} \tag{2.2.25}$$

2.2.5 Non-Cartesian k-Space Encoding

In Section 2.2.3.2, the notion that k-space can be traversed in different ways was introduced, by comparing sequential (Figure 2.2.7A) and centric (Figure 2.2.7B) trajectories. The entire k-space grid can be acquired following one RF pulse by using a ‘zig-zag’ trajectory (Figure 2.2.7C), which is used in echo planar imaging (EPI) and is achieved with a frequency encoding gradient of alternating polarity and intermittent low-amplitude phase encoding gradient pulses, which provide a step-wise increase along the k_y axis. More efficient k-space sampling can be achieved with non-Cartesian encoding trajectories; trajectories where the acquired k-space points do not lie on a rectangular matrix. In this section we discuss two of the most popular non-Cartesian encoding schemes, radial (Figure 2.2.7D) and spiral (Figure 2.2.7E), and the procedure for reconstructing the resulting k-space data.

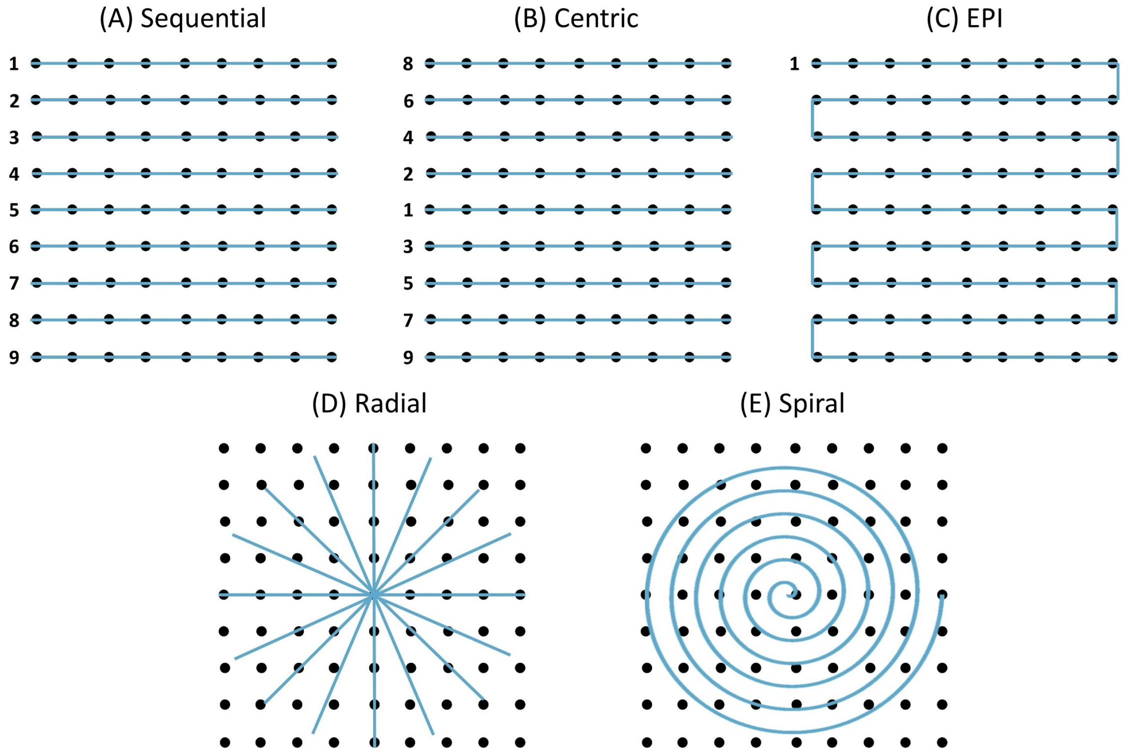

Figure 2.2.7: (A-C) Cartesian and (D-E) non-Cartesian k-space encoding trajectories. For the Cartesian trajectories, the numbers indicate the order of acquisition.

2.2.5.1 Radial Imaging

In radial encoding, frequency encoding gradients are applied along multiple directions simultaneously (x and y for 2D, x, y and z for 3D). A pulse sequence diagram for 3D radial imaging is shown in Figure 2.2.8A. For 2D radial imaging, this results in a gradient vector $\boldsymbol{G}$ which has strength:

$$G = \sqrt{G_x^2 + G_y^2} \tag{2.2.26}$$

and points along an angle θ:

$$\theta = \tan^{-1}\left(\frac{G_y}{G_x}\right) \tag{2.2.27}$$

Using the definition of the k-space vector in Equation 2.2.9, analogous equations can be used for k:

$$k = \sqrt{k_x^2 + k_y^2} \tag{2.2.28}$$

$$\theta = \tan^{-1}\left(\frac{k_y}{k_x}\right) \tag{2.2.29}$$

such that

$$k_x = k\cos\theta \tag{2.2.30}$$

$$k_y = k\sin\theta \tag{2.2.31}$$

The vector k can be thought of as a radial ‘spoke’ which originates at either the k-space origin (centre-out) or edge of k-space (diameter) and terminates at the opposing edge of k-space. The centre of k-space is inherently oversampled, which make radial sequences robust to motion artefacts.

The angular separation between two consecutive k-space spokes ($\Delta\theta$) and the separation between the samples along the spokes (Δk_r) are limited by the Nyquist sampling limit. The maximum angular separation to fulfil this condition is:

$$\Delta\theta = \frac{1}{k_{max}L} \tag{2.2.32}$$

where L is the FOV. Similarly, the maximum sample separation is:

$$\Delta k_r = \frac{1}{L} \tag{2.2.33}$$

Therefore, the minimum required number of radial spokes and samples are:

$$n_\theta = \frac{\pi}{\Delta\theta} \geq \pi k_{max}L \tag{2.2.34}$$

$$n_r = \frac{2k_{max}}{\Delta k_r} \geq 2k_{max}L \tag{2.2.35}$$

Whereas for Cartesian imaging, sampling at a rate below the Nyquist limit results in aliasing, for radial imaging, this manifests as radial streaking artefacts.

2.2.5.2 Spiral Imaging

Time-varying gradients can be applied in the x and y directions to create a spiral trajectory (Figure 2.2.8B). A popular choice is the Archimedean spiral, where the radius is defined by:

$$r = a\theta \tag{2.2.36}$$

in polar coordinates, where a is a constant. Spiral encoding provides efficient k-space coverage and short echo times. It can be implemented in a single-shot or interleaved multi-shot fashion. For multi-shot acquisitions, several spiral interleaves are acquired, with each shot rotated by an angle of $\pm 2\pi/N_{shot}$ from the previous shot, where N_{shot} is the number of interleaves.

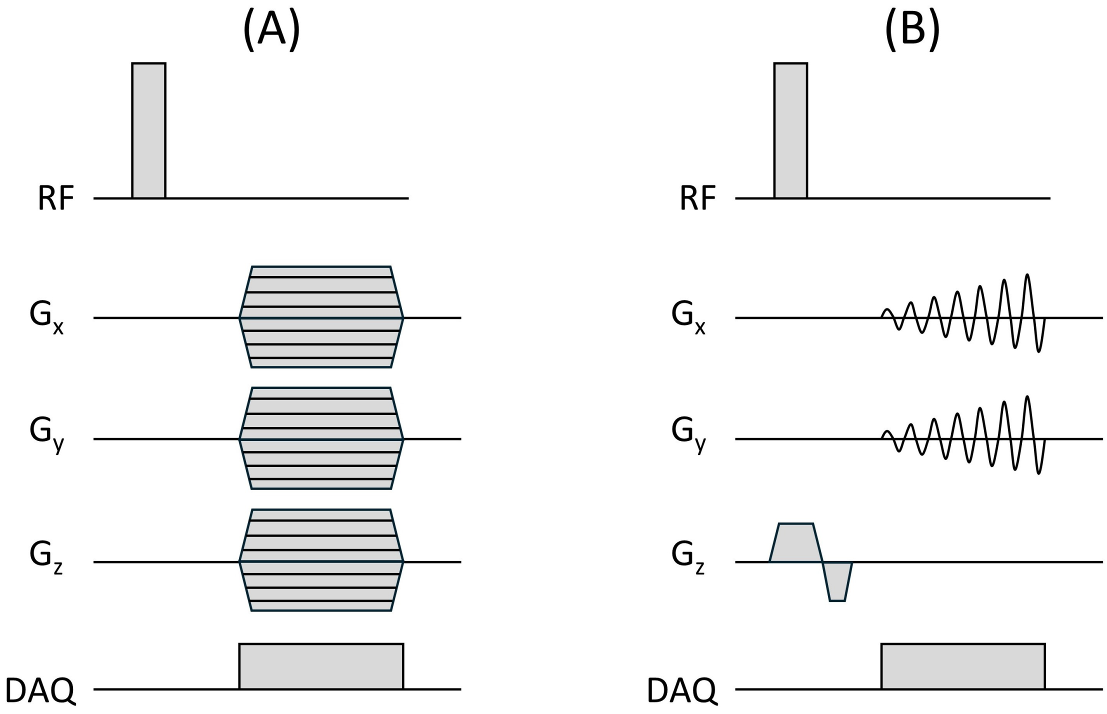

Figure 2.2.8: Pulse sequence diagrams for (A) 3D radial and (B) stack-of-spirals (2D spiral with slice selection) imaging. The gradients in the radial pulse sequence diagram have been drawn to represent their varying amplitudes (not phase-encoding).

The spirals originate from the centre of k-space and trace out a path defined by:

$$k_x = \frac{N_{shot}}{2\pi L}\theta \sin\theta \tag{2.2.37}$$

$$k_y = \frac{N_{shot}}{2\pi L}\theta \cos\theta \tag{2.2.38}$$

The gradient waveforms are:

$$G_x = \frac{N_{shot}}{\gamma L}\frac{d\theta}{dt}(\sin\theta + \theta\cos\theta) \tag{2.2.39}$$

$$G_y = \frac{N_{shot}}{\gamma L}\frac{d\theta}{dt}(\cos\theta + \theta\sin\theta) \tag{2.2.40}$$

A limitation of spiral acquisitions is that they are susceptible to image blurring as a result of off-resonance effects. B_0 inhomogeneities, susceptibility effects and chemical shifts cause variations in the Larmor frequency. Whilst for Cartesian imaging, these frequency offsets result in a chemical shift artefact, in spiral imaging this causes blurring because the trajectory is changing in both the x and y directions simultaneously. It is possible to correct for some of this effect by mapping the B_0 inhomogeneity prior to imaging and incorporating this into the reconstruction algorithm.

2.2.6 Spectral Encoding

In Section 2.1.4, the concept of chemical shift was introduced. Nuclei in different chemical environments resonate at different frequencies, which are superimposed in the FID signal and can be decoded using the Fourier transform. The spectral peaks revealed in the Fourier transformed signal allow for quantification of chemical concentrations. Spatial localisation of the spectral signals can be achieved using phase encoding. In conventional chemical shift imaging (CSI), phase encoding gradients are used to modulate the signal phase and amplitude before detecting the FID. Frequency encoding is generally not used during the signal detection period so that there is complete separation of spatial and spectral encoding; otherwise it becomes difficult to disentangle the spatial and chemical shift information [53]. CSI can be implemented using different excitation and signal generation methods, for example FID or spin echo sampling. A limitation of CSI is the long imaging time. For 3D CSI, phase encoding is required in three directions, resulting in an scan time of $\mathrm{TR} \cdot N_x N_y N_z$, where N is the matrix size in each direction.

2.2.6.1 Echo Planar Spectroscopic Imaging

An alternative method for performing both spectral and spatial encoding is echo planar spectroscopic imaging (EPSI) [53]. In this technique, spectral information and one spatial dimension are encoded simultaneously in a single readout, reducing the imaging time. Conventional phase encoding can be used in the other two spatial dimensions. An example of an EPSI pulse sequence is shown in Figure 2.2.9. A frequency encoding readout gradient with rapidly alternating polarity is used, which results in the zig-zag k-space trajectory shown in Figure 2.2.7C and creates a train of gradient echoes. The signal is sampled at the gradient plateaus, which are separated by the dwell time, t_d. This was defined in Equation 2.2.12 and determines the spectral bandwidth. t_d must be short enough to capture the full spectral bandwidth and avoid aliasing. The spectral resolution (Δf) is determined by the inverse of the acquisition time, which must be long enough to resolve all spectral peaks, whilst also minimising T_2^* signal decay [54]. The spectral and spatial information are convolved and so EPSI reconstruction involves unfolding the chemical shift aliasing using chemical shift modelling.

In practise, the echoes from the positive and negative gradients are not equivalent, because of asymmetries in gradient switching and signal distortions from eddy currents. To avoid ghosting artefacts, the odd and even echoes should be reconstructed separately, or these issues can be avoided by using a flyback gradient scheme, where data are only acquired for the positive gradients. Short gradient pulses with negative polarity and high slew rate are used for refocusing.

The decrease in scan time achieved with EPSI comes at the cost of decreased SNR compared to CSI, because the increase in the acquisition bandwidth introduces noise which is proportional to $\sim \sqrt{BW}$. The SNR of an EPSI experiment can be improved by using signal averaging [53].

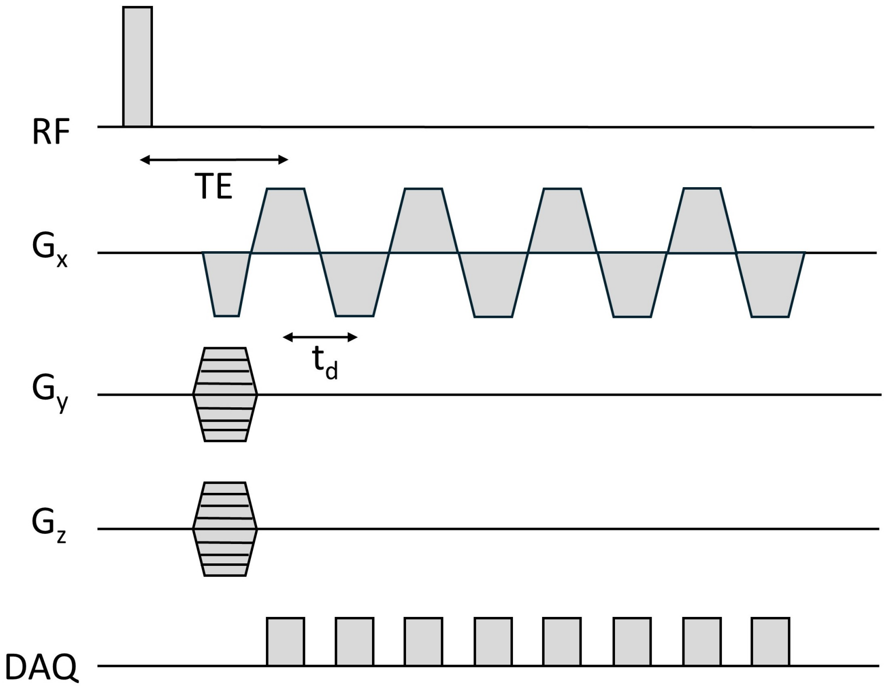

Figure 2.2.9: EPSI pulse sequence diagram: an oscillating trapezoid readout gradient (G_x) with a period of t_d is used to produce a train of gradient echoes. Phase encoding is applied in the y and z directions. Data is acquired during the readout gradient plateaus.

2.2.7 Reconstruction of Non-Cartesian k-Space Data

In conventional Cartesian MRI, images are reconstructed using the inverse discrete Fourier transform. However, non-Cartesian data are non-uniform and so must first be resampled to a Cartesian grid prior to Fourier transformation. The most popular way to resample the data is with a 'gridding' algorithm [55]. The principle of gridding is presented in Figure 2.2.10; each data point on the non-Cartesian k-space trajectory is convolved with a gridding kernel, then mapped onto a Cartesian grid, before being reconstructed with an inverse discrete Fourier transform.

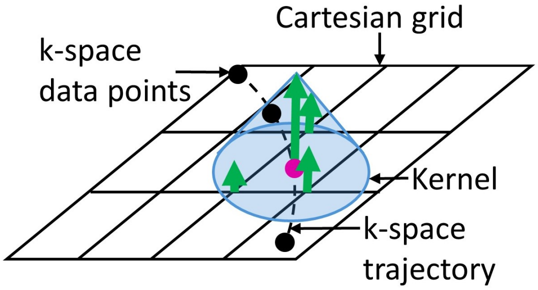

Figure 2.2.10: Schematic to illustrate the principle of gridding. The dashed line shows the k-space trajectory on which the data (circles) are sampled. Each data point is convolved with a gridding kernel (blue cone), which is shown here for the pink data point, and the convolution evaluated at the adjacent Cartesian grid points (green arrows).

Mathematically, this process can be written:

$$M_C(k_x, k_y) = [(M(k_x, k_y)S(k_x, k_y)) \otimes C(k_x, k_y)] \times R(k_x, k_y) \qquad (2.2.41)$$

where ⊗ represents convolution, $M_C(k_x, k_y)$ is the regridded Cartesian data, $M(k_x, k_y)$ is the original non-Cartesian data, $S(k_x, k_y)$ is a sampling function:

$$S(k_x, k_y) = \sum_i \delta(k_x - k_{x,i}, k_y - k_{y,i}) \qquad (2.2.42)$$

$C(k_x, k_y)$ is the gridding kernel and $R(k_x, k_y)$ is the Cartesian grid:

$$R(k_x, k_y) = \sum_i \sum_j \delta(k_x - i, k_y - j) \qquad (2.2.43)$$

The Fourier transform of Equation 2.2.41 is:

$$m_C(x, y) = [(m(x, y) \otimes s(x, y))c(x, y)] \otimes r(x, y) \qquad (2.2.44)$$

The gridding kernel acts as an apodization function in the spatial domain. The ideal $c(x, y)$ apodization function would be a rectangular function, with extents $\pm\frac{L}{2}$, where L is the FOV. However, the Fourier transform of a rectangular function is a sinc function with infinite extents, which is not computationally possible. Instead, a windowed sinc function can be used for the gridding kernel, but this will result in some loss of FOV from increased apodization and side lobes that are aliased back into the image. Another possibility is the Kaiser-Bessel gridding kernel, which was found to be preferable by Jackson et al [56] because it is easy to compute and minimises aliasing [57]. Both apodization and aliasing can be reduced significantly by choosing to resample to a Cartesian grid with greater density than the underlying k-space data (in other words, by oversampling). Apodization can also be corrected for by dividing the image by $c(x, y)$.

Non-Cartesian sampling schemes use non-uniform sampling density, which must be compensated for before gridding to achieve accurate interpolation and reduce reconstruction error. This is corrected for by multiplying the sampling function, $S(k_x, k_y)$ by a density compensation function (DCF). For simple k-space trajectories the DCF can be calculated analytically using the pattern's geometry. For example, for uniform radial sampling with N spokes, the DCF is given by:

$$DCF = \begin{cases} \frac{2\pi}{N}(\Delta k_r)^2 \frac{1}{8} & \text{if } (k_x, k_y) = 0 \\ \frac{2\pi}{N}(\Delta k_r)^2 n & \text{otherwise} \end{cases} \qquad (2.2.45)$$

where n is an integer, denoting the number of Δk_r intervals that a sample lies from the origin.

For more complicated trajectories, numerical approaches can be used to estimate the DCF by calculating the area associated with each data sample, for example with the Voronoi algorithm [58].

2.2.8 Compressed Sensing

In this section, an alternative image reconstruction method is described: 'compressed sensing' (CS). Most natural images are compressible, meaning that they can be described using far fewer samples than their number of pixels, without significant loss of information. The principle of CS is that, since accurate image reconstruction can be achieved without these additional data, they do not need to be acquired. In MRI, CS can be used to reconstruct images using a small subset of k-space, making it possible to decrease image acquisition time and reduce the number of RF pulses needed. This is useful for a wide range of MRI applications, such as dynamic contrast enhanced (DCE) MRI, angiography and cardiac imaging [59]. CS is well-suited to hyperpolarised gas lung imaging, where image acceleration techniques are particularly useful since most imaging takes place during breath hold.

There are three main conditions for CS reconstruction, which are described in more detail in the following sections. The first condition is that the image data must be sparse in the image domain or a known transform domain. Secondly, in the sparse domain, the artefacts from undersampling must be incoherent; this is achieved with random undersampling. Finally, a non-linear reconstruction must be used to impose sparsity and data consistency. Figure 2.2.11 illustrates these concepts.

2.2.8.1 Sparsity

A sparse image is one which has relatively few non-zero coefficients in a mathematical transform domain. Some MR images, such as angiograms, are naturally sparse in the image domain, i.e. they have few non-zero valued pixels. More complicated images may not be sparse in the image domain but can be made sparse by applying a sparsifying transform. Examples of such transforms include the discrete cosine transform, which is used in JPEG compression of natural images, the wavelet transformation and the finite-difference transform. In Figure 2.2.11, a wavelet transform is used as the sparse transform.

2.2.8.2 Random Undersampling

Reconstructing undersampled k-space data leads to undersampling artefacts which impair image quality. Uniform undersampling can lead to coherent artefacts, such as aliasing, but with random undersampling, the artefacts are incoherent and resemble random noise. These can be removed using thresholding or smoothing algorithms. Truly random k-space sampling is difficult to implement and may be impractical due to hardware limitations (since it requires rapid gradient switching). Also, since most of the image contrast is contained within the low-frequency k-space data, it can be beneficial to acquire more samples from the central portion of k-space, in order to improve the image SNR. Variable density pseudo-random trajectories which undersample the edges of k-space more than the centre can be used to maximise incoherence and SNR [60].

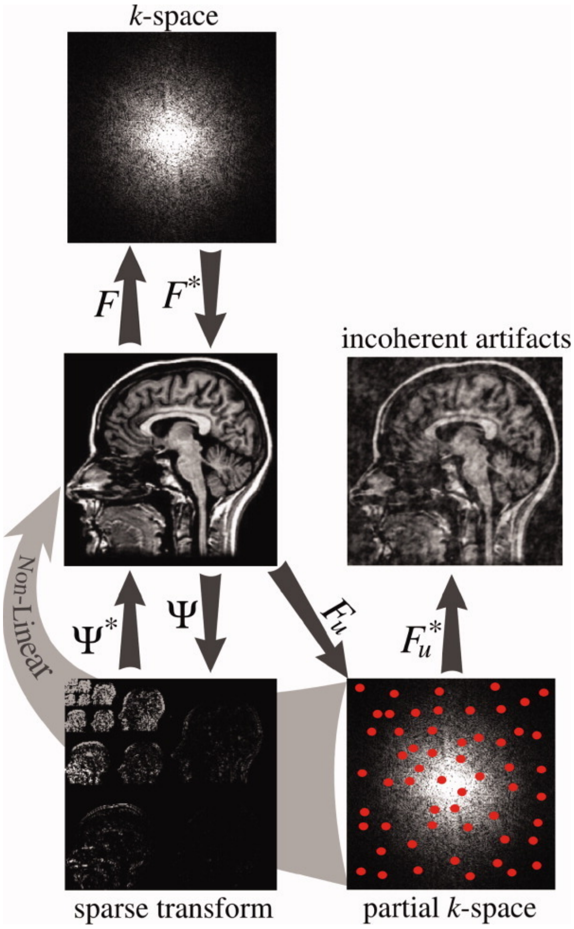

Figure 2.2.11: Illustration of the main concepts of CS. Reconstructing the randomly undersampled k-space data using the inverse Fourier transform (F^*) results in incoherent undersampling artefacts. Transforming these data to a sparse representation and then applying a non-linear reconstruction yields an image which has good agreement with the fully-sampled image and does not contain undersampling artefacts. Copyright 2007 Wiley, reproduced with permission from [60].

2.2.8.3 Non-Linear Reconstruction

The final component of CS is the non-linear reconstruction. After transforming the image to the sparse representation and applying a denoising algorithm, the image is transformed back to k-space and compared to the original acquired k-space with a cost function. This iterative process is repeated until the final image is reached, which should have good fidelity with the original but minimal undersampling artefacts.

This reconstruction can be written as the following constrained optimisation problem:

$$\text{minimise} \quad ||\Psi x||_1 \tag{2.2.46}$$

$$\text{subject to} \quad ||F_u x - y||_2 < \epsilon \tag{2.2.47}$$

where Ψ is the sparsifying transform, x is the reconstructed image, F_u is the undersampled Fourier transform, y is the acquired k-space data and ϵ is a parameter that controls the fidelity of the reconstructed image to the acquired data. The objective function (minimising $||\Psi x||_1$) promotes sparsity by using the l_1 norm, because this places a larger penalty on having many small coefficients than having a few large coefficients. Data consistency is enforced by the condition $||F_u x - y||_2 < \epsilon$, with ϵ usually set to below the expected noise level.

Equation 2.2.46 can also be expressed as an unconstrained problem:

$$x = \text{argmin}_x ||F_u x - y||_2 + \lambda_1 ||\Psi x||_1 \tag{2.2.48}$$

where λ_1 is the regularisation weighting, which balances data sparsity and fidelity. Other sparsifying transforms can be used for further regularisation by adding additional $\lambda||\Psi x||_1$ terms. There are several methods for solving Equation 2.2.48. Two of the most common methods are the Alternating Direction Method of Multipliers (ADMM) algorithm [61] and the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [62].

2.3 Principles of Hyperpolarised ^{129}Xe Lung MRI

2.3.1 Properties of ^{129}Xe

Xe is a noble gas occurring naturally in air at 0.087 ppm, with a ^{129}Xe isotope abundance of 26.44% [63]. ^{129}Xe is suitable as an MRI contrast media for several reasons; it is inert, non-toxic, well-tolerated by healthy and diseased subjects [64] and has a nuclear spin of $\frac{1}{2}$ [65]. These features are also shared by helium-3 (^{3}He), which used to be more commonly used for MRI than ^{129}Xe [66], but has become prohibitively expensive to acquire leading to an almost complete shift to ^{129}Xe for in vivo applications in the last decade [67]. Some magnetic properties of ^{129}Xe and ^{3}He are given in Table 2.3.1, in comparison to ^{1}H.

Table 2.3.1: Properties of ^{1}H, ^{3}He and ^{129}Xe. Data from [68].

Property	^{1}H	^{3}He	^{129}Xe
Gyromagnetic ratio, γ (MHzT^{-1})	42.58	-32.43	-11.78
Nuclear spin, I	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
Magnetic moment, μ/μ_N	4.84	-3.69	-0.78

A further advantage of ^{129}Xe over ^{3}He, which is particularly advantageous for functional MRI, is its solubility in various somatic compounds. ^{129}Xe dissolves in the pulmonary parenchymal tissue, blood plasma and RBCs, with Ostwald solubility coefficients $\sim$ 0.1, 0.09 and 0.2 respectively [63]. Due to its large and easily distorted electron cloud, ^{129}Xe is a sensitive probe of its local chemical environment and displays different chemical shifts depending on its surroundings. A summary of ^{129}Xe Ostwald solubility coefficients and chemical shifts in biologically-relevant solvents is given in Table 2.3.2.

When ^{129}Xe gas is inhaled, a small amount passes from the alveolus through the tissue barrier, dissolving into the capillary plasma and entering the RBCs (Figure

Table 2.3.2: ^{129}Xe Ostwald solubility coefficients and corresponding chemical shifts (expressed in parts per million (ppm) from the gas-phase resonance at 0 ppm) in somatic substances. Perfluorooctyl bromide emulsion is used as a blood substitute.

Solvent	Ostwald Solubility Coefficient	Chemical Shift (ppm)
Blood plasma	0.103 [69] 0.091 [71] 0.094 [73]	194 [70] 197 [72]
Distilled water	0.083 [74, 71]	190 [70]
Erythrocytes (RBCs)	0.20 [69] 0.19 [71] 0.27 [73]	216 [70] 222 [72]
Olive oil	1.83 [73] 1.79 [71]	198 [75]
Perfluorooctyl bromide emulsion	1.20 [76]	106 [76]
Rat adipose tissue	1.715 [73]	191 [76] [77]
Saline (0.9% NaCl)	0.093 [69] 0.078 [74]	194 [78]

2.3.1A). This process is driven by diffusion, from the concentration gradient between the gas space and capillary blood [79]. The different ^{129}Xe environments in the lung lead to a number of separate spectral peaks, as shown in Figure 2.3.1B. The large gas-phase peak is used as the reference, at a chemical shift of 0 ppm by definition. The other two peaks are from ^{129}Xe dissolved in the lung tissue and blood plasma compartments at 197 ppm (the membrane resonance) and the RBC compartment at 218 ppm. There is a continuous chemical exchange of ^{129}Xe between the RBCs and the capillary plasma, with a chemical exchange rate constant of ∼12 ms [80], which leads to exchange-broadened peaks.

^{129}Xe has a T_1 which is of the order of hours in a stable magnetic field and oxygen-free environment, but around 20 s in the lungs due to dipole-dipole coupling between the ^{129}Xe and paramagnetic oxygen [81]. In the blood, ^{129}Xe has a T_1 of a few seconds; the value of which is dependent on the blood oxygenation. Norquay et al. measured T_1 = 2.2 s in deoxygenated blood (oxygen saturation, sO_2, = 0.06) and T_1 = 7.8 s in fully oxygenated blood (sO_2 = 1.00) [82]. This relatively long T_1 has allowed inhaled ^{129}Xe which has dissolved into the circulating blood to be imaged in the brain [83, 84, 85] and kidneys [86]. The susceptibility gradients that occur as a result of the many air-tissue boundaries in the lung result in a very short ^{129}Xe T_2^* values. For ^{129}Xe in the alveolar airspaces, T_2^* is ∼25 ms for partially inflated lungs and ∼50 ms at total lung capacity at 1.5 T [87]. The T_2^* is even shorter for dissolved-phase ^{129}Xe, as a result of chemical exchange broadening: T_2^* ∼2.2 ms at 1.5 T [36] and ∼1.1 ms at 3 T [88].

The ^{129}Xe atoms form weak Van der Waals bonds with haemoglobin when dissolved in blood. The level of blood oxygenation affects the chemical shift of the ^{129}Xe RBC peak, which has been shown to increase non-linearly with sO_2 level by Wolber et al. [89] and by Norquay et al. [90]. Norquay et al. observed a ∼ 5 ppm difference in ^{129}Xe chemical shift between fully deoxygenated and fully

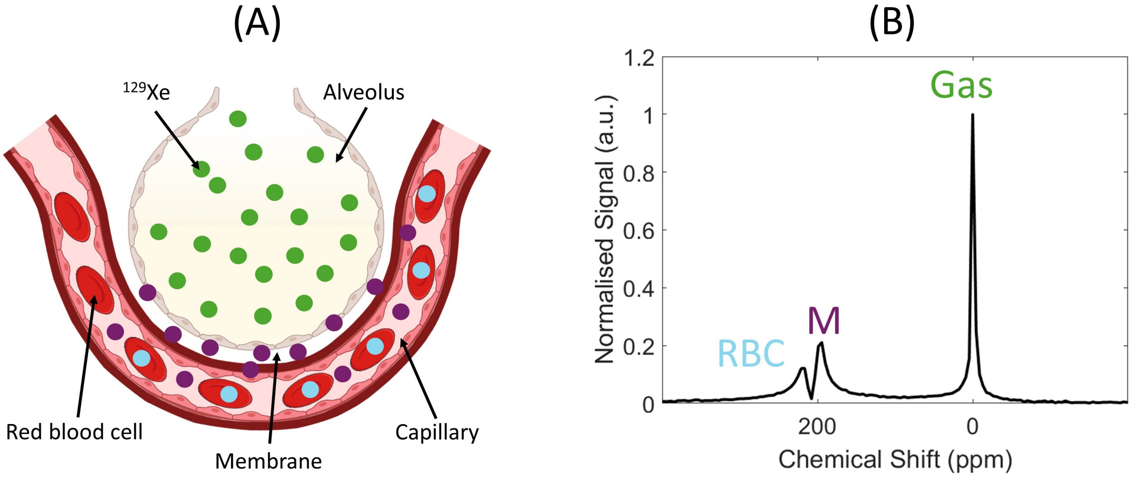

Figure 2.3.1: (A) Inhaled ^{129}Xe atoms (green) enter the alveoli, where they diffuse with random Brownian motion. A small proportion of ^{129}Xe atoms dissolve into the alveolar tissue membrane and capillary blood plasma (purple) and RBCs (blue). (B) This results in three distinct spectroscopic peaks, from the ^{129}Xe in the alveolar airspace (0 ppm), membrane and plasma (197 ppm) and RBCs (218 ppm). This image was partly generated using Biorender.

oxygenated blood and have demonstrated the feasibility of using this shift to measure pulmonary blood oxygenation. The non-linear relationship between chemical shift and sO_2 has been postulated to be because of an oxygen-dependent haemoglobin affinity, driven by conformational changes brought about by oxygen binding [89].

^{129}Xe has a self-diffusion coefficient of D_0 = 0.06 cm^{2}s^{-1} [91], a diffusion coefficient in air of D_a = 0.14 cm^{2}s^{-1} [92] and of D_g = 0.03 - 0.04 cm^{2}s^{-1} in the airspace of a healthy lung [30, 27], but the diffusion coefficient in lung tissue and blood is not well known. The currently accepted value of the diffusion coefficient of ^{129}Xe dissolved in lung tissue and blood is 3.3 x 10^{-6} cm^{2}s^{-1}. However, this comes from a single animal study [93] which has not yet been validated either in humans or in vitro.

2.3.2 Spin Exchange Optical Pumping

The low gyromagnetic ratio and spin density of Xe make the NMR signal of thermally polarised ^{129}Xe too low for in vivo lung imaging at clinical field strengths. To overcome this, a process called spin exchange optical pumping (SEOP) can be used to increase the ^{129}Xe nuclear spin polarisation (and MR signal) by up to five orders of magnitude compared to the thermal Boltzmann polarisation defined in Equation 2.1.11.

Although other methods exist for generating hyperpolarised ^{129}Xe, such as Dynamic Nuclear Polarisation (DNP) [94], here we focus solely on SEOP, which is the most commonly used technique. In the first stage of SEOP, vaporised alkali metal atoms are optically pumped with resonant laser light to polarise the valence electron spins. Figure 2.3.2A shows the atomic energy level diagram for rubidium (Rb), the alkali metal typically used for ^{129}Xe SEOP, due to its high vapour pressure which allows for lower operating temperatures [95]. Alkali metals have a single valence electron which occupies the $^2S_{\frac{1}{2}}$ energy sublevel, with orbital angular momentum $l = 0$. When in the presence of a magnetic field, the energy levels undergo Zeeman splitting, whereby the S and P energy levels are additionally split into two sublevels with spin angular momentum $m = \pm\frac{1}{2}$ (S) and $m = \pm\frac{3}{2}$ (P). By exciting the ground state electrons with left-circularly (σ^+) polarised photons resonant with the D_1 transition ($^2S_{\frac{1}{2}}$ to the lowest $^2P_{\frac{1}{2}}$ excited state), spin angular momentum is transferred to the valence electrons. For Rb, a laser wavelength of 794.77 nm is used to drive this transition, which leads to a change in the spin angular momentum quantum number of one unit, $\Delta m = +1$. According to quantum selection rules, the excited electrons can only relax into the $m = +\frac{1}{2}$ energy sublevel, leading to an accumulation of electrons in this ground state quantum level. This process is illustrated in Figure 2.3.2B.

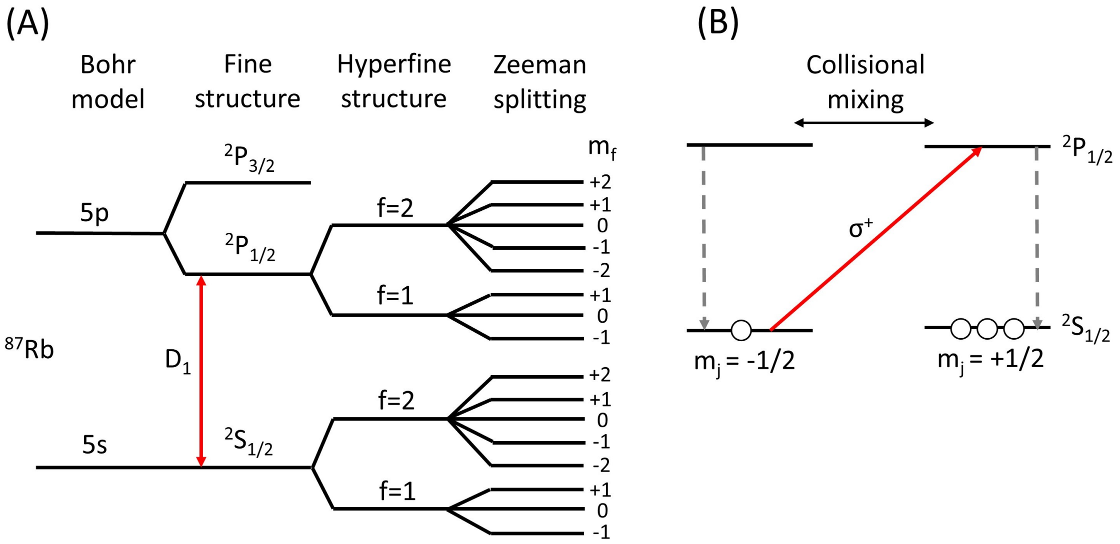

Figure 2.3.2: (A) ^{87}Rb energy levels, including the Zeeman splitting that occurs in the presence of a magnetic field. (B) Optical pumping of Rb with circularly polarised resonant light. 'Collisional mixing' refers to the process whereby spin angular momentum is redistributed between the ground state energy levels as a result of collisions between the excited Rb atoms.

The second stage of SEOP is spin exchange, whereby the electronic Rb spin polarisation is transferred to the ^{129}Xe nuclei through collisional spin exchange. At the high pressures typically used for SEOP, Rb-Xe binary collisions dominate over van der Waals collisions (Figure 2.3.3). Using pressures of more than one atmosphere also serves to broaden the narrow Rb D_1 transition and increase the optical pumping efficiency.

SEOP takes place in an optical pumping cell, which contains a pool of Rb, a low density Xe gas mix, and a buffer gas mix of ^{4}He and N_2, which is needed to reduce the probability of Rb-Xe collisions which destroy the Rb electron polarisation. N_2 is also acts as a quenching gas, reducing Rb depolarisation from radiation trapping - the reabsorption of unpolarised light radiated from Rb during relaxation.

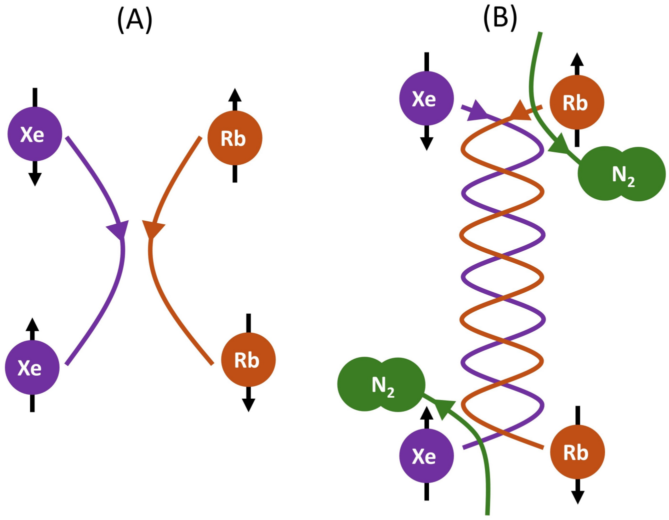

Figure 2.3.3: Spin exchange between Xe and Rb atoms occurs through two mechanisms: (A) binary collision, where the spin polarisation is transferred via the Fermi-contact hyperfine interaction and (B) the formation of temporary Rb-Xe van der Waals molecules from three-body collisions between Xe, Rb and N_2 molecules.

^{129}Xe SEOP can be performed in either ‘stopped-flow’ (SF) or ‘continuous-flow’ (CF) mode. In SF-SEOP, also known as ‘batch mode’, the hyperpolarised ^{129}Xe dose is dispensed directly from the optical pumping cell. Higher hyperpolarised Xe production rates can typically be achieved by using SF mode, although there may be a trade-off between production rate and the ^{129}Xe polarisation level [96]. CF-SEOP involves continuously flowing ^{129}Xe through the optical pumping cell at a low rate, such that sufficient polarisation occurs whilst the nuclei are within the cell. The freezing temperature of Xe is greater than that of the buffer gases, so by passing the gases through a ‘cold finger’ (a glass spiral submerged in liquid nitrogen) after they exit the optical pumping cell, the ^{129}Xe can be cryogenically separated from the buffer gases. Once the desired volume of hyperpolarised ^{129}Xe has been accumulated (approximately 18 minutes for 1 L), the frozen ^{129}Xe is sublimated by submerging the cold finger in warm water and the acquired gas dispensed into a plastic Tedlar bag for delivery to a patient.

Specific technical details of the polariser system used for the experiments described in this thesis are provided in Section 4.1 of Chapter 4.

2.3.3 Considerations for Hyperpolarised ^{129}Xe MRI

Hyperpolarised ^{129}Xe lung MRI has several unique properties which differ from standard proton MRI and lead to some additional considerations, which are discussed in this section.

One of the principle distinctions between hyperpolarised gas MRI and ^{1}H MRI is that, whereas for ^{1}H MRI, the MR signal can undergo excitation and relaxation perpetually, for ^{129}Xe, the signal enhancement is non-renewable. The ^{129}Xe hyperpolarisation is finite and will undergo longitudinal relaxation back to thermal polarisation, which is insufficient for MR imaging, as discussed in the previous section. Therefore, for hyperpolarised ^{129}Xe MRI, T_1 represents the time constant of polarisation decay. This decay is accelerated by the application of RF pulses. To conserve the non-renewable hyperpolarisation, low flip angle spoiled gradient echo (SPGR) and balanced steady-state free precession sequences (bSSFP) are typically used for gas-phase imaging.

The longitudinal magnetisation after the n^{th} RF pulse of flip angle θ in an SPGR acquisition is given by the following equation:

$$M_z(n) = M_0 \cos^{n-1}\theta \tag{2.3.1}$$

and the transverse magnetisation after the n^{th} RF pulse is given by:

$$M_{xy}(n) = M_0 \exp\left(-\frac{(n-1)TR}{T_1}\right) \sin\theta \cos^{n-1}\theta \tag{2.3.2}$$

The dependence of the magnetisation decay on the RF pulse number imposes a k-space filter, the shape of which is determined by the ordering of the phase encoding. Higher SNR can be achieved using a centric rather than sequential Cartesian trajectory, because less filtering is applied to the low frequency k-space data. However, the high frequency data experience more signal decay, so this comes at the expense of loss of fine structure detail.

Most hyperpolarised lung MRI is performed at breath hold, to avoid motion artefacts from respiratory movement. This has the added benefit of minimising the amount of O_2 in the lungs, and therefore the T_1 signal decay.

The numerous air-tissue interfaces within the lungs introduce magnetic field inhomogeneities related to the susceptibility difference between the lung tissue (diamagnetic) and oxygen (paramagnetic). This can lead to susceptibility artefacts and image degradation, but can be corrected for in ^{129}Xe gas-phase images by acquiring B_0 maps [97].

2.3.3.1 Dissolved-phase ^{129}Xe MRI

Imaging the ^{129}Xe dissolved in the lung tissue and blood poses several additional challenges. The transverse relaxation time, T_2^*, is very short for dissolved ^{129}Xe, which necessitates imaging with a short TE to limit signal loss during readout. Imaging the dissolved ^{129}Xe requires specialised frequency-selective RF pulses such that only the dissolved-phase ^{129}Xe resonance is excited and off-resonance excitations are minimised. The dissolved-phase signal makes up only ~2% of the gas-phase signal, due to the smaller volume of the tissue and blood compartments compared to the airspace and low solubility, so even a small excitation of the gas resonance can result in a relatively large signal [98]. To avoid gas-phase signal contamination, narrow bandwidth pulses are used, but these must be designed carefully to balance frequency selectivity with a short pulse width due to the short T_2^* [88]. The small proportion of dissolved-phase ^{129}Xe when compared to the gas-phase means that dissolved-phase imaging suffers from inherently low SNR [99]. To boost this, a high flip-angle can be applied selectively to the dissolved-phase resonance. The depleted hyperpolarisation after excitation is replenished by gas exchange from ^{129}Xe in the alveolar airspace.

2.3.3.2 Radio-frequency coils

Specialised RF coils tuned to the ^{129}Xe Larmor frequency are required for hyperpolarised ^{129}Xe MRI. Either rigid birdcage or flexible vest coils can be used for this purpose. Birdcage coils tend to have better B_1 homogeneity but vest coils are more sensitive and efficient due to their higher filling factor [100]. For clinical lung imaging, it is often necessary to acquire both ^{129}Xe and proton images at the same lung inflation state, which can be achieved using multi-tuned RF coils or using a dedicated ^{129}Xe coil in conjunction with the proton body coil of the scanner. The ^{129}Xe coil must be electrically isolated from the proton resonance. Increased sensitivity and accelerated acquisition are possible by using a birdcage transmit coil with a multi-channel receiver array, which allows for parallel imaging [101].

Chapter 3

Lung Physiology and Gas Exchange Measurement

The primary function of the lungs is gas exchange: the diffusion of oxygen (O_2) from the air into the blood and the removal of carbon dioxide (CO_2). This chapter starts with an introduction to lung anatomy and physiology, which also covers some common tests of pulmonary function. The majority of this chapter is then focused on reviewing hyperpolarised ^{129}Xe MRI/S techniques for measuring gas exchange.

3.1 Lung Anatomy and Physiology

The lungs have a lobar anatomical structure; the right lung consists of three lobes (superior, middle and inferior) whereas the left lung only has two lobes (superior and inferior), because of the presence of the heart. The airways are formed from a series of branching tubes which start from the trachea. The trachea divides into the left and right bronchi, which in turn bifurcate. This subdivision of airways continues for 23 generations (Figure 3.1.1). The first 14 airway generations are called the conducting airways, which lead inspired air into the transitional bronchioles and respiratory airways, which make up the majority of the lung volume (~3 L at rest).

The volume of gas in the lungs during normal (tidal) breathing and at maximum inspiration/expiration can be measured by PFT. These are important functional measurements for the diagnosis and monitoring of lung disease, including differentiating between restrictive and obstructive disease [102]. Figure 3.1.2 depicts the lung volumes and capacities which can be measured by PFT. The volumes written in blue can be measured using spirometry. There are two different types of spirometry equipment: closed-circuit spirometers, which collect and measure the volume of exhaled air, and open-circuit spirometers, which do not collect the air and measure airflow instead, for example with a turbine flow meter [103]. In a spirometry exam, the patient starts by breathing normally, allowing the tidal volume to be measured. Then, they are instructed to inhale to maximal inspiration and exhale to maximum exhalation. The volume of exhaled gas is called the vital capacity, which is not equivalent to TLC, because some gas remains in the conducting airways. This residual volume (RV) of gas cannot be measured with spirometry and is instead calculated from the functional residual capacity (FRC). This is the volume of air which remains in the lungs after a normal expiration and can be measured using body plethysmography or gas dilution methods [102]. TLC is calculated from the sum of FRC and the inspiratory capacity. Spirometry is also used to measure the forced expiratory volume in one second (FEV_1). It is typical for PFT results to be reported as predicted values or z-scores, which are derived from reference equations based on subject demographics such as age and sex.

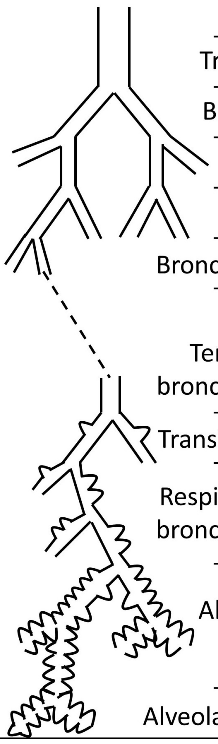

		G	Diameter (cm)	Length (cm)	Number	Surface area (cm^2)
Conducting airways	Trachea	0	1.80	12.0	1	2.54
	Bronchi	1	1.22	4.76	2	2.33
		2	0.83	1.90	4	2.13
		3	0.56	0.76	8	2.00
	Bronchioles	4	0.45	1.27	16	2.48
	Terminal bronchioles	5 - 14	0.35 - 0.074	1.07 - 0.23	32 - 1.6×10^4	3.11 - 69.4
Respiratory airways	Transitional	15	0.066 – 0.050	0.20 - 0.117	3.3×10^4 - 2.6×10^5	113.0 - 534.0
	Respiratory bronchioles	16 - 18				
	Alveolar ducts	19 - 22	0.047 - 0.041	0.099 - 0.059	5.2×10^5 - 4.2×10^6	944.0 - 5880.0
	Alveolar sacs	23	0.041	0.050	8.4×10^6	1.2×10^4

Figure 3.1.1: The Weibel model of the airways in the human lung, adapted from [104], along with morphometric data from Weibel [105]. 'G' refers to the airway generation.

The lungs are highly specialised to maximise the diffusion of inspired oxygen into the bloodstream. This is achieved by the division of the lung into about 500 million acinar structures called alveoli, which provides a very large surface area (∼100 m^2) for gas exchange to take place [102]. Beyond the 16th airway generation (respiratory bronchioles), where the alveoli first appear, the number of alveoli increases with each generation until the final generation, which is entirely made up of alveolar sacs. The alveoli have a mean diameter of ∼250 μm and a volume of ∼4 x 10^6 μm^3 [106]. They are wrapped with capillaries, which form a dense network of blood vessels, equivalent to an almost continuous sheet of blood. This is very efficient for gas exchange. The alveolar airspace is segregated from the capillaries by a tissue barrier which is around 1 μm thick in the healthy lung, made of epithelial and endothelial cells, separated by an interstitial space [104]. For inspired gas to enter the blood, it must travel through the tissue barrier, which occurs by passive diffusion along the partial pressure gradient. Once the gas has crossed the blood-gas barrier, it dissolves in the blood and is transported through the pulmonary circulation, which also has a branching structure. Starting at the main pulmonary artery (MPA), the blood vessels split into smaller and smaller branches until they join the pulmonary capillary bed. The diameter of the capillaries in the capillary bed network range between 7 - 10 μm, which is around the size of an RBC [102]. The capillary bed is connected to the smallest veins, known as venules, which converge to form larger vessels and eventually join the main pulmonary vein (MPV).

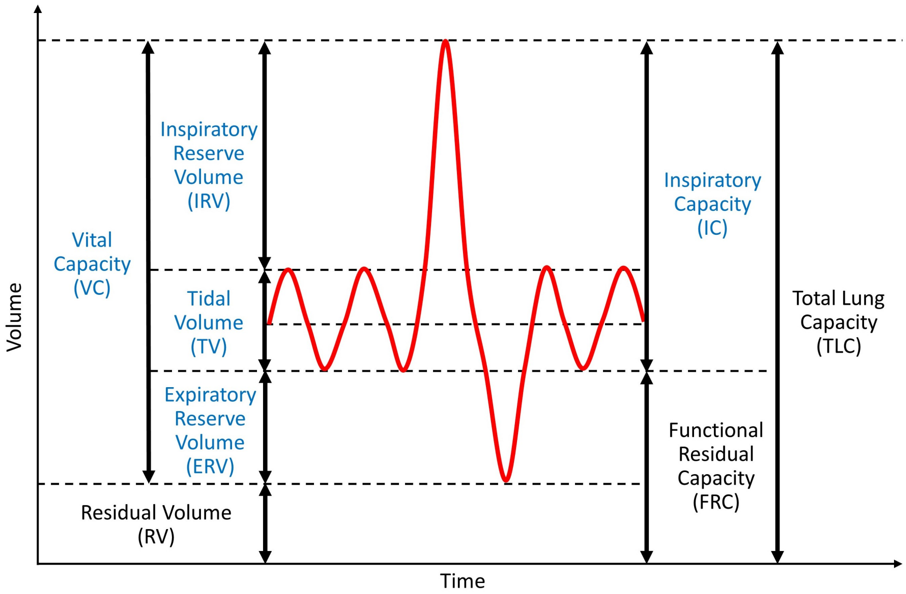

Figure 3.1.2: Diagram of the different lung volumes and capacities. The red curve shows the change in lung volume over time during tidal breathing, forced inhalation to total lung capacity and exhalation to expiratory reserve volume, as would take place in a typical spirometry exam. This manoeuvrer allows the lung volumes and capacities shown in blue to be measured. Residual volume, functional residual capacity and total lung capacity can be measured using alternative techniques, such as body plethysmography.

3.2 Conventional Measures of Gas Exchange

3.2.1 Transfer Factor of the Lung for Carbon Monoxide (TL_{CO})

The amount of gas which diffuses across a tissue barrier is directly proportional to the area of the tissue and inversely proportional to its thickness, from Fick's law [102]. The combined effect of surface area and thickness, as well as the diffusion of the gas involved, can be described by the diffusing capacity of the lungs, D_L. This is related to the amount of gas transferred, $\dot{V}$, and the difference in partial pressure between the alveolar (P_A) and capillary blood (P_C):

$$D_L = \frac{\dot{V}}{P_A - P_C} \tag{3.2.1}$$

The carbon monoxide (CO) transfer factor (TL_{CO}), also called the diffusing capacity for carbon monoxide (DL_{CO}), represents the whole-lung average efficiency of gas transfer from the alveoli to the bloodstream. TL_{co} can be measured in a common non-invasive test of gas transfer, first described by Ogilvie et al. in 1956 [107]. CO is used for this purpose because it is diffusion-limited: the strong bonds that it forms with haemoglobin mean that the CO alveolar capillary partial pressure gradient increases only very slightly with CO uptake by the RBCs. Instead, CO uptake is limited by the diffusion rate across the tissue barrier [102]. The change in capillary partial pressure for a diffusion-limited gas such as CO is shown against that for a perfusion-limited gas in Figure 3.2.1. For a perfusion-limited gas, there is a large partial pressure gradient at the beginning of the alveolar capillary which drives the gas into the capillary. Due to the low affinity for haemoglobin, the gas remains in the capillary blood plasma producing a partial pressure and an equilibrium is quickly reached. There is no longer a partial pressure gradient so there is no more net diffusion and the only way to increase the amount of gas dissolved is to increase perfusion.

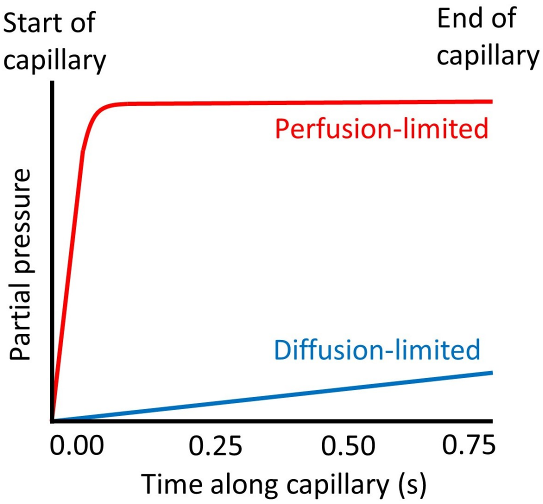

Figure 3.2.1: Capillary partial pressure as a function of time along the capillary for a perfusion-limited gas (red) such as nitric oxide and a diffusion-limited gas (blue) such as CO or NO.

TL_{CO} is found using Equation 3.2.1, where $P_{C,CO}$ can be neglected due to its high affinity for haemoglobin [102], giving:

$$T_{LCO} = \frac{\dot{V}_{CO}}{P_{A,CO}}. \tag{3.2.2}$$

To measure TL_{CO}, the inspired and expired concentrations of CO are measured for a single breath hold and the amount of gas diffused found from the difference between the concentrations [108]. The patient inhales a test gas containing 0.3% CO and 0.3% tracer gas, such as helium, and holds their breath at TLC for ten seconds [109]. The dilution of the tracer gas is used to calculate the number of accessible alveolar units (V_A), and the rate of disappearance of CO from the alveolar gas gives the gas exchange efficiency per unit (K_{CO}). The product of these measurements yield TL_{CO}.

TL_{CO} is used as a clinical measure of gas exchange in the diagnosis and assessment of lung disease [109, 110, 111], as well as a marker in drug trials [112]. Despite its widespread use, there are several limitations to using TL_{CO} to measure gas exchange function. It is a global measure with no regional information and is dependent on several other factors such as patient effort, testing conditions, anaemia, and body weight [113]. TL_{CO} is reduced by the presence of PVD, emphysema and ILD, and so interpreting a low TL_{CO} can sometimes be challenging [114].

Using nitric oxide (NO) as an alternative to CO may avoid some of the technical problems surrounding TL_{CO} measurement; for example it is much less affected by haemoglobin concentration and ambient pressure [115]. Like CO, NO is a diffusion-limited gas, but NO reacts with haemoglobin at a much faster rate than CO (around 300 times faster in vitro [116]). This means that TL_{NO} is more sensitive to the diffusive properties of the tissue barrier than TL_{CO} and so might be more representative of gas exchange impairment from tissue thickening. Recent works have explored the use of the TL_{NO}/TL_{CO} ratio to distinguish between PVD and ILD. Where both TL_{NO} and TL_{CO} are reduced, an increase in their ratio signifies pulmonary vascular disease, whereas a decrease indicates fibrosis or ILD [116]. However, this method has not been fully validated and Degano et al. found that the TL_{NO}/TL_{CO} ratio did not improve the identification of PH in patients with SSc [117].

3.2.2 Ventilation-Perfusion Imaging

How efficient a lung is at exchanging gas is determined by the ratio of alveolar ventilation to the perfusion of deoxygenated blood in the capillary ($\dot{V}_A/\dot{Q}$) [118]. Optimal efficiency for a single gas exchange unit is at $\dot{V}_A/\dot{Q} = 1$ and $\dot{V}_A/\dot{Q}$ mismatches can be characteristic of disease, for example reduced $\dot{V}_A/\dot{Q}$ ratio in ILD patients [119, 120]. Several methods exist to measure the $\dot{V}_A/\dot{Q}$ ratio, including the Multiple Inert Gas Elimination technique (MIGET) [121], nuclear medicine [122, 123, 124, 125], CT [126], and more recently, ^{1}H MRI and MRI with hyperpolarised noble gas [118]. Each method comes with its own limitations, such as the lack of spatial localisation in MIGET, the use of ionising radiation in nuclear methods and gadolinium in DCE-MRI for imaging perfusion [127].

3.3 Measuring Gas Exchange with Hyperpolarised ^{129}Xe MRI/S

In this section, we review the literature on dissolved-phase ^{129}Xe MRI/S techniques used for measuring pulmonary gas transfer. For an overview of ^{129}Xe ventilation imaging and ^{129}Xe DW-MRI, the reader is referred to the following review papers: [63, 128].

3.3.1 Spectroscopy

Hyperpolarised ^{129}Xe MRS was first carried out in human blood [129, 80] and rodents [130, 131, 132], before advances in polarisation technology enabled the first human results in the chest and brain in 1997 [133]. These early results were hampered by low spectral resolution due to the limited polarisation levels achievable at the time, but were nonetheless able to resolve the two dissolved-phase peaks in the lungs.

Nowadays, ^{129}Xe gas exchange is generally measured using imaging techniques, which provide insight into the regional gas transfer distribution and heterogeneity. Whole-lung ^{129}Xe spectroscopy is often incorporated into imaging sequences, in order to measure spectral parameters such as the ^{129}Xe resonant frequencies, needed in the imaging reconstruction. However, ^{129}Xe spectroscopy is also an interesting technique in its own right, and can be used to derive whole-lung gas exchange ratios. For example, Kaushik et al. used ^{129}Xe spectroscopy to measure RBC:M in patients with IPF, finding reduced values which were strongly correlated with $\mathrm{TL_{CO}}$ [134]. This is a useful tool for particular clinical applications where dissolved-phase imaging may be challenging or infeasible, such as in neonatal studies. Stewart et al. used ^{129}Xe spectroscopy to detect reduced RBC:M in nine neonates with bronchopulmonary dysplasia recruited from a neonatal intensive care unit [135].

Several useful spectral parameters can be inferred from ^{129}Xe spectroscopy, which are sensitive to different aspects of lung physiology. The chemical shift of the RBC peak is dependent on blood oxygenation level [90] and has been found to be reduced in patients with ILD and COPD [136, 134, 137, 138, 139]. The full-width at half-maximum (FWHM) of the peaks is equivalent to the inverse of the T_2^* of each signal. This has been proposed to be sensitive to the rate of chemical exchange of ^{129}Xe between the tissue, blood plasma and RBCs [140].

^{129}Xe spectra can be acquired dynamically, typically by repeatedly acquiring FIDs with a short TR of ∼20 ms [137]. This offers insight to temporal variations in the spectral parameters and has revealed periodic fluctuations in the amplitude of the RBC peak. The amplitude of these oscillations has been shown to be sensitive to the presence of cardiopulmonary disease, and is reviewed in more detail in Section 3.5. Oscillations have also been observed in the RBC frequency shift of patients with IPF [37]. Time-resolved ^{129}Xe spectroscopy with a variable TR can be used to characterise uptake into the alveolar septa; this is discussed in the next section.

3.3.1.1 Chemical Shift Saturation Recovery

To measure the ^{129}Xe exchange dynamics, a spectroscopy technique called chemical shift saturation recovery (CSSR) can be used. This was pioneered by Ruppert et al. [40] and initial human results were reported in 2008 by Patz et al. [41]. CSSR provides a means to sensitise ^{129}Xe MRS to gas exchange. Unlike CO and NO, the gases used to measure gas exchange in PFT, ^{129}Xe is not diffusion-limited and instead is limited by perfusion. This is due to the weak interaction between ^{129}Xe and haemoglobin which allows ^{129}Xe to quickly saturate the alveolar septum and eliminates the alveolar-capillary partial pressure gradient [134]. The CSSR sequence involves acquiring consecutive FIDs following repeated 90° RF pulses, which selectively excite the dissolved-phase, applied with a varying time delay. A spoiler gradient is used such that the dissolved-phase magnetisation is destroyed at the start of each TR. This is replenished by ^{129}Xe in the gas-phase, which first crosses the tissue barrier and then enters the pulmonary capillary plasma and RBCs. This allows ^{129}Xe to measure diffusion limitation; the dissolved-phase signal depends on how much gas diffuses into the dissolved-phase over the delay time and if there is interstitial thickening, the RBC signal will decrease whilst the M signal will increase. A typical CSSR pulse sequence is shown in Figure 3.3.1A. Examples of ^{129}Xe uptake curves are shown in Figure 3.3.1B along with the MRI signal, where F(TR), which is the dissolved-phase signal normalised by the gas-phase signal, is plotted against TR. Clinically relevant parameters, such as the alveolar septal thickness (h), alveolar surface area to volume ratio (S/V) and capillary transit time (τ), can be estimated by fitting the gas uptake data to mathematical models. The most common models are the Patz model and MOXE, which are described in Section 3.4.

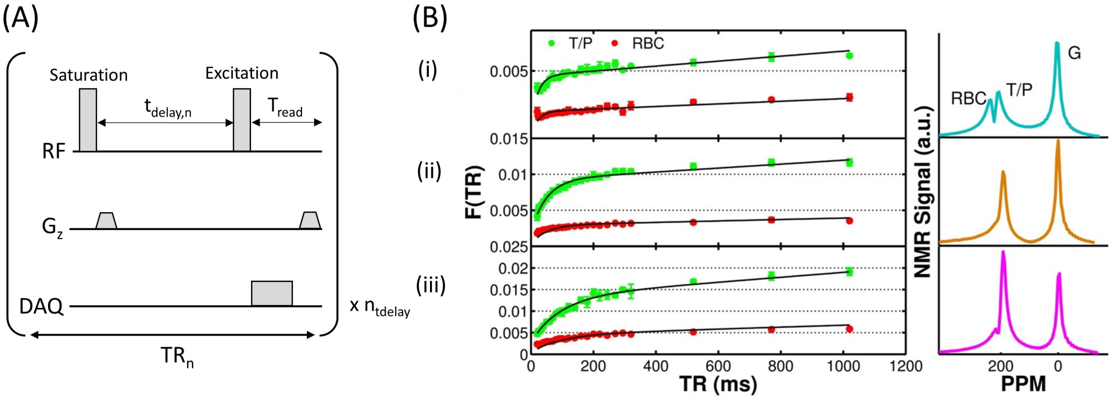

Figure 3.3.1: (A) A typical CSSR pulse sequence, where the first 90° pulse saturates the dissolved-phase magnetisation and the second generates the FID after a time t_{delay}. This is repeated with increasing t_{delay}. T_{read} is the readout time. (B) Representative ^{129}Xe uptake curves for (i) a healthy volunteer, (ii) a SSc patient, (iii) an IPF patient. The signal first increases exponentially because of the increasing amount of gaseous ^{129}Xe which can diffuse into the dissolved-phase as the delay time is increased [63], then plateaus as the blood is saturated and increases with time due to the effect of blood flow effect (see Figure 3.4.2). The solid black lines represent a fit to the MOXE. The right panel shows the corresponding spectra, where differences in the dissolved-phase signal between the healthy and ILD subjects are apparent. Copyright 2014 Wiley, reproduced with permission from [39].

The clinical applications of CSSR have been demonstrated in several different lung diseases and it has been shown to be a reproducible method for quantifying h [141]. Patz et al. were the first to show that CSSR could detect septal thickening in ILD patients, even in mild disease, in 2011 [142]. A shortcoming of this work is that the low field strength used (0.2 T) meant that the dissolved-phase signal could not be separated into the M and RBC spectral peaks, potentially causing errors in the subsequent analysis. Stewart et al. found evidence of alveolar septal tissue thickening in IPF and SSc patients using CSSR at 1.5 T [39]. This work also offered the first comparison of CSSR metrics to gold standard clinical methods; TL_{CO} and DCE-MRI. A significant negative correlation was found between h and TL_{CO}, although no significant correlation was found between τ and similar metrics from DCE-MRI.

CSSR-derived h has been shown to be sensitive to tissue damage in COPD [38]. Another study found increased age-corrected h in both patients with COPD and healthy smokers, compared to healthy never-smokers [143]. The relationship between age and h was found using ordinary least-squares regression for the group of 21 healthy volunteers and this was used to derive a reference equation. None of the metrics from the other methods examined (PFT, six minute walk test and ^{129}Xe DW-MRI) were able to differentiate the healthy smokers from the healthy non-smokers.

Recently, CSSR was used to help elucidate the disease mechanisms in the so-called ‘lung phenotype’ of idiopathic PAH [144]. This phenotype is characterised by a reduced TL_{CO} and predominantly affects male patients with a significant smoking history. RBC:M and CSSR-derived S/V were reduced in this patient group when compared to idiopathic PAH patients with preserved TL_{CO}. This is consistent with the hypothesis that idiopathic PAH with reduced TL_{CO} is caused by pulmonary capillary loss, along with early diffuse emphysema.

Another application of CSSR is to study the effects of air pollution on the lung microstructure and function. An increased h was detected in six rats who had been exposed to fine particulate matter ($PM_{2.5}$) in comparison to six healthy control rats [145]. The measurements of h correlated well with those from histology. This demonstrates the potential of CSSR to detect early lung changes resulting from air pollution exposure, which may prove useful for future human studies.

CSSR and ^{129}Xe DW-MRI can be performed within the same breath hold, which reduces cost, scanning time and errors resulting from lung inflation differences between different acquisitions [146]. Xie et al. developed a combined sequential CSSR and DW imaging sequence [147]. This consisted of a CSSR data acquisition, followed by 2D multi-slice DW-MRI and finally a short spectral acquisition to obtain FIDs for correcting gas/dissolved-phase flip angle differences. Compressed sensing was used to reduce the scan time of the DW-MRI sequence block, resulting in a total scan duration of 10 s. This method was tested in healthy volunteers and patients with COPD and was able to detect alveolar tissue loss in the COPD patients. Alternatively, diffusion-weighted CSSR (DWCSSR) [146] takes advantage of the multiple delay times in CSSR, utilising these empty sequence blocks to perform DW imaging repetitions. So far, this method has only been assessed in healthy rats. In comparison to acquiring CSSR and DW-MRI sequentially, DWCSSR required a shorter scan time (4 s vs 6 s) and the gaseous ^{129}Xe signals of both CSSR and DW-MRI were relatively higher.

Standard CSSR sequences do not feature any spatial encoding, which limits their regional sensitivity because any parameters extracted are global averages. Both regional and temporal gas exchange information can be obtained by performing spatial encoding after each CSSR delay time. However, to maintain a tolerable breath hold time, fewer delay times can be used, resulting in fewer data points and potentially less accurate fitting. Following their previous work in rats

[99, 148], Doganay et al. used three-point IDEAL imaging (Iterative Decomposition of Water and Fat with Echo Asymmetry and Least-Squares Estimation; see Section 3.3.2.3) and a multi-interleave spiral k-space sampling scheme to acquire 2D images of the gas, RBC and M signals at seven different time points in humans [149]. They found evidence of decreased gas transfer surface area in patients with emphysema in comparison to healthy individuals. Zanette et al. developed a 2D four-point multi-interleave spiral-IDEAL sequence for measuring gas exchange at five different time points in healthy volunteers [150]. The RBC and M signals were fit to a theoretical diffusion model (see Section 3.4) on a voxelwise basis to yield maps of the derived parameters. Using two-fold undersampling and parallel imaging, Zanette et al. were able reduce the length of the required breath hold from 11 s to only 6 s. However, the accelerated sequence resulted in worse agreement between the data and the diffusion model. A four-point spiral-IDEAL sequence has also been reported to measure the time-course of ^{129}Xe RBC and brain tissue signals the rat brain [151].

Kern et al. used a multi-channel coil and the SPectral Localization Achieved by Sensitivity Heterogeneity (SPLASH) method [152, 153] to perform regional CSSR in a small study of healthy volunteers and COPD patients [154]. Ventilation images were taken within the CSSR breath hold and used to section the lung into eight regions of interest (ROI) and 16 sub-ROI, to match the 16 receive elements of the coil. The spatial heterogeneity of the receiver coil elements were then used to resolve the signals from different compartments. The CSSR parameters showed good reproducibility (except for τ in the COPD patients), and were in the expected range from literature. A significant difference was found for the S/V and RBC:M ratios between the anterior and posterior of the lung in healthy volunteers, which is most likely a gravitational effect due to the patients' supine position. This was not observed in COPD patients, possibly due to hyperinflation of the lung tissue. However, these results may be biased due to the exclusion of very small delay times (and therefore sensitisation to the very thin lung tissue which saturates in this time), because of hardware limitations. SPLASH also assumes a homogeneous signal from each region of interest, which might not be true in diseased lungs.

Another paper from Kern et al. uses a different approach for regional CSSR [155]. They utilised the CSSR delay time by adding low-flip angle acquisitions which do not significantly alter the overall magnetisation, in a Look-Locker-like scheme [156]. Therefore, multiple measurements could be made before the dissolved-phase magnetisation is destroyed by the 90° RF pulse. To further accelerate the imaging time, an undersampled stack-of-stars trajectory was used with low-rank sparse matrix decomposition to reduce artefacts [157]. The images were separated into M and RBC compartments using the 1-point Dixon method (Section 3.3.2.2) and the M compartment of each voxel was then fitted to the Patz model to create regional parameter maps. A strong correlation was found between the global average of the parameters found from regional CSSR and the parameters found from spectroscopic CSSR. However, one limitation of this research is that the two spectroscopic measurements were carried out with different lung inflation, meaning that they cannot be compared directly, and there was no age-matching.

Regional CSSR can also be performed by repeating the CSSR sequence block several times, acquiring one line of k-space at each repeat. This approach is known as CSI-CSSR because it combines CSSR gas exchange measurement with CSI spatial localisation. A limitation of this method is that repeating the sequence for several frequency encodes and delay times is time-consuming and restricts the achievable image resolution. CSI-CSSR with spatial encoding in one dimension (15 s breath hold time) has been employed to assess regional variations in gas exchange dynamics in the anterior to posterior direction in humans [158]. Ruppert et al. implemented a 2D CSI-CSSR sequence, but this necessitated the use of multiple breath holds and, for this reason, has only been tested in rats [159]. Acquiring spectra at each phase encode, rather than absolute signals, allows additional parameters such as the T_2^* to be mapped regionally. Ruppert et al. found that, whereas in healthy rats all metrics were comparable in both lungs, there were several differences between the lungs of a rat with radiation induced lung injury (RILI) in the right lung [159]. CSI-CSSR was able to detect increased h, M:Gas, M T_2^* and RBC T_2^* and decreased RBC:Gas and RBC:M in the irradiated lung.

3.3.2 Imaging

Regional gas exchange information can be provided by spectroscopic MRI techniques. As mentioned in Section 2.3.3.1, the dissolved-phase ^{129}Xe signal is inherently low, but can be maximised by imaging with a short TR and reasonably large flip angle (compared to the gas-phase, where a small flip angle is used to conserve the longitudinal magnetisation). Using a high flip angle to excite the dissolved-phase is possible because the dissolved-phase magnetisation is being constantly replenished by gaseous ^{129}Xe diffusing through the tissue barrier. Early ^{129}Xe dissolved-phase imaging results were reported in 2010. Cleveland et al. used a radially encoded imaging sequence to acquire regional dissolved-phase images of the healthy human lung [160]. Mugler et al. simultaneously imaged gas and dissolved-phase ^{129}Xe using a Cartesian-encoded sequence which exploited the chemical shift artefact [161]. However, neither of these works separated the dissolved-phase signal into the RBC and M compartments.

3.3.2.1 Chemical Shift Imaging

The simplest dissolved-phase imaging technique which separates the dissolved-phase signal into its two compartments is CSI (described in Section 2.2.6), a FID based approach which repeats spectroscopy measurements at each Cartesian phase encoding step [162]. CSI has been used to show reduced RBC:M in the IPF lung compared to healthy volunteers, which correlated well with TL_{CO} [138].

Although slower than other imaging methods, an advantage of CSI is that it allows full characterisation of the ^{129}Xe spectral peaks, and therefore provides valuable additional information. As mentioned in Section 3.3.1, the chemical shift of the RBC spectral peak is correlated with blood oxygenation and is reduced in some patient groups. Using CSI, regional maps of the dissolved-phase ^{129}Xe chemical shifts and T_2^* can be generated [138, 139]. Figure 3.3.2 shows an example of the RBC chemical shift maps for a healthy volunteer, an IPF patient, a CF patient

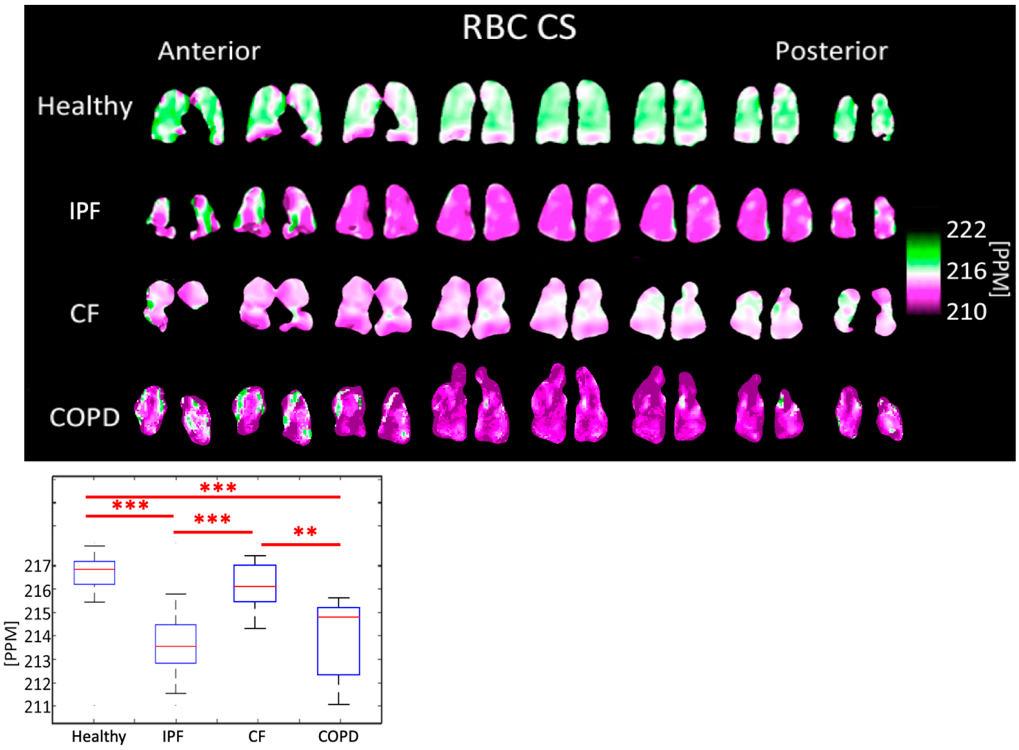

Figure 3.3.2: Maps of RBC chemical shift (top) and boxplots of the whole-lung CSI averages (bottom) for each patient group. ($p < 0.001$ = ***, $p < 0.01$ = **). Copyright Guan et al. 2022, licensed under Creative Commons Attribution CC-BY, reproduced with permission from [139].

and a COPD patient from Ref. [139]. The whole-lung averages of these maps were significantly different; RBC chemical shift was reduced in patients with IPF and COPD, when compared to healthy volunteers. Changes in the M and RBC T_2^* relative to healthy values were also observed in this group of patients. The M T_2^* was longer in IPF patients, but shorter in CF patients, and the RBC T_2^* was increased in both IPF and COPD patients. The physiological reason for these changes is not well-understood, but has been hypothesised to be related to the rate of ^{129}Xe chemical exchange between the capillary blood plasma and RBCs [140].

3.3.2.2 Dixon Method

The Dixon spectroscopic imaging method was developed for separating fat and water signals. This is done by introducing a 90° phase difference between the fat and water signals to produce in-phase and out-of-phase images, from which fat-only and water-only images can be constructed [163]. This approach can be adapted to separate the RBC and M signals from ^{129}Xe MRI and to produce 3D images and image-derived gas exchange ratios, which correlate well with results from spectroscopy [31]. Early pre-clinical work was limited by the assumption of B_0 homogeneity across the lungs [31], but this has since been corrected for with gas-phase image acquired phase maps [32].

The Dixon method is the most widely-used technique for dissolved-phase ^{129}Xe

MRI, and is typically implemented with radial k-space encoding and a single TE value (‘1-point Dixon’) to introduce a 90° phase difference between the RBC and M peaks [164]. Phase-shifting the acquired image therefore restricts the RBC and M signals to the real and imaginary receiver channels respectively, allowing their separation. A pulse sequence diagram for radial 1-point Dixon imaging is shown in Figure 3.3.3.

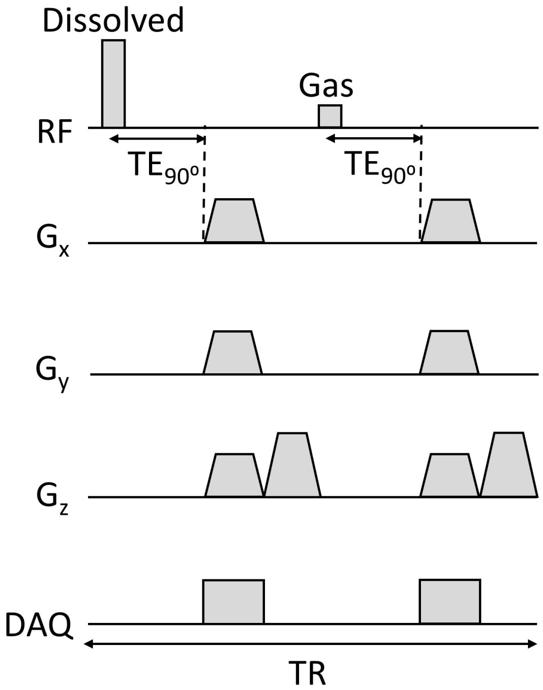

Figure 3.3.3: 3D radial 1-point Dixon pulse sequence diagram: an interleaved acquisition of dissolved and gas signal is achieved using frequency-selective RF pulses, with typical flip angles of <1° on the gas-phase resonance and 20° on the dissolved-phase resonance. The TE is chosen such that a 90° phase shift is introduced between the RBC and M signals ($TE_{90°}$). Gradients are applied in each direction for 3D radial k-space encoding and a spoiler gradient is applied in the z-direction at the end of each interleaf.

Dissolved ^{129}Xe Dixon imaging is repeatable [165] and sensitive to gas exchange impairment in a range of lung diseases, including post-COVID-19 lung disease [166, 167], COPD [37] and pulmonary vascular disease [17, 37]. In ILD patients, Dixon imaging can reveal RBC signal deficits, such as in Figure 3.3.4, which correspond to fibrosis seen with CT [32]. In Ref. [33], a poor correlation was found between RBC:M and CT fibrosis score but a strong correlation was found with TL_{CO}, suggesting that ^{129}Xe MRI may be able to detect early fibrosis not visible with CT. Significant differences have been detected in mean M uptake and RBC:M ratio in patients with IPF [168, 33], Non-Specific Interstitial Pneumonia (NSIP) [169] and chronic hypersensitivity pneumonitis [136], compared to healthy volunteers. Dixon-based gas exchange imaging has been used to assess longitudinal changes in lung function in patients with IPF [170]. Reduced RBC:M at baseline was associated with disease progression one year later.

This method is well-tolerated by paediatric patients and its feasibility has been demonstrated in healthy children and children with bone marrow transplantation, bronchopulmonary dysplasia, CF and childhood interstitial lung disease [171, 172]. A novel application of ^{129}Xe Dixon imaging is to use the gas exchange maps to guide functional avoidance treatment planning in radiation therapy [173]. Rankine et al. demonstrated the potential of ^{129}Xe ventilation and gas exchange imaging to reduce radiation dose to high gas exchanging regions of the lung.

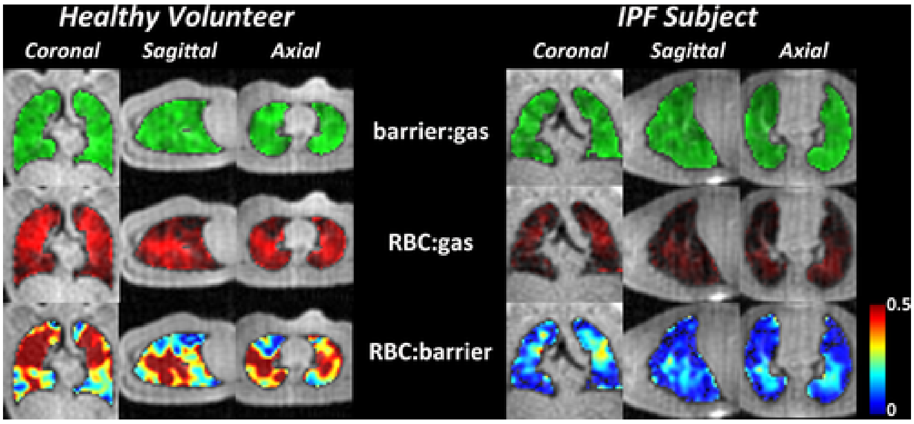

Figure 3.3.4: Gas, M and RBC ratio images from a healthy volunteer and a person with IPF. The IPF patient shows reduced ^{129}Xe RBC uptake. Copyright 2015 Wiley, reproduced with permission from [32].

Dixon imaging has been combined with ventilation imaging in a single breath hold. Niedbalski et al. interleaved radial 1-point Dixon imaging with 3D spiral encoding of the gas-phase ^{129}Xe signal [174]. The VDP and gas exchange ratios from the single breath method showed good agreement with those from dedicated ventilation and gas exchange imaging.

The 1-point Dixon method is beneficial for reducing scan time but suffers from gas-phase signal contamination, from imperfect frequency selection [175]. Contamination can lead to errors in separating the dissolved-phase resonances which in turn leads to image artefacts [176]. However, methods have been developed to remove the gas-phase contamination retrospectively using the gas-phase images, either using a second echo [176] or, more recently, requiring no additional echoes [175]. Another limitation comes from the assumption that the phase shift between the RBC and M signals during the acquisition time is negligible; however Collier et al. estimate this phase shift to be $\sim 10°$ [35] and a non-negligible phase shift during readout was observed in Ref. [32]. This phase evolution limits the image quality, causing blurring. A solution to this problem is to use multi-point Dixon imaging [177, 178] or multi-echo spectroscopic imaging, which is described in the next section.

3.3.2.3 Multi-Echo Spectroscopic Imaging

Similar to the Dixon method, multi-echo spectroscopic imaging (MESI) was originally developed as a fat-water separation method [179], but can be applied to separate other chemical species, such as the three ^{129}Xe resonances. Figure 3.3.5 shows a pulse sequence diagram for this technique. Several echoes are acquired, which induces time-dependent phase shifts due to the chemical shift between the ^{129}Xe gas, M and RBC signals. This creates a system of linear equations which are solved to return the signal distribution of each of the three resonances. As there are three unknowns, a minimum of three linear equations, and therefore three echoes, is required to separate the gas, M and RBC signals. A fourth echo is sometimes acquired, which is favourable for improving SNR and can be used to remove gas-phase signal contamination from the dissolved-phase images [35]. The system of linear equations is solved using iterative model-based reconstruction, such as IDEAL [179, 180]. An early challenge of multi-point ^{129}Xe imaging was the need to obtain at least three images per TR, whilst minimising signal decays arising as a result of the very short T_2^* of dissolved-phase ^{129}Xe [32]. This was solved by using centre-out radial [36] or spiral k-space encoding [99, 181, 150] which are more efficient than Cartesian encoding.

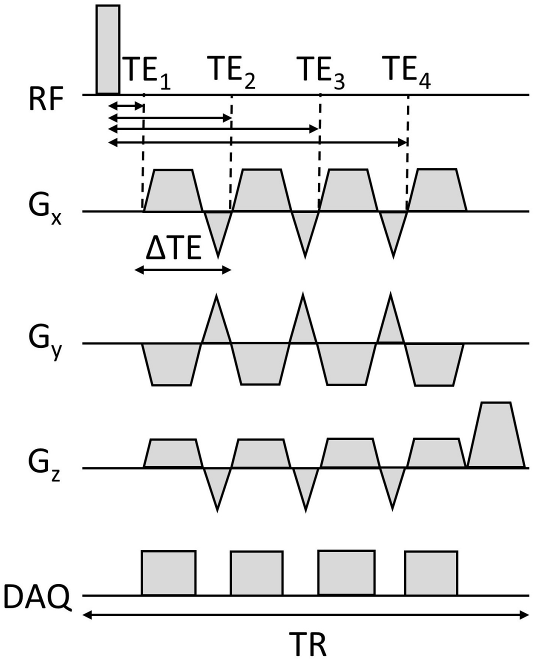

Figure 3.3.5: 3D radial MESI pulse sequence diagram: after excitation with a frequency-selective RF pulse (applied to the dissolved-phase ^{129}Xe resonance, with minimal excitation of the gas-phase ^{129}Xe resonance), k-space data are acquired at four different TEs, which are separated by an echo time spacing ΔTE. The gradient waveforms (not drawn to scale) vary in amplitude each TR to encode the full 3D radial k-space.

The noise performance of MESI is dependent on the proportion of each species in a given voxel and the position of the echoes [179]. Therefore, the choice of TEs is important to maximise the SNR, whilst minimising T_2^* decay [150]. Equally spaced echoes are generally used for practical reasons, because implementing a variable TE is more challenging on most MRI systems. The echo spacing, ΔTE, can be optimised using the NSA approach, which describes the SNR performance of the given ΔTE [180]. This is described in more detail in Section 4.2.3.

Qing et al. developed a 3D radial multi-point dissolved-phase ^{129}Xe spectroscopic imaging method which acquired three echoes following an RF pulse centred on the dissolved-phase resonance and two echoes following RF excitation of the gas-phase, to permit calculation of a reference B_0 map [36]. They used a hierarchical IDEAL reconstruction, which is a multi-resolution approach that allows for T_2^* decay estimation in the spectral decomposition [182, 183]. Using this technique, significantly lower M:Gas and RBC:Gas ratios were found for smokers and asthmatics than for healthy volunteers and the gas exchange ratios correlated with global spectroscopy measurements [36]. It has also been used to probe lung physiology changes in patients with COPD [184, 185]. A significant correlation between RBC:Gas from multi-point imaging and % predicted TL_{CO} was found, as well as between M:Gas and the % emphysema quantified by CT [185].

Qing's method has shown potential sensitivity to early interstitial lung abnormalities in healthy young electronic-cigarette users [186]. In a pilot study of five subjects, M:Gas was reduced and RBC:M was increased, compared to age-matched healthy controls, whereas RBC:Gas and TL_{CO} were preserved. Interestingly, similar results were found in a group of young asthma patients [187]. This seems to indicate tissue loss combined with increased perfusion, which may be due to airway remodelling or hyperemic capillaries.

Kammerman et al. used an interleaved 3D radial acquisition to perform ^{129}Xe gas exchange imaging, in which the RF transmit and receive frequencies were alternated between the RBC and gas resonances every TR/2 [188]. Four echoes were acquired following each excitation. An iterative model-based reconstruction was used, which was based on the IDEAL method and included estimates of T_2^*, frequency shifts and field inhomogeneity maps. A decrease in the transverse relaxation rate of the M and RBC compartments was measured in IPF patients, possibly indicating reduced rates of ^{129}Xe exchange between the airspaces, membrane and blood as a result of decreased diffusion and/or perfusion.

A four-point 3D radial dissolved-phase spectroscopic imaging sequence has been developed by Collier et al. from our research group [35, 189], which is used throughout this thesis and is described in more detail in Section 4.2.3. This sequence uses a flyback radial acquisition, which means that data were acquired from the centre of k-space along the radii, but not from the edge of k-space back to the centre. Whereas other multi-point methods relied on a low flip angle to avoid gas-phase contamination, Collier et al. used the additional fourth echo to remove the gas-phase contamination. Prior knowledge of the ^{129}Xe resonant frequency shifts obtained from calibration spectra is used with a least-squares algorithm to reconstruct the M, RBC and gas images and ratio maps. This method has been used to detect gas exchange impairment in patients with IPF, asthma, COPD and COVID-19 [35, 190, 191, 192, 193, 194].

As mentioned in Section 3.3.1.1, regional and temporal measurements of ^{129}Xe gas exchange dynamics can be made using time-series CSSR MESI [99, 148, 149, 151, 150].

3.3.2.4 Xenon Polarisation Transfer Contrast

The Xenon Polarisation Transfer Contrast (XTC) method overcomes the challenge of the small proportion of dissolved-phase ^{129}Xe by using the much larger gas-phase signal to indirectly image gas exchange [79]. In the basic XTC sequence, the gas-phase is imaged twice in one breath-hold with a gradient echo pulse sequence. In between the two acquisitions, a series of either saturation (90°) or inversion (180°) RF pulses centred on the dissolved-phase are applied. In a separate breath-hold, the sequence is repeated with the RF pulses centred off-resonance to control for T_1 relaxation, perfusion effects and depolarisation due to the RF pulses. The difference in depolarisation between the two images is then due to ^{129}Xe transfer between the gas space and lung parenchyma. Ruppert et al. used animal models to optimise the pulse sequences and improve sensitivity, allowing quantitative information to be extracted [195, 196]. XTC has also been implemented with a 3D acquisition and temporal resolution by repeating with multiple delay times (MXTC) [197, 198]. Two lung function parameters can be derived from MXTC, relating to tissue to alveolar-volume ratio and alveolar septal thickness. The alveolar septal thickness parameter has been found to be elevated in a small number of asthma and COPD patients [197].

The need for multiple breaths in XTC can introduce errors. To address this limitation, a single-breath technique has since been developed [41, 199]. When a 90° saturation pulse is used, XTC yields very similar results to CSSR, allowing regional measurement of S/V [199]. XTC and CSSR show good agreement at 0.2 T, although the information available from XTC is limited because it does not resolve the dissolved-phase signal into its RBC and M compartments [41]. Amzajerdian et al. adapted the XTC method by using narrow bandwidth RF pulses to only image one of the compartments, generating RBC or M specific depolarisation maps [200]. This increases the scanning time, but allows ratio maps to be made and significant differences were found between healthy and diseased lungs. However, the XTC results are flip angle and delay time dependent, so cannot be easily compared with ratios calculated from other imaging techniques. Another avenue explored in Ref. [200] is a free-breathing modification of XTC. This may be better tolerated by diseased patients and reduce the anaesthetic effects of ^{129}Xe gas, but also involves a more complicated gas delivery and analysis. Recently, the capability of MXTC to simultaneously evaluate ventilation and gas exchange was demonstrated [201].

3.4 Analytical Models of Gas Exchange

Modelling pulmonary gas exchange allows quantitative information regarding gas transfer and lung morphometry to be extracted from experimental data acquired with CSSR. Currently, analytical models utilising the 1D diffusion equation (3.4.1) are most widely used for modelling gas exchange imaging, although some numerical models have been developed (see Appendix B).

$$\frac{\partial M_d(x,t)}{\partial t} = D\frac{\partial^2 M_d(x,t)}{\partial^2 x}, \tag{3.4.1}$$

where M_d and D are the dissolved-phase ^{129}Xe magnetisation density and diffusion coefficient respectively. The magnetisation density can be treated as equivalent to ^{129}Xe concentration, if complete saturation of the dissolved-phase is achieved at t = 0 and the only process affecting the dissolved-phase magnetisation is assumed to be the diffusion of ^{129}Xe into the dissolved-phase from the gas-phase, i.e. M$_0$[Xe] → [Xe].

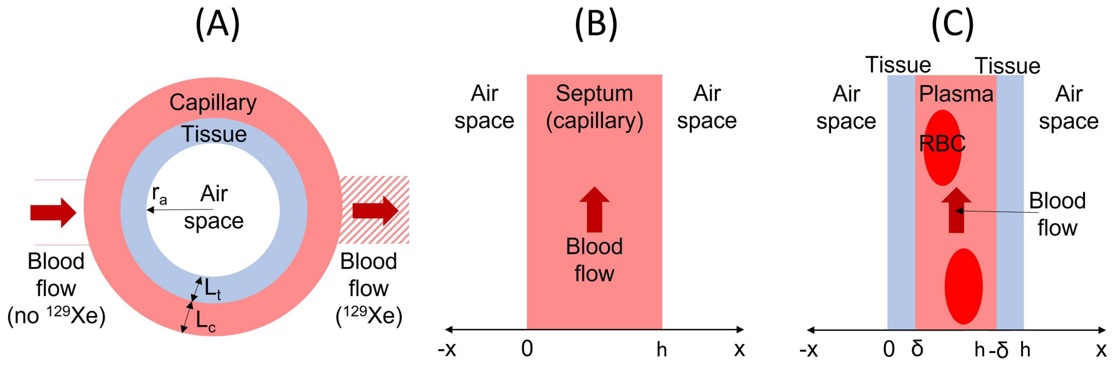

Figure 3.4.1: The geometries for the analytical models of gas exchange: (A) Månsson model, (B) Patz model and (C) MOXE, where r$_a$ is the alveolar radius, L$_t$ and L$_c$ are the tissue and capillary compartment thicknesses respectively, *h* is the septal thickness and *δ* is the thickness of the tissue compartment.

3.4.1 Månsson Model

Hyperpolarised ^{129}Xe gas transfer was first modelled by Månsson et al. using a radially-symmetric model which included three compartments: gas, tissue and capillary blood [202], as shown in Figure 3.4.1A. The dissolved-phase magnetisation increases with time as gaseous ^{129}Xe diffuses across the tissue barrier, then plateaus as the capillary blood is saturated. If the blood was stationary, the magnetisation and hence signal would remain constant, as demonstrated by the dashed line in Figure 3.4.2A; instead it continues to increase due to perfusion (solid line in Figure 3.4.2A) [40, 63]. This is referred to as the 'blood flow effect'. The flow of fresh blood into the gas exchange region means that more ^{129}Xe can dissolve, which increases the total signal because blood outside and downstream of the gas exchange region contributes to the signal as long as it is still within the ^{129}Xe coil and TR $< T_1$. The Månsson model captures this behaviour with the following equation, which is fit to the RBC and M signals separately:

$$S(t) = S_0\left(1 - e^{-\frac{t}{T}}\right) + S_1 t, \tag{3.4.2}$$

where S_0 is the value of the linear component at t = 0, T is the time constant of the exponential component and S_1 is the gradient of the linear component. From the values of S_0 and S_1, the tissue thickness, capillary thickness, haematocrit, perfusion and mean transit time can be found.

3.4.2 Patz Model

A 1D slab geometry can also be used to model the gas transfer [203, 142]. The Patz model [142], shown in Figure 3.4.1B, treats the septum between two alveoli as consisting of a capillary compartment with plug blood flow perpendicular to the

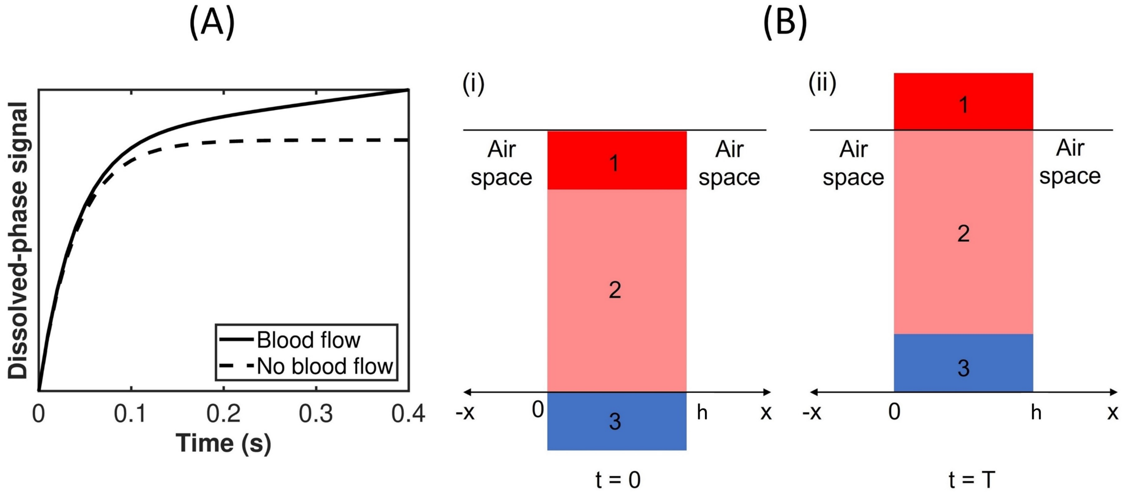

Figure 3.4.2: (A) The dissolved-phase signal simulated using the Månsson model, which first increases exponentially and then linearly with time, because of the blood flow effect: (B)(i) At t = 0, ^{129}Xe dissolves in the blood plasma and RBCs in regions 1 and 2. Region 3 contains no ^{129}Xe. (ii) At time T, blood in region 1 still contains ^{129}Xe, but flows out of the gas exchange region. Region 2 is saturated, so no more ^{129}Xe dissolves here. Fresh blood in region 3 flows into the gas exchange region and more ^{129}Xe can dissolve in this region. Adapted from Ref. [142].

diffusion direction. An alternative method is used to deal with the dissolved ^{129}Xe which flows out of the gas exchange region to the method used in the Månsson model. The blood is separated into the three regions shown in Figure 3.4.2B and the signal contribution from each found separately. This model can be used to extract three clinically relevant parameters; S/V, h and τ, the latter of which is the average time a RBC spends in the gas exchange region. The set of equations defining the Patz model are:

$$F(t) = F_0 + \frac{\lambda h}{2}\frac{S}{V}\left(\frac{\tau - t}{\tau}\right)f(q) + \lambda h\frac{S}{V}\left[\frac{t}{\tau} + \frac{8h^2}{D\pi^4\tau}g(q)\right], \tag{3.4.3}$$

where

$$f(q) = 1 - \sum_{n,odd}\frac{8}{\pi^2 n^2}\exp(-q\pi^2 n^2), \tag{3.4.4}$$

$$g(q) = \sum_{n,odd}\frac{1}{n^4}(\exp(-q\pi^2 n^2) - 1) \tag{3.4.5}$$

and

$$q = \frac{Dt}{h^2}. \tag{3.4.6}$$

F_0 is an offset term, which accounts for imperfect dissolved-phase saturation, λ is the xenon Ostwald solubility and D is the dissolved ^{129}Xe diffusion coefficient.

3.4.3 MOXE

Chang's Model of Xenon Exchange (MOXE) builds upon the Patz model to separate the dissolved-phase into the RBC and tissue compartments (Figure 3.4.1C) [204]. Two coupled equations are used to model the M and RBC ^{129}Xe signal (see [204]), from which S/V, h, τ and haematocrit can be estimated. The MOXE improves on the Månsson model by including the S/V ratio and treating the M and RBC signals as being correlated.

Stewart et al. used the three models described above to estimate S/V, h and τ from ^{129}Xe CSSR data [39]. The goodness of fit to the experimental data was found to be better for the Patz model and MOXE, while the Månsson model was hampered by its mono-exponential nature, resulting in less accurate fitting. The MOXE model appeared to be the most accurate, although the high number of fit parameters that are dependent on one another sometimes resulted in unrealistic fit parameters.

3.4.4 Other Models of Gas Exchange

The Patz model and MOXE are the most widely used analytical models of ^{129}Xe gas exchange, however, other models have been explored in recent years. Kern et al. have used a generalisation of the Patz model, which allows for a different diffusivity in the membrane and blood compartments, to assess membrane permeability (κ) [205]. This model enabled the detection of reduced κ in patients with PH when compared to healthy volunteers. In principle, Kern's model can also be used to estimate the capillary diameter, however in practise, noise in the CSSR uptake data meant that to determine κ, the capillary diameter had to be fixed.

A kinetic model of ^{129}Xe exchange has been developed by Taylor et al., which uses a similar approach to tracer kinetic modelling of DCE-MRI [206]. This model treats the lung as consisting of three compartments: alveolar, tissue and blood, which have ^{129}Xe concentrations C_a, C_t and C_b respectively. ^{129}Xe diffuses from the tissue and blood compartments from the alveolus with rate constants k_1 and k_2. A venous exit function is used to model the blood concentration leaving the lung. An advantage of this model is that blind estimation techniques are used to extract the kinetic parameters, without requiring separate M and RBC spectroscopic peaks. This is useful for characterising gas exchange in mice and rabbits in pre-clinical research, because these species generally do not exhibit separate M and RBC peaks, and for low-field MRI, where the frequency shift between the peaks is reduced [207].

3.5 Cardiogenic Oscillations of the ^{129}Xe Red Blood Cell Resonance

It is possible to extract information about pulmonary microvascular function from the ^{129}Xe dissolved in the RBCs using dynamic spectroscopy and gas exchange imaging. Periodic oscillations in the ^{129}Xe RBC signal were first observed by Venkatesh et al. [208] and have since been demonstrated to be a potential biomarker of cardiopulmonary disease [209]. Comparison to pulse oximetry and ECG has demonstrated that the frequency of the oscillations matches the heart rate [37, 90, 137] and that the RBC signal maxima and minima occur at ventricular systole and diastole respectively [210]. The ^{129}Xe RBC signal is localised to ^{129}Xe dissolved in the RBCs in the pulmonary capillary bed, due to the choice of TR and flip angle. A short (∼15 ms) TR is used, which is much smaller than the ∼750 ms RBC transit time, and a relatively large (∼20°) flip angle ensures that magnetisation from ^{129}Xe that has travelled into the larger vasculature is destroyed (and is not replenished by gas transfer from the alveoli). Therefore, the source of the oscillations has been ascribed to fluctuations in capillary blood volume over the cardiac cycle, related to the propagation of the cardiac pulse wave to the pulmonary capillaries, as illustrated in Figure 3.5.1.

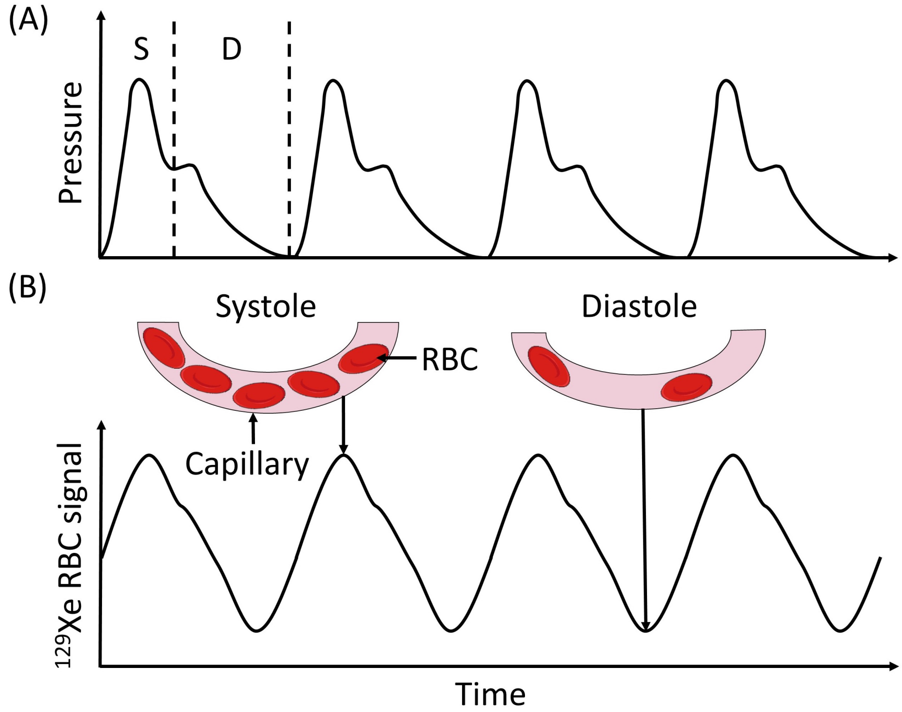

Figure 3.5.1: Illustration of the proposed origins of the ^{129}Xe RBC signal oscillations. (A) The pumping motion of the heart generates compression and relaxation pressure waves, which superimpose to create the cardiac pulse wave [211]. The pulse wave propagates through the arterial network to the pulmonary capillaries, driving changes in blood pressure, flow and volume. (B) At systole (S), the heart muscles contract and deoxygenated blood is pushed into the pulmonary capillaries. The increase in the volume of capillary blood means that there are more RBCs available for the dissolved ^{129}Xe atoms to form bonds with, and hence the ^{129}Xe RBC signal increases, reaching a maximum at end-systole. At diastole (D), the heart muscles relax, and the capillary blood volume, and therefore the ^{129}Xe RBC signal, reach a minimum. Waveforms not drawn to scale.

The RBC oscillations can be detected with whole lung dynamic spectroscopy, by acquiring FIDs at ∼ 15 ms intervals, characterising the Fourier transformed time-domain data with spectral lineshapes, and quantifying the temporal variation in the amplitude of the RBC peak [137]. Signal oscillations have also been observed in CSSR uptake curves [212] and can be extracted from the RBC signal from the centre of 3D radial k-space (k_0) from gas exchange imaging [35, 213]. This process (illustrated in Figure 3.5.2) is advantageous because it does not require a separate dedicated scan.

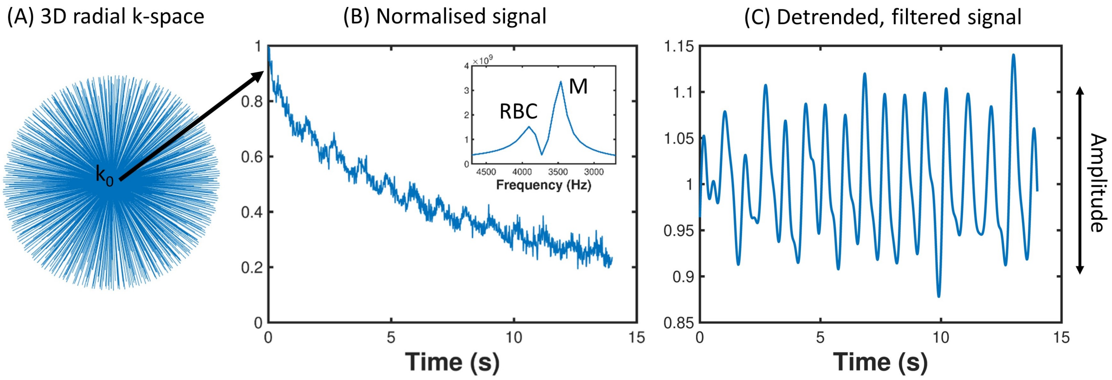

Figure 3.5.2: The RBC oscillations are encoded in radial dissolved-phase ^{129}Xe imaging: (A) the RBC signal from the centre of 3D radial k-space (k_0) is (B) normalised and (C) corrected for T_1 and RF pulse induced decay, to isolate the signal oscillations. The oscillation amplitude is defined as the peak-to-peak difference.

The amplitude of the RBC signal oscillations is influenced by both pre-capillary and post-capillary vessel impedance, and changes in capillary blood volume. In diseases affecting the pre-capillary bed vasculature, such as PAH, oscillation amplitude is decreased, whereas in diseases affecting the post-capillary bed network, oscillation amplitude is increased. Table 3.5.1 summarises some of the cardiopulmonary diseases for which RBC oscillation mapping has been able detect amplitude changes.

A physiological model was recently proposed, using an electric circuit analogy, to help better understand the competing contributions of vessel impedance and blood volume to the oscillation amplitude [214]. The electric circuit model is shown in Figure 3.5.3. The heart was modelled as a current source driving blood flow Q through the pulmonary vasculature. Pulmonary vascular resistance (PVR) was modelled as a series of three impedances, each with a resistance and compliance; these represent the arterial, capillary and venous contributions. Including vessel compliance in the model allows for the variation in capillary blood volume that drives the RBC signal oscillations. In Costelle et al.'s framework, the oscillation amplitude, α is given by:

$$\alpha = \frac{C_C}{\bar{V_C}} \Delta Q (Z_C + Z_V), \tag{3.5.1}$$

where C_C is the capillary compliance, $\bar{V_C}$ is the mean capillary blood volume, ΔQ is the change in blood flow between systole and diastole, and Z_C and Z_V are the capillary and venous impedances respectively. This model was applied to a group of 129 non-PH lung disease patients and used to estimate the contributions to PVR from the arterial, capillary and venous circulation, which were found to be 70%, 11% and 19%, respectively. In a second group of subjects who were being evaluated for PH, RBC oscillation amplitude was corrected for estimated pulmonary blood volume and the corrected amplitude was found to correlate with RHC-measured PVR.

Table 3.5.1: Summary of how the RBC oscillation amplitude is affected by different pulmonary disorders. LHF = left heart failure.

Disease	Oscillation amplitude	Explanation	Reference(s)
IPF	Increased	A preserved cardiac stroke volume is delivered to a reduced capillary bed, causing larger relative changes in blood volume.	Wang et al. [37], Bier et al. [137]
NSIP	Increased	A preserved cardiac stroke volume is delivered to a reduced capillary bed, causing larger relative changes in blood volume.	Mummy et al. [169]
LHF	Increased	High post-capillary impedance increases the capillary blood volume and causes pooling in the capillary bed.	Wang et al. [37]
COPD	Decreased	Emphysematous lung tissue destruction.	Wang et al. [37]
PAH	Decreased	High pre-capillary impedance.	Wang et al. [37]
CTEPH	Decreased	High pre-capillary impedance as a result of vessel occlusion.	Lu et al. [215]

A limitation of the whole-lung RBC oscillation amplitude as a potential biomarker of cardiopulmonary disease is that it lacks spatial sensitivity. This makes it difficult to detect oscillation abnormalities in patients with combined pre- and post-capillary PH and patients with both pre-capillary PH and RBC transfer defects [215]. To measure the RBC oscillations regionally, a technique was developed by Niedbalski et al. which uses a keyhole reconstruction [216] to spatially resolve the oscillation amplitude [213]. This method is discussed in detail in Chapter 7. The mean oscillation amplitude obtained from the maps was significantly higher in patients with IPF and lower in patients with PAH in comparison to healthy volunteers.

This work has since been followed up by Lu et al., who used digital phantom simulations to optimise Niedbalski's method [215]. They evaluated their method in healthy volunteers to establish healthy reference values and in patients with CTEPH to assess microvascular function both before and after pulmonary endarterectomy (PEA). No significant differences were found in the average oscillation amplitude before and after PEA, but the number of lung voxels which were classified as low oscillation amplitude decreased following PEA.

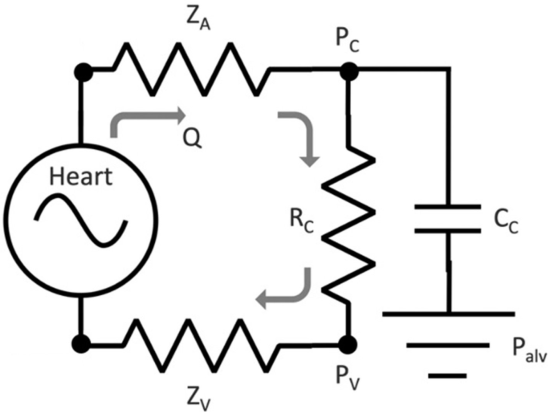

Figure 3.5.3: Costelle et al.'s electric circuit model of the pulmonary circulation. The heart is modelled as a current source with variable blood flow rate Q, which is higher at systole than diastole. The PVR is divided into arterial, capillary and venous impedances (Z_A, Z_C and Z_V, respectively) connected in series. The capillary bed has a resistance R_C and a capacitance C_C, which represents compliance and allows for blood storage between systole and diastole. C_C is proportional to the pressure drop across the capacitor, which is the difference between the capillary bed pressure P_C and the alveolar pressure P_{alv}, which acts as 'ground'. The pressure the other side of the capillary bed is the venous pressure, P_V. Copyright Costelle et al. 2025, licensed under Creative Commons Attribution CC-BY 4.0, adapted with permission from [214].

Chapter 4

Overview of Experimental Methods

This chapter consists of an overview of the MRI hardware and both the theoretical background and methodological details for the lung imaging pulse sequences used throughout this thesis.

4.1 MRI Hardware

All experiments described in this thesis were performed on either a 1.5 T GE HDx or Artist whole-body clinical scanner with a flexible transmit-receive vest coil (Clinical MR Solutions) tuned to the ^{129}Xe Larmor frequency.

^{129}Xe doses were hyperpolarised to 10 - 30% polarisation using a SEOP polariser (POLARIS, University of Sheffield), which is described in [217] and shown in Figure 4.1.1. SEOP took place within a 3530 mL (7.5 cm diameter, 80 cm length) glass optical pumping cell, pressurised to 2 bar. The cell is contained within a ceramic oven, which was heated to 125°C to vaporise the Rb, and a Helmholtz coil pair, which generated the magnetic field which induces Zeeman splitting. A gas mixture of 3% Xe, 10% N_2 and 87% ^{4}He was flowed through the cell at a rate of 2000 sccm. Rb optical pumping was achieved using a 150 W laser (BrightLock 200 W, QPC, CA, USA) and a wavelength of 794.77 nm.

Unless stated otherwise, subjects inhaled doses of 86% enriched hyperpolarised ^{129}Xe from a Tedlar bag, starting from FRC. The ^{129}Xe dose volumes for ventilation and gas exchange imaging were generally chosen based on participant height, according to the dose chart shown in Table 4.1.1. The rationale for these volumes is described in Smith et al. [218]: literature values of predicted FRC and TLC volumes in children and adults were used to calculate the FRC volumes such that patients would not be near their TLC when they inhaled the FRC dose. The FRC dose was chosen such that the total lung volume (FRC + bag volume) would be approximately 60% of TLC. This approach of adjusting the dose volume is more simplistic than the currently recommended approach [164] of using subject-specific measured or predicted TLC or FRC volumes to calculate the individual dose volume, but was consistent across all data used in this thesis.

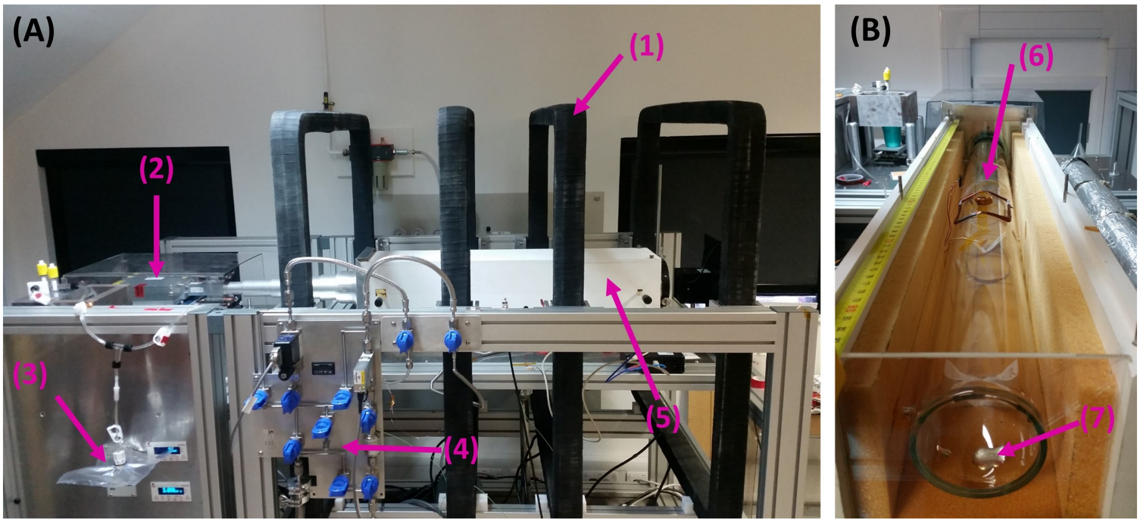

Figure 4.1.1: (A) Photo of the ^{129}Xe SEOP polariser used in this work. (1) Helmholtz coils, (2) laser diode array, (3) Tedlar bag, (4) gas manifold, (5) oven. (B) The oven with the lid off, showing the (6) SEOP cell enclosed within and (7) the pool of rubidium.

Table 4.1.1: ^{129}Xe dose chart for ventilation and gas exchange xenon imaging according to patient height.

Height (cm)	Ventilation Imaging	Gas Exchange Imaging
> 160	500 ml ^{129}Xe 500ml N_2	1000 ml ^{129}Xe 0ml N_2
150 - 160	450 ml ^{129}Xe 350ml N_2	800 ml ^{129}Xe 0ml N_2
140 - 150	400 ml ^{129}Xe 250ml N_2	650 ml ^{129}Xe 0ml N_2
130 - 140	350 ml ^{129}Xe 150ml N_2	500 ml ^{129}Xe 0ml N_2
120 - 130	300 ml ^{129}Xe 100ml N_2	400 ml ^{129}Xe 0ml N_2

4.2 MR Pulse Sequences for Hyperpolarised ^{129}Xe Lung Imaging

4.2.1 Calibration

Prior to ^{129}Xe imaging or spectroscopy, it is necessary to calibrate both the transmit gain (TG), so that the desired flip angles are reached, and the transmit/receive RF frequency. These are dependent on the subject's coil loading and the local magnetic environment in the lungs. Whereas for proton imaging, an automatic prescan procedure can be used to do this, for multi-nuclear MRI, a separate calibration scan is usually required. Schulte et al. devised a calibration method based on the Bloch-Siegert effect [219]: where an off-resonant RF field induces a phase shift proportional to the square of the RF field amplitude in spins but does not cause them to precess [220]. An on-resonance excitation pulse excites the

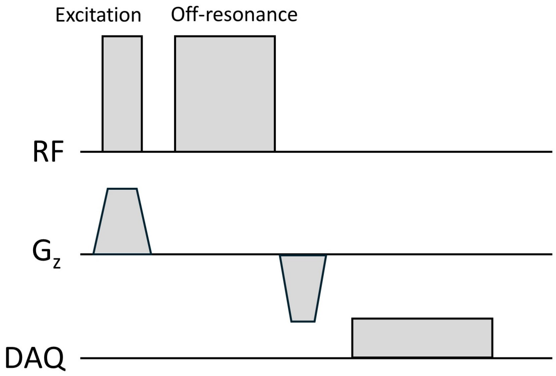

Figure 4.2.1: The pulse sequence used for Bloch-Siegert transmit gain calibration. A slice-selective excitation pulse and gradient are followed by an off-resonance pulse. Slice-refocusing is performed with a second gradient of opposite polarity to the first.

spins into the transverse plane and is followed by an off-resonance pulse which induces the Bloch-Siegert phase shift:

$$\phi_{BS} = \int_0^T \omega_{BS}(t)\mathrm{d}t = \int_0^T \frac{[\gamma B_1^+(t)]^2}{2(\omega_{RF} + \omega_{B0})}\mathrm{d}t \tag{4.2.1}$$

where ω_{BS}, ω_{RF} and ω_{B0} are the frequencies of the Bloch-Siegert precession, the off-resonance excitation pulse and the B_0 inhomogeneity-induced precession, respectively. T is the time period of the off-resonance pulse and γ is the gyromagnetic ratio of the nucleus. ϕ_{BS} is dependent on the magnitude of the transmit RF field, $B_1{}^+$. To rule out any other phase deviation, for example from eddy currents, the off-resonance excitation is repeated at $\pm\omega_{RF}$ and ω_{BS} is half the phase difference between the two acquisitions.

The pulse-acquire sequence used to measure ϕ_{BS} is shown in Figure 4.2.1. This sequence was repeated for $\pm\omega_{RF}$. 32 FIDs were acquired, with TR = 75 ms, flip angle = 10°, 1024 spectral points, bandwidth = 20 kHz and initial guesses TG = 150 and on-resonant frequency = 17,676,450 Hz were used. A dose of ∼700 ml gas acquired straight from the optical pumping cell (i.e. the ^{129}Xe was not cryogenically separated from the buffer gases) was used.

ϕ_{BS} was then calculated from the time-domain FID data (d) acquired from the $\pm\omega_{RF}$ scans:

$$\phi_{BS} = \frac{1}{2}\arctan\left[\frac{\Im(d^*_{-\omega RF} \cdot d_{+\omega RF})}{\Re(d^*_{-\omega RF} \cdot d_{+\omega RF})}\right] \tag{4.2.2}$$

where d^* is the complex conjugate of d. The required change in TG for an excitation of flip angle θ is then given by:

$$\Delta\mathrm{TG} = -20\log_{10}\left(\frac{B^+_{1,peak}}{B^+_{1,desired}}\right) \tag{4.2.3}$$

where $B_{1,desired}^{+}$ is the peak RF amplitude for an excitation of flip angle θ and $B_{1,peak}^{+}$ is the peak off-resonance transmit RF field, defined as:

$$B_{1,peak}^{+} = \sqrt{\frac{\phi_{BS}}{K_{BS}}} \tag{4.2.4}$$

where K_{BS} is a phase shift constant which is dependent on the RF pulse shape, ω_{RF} and T.

The centre frequency, f_0, can be calculated from the same data using the centre-of-gravity around the ^{129}Xe spectral peak.

4.2.2 Ventilation Imaging

No prospective ventilation imaging was performed as part of this thesis, however, ventilation images and metrics were analysed retrospectively in Chapters 5 and 8. These were acquired using a 3D bSSFP sequence [97]. Briefly, bSSFP uses phase encoding along the y and z directions with balanced gradients and a short TR such that TR $< T_2^*$ of the gaseous ^{129}Xe ($\sim$25 ms in partially-inflated lungs [87]), such that some ^{129}Xe magnetisation remains at the end of each phase encode. This magnetisation is not destroyed by spoiling, as in SPGR, but instead is ‘recycled’ at the next phase encode, permitting a higher SNR. This leads to the generation of a steady-state transverse magnetisation.

SSFP imaging was performed at end-inspiratory tidal volume with a maximum dose of 0.5L ^{129}Xe and 0.5L nitrogen (volume adjusted according to height; see Table 4.1.1), FOV = 40 cm, phase FOV = 0.8 or 1.0, 22 - 26 coronal slices with 10 mm slice thickness in the anterior-posterior direction, 256 x 256 reconstructed matrix size (except for the CTEPH patients in Chapter 8, for whom a 512 x 512 reconstructed matrix size was used), flip angle = 10°, TE/TR $\sim$ 1.4/4 ms. In Chapter 5, the patients with asthma and/or COPD were scanned with TE/TR = 2.1/6.5 ms.

Co-registered ^{1}H anatomical images were acquired in order to derive the total lung volume (TLV). Imaging was performed at breath hold following the inhalation of a bag of room air, in order to achieve the same lung inflation as the ^{129}Xe imaging. An SPGR sequence was used, with the same FOV as the ^{129}Xe sequence, but double the number of slices, TE/TR = 0.6/1.4 ms and flip angle = 5°.

From these two images, the ventilation defect percentage (VDP) and ventilated volume (VV) were derived, which are defined as:

$$VDP = 1 - \frac{VV}{TLV} \mathrm{x} 100\% \tag{4.2.5}$$

4.2.3 Gas Exchange Imaging

As discussed in Section 3.3.2, there are several methods for ^{129}Xe gas exchange imaging. In this thesis, a 3D radial MESI sequence with four echoes [35, 189], was used to acquire gas exchange images. This technique is based on a previous method for chemical shift ^{13}C metabolic imaging [221].

4.2.3.1 Spectroscopic Image Encoding

The spectral encoding of the ^{129}Xe gas, M and RBC resonances is achieved using an acquisition scheme with four gradient echoes. The time between the m^{th} excitation and the acquisition of the n^{th} k-space data point at position k_n is different for each echo, and is defined as:

$$\tau_{m,n} = m\Delta\mathrm{TE} + t_n, \tag{4.2.6}$$

for the m^{th} echo, where ΔTE is the echo time spacing. After discretisation of the spatial resonance frequencies of the q compartments, ω_q, and spatial dimension $\boldsymbol{r_p}$, the signal acquired at time $\tau_{m,n}$ can be written:

$$y_{m,n} = \sum_{p,q} \exp(i\boldsymbol{k_n}\boldsymbol{r_p})\exp\left(i\omega_q - \frac{1}{T^*_{2,q}}\tau_{m,n}\right)x_q(\boldsymbol{r_p}) \tag{4.2.7}$$

where $x_q(\boldsymbol{r_p})$ is the image space distribution of the q^{th} resonance. The effects of magnetic field inhomogeneities and T_1 and RF pulse induced decay are not included in this model. Neglecting the T_2^* decay during the acquisition time of each echo, the signal can be represented as:

$$y_{m,n} = \sum_q E_{m,q} e^{i\omega_q t_n} \xi_q(\boldsymbol{k_n}) \tag{4.2.8}$$

where $\xi_q(\boldsymbol{k_n})$ represents the k-space distribution of the q^{th} resonance:

$$\xi_q(\boldsymbol{k_n}) = \sum_p e^{i\boldsymbol{k_n}\boldsymbol{r_p}} x_q(\boldsymbol{r_p}) \tag{4.2.9}$$

and $E_{m,q}$ is the chemical shift encoding matrix:

$E_{m,q} = \exp\left(i\omega_q - \frac{1}{T^*_{2,q}}\mathrm{TE}_m\right).$

$y_{m,n}$ can be described as a linear combination of the $\xi_q(\boldsymbol{k_n})$ for each resonance, each with a weighting given by $E_{m,q}$. To reconstruct the k-space data, the first step is matrix inversion of $E_{m,q}$ to reconstruct each $\xi_q(\boldsymbol{k_n})$:

$$\xi_q(\boldsymbol{k_n}) = e^{-i\omega_q t_n}(E^{\dagger} y_n)_q \tag{4.2.10}$$

where † represents the Moore-Penrose pseudo-inverse. Then, the $\xi_q(\boldsymbol{k_n})$ for each resonance are reconstructed spatially using gridding (see *Image Reconstruction* below). More accurate estimation of the spectral signals can be achieved by pre-conditioning E using prior knowledge of of the resonant frequencies ω_q and $T^*_{2,q}$ of the resonances obtained from calibration spectra. The matrix E needs to be well-conditioned to minimise noise amplification resulting from the chemical shift inversion in Equation 4.2.10 [180]. This can be done by optimising the choice of ΔTE such that the system of equations are as independent as possible (i.e. the matrix condition number is minimised). As mentioned in Section 3.3.2.3, the optimisation is typically performed using the NSA approach, which describes the SNR as a function of ΔTE. The NSA is calculated as:

$$NSA_q = \frac{1}{(E^H E)^{-1}_{q,q}} \tag{4.2.11}$$

for each resonance, where H denotes the complex conjugate transpose. The optimal ΔTE is determined by maximising the combination of NSA_q values, whilst taking into account T_2^* signal decay.

4.2.3.2 Image Acquisition

In this thesis, ^{129}Xe gas exchange imaging data were both retrospectively analysed and prospectively acquired. Some of the retrospective data were acquired with different imaging parameters because an older version of the sequence was used. The original implementation of the sequence is described in Ref. [35]. The updated sequence, used for most of the retrospective data and all prospective data in this thesis, is described in Ref. [189]. As well as the differences in some parameters, which are listed in Table 4.2.1, there were some other changes between the two versions. The MESI pulse sequence diagram shown in Figure 3.3.5 represents the newer version.

In the first version of the sequence, a spoiled interleaved RF excitation of gas-phase and dissolved-phase ^{129}Xe was used. The RF pulse was an amplitude-modulated composite pulse, designed to selectively excite the dissolved-phase with minimal excitation of the gas-phase magnetisation [98]. This was replaced by a frequency-tailored RF pulse with 1% excitation of the gas-phase in the updated sequence, which removed the need for interleaved gas and dissolved-phase excitation. Another difference between the two versions was that, whereas the original sequence necessitated a separate acquisition to collect calibration spectra, these were integrated into the start of the updated sequence. In version one of the sequence, a whole-lung spectrum was acquired using a separate breath hold and 600 ml of ^{129}Xe. The other sequence parameters were: bandwidth = 8 kHz, 512 points, 200 averages, TR = 74 ms, flip angle = 22°. At the beginning of the sequence, 20 dummy RF pulses were implemented to deplete the downstream signal from ^{129}Xe in the pulmonary veins. In version two, FIDs were acquired for the 20 dummy pulses and used as the calibration spectra to precondition the chemical shift encoding matrix.

Both sequences used a 3D radial k-space trajectory and had a total scan time of ∼14 s. Fewer radial spokes were acquired within this time for the original sequence in comparison to the updated sequence, because of the interleaved acquisition scheme.

Table 4.2.1: Comparison of imaging parameters between the two versions of the MESI gas exchange imaging sequence.

	Version 1 [35]	Version 2 [189]
FOV (cm)	40	
Acquired (reconstructed) matrix size	20x20x20 (32x32x32)	
Bandwidth (kHz)	31.25	
ΔTE (ms)	0.7	
TE_1 (ms)	0.571	
Dissolved-phase flip angle (°)	40	22
Gas-phase flip angle (°)	0.7	0.2
RF pulse width (ms)	0.8	1.2
TR (ms)	40	15
$TR_{90°,equiv}$ (ms)	171	206
n_{spokes}	332	934
$n_{samples}$	11	13
Scan time (s)	14	

$TR_{90°,equiv}$ is a parameter derived from the principle of flip angle/TR equivalence, which is the idea that different combinations of TR and flip angle can sample the same part of the dissolved ^{129}Xe uptake curve [222]. For any TR and flip angle combination, $TR_{90°,equiv}$ is the TR which would result in an equivalent measurement for a 90° dissolved-phase flip angle, and is defined as:

$$TR_{90°,equiv} = \frac{TR}{1 - \cos\theta} \tag{4.2.12}$$

where θ is the dissolved-phase flip angle. This relationship was derived empirically from experiments in rabbits [223].

4.2.3.3 Image Reconstruction

The calibration spectra were used to estimate subject-specific resonant frequencies and T_2^* of the ^{129}Xe dissolved in the alveolar airspace, membrane, and capillaries. This was done with a triple Lorentzian fit in the frequency domain [134]. The gas, RBC and M resonances were separated in k-space using the matrix inversion described in Equation 4.2.10 and prior knowledge of the chemical shifts and T_2^* of the resonances obtained from the calibration spectra [35, 54]. The 3D radial reconstruction was performed using gridding, with an analytical density compensation function to correct for the non-uniform sampling density of the radial trajectory (Section 2.2.7). The gridding kernel was chosen using the optimal Kaiser-Bessel convolution formula [57]. The kernel had an oversampling ratio of 1.2 and width of 8.

4.3 Clinical Studies

Data from several clinical studies were analysed as part of this thesis, which are summarised in Table 4.3.1.

Table 4.3.1: Details of the clinical studies which feature in this thesis.

Study Name	Disease(s)	REC Number/Ethics	Chapters featured in
ADPro	Asthma, COPD	16/EM/0439	5
Clinical Referral	CTEPH	STH18877	8
EXPLAIN	Long-Covid and healthy	21/SC/0398	5, 7, 8
HRUK	SSc, SSc-PAH	12/NE/0355	9
MICHAL	Healthy	17/LO/0725	6, 7, 8
MURCO/PC-XMAS	Post-COVID-19 lung disease	9/LO/1115	5, 6, 8
SUMMER	COPD	22/NW/0009	6
PHOENIX	PAH	23/NE/0067, 12/NE/0355	9
UoS Healthy Volunteers	Healthy	UOS030529, UOS052024	6, 7, 8, 9

Chapter 5

A Framework for Modelling Whole-Lung and Regional TL_{CO} Using Hyperpolarised ^{129}Xe Lung MRI

5.1 Introduction

Pulmonary gas exchange function is usually evaluated using the TL_{CO} PFT, which was described in Section 3.2.1. Although widely used [110], TL_{CO} measurement is a breathing test, measured at the mouth, and therefore lacks regional gas exchange information. Hyperpolarised ^{129}Xe lung MRI is an alternative method to quantify pulmonary gas transfer and the sensitivity of the ^{129}Xe gas exchange ratios, RBC:M, RBC:Gas and M:Gas has been demonstrated in COPD [38, 37, 212], ILD [35, 169, 33] and post-COVID-19 lung disease [193, 194, 190]. However, these gas exchange ratios lack a well-defined conventional physiological interpretation.

There are parallels between the measurements made with ^{129}Xe MRI and the constituent components of TL_{CO}, illustrated in Figure 5.1.1. A model has been proposed by Wang et al. to exploit these similarities by using features derived from ^{129}Xe ventilation and gas exchange imaging, along with V_A and K_{CO} from PFT, to predict TL_{CO} using linear regression [224]. In this model, M:Gas and RBC:Gas are used to estimate K_{CO} via linear regression coefficients. To convert K_{CO} to TL_{CO}, V_A is estimated from its linear regression with VV. This approach allows the underlying components of TL_{CO} to be estimated, which may clarify the interpretation of a reduced TL_{CO} and reveal signatures of different cardiopulmonary diseases. However, the generalisability of this model is not clear, because the same data (from 41 healthy subjects and 101 lung disease patients) were used to both train the model and test its performance, leading to potentially biased results. Most commonly-used TL_{CO} prediction models are based on participant demographics such as age and sex [225, 226, 227, 228]. Indeed, the model-based prediction of TL_{CO} from ^{129}Xe MRI metrics may be improved by considering age and sex, as both affect ^{129}Xe MRI gas exchange metrics

[229, 230, 189]. Furthermore, previous work [224] used the whole-lung average metrics from ^{129}Xe imaging to predict whole-lung TL_{CO}, which does not utilise the regional information offered by imaging. Here we propose that inputting the imaging maps to the predictive models instead would allow for regional visualisation of TL_{CO}. Prediction of TL_{CO} may be further improved by using machine learning regression algorithms instead of classical least-squares based methods.

The objectives of this work were therefore to: 1) evaluate Wang et al.'s model [224] in a large cohort of asthma and COPD patients with a rigorous testing and training group validation strategy, and 2) build upon this model to predict both whole-lung and regional maps of TL_{CO} using ^{129}Xe imaging and participant demographic data with machine learning algorithms.

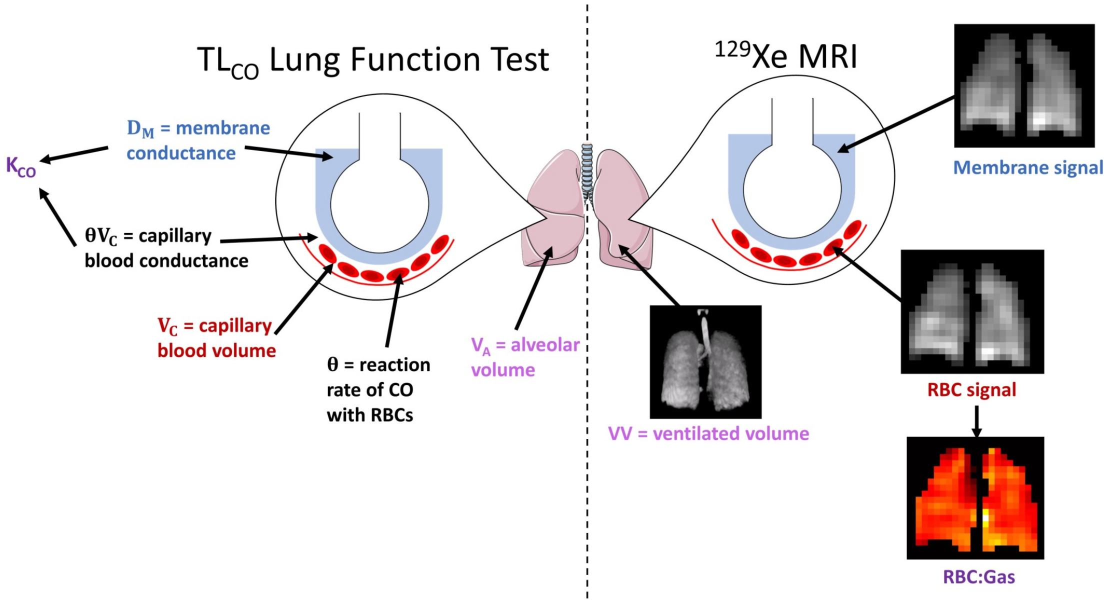

Figure 5.1.1: Schematic of the parallels between the underlying physiology measured by the TL_{CO} lung function test and ^{129}Xe MRI. Like the membrane conductance (D_M), the ^{129}Xe membrane signal is dependent on the surface area and thickness of the alveolar membrane. The ^{129}Xe RBC signal is influenced by both the gas exchange across the alveolar membrane and the capillary perfusion, so can be linked to the capillary blood volume, V_C. RBC:Gas measures the transfer of gas from the alveoli, across the alveolar membrane and into the RBCs, so is analogous to the transfer coefficient, K_{CO}. VV is the volume of the lung where ^{129}Xe signal is detected, which is alike to the alveolar volume, V_A. (This figure was partly generated using Servier Medical Art, provided by Servier, licensed under a Creative Commons Attribution 3.0 unported license.)

5.2 Theory

TL_{CO} represents the overall conductance of CO from the alveolar gas to the pulmonary capillary blood and is made up of two components:

$$\frac{1}{TL_{CO}} = \frac{1}{D_M} + \frac{1}{\theta V_C} \tag{5.2.1}$$

where D_M is the membrane conductance and θV_C is the capillary blood conductance: θ is the reaction rate of CO with the RBCs and V_C is the capillary blood volume [231]. $\mathrm{TL_{CO}}$ also depends on the volume of the alveoli available for gas exchange, $\mathrm{V_A}$:

$$\mathrm{TL_{CO}} = \mathrm{K_{CO}V_A} \tag{5.2.2}$$

where $\mathrm{K_{CO}}$ is the CO transfer coefficient [109].

Wang et al.'s model [224] uses the following linear regression equations to predict whole-lung $\mathrm{TL_{CO}}$, where k_V, k_M and k_R are coefficients found by fitting the model to measured data:

$$\mathrm{V_A} = k_V \mathrm{VV} \tag{5.2.3}$$

$$D_M = k_M \mathrm{M:Gas_{rel}V_A} \tag{5.2.4}$$

$$\theta V_C = k_R \mathrm{RBC:Gas_{rel}V_A} \tag{5.2.5}$$

^{129}Xe MRI metrics; VV is the lung ventilated volume from ventilation imaging [232] and $\mathrm{M:Gas_{rel}}$ and $\mathrm{RBC:Gas_{rel}}$ are the dissolved ^{129}Xe signal ratio values divided by a healthy reference value (see *Methods*). Equations 5.2.4 and 5.2.5 are then substituted into Equation 5.2.1:

$$\frac{1}{\mathrm{K}} = \frac{1}{k_M \mathrm{M:Gas_{rel}}} + \frac{1}{k_R \mathrm{RBC:Gas_{rel}}} \tag{5.2.6}$$

Here, the 'CO' has been dropped from $\mathrm{K_{CO}}$ to specify that this is the predicted value from ^{129}Xe MRI metrics. The predicted K and $\mathrm{V_A}$ values are then multiplied to give the predicted transfer factor, TL.

5.3 Methods

5.3.1 Subject Details

Models were trained on data from the Advanced Diagnostic Profiling (ADPro) substudy of the NOVEL observation longiTudinal studY (NOVELTY; ClinicalTrials.gov identifier: NCT02760329) of patients with asthma and/or COPD [233, 234]. Patients were recruited from primary care and had a mix of physician-assigned disease severity (mild, moderate or severe). The study involved two visits, at which participants underwent PFTs and MRI, with the second visit 12 ± 2 months after the first. At visit one, ^{129}Xe ventilation and gas exchange imaging and PFTs were performed post-bronchodilator, whereas at visit two, ^{129}Xe imaging and PFTs were carried out both pre- and post-bronchodilator. Models were trained using the visit one dataset (n = 165; 42 mild, 76 moderate and 47 severe). Participants who had missing PFT or imaging data (due to missed appointments or scanner failure) were excluded and so the final training group consisted of 150 participants (Table 5.3.1). 123/150 participants had pre-bronchodilator and 127/150 participants had post-bronchodilator ^{129}Xe imaging and PFT data at visit two.

Model validation was carried out on a separate cohort of 42 participants, who differed from the training group in diagnosis. 19/42 participants were part of the healthy control group of the Hyperpolarised Xenon Magnetic Resonance PuLmonary Imaging in PAtIeNts with Long-COVID (EXPLAIN) study [235]. The remaining 23 participants had been hospitalised due to COVID-19 (the MURCO: MUlti-nuclear MRI in COVID-19 study [190]) and their PFT and ^{129}Xe MRI data was from 1 - 12 months after hospital admission (median 6 months).

In accordance with Ref. [224], RBC:Gas and M:Gas were normalised by healthy reference values, which were taken from a previous study [35]. The subject demographics for this group are given in Table 5.3.2.

5.3.2 MRI Acquisition and Pulmonary Function Testing

Hyperpolarised ^{129}Xe ventilation and gas exchange imaging were performed as described in Chapter 4. The training data and the healthy volunteer data used for normalisation (Table 5.3.2) were acquired using the original implementation of our spectroscopic dissolved-phase imaging sequence, which used TR = 40 ms and a flip angle of 40° centred on the dissolved ^{129}Xe resonance, as in [35]. The validation data were acquired with the updated version of this sequence, which used TR = 15 ms and dissolved-phase flip angle = 22° [189]. Additional anatomical imaging was performed using a UTE sequence [20]. Measurement of TL_{CO}, K_{CO} and V_A were performed using a Vyaire PFT Pro (Vyaire Medical, Inc., Basingstoke, UK) and in accordance with international guidelines [236].

Table 5.3.1: Subject demographics for the testing and validation groups. Normally distributed variables are given as mean ± standard deviation, whereas non-normally distributed variables are given as median (range).

	All Asthma & COPD	Asthma	COPD	Asthma + COPD	Post-COVID-19 Hospitalisation	Healthy
Number (females)	150 (75)	76 (39)	24 (17)	50 (19)	23 (3)	19 (13)
Age (years)	60.6 (21.4 - 82.2)	54.2 ± 13.8	66.5 ± 8.4	63.6 ± 10.5	62 (42 - 79)	40.9 ± 9.5
Weight (kg)	80.7 ± 17.4	83.7 ± 17.0	68.4 ± 13.5	82.0 ± 17.5	97.0 ± 16.0	76.8 ± 14.5
Height (cm)	168.5 ± 10.4	169.0 ± 9.8	162.0 ± 10.4	170.8 ± 10.3	173.4 ± 9.1	170.5 ± 11.3
FEV_1 (z-score)	-0.65 (-4.49 - 2.66)	-0.10 (-3.84 - 2.66)	-1.89 ± 1.60	-1.17 ± 1.16	-0.70 ± 1.00	0.13 ± 0.75
TL_{CO} (mmol min^{-1} kPa^{-1})	7.56 ± 2.55	8.48 ± 2.18	5.16 ± 2.46	7.01 (3.61 - 13.86)	6.00 ± 1.68	8.84 ± 2.19
TL_{CO} (z-score)	-0.07 (-6.37 - 3.83)	0.29 ± 1.14	-2.01 ± 2.28	-0.65 ± 1.30	-2.04 ± 1.37	0.33 ± 0.84
K_{CO} (mmol min^{-1} kPa^{-1} L^{-1})	1.38 ± 0.32	1.53 ± 0.24	1.10 ± 0.35	1.27 ± 0.28	1.30 ± 0.17	1.52 ± 0.16
K_{CO} (z-score)	-0.22 (-4.83 - 3.00)	0.30 ± 1.07	-1.63 ± 1.86	-0.67 ± 1.29	-0.52 ± 0.78	0.09 ± 0.68
VA (L)	5.45 ± 1.27	5.55 ± 1.16	3.99 (2.76 - 7.81)	5.72 ± 1.22	4.58 ± 0.87	5.82 ± 1.33
VV (L)	4.47 ± 0.86	4.46 (2.08 - 7.28)	4.03 ± 0.74	4.60 (2.96 - 7.18)	3.80 ± 0.43	4.39 ± 0.74
M:Gas	0.0091 (0.0051 - 0.0157)	0.0099 ± 0.0021	0.0076 ± 0.0017	0.0090 ± 0.0024	0.0113 (0.0089 - 0.0125)	0.0087 ± 0.0012
RBC:Gas	0.0028 (0.0012 - 0.0068)	0.0031 (0.0017 - 0.0068)	0.0017 (0.0012 - 0.0034)	0.0025 (0.0013 - 0.0058)	0.0021 (0.0013 - 0.0039)	0.0032 ± 0.0006

Table 5.3.2: Details of the healthy volunteer cohort used to normalise the gas exchange ratios in Model 1.

Number (females)	10 (2)
Age (years)	38 (25 - 65)
M:Gas	0.0075 ± 0.0018
RBC:Gas	0.0036 ± 0.0013

5.3.3 TL_{CO} Prediction Models

Three TL prediction models were evaluated:

1. Physiology-based linear regression: Equations 5.2.3 and 5.2.6 were fitted on the training data, first using the values of k_V, k_M and k_R from [224] (model 1a) and separately using values found from a least squares solver to minimise the mean squared error (MSE) and best fit our training data (model 1b) in Matlab (version R2022a, Mathworks, Natick, MA). As in [224], RBC:Gas and M:Gas were normalised by healthy reference values (Table 5.3.2).

2. Multivariable linear regression: features were chosen by first identifying correlated variables with a Spearman's correlation matrix of possibilities (VV, M:Gas, RBC:Gas, age, sex, height and weight). Strongly correlated variables were removed, to avoid multi-collinearity. Separate prediction equations were then formed for K and V_A by testing the predictive power of linear combinations of the remaining variables. Model fitting was performed with a linear regression solver from the scikit-learn Python toolbox (Python 3.9.12) [237].

3. Random forest regression: this is an ensemble machine learning algorithm which combines predictions from many uncorrelated decision trees to output a prediction or classification [238]. Two regression models were trained using scikit-learn [237], to predict K and V_A separately, using the features identified from linear regression modelling. Tree splitting was based on minimisation of the MSE and model hyperparameters (Table 5.3.3) were tuned using a grid search.

For both models 2 and 3, RBC:Gas and M:Gas were not normalised by the healthy reference values. Two other machine learning algorithms, polynomial regression and support vector machine regression, were also explored in the preliminary stages of this work, but were not pursued because initial results showed worse performance than linear regression.

Table 5.3.3: Random forest regression model parameters.

	Maximum tree depth	Number of estimators
V_A	4	500
K	3	200

5.3.4 Model Training and Validation

The three models were initially trained using five-fold cross-validation: the training data were split randomly into five folds of 30 participants (stratified such that each group contained approximately the same proportion of asthma, COPD and combined asthma and COPD patients), and the models were fitted on four of the folds, with the remaining fold used to test the fit performance. This was repeated five times so that each fold acted as the testing data once. The model with the lowest MSE across the cross-validation folds was chosen as the final model and this was subsequently retrained on the entire training data and evaluated by applying to the validation set and ADPro visit two data.

5.3.5 Regional TL Mapping

To derive regional maps of TL, the final trained K and $\mathrm{V_A}$ models were applied to each voxel of the RBC:Gas map and a 'relative' ventilation map determined from the gas-phase images of the gas exchange acquisition (Figure 5.3.1). The maps were masked by applying a noise threshold to the M signal images. Ventilation distribution maps were found using Equation 5.3.1 [239]:

$$V_r = \mathrm{VV}\frac{I_r N_{voxel}}{I_{total}}, \tag{5.3.1}$$

where V_r is the relative ventilation at position $r = (x, y, z)$, I_r is the gas signal intensity at position r, N_{voxel} is the total number of ventilated voxels (i.e. voxels in the lung mask) and I_{total} is the total gas signal intensity. This intermediate step was required so that the input to the $\mathrm{V_A}$ model (V_r) had the same units and order of magnitude as the VV data that the model was trained on. So that the $\mathrm{V}_{\mathrm{A},r}$ maps had a clearer physiological meaning (ventilated volume per voxel), the initial $\mathrm{V_A}$ model output was scaled by N_{voxel}. The whole-lung predictions can be recovered from the regional maps by summing the fractional TL and $\mathrm{V_A}$ values (TL_r and $\mathrm{V}_{\mathrm{A},r}$) and averaging the rate K (K_r):

$$\mathrm{TL} = \sum_r \mathrm{TL}_r \tag{5.3.2}$$

$$\mathrm{V_A} = \sum_r \mathrm{V}_{\mathrm{A},r} \tag{5.3.3}$$

$$\mathrm{K} = \bar{\mathrm{K}}_r \tag{5.3.4}$$

In order to aid the interpretation of parametric maps, it is often helpful to perform binning based on a healthy distribution so that regions that are elevated or reduced relative to the healthy value can be identified. However, the healthy volunteers included in this work were limited in number and not age-matched to the patients, so it was not appropriate to define a healthy TL_r distribution from their TL maps. PFT results are typically reported as z-scores and/or the percentage of predicted values. Therefore, in a regional analogue to this, the Global Lung Initiative (GLI) predicted TL_{CO} values [227] were used to perform subject-specific binning of the TL maps. To do this, the whole-lung GLI predicted value was divided by N_{voxel} in order to estimate the TL_r value for each voxel: $TL_{pred,r}$ (assuming a homogeneous gas exchange distribution). The bin edges were chosen based on percentiles of $TL_{pred,r}$: [0, 0.25, 0.5, 0.75, 1, 1.25, 1.4, 1.75]*$TL_{pred,r}$ and the binning colours were chosen such that bins within the range of normal TL_{CO} values (75% - 140% predicted [109]) were shown in different shades of green. Binning was carried out using the Python `np.digitize` function.

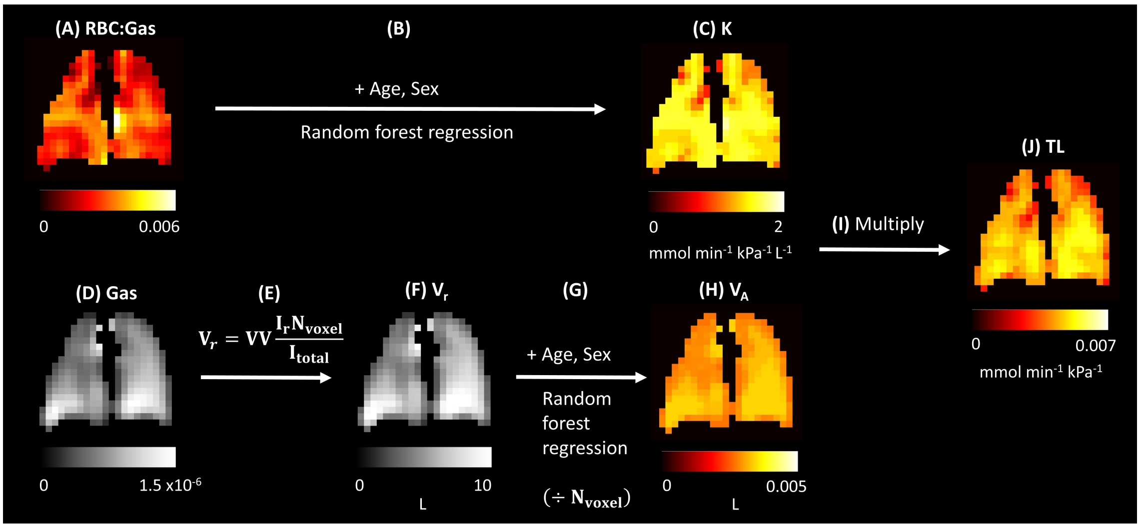

Figure 5.3.1: Regional random forest model: information on regional RBC uptake and gas signal distribution from dissolved ^{129}Xe imaging was utilised to produce regional maps of TL. The K prediction pipeline was applied to every voxel of the RBC:Gas map (A), along with patient age and sex (B), to output a map of predicted K (C). For the prediction of V_A, an extra step (Equation 5.3.1) was required in order to convert the gas signal map (D) into a map of ventilation distribution (V_r), which had the required units of litres and order of magnitude. (E) This involved finding the signal intensity of each pixel (I_r), dividing this by the mean signal intensity (I_{total}/N_{voxel}) and multiplying this fraction by VV from ^{129}Xe ventilation imaging to produce the V_r map (F). (G) The V_A random forest prediction was then applied to each voxel of this map, along with patient age and sex. The resulting map was renormalised by N_{voxel} so that it represented the ventilation per voxel and summed to give predicted V_A (H). (I) This was then multiplied with the K map to obtain a map of TL (J).

5.3.6 Statistical Analysis

The predictive power of the models was assessed by calculating the MSE and R^2 of the linear fit of predicted and measured TL_{CO}, the mean absolute error (MAE) between the predicted and measured values, and their Bland-Altman bias. Normality of variables was tested for using Shapiro-Wilks tests. Correlations were assessed using the Pearson/Spearman correlation coefficient for normally/non-normally distributed variables. To compare the correlations between different pairs of variables, Hittner's z [240] and Zou's confidence interval [241] for dependent, overlapping correlations were calculated using the `cocor` R package [242]. All statistical analysis was carried out with RStudio (R version 4.3.0).

5.4 Results

5.4.1 Validation of Physiological Model

5.4.1.1 Model 1a

The model using linear regression coefficients from Wang et al., given in Table 5.4.1, did not fit our training data well (Table 5.4.2; MAE = 2.66 mmol min^{-1} kPa^{-1}).

5.4.1.2 Model 1b

The prediction accuracy was improved by refitting the values of k_V, k_M and k_R on our dataset, which resulted in a MAE value of 1.24 ± 0.15 mmol min^{-1} kPa^{-1} across the five cross-validation folds. The refitted values are shown in Table 5.4.1. Linear regression and Bland-Altman analysis of the measured and predicted TL_{CO} values for each model are shown in Figure 5.4.1.

Table 5.4.1: Linear regression coefficients and confidence intervals/standard deviation from the physiological model.

	Wang et al.	This work
k_V	1.47 [1.42, 1.52]	1.21 ± 0.01
k_M (mmol min^{-1} kPa^{-1}L^{-1})	3.55 [2.88, 4.55]	4.51 ± 0.26
k_R (mmol min^{-1} kPa^{-1}L^{-1})	4.55 [3.82, 5.59]	2.97 ± 0.13

Table 5.4.2: Evaluation of the four transfer factor prediction models on the training data. Model 1a = physiological model with coefficients from Wang et al. [224], 1b = physiological model with refitted coefficients, 2 = multivariable linear regression, 3 = random forest regression. For models 1b, 2 and 3, the mean and standard deviation across the five cross-validation folds is given. LOA = limits of agreement.

Model	MSE (mmol min^{-2} kPa^{-2})	R^2	MAE (mmol min^{-1} kPa^{-1})	Bias [LOA] (mmol min^{-1} kPa^{-1})
1a	3.49	0.470	2.66	2.41 [-1.56, 6.39]
1b	1.63 ± 0.31	0.604 ± 0.028	1.24 ± 0.15	-0.14 [-3.25, 2.97]
2	1.16 ± 0.23	0.739 ± 0.069	1.01 ± 0.06	-0.02 [-2.55, 2.52]
3	1.13 ± 0.24	0.744 ± 0.063	0.995 ± 0.129	0.02 [-2.50, 2.54]

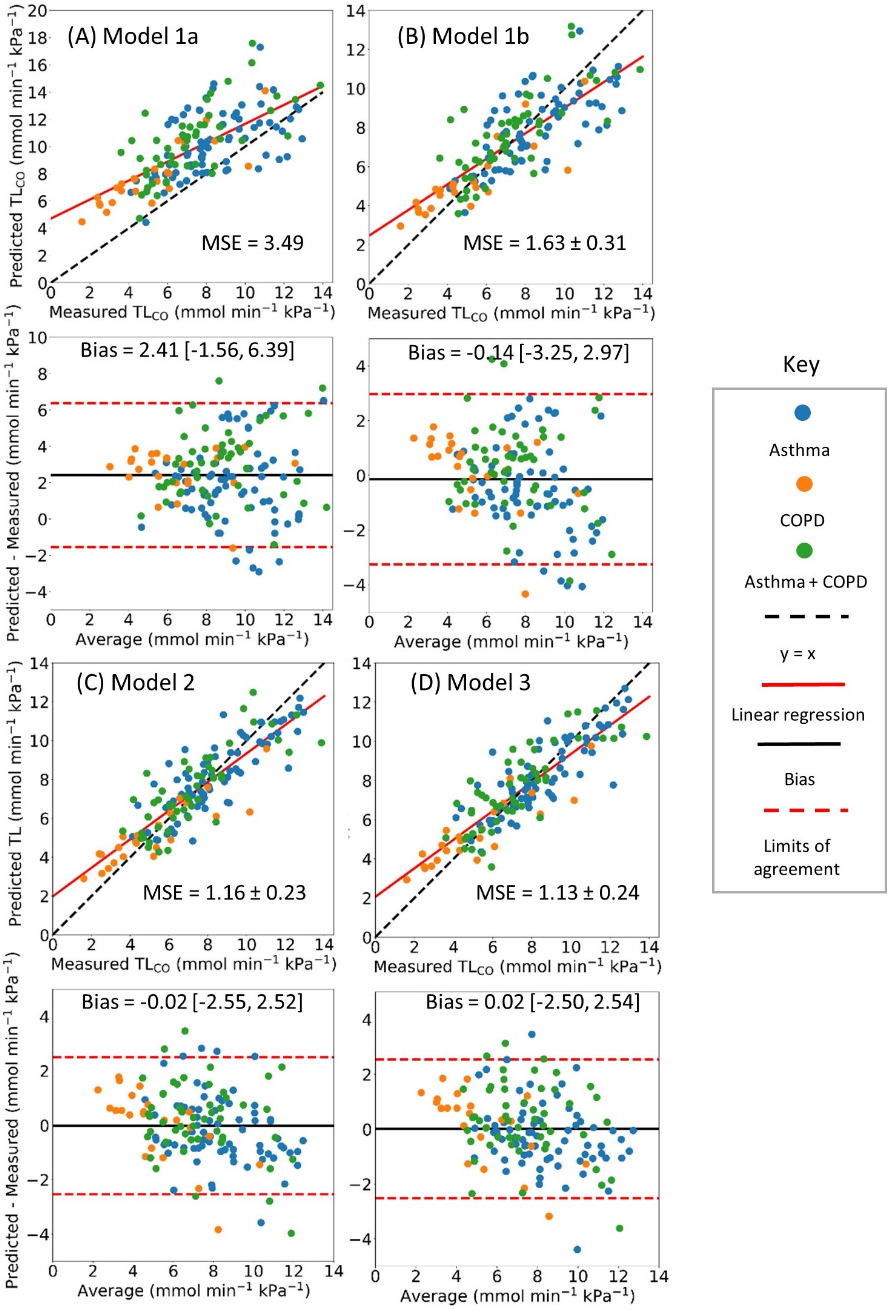

Figure 5.4.1: Evaluation of the three prediction models on the training data: (A) physiology-based model with coefficients from Wang et al. [224] and (B) coefficients refitted on our training data, (C) multivariate linear regression and (D) random forest regression, via linear regression and Bland-Altman plots of the measured and predicted TL$_{CO}$ values from the respective models.

5.4.2 Model Development and Validation

5.4.2.1 Model 2

The correlation matrix showed strong correlations between RBC:Gas and M:Gas (Spearman's ρ = 0.57), age and height (ρ = -0.75) and VV and height (ρ = 0.62). Therefore, M:Gas and height were not included as features in the multivariable linear regression model because of their dependence on other variables. The final predictive equations were:

$$\mathrm{K} = a_1 \frac{1}{\mathrm{RBC : Gas}} + a_2\mathrm{Age} + a_3\mathrm{Sex} + a_4 \tag{5.4.1}$$

and

$$\mathrm{V_A} = b_1\mathrm{VV} + b_2\mathrm{Age} + b_3\mathrm{Sex} + b_4, \tag{5.4.2}$$

where Sex = 0 for males and 1 for females. The fitted values of $a_{1,2,3,4}$ and $b_{1,2,3,4}$ can be found in Table 5.4.3. Including participant age and sex as predictors improved the prediction of $\mathrm{TL_{CO}}$ compared to model 1b (MAE = 1.01 ± 0.06 mmol min^{-1}kPa^{-1}).

Table 5.4.3: The coefficients of the multivariate linear regression model (Equations 5.4.1 and 5.4.2), given as the mean ± standard deviation across the five cross-validation training folds.

	K	$\mathrm{V_A}$
VV	N/A	b_1 = 0.6378 ± 0.0281
1/RBC	a_1 = -0.0015 ± 0.0001	N/A
Age	a_2 = -0.0018 ± 0.0004	b_2 = -0.0245 ± 0.0038
$\mathrm{Sex_F}$	a_3 = 0.0740 ± 0.0203	b_3 = -1.3752 ± 0.0650
constant term	a_4 = 2.0500 ± 0.0371	b_4 = 4.75 ± 0.35

5.4.2.2 Model 3

The random forest regression model performed slightly better than linear regression (MAE = 0.995 ± 0.129 mmol min^{-1} kPa^{-1}) so was chosen as the final model. The most significant predictor of K was RBC:Gas, whereas sex was the highest ranking predictor of $\mathrm{V_A}$ (Figure 5.4.2). This might be because sex is strongly correlated with height, which in turn is a significant predictor of lung volume [37]. The model also performed well when applied to the validation data (Figure 5.4.3: MAE = 0.840 mmol min^{-1} kPa^{-1}, MSE = 0.647 mmol min^{-2} kPa^{-2}, R^2 = 0.828). For completeness, models 1a, 1b and 2 were also evaluated on the validation data and the performance of model 2 was found to be slightly better than that of model 3: MAE = 0.744 mmol min^{-1} kPa^{-1}, MSE = 0.452 mmol min^{-2} kPa^{-2}, R^2=0.877.

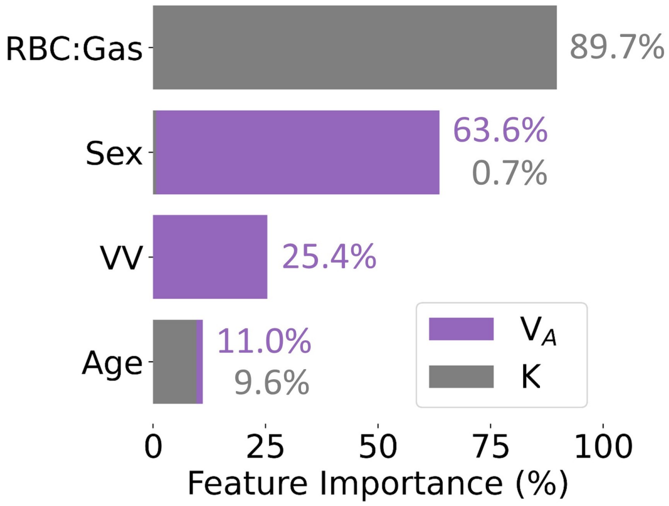

Figure 5.4.2: Ranked importance of the prediction variables for the random forest regression model. The purple bars represent the contribution of the variables to the V_A prediction and the grey bars represent the contribution to the K prediction.

5.4.2.3 Correlations

Both RBC:Gas and RBC:M were significantly correlated with measured TL_{CO}, with Spearman's ρ values of 0.71 and 0.64, respectively. To determine whether there was a stronger relationship between measured TL_{CO} and modelled TL than there was between measured TL_{CO} and RBC:Gas/RBC:M, the correlation coefficients were compared using statistical tests. The results are summarised in Table 5.4.4. The correlation between measured TL_{CO} and predicted TL was stronger than the correlation between measured TL_{CO} and RBC:M for models 1b, 2 and 3. For models 2 and 3, the predicted TL correlation was also stronger than the RBC:Gas correlation, whereas for model 1b, there was not a significant difference between the predicted TL and RBC:Gas correlation coefficients.

Table 5.4.4: Comparison of the correlation between measured TL_{CO} and RBC:Gas/RBC:M and the correlation between measured TL_{CO} and the predicted TL from each model. A significant *z* value indicates a significant difference between the correlation coefficients, as does a Zou confidence interval (CI) which does not include 0 in its range.

		RBC:Gas		RBC:M	
Model	Pearson's *r*	Hittner's *z*	Zhou's CI	Hittner's *z*	Zhou's CI
1b	0.78	-1.18, ($p = 0.24$)	[-0.13, 0.03]	-2.92 ($p = 0.004$)	[-0.21, -0.04]
2	0.86	-4.40 ($p < 1 \times 10^{-4}$)	[-0.20, -0.07]	-5.73 ($p < 1 \times 10^{-4}$)	[-0.29, -0.13]
3	0.86	-5.64 ($p < 1 \times 10^{-4}$)	[-0.29, -0.12]	-4.23 ($p < 1 \times 10^{-4}$)	[-0.20, -0.07]

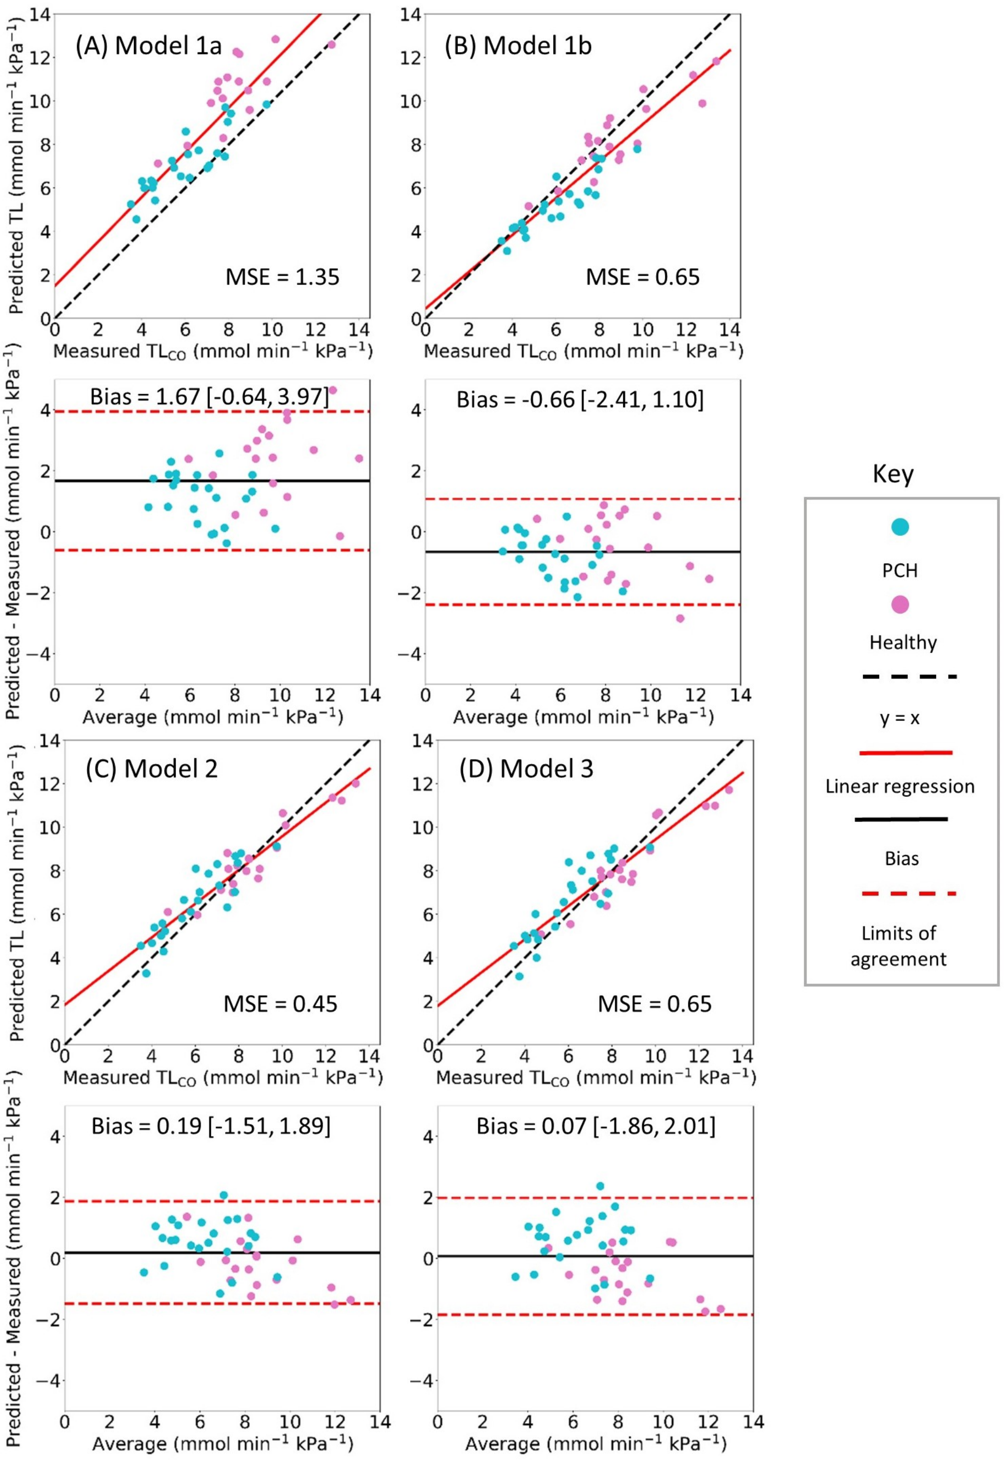

Figure 5.4.3: Evaluation of the three prediction models on the validation data: (A) model 1a: physiology-based model with coefficients from Wang et al. [224] and (B) model 1b: coefficients refitted on our training data, (C) model 2: multivariate linear regression and (D) model 3: random forest regression, via linear regression and Bland-Altman plots of the measured and predicted TL_{CO} values. PCH = post-COVID-19 hospitalisation.

5.4.3 Regional TL Mapping

Example predicted TL maps demonstrating differences between four participants with different pathology, but similar TL_{CO} z-scores, and one healthy volunteer are shown in Figure 5.4.4. With the addition of regional gas exchange information, the performance of the random forest model was not improved when compared to the whole-lung model, but was still reasonably good: MAE = 1.18 mmol min^{-1} kPa^{-1}, MSE = 1.26 mmol min^{-2} kPa^{-2}, R^2 = 0.736. It also performed well when applied to visit two data from the same study, for both pre-bronchodilator (MAE = 1.30 mmol min^{-1} kPa^{-1}, MSE = 1.17 mmol min^{-2} kPa^{-2}, R^2 = 0.747) and post-bronchodilator (MAE = 1.14 mmol min^{-1} kPa^{-1}, MSE = 1.13 mmol min^{-2} kPa^{-2}, R^2 = 0.756) data. Linear regression and Bland-Altman plots comparing the measured, whole-lung random forest predicted and regional random forest predicted TL values for both visits are shown in Figure 5.4.5. Example TL-constituent K_{CO} and V_A maps for one participant with both asthma and COPD are shown in Figure 5.4.6. This person had normal measured TL_{CO} (z-score = 0.48), but their maps reveal regional heterogeneity in both ventilation and gas transfer.

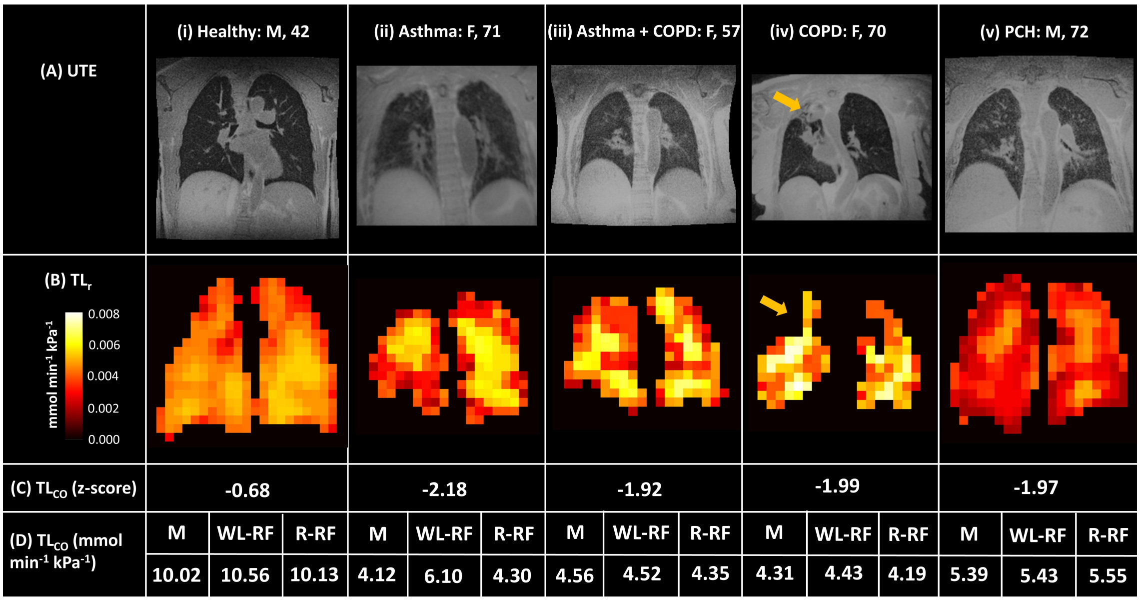

	(i) Healthy: M, 42			(ii) Asthma: F, 71			(iii) Asthma + COPD: F, 57			(iv) COPD: F, 70			(v) PCH: M, 72		
(C) TL_{CO} (z-score)	-0.68			-2.18			-1.92			-1.99			-1.97		
(D) TL_{CO} (mmol min^{-1} kPa^{-1})	M	WL-RF	R-RF	M	WL-RF	R-RF	M	WL-RF	R-RF	M	WL-RF	R-RF	M	WL-RF	R-RF
	10.02	10.56	10.13	4.12	6.10	4.30	4.56	4.52	4.35	4.31	4.43	4.19	5.39	5.43	5.55

Figure 5.4.4: (A) UTE lung structure images and (B) random forest regression-predicted TL maps for five participants and their diagnosis, sex, age and (C) TL_{CO} z-score at visit one. The patient shown in (A iv) had a lack of ^{129}Xe signal in the upper right lung due to underlying structural changes (yellow arrows). (D) The measured and estimated TL values for each participant are given, where 'WL-RF' refers to the value from the whole-lung random forest model and 'R-RF' refers to the value from the sum of TL_r from the regional random forest model over all voxels. PCH: post COVID-19 hospitalisation.

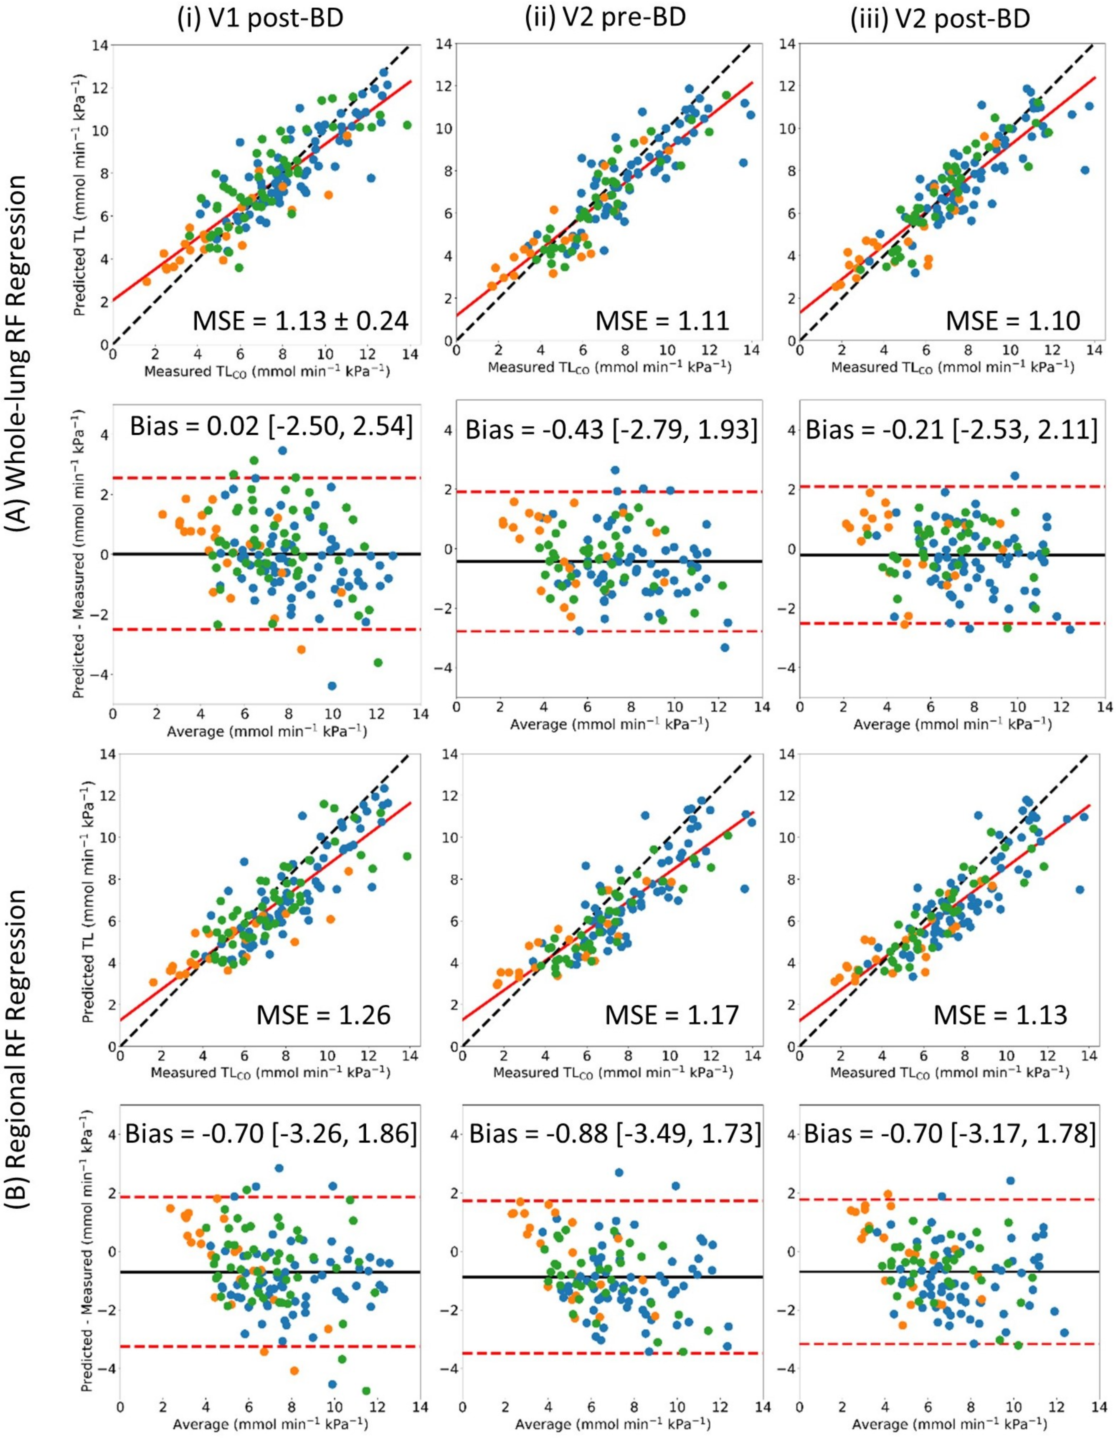

Figure 5.4.5: Evaluation of the (A) whole-lung random forest regression model and (B) regional random forest regression model on data from the ADPro asthma and COPD study for (i) visit 1 (training data), (ii) visit 2, pre-bronchodilator and (iii) visit 2, post-bronchodilator. The first and third rows show linear regression and the second and fourth rows show Bland-Altman plots of the measured and predicted TL$_{CO}$ values.

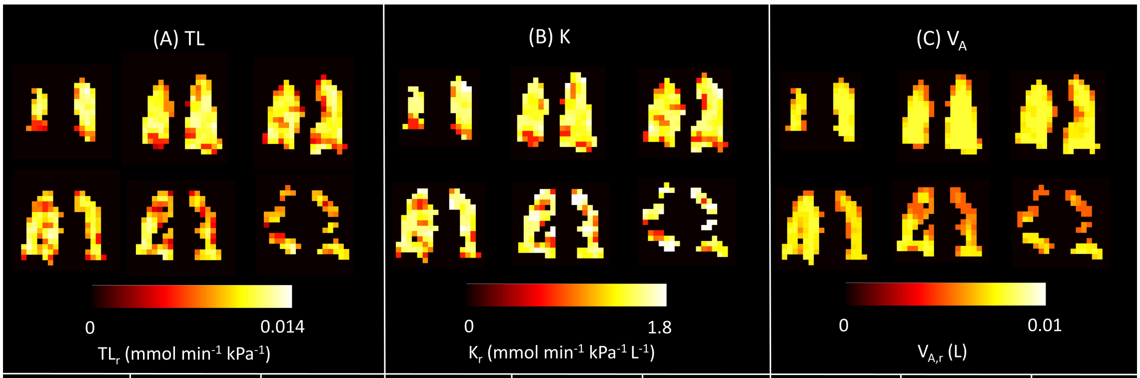

Measured	WL-RF	R-RF	Measured	WL-RF	R-RF	Measured	WL-RF	R-RF
7.01	7.08	6.02	1.73	1.52	1.42	4.05	4.65	4.29

Figure 5.4.6: Random forest regression modelled (A) TL, (B) K and (c) V_A maps for a 57-year-old female patient with both asthma and COPD for six lung slices, plus the modelled and measured whole-lung values at visit one. Although this patient had a normal TL_{CO} (as measured by PFT), their TL map indicates reduced gas exchange. The K map shows a heterogeneous gas transfer rate which does not match the ventilation distribution. (The TL map is the product of the K_r and $V_{A,r}$ maps.)

Figure 5.4.7 shows the predicted TL, K and V_A maps for a single slice for a participant with both asthma and COPD before and after bronchodilator administration. The maps show an increase in the lung mask area post-bronchodilator, but the gas signal intensity per voxel decreases because the same amount of gas is distributed over more voxels due to the opening of previously unventilated airways. Changes in the regional distribution of TL, K and V_A following bronchodilator can be assessed by considering the predicted values for each lung slice (Figure 5.4.7 bottom panel).

The binned TL maps for five central lung slices along with the corresponding maps of RBC:M and TL are shown for a patient with both asthma and COPD in Figure 5.4.8 and a patient who was hospitalised post-COVID-19 infection in Figure 5.4.9. Both patients had a % predicted TL_{CO} of ∼70%, but there are differences in the distribution of TL values, which are more stark in the binned TL maps. For comparison, the binned TL maps for a healthy volunteer (% predicted TL_{CO} = 90%) are shown in Figure 5.4.10.

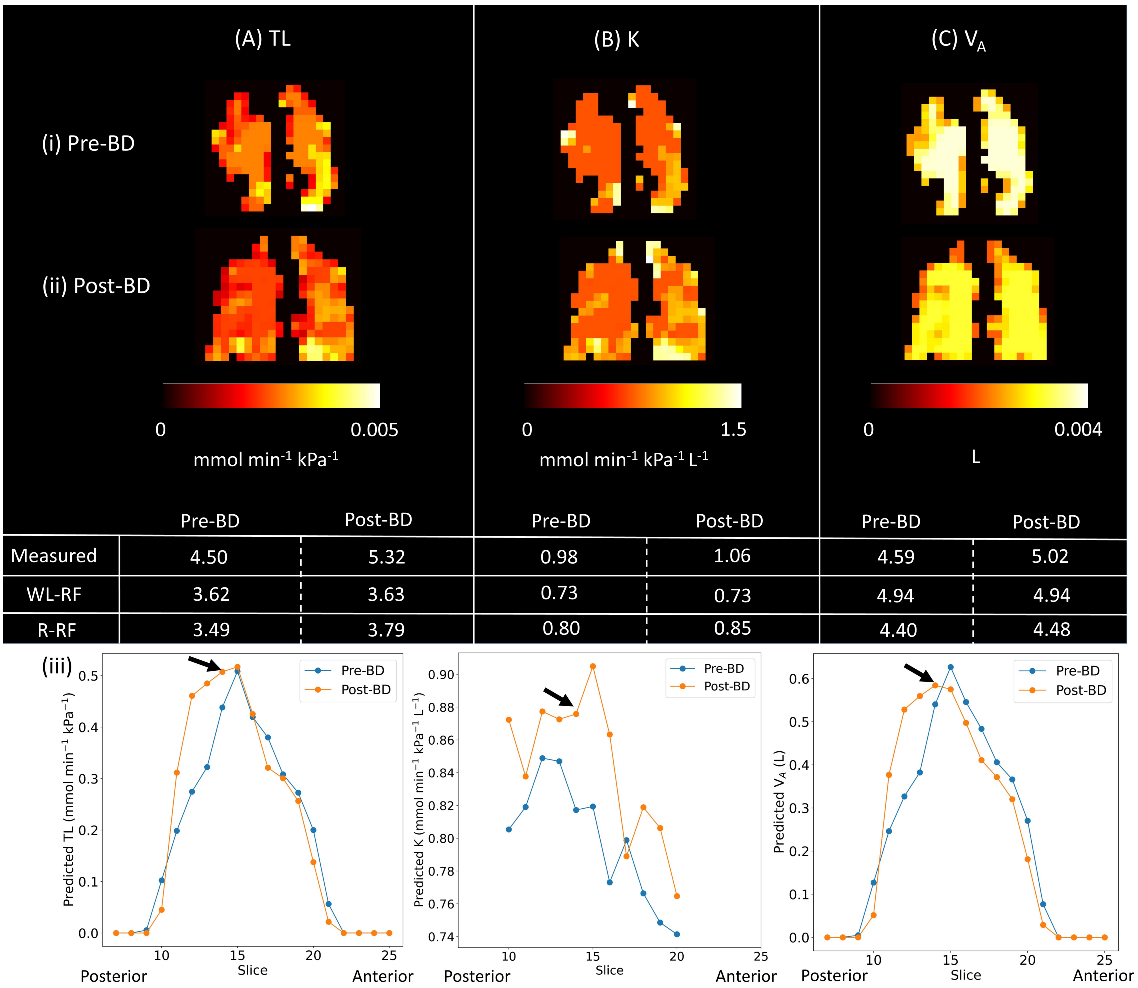

	TL Pre-BD	TL Post-BD	K Pre-BD	K Post-BD	V_A Pre-BD	V_A Post-BD
Measured	4.50	5.32	0.98	1.06	4.59	5.02
WL-RF	3.62	3.63	0.73	0.73	4.94	4.94
R-RF	3.49	3.79	0.80	0.85	4.40	4.48

Figure 5.4.7: Modelled parametric maps of (A) TL, (B) K and (C) V_A for a single lung slice for a 61-year-old female with both asthma and COPD, (i) pre- and (ii) post-bronchodilator. In (iii), the sum (mean for K_r) for each slice (posterior to anterior) is plotted, showing the change in distribution following bronchodilator administration. The lung slice shown in (i) and (ii) is indicated with an arrow. These plots show that $V_{A,r}$ and TL_r increase in the posterior lung following bronchodilation, whilst slightly decreasing in the central and anterior slices.

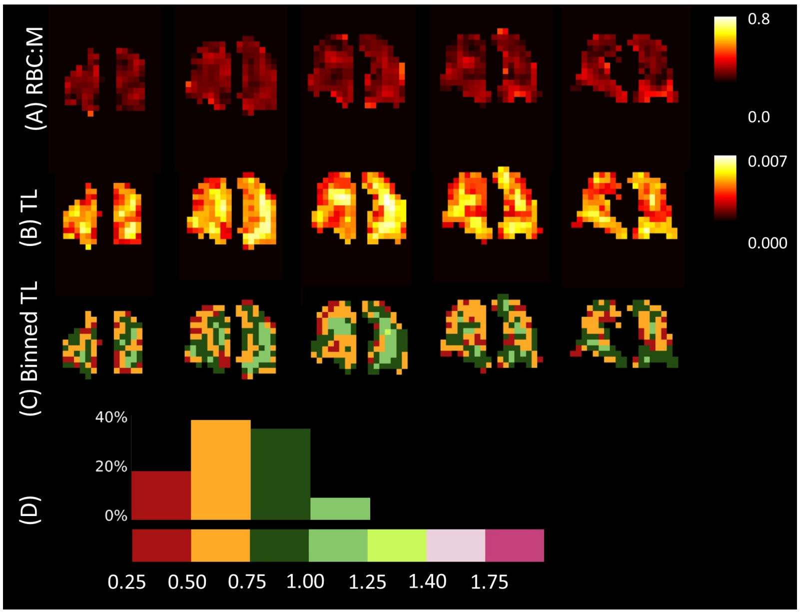

Figure 5.4.8: TL binning for a 57 year old female patient with both asthma and COPD and a TL_{CO} z-score of -1.91. For five central lung slices, maps of (A) RBC:M, (B) TL and (C) TL binned according to percentage of predicted TL per voxel, are shown. In (D), the binning edges and colours are shown along with a histogram showing the proportion of voxels in each bin for this subject.

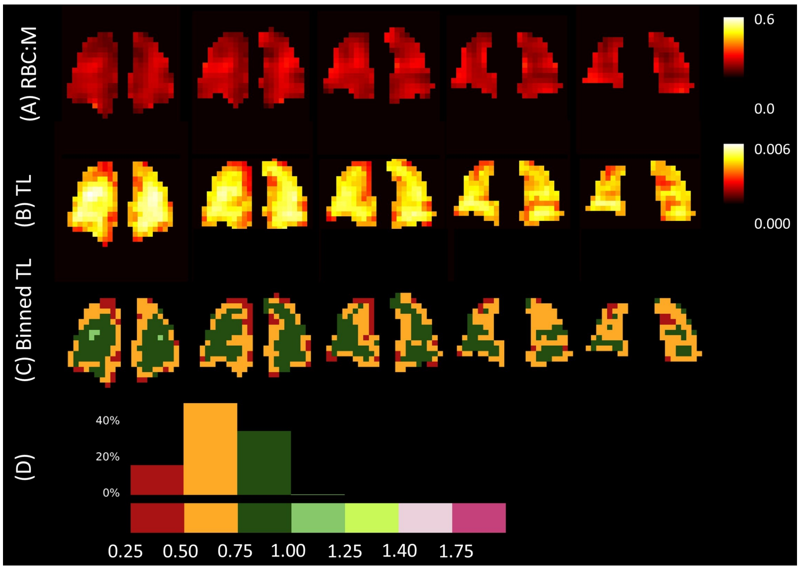

Figure 5.4.9: TL binning for a 59 year old male patient who was hospitalised following COVID-19 infection and had a TL_{CO} z-score of = -2.23: (A) RBC:M, (B) TL, (C) TL binned according to percentage of predicted TL per voxel, (D) histogram showing the proportion of voxels in each bin for this subject.

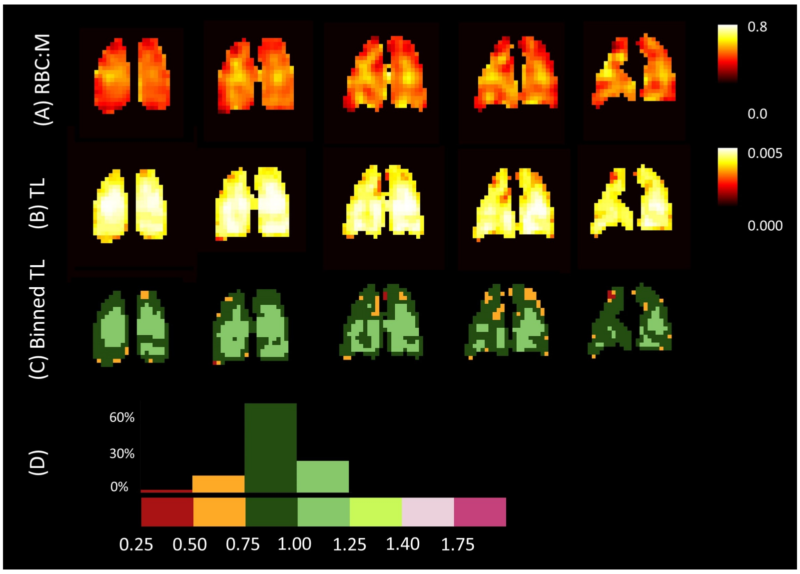

Figure 5.4.10: TL binning for a 42 year old healthy man who had a TL_{CO} z-score of = -0.68: (A) RBC:M, (B) TL, (C) TL binned according to percentage of predicted TL per voxel, (D) histogram showing the proportion of voxels in each bin for this subject.

5.5 Discussion

In this work, we have demonstrated that random forest regression modelling improves prediction of TL from ^{129}Xe MRI metrics when compared to a physiology-based model. We then go on to apply this model voxelwise to create maps of TL, K and V_A. This approach provides a valuable means to visualise these clinical lung physiology metrics at a regional level, allowing the contributions from the ventilation distribution and gas uptake rate to be examined, and links measures from ^{129}Xe gas exchange imaging to well-established physiological quantities.

5.5.1 Model Development and Performance

Three models for predicting TL from ^{129}Xe MRI data were evaluated. The coefficients from the previously proposed physiology-based model [224] did not fit our data well, likely due to differences in MRI acquisition strategies/parameters and participant disease aetiologies. The previous model was trained on data acquired with a 1-point Dixon method dissolved-phase imaging sequence at both 1.5 T and 3 T from healthy participants and those with obstructive, restrictive, and pulmonary vascular lung disease. It is also possible that, by not using separate testing and training groups or cross-validation, the coefficients in [224] were reached by finding a sub-optimal local minimum of MSE. In contrast to [224], we found $k_M > k_R$, which suggests a greater contribution to the total impedance from the membrane than the capillary blood, or ventilation-perfusion mismatch. For a healthy individual, RBC:Gas$_{rel}$ $\sim$ M:Gas$_{rel}$ $\sim$ 1, thus, D_M = k_M V_A and θV_C = k_R V_A. For the coefficients found in this work, this gives a ratio of $\frac{1}{D_M}$: $\frac{1}{\theta V_C}$ of 2:3. The membrane conductance contributes $\sim$ 40% to the total impedance and the capillary blood volume contributes $\sim$ 60%. Although different to Ref. [224], which found a $\sim$ 44 % contribution from the capillary blood volume, our result of 60% is more similar to estimates from reference equations, which are around 75% [228].

In models 2 and 3, machine learning was used to try and improve prediction accuracy. The model features were a combination of ^{129}Xe MRI and demographic variables, chosen to minimise multi-collinearity and maximise predictive power. Where two variables were dependent on one another, a choice had to be made over which variable to include in the model. RBC:Gas was chosen to be included over M:Gas because this measurement reflects the ^{129}Xe pulmonary gas transfer from alveoli to bloodstream, i.e. RBC:Gas measures ^{129}Xe transfer across the membrane barrier, but M:Gas only measures increased membrane signal. As highlighted in Figure 5.1.1, RBC:Gas is analogous to K_{CO}. Sex was chosen over height for the following reasons:

1. To follow previous work on the age and sex dependence of ^{129}Xe gas exchange ratios [189, 230].
2. The ^{129}Xe dosing was chosen based on height, so some adjustment for height is already made in the ^{129}Xe MRI protocol.
3. There are other physiological differences between males and females which affect gas exchange, such as haemoglobin levels and effects of the menstrual cycle [243].
4. Height was strongly correlated with both sex and VV.

The features for the random forest model were chosen to be the same as the features in the linear regression model. However, with random forest there are no multi-collinearity concerns so model performance with the addition of height as a feature was also tested. Including height was found to have no improvement on the model performance (MSE = 1.13 ± 0.24 mmol min^{-2} kPa^{-2}, MAE = 0.995 ± 0.129 mmol min^{-1} kPa^{-1} without height, MSE = 1.13 ± 0.54 mmol min^{-2} kPa^{-2}, MAE = 1.00 ± 0.16 mmol min^{-1} kPa^{-1} with height). The standard deviation between folds was higher when height was included, suggesting that this feature is unnecessary and is leading to overfitting. Ultimately, there is no definitive answer about which variables to include in machine learning prediction models such as ours. Factors that have a significant impact on lung function, such as smoking history, were not included in our models because they were not available, but would very likely improve prediction accuracy. With more variables in the prediction models, we would expect to see improved fitting of the training data, however with these additional degrees of freedom there is also higher risk of overfitting and the model not performing well on validation data.

The addition of participant age and sex in the multivariable linear and random forest regression models further improved prediction accuracy when compared to the physiology-based model. Age appeared in both the K and V_A prediction, suggesting TL_{CO} has a second order age dependence. An age^2 dependence represents an accelerated loss of lung function with age and is also found in Munkholm's TL_{CO} prediction equations [228].

The correlations between measured TL_{CO} and TL from models 1b, 2 and 3 were significantly stronger than the correlation between TL_{CO} and RBC:Gas. This demonstrates the relative insensitivity of ^{129}Xe gas exchange MRI to obstructed regions of the lung and the advantage of combining RBC:Gas with information from ventilation imaging in our models. All models had small standard deviation across the five cross-validation folds, demonstrating good accuracy and minimal over-fitting, but the best performing model for the training data was random forest regression. This may be because the model is able to account for non-linear relationships [238]; whereas Equations 5.2.3 and 5.4.2 assume a linear relationship between VV (acquired at FRC + bag volume) and V_A (acquired at TLC), the lung volume dependent pathophysiology in obstructive lung disease means that the relationship is likely to be more complicated.

5.5.2 Regional TL Mapping

Beyond improved modelling, a further novel aspect of this work is the application of the random forest prediction model to produce spatially localised TL, K and V_A parametric maps, allowing for regional visualisation of otherwise global metrics. Combining the information from dissolved-phase and ventilation ^{129}Xe MRI, along with patient age and sex, into a parametric map with defined units may provide a way for respiratory physicians to more easily interpret ^{129}Xe gas exchange MRI. Binning the TL maps according to the patient-specific predicted TL_{CO} per voxel from the GLI reference values [227] further increases their comprehensibility, by highlighting areas that are outside of the normal range of values. This is a regional analogue to presenting the whole-lung TL_{CO} values as a percentage of the predicted value.

TL mapping could assist in the phenotyping of patients and in assessing longitudinal changes and treatment response, especially for patients with both gas exchange limitation and ventilation heterogeneity. One of the key assumptions in the TL_{CO} PFT is that the inspired test gas is homogeneously distributed throughout the lungs. However, in obstructive lung disease, ventilation is heterogeneous and so this assumption may not be appropriate and can lead to an underestimation of TL_{CO} [244]. Contrarily, our models use regional gas exchange measurements from dissolved ^{129}Xe MRI to estimate TL, which are intrinsically sensitive to the distribution of ^{129}Xe gas in the alveolar spaces within the lungs. This may explain why the models consistently overestimated TL for COPD patients with low measured TL_{CO} values and why there is an increase in the percentage difference between measured and modelled TL_{CO} as the FEV_1 z-score decreases in Figure 5.5.1A. It has been suggested that in obstructive lung disease patients with a heterogeneous ventilation distribution, it is more appropriate to use the TLC from body plethysmography instead of V_A in the calculation of TL_{CO}. As shown in Figure 5.5.1B, there was a positive correlation between V_A and TLC for the patients in the training data cohort, although TLC exceeded V_A for some COPD and combined asthma and COPD patients. Replacing the measured TL_{CO} value with the product of K_{CO} and TLC for patients where the difference between V_A and TLC was more than 15% lead to less of an overestimation by the random forest model at low TL_{CO} values, although overall model performance was not improved (Figure 5.5.1C, MAE = 1.25 ± 0.28 mmol min^{-1} kPa^{-1}).

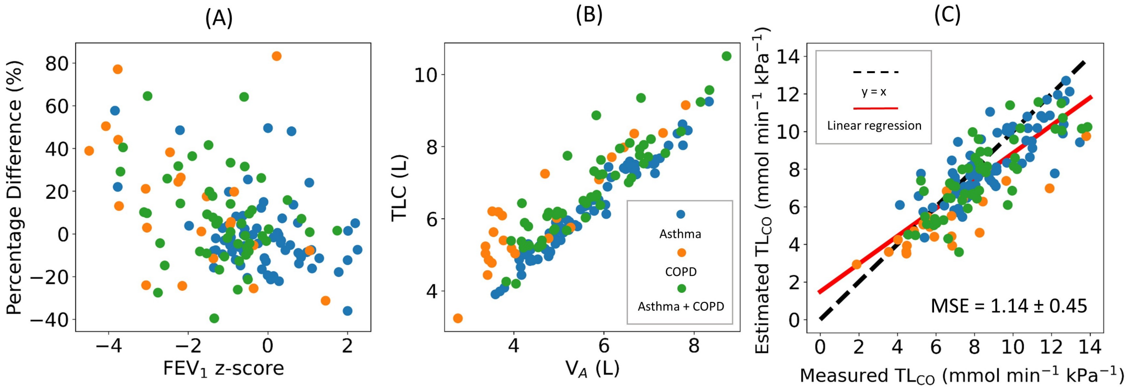

Figure 5.5.1: (A) As the FEV$_1$ z-score decreases and therefore the amount of obstruction increases, the random forest model overestimates TL_{CO} to a greater degree, possibly because measured TL_{CO} is underestimated due to the inhomogeneous ventilation distribution. (B) The relationship between TLC and V_A. (C) Linear regression of the measured and random forest predicted TL_{CO} values, where measured TL_{CO} was calculated from the product of K_{CO} and TLC.

TL was often underestimated for participants with high measured TL_{CO}. This may be due to the larger error on predicted TL (δTL) for higher TL_{CO} values. TL was found from the product of the predicted V_A and K values/maps and so the error on these components was propagated through to TL according to the error propagation formula:

$$\delta \mathrm{TL} = \mathrm{TL}\sqrt{\left(\frac{\delta \mathrm{V_A}}{\mathrm{V_A}}\right)^2 + \left(\frac{\delta \mathrm{K}}{\mathrm{K}}\right)^2} \tag{5.5.1}$$

5.5.3 Model Generalisability and Limitations

The models in this work were trained solely on patients with obstructive lung disease. Although we validated the models using data from healthy people and post-COVID-19 patients, their application to patients with restrictive lung disease, PVD and patients from a different site has not been explored. There are still considerable differences between the ^{129}Xe MRI sequences and acquisition parameters used between sites, which may limit the application of our models to external data, however efforts towards harmonisation are being made through the ^{129}Xe MRI Clinical Trials Consortium (^{129}Xe-CTC) [164]. Hence, it is possible that some retraining could be required to tune the model parameters for external data, but we anticipate that the models themselves and the principle of using regression modelling to map TL should be generalisable. It would be interesting to test the performance of our models on Wang's dataset, to test their generalisability, however this is not possible without access to their data. Since their model did not fit our data well, we could speculate that the same might be true of our models on their data, because of the differences in dissolved ^{129}Xe imaging sequence (MESI vs 1-point Dixon) and patient cohorts used. However, we made efforts to reduce overfitting and increase the generalisability of our models by using a rigorous five-fold cross-validation training strategy with training data from a heterogeneous population with a large range of TL_{CO} values (1.60 - 13.86 mmol

min^{-1} kPa^{-1}). Therefore, we might expect that our models are more generalisable than Wang's model, which did not separate the patient data into testing, training and validation groups.

Although at present we do not have access to a complete dataset from a different site with which to test our models (we hope that this may be possible in the future through our involvement in the ^{129}Xe-CTC), the models performed well on the validation data, suggesting generalisability to other lung diseases. These data were from two different clinical studies, with different patient groups and acquisition parameters to the training data. The differences in TR and flip angle used mean that there is likely to be some bias between the datasets, because the combination of these two parameters determines where on the dissolved ^{129}Xe uptake curve the measurement samples [223]. When the dissolved ^{129}Xe imaging sequence was updated, the impact of the new sequence parameters was evaluated in six healthy volunteers. The average difference in mean RBC:Gas between the two sequences was $\sim$15%. This variability is similar to the repeatability of RBC:Gas and does not seem to have limited the applicability of our models to the validation data.

Inherent differences between ^{129}Xe MRI and TL_{CO} measurement, such as the diffusivity and solubility of the gases used, lung volumes and body positions, may limit the predictive power of the models. Both the TL_{CO} PFT and ^{129}Xe MRI have their own limitations, such as the dependence of TL_{CO} on several factors such as patient effort, testing conditions, anaemia and body weight [113]. Gas exchange ^{129}Xe MRI ratio values have quite poor reproducibility and are dependent on factors such as lung inflation [245] and haematocrit [246].

Although the random forest and linear regression models improved the prediction of TL, with these models the ability to estimate the membrane and capillary blood conductances is lost. Another limitation of our work is the requirement for both ventilation and gas exchange ^{129}Xe imaging, which necessitates two separate acquisitions with up to 1.5 L of xenon. At present it is not practical to calculate VV from the gas exchange imaging gas-phase image, because the image resolution (1.25 cm^3 voxel size) is too low for accurate co-registration with the anatomical ^{1}H image. However, image acceleration techniques such as compressed sensing (see Chapter 6) and rapid spiral k-space encoding gradients may make a combined ventilation and gas exchange imaging sequence feasible [174]. As the mainstay of ^{129}Xe MRI, ventilation imaging is part of the vast majority of our ^{129}Xe protocols so this data is usually available alongside gas exchange imaging. The ability to combine information from two different ^{129}Xe scans into one easily interpreted TL map is a strength of our model.

5.6 Conclusions

TL_{CO} can be modelled from ^{129}Xe MRI metrics for patients with obstructive lung disease. Prediction accuracy was improved compared to previous work by using which used patient age, sex and ^{129}Xe MRI-derived VV and RBC:Gas as prediction variables. For the training data, the best performing model was a random forest regression model, although linear regression performed better for the validation data consisting of healthy volunteers and post-COVID-19 hospitalisation patients. The ability of the random forest model to generate TL maps presents a useful tool for visualisation and interpretation of regional TL limitation and should help facilitate the clinical translation of ^{129}Xe gas exchange MRI.

Chapter 6

Compressed Sensing Reconstruction for High-SNR, Rapid Dissolved ^{129}Xe Gas Exchange MRI

6.1 Introduction

Dissolved ^{129}Xe MRI provides sensitivity to gas transfer limitation in lung diseases such as ILD, asthma, and COPD [32, 33, 34, 35, 36, 37, 38, 39]. However, the technique is limited by low signal ($\sim$2% that of the gaseous-phase ^{129}Xe signal in the lungs), resulting from the low solubility of xenon in the alveolar membrane and the small volume of the dissolved-phase compartment when compared to the alveolar airspace [63]. This is in part overcome by using large dissolved-phase flip angles, which is possible because of the rapid exchange of the dissolved-phase magnetisation with the gaseous ^{129}Xe reservoir. Despite the above, achieving sufficient dissolved-phase ^{129}Xe SNR can still be challenging, especially in lung disease patients who have inherently lower RBC signal due to decreased gas transfer and/or reduced capillary perfusion, and those who may also struggle to complete the $\sim$15 s duration breath-hold [174].

Compressed sensing (CS) was described in Section 2.2.8 and involves using the transform sparsity of images and non-linear optimisation to reconstruct randomly undersampled k-space data [60]. The denoising properties of CS can also help improve apparent image SNR and enable reduced scan time for both Cartesian [26, 27, 247, 248, 249] and non-Cartesian [250] ^{129}Xe imaging. CS reconstruction has recently been explored for radial gas exchange imaging data [250], but its effect on the key clinical metrics of gas exchange imaging; the RBC:M, RBC:Gas and M:Gas signal ratios, has yet to be studied.

^{129}Xe imaging is typically performed with enriched Xe (>85% ^{129}Xe isotopic abundance), due to its increased concentration of the spin-$\frac{1}{2}$ isotope of interest when compared to natural abundance (NA, 26% ^{129}Xe) [97]. NA Xe is typically an order of magnitude cheaper than enriched [164] and, together with recent progress in gas polarisation [217], this provides motivation for NA ^{129}Xe imaging. The feasibility of NA ^{129}Xe gas-phase ventilation imaging has been established

[97], but dissolved ^{129}Xe imaging with NA Xe has yet to be demonstrated.

In this work, the gas exchange ratios from both CS and a conventional gridding [57] reconstruction are compared retrospectively in healthy volunteers, patients with COPD and patients hospitalised following COVID-19 infection. We then assess the feasibility of imaging dissolved NA ^{129}Xe with CS reconstruction in four healthy volunteers.

6.2 Theory

Non-Cartesian k-space data were reconstructed with CS by solving a non-linear regularised optimisation problem [60]:

$$x = \mathrm{argmin}_x ||Ax - y||_2 + \lambda_1 ||\Psi x||_1 + \lambda_2 ||Tx||_1 \tag{6.2.1}$$

where x is the reconstructed image, y is the acquired k-space data, $\lambda_{1,2}$ are regularisation weights, Ψ is the sparsity operator and T is the finite difference transform for image total variation. The ADMM algorithm [61] was used to solve Equation 6.2.1, with $\Psi = I$ (identity) because hyperpolarised gas lung images have natural sparsity in the image domain, due to lack of background signal from surrounding tissue [26, 27]. A is the non-uniform Fourier transform operator:

$$A = PF \tag{6.2.2}$$

$$y = Ax + \nu \tag{6.2.3}$$

where P is the sampling density, F is the Fourier transform and ν represents measurement noise. In non-Cartesian gridding reconstruction, x is calculated by convolving the radial k-space data with a gridding kernel to resample to a Cartesian grid, then applying the inverse fast Fourier transform [35, 57, 54].

6.3 Methods

One dissolved ^{129}Xe spectroscopic imaging dataset (healthy female, 41 years) was retrospectively analysed to optimise the CS reconstruction parameters, as in previous work [26, 247]. Images were acquired with the 3D radial MESI sequence described in Section 4.2.3 and in [189].

Image reconstruction and analyses were carried out in Matlab (version 2022a, Mathworks, Natick, MA) on a Windows PC (Intel® Core™ i7-4790 processor). Conventional gridding reconstruction was carried out as described in Section 4.2.3. CS reconstruction was performed using the Berkeley Advanced Reconstruction Toolbox (BART) [251] `pics` function to solve Equation 6.2.1 for each ^{129}Xe signal. The values of λ_1, λ_2 and the acceleration factor (AF) for random undersampling were chosen empirically using a grid search. Specifically, the

CS reconstruction was performed many times with different combinations of λ_1 (ranging from 0.001 - 0.01), λ_2 (ranging from 0.0001 - 0.001) and AF (ranging from 1 - 4). The fully sampled data consisted of 934 angularly-equidistant radial spokes which were acquired in a random order. Retrospective CS reconstructions were performed using all data (AF = 1) and with AFs of 2, 3 and 4 by using only the first 467, 311 and 234 radial spokes respectively. Figure 6.3.1 shows the radial k-space spokes for each undersampling factor. The CS-reconstructed dissolved ^{129}Xe ratio maps were compared with the gridding-reconstructed maps and the final parameters were chosen based on the combination that minimised the normalised mean absolute error (NMAE) between the ratio maps. NMAE was calculated from the masked images, where masking was carried out by applying a noise threshold to the membrane image from gridding reconstruction, and was defined as:

$$\mathrm{NMAE} = \frac{\frac{1}{n}\sum_i^n |y_i - \hat{y}_i|}{\bar{y}} \mathrm{x}100\% \tag{6.3.1}$$

where n is the number of voxels within the lung mask, y_i and $\hat{y}_i$ are the ratio map value for the i^{th} voxel from gridding and CS reconstruction respectively, and $\bar{y}$ is the mean of the gridding ratio map.

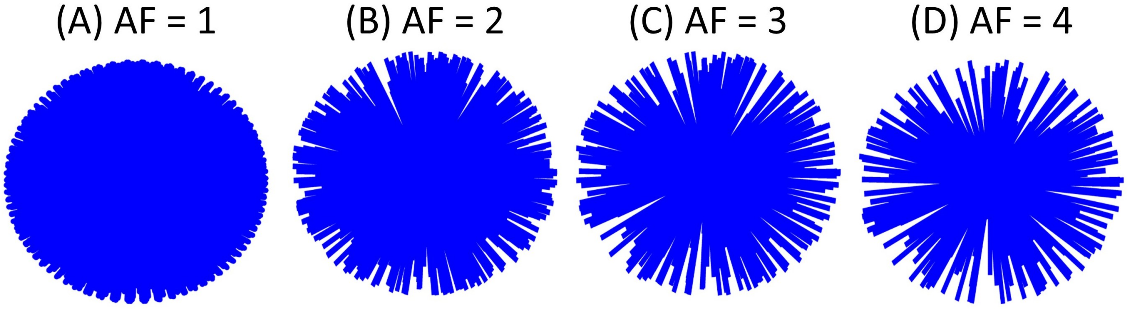

Figure 6.3.1: Radial k-space trajectories with different acceleration factors (AF): (A) fully-sampled, (B-D) AF = 2 - 4.

To evaluate the fidelity of the CS reconstruction, metrics including the NMAE, coefficient of variation (CV) and the pixelwise linear regression of both the ratio maps and the raw images were assessed. SNR was calculated as the ratio of the mean absolute signal within the lung mask to the standard deviation of the real signal from the first three and last three slices (outside of the lung).

The optimised CS and conventional gridding reconstructions were both evaluated retrospectively for a cohort of 41 subjects: 13 healthy volunteers, 5 COPD patients and 23 patients hospitalised post-COVID-19 infection (PCH), 14 of whom had residual lung abnormalities on CT or structural lung MRI three months after hospital admission (PCH-RLA). These participants were imaged using the same parameters as described above and in Section 4.2.3, and data were retrospectively undersampled. The mean ratio and CV values from both reconstructions were compared using linear regression and Bland-Altman analysis. Image fidelity was evaluated with NMAE and pixelwise linear regression of the ratio maps. To compare the regional trends in the ratio maps from gridding and CS, the mean and standard deviation of RBC:M, RBC:Gas and M:Gas in each slice in the anterior-posterior direction were calculated.

Four healthy volunteers were then prospectively imaged (with the fully-sampled sequence) with 0.8 - 1 L hyperpolarised NA ^{129}Xe (dose equivalent volume (DEV) [164] of ∼80 ml for the 1 L dose, with 30% polarisation) on a 1.5T GE Artist scanner to assess the feasibility of NA dissolved ^{129}Xe CS imaging. 3/4 participants additionally underwent dissolved imaging with the same volume of enriched ^{129}Xe (DEV = 260 ml) in the same scanning session and these data were added to the healthy volunteer retrospective cohort for the subsequent analysis.

6.4 Results

6.4.1 CS Optimisation

To balance image fidelity and noise suppression, λ_1 = 0.003 and λ_2 = 0.0003 were chosen for the optimised CS implementation via grid search. These parameters provided good agreement between the gridding and CS-reconstructed mean ratio values at each AF (Figure 6.4.1, Table 6.4.1) and worked well over a range of other lung images tested. As the AF increased from 1 to 4, NMAE increased (RBC:M: 3.5% - 6.9%, RBC:Gas: 5.2% - 8.3%, M:Gas: 4.2% - 4.8%) and CV decreased (RBC:M: 0.17 - 0.16, RBC:Gas: 0.29 - 0.22, M:Gas: 0.28 - 0.20), as expected from the effect of CS smoothing and loss of information. Linear regression of the normalised pixelwise signal showed high R^2 values, which decreased with increasing AF (RBC: 0.99 - 0.91, M: 0.98 - 0.90, Gas: 0.98 - 0.87). As a compromise between image fidelity and scan time reduction, AF = 2 was chosen for the optimised CS implementation. With these optimised parameters, the distribution of the gas exchange ratios was preserved, resulting in similar maps and histograms (Figure 6.4.2). The CS reconstruction time was similar to that of gridding (<5 minutes).

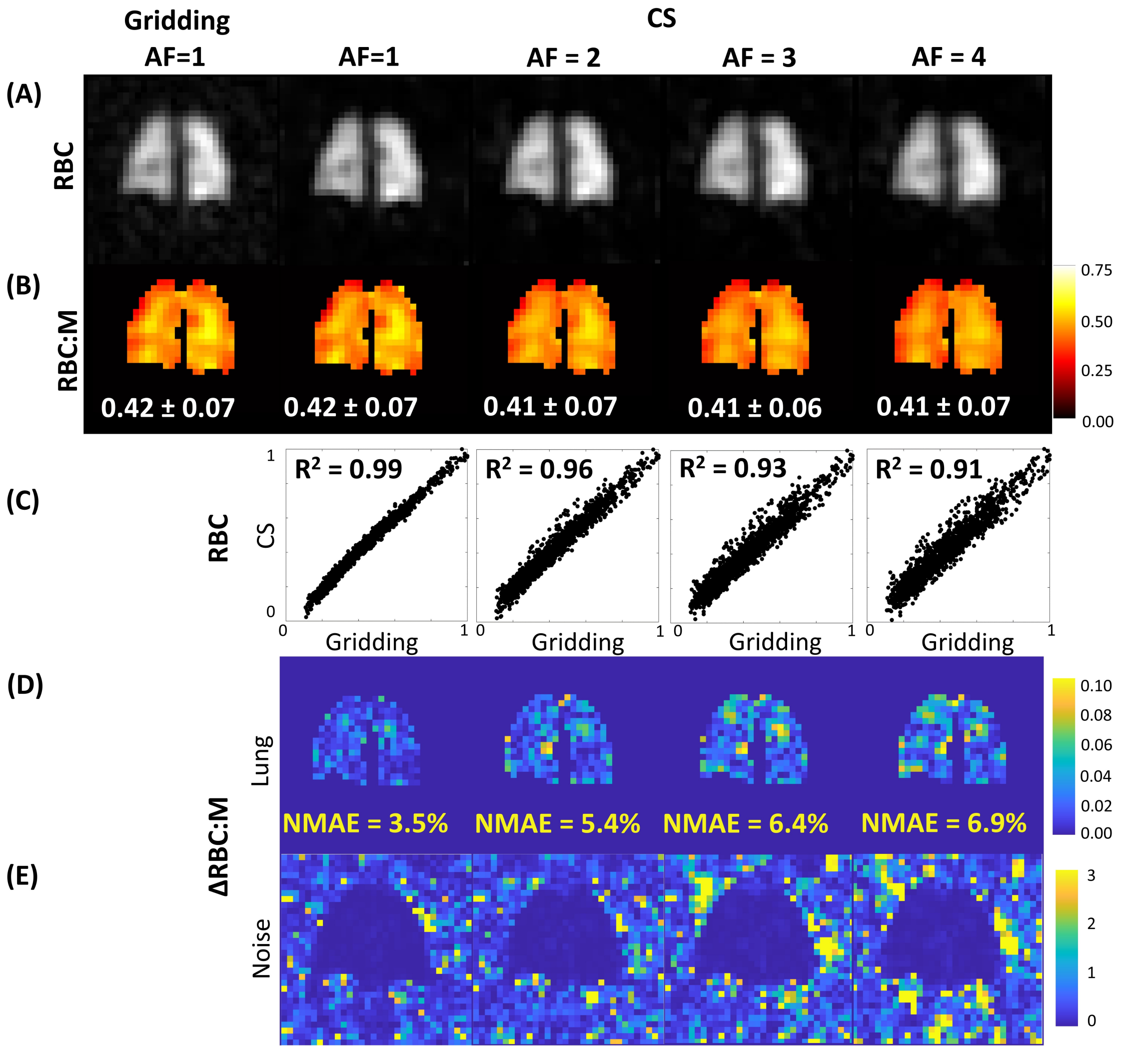

Figure 6.4.1: (A) RBC signal and (B) RBC:M from conventional gridding reconstruction (100% sampling) and CS reconstruction with AF = 1 - 4 and (C) the corresponding linear regression of the normalised pixelwise RBC signal within the lung mask. The absolute difference between the RBC:M maps from gridding and CS reconstruction are shown for each AF, for (D) inside and (E) outside the lung mask.

Table 6.4.1: CS reconstruction parameters for the healthy volunteer dataset used in the optimisation process for AF = 1, 2, 3 and 4. The mean ratio values and CV from conventional gridding reconstruction are included for comparison.

		RBC:M	RBC:Gas	M:Gas
Mean	Gridding	0.42 ± 0.07	0.0040 ± 0.0010	0.0098 ± 0.0021
	CS AF = 1	0.42 ± 0.07	0.0041 ± 0.0012	0.0100 ± 0.0028
	CS AF = 2	0.41 ± 0.07	0.0041 ± 0.0011	0.0099 ± 0.0025
	CS AF = 3	0.41 ± 0.06	0.0041 ± 0.0010	0.0099 ± 0.0022
	CS AF = 4	0.41 ± 0.07	0.0041 ± 0.0009	0.0099 ± 0.0020
CV	Gridding	0.16	0.24	0.22
	CS AF = 1	0.17	0.29	0.28
	CS AF = 2	0.17	0.26	0.25
	CS AF = 3	0.16	0.23	0.22
	CS AF = 4	0.16	0.22	0.20
NMAE (%)	CS AF = 1	3.5	5.2	4.2
	CS AF = 2	5.4	6.2	4.2
	CS AF = 3	6.4	7.3	4.4
	CS AF = 4	6.9	8.3	4.8
R^2	CS AF = 1	0.92	0.87	0.88
	CS AF = 2	0.81	0.78	0.82
	CS AF = 3	0.73	0.83	0.89
	CS AF = 4	0.69	0.79	0.90
		RBC	M	Gas
R^2	CS AF = 1	0.99	0.98	0.98
	CS AF = 2	0.96	0.95	0.95
	CS AF = 3	0.93	0.92	0.91
	CS AF = 4	0.91	0.90	0.87

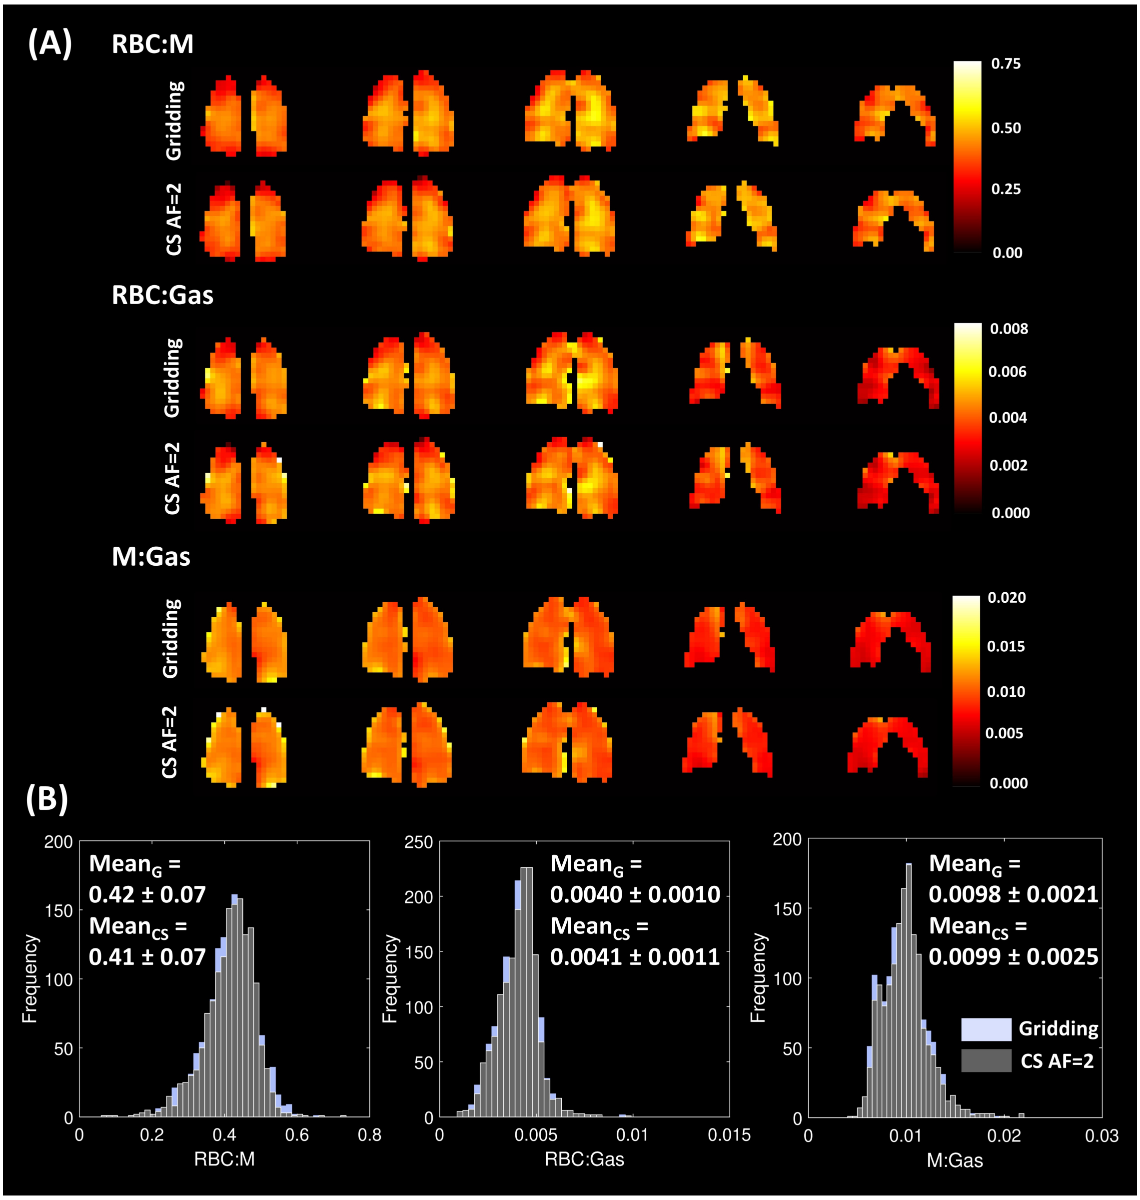

Figure 6.4.2: (A) Ratio maps from five slices of the lung for the healthy dataset used to optimise the CS parameters and (B) histograms of the ratio maps from gridding and CS reconstruction (AF = 2) for RBC:M, RBC:Gas and M:Gas.

6.4.2 Retrospective Analysis

Patient demographics and key results are shown in Table 6.4.2 and example ratio maps are shown in Figure 6.4.3. The mean gas exchange ratio values were preserved with the CS reconstruction (Figure 6.4.4A); the R^2 values for each ratio were high (0.99) and Bland-Altman analysis showed evenly spread residuals. The NMAE values between the gridding and CS derived mean ratios were 3.8% for RBC:M, 3.3% for RBC:Gas and 1.5% for M:Gas. The CV values from gridding and CS reconstruction were similar for RBC:M, RBC, M and Gas; linear regression analysis showed high R^2 values (> 0.91) and the Bland-Altman biases were small (Figure 6.4.4B). For RBC:Gas and M:Gas, CS CV was significantly higher than that of gridding (Wilcoxon signed rank exact test: $p < 0.001$). Lower R^2 values of

0.73 and 0.30 for RBC:Gas and M:Gas respectively resulted from a small number of outliers (driven by low signal pixels at the periphery of the lungs with very high ratio values).

Table 6.4.2: Subject demographics and the RBC SNR, NMAE and adjusted R^2 of the pixelwise linear regression for the CS reconstruction with AF = 2. Normally distributed variables (as determined with Shapiro-Wilk tests) are given as mean ± standard deviation, whereas non-normally distributed variables are given as median (range). Gas exchange ratio values are from gridding reconstruction.

		Healthy	COPD	PCH	PCH-RLA
n (female)		16 (5)	5 (3)	9 (2)	14 (2)
Age (years)		40 ± 13	60 ± 9	57 ± 9	66 ± 9
RBC:M		0.41 ± 0.08	0.23 ± 0.05	0.21 (0.17 - 0.36)	0.19 ± 0.06
RBC:Gas		0.0040 ± 0.0011	0.0019 ± 0.0004	0.0025 ± 0.0006	0.0017 ± 0.0005
M:Gas		0.0098 ± 0.0024	0.0084 ± 0.0025	0.0115 (0.0088, 0.0176)	0.0097 ± 0.0016
RBC SNR	Gridding	9.2 ± 3.3	2.1 (2.0 - 7.4)	4.9 (3.6 - 8.7)	3.4 (1.9 - 12.6)
	CS AF = 2	27.0 (10.2 - 100.0)	2.7 (2.3 - 18.9)	9.3 ± 2.4	5.3 (2.1 - 28.4)
NMAE (%)	RBC:M	9.9 ± 3.3	27.6 ± 9.6	17.2 ± 3.9	19.5 ± 9.8
	RBC:Gas	10.5 ± 3.2	33.6 (12.7 - 33.9)	18.8 ± 4.0	21.5 ± 9.5
	M:Gas	5.6 ± 1.3	8.9 ± 1.8	5.8 ± 1.0	6.3 ± 2.0
R^2	RBC:M	0.71 ± 0.11	0.40 (0.30 - 0.81)	0.61 ± 0.08	0.68 ± 0.18
	RBC:Gas	0.78 (0.56 - 0.88)	0.54 ± 0.18	0.62 ± 0.06	0.61 ± 0.15
	M:Gas	0.83 (0.40 - 0.94)	0.75 ± 0.02	0.83 ± 0.07	0.72 ± 0.17

The pixelwise linear regression between the gridding and CS derived ratio maps showed good agreement: R^2 = 0.65 ± 0.16 for RBC:M, 0.65 ± 0.14 for RBC:Gas and 0.78 (0.40 - 0.94) for M:Gas, across the 44 subjects. Regional trends in the signal ratios were also preserved with the CS reconstruction, as shown in Figure 6.4.5 for a healthy participant and a PCH-RLA patient for the anterior-posterior gravitational gradient.

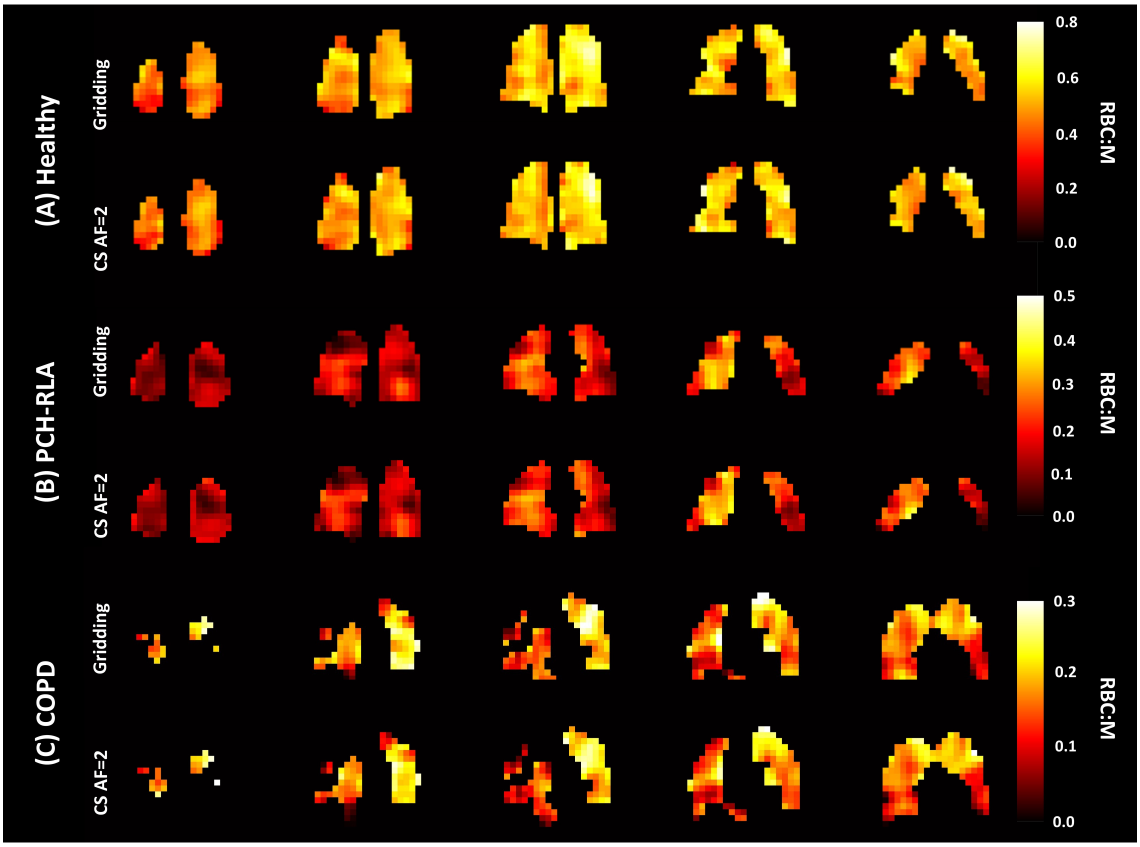

Figure 6.4.3: RBC:M maps from gridding and CS reconstruction for five lung slices (posterior to anterior) for (A) a healthy participant (NMAE = 8%), (B) a PCH-RLA patient (NMAE = 11%) and (C) a patient with COPD (NMAE = 11%). For the two lung disease patients, regions of RBC transfer defect are preserved in the CS maps.

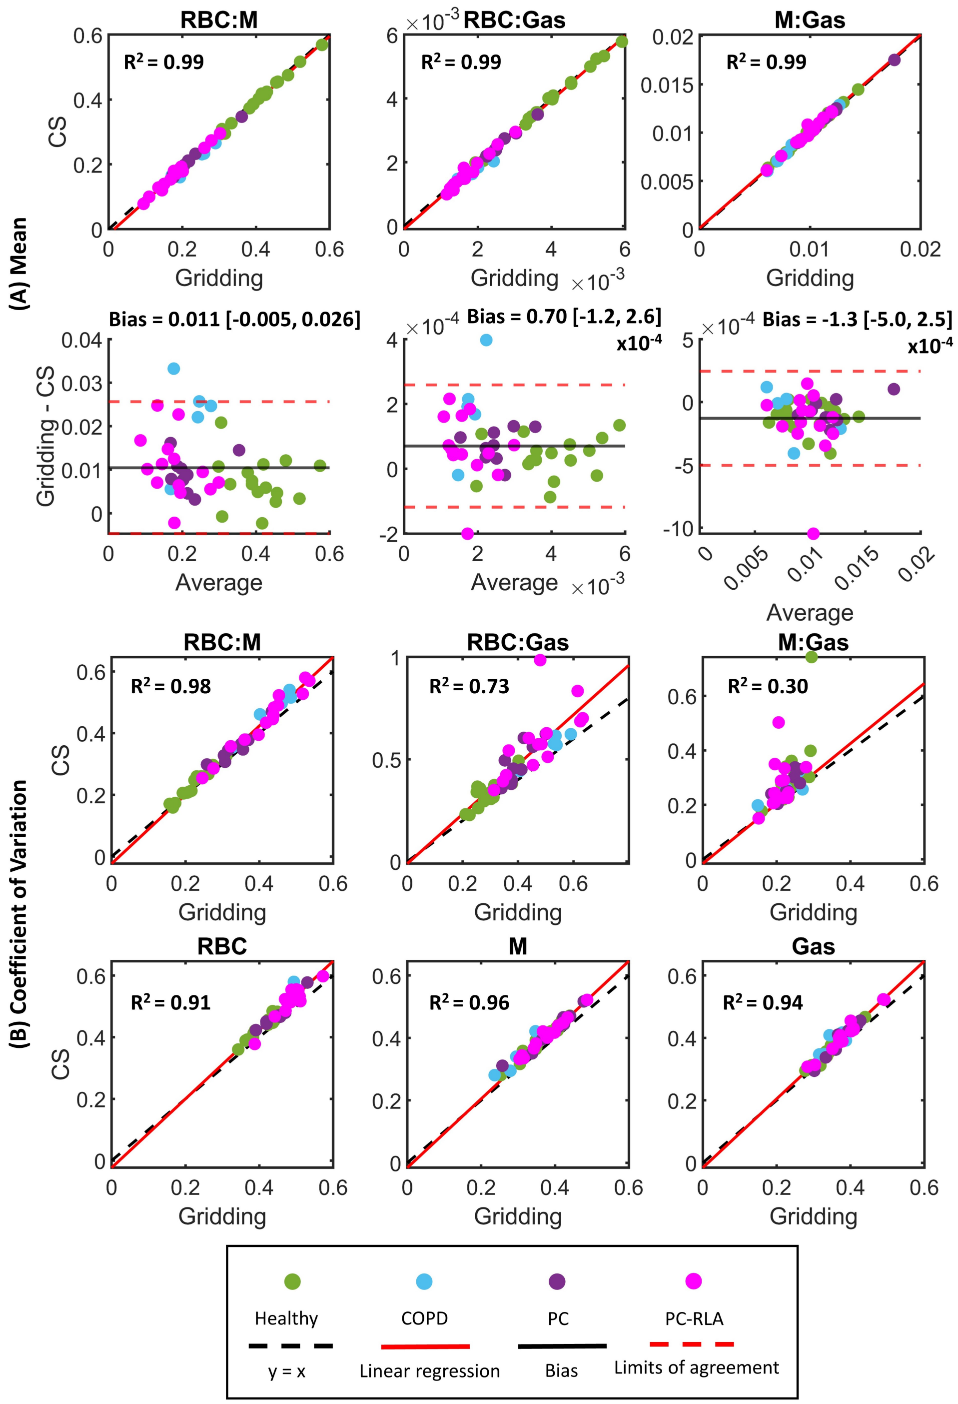

Figure 6.4.4: (A) Linear regression and Bland-Altman plots of whole-lung mean for RBC:M, RBC:Gas, and M:Gas, comparing gridding and CS with AF = 2. (B) Linear regression plots of coefficient of variation (CV) for RBC:M, RBC:Gas, M:Gas, and the RBC, M, and Gas signals.

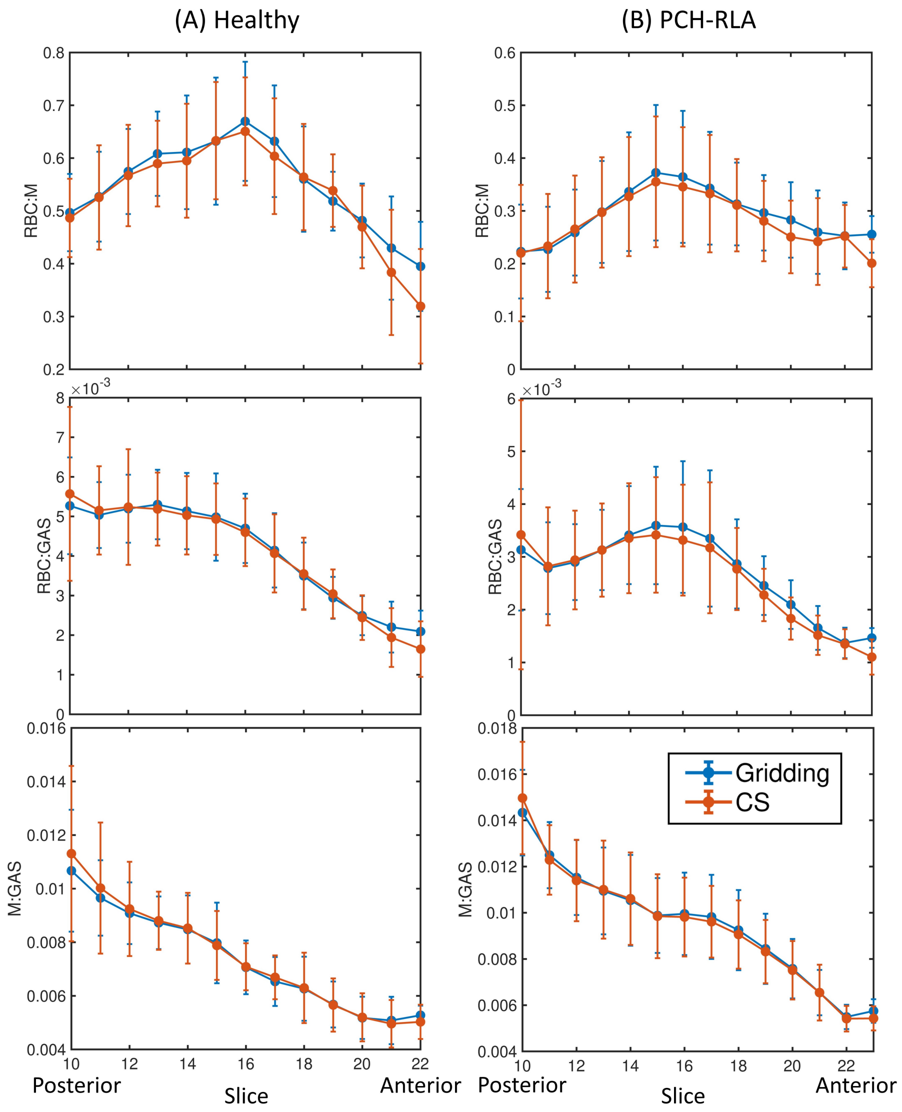

Figure 6.4.5: Comparison of the anterior-posterior regional trends in RBC:M, RBC:Gas and M:Gas for gridding and CS reconstruction for (A) a 30 year old male healthy volunteer and (B) a 62 year old male patient who was hospitalised following COVID-19 infection and had residual lung abnormalities. The points show the mean value over the slice and the error bars represent the standard deviation. Connecting lines have been drawn to make the regional trends clearer.

The NMAE values were higher in the retrospective cohort than from the CS simulations, and a median NMAE of 15.0% was found for RBC:M across data from all subjects. The NMAE values for RBC:M, RBC:Gas and M:Gas were observed to scale non-linearly with the dissolved ^{129}Xe gridding SNR values (Figure 6.4.6A-D). In particular, when gridding RBC SNR dropped below ∼5, NMAE increased sharply; overall the relationship was found to be best explained by a power law:

$$\mathrm{NMAE} = a\mathrm{SNR}^{k} \tag{6.4.1}$$

where a is a scaling constant and $k < 0$ is the exponent of the power law.

CS suppressed noise such that RBC SNR increased on average by a factor of

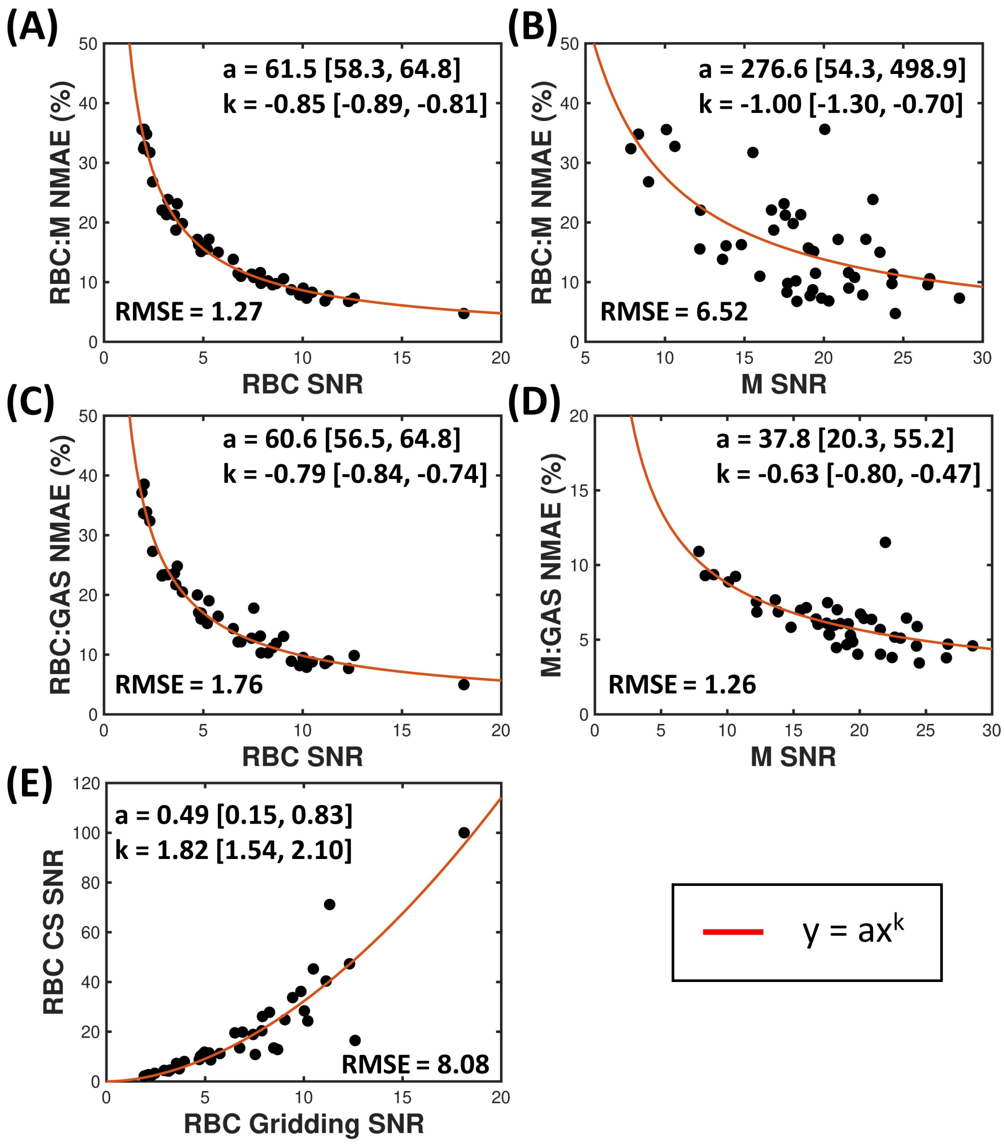

Figure 6.4.6: A non-linear power law relationship was found between the NMAE of the CS ratio maps, relative to the gridding ratio maps, and the dissolved ^{129}Xe gridding SNR (A-D) and the SNR of the RBC image derived from gridding and CS (E). The fitted power law parameters, a and k, are given along with their 95% confidence intervals. RMSE = root mean squared error.

two for the retrospective analysis cohort. A three-fold SNR increase was found for the healthy volunteers CS images. The relationship between CS and gridding RBC SNR was non-linear, and followed a power law with k = 1.82, as shown in Figure 6.4.6E.

6.4.3 Natural Abundance ^{129}Xe

Table 6.4.3 shows the subject demographics and CS results for the healthy participants who were imaged with NA ^{129}Xe. As expected from the lower isotopic concentration, the SNR of the NA dissolved Xe images was worse than those acquired with enriched ^{129}Xe (Figure 6.4.7). For the three participants who had enriched and NA dissolved ^{129}Xe imaging in the same session, the ratio of enriched to NA gridding RBC image signal was 3.0 (compared to a ∼3.3 difference in ^{129}Xe concentration). CS was successful at suppressing noise in the NA ^{129}Xe images; at AF = 2, RBC SNR was increased three-fold compared to gridding reconstruction. The mean ratio values were similar between the gridding (RBC:M = 0.41 ± 0.04, RBC:Gas = 0.0055 ± 0.0010, M:Gas = 0.014 ± 0.004) and CS AF = 2 reconstructed maps (RBC:M = 0.41 ± 0.04, RBC:Gas = 0.0056 ± 0.0010, M:Gas = 0.014 ± 0.004) and the R^2 and NMAE values were comparable to those from enriched ^{129}Xe imaging, although the M:Gas NMAE was higher (see Table 6.4.2).

Table 6.4.3: Subject demographics and metrics for comparing the gridding and CS (AF = 2) reconstructions for the healthy subjects who underwent NA dissolved ^{129}Xe imaging; three of these subjects are in common with the healthy group. Normally-distributed data are displayed as mean ± standard deviation and non-normally distributed data are displayed as median (range).

		Healthy NA
n (female)		4 (2)
Age (years)		33 ± 8
RBC:M		0.41 ± 0.04
RBC:Gas		0.0055 ± 0.0010
M:Gas		0.0138 ± 0.0036
RBC SNR	Gridding	5.9 ± 1.7
	CS	19.8 ± 10.3
NMAE (%)	RBC:M	15.0 ± 4.2
	RBC:Gas	17.4 ± 7.7
	M:Gas	10.3 ± 7.5
R^2	RBC:M	0.54 ± 0.02
	RBC:Gas	0.62 (0.19 - 0.68)
	M:Gas	0.83 (0.29 - 0.88)

The top panel of Figure 6.4.7 shows the enriched and NA RBC images for a healthy male participant. The NA images exhibited more heterogeneity than the enriched image, with CV = 0.32 (gridding), 0.37 (CS AF = 1) and 0.35 (CS AF = 2), compared to CV = 0.23 from enriched ^{129}Xe imaging (gridding). In a second healthy volunteer (Figure 6.4.7 bottom panel), the NA RBC gridding image SNR (6.9) was higher than that of the first healthy volunteer (4.8) and there was a closer resemblance between the enriched and NA xenon spectroscopic images. The NA images were again more heterogeneous than the enriched image: CV = 0.20 (gridding), 0.26 (CS AF = 1) and 0.23 (CS AF = 2), compared to CV = 0.15 from enriched ^{129}Xe imaging (gridding).

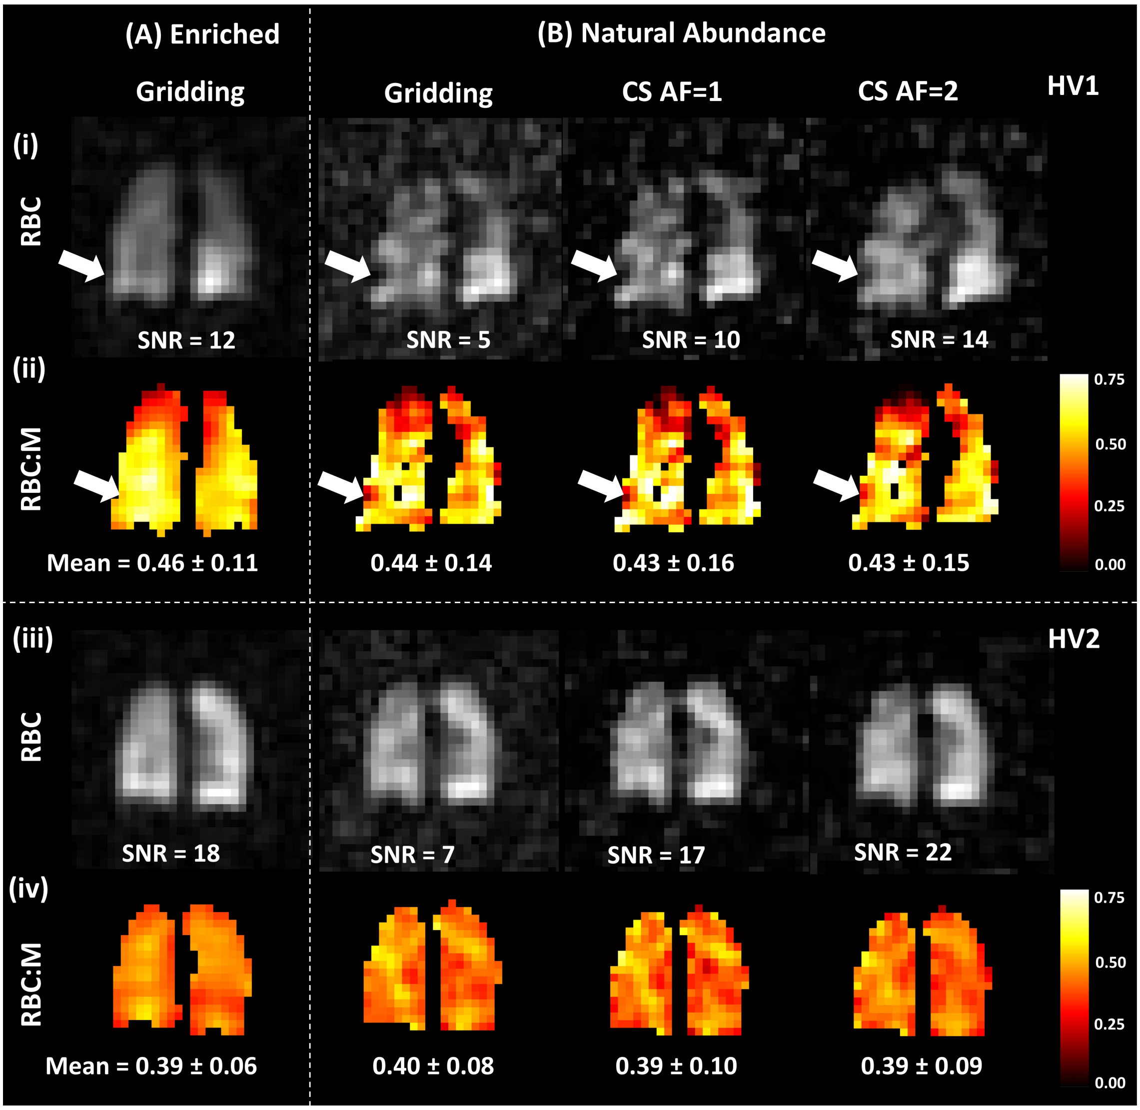

Figure 6.4.7: (A) Enriched and (B) NA dissolved ^{129}Xe (i and iii) RBC signal and (ii and iv) RBC:M ratio maps for a central lung slice in two healthy volunteers (HV). The mean ± standard deviation of the RBC:M ratio maps are given below. The NA signal images and ratio maps exhibited greater heterogeneity than those from enriched ^{129}Xe imaging. For HV1, there are areas of very low/no signal, for example in the lower right lung (white arrow), which appear as RBC transfer defects in the RBC:M maps and are not present in the enriched ^{129}Xe image.

6.5 Discussion

We have presented a CS reconstruction method for dissolved ^{129}Xe spectroscopic imaging data, using the open-source BART software [251]. We chose to optimise the CS parameters based on the gas exchange ratios maps of one healthy dataset, in addition to the RBC, M and gas signals, because it is the ratios that are ultimately used in current approaches of clinical interpretation. Although the optimal parameters may be subject-dependent, in our experience the changes are small. We empirically selected AF = 2 for our optimised CS reconstruction because this provided the best balance between image fidelity and scan time reduction. With this AF, only 467 radial spokes are required and so the breath hold could be reduced to 7 s in a prospectively undersampled implementation, compared to the 14 s used in our centre routinely [35, 189]. A shorter scan time should be more comfortable for dyspnoeic and paediatric patients.

SNR increased with increasing AF for both the optimisation and NA reconstructions, despite the greater level of undersampling. This is likely related to the temporal ordering of the undersampled spokes. Figure 6.5.1 compares the SNR from gridding and CS reconstruction when the spokes are sampled consecutively and randomly in time. Using consecutive spoke ordering (in time) increased the SNR relative to that found using random temporal ordering. This is because, with consecutive ordering, the latter 1/2, 2/3 or 3/4 of the radial spokes which have undergone more T_1 and RF pulse induced depolarisation of the gas-phase signal reservoir [160], are discarded. As more spokes are discarded, the average signal of the spokes selected for reconstruction increases and this partly explains why SNR is seen to increase with AF. The high SNR seen with increased AF also has a contribution from the nonlinear noise suppression of the compressed sensing algorithm [60]. However, the high SNR values at AF = 4 are not ‘true’ SNR as per its standard definition, and may be due to over-smoothing and/or reconstruction errors from heavy undersampling and from the conventional standard deviation-based approach to calculate the SNR. This may not be appropriate for highly accelerated CS images where the noise appears as a ‘flat’ background.

Both the gridding and CS images were masked based on the gridding M image, however this occasionally led to outliers in the ratio maps. Some low signal voxels at the periphery of the gridding masks may have been discarded as noise in the CS reconstruction but were still included in the ratio maps. In a prospective CS implementation, masking would be performed using the CS M images instead, alleviating this issue.

CS reconstruction improved SNR by suppressing image noise. Our results are in agreement with those of Plummer et al. [250], who found an increase in the overall (RBC + M) dissolved-phase ^{129}Xe SNR with CS and good preservation of median dissolved-phase to gas signal ratio in healthy volunteers and patients who had undergone a hematopoietic stem cell transplant. Here, we have gone further by separating the dissolved-phase ^{129}Xe signal into the discrete RBC and M contributions, and examining the effect of the CS reconstruction on the gas exchange ratios in healthy participants and patients with lung disease.

Although the increase in RBC SNR was lower for the lung disease patients than the three-fold increase seen for the healthy volunteers, the healthy participants had reasonably high gridding SNR, whereas this was low (<5) for many of the patients. The CS SNR enhancement was dependent on the input (gridding) SNR, as shown in Figure 6.4.6E . When the input SNR is higher, there is increased separation between the signal in the lung region and noise in the image domain representation, providing a better ground truth for the iterative CS reconstruction and leading to better denoising performance. There were two potential outliers in Figure 6.4.6E, where the data points were distanced from the power law curve (however they still followed the general pattern). The reasons for these outliers are unclear. The data point where the CS RBC SNR was higher than expected (gridding SNR = 11.2, CS SNR = 71.2) came from a healthy volunteer with particularly high mean RBC:M ratio (0.52; mean RBC:M for healthy volunteers = 0.41). The data point where the CS RBC SNR was lower than expected (gridding SNR = 12.6, CS SNR = 16.5) came from a PC-RLA patient. Although the gridding SNR was high for this subject, there were some artefacts/non-Gaussian noise in the image which perhaps the CS denoising algorithm struggled to remove with thresholding. Whilst outside the scope of this work, a comparison of the denoising effect of CS with novel denoising techniques such as tensor Marchenko-Pastur principle component analysis [252], Global Local Higher Order Singular Value Decomposition [253] and deep learning [254, 255] would be an interesting future avenue of research.

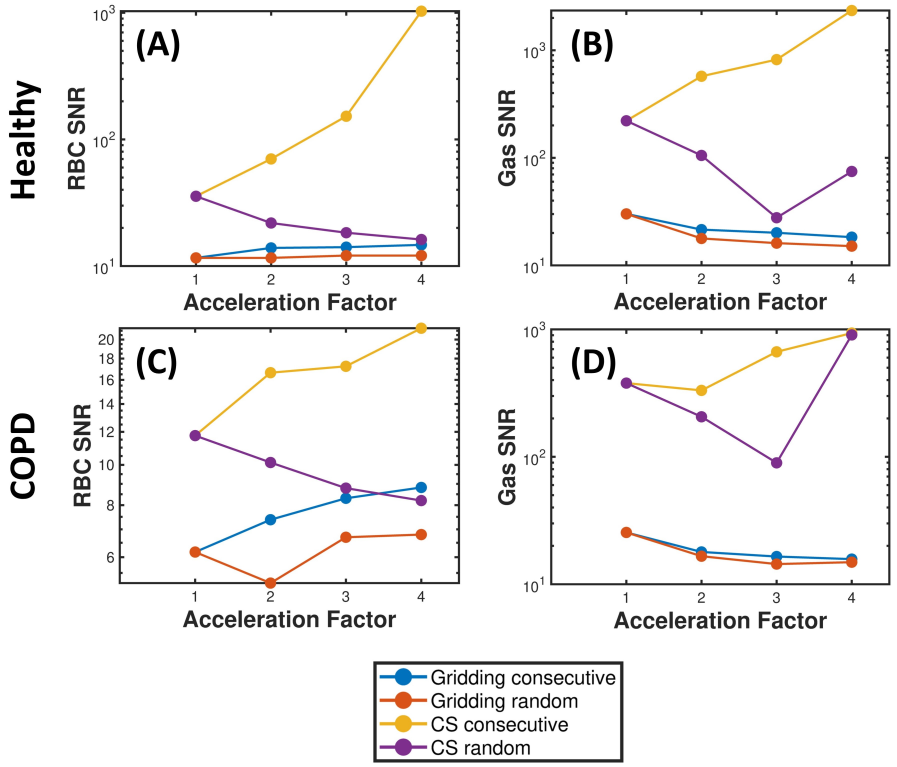

Figure 6.5.1: SNR of the signal from ^{129}Xe dissolved in the RBCs and the gaseous-phase ^{129}Xe at AF = 1, 2, 3 and 4 for (A-B) a healthy participant and (C-D) a patient with COPD. Two different sampling patterns were compared for both CS and gridding: consecutive temporal undersampling of the chosen spokes (taking the first half/third/quarter) and random temporal spoke undersampling (taking spokes from throughout the acquisition to 'average out' the polarisation decay). Note the logarithmic y-axis scale.

For patients with low dissolved-phase ^{129}Xe signal, although CS reconstruction provided a beneficial increase in SNR, there was an associated increase in NMAE between the CS and gridding ratio maps; for RBC gridding image SNR <5, errors for RBC:M rose above 15%. A similar relationship between SNR and reconstruction error was found by Lazarus et al. [256]. Nevertheless, in cases of such low SNR, it is not appropriate to use the gridding images as the ground truth, because comparison of signal with noisy images is not reliable. Given the lack of an alternative ground truth and therefore the uncertainty in the CS algorithm performance when signal is low, we cannot reliably conclude that the CS reconstruction is capable of recovering RBC data with raw SNR <5. This highlights the limitations of retrospective denoising and the importance of techniques to prospectively increase dissolved-phase ^{129}Xe SNR, such as increasing the polarisation, dose/^{129}Xe fraction and using highest sensitivity RF coils where feasible. Even for higher SNR cases, NMAE values were sometimes reasonably high. For five subjects with gridding RBC SNR ≥ 8, the NMAE of RBC:M was $\geq 10\%$. However, it is important to consider the NMAE values in the context of the repeatability of the gas exchange ratio measurements themselves. In a previous work [189], we evaluated the intra- and inter-session repeatability of these metrics, by scanning eight healthy volunteers at three time points - two scans ∼5 minutes apart and a third scan ∼4 hours earlier or later (see Section 7.2.2). The CV across the three scans for the mean ratio values were (4.5 ± 0.6)% for RBC:M, (7.0 ± 2.8)% for RBC:Gas and (5.8 ± 2.1)% for M:Gas respectively for eight healthy volunteers. The NMAE values between the ratio maps acquired in the same scanning session were similar to the NMAE found between the gridding and CS maps: (12.9 ± 2.2)% for RBC:M, (16.1 ± 3.5)% for RBC:Gas and (9.5 ± 4.0)% for M:Gas. Similarly, the inter-session CV values for the mean ratios have been reported to be 11-20% by Hahn et al. [165] in 13 healthy volunteers scanned at two time points one month apart with the 1-point Dixon method [32, 31]. In a study of 18 healthy participants and 14 with COPD, Garrison et al. found intra-session CV values of 7-13% [245]. Therefore, here, for SNR >5, the NMAE values of the ratio maps are similar to their repeatability metrics.

NA dissolved ^{129}Xe imaging had not previously been explored due to the three-fold lower signal strength of NA when compared to enriched ^{129}Xe. Here we have demonstrated the feasibility of dissolved NA ^{129}Xe imaging for the first time in healthy subjects. We used advancements in polarisation technology and the denoising effect of CS to achieve an acceptable RBC image SNR from the same volume of NA ^{129}Xe as was used for enriched dissolved ^{129}Xe imaging. A limitation of this work is that NA imaging was only acquired in a small number of healthy participants who did not have any gas exchange limitation. Further prospective evaluation in increased subject numbers, including patients with lung disease, is required to fully evaluate the differences between enriched and NA dissolved ^{129}Xe imaging and to assess the feasibility of this approach when gas transfer and/or RBC SNR is low. These initial results suggest that low-cost NA dissolved Xe imaging is feasible in healthy volunteers, although might be challenging for lung disease patients who have reduced gas transfer.

This work did not include any prospectively undersampled data, but since the two-fold radial undersampling was carried out by removing the last half of the spokes, this is equivalent to a prospective implementation with half the scan duration. Whilst for gas-phase ^{129}Xe imaging, the reduction in the required number of RF pulses achieved with compressed sensing means that the signal strength can be increased by using higher flip angles [63, 257], this is not applicable for dissolved-phase ^{129}Xe imaging. For dissolved-phase imaging, a change in flip angle needs to be accompanied by a change in TR in order to maintain the same $\mathrm{TR}_{90°,\mathrm{equiv}}$ [223] and hence underlying dissolved ^{129}Xe magnetisation distribution. Therefore, SNR would not improve and there would be no change in signal behaviour upon performing a prospectively undersampled compressed sensing implementation for our sequence. A benefit of our retrospective approach is that it makes it possible to retrieve usable images from datasets where the subject failed to hold their breath for the entire 14 s or where there was motion in the second half of the scan, which can typically be identified from the data itself. This should help to reduce the number of failed and repeated scans.

Further improvement to image quality may be possible by incorporating prior knowledge into the CS reconstruction. T_1 and RF pulse-induced decay can result in image blurring and coherent undersampling artefacts, which can be reduced by accounting for signal decay in the iterative reconstruction. Plummer et al. incorporated global magnetisation decay directly into the Fourier operator in their CS reconstruction, which resulted in improved image sharpness and SNR for 2D/3D spiral ventilation and 3D radial gas exchange MRI [250]. Ajraoui et al. found that incorporating prior knowledge of either magnetisation decay or lung structure (via the inverse of a registered proton image) allowed for the use of higher undersampling factors and reduced image blurring in hyperpolarised ^{3}He lung imaging [258].

6.6 Conclusions

CS reconstruction of dissolved ^{129}Xe spectroscopic imaging improved image SNR and enabled reduced scan time, whilst preserving mean ratio values. This benefits patients with breathlessness and/or low gas transfer, allowing for faster gas exchange imaging. Therefore, CS reconstruction may be suitable as a replacement for conventional gridding in our dissolved ^{129}Xe image reconstruction pipeline, although caution should be taken with patient groups with very low RBC SNR, such as COPD and ILD. Preliminary results with NA dissolved ^{129}Xe imaging in healthy participants suggest that reduced cost NA gas exchange imaging could be made feasible using CS.

Chapter 7

Mapping the Amplitude and Phase of Dissolved ^{129}Xe Red Blood Cell Signal Oscillations with Keyhole Spectroscopic Lung Imaging

7.1 Introduction

The ^{129}Xe RBC signal oscillation was introduced in Section 3.5; here we reiterate the main concepts.

In dissolved-phase ^{129}Xe spectroscopy of the lungs, the M and RBC signals decay over the 10-15 s duration breath hold, due to RF pulse-induced depolarisation and T_1 relaxation of the gas-phase signal, which acts as a magnetisation 'reservoir' replenishing dissolved-phase signal (gas-phase T_1 $\sim$20 s) [81]. Additionally, the RBC signal is periodically modulated by the heartbeat [212, 208, 90]; these RBC oscillations have been suggested to originate from changes in the capillary blood volume over the cardiac cycle [37]. The underlying ^{129}Xe RBC signal is also dependent upon the diffusion rate of ^{129}Xe across the alveolar membrane and the haematocrit within the alveolar capillary bed [32, 259]. The cardiogenic RBC oscillations revealed in whole-lung ^{129}Xe spectroscopy provide a potential means to monitor blood flow in the pulmonary microvasculature, which is often affected by chronic lung disease [260, 261]. The ^{129}Xe spectroscopy-derived RBC oscillation has been shown to be sensitive to disease state (see Table 3.5.1): the oscillation amplitude is increased in IPF, NSIP and LHF patients but decreased in COPD and PAH patients when compared with healthy subjects [169, 35, 37, 137, 209]. Furthermore, the amplitude of the RBC oscillations may be able to differentiate between PH subtypes [209].

Imaging microvascular function directly is difficult due to the small vessel size and the effects of cardiac and respiratory motion. Larger pulmonary blood vessels' form and function can be imaged using CT, MRI and echocardiography, whilst microvascular perfusion can be deduced indirectly from DCE-MRI [261, 262, 263]. MRI can also be used to measure cardiac pulse wave signals; the temporal variations in blood pressure, flow velocity and volume flow rate through the arterial network, caused by changes in cardiac ejection rate which generate compression and relaxation waves [211]. The pulse wave velocity can be measured using phase-contrast or velocity-encoded MRI and used to infer artery stiffness [264, 265]. Pulmonary pulse transit time (pPTT) can be extracted from phase-resolved functional lung imaging (PREFUL), and measures the time taken for the pulse wave to travel from the central pulmonary artery to the microvasculature [266, 267]. However, unlike dissolved ^{129}Xe spectroscopy, these techniques do not capture the gas exchange dynamics in the pulmonary capillary bed.

Recently, Niedbalski et al. proposed a technique to spatially resolve the ^{129}Xe RBC oscillations from 3D radial dissolved-phase imaging [213]. The authors used the inherent oversampling at the centre of radial k-space (k_0) to obtain dynamic signal information via a post-acquisition ‘keyhole’ reconstruction [216]. By binning the k_0 data according to RBC signal amplitude, reconstructing images from the radial spokes in ‘low’ and ‘high’ bins and finding the difference between the resulting ‘low’ and ‘high’ keyhole images, regional RBC signal oscillations were mapped in a cohort of healthy volunteers, patients with IPF and with PAH. This method has since been further optimised using digital phantom simulations by Lu et al. [215] and applied to patients with CTEPH.

In previous work, a 1-point Dixon spectroscopic imaging approach has been used to separate the RBC and M signals [32, 31]. However, the 90° phase shift between the RBC and M signals evolves during readout, leading to image blurring. An alternative spectroscopic imaging method to differentiate the dissolved-phase signal is a MESI acquisition with [36, 188] or without [35, 189] IDEAL [179].

In this chapter, we adapted the keyhole oscillation mapping method from Niedbalski et al. to spatially resolve the RBC oscillations in dissolved-phase ^{129}Xe lung spectroscopic images acquired using our 3D radial MESI acquisition [189] in 28 healthy volunteers. In previous work [213, 215], the regional RBC oscillations were assumed to be in phase across the lung and only two keyhole images, corresponding to maximum and minimum RBC signal, were reconstructed. We hypothesise that the phase of the RBC oscillation is also spatially dependent, reflecting lung and PVD heterogeneity. We show that the previous assumption of constant phase leads to a reduced and sometimes negative local RBC oscillation amplitude. To address this, we introduce a sliding window (SW) technique to reconstruct additional keyhole projections, to map not only the amplitude, but also the phase of the regional RBC oscillations. We present the quantitative mapping of this phase as a novel means to probe the effects of the cardiac pulse wave in the pulmonary capillary bed.

7.2 Methods

7.2.1 Subject Details

Data from 28 healthy volunteers (12 female, 16 male, average age = 38.8 ± 11.1 years) were retrospectively analysed in this work.

7.2.2 MRI Acquisition

Imaging was performed on 1.5 T whole-body GE HDx (n = 24) and 1.5 T GE Artist (n = 4) clinical scanners, with 0.8 - 1L (depending on height; see Table 4.1.1) doses of enriched hyperpolarised ^{129}Xe, as described in Chapter 4. Dissolved-phase images were acquired using the four-echo 3D radial MESI sequence described in Section 4.2.3 [189]. To assess the inter- and intra-scan repeatability of the proposed RBC oscillation phase mapping method, eight of the healthy volunteers (three female, five male, ages 23 - 41) underwent additional imaging: once in the morning and once in the afternoon, with 3.5 - 5 hours between the two sessions. In one of the sessions the volunteers were scanned twice, with approximately five minutes between the scans.

7.2.3 Data Analysis

Image pre-processing, reconstruction and analysis were carried out in MATLAB (version 2022a, Mathworks, Natick, MA). The 3D radial spectroscopic data were reconstructed as described in Section 4.2.3 and the global RBC oscillation amplitude was evaluated as follows:

Starting from the k_0 signals from the spectrally reconstructed M and RBC k-space (Figure 7.2.1A, B), the M k_0 signal was normalised by its mean value and fit to a biexponential decay function (Figure 7.2.1C). To help detect the oscillations, the RBC k_0 signal was normalised by its mean value and corrected for RF and T_1 decay by multiplying by the inverse of the fit to the M k_0 signal (Figure 7.2.1D). A band-pass filter of 0.5 - 2.5 Hz was used to smooth the signal (Figure 7.2.1E) and a peak detection algorithm was used to identify the minima and maxima (red triangles and blue squares in Figure 7.2.1F and black triangles in Figure 7.2.2A). The mean k_0 peak-to-peak amplitude of the oscillations (α_{k0}) was found from the difference between the mean of the maxima and minima in the first ~7 s of the breath hold where the SNR was highest, which was then multiplied by 100%. Heart rate was estimated from the frequency of these oscillations. The first second of acquisition (~60 projections) was generally excluded from the analysis due to transient behaviour of the magnetisation, resulting from RF-induced depolarisation of the post-capillary signal from ^{129}Xe in the pulmonary veins [90]. In some cases, where SNR decreased at the end of the data acquisition period such that the final oscillations became indistinguishable from noise, the last second of data acquisition was also excluded.

We define the RBC signal (after correction for RF and T_1 decay) at position $r = (x, y, z)$ and time t as:

$$S(r,t) = \frac{\alpha(r)}{2} f\left(\omega t + \Phi(r)\right), \tag{7.2.1}$$

where $\alpha(r)$ is the spatially dependent peak-to-peak ^{129}Xe RBC signal amplitude and $f\left(\omega t + \Phi(r)\right)$ is a periodic function with angular frequency ω (heartbeat frequency) and spatially dependent phase $\Phi(r)$.

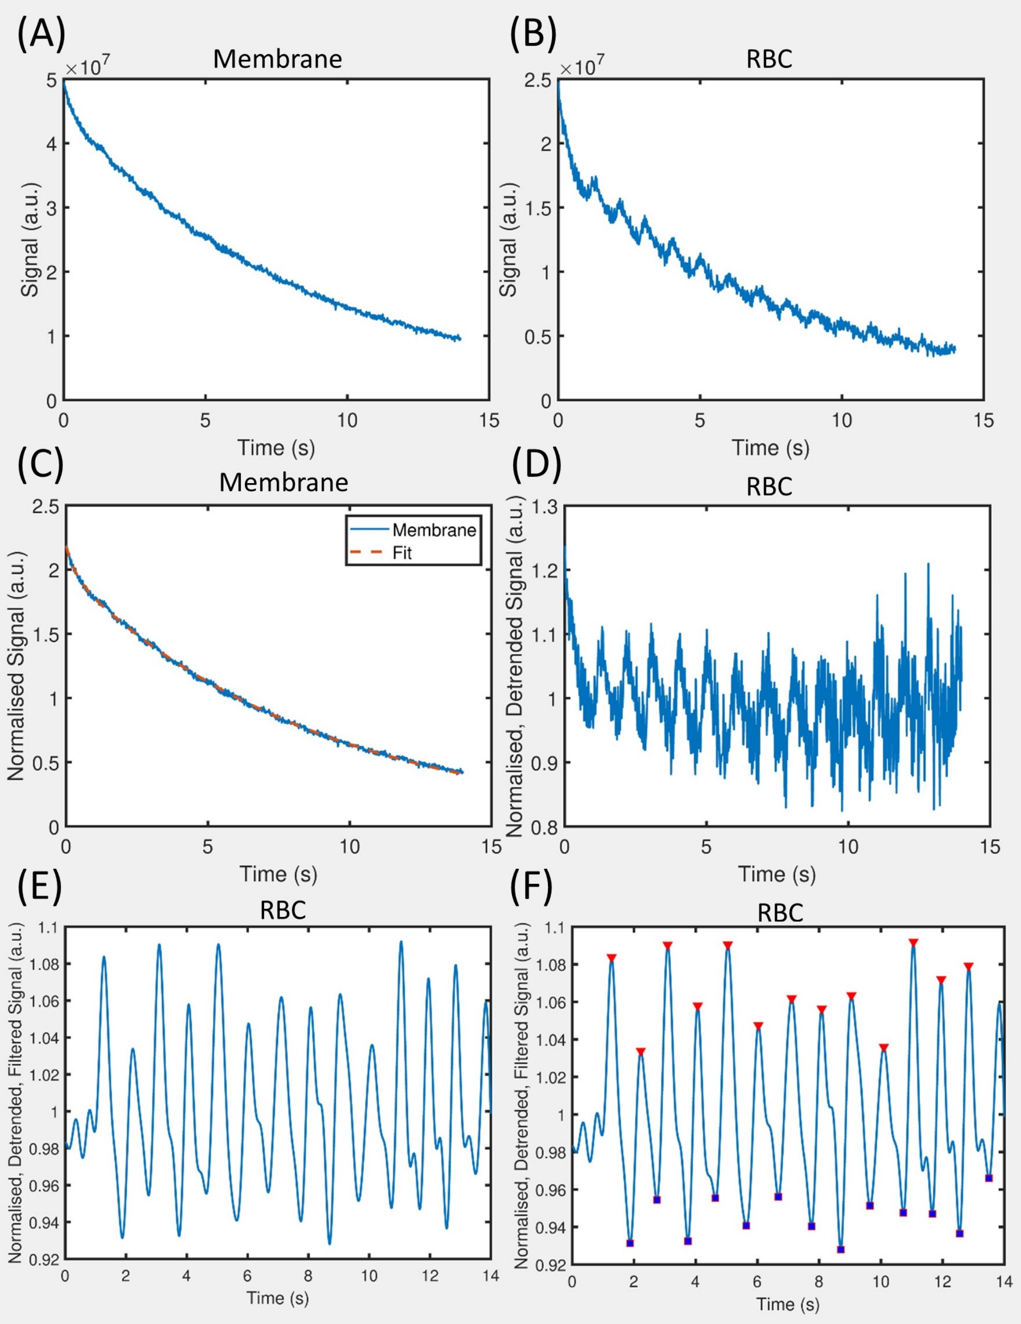

Figure 7.2.1: Method to isolate the ^{129}Xe RBC signal oscillations for oscillation mapping: the (A) membrane and (B) RBC k_0 signals were first normalised by their mean values. The membrane signal was fit to a biexponential decay function (C) and this was used to detrend the RBC signal (D). The detrended RBC signal was filtered (E) and a peak detection algorithm used to find the maxima and minima (F).

To map the RBC oscillations regionally, two keyhole reconstruction methods were implemented and compared:

7.2.3.1 Method 1: 'Two-Key' RBC Oscillation Mapping

RBC oscillation maps were calculated from the difference in signal amplitude in the images corresponding to low and high keyhole data, as described in [213]. After pre-processing with the steps described above, the k_0 projections adjacent to the minima and maxima were binned into low (purple circles in Figure 7.2.2A) and high (yellow circles in Figure 7.2.2A). The number of projections in each bin was approximately 20% of the total number of projections and a keyhole radius of six points was used, as shown in Figure 7.2.2B. These selection criteria were chosen to maximise the radius of the key and therefore the oscillation mapping fidelity, whilst minimising undersampling. The keys were then inserted separately into the high frequency 'keyhole' data using the last seven points from all spokes (blue in Figure 7.2.2B, C). To account for the fact that the k-space sampling was no longer uniform, an iterative numerical DCF was used [268]. Chemical shift separation and image reconstruction were then carried out for both the high and low keyhole data, resulting in two sets of ^{129}Xe images corresponding to the 'high' and 'low' RBC signal (S_{high} and S_{low}, respectively). The Two-Key RBC oscillation amplitude (α_{2-Key}) map was calculated from the pixelwise difference between S_{high} and S_{low}, divided by their mean and multiplied by 100%:

$$\alpha_{2-Key}(r) = \frac{S_{high}(r) - S_{low}(r)}{S_{high}(r) + S_{low}(r)} \text{x} 200\% \tag{7.2.2}$$

This normalisation ensures that: the oscillation amplitude is dimensionless, the amplitude at each pixel is normalised by its mean value and the maps are normalised for regional coil sensitivity that is intrinsic to all signals, $S(r)$, in both denominator and numerator.

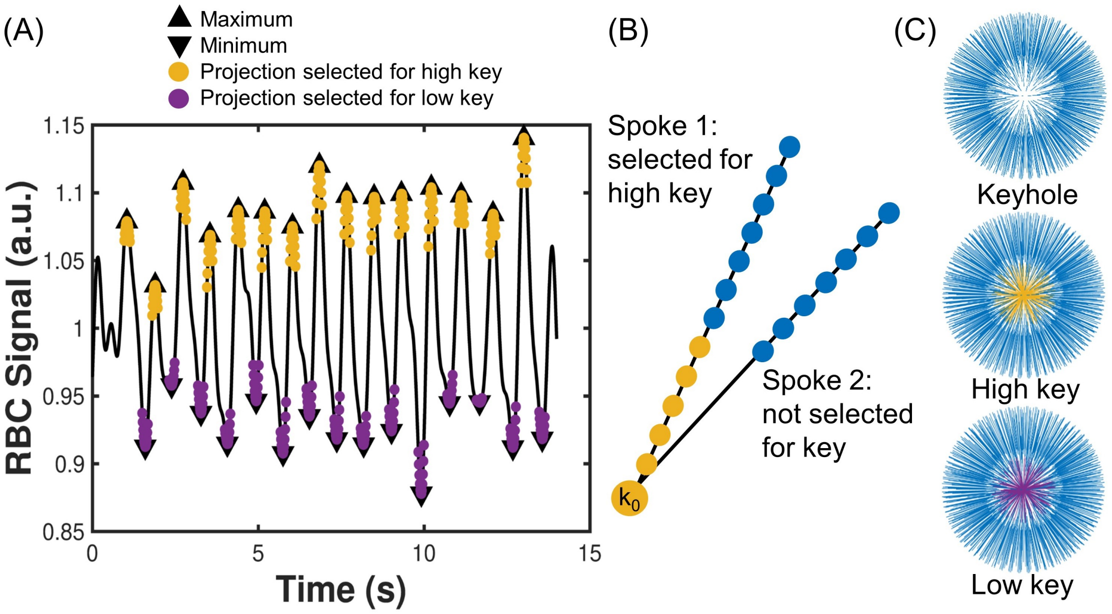

Figure 7.2.2: ‘Two-Key’ RBC oscillation mapping. (A) A peak detection algorithm is applied to the detrended, normalised, and filtered k_0 RBC signal (black line) to identify the maxima and minima (black triangles). Each point corresponds to the k_0 signal for a given spoke in the 3D radial trajectory. The projections adjacent to the extrema are selected for the high key (yellow) or low key (purple). (B) The high-frequency data (points seven to thirteen along the spoke, shown in blue) of every projection are included in the keyhole. Points one to six are only included for projections that have been selected for the key (yellow points). (C) The 3D radial k-space data for the keyhole alone and the keyhole plus the high key, and the low key.

7.2.3.2 Method 2: ‘Sliding Window’ RBC Oscillation Phase Mapping

To account for regional phase differences in the RBC signal oscillation, a SW keyhole reconstruction was implemented. The k_0 maxima were the same as with the Two-Key method and the neighbouring projections binned as before. Then, multiple keyhole reconstructions were performed, with the chosen projections stepped forward by one projection per reconstruction. Figure 7.2.3A shows the k_0 projections chosen for four of the keyhole reconstructions and Figure 7.2.3B shows the corresponding keyhole images for a central slice of the lung. The total number of keyhole reconstructions, N, was determined by the mean number of projections between adjacent maxima, such that one cardiac cycle was sampled. For example, for a heart rate of 70 bpm, the duration of each beat/cycle is ∼0.9 s and so, for TR = 15 ms, 60 projections are acquired in this time and therefore 60 keyhole reconstructions are performed. The RBC signal formed from each keyhole reconstruction at each voxel ($S_{r,j}$) was smoothed using a Gaussian-weighted moving average filter with a window length of 10. For each voxel, the maximum and minimum signals across the N reconstructed images, referred to as S_{max} and S_{min}, were found (Figure 7.2.3C). The SW oscillation amplitude (α_{SW}; Figure 7.2.3D) was then calculated from the difference between S_{max} and S_{min}, normalised by the mean signal of all N keyhole images (S_{mean}) and multiplied by 100%:

$$\alpha_{SW}(r) = \frac{S_{max}(r) - S_{min}(r)}{S_{mean}(r)} \text{x}100\% \tag{7.2.3}$$

The index, j, of the keyhole image where S_{max} was found, denoted j_{max}, provides information about the regional relative oscillation phase. This phase, denoted ϕ, is not equivalent to the absolute oscillation phase Φ, defined in Equation 7.2.1. Instead, it is defined as the difference between the local oscillation phase and the phase of the α_{k0} oscillation:

$$\phi(r) = \Phi(r) - \Phi_{k0} \tag{7.2.4}$$

If j_{max} = 1, then the oscillation is in phase with the global k_0 oscillation and ϕ = 0. $\phi = \pi$ signifies that the k_0 projections selected by the SW are π out of phase with the k_0 projections selected for the j = 1 keyhole image. To convert j_{max} to ϕ according to these definitions, the following piece-wise equation was used:

$$\phi = \begin{cases} \pi\frac{j_{max}-1}{j_{\pi}-1} & \text{if } j_{max} \leq j_{\pi} \\ \pi\frac{j_{max}-(\mathrm{N}-1)}{\mathrm{N}-j_{\pi}} & \text{if } j_{max} > j_{\pi} \end{cases} \tag{7.2.5}$$

By converting j_{max} to ϕ, the regional phase differences of the RBC signal oscillation can be quantified and visualised.

7.2.4 Image Analysis

Reconstructed images were masked by applying a noise threshold to the membrane signal images. Regions of interest were created to analyse RBC oscillation differences in the left, right, upper, lower, anterior, posterior, central, and peripheral lung. To define the central and peripheral lung, an ellipsoid was created at the centre of the lung mask. The radii of the ellipsoid was 0.79 times the original mask radii in each dimension, which was chosen such that the central and outer masks had volumes equal to half the total mask volume ($0.79 \sim 0.5^{1/3}$) [35]. The average α_{2-Key}, α_{SW} and ϕ within the whole lung mask and eight regional masks were calculated for each subject.

7.2.5 Statistical Analysis

Statistical analysis was performed using RStudio (version 2023.03.1, R version 4.3.0). Normality of variables was determined with Shapiro-Wilk normality tests. Correlations between α_{k0}, α_{2-Key}, α_{SW} and heart rate were assessed using the Pearson correlation coefficient for normally distributed variables and Spearman's correlation coefficient for non-normal variables. Differences between variables were tested for with paired Student's t-tests or Wilcoxon signed-rank tests. A significance level of $p < 0.05$ was used for all tests. To assess repeatability, the following metrics were used: bias (mean difference between two scans), % difference (mean absolute percentage difference between two scans) and CV (across all three scans).

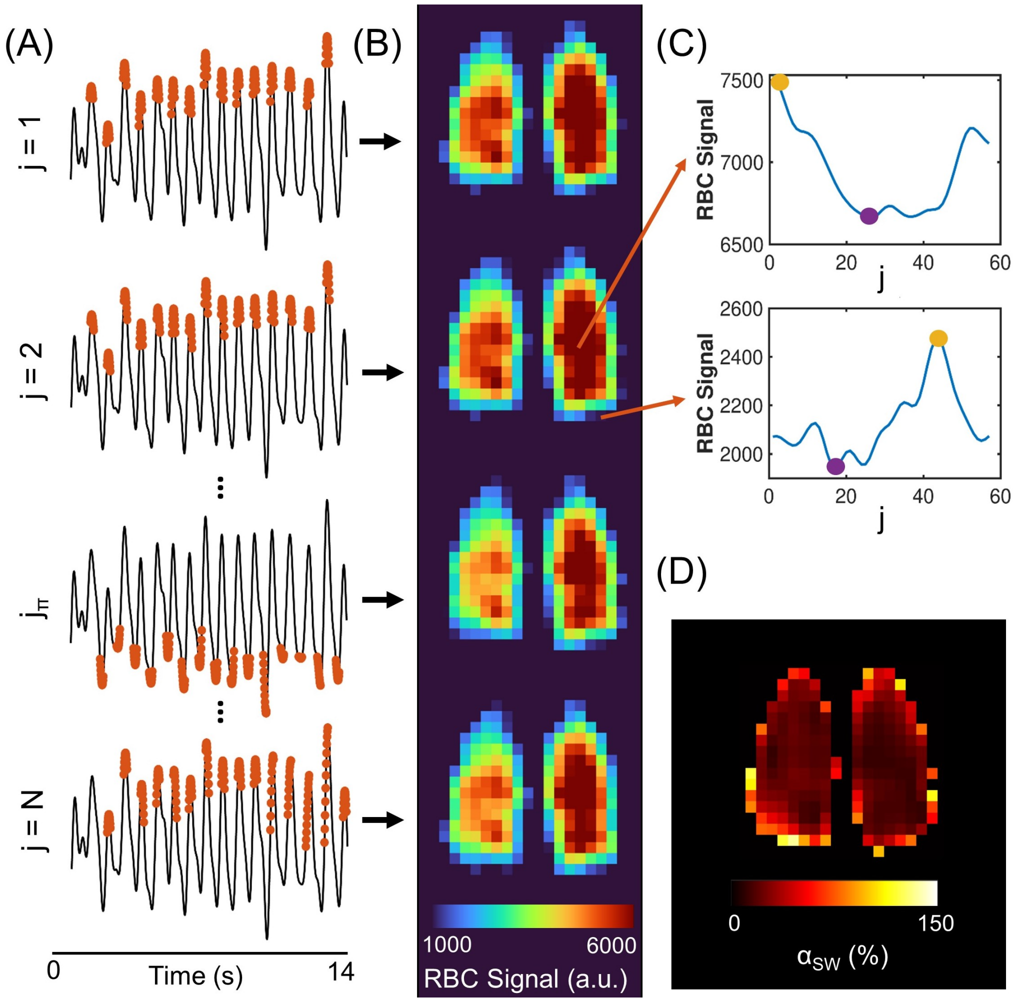

Figure 7.2.3: Overview of the ‘Sliding Window’ RBC oscillation phase mapping method. (A) k_0 RBC signal maxima are identified using a peak finding algorithm and the adjacent projections (red) are included in the first key (j = 1), which is combined with the high frequency k-space data in a keyhole reconstruction. The chosen projections are then stepped forward by one to form the second key (j = 2) and produce the second keyhole image, and then again for j = 3 to produce a third image, and so on. j_π corresponds to the key where the chosen projections are approximately in anti-phase with the projections from j = 1, i.e. distributed around the k_0 minima. j = N is the final key and should correspond to the projections from j = 1 shifted by one step to the left, however this is generally not the case because of the unequal number of projections per wave cycle. (B) Keyhole images are produced for each key. (C) For every pixel, the minimum (purple circle) and maximum (yellow circle) RBC signal values as a function of key number (j) are found. (D) The oscillation amplitude at each pixel (α_r) is found from the difference between the maximum and minimum RBC signal value across all keyhole images (S_{max} and S_{min}), normalised by the mean value across all keyhole images (S_{mean}) and multiplied by 100%. This results in a phase-corrected α_{SW} map.

7.3 Results

To ensure clear RBC k_0 oscillations, only subjects with an RBC image SNR of above 4.5 were selected for this work. Cases with an SNR of between 4.5 and 5.5 (i.e. on the borderline of the Rose criterion [269]) were visually inspected prior to filtering for discernible RBC k_0 oscillations before inclusion. This cut-off was chosen empirically based on preliminary analysis of the relationship between RBC SNR and the percentage difference between α_{k0} and α_{2-Key}. As a result, 28/42 healthy datasets were included in this work. For each subject, keyhole RBC images were reconstructed for low and high RBC k_0 signal and for each key of the SW method with minimal undersampling; across all subjects the median of the maximum value of the iterative DCF was 1.46. The total keyhole mapping reconstruction and processing time was ∼5 minutes for both methods on a Windows PC with an Intel® Core™ i7-4790 processor.

The SW keyhole RBC signal revealed regional variations in oscillation phase as well as amplitude, whereas the Two-Key method, by definition, only provides information on oscillation amplitude. Example maps for three healthy volunteers are shown in Figure 7.3.1. Table 7.3.1 shows the subject demographics and a summary of the RBC oscillation mapping results. On an individual subject basis, the α_{2-Key} maps were normally distributed, hence the mean was used, whereas the α_{SW} and ϕ maps were not normally distributed, so the median value from the map was used. Table 7.3.1 summarises the inter-subject means/medians.

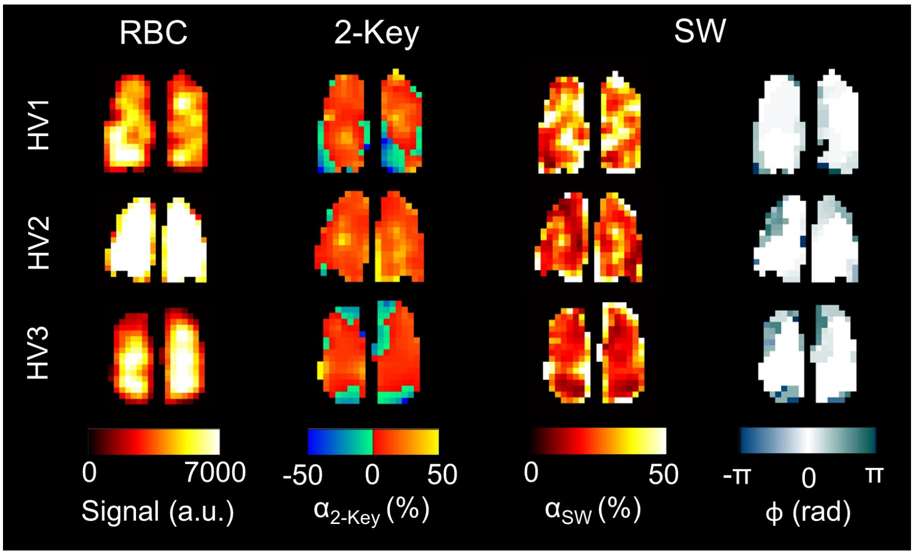

Figure 7.3.1: Maps of RBC signal (reconstructed from all k-space data), α_{2-Key}, α_{SW} and ϕ for three healthy volunteers (HV). Colour map limits have been set to aid visualisation.

All of the α_{2-Key} maps exhibited some regions of negative oscillation amplitude (RBC signal greater in the low-key image than the high-key image). These areas were positive in the α_{SW} maps and correspondingly, CV_{SW} was significantly lower than $CV_{2\text{-Key}}$ ($p < 0.001$). The distribution of the healthy volunteer α_{2-Key}, α_{SW} and ϕ values across all pixels and subjects are shown in Figure 7.3.2. α_{2-Key} was normally-distributed, with a mean value of 14%, but α_{SW} and ϕ were non-normally distributed, with distribution parameters shown in Table 7.3.2.

Table 7.3.1: Subject demographics and inter-subject α mapping results. Normally distributed variables are given as mean ± standard deviation and non-normally distributed variables are given as median (range).

n (female)	28 (12)
Age (years)	38.8 ± 11.1
RBC:Gas	0.0040 ± 0.0011
RBC:M	0.41 (0.31 - 0.58)
α_{k0}	15 (10 - 27)
α_{2Key}	14 ± 3
α_{SW}	29 ± 3
ϕ	0.27 ± 0.19
$CV_{2\text{-Key}}$	1.4 ± 0.3
CV_{SW}	0.66 ± 0.06
CV_{ϕ}	0.85 ± 0.05

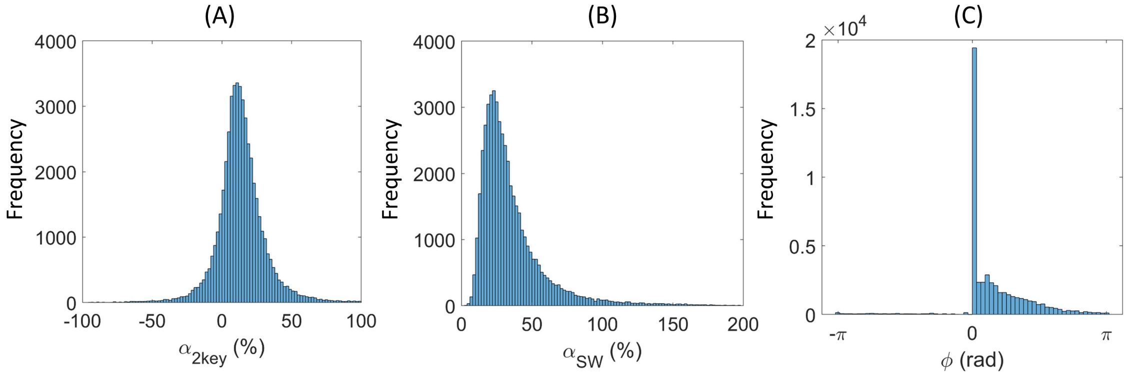

Figure 7.3.2: Histograms of (A) α_{2-Key}, (B) α_{SW} and (C) ϕ values across all lung pixels for all 28 healthy volunteers.

Table 7.3.2: Distribution characteristics for the maps of α_{2-Key}, α_{SW} and ϕ in healthy volunteers.

	α_{2-Key} (%)	α_{SW} (%)	ϕ (rad)
Mean	14	36	0.48
SD	29	24	0.82
Median	13	29	0.27
[Min, Max]	[-178, 183]	[2, 252]	[-3.14, 3.14]
IQR	13	21	0.87

The means of the α_{2-Key} maps showed significant correlation with the mean α_{k0} values (Spearman's ρ = 0.60, p = 0.001), but they were significantly different ($p = 2 \times 10^{-4}$). Median α_{SW} was significantly correlated with both α_{k0} (ρ = 0.50, p = 0.008) and the mean α_{2-Key} (Pearson's r = 0.40, p = 0.04). The correlation between α_{SW} and α_{2-Key} was evaluated regionally, by comparing the amplitude values from each method for each pixel within the lung mask. Plotting the pixelwise correlation resulted in v-shaped graphs with an inflection point at 0, for example in Figure 7.3.3A, because negative values of α_{2-Key} correspond to positive α_{SW} values. Taking the absolute values of α_{2-Key} resulted in a positive linear relationship (Figure 7.3.3B). Subjects with higher ϕ values had a weaker correlation between α_{2-Key} and α_{SW}, possibly because they had an increased number of voxels with $\alpha_{2-Key} < 0$, and the Spearman's ρ of the α_{SW}-α_{2-Key} correlation decreased with increasing root mean square (RMS) ϕ (Figure 7.3.3C). α_{SW} and ϕ were both significantly negatively correlated with estimated heart rate (ρ = -0.43, p = 0.02 and ρ = -0.47, p = 0.01, respectively), but there was no significant correlation between heart rate and α_{k0} or α_{2-Key}. The relationships between heart rate and the oscillation mapping metrics are summarised in Figure 7.3.4. No dependence on age was found for α_{k0} (ρ = 0.25, p = 0.20), α_{2-Key} (ρ = 0.14, p = 0.47), α_{SW} (ρ = 0.18, p = 0.37) or ϕ (ρ = -0.10, p = 0.61). No significant differences in oscillation amplitude or phase were found between males and females.

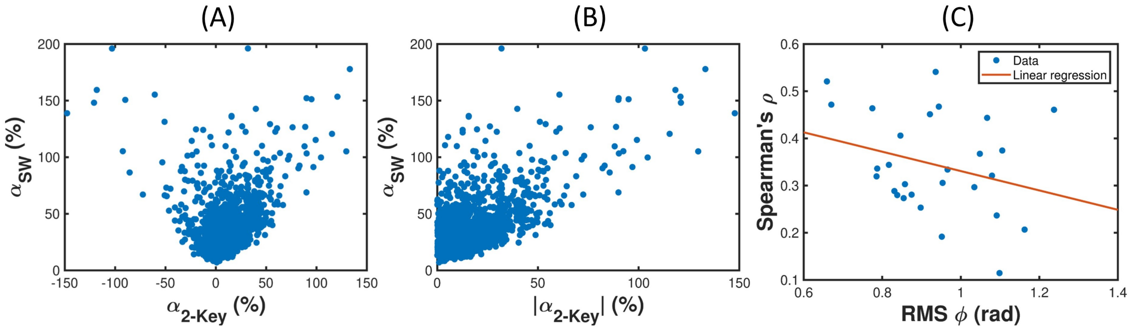

Figure 7.3.3: Pixelwise correlation of (A) α_{2-Key} and α_{SW} and (B) absolute α_{2-Key} and α_{SW} for a healthy participant, who had RMS ϕ = 0.82 rad. (C) Spearman's correlation coefficient for the pixelwise correlation between α_{2-Key} and α_{SW} for all healthy participants, plotted against RMS ϕ.

Several significant regional trends were identified for α_{SW}, which are shown in Figure 7.3.5A. In summary, α_{SW} was observed to increase from the upper to lower, left to right, posterior to anterior and central to peripheral lung. ϕ was significantly greater in the peripheral lung than the central lung region, but no other significant regional trends were observed (Figure 7.3.5B). Same-session and same-day repeatability of α_{k0}, mean α_{2-Key}, median α_{SW} and median ϕ were investigated in eight healthy subjects. The RBC oscillation and phase maps across the three scans for one of the healthy volunteers are shown in Figure 7.3.6. The results are summarised in Table 7.3.3 and Bland-Altman plots for α_{k0} and α_{SW} are shown in Figure 7.3.7. One of the intra-session data sets had to be discarded due to a technical fault at the scanner. The intra-session bias was smallest for α_{2-Key}, compared to α_{k0} and α_{SW}, whereas the inter-session bias was smallest for α_{k0}. To account for the increased value of α_{SW} compared to α_{k0} and α_{2-Key}, the mean absolute percentage difference between α values was calculated. This was smallest for α_{SW} both between sessions and within the same session, with an average variation of less than 10%. Mean CV across all three scans was also smallest for α_{SW} (0.07 ± 0.04).

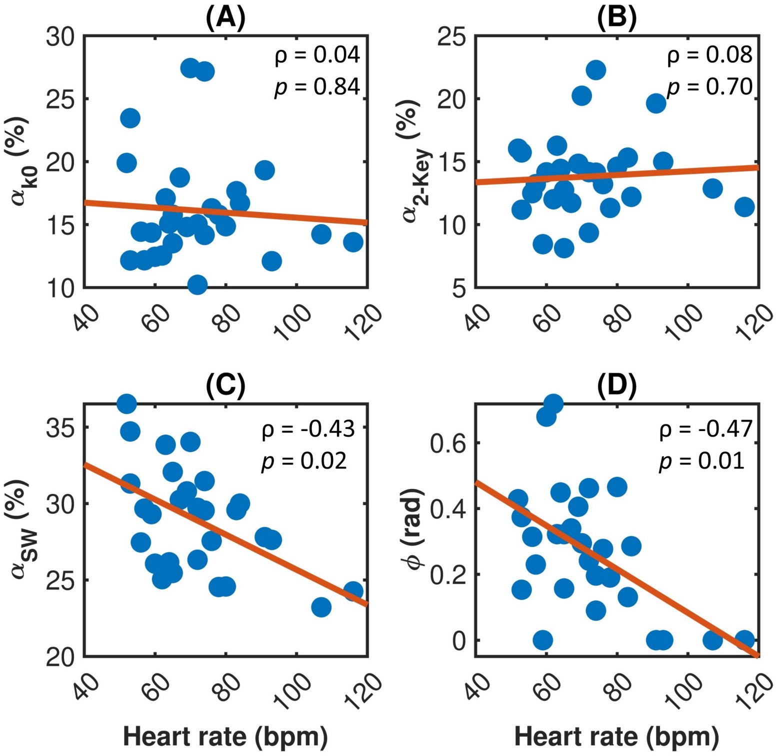

Figure 7.3.4: No significant dependence on heart rate was found for (A) α_{k0} or (B) α_{2-Key}. A significant negative correlation was found between heart rate and (C) α_{SW} and (D) ϕ. Spearman's ρ values and p-values are given in the top right of each plot.

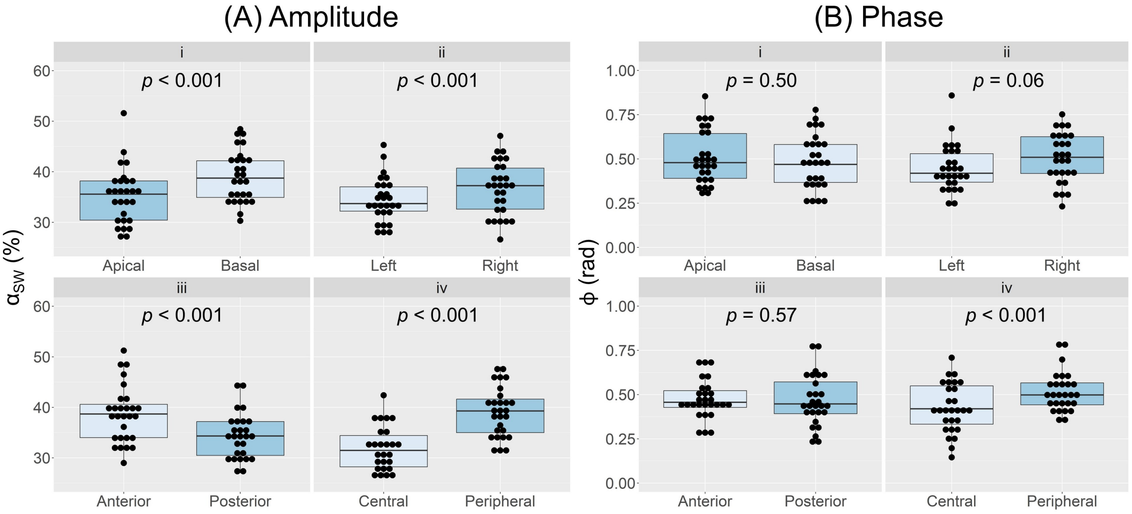

Figure 7.3.5: Comparison of (A) α_{SW} and (B) ϕ between the (i) apical/basal, (ii) left/right, (iii) anterior/posterior and (iv) central/peripheral regions of the lung, with the corresponding p-values from paired t-tests and Wilcoxon signed-rank tests.

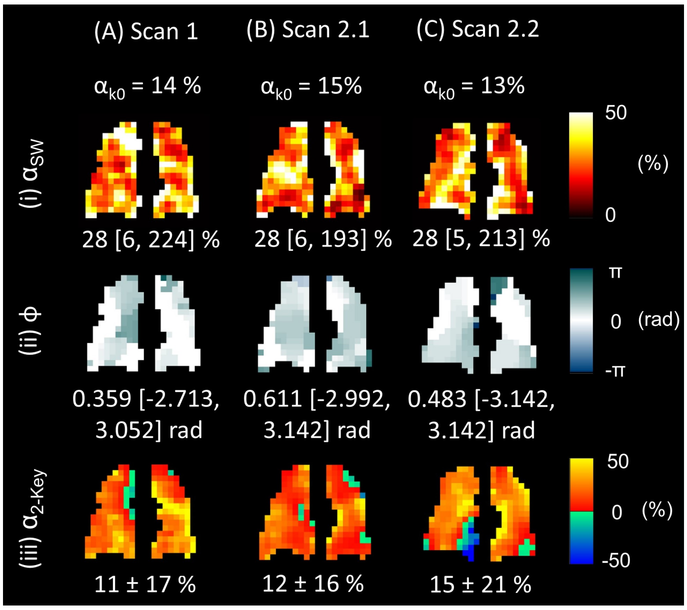

Figure 7.3.6: The amplitude and phase maps for a 41 year old male volunteer acquired at different time points: (A) one scan in the morning and (B - C) two scans in the afternoon, approximately five minutes apart.

Table 7.3.3: Intra- and inter-session repeatability of α_{k0}, mean α_{2-Key}, median α_{SW} and RMS ϕ. RMS ϕ was used instead of median ϕ to avoid zero division problems in the calculation of percentage difference. Bias = mean difference, LOA = limits of agreement/95% confidence interval, %Diff = mean absolute percentage difference, CV = coefficient of variation across all three scans. For the subject where one scan failed, data from a previous scan from three months prior was included to calculate 3-scan CV.

	Intra-session		Inter-session		
	Bias [LOA] (%)	%Diff (%)	Bias [LOA] (%)	%Diff (%)	(3-scan) CV
α_{k0}	-1.09 [-5.37, 3.81]	13.3 ± 11.8	0.07 [-4.53, 4.66]	10.9 ± 10.9	0.11 ± 0.04
α_{2-Key}	-0.14 [-4.09, 3.82]	11.8 ± 9.1	-0.48 [-6.56, 5.60]	22.8 ± 12.2	0.16 ± 0.05
α_{SW}	-0.68 [-5.26, 3.91]	5.9 ± 8.2	-0.65 [-6.24, 4.95]	8.6 ± 6.3	0.07 ± 0.04
ϕ	-0.04 [-0.53, 0.45]	21.4 ± 21.3	0.10 [-0.27, 0.46]	15.9 ± 10.6	0.17 ± 0.09

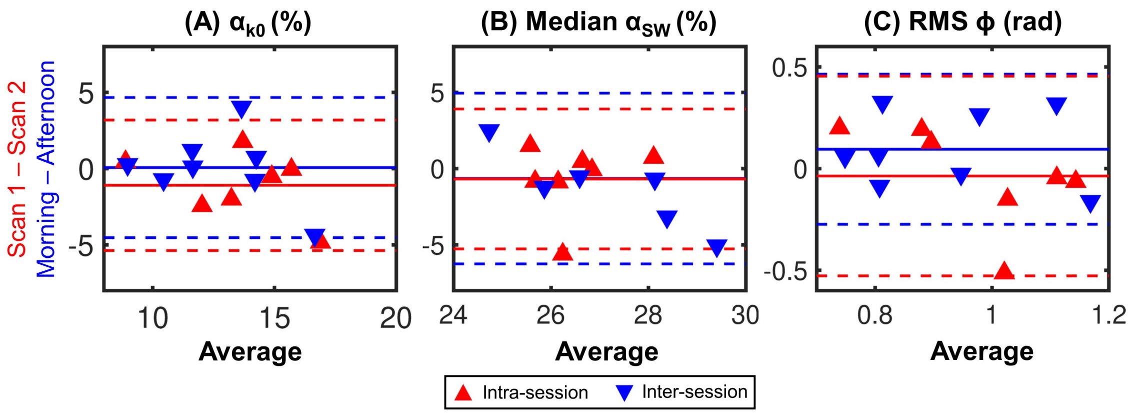

Figure 7.3.7: Bland-Altman plots of the intra- and inter-session repeatability of (A) α_{k0}, (B) median α_{SW} and (C) RMS ϕ.

7.4 Discussion

Oscillations of the ^{129}Xe RBC signal originating from changes in the capillary blood volume over the cardiac cycle can be spatially resolved retrospectively from dissolved ^{129}Xe spectroscopic imaging using keyhole reconstruction. Adapting Niedbalski's method [213] for our MESI acquisition and data acquired therewith, we found a mean α_{2-Key} of (14 ± 3)% in 28 healthy subjects. Mean α_{2-Key} was correlated with mean α_{k0}, although the values tended to be smaller. This difference may be because the low- and high-key images were generated using approximately six projections either side of the signal extrema, whereas α_{k0} was calculated as the peak-to-peak amplitude.

We further adapted the previously published method to estimate and correct for regional phase differences of the RBC oscillation. By adopting a SW approach to select the k_0 projections and repeating the keyhole reconstruction for each selection, we were able to resolve the pixelwise RBC signal evolution with time. The oscillation amplitudes were calculated individually for each pixel, without the assumption that the oscillation is in phase with the whole lung k_0 oscillation. Using this method, areas of physiologically unrealistic negative oscillation amplitude in the Two-Key maps became positive in the SW maps due to their regional phase correction. By converting the (key) index of the keyhole from which the first maximum originated for each pixel (j_{max}) into phase, it was possible to create oscillation phase maps. Phase differences relative to the k_0 oscillation occurred mostly in the peripheral lung and, furthermore, these regions qualitatively correlated with areas of negative oscillation amplitude from the Two-Key oscillation maps. This corroborates our hypothesis that negative oscillation amplitudes produced in the Two-Key method are caused by phase differences. Phase differences are thought to originate from effects of the cardiac pulse wave, which will reach regions of the capillary bed at different times due to different distances from the heart. For a typical heart rate of 70 bpm, the maximum ϕ value of $\pm\pi$ corresponds to a delay of ∼400 ms, which is of the order of the whole-lung average conduction time of the cardiac pulse from the pulmonary valve to the capillary bed (120 - 180 ms) [270, 271]. Other regional variations in phase may result from cardiac pulse wave reflections due to impedance mismatch at bifurcations, ineffective vascular coupling related to focal lung disease or PVD, or variations in blood flow velocity with vessel narrowing or change in vessel wall stiffness or compliance.

The replacement of negative oscillation values with positive values in the SW maps explains why the average regional α_{SW} was higher than α_{k0} or α_{2-Key}. In addition, the distribution of α_{SW} values was not normally distributed and was found to be positively skewed. The range of SW oscillation values was large and some pixels at the periphery of the lung had oscillations of $> 200\%$, which could be because of blurring or partial volume effects. These pixels tended to have very small RBC signal values (used as the denominator in the normalisation process) and it is possible that the corresponding large oscillation values were partly due to noise. Across all subjects, median α_{SW} was significantly correlated with both α_{k0} and mean α_{2-Key}. The voxel-wise correlation between α_{SW} and α_{2-Key} showed higher ρ values for lungs with fewer RBC oscillation phase differences (lower RMS ϕ). This is because, whilst α_{SW} is calculated independently of phase, α_{2-Key} decreases when the local oscillation is out of phase with the whole-lung k_0 oscillation.

Several regional trends were observed in α_{SW} in healthy volunteers. The anterior-posterior, centre-peripheral, and left-right gradients reflect the reverse of the RBC signal trends, which is explained by the pixelwise normalisation used to calculate α_{SW}. No significant difference was found between the RBC signal at the base and the apex of the lung; however, a significant decrease was found in α_{SW} from base to apex. Ventilation at the base of the lungs is increased in the supine position [272]; in healthy volunteers with associated V/Q matching, there may also be an increase in perfusion which may explain the reduced α_{SW}. To directly compare with the results of Niedbalski et al. [213] we also evaluated the regional trends of α_{2-Key} when a normalisation by the whole-lung mean RBC signal was used. A significant increase was observed from the anterior to the posterior of the lung due to gravitational effects, similar to that work. No significant changes were found between the base and apex, or the left and right lung, but a significant decrease was found between the central and peripheral lung, also in agreement with Ref. [213].

α_{SW} exhibited a moderate (ρ = -0.43) significant (p = 0.02) negative correlation with the (RBC-signal derived) heart rate for the healthy volunteers. This may be explained by considering the effect of the heart rate on pulmonary blood flow. With increased heart rate, the heart spends relatively less time in diastole per beat if the stroke volume, PVR and compliance remain the same [273]. A higher blood flow is maintained throughout the cardiac cycle and blood flow pulsatility is reduced. This effect is propagated to the pulmonary capillaries and so the relative change in capillary blood volume, and hence α_{SW}, is decreased. Significant correlations were not seen with heart rate and either α_{k0} or α_{2-Key}. As far as we are aware, a relationship between ^{129}Xe RBC oscillation amplitude and heart rate has not been reported previously; this warrants further investigation in healthy volunteers and patient groups.

No dependence on participant age was found for any of the oscillation mapping metrics, however, the number of older volunteers included in this work was limited, with only 5/28 subjects over the age of 50. The inclusion of increased numbers of older healthy volunteers may be necessary to fully evaluate the age relationship. Costa et al. used Niedbalski's oscillation mapping method and 1-point Dixon method data to evaluate the age-dependence of the oscillation amplitude in healthy volunteers aged 22 to 73 [274]. They did not find a correlation between amplitude and age, which also follows the observations of Niedbalski [213] and Lu [215]. Collier et al. found no dependence of α_{k0} on age [189], in line with our findings, but Mummy et al. reported an increase in α_{k0} with age [230]. We did not find any sex differences in the oscillation amplitudes and phase, but Costa et al. found an increased mean oscillation amplitude in healthy females when compared with healthy males, using Niedbalski's oscillation mapping method. Females tend to have lower levels of haemoglobin than males [275], and hence lower ^{129}Xe RBC:M ratios [259]. With our sliding window method and pixelwise normalisation, the oscillation amplitudes were calculated relative to their 'equilibrium' value and so any baseline differences in RBC signal from reduced haemoglobin concentration in females are unlikely to have affected the oscillation amplitude. In contrast, Costa et al. normalised the oscillation amplitudes by the mean whole-lung RBC signal, therefore a lower mean signal in females from sex-related differences in haemoglobin levels may be responsible for the higher oscillation amplitude observed in this group.

In a subgroup of eight healthy volunteers, we demonstrated that RBC oscillation phase mapping is repeatable between scans, with a smaller mean intra-session bias (-0.68%) for α_{SW} than that of α_{k0} (1.09%). The α_{SW} bias between separate examinations on the same day (-0.65%) was similar to the intra-session bias but greater than that of α_{k0} and α_{2Key}. However, α_{SW} is larger on average than α_{k0} and α_{2-Key} and when the mean absolute percentage difference and CV across all three scans were compared, these were smallest for α_{SW} (8.6% and 0.07). This may be because α_{SW} is independent of phase, so one source of variation that might occur between scans was removed. The LOA found for α_{SW} and ϕ in the Bland-Altman analysis were quite large, which might indicate that the precision of our method is low. The LOA for α_{SW} ($\sim \pm$ 5%) were similar to those of α_{k0}, but, given the greater values of α_{SW}, were relatively smaller (LOA $\sim$ 17% of the mean α_{SW} value, compared to LOA $\sim$ 33% of the mean α_{k0} value). Testing with a larger group of subjects - including patients with a wider range of pulmonary diseases - is required to fully assess the repeatability and sensitivity of this method.

7.4.1 Comparison with Previous Work

Our value of α_{2-Key} = (14 ± 3)% in healthy volunteers was higher than the value of 8.7% found by Niedbalski et al. [213], although the mean α_{k0} was also higher for our subjects; 15% when compared to 10%. The difference in α_{k0} may be because in Ref. [213], a sinusoidal fit was used to calculate the amplitude whereas we used a peak detection algorithm, due to the underlying signal being non-sinusoidal in shape [210]. The latter method was recently shown to return higher α values using the same acquisition method as Niedbalski [276]. In addition, the mean α_{2-Key} is likely higher due to differences in the normalisation methods used. We used a pixelwise normalisation whereby the pixelwise difference between S_{high} and S_{low} were normalised by their mean value, whereas Niedbalski [213] normalised by the whole-lung mean of the fully-sampled RBC image. We chose our approach for two reasons: firstly, the RBC signal itself is spatially varying and subject to regional bias from B_1 inhomogeneity, therefore a regional normalisation is more appropriate to distinguish trends in oscillation amplitude and phase from the variations in mean RBC signal itself; secondly, the fully-sampled RBC image includes all acquired spokes and thus has higher signal compared with the keyhole images, therefore normalisation by the fully-sampled image results in lower oscillation amplitudes. In other words, we chose the normalisation image to undergo reconstruction with the same keyhole sampling pattern as the high and low images so that the SNR was comparable.

Lu et al. found an α_{2-Key} distribution with a mean of 5 ± 3% across 17 healthy participants [215]. Although the mean value of α_{k0} for these subjects is not mentioned, in a subsequent work by the same authors [277] mean α_{k0} was ∼15% whilst mean α_{2-Key} was ∼6% for 21 healthy subjects. Therefore, the cause of this lower α_{2-Key} may be related to the authors' keyhole reconstruction methods rather than the participants studied, or preprocessing steps. As in our work, Lu et al. used a regional normalisation for the α_{2-Key} maps, but they divided the difference between the keyhole images by the fully-sampled RBC image. As noted in the previous paragraph and demonstrated in Figure 7.4.1, using the fully-sampled image rather than a keyhole sampled image can lead to lower oscillation amplitude values. Another reason for this low mean α_{2-Key} might be the choice of key radius. Here, we followed Niedbalski's approach and chose a key radius that led to 50% undersampling at the key edge, whereas Lu et al. used a smaller key radius to minimise undersampling further. Based on digital phantom simulations, which showed that a key radius of 0.14 times the maximum k-space radius was optimal for resolving oscillation defects and minimising heterogeneity, they used 9/64 points from each spoke instead of the 12/64 points which would correspond to Niedbalski's method. Using a reduced key radius reduces the sensitivity to differences in the RBC signal between the high and low key images, because more of image contrast remains the same between the high and low key k-space patterns. To illustrate this, we performed our Two-Key reconstruction with a reduced key radius for one healthy volunteer and found that the mean α_{2-Key} decreased with decreasing key radius (Figure 7.4.2).

The choice of key radius is a balance between maximising the sensitivity to signal oscillations and minimising the radial undersampling - it is not clear which of these should be prioritised. With Lu's approach, although the healthy volunteer α_{2-Key} values were reduced, significant differences were still identified between the volunteers and patients with CTEPH. With less undersampling, Lu's approach may be more reliable, however radial sampling is already robust to undersampling and we have not seen any evidence of streaking artefacts in the low and high key images. A benefit of Niedbalski's choice of key radius is that there is better agreement between mean α_{2-Key} and α_{k0}, which is a more well-established variable and has been reported in several studies and patient groups [169, 35, 37, 137, 209].

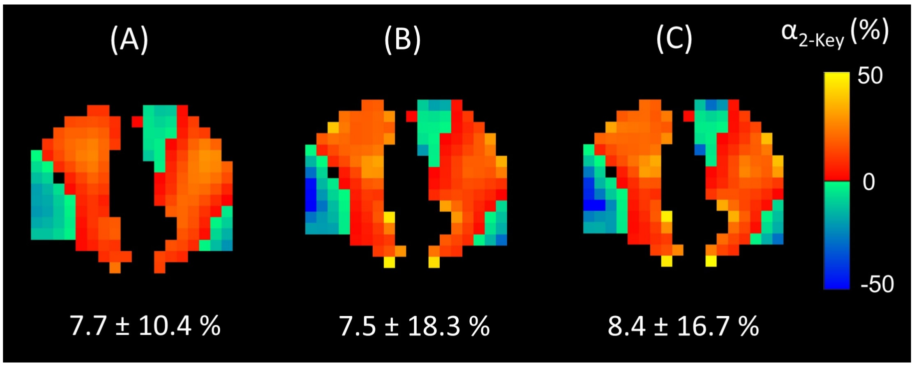

Figure 7.4.1: Comparison of the effect of normalisation choice on the α_{2-Key} maps (shown for one lung slice) and the mean value (across all slices) for a healthy volunteer. The difference between the high key and low key reconstructed signal was normalised by (A) the mean of the fully-sampled RBC signal, as in [213] (B) the fully-sampled RBC signal, as in [215] and (C) the mean of the low and high key RBC signal, as in this work.

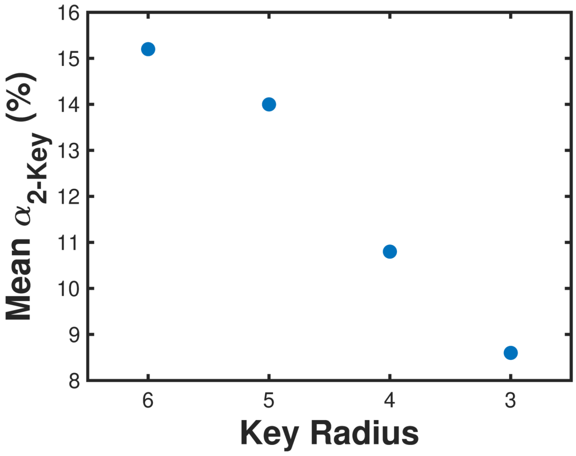

Figure 7.4.2: The value of mean α_{2-Key} for one healthy volunteer when the radius of the high and low keys was decreased from six points (as in this work) to three points.

Both Niedbalski and Lu observed regions of negative amplitude in their α_{2-Key} maps, which were treated as 'less healthy' than positive values in the linear binning of the healthy volunteer distribution. For example, a voxel with α_{2-Key} = -10% is 'worse' than α_{2-Key} = 0%. However, we observed phase differences and negative α_{2-Key} values in all of the healthy participants studied, suggesting that the phase differences are not necessarily always associated with pathology. ϕ was significantly higher in the periphery of the lungs than the centre, suggesting that the pulse wave becomes out of phase with the wave in the central lung due to the path difference. Hence, treating areas of the lung with negative α_{2-Key} as 'unhealthy' may lead to incorrectly considering areas of the lung to be diseased.

7.4.2 Limitations

The primary limitation of our method is the requirement for high SNR (>4.5) RBC images and clear cardiogenic RBC signal oscillations. This may limit the application of the method to lung disease patients who struggle to inhale the full ^{129}Xe dose or complete the 14 s breath hold due to their symptoms, or who might have inherently lower RBC signal due to reduced gas transfer. The minimum SNR condition was not met in a high proportion of the healthy volunteer data inspected, which was primarily due to a temporary dip in polariser/RF coil performance.

Another limitation is that the frequency and waveform of the RBC k_0 oscillation were not constant over the breath hold, because of SNR and heart rate changes. As j approaches N, different sections of each cardiac cycle waveform are sampled because the projections selected by the SW become ‘out of step’ with each other. This effect can be seen in the projections chosen for the j = N keyhole in Figure 7.2.3A. Decreasing the TR of the imaging sequence, whilst maintaining the $TR_{90°,equiv}$ with a corresponding change in flip angle [223], may help to mitigate this effect by reducing the duration of the breath hold. A shorter TR would also be advantageous for increasing the temporal resolution of the k_0 signal, but this might be unattainable for our four-echo sequence.

Finally, it is challenging to validate our RBC oscillation phase mapping method because there is a lack of well-established methods imaging methods to quantify the function of the pulmonary microvasculature. In Chapter 8, some initial attempts to validate our method are made.

7.5 Conclusions

Cardiogenic oscillations of the ^{129}Xe RBC signal can be mapped retrospectively from 3D radial MESI using a keyhole reconstruction scheme. This work builds upon previous methodology by sampling regional phase differences in RBC oscillation using a SW keyhole reconstruction. This approach allows for the oscillation phase to be regionally estimated, which may provide a means to detect the effects of the cardiac pulse wave in the pulmonary microvasculature and its alteration in cardiopulmonary disease. The focus of the present chapter was the methodology of the RBC oscillation phase mapping technique; in the next chapter we apply our method to patients with lung disease and compare the resulting oscillation maps with other modalities such as quantitative DCE-MRI.

Chapter 8

Preliminary Evaluation of ^{129}Xe Red Blood Cell Signal Oscillation Amplitude and Phase Mapping in Patients with Lung Disease

8.1 Introduction

The microvasculature plays a crucial role in the pulmonary circulation and is often damaged in lung disease, but is challenging to image directly. Vessel destruction, microthrombi, microemboli, emphysema, and fibrosis can increase morbidity in pulmonary disorders and eventually lead to progressive right heart failure and death [278]. For example, patients with end-stage IPF often develop PH and an elevated pulmonary arterial pressure [279]. In COPD, PH is frequently concomitant and blood flow to the microvasculature has been shown to be reduced by up to 52% in severe COPD [261]. A significant risk of thrombotic conditions has been associated with COVID-19, including microvascular thrombosis [280], and histologic analysis of patients who died from COVID-19-associated respiratory failure showed widespread alveolar capillary microthrombi [281]. Hyperpolarised ^{129}Xe MRI has showed a reduction of RBC:M in post-COVID-19 patients, which could be explained by micro embolic/thrombolic obstruction of capillaries [282]. In cases of severe COVID-19 infection, long-term lung damage has been observed, including fibrotic and non-fibrotic abnormalities and alveolar damage as a result of viral pneumonia [190].

Persistent dyspnoea is one of the most commonly reported pulmonary symptoms of ‘long-Covid’, a chronic condition affecting one or more organ systems that occurs after COVID-19 infection and is present for at least three months [283]. The mechanisms responsible for long-Covid breathlessness are not well-understood and clinical tests often return normal values for non-hospitalised long-Covid patients [282]. Gas exchange abnormalities have been reported in these patients, although the evidence is limited. In a study published by Grist et al., non-hospitalised long-Covid patients were found to have abnormally low RBC:M compared to healthy participants at more than six months after infection [194].

The RBC:M ratio is affected by both capillary perfusion and the diffusion of ^{129}Xe across the alveolar membrane, so this reduction may be due to reduced blood volume resulting from microvascular abnormality, thickening of the alveolar membrane, or a combination of both. There has been some indication of microvascular damage in long-Covid, for example, in a study of a small group of post-COVID-19 patients, pulmonary microvascular perfusion was significantly reduced compared to a group of healthy volunteers of a similar age, despite 90% of patients having normal TL_{CO} values [284].

Quantifying the extent of microvascular damage is particularly important in CTEPH. This disease is a rare complication of pulmonary embolism, which has an incidence of 0.9 - 39 patients per million [285]. CTEPH is characterised by occlusions in the pulmonary arteries following pulmonary embolism, along with small vessel disease [286]. It can be fatal if left untreated but is potentially curable through PEA, which is a complex surgery to remove thromboembolic material from the proximal vessels [287]. Patient operability is decided by a multidisciplinary team, and considers multiple factors, such as clot location, haemodynamic measurements and co-morbidities [287]. Obstructive lesions in the main, lobar or segmental arteries are more surgically accessible and easier to remove than distal thrombi. Widespread distal disease may render the patient unsuitable for surgery, or be responsible for persistent PH despite successful surgery. Up to 50% of post-PEA patients experience residual PH, which is associated with increased morbidity and mortality [288, 289, 290]. Assessment of the extent of microvascular disease is typically based on measurement of PVR, but this does not differentiate between vessel occlusion and small vessel disease. Clinical decision-making and prediction of postoperative outcomes could be improved by methods to measure microvascular damage directly, which are currently limited.

The hyperpolarised ^{129}Xe RBC oscillation amplitude provides insight to pulmonary haemodynamics at the capillary level, and oscillation maps obtained using keyhole reconstruction may help to identify regional microvascular abnormalities. Niedbalski et al. found a higher mean oscillation amplitude in 12 IPF patients and a reduced mean oscillation amplitude in 11 PAH patients compared to healthy volunteers, using their implementation of Two-Key oscillation mapping [213]. These results are in accordance with previous spectroscopy results [37]. Oscillation mapping was able to identify areas of low oscillation amplitude in CTEPH patients, which improved after PEA surgery, although there was no statistically significant change in the mean oscillation amplitude [215]. The CTEPH patients were found to have lower mean oscillation amplitude than healthy volunteers. Although there is a lack of literature on RBC oscillation mapping in post-COVID-19 patients, this method could be useful for identifying fibrotic and microvascular damage which is difficult to detect otherwise.

In Chapter 7 we adapted Niedbalski's Two-Key oscillation mapping method for our dissolved ^{129}Xe MESI sequence and extended it to map oscillation phase in addition to amplitude. Our sliding window keyhole reconstruction revealed regional phase differences in the RBC oscillations which were not captured when performing Two-Key keyhole reconstruction. This regional phase information may reflect the haemodynamic effect of the cardiac pulse wave in the pulmonary microvasculature. Alterations in vessel stiffness and compliance, remodelling and the presence of blood clots all affect pulse wave transmission and blood flow pulsatility, which may be detectable through RBC oscillation phase mapping. However, further work is needed to validate our phase mapping technique and to benchmark RBC oscillation phase against established methods.

One such method that might serve as a useful benchmark for RBC oscillation mapping is DCE-MRI, which allows for regional quantification of various parameters pertaining to tissue haemodynamics [291]. In DCE-MRI, images are acquired before, during and after the arrival of a Gadolinium-based contrast agent (CA). Gadolinium ions are strongly paramagnetic, and so shorten the T_1 values in the tissues where the CA is present. When the paramagnetic particles disperse through the tissue, there is an associated change in MR signal intensity, dependent on the local concentration of CA and factors such as tissue vascularisation, vessel permeability and perfusion. Parametric maps can be derived from analysis of the temporal signal enhancement after CA administration, by using the signal-time curves (semi-quantitative analysis [292]) or fitting to mathematical models (quantitative analysis, which also requires T_1 mapping) [293]. Blood volume, blood flow and mean transit time (the average transit time for a tracer particle to pass through the pulmonary capillary bed) can be estimated from quantitative DCE-MRI analysis. From semi-quantitative analysis, characteristics of the signal-time curve such as time-to-peak (TTP; the time taken for the CA concentration to reach its maximum value), time of arrival and FWHM (the width of the signal-time curve at half its maximum signal intensity) can be extracted. In recent years, the application of DCE-MRI to lung imaging has grown, and pulmonary perfusion has been measured using DCE-MRI in diseases such as IPF, where a significant increase in FWHM in comparison to healthy volunteers was found [260].

To further elucidate the physiological meaning of RBC oscillation phase, evaluation of this metric in relevant patient groups is required. In this chapter, we apply the method of oscillation amplitude and phase mapping to a group of post-COVID-19 and a small group of CTEPH patients and compare the results to those of the healthy volunteer group from the previous chapter. We investigate whether RBC oscillation amplitude and phase mapping can reveal any abnormalities in non-hospitalised long-Covid patients with dyspnoea. As well as evaluating our method in these patients, we aim to validate oscillation phase mapping through comparison to other imaging modalities such as DCE-MRI. In particular, we compare the oscillation phase (ϕ) to the DCE-MRI-derived TTP, which we hypothesise should show some spatial accordance, because both metrics are dependent on the time taken for a CA to reach its maximum concentration and so are sensitive to similar underlying physiology. With the aim of further understanding the origins of RBC oscillation phase differences in patients with CTEPH, a computational model of the pulmonary vasculature [294] is used to perform patient-specific modelling of capillary blood flow.

8.2 Theory: Computational Model of Pulmonary Circulation

The pulmonary vasculature and its response to pathology associated with pulmonary hypertension can be simulated using numerical modelling. Such a model has been developed by Ebrahimi et al. [294, 295, 296], which builds on work by the research team at the University of Auckland Bioengineering Institute [297, 298, 299]. Here, we summarise the key details of the model; for a full description, the reader is directed to the aforementioned references.

8.2.1 Model Geometry

The pulmonary circulation model is based on three levels of blood vessel: macrovascular (arteries and veins), acinar (arterioles and venules) and capillary bed. The macrovascular structure is derived using high-resolution micro-CT lung images and a volume filling branching algorithm [297, 298]. This algorithm creates branching artery structures that are represented by a start point, end point and radius and end at ∼ 32,000 terminal blood vessels. For simplicity, the venous structure is assumed to follow the arterial structure. Below this level, a nine-generation branching model [300] is used to generate discrete arterioles and venules in a ladder-like model, with multiple connections to the capillaries, which are modelled as recruitable sheets [301, 302, 303]. A schematic of the pulmonary circulation model is shown in Figure 8.2.1.

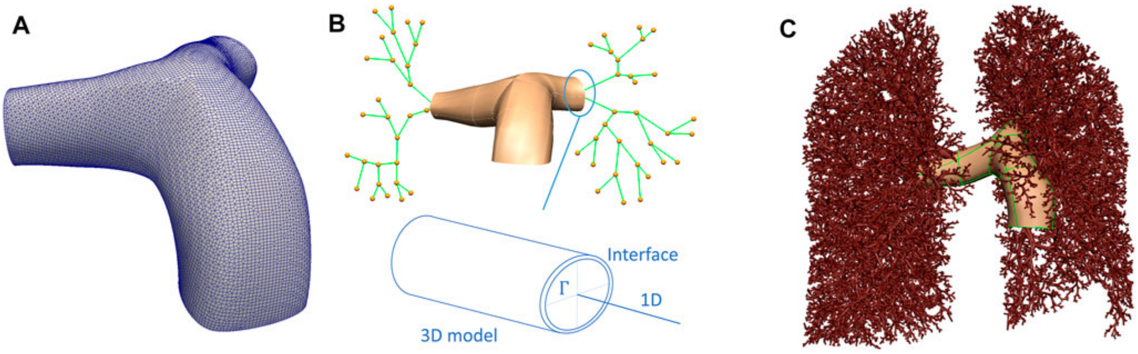

Figure 8.2.1: Pulmonary circulation model: (A) 3D model mesh derived from CT imaging, (B) depiction of how the 3D model and 1D model are connected at the interface Γ and (C) the whole lung model. Copyright 2022 Ebrahimi et al., reproduced with permission from [296].

8.2.2 Simulating Blood Flow and Vessel Compliance

The pulmonary circulation model consists of both a steady-state and dynamic component. An electric circuit analogue model is used to simulate static blood flow through the branching vessel structure [299]. Ohm's law defines electric potential difference as the product of current and resistance. In the electric circuit analogue, each vessel is represented as an element with a resistance R, flow Q (analogous to current) and compliance C (analogous to capacitance). The pressure drop across the vessel, ΔP, is analogous to potential difference. Thus, Ohm's law becomes:

$$\Delta P = QR \tag{8.2.1}$$

The pressure drop in an incompressible and Newtonian fluid in laminar flow through a cylindrical tube is given by Poiseuille's law:

$$\Delta P = \frac{8\mu L Q}{\pi r^4}, \tag{8.2.2}$$

where μ is the fluid viscosity, L is the length of the tube and r is the radius. Therefore,

$$R = \frac{8\mu L}{\pi r^4} \tag{8.2.3}$$

To incorporate vessel compliance into the model, to simulate distensible vessels, radius becomes a function of pressure:

$$r = r_0\left(1 + \frac{3r_0}{4Eh_v}P_{tm}\right) = r_0\left(1 + CP_{tm}\right), \tag{8.2.4}$$

where r_0 is the unstrained radius, E is the Young's modulus of the vessel wall, h_v is the vessel wall thickness and P_{tm} is the transmural pressure (the difference in blood pressure and air pressure in the surrounding alveolar tissue). In the capillaries, blood flow is simulated using Fung's sheet flow theory [303], which allows for a change in capillary diameter across the capillary sheet. To account for capillary compliance, the capillary sheet height, H, is permitted to vary with transmural pressure:

$$H = H_0 + C_c P_{tm}, \tag{8.2.5}$$

where H_0 is the unstrained capillary sheet height and C_c is the capillary compliance. Capillary flow is given by

$$Q_c = \frac{SA}{\mu_c f l_c^2}\int H^3 dP_{tm}, \tag{8.2.6}$$

where SA is the capillary surface area, μ_c is capillary blood viscosity, f is a constant term and l_c is the average route length between arteriole and venule in the capillary sheet. The Poiseuille equation is modified for gravity by adding a pleural pressure term and solved iteratively for every vessel segment's pressure and flow using a sparse matrix solver.

The dynamic part of the solution is based on the Navier-Stokes equations, which give the following equations for position and time-dependent pressure, $p(x,t)$, and flow, $q(x,t)$:

$$\frac{\partial^2 p(x,t)}{\partial t^2} = c_0^2 \frac{\partial^2 p(x,t)}{\partial x^2}, \tag{8.2.7}$$

$$\frac{\partial^2 q(x,t)}{\partial t^2} = c_0^2 \frac{\partial^2 q(x,t)}{\partial x^2}, \tag{8.2.8}$$

where c_0 is the wave speed constant. The wave equations are solved using the approach of Duan and Zamir [304], who defined each extra-acinar segment with an admittance (inverse of impedance). Wave reflections occur when vessels bifurcate; in the arteries, the vessels diverge, whereas in the veins, the vessels converge, and so the wave transmission in the pulmonary network consists of both forward travelling and backward travelling waves.

In the capillary bed sheet, Fung's theoretical formulation for pulmonary microvascular impedance under dynamic conditions is applied [303]. The local average sheet thickness is resolved into three components:

$$\bar{h}_s(x,y,t) = h_{s,SI}(x,y) + e^{i\omega t}H(x,y) + \Phi(x,y), \tag{8.2.9}$$

where $h_{s,SI}$ is the steady-state sheet thickness, assuming impervious walls, $e^{i\omega t}H(x,y)$ is a pulsation term describing oscillations of amplitude $H(x,y)$, and $\Phi(x,y)$ is a permeation term. From this equation, ordinary differential equations can be derived and either solved analytically (by assuming a constant sheet height on both the arteriole and venule sides [299]) or, as in this work, numerically, using a finite difference approach [295].

The final model solution is given by the superposition of the static and dynamic flow solutions.

8.2.3 Boundary and Initial Conditions

The model inlet is the MPA and the outlet is the pulmonary veins. At the inlet, the cardiac output is modelled using a cosine Fourier series incident flow wave:

$$q(t) = q_0 + \sum_{n=1}^{N} A_n \cos(\omega_n t + \phi_n), \tag{8.2.10}$$

where q_0 is a constant inflow of 4.8 L min^{-1}, ω_n is the frequency, A_n is the amplitude and ϕ_n is the phase of the n^{th} harmonic. Values of these parameters for the first eight terms are given in Ref. [294]. A mean left atrium pressure of 5 mm Hg is used as the outlet condition and a boundary condition of zero reflection at the left atrium is imposed.

8.2.4 Modelling Pathology

The model can be perturbed in order to simulate lung pathology. To implement the vascular pathology specific to CTEPH, a combination of occlusion and remodelling is applied. Occlusions are simulated by reducing the diameter of arteries until there is no blood flow through them. Remodelling is simulated according to the five grades defined by Heath and Edwards through histological analysis of PAH [305], via the application of three factors to Equation 8.2.4 and a pruning factor, p_f:

$$r = n_f r_0 \left(1 + C_h C_f C_0 P_{tm}\right) \tag{8.2.11}$$

1. Hypertrophy and hyperplasia (C_h): a reduction in compliance due to hypertrophy and hyperplasia of the smooth muscle cells in the pulmonary arteries is observed in the early stages of remodelling.
2. Fibrosis (C_f): a reduction in compliance due to fibrosis is observed for remodelling of grade 3 and above.
3. Proliferation of the intima (n_f): narrowing of the vessel lumina resulting from proliferation of endothelial cells in the intima occurs in grade 2 remodelling and reaches a limit in grade 4 remodelling.
4. Vessel pruning (p_f): in vessels where there is proliferation of the intima, lumen narrowing leads to occlusion in stage 2 remodelling and beyond.

The diameter range of the vessels affected by these factors is dependent on the stage of remodelling [305, 295]. The factors are defined according to the following functions, where RM is a continuous variable in the range 1 - 100, which is based on interpolating Heath and Edwards' five remodelling grades:

$$C_h(RM) = \begin{cases} \frac{1}{6} & \text{if } RM \geq 60 \\ -\frac{1}{60}RM + \frac{7}{6} & \text{if } RM < 60 \end{cases} \tag{8.2.12}$$

$$n_f(RM) = \begin{cases} 1 & \text{if } RM \leq 20 \\ -\frac{9}{800}RM + 1.225 & \text{if } 20 < RM < 60 \\ 0.55 & \text{if } RM \geq 60 \end{cases} \tag{8.2.13}$$

$$p_f(RM) = \begin{cases} 0 & \text{if } RM \leq 20 \\ \frac{1}{160}RM - \frac{1}{8} & \text{if } RM > 20 \end{cases} \tag{8.2.14}$$

$$C_f(RM) = \begin{cases} 1 & \text{if } RM \leq 50 \\ -\frac{1}{60}RM + \frac{11}{6} & \text{if } RM > 50 \end{cases} \tag{8.2.15}$$

8.3 Methods

8.3.1 Subject Details

As in the previous chapter, an SNR threshold of 4.5 was applied to the RBC images of patients with non-zero whole-lung oscillations to select the data for this study. After excluding 22 data sets due to low SNR, the COVID-19 cohort consisted of 54 patients: 10 subjects who were imaged at either six or twelve months following hospitalisation due to COVID-19 pneumonia (who were part of the MURCO study [190]) and 44 subjects who were not hospitalised but had ongoing symptoms and were diagnosed with long-Covid (and were part of the EXPLAIN study [235]). The COVID-19 patients were divided into four sub-groups, shown in Table 8.3.1.

Table 8.3.1: Definitions and group sizes for the four post-COVID-19 patient cohorts. Demographic data are provided in the results (Section 8.4).

Group	n	Description
PCH	5	Patients hospitalised following COVID-19 pneumonia
PCH-RLA	5	Patients hospitalised following COVID-19 pneumonia who had residual lung abnormalities
PC	15	Patients not hospitalised with COVID-19 who had ongoing symptoms post-infection but symptoms did not include breathlessness
PC-B	29	Patients not hospitalised with COVID-19 who had ongoing symptoms including breathlessness post-infection

Five of the patients hospitalised with COVID-19 had residual lung abnormalities (PCH-RLA) present at 12 weeks after hospital discharge, identified on CT imaging by a chest radiologist. Follow-up CTs (mean time since discharge: 6.6 ± 2.7 months) identified ground glass opacities in 5 of 5 subjects, reticulation in 4 of 5 subjects, and fibrotic-like changes in 2 of 5 subjects. The PC and PC-B patients did not have any abnormalities identified on CT imaging or prior respiratory disease. Five CTEPH patients were studied. After excluding one patient due to low RBC SNR, the CTEPH patient cohort included four patients, all of whom were identified as having pre-capillary disease via RHC.

8.3.2 Image Acquisition

Dissolved ^{129}Xe imaging was carried out using a MESI sequence [189], as described in Section 4.2.3. The 44 post-COVID-19 patients were imaged on a 1.5 T GE HDx scanner, whereas the four CTEPH patients were scanned using a 1.5 T GE Artist scanner. ^{129}Xe gas was hyperpolarised and delivered as described previously. All subjects underwent ^{129}Xe ventilation imaging and paired anatomical proton imaging, acquired as described in Section 4.2.2, although these images were only used for image registration purposes and not analysed in this chapter. Semi-quantitative DCE lung perfusion imaging was performed in two of the CTEPH patients using a 3D SPGR sequence. Images were obtained at inspiratory breath hold (at a lung volume comfortable to the patient) following the injection of 0.05 ml/kg of Gadovist (Bayer) administered at 4 ml/s with a saline flush of 20 ml at the same rate. 40 or 44 time-frame images were acquired, with approximately 0.5 s per volume. Other sequence parameters were: matrix size = 200 x 80, bandwidth = ±125 kHz, TE = 0.6 ms, TR = 1.7 ms and flip angle = 20°. The CT lung images used in this chapter were acquired using an Aquilion ONE system (Canon Medical Systems) with the following parameters: matrix size = 512 x 512, slice thickness = 0.5 mm (PCH-RLA patients) or 1 mm (CTEPH patients), single collimation width = 0.5 mm, spiral pitch factor = 0.813, x-ray tube current = 280 - 656 mA at peak kilovoltage = 100 kV or 120 kV.

8.3.3 Data Analysis

RBC oscillation mapping was performed using the methods described in Chapter 7. For the COVID-19 cohort, only SW oscillation phase mapping was performed, whereas the CTEPH patient data were analysed using both the SW and Two-Key oscillation mapping methods, in order to compare the two.

To quantitatively compare the oscillation amplitude and phase maps with the parametric maps from DCE-MRI for the two CTEPH patients who had perfusion imaging, the maps first needed to be co-registered. To do this, the ^{129}Xe ventilation images, which were already registered to an anatomical proton image, were used as an intermediary. The procedure for performing the perfusion registration and analysis is illustrated in Figure 8.3.1.

All images were preprocessed to ensure that they had the same orientation and resolution. The ^{129}Xe-paired anatomical images were resized from a 512 x 512 x 28 matrix to a 256 x 256 x 30 matrix size to match the perfusion images. The ventilation images were downsampled to a 64 x 64 x 64 matrix, and the gas exchange and oscillation mapping images were extrapolated to the same resolution. The time-series perfusion images were co-registered to the initial time-point image (Figure 8.3.1A). The initial time-point image was then masked and registered to the masked anatomical image (Figure 8.3.1B). This transform was applied to the rest of the time-series perfusion images and to the perfusion mask (Figure 8.3.1C). The gas-phase image from dissolved ^{129}Xe gas exchange imaging was registered to the ventilation image and this transform applied to the M and RBC images, mask and the α_{2-Key}, α_{SW} and ϕ maps. Registrations were performed using the Advanced Normalization Tools (ANTs) C++ library [306]. The anatomical and ventilation images were the ‘fixed’ images and the perfusion and gas-phase images were the ‘moving’ images which were warped during registration. Three registration algorithms were used, applied in the following order: rigid, affine and SyN. The rigid transformation only rotates and translates the moving image, whereas the affine step can additionally apply shearing and scaling. SyN is based on a bi-directional diffeomorphism deformation model, which has been found to be one of the most accurate non-linear registration algorithms in the brain [307]. For each algorithm, the registration was repeated iteratively at different resolution levels, from low to high.

Following registration, semi-quantitative perfusion analysis [292] was performed using the warped perfusion images to create maps of pulmonary blood volume (PBV), pulmonary blood flow (PBF), mean transit time (MTT) and TTP (Figure 8.3.1D). A linear relationship between the first-pass MR signal intensity and the CA concentration in the voxel was assumed. A region of interest was drawn in the MPA in a central lung slice of the the peak perfusion image. From this, the arterial input function (AIF), representing the passage of the bolus of contrast agent through the artery, was extracted. MTT was calculated as the first moment of the MR signal intensity-time curve:

$$MTT = \frac{\int t(S(t) - S_0)dt}{\int (S(t) - S_0)dt}, \tag{8.3.1}$$

where $S(t)$ is the measured signal as a function of time and S_0 is the baseline signal intensity. The regional PBV was calculated from the area of the MR signal intensity-time curve for a ROI, normalised by the integrated AIF. Using the central volume principle, PBF was calculated with the following equation:

$$PBF = \frac{PBV}{MTT} \tag{8.3.2}$$

TTP was determined from the time taken for the signal enhancement curve to reach its maximum value. Only the PBV and TTP maps were considered in the further analysis.

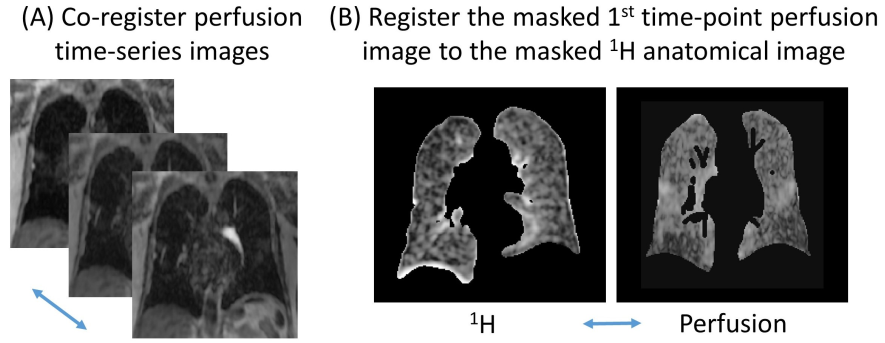
(A) Co-register perfusion time-series images

(B) Register the masked 1st time-point perfusion image to the masked ^{1}H anatomical image

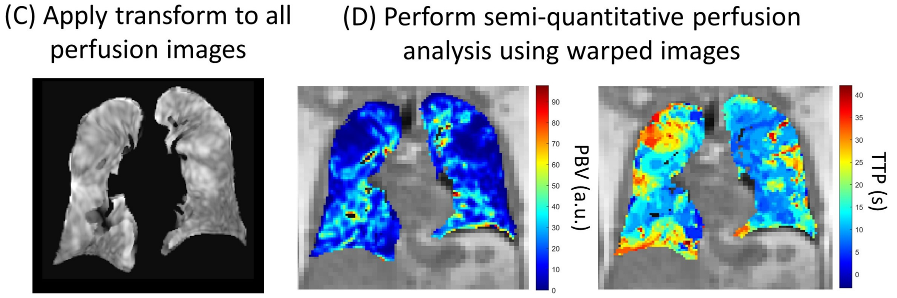
(C) Apply transform to all perfusion images

(D) Perform semi-quantitative perfusion analysis using warped images

Figure 8.3.1: Workflow used to produce the semi-quantitative perfusion maps. (A) The time-series perfusion images were registered to each other. (B) The first time-series perfusion image was masked and registered to the masked ^{129}Xe ventilation paired anatomical image. (C) This transformation was applied to all time-series perfusion images and (D) the warped images were analysed using tracer kinetic theory to create parametric maps, for example of pulmonary blood volume (PBV) and time-to-peak (TTP).

The co-registered parameter maps were masked by multiplying the warped masks from the perfusion and gas exchange image analysis. To evaluate potential differences between the ^{129}Xe gas exchange ratios and RBC oscillation mapping metrics in poorly-perfused and well-perfused lung regions, an additional mask representing regions of perfusion defect was created. A threshold was applied to a map of normalised peak perfusion, $S_{peak,norm}$, defined as the maximum of the normalised difference between the perfusion signal at each time point and the baseline signal:

$$S_{peak,norm} = max\left(\frac{S(t) - S_0}{S_0}\right) \tag{8.3.3}$$

Perfusion defects were defined as voxels with values of $S_{peak,norm}$ less than 20% of the 75% percentile of all voxels within the lung mask. This thresholding method is similar to approaches from literature [308], with the choice of threshold level based on previous empirical results from our group from COVID-19 patients and visual comparison of the resulting defect mask and PBV maps (as shown in Figure 8.3.2).

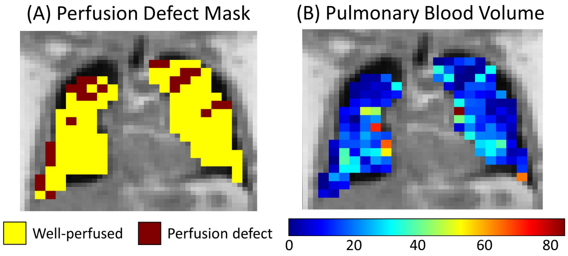

Figure 8.3.2: (A) Perfusion defect mask for one lung slice, where yellow represents well-perfused lung and red represents perfusion defects, and (B) the corresponding PBV map. The threshold level was set empirically as 20% of the 75% percentile of the peak perfusion.

8.3.4 Statistical Analysis

Normality of variables was tested for with Shapiro-Wilks tests. Due to the small size of some of the patient groups, non-parametric tests were used to test for differences between the oscillation mapping metrics. For the COVID-19 cohort, differences between subgroups and the healthy volunteer group from Chapter 7 were assessed using Kruskal-Wallis tests and post-hoc Dunn tests with a Benjamin-Hochberg multiple comparisons correction. Other differences were tested for using Student's t-test or Wilcoxon rank sum tests for normally/non-normally distributed variables. Correlations between variables were evaluated using Pearson's correlation coefficient for normally distributed variables and Spearman's correlation coefficient for non-normal variables. A significance level of $p<0.05$ was used for all tests.

8.3.5 Computational Modelling of the Pulmonary Vasculature

Blood flow simulations were carried out in Python by Dr Behdad Shaarbaf Ebrahimi from the University of Auckland Bioengineering Institute, using the model described in Section 8.2, which is implemented as a Fortran library (available at https://github.com/LungNoodle/lungsim).

The model was modified to reflect patient-specific pathology. The lungs, large arteries and veins were automatically segmented from the CT images of each

CTEPH patient and used to create finite element meshes. The 1D arterial tree models were then generated as described previously, with the MPA and MPV radii calculated from CT. The locations of proximal occlusions were deduced from the CT images and corresponding radiological reports. It was assumed that the CT pixel intensity was proportional to the relative blood flow, in order to identify regions of low perfusion [295]. Occlusions were included in the model by reducing the diameter of the arteries feeding the under-perfused subsegments until blood flow through them was negligible. Remodelling was simulated by reducing vessel compliance and radius and by pruning vessels according to the stages of remodelling defined in Section 8.2 until the model returned the mean pulmonary artery pressure (mPAP) and PVR that matched the measured values from RHC.

The simulation results from Dr Ebrahimi consisted of blood flow values for each node at 11 timepoints: 0 ms, 8 ms, 16 ms, 24 ms, 32 ms, 40 ms, 48 ms, 56 ms, 64 ms, 72 ms and 80 ms. To convert blood flow into variables comparable to RBC oscillation amplitude and phase, we used two different methods:

1. Two-Key Method Analogue: the times at which the whole-lung average blood flow reached its minimum and maximum values were found, denoted t_{min} and t_{max}. For each node, the blood flow amplitude, $\Delta F_{2-Key,i}$, was calculated as the difference between the blood flow at that node, F_i, at t_{min} and t_{max}:
$$\Delta F_{2-Key,i} = F_i(t_{max}) - F_i(t_{min}) \tag{8.3.4}$$
2. Sliding Window Method Analogue: $\Delta F_{SW,i}$ was found from the difference between the minimum and maximum blood flow at each node:
$$\Delta F_{SW,i} = max(F_i) - min(F_i) \tag{8.3.5}$$
A phase lag time was defined as the difference between t_{max} and the time at which each node reached its maximum value:
$$t_{lag} = t_{i,max} - t_{max} \tag{8.3.6}$$

8.4 Results

8.4.1 RBC Oscillation Mapping in Post-COVID-19 Patients

No significant differences were found between α_{k0} or the SW oscillation mapping metrics of the healthy volunteer group and the COVID-19 subgroups (Table 8.4.1). There were significant differences in age, RBC:M and RBC:Gas between the groups; the details of the differences are shown in Table 8.4.2. Boxplots comparing the key metrics between the groups are shown in Figure 8.4.1. The overlaid jitter plots expose the distribution of data points, which are quite widely spread, particularly for the PC and PC-B groups.

Several significant correlations were found between the oscillation mapping metrics when all COVID-19 groups were combined. Figures 8.4.2 and 8.4.3 summarise these significant correlations, including the correlation coefficients and

Table 8.4.1: Subject details and oscillation phase mapping results for the post-COVID-19 patients, along with the healthy volunteer group from Chapter 7. The final column shows the *p*-value from the Kruskal-Wallis tests, where an asterisk denotes a significant result.

	Healthy	PCH	PCH-RLA	PC	PC-B	*p*-value
n (female)	28 (12)	5 (0)	5 (0)	15 (5)	29 (19)	-
Age (years)	38.8 ± 11.1	58.8 ± 9.7	64.6 ± 9.7	47.3 ± 12.9	40 (18 - 55)	0.0002*
RBC:Gas	0.0040 ± 0.0011	0.0029 ± 0.0008	0.0024 ± 0.0004	0.0035 ± 0.0009	0.0033 ± 0.0010	0.006*
RBC:M	0.41 (0.31 - 0.58	0.29 ± 0.09	0.23 ± 0.05	0.34 ± 0.09	0.39 ± 0.08	0.0004*
α_{k0} (%)	15 (10 - 27)	17 ± 5	25 ± 8	19 ± 7	17 ± 4	0.08
α_{SW} (%)	29 ± 3	29 ± 3	33 ± 4	30 (25 - 38)	29 ± 3	0.23
ϕ (rad)	0.27 ± 0.19	0.20 ± 0.17	0.33 ± 0.19	0.34 (0 - 0.66)	0.15 (0 - 0.84)	0.28
CV_{SW}	0.66 ± 0.06	0.67 ± 0.04	0.69 ± 0.07	0.66 ± 0.08	0.63 ± 0.06	0.38
CV_{ϕ}	0.85 ± 0.05	0.84 ± 0.05	0.83 ± 0.04	0.86 ± 0.05	0.86 ± 0.07	0.50

Table 8.4.2: Summary of significant differences between the post-COVID-19 patient subgroups and healthy volunteers, assessed with post-hoc Dunn's tests.

Variable	Group 1	Group 2	*p*-value
Age	Healthy	PCH	0.006
	Healthy	PCH-RLA	0.0002
	PC-B	PCH	0.02
	PC-B	PCH-RLA	0.005
RBC:M	Healthy	PCH	0.04
	Healthy	PCH-RLA	0.002
	PC-B	PCH-RLA	0.005
RBC:Gas	Healthy	PCH-RLA	0.01

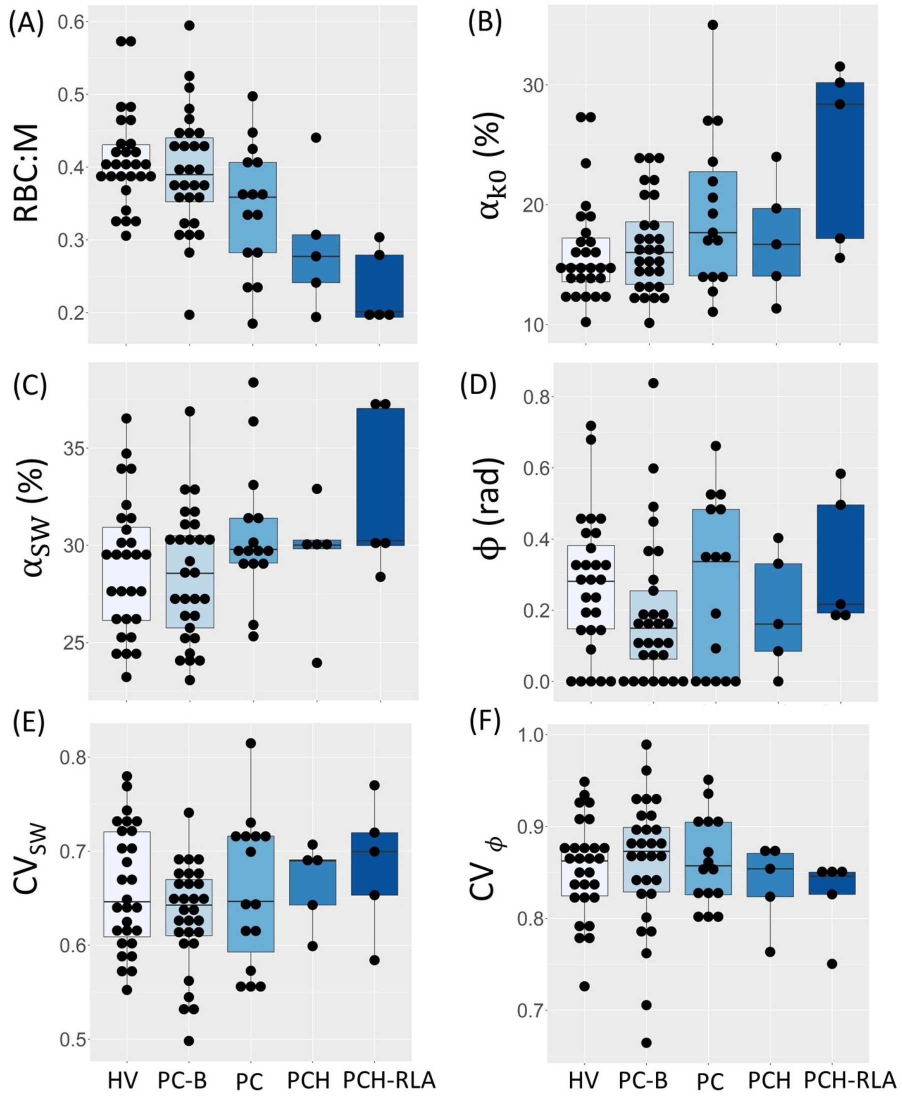

Figure 8.4.1: Boxplots comparing the key oscillation phase mapping metrics between the healthy volunteers, non-hospitalised post-COVID-19 patients with and without breathlessness and the hospitalised post-COVID-19 patients with and without residual lung abnormalities: (A) RBC:M, (B) α_{k0}, (C) α_{SW}, (D) ϕ, (E) $\mathrm{CV_{SW}}$ and (F) CV_{ϕ}.

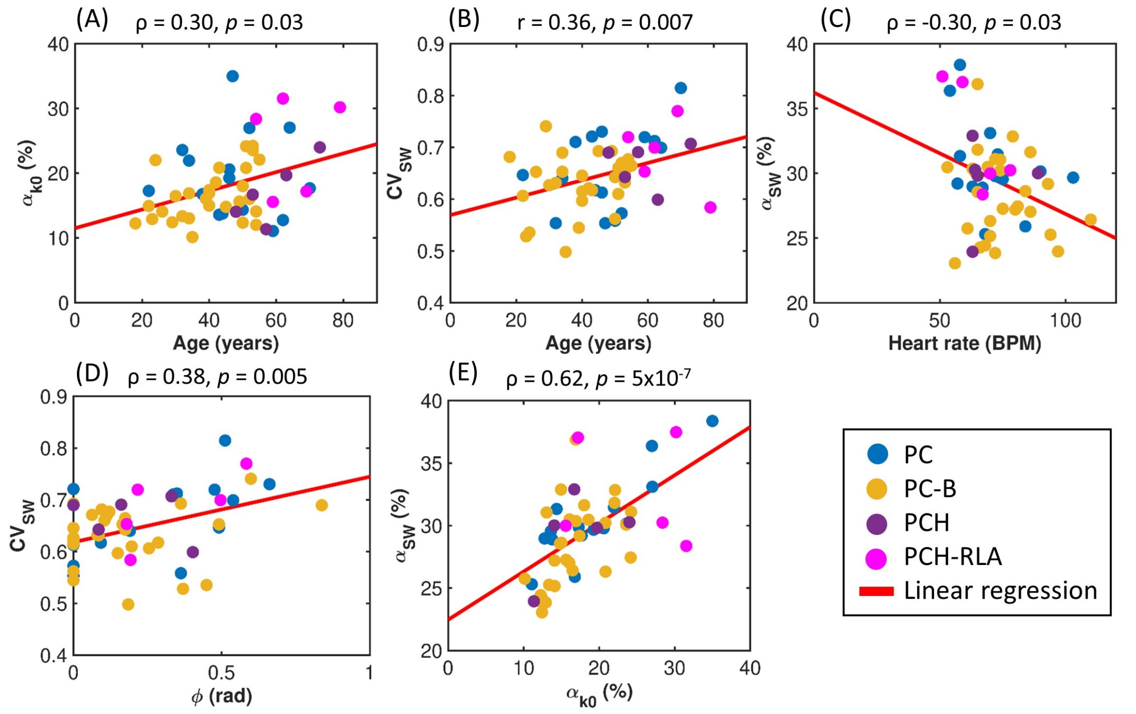

Figure 8.4.2: Significant correlations between RBC oscillation mapping metrics, age and heart rate for the post-COVID-19 patient cohort.

p-values. Both α_{k0} and CV_{SW} were significantly correlated with age, but no significant correlations were found with age and α_{SW} or ϕ. As with the healthy volunteers (see Section 7.3), a negative correlation was observed between α_{SW} and heart rate, but ϕ and heart rate were not significantly correlated. Also like the healthy volunteer group, median α_{SW} was significantly correlated with α_{k0}. Weak negative correlations were found between α_{k0} and RBC:M and RBC:Gas. Patients with worse gas transfer (lower RBC:M and RBC:Gas) had more heterogeneous α_{SW} maps. CV_{SW} was significantly correlated with ϕ and patients with reduced RBC:M also had increased median ϕ. Measurements of TL_{CO} and haemoglobin were available for some, but not all, patients. In the 50 patients who had TL_{CO} measurements, there was only a significant correlation with α_{k0} (ρ = -0.39, p = 0.005). 30/44 long-Covid patients had haemoglobin measurements; no significant correlations between haemoglobin and the oscillation metrics were found.

Significant regional trends were found for both α_{SW} and ϕ, which are shown in Figure 8.4.4. As with the healthy volunteer results in Chapter 7 (see Figure 7.3.5), α_{SW} was higher in the lower lung compared to upper, right compared to left, anterior compared to posterior and peripheral to central lung. ϕ was significantly elevated in the peripheral lung compared with the central lung, as was found for the healthy volunteers, but was also higher in the right lung than the left lung, and in the anterior than the posterior, whereas no other significant trends in ϕ were found for the healthy volunteers. Figure 8.4.5 shows the regional trends for just the PCH and PCH-RLA patients; these were similar to the regional trends found in the healthy volunteers, except there was not a left-right trend in α_{SW} and there was a left-right trend in ϕ.

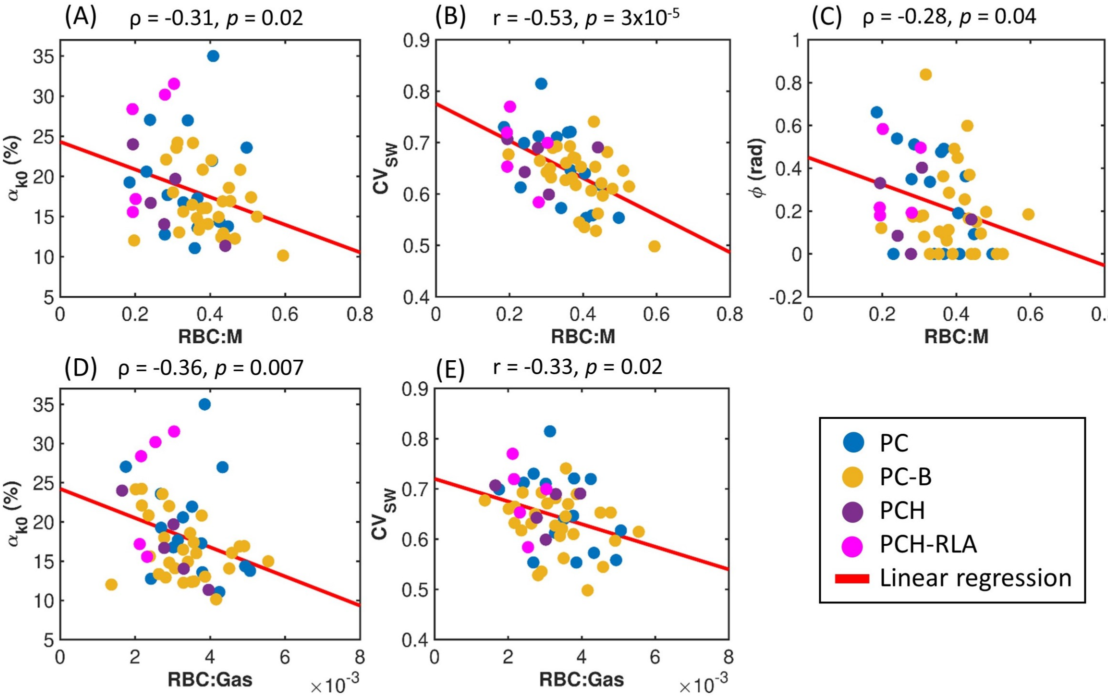

Figure 8.4.3: Significant correlations between RBC oscillation mapping and gas exchange imaging metrics for the post-COVID-19 patient cohort.

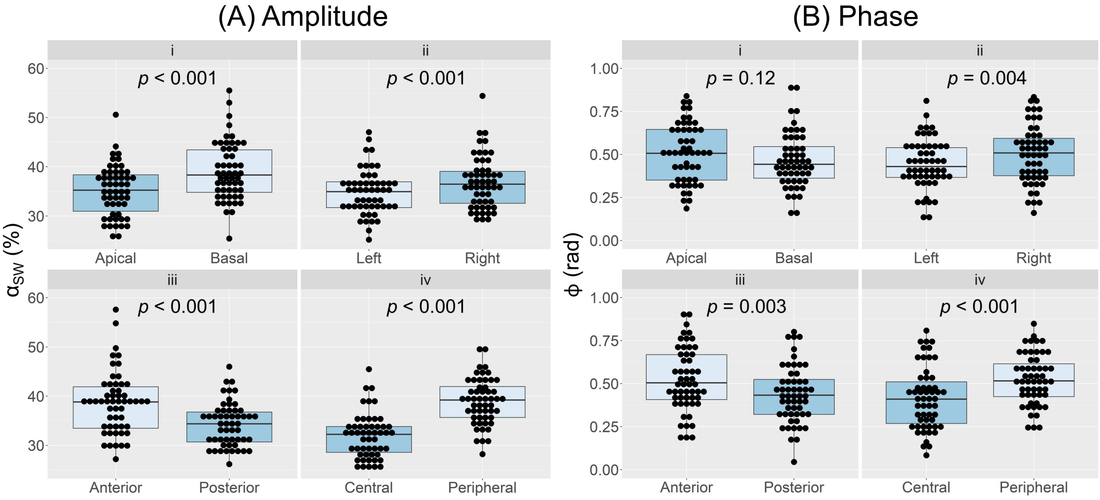

Figure 8.4.4: Comparison of (A) α_{SW} and (B) ϕ between the (i) apical/basal, (ii) left/right, (iii) anterior/posterior and (iv) central/peripheral regions of the lung for the COVID-19 patient cohort, with the corresponding *p*-values from paired t-tests and Wilcoxon signed-rank tests.

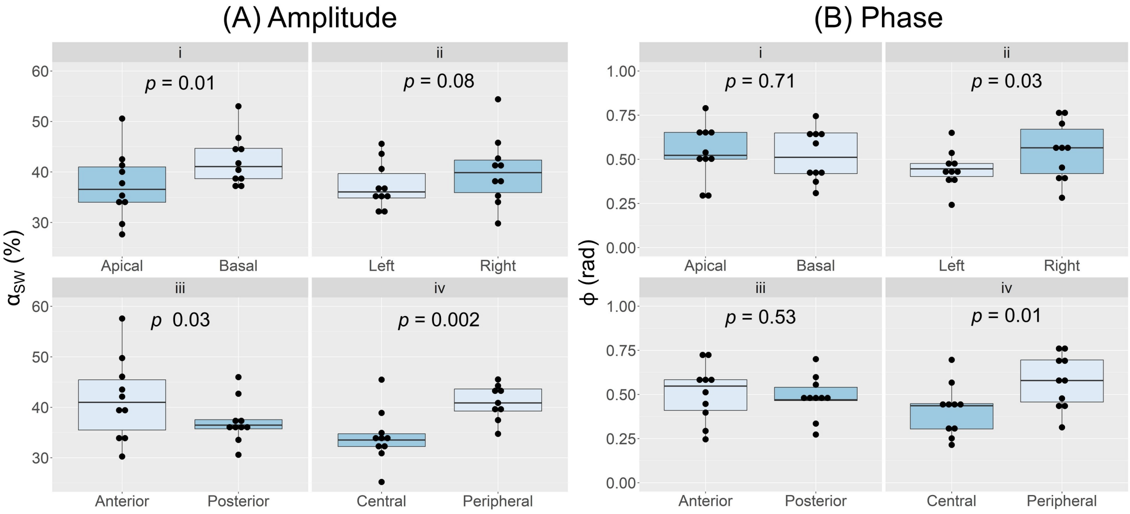

Figure 8.4.5: Comparison of (A) α_{SW} and (B) ϕ between the (i) apical/basal, (ii) left/right, (iii) anterior/posterior and (iv) central/peripheral regions of the lung for the post-COVID-19 hospitalisation patient cohort, with the corresponding p-values from paired t-tests and Wilcoxon signed-rank tests.

8.4.1.1 Comparison with CT

In some of the PCH-RLA patients, regions of increased opacity on CT approximately corresponded to regions of increased ϕ in the oscillation phase maps. Figure 8.4.6 shows the CT image for two of the PCH-RLA patients and their oscillation phase and amplitude maps.

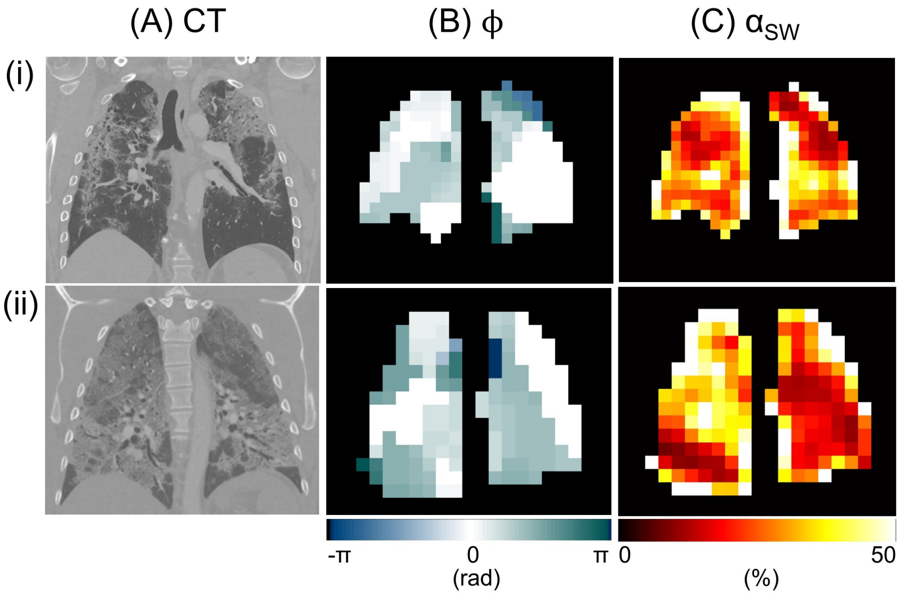

Figure 8.4.6: (A) CT images and RBC oscillation (B) phase and (C) amplitude maps for (i) a 59 year old male PCH-RLA patient who had ground glass opacities and reticulation and (ii) a 54 year old male PCH-RLA patient who had ground glass opacities, reticulation and fibrosis.

The CT image of the first patient (top row) shows increased opacity in the upper left lung, which qualitatively corresponds to a region of increased phase difference. The second patient (bottom row) has increased opacity in the lower lungs, which corresponds to increased phase difference and reduced oscillation amplitude in the right lung.

8.4.2 RBC Oscillation Mapping in CTEPH Patients

Patient demographics, haemodynamic measurements and RBC oscillation mapping results for each of the four CTEPH patients are shown in Table 8.4.3, along with their DCE-MRI results where applicable. Two of the patients had received previous surgical treatment for their disease, but were referred for imaging due to a worsening of their symptoms on exertion and possible presence of residual disease.

Table 8.4.3: Demographic data and MRI results for patients with CTEPH. For CTEPH04, some clinical information was unknown. BPA = balloon pulmonary angioplasty.

	CTEPH01	CTEPH02	CTEPH03	CTEPH04
Age (years)	54	71	66	61
Sex	Male	Male	Male	Male
WHO functional class	II	III/IV	III	Unknown
Comorbidities	COPD	Type 1 respiratory failure	Obstructive sleep apnea	Pneumonia (January 2024)
EmPHasis-10 score [309]	35	40	13	N/A
Prior PEA/BPA?	PEA	No	BPA	No
mPAP (mm Hg)	32	58	54	36
PVR (WU)	3.2	10.1	10.6	4.1
RBC:Gas	0.0019	0.0023	0.0029	0.0028
RBC:M	0.17	0.28	0.27	0.30
α_{k0} (%)	19	8	16	17
α_{2Key} (%)	10	7	13	11
α_{SW} (%)	36	32	37	28
Median ϕ (rad)	0.14	0.33	0.36	0.12
RMS ϕ (rad)	1.29	1.34	1.24	1.11
CV_{2Key}	3.9	3.5	2.0	2.3
CV_{SW}	0.74	0.76	0.76	0.76
CV_{ϕ}	0.94	0.93	0.94	0.88
PBV (a.u.)	N/A	N/A	16.2	16.5
TTP (s)	N/A	N/A	13.4	6.3

A comparison of key oscillation mapping metrics from this group with the healthy volunteer group from Chapter 7 is shown in Table 8.4.4. α_{k0}, α_{SW} and ϕ were similar to those of the healthy volunteers, but the α_{2-Key}, α_{SW} and ϕ maps were more heterogeneous, with significantly elevated CV values. Mean α_{2-Key} was lower in this patient group than for the healthy volunteers. Figure 8.4.7 shows the RBC:M, α_{2-Key}, α_{SW} and ϕ maps for six slices of the lung for one of the CTEPH patients. There is substantial heterogeneity in the phase maps and in slice 6, the left and right lung appear to be in anti-phase with each other. Areas of RBC transfer defect corresponded to regions of increased phase difference.

Table 8.4.4: Comparison of RBC oscillation mapping results between healthy participants and patients with CTEPH. Differences between the groups were tested for using Wilcoxon rank sum tests.

	Healthy	CTEPH	*p* value
Age (years)	38.8 ± 11.1	63.0 ± 7.3	0.0036*
RBC:Gas	0.0040 ± 0.0011	0.0025 ± 0.0005	<0.0001*
RBC:M	0.41 (0.31 - 0.58)	0.25 ± 0.06	<0.0001*
α_{k0} (%)	15 (10 - 27)	15 ± 5	0.8047
α_{2-Key} (%)	14 ± 3	10 ± 3	0.0286*
α_{SW} (%)	29 ± 3	33 ± 4	0.0541
Median ϕ (rad)	0.27 ± 0.19	0.24 ± 0.13	0.8416
CV_{2-Key}	1.4 ± 0.3	2.9 ± 0.9	0.0001*
CV_{SW}	0.66 ± 0.06	0.76 ± 0.01	0.0039*
CV_{ϕ}	0.85 ± 0.05	0.92 ±0.03	0.0134*

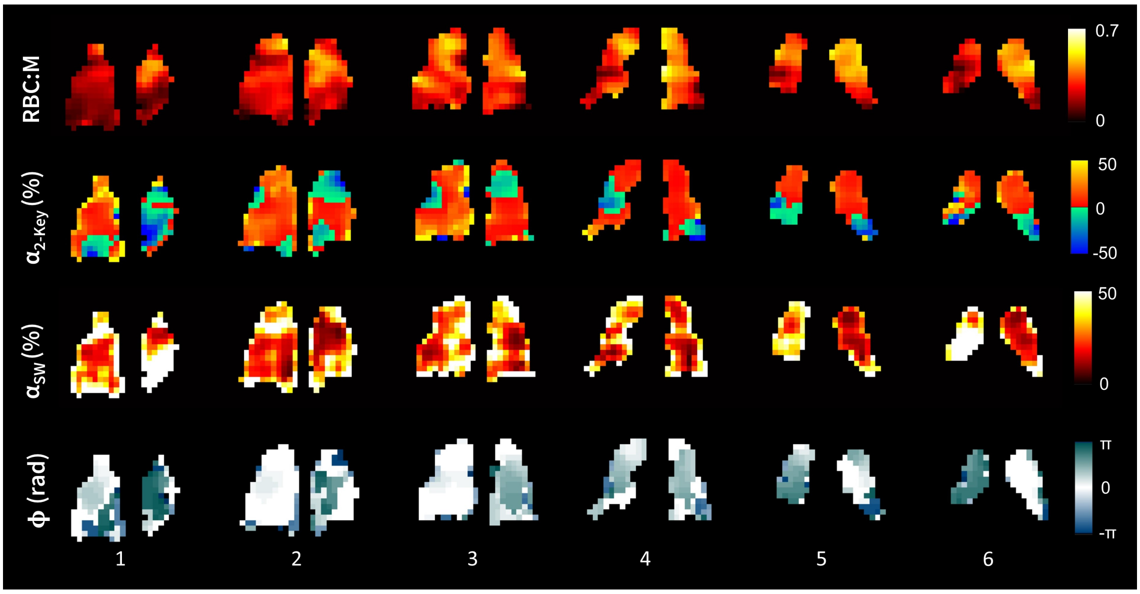

Figure 8.4.7: Maps of RBC:M, α_{2Key}, α_{SW} and ϕ for patient CTEPH02, shown for six slices of the lung, where 1 - 6 represents the posterior to anterior direction.

8.4.2.1 Comparison with DCE-MRI

In the two CTEPH patients who had DCE imaging (CTEPH03 and CTEPH04), areas of reduced perfusion in the PBV maps corresponded to reduced gas transfer and increased phase difference, as demonstrated in Figure 8.4.8. Phase was also increased in areas without reduced perfusion. There was a significant reduction in RBC:M, RBC:Gas and M:Gas within the perfusion defect mask when compared with the rest of the lung for both patients. α_{SW} and abs(ϕ) were found to be significantly higher in the areas of reduced perfusion. In one patient, no difference was found for α_{2-Key} between the well-perfused and poorly-perfused lung regions, and in the second patient a decrease in α_{2-Key} was found in the perfusion defect mask. Boxplots comparing the gas exchange ratios and oscillation metrics for voxels within the low/well-perfused lung masks for each patient are shown in Figures 8.4.9 and 8.4.10. There were several outliers in the boxplots for α_{2-Key} and α_{SW}, which reflect the long-tailed distributions also seen in the healthy volunteer oscillation amplitude distributions (Figure 7.3.2).

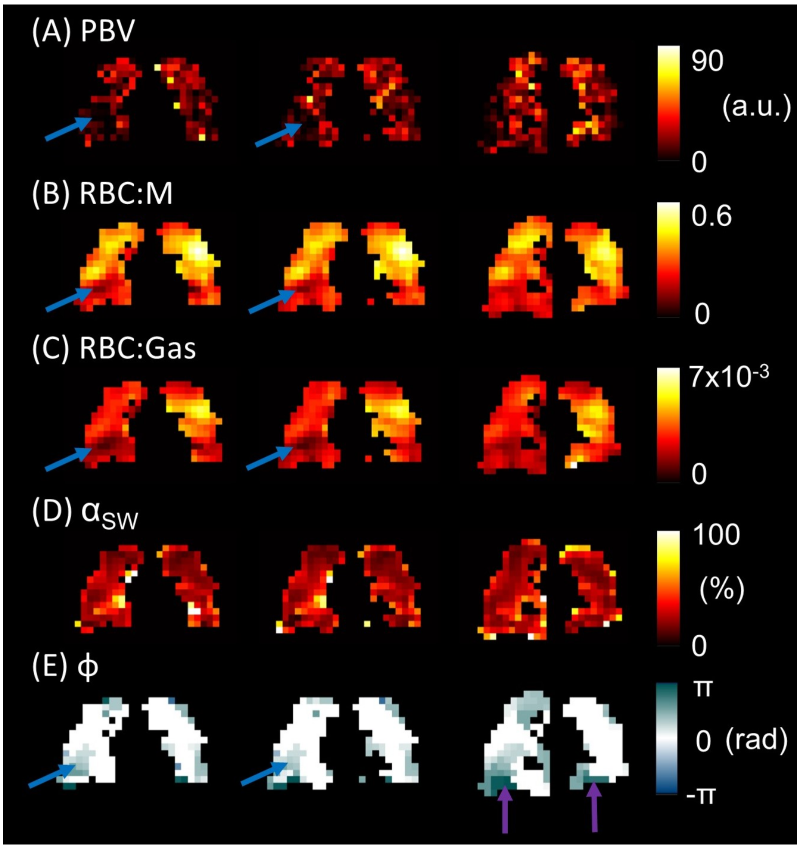

Figure 8.4.8: Maps of (A) PBV from DCE-MRI, (B) RBC:M and (C) RBC:Gas from ^{129}Xe gas exchange MRI and (D) α_{SW} and (E) ϕ from RBC oscillation mapping, shown for three central lung slices, for patient CTEPH04. Areas of reduced perfusion corresponded to decreased RBC:M and RBC:Gas ratios and increased phase, as highlighted by the blue arrows. Phase was also increased in areas of the lung that did not correspond to reduced PBV, RBC:M or RBC:Gas (purple arrows).

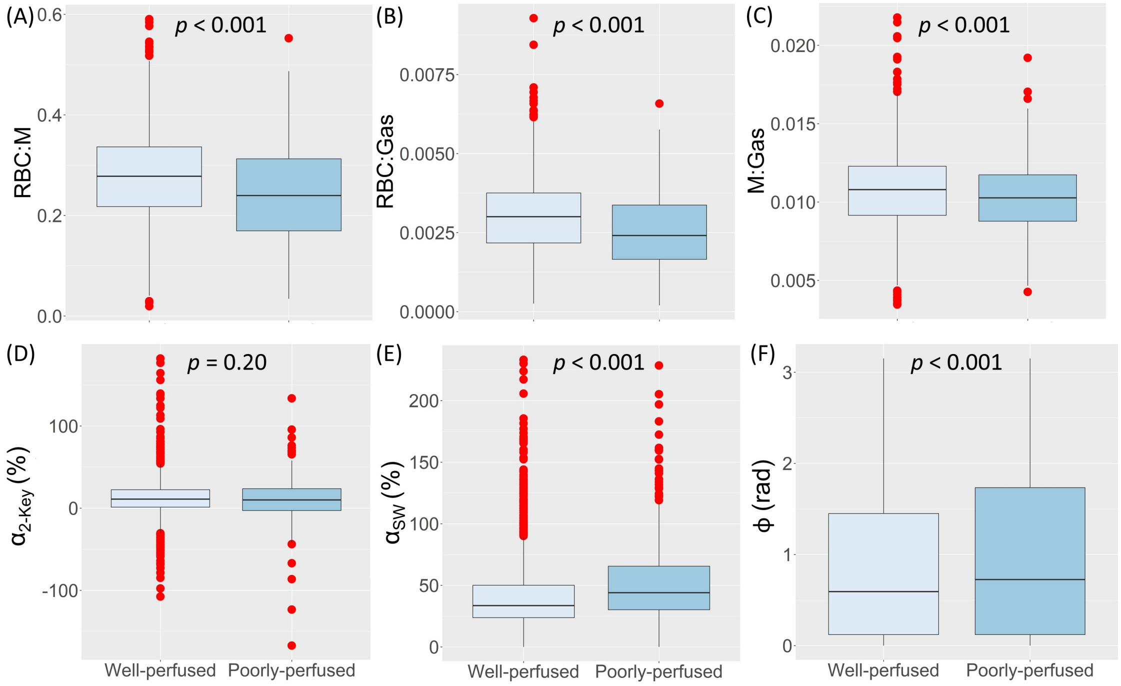

Figure 8.4.9: Comparison of (A) RBC:M, (B) RBC:Gas, (C) M:Gas, (D) α_{2-Key}, (E) α_{SW} and (F) ϕ between the well-perfused voxels and poorly-perfused voxels (defined based on thresholding the peak perfusion signal) for patient CTEPH03. Red circles indicate outliers.

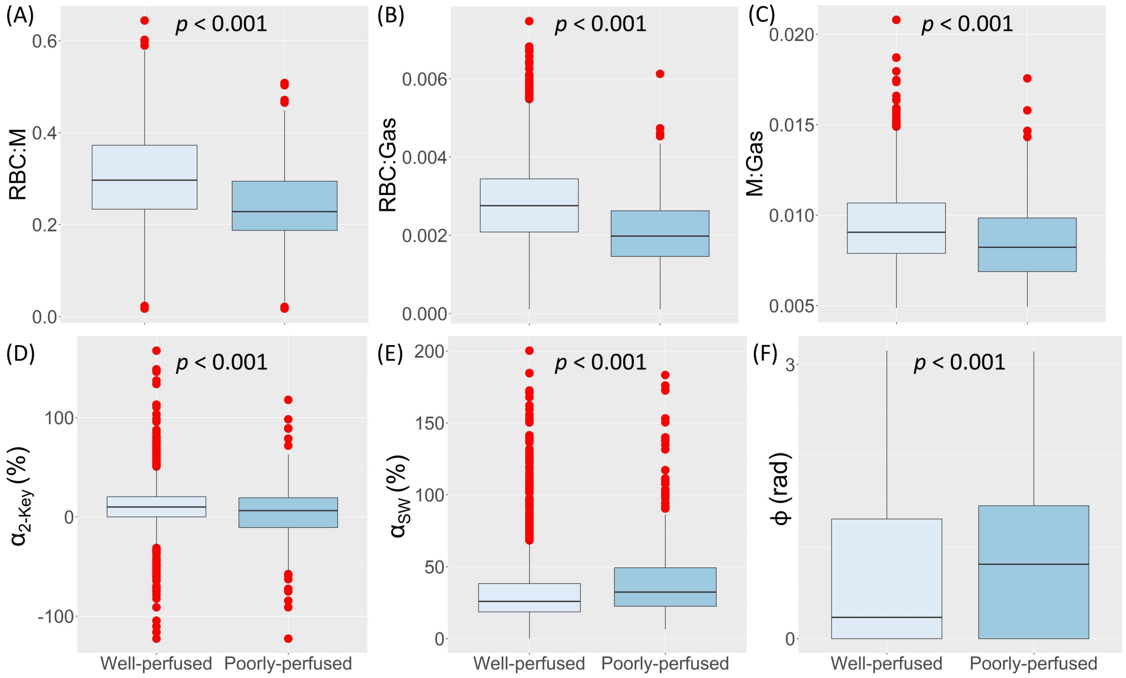

Figure 8.4.10: Comparison of (A) RBC:M, (B) RBC:Gas, (C) M:Gas, (D) α_{2-Key}, (E) α_{SW} and (F) ϕ between the well-perfused voxels and poorly-perfused voxels for patient CTEPH04. Red circles indicate outliers.

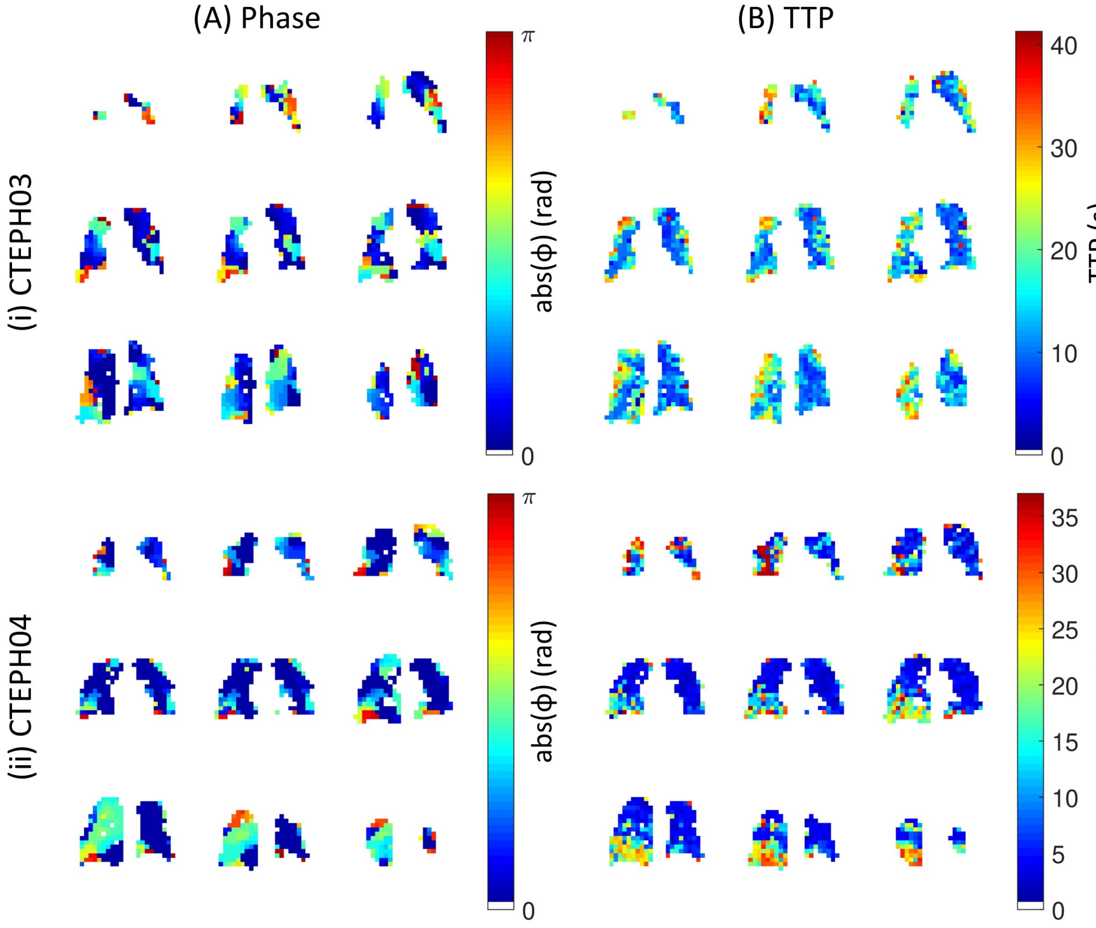

Figure 8.4.11: Comparison of (A) ϕ and (B) TTP from DCE-MRI for two patients with CTEPH.

Figure 8.4.11 compares the maps of ϕ and TTP for both CTEPH patients. To facilitate qualitative comparison of the maps, the same colour map was used for both metrics and absolute ϕ was used so that both parameters had a lower limit of zero. The periodic nature of ϕ means that voxels with a phase difference of π and $-\pi$ are equivalent, and correspond to the maximum phase delay, which physiologically relates to maximum TTP. In both patients, there were some visual similarities between the ϕ and TTP maps, with areas of increased ϕ roughly corresponding to areas of increased TTP.

8.4.2.2 Comparison with Computational Model

Recent CT scans were not available for two of the four CTEPH patients, so modelling was only performed for CTEPH02 and CTEPH04, who had CT imaging on the same day as MRI. The remodelling stages that replicated the haemodynamic measurements from RHC were $RM = 49$ for CTEPH02 and $RM = 41$ for CTEPH04. The whole-lung average blood flow for one of the simulated data sets is shown in Figure 8.4.12. From this graph, the times t_{min} and t_{max} were identified and ΔF_{2-Key}, ΔF_{SW} and t_{lag} were calculated as defined in Equations 8.3.4, 8.3.5 and 8.3.6. Representative slices for each of the model-derived maps are shown in Figure 8.4.13 and the mean map values are given in Table 8.4.5, along with the values of the remodelling factors. As with α_{2-Key}, ΔF_{2-Key} exhibited some negative values resulting from the signal at t_{min} exceeding the signal at t_{max}.

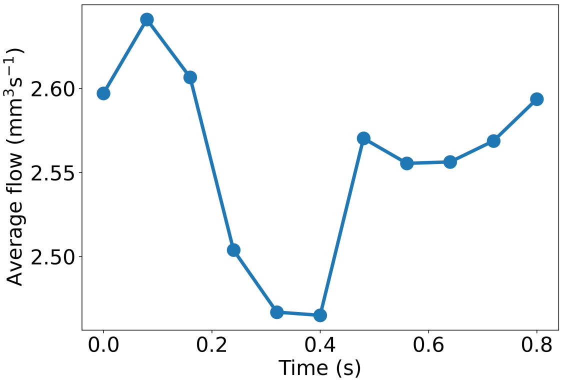

Figure 8.4.12: Blood flow at each time point, averaged over all nodes, from the computational model.

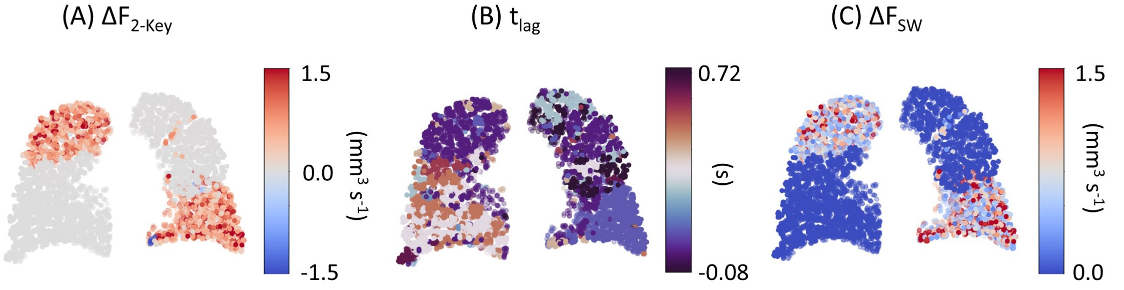

Figure 8.4.13: Maps of (A) ΔF_{2-Key}, (B) t_{lag} and (C) ΔF_{SW} found from the computational model for a patient with CTEPH (CTEPH04). Regions of the lung with increased t_{lag} heterogeneity correspond to regions with decreased blood flow amplitude.

The simulated phase lag maps showed extensive heterogeneity, with greater CV values than the ^{129}Xe RBC phase maps (1.45 and 0.99 compared to 0.93 and 0.88). The simulated phase lag and RBC oscillation phase maps for eight similar lung slices from patient CTEPH02 are shown in Figure 8.4.14. There are some similarities in the regional distribution of phase values, for example the increased t_{lag} in the right lung compared to the left lung and in the lower left lung in the second and third slices. However, there are also areas of increased t_{lag} that do not correspond to an increased oscillation phase and vice versa.

Table 8.4.5: Remodelling factors, blood flow amplitude and phase lag map results from the computational pulmonary circulation model.

	CTEPH02	CTEPH04
RM	49	41
C_h	0.350	0.483
n_f	0.674	0.764
p_f	0.181	0.131
C_f	1	1
ΔF_{2-Key} ($\mathrm{mm^3s^{-1}}$)	0.03 (-1.62 - 5.18)	0.00 (-9.63 - 7.30)
$\mathrm{CV}_{\Delta F_{2-Key}}$	1.36	1.54
ΔF_{SW} ($\mathrm{mm^3s^{-1}}$)	0.09 (0.00 - 5.18)	0.00 (0.00, 13.71]
$\mathrm{CV}_{\Delta F_{SW}}$	1.38	1.59
t_{lag} (s)	0.08 (0.00 - 0.72)	0.08 (0.00 - 0.72)
$\mathrm{CV}_{t_{lag}}$	1.45	0.99

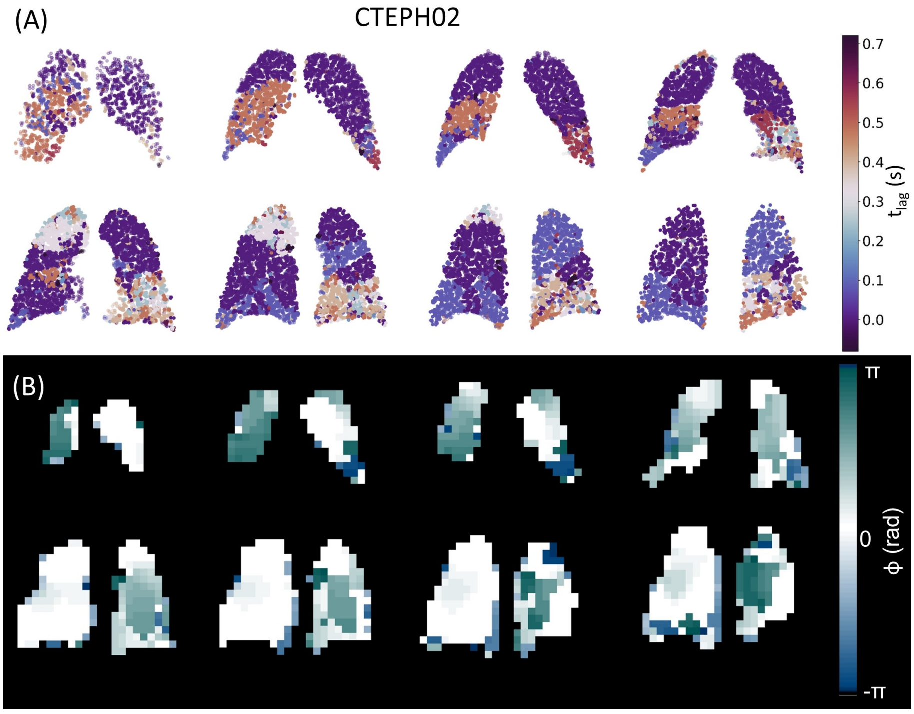

Figure 8.4.14: Comparison of (A) phase lag from the computational blood flow model and (B) RBC oscillation phase for eight matched lung slices for patient CTEPH02.

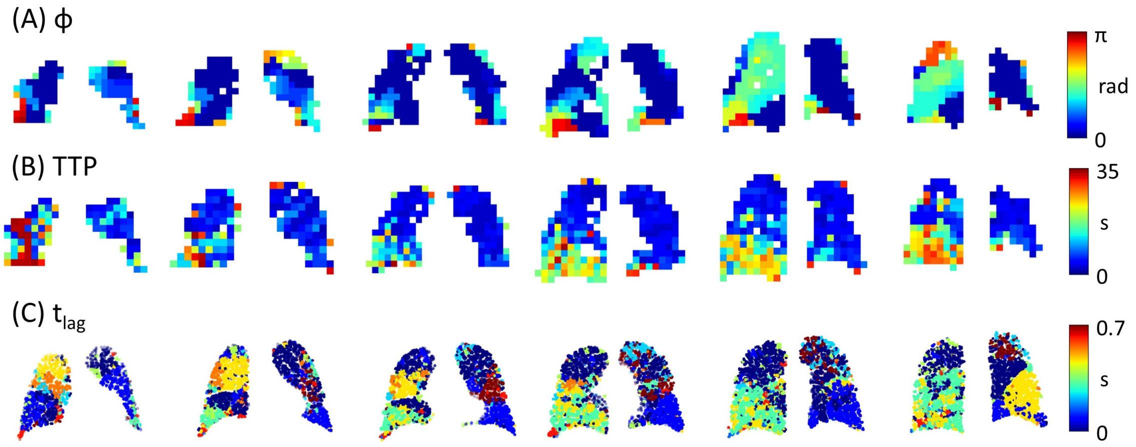

Figure 8.4.15: Comparison of (A) t_{lag} from the computational blood flow model, (B) TTP from DCE-MRI and (C) RBC oscillation phase for six matched lung slices for patient CTEPH04.

Figure 8.4.15 compares the simulated phase lag maps with RBC oscillation phase and TTP from DCE-MRI. Here, the maps of absolute phase which were registered to the DCE-MRI images are displayed. All maps show increased phase/phase lag in the lower right lung. This correlates to the location of occlusions identified in the patient's radiology report, which diagnosed a mural thrombus in the distal right MPA and occluded main and segmental branches of the lower right lobe.

8.5 Discussion

In this chapter, we applied the RBC oscillation mapping method developed in Chapter 7 to two different patient groups and described the initial progress towards validating our method against DCE-MRI and a numerical model of the pulmonary circulation.

8.5.1 RBC Oscillation Mapping in Post-COVID-19 Patients

Applying RBC oscillation phase mapping to a large but heterogeneous group of patients experiencing ongoing symptoms post-COVID-19 infection did not reveal any significant differences in oscillation amplitude or phase between the subgroups and the healthy volunteer group. There is no previous literature on the RBC oscillation amplitude in COVID-19 patients so it is unclear whether this is related to a lack of sensitivity of our method or because these patients do not have significant microvascular damage. The PC and PC-B patients included in this work were a subset of patients from the EXPLAIN study. We found no evidence of gas exchange impairment in these patients, which is consistent with the findings of the EXPLAIN study [235], although at odds with a previous study [194]. The absence of any evidence of microvascular abnormalities in the PC and PC-B patients is supported by measurements of coagulopathy blood biomarkers in these patients [310]. Blood testing did not find any difference between levels of prothrombotic blood biomarkers of between the groups, however, levels of serum amyloid-A were elevated in the PC and PC-B patients compared to healthy controls, suggesting ongoing inflammation.

The hospitalised COVID-19 patients had more severe disease than the non-hospitalised patients, and a significant reduction in gas exchange compared to the healthy volunteers was seen. This is in line with other literature reports of abnormally low RBC:M in patients who were hospitalised following COVID-19 infection [311, 193, 166]. Although the group size was small (n = 5), our preliminary observations show that mean α_{k0}, α_{SW} and CV_{SW} appear elevated in the PCH-RLA patients when compared with the healthy volunteers, although not significantly. However, Figure 8.4.1 shows that the oscillation mapping results were quite heterogeneous in this patient group and for α_{k0}, α_{SW} and ϕ, the PCH-RLA data points appear to be clustered into two groups. All five PCH-RLA patients had ground glass opacities identified on CT, and we have seen evidence of some concordance between these areas of increased opacity on the CT images with areas of increased phase difference, such as in Figure 8.4.6. The two PCH-RLA patients who had pulmonary fibrosis had elevated α_{k0} values of 28% and 30%. Whole-lung oscillation amplitude has been reported to be increased in patients with IPF [35, 37, 137]. This is because parts of the pulmonary capillary bed are destroyed by fibrosis but cardiac output is preserved, so the same blood volume is delivered to a reduced capillary bed causing a larger relative change between systole and diastole [215]. One of the fibrosis patients also had high median α_{SW} (37%), possibly reflecting regions of ILD-like fibrosis, but the other fibrosis patient had median α_{SW} = 30%, which is more similar to the healthy volunteer values, and neither of the fibrosis patients showed an increase in median ϕ. With only five PCH-RLA patients, the conclusions that can be drawn from these preliminary data are limited.

When considered in aggregate, the relatively large number of post-COVID-19 patients included in this chapter allowed us to assess correlations of the oscillation mapping metrics with demographic and ^{129}Xe gas exchange imaging data. A significant dependence on age was found for α_{k0} in the COVID-19 patients, which was not found for the healthy volunteers in the previous chapter. This may be related to the older age of the more severe, hospitalised COVID-19 patients in comparison to the non-hospitalised patients. Nonetheless, the correlation was preserved in the subgroup of PC-B patients ($\rho = 0.40$, $p = 0.03$) and the heterogeneity of the α_{SW} maps also increased with age, which may suggest that there is greater disease burden in older people. The negative correlation that we observed between heart rate and α_{SW} can be explained by the reduction in blood flow pulsatility and hence relative change in capillary blood volume with higher heart rate, which was discussed in Chapter 7.

The increase in α_{k0} with decreasing RBC:M and RBC:Gas alludes to a worsening of gas exchange related to damage at the capillary bed or post-capillary bed level, which causes greater relative changes in capillary blood volume over the cardiac cycle and thus higher oscillation amplitude. α_{SW} was not correlated with either RBC:M or RBC:Gas, but CV_{SW} was significantly correlated with both ratios, which suggests that the heterogeneity of the oscillation amplitude maps may be a more sensitive biomarker of microvascular damage than the median amplitude values. No correlations were found between TL_{CO} and the oscillation mapping metrics. This may imply that, although the oscillations are related to the ^{129}Xe RBC signal, they are not directly sensitive to diffusing capacity, possibly because the oscillation amplitudes were defined relative to the baseline RBC signal. This also explains why oscillation amplitude was not affected by haematocrit, unlike the ^{129}Xe gas exchange ratios, for which a haematocrit correction has been suggested [259].

The regional trends identified in α_{SW} were similar to those found for healthy volunteers and most likely are dominated by the trends in the inverse of RBC signal, used as the normalisation. Interestingly, further regional trends in ϕ were observed in addition to the central-peripheral trend detected in the healthy volunteers. The increased ϕ values in the right lung when compared to the left lung might be due to the path difference between the right-side vasculature and the left-side, where the cardiac pulse wave originates. The cause of the apparent anterior-posterior trend is less clear and this trend was not present in the subgroups of PCH and PCH-RLA patients.

In summary, our results are inconclusive and do not provide any strong evidence of microvascular abnormalities in post-COVID-19 patients. No previous literature on oscillation mapping exists for these patient groups and there is still much that is unknown about the mechanisms of COVID-19 lung disease and, in particular, long-Covid. Consequently, this patient group is limited as a testbed for our method. Future opportunities to continue this investigation are discussed in Section 8.5.3.

8.5.2 RBC Oscillation Mapping in CTEPH Patients

A more appropriate patient group for testing the sensitivity of our method is CTEPH, a disease with known blood flow impairment resulting from vascular occlusion following pulmonary embolism [288]. Although average α_{k0} for the CTEPH group was the same as that of the healthy group, the oscillation amplitude and phase maps both revealed significant heterogeneity. This demonstrates the potential of oscillation mapping to identify microvascular abnormalities which might otherwise be lost in the whole-lung average oscillation amplitude. The ability of the sliding window keyhole reconstruction to estimate RBC oscillation phase may be particularly useful in this patient group. The observed phase differences in the ^{129}Xe RBC signal oscillations may represent a haemodynamic response at the capillary level to impedance to flow from vascular thrombi in the larger vasculature. The sensitivity of ^{129}Xe RBC oscillation mapping to pulmonary capillary haemodynamics may provide a useful tool for identifying arterial remodelling and indicating the severity of microvascular disease. This is important for clinical decision making because patients with extensive microvascular disease do not respond as well to PEA surgery [312]. PEA is an intense operation that can only be performed at specialist centres, has a risk of complications and a long recovery time [287], so there is motivation to develop methods and models which can help predict post-surgical outcomes in order to select candidates for surgery.

Our results are similar to those of Lu et al., who measured oscillation amplitude in a group of ten CTEPH patients using their implementation of Two-Key oscillation mapping [215]. Akin to our finding of normal α_{k0} values in the CTEPH patients but heterogeneous α_{SW}, α_{2-Key} and ϕ maps, Lu et al. demonstrate a case where a CTEPH patient had whole lung oscillations which were in the healthy range, but their oscillation maps revealed regions of both decreased and increased amplitude. This patient had combined pre-capillary and post-capillary PH based on their RHC measurements and so this highlights a particular advantage of being able to spatially resolve the oscillations. Six of the CTEPH patients in that work were scanned both at baseline and following PEA and the proportion of the lung categorised as defect and low oscillation amplitude (based on a healthy reference population) was found to decrease from 37±13% to 23±14% in these patients. As discussed in detail in Chapter 7, Lu's method did not account for regional variations in oscillation phase. The oscillation amplitudes that they categorised as defect and low based on the healthy volunteer distribution are those with values $\alpha \leq 0$. We believe that these negative amplitudes are caused by increased oscillation phase difference, as we reasoned in the previous chapter. As such, the observed decrease in the number of voxels in the low oscillation and oscillation defect bins may actually reflect an decrease in oscillation phase rather than an increase in oscillation amplitude. Lu et al. did not find a significant change in mean RBC oscillation amplitude after PEA, which might be because they were not measuring amplitude independently of oscillation phase. Pöhler et al. showed that maps of PREFUL pPTT, which is conceptually similar to ϕ, were more homogeneous in CTEPH patients after PEA [267]. We predict that our metrics of ϕ and CV_{ϕ} may be sensitive to detecting changes post-PEA, although currently we do not have any pre- and post-surgery ^{129}Xe MRI data with which to test this hypothesis.

Validation of new reconstruction methods and biomarkers through comparison to clinical gold standards is important to assess their efficacy and reliability. Here we presented preliminary results comparing RBC oscillation phase mapping to DCE-MRI which, as far as we are aware, has not been attempted before. Although our results are limited to only two CTEPH patients, these initial data are promising and the significantly elevated ϕ in regions of the lung with low perfusion increases our confidence that ϕ has a physiological basis and is sensitive to blood flow limitation. Interestingly, we also saw higher α_{SW} values within the perfusion defect mask than in the well-perfused regions. All of the CTEPH patients were diagnosed with pre-capillary disease based on their RHC measurements and pre-capillary PH has previously been associated with decreased whole-lung oscillation amplitude due to increased impedance to flow in the arterioles [37]. However, this effect may be obscured by a concomitant increase in oscillation amplitude resulting from a reduction in mean capillary blood volume [214]. The RBC signal oscillation is assumed to be caused by variations in capillary blood volume, and so increases when a preserved cardiac stroke volume is delivered to a reduced capillary bed. RBC:M, RBC:Gas and M:Gas were also significantly reduced in areas of perfusion defect and so it seems likely that the increased α_{SW} in these areas is due to a decrease in capillary blood volume, which is related to the RBC signal that was used as the normalisation in the oscillation amplitude calculation. Consequently, we expected to see an inverse relationship between α_{SW} and DCE-MRI-derived PBV. No clear relationship between the two maps was observed, and further work is needed to quantitatively assess this relationship.

Inspecting the co-registered ϕ and TTP maps showed some spatial similarities, which was expected because both metrics are detecting the same underlying process; the delay in a CA reaching its maximum value. Regions of increased ϕ and TTP in the lower right lung of patient CTEPH04 corresponded to the location of occlusions identified on CT. Although ϕ was significantly greater in regions of low lung perfusion, there were regions of the lung that were well-perfused but corresponded to increased ϕ values. These preliminary results hint at the potential ability of RBC oscillation phase mapping to identify vascular remodelling in unoccluded lung regions, which is of great clinical importance for patients with CTEPH.

A limitation in comparing the ^{129}Xe MRI-derived gas exchange ratio and oscillation maps is that the image registration was challenging, requiring an intermediate step where the images were registered to ^{129}Xe ventilation imaging. Differences in lung inflation level between the different scans may have lead to mis-registration at the diaphragm, as shown in Figure 8.5.1A. The low resolution of the gas exchange images (1.25 cm^3 voxel size) also made registration less reliable (Figure 8.5.1B). Due to the potential unreliability of the registration, pixelwise correlations between the maps were avoided. Another limitation to our DCE analysis is the thresholding method used to define the perfusion defect mask. There are a variety of methods to threshold perfusion images for this purpose (or for defining perfusion defect percentage). Uniform thresholding methods are popular, although more advanced methods such as Otsu's method and texture-based methods have also been employed [308]. We chose our threshold value based on previous literature values, empirical results based on COVID-19 patient data and by visual inspection. Our chosen threshold seemed to work well for the CTEPH patients, however, it did not take into account the known physiological anterior-posterior gradient which affects perfusion in the supine position. The masking may be improved by using an adaptive threshold which varies slice-by-slice to adjust for this gradient [313].

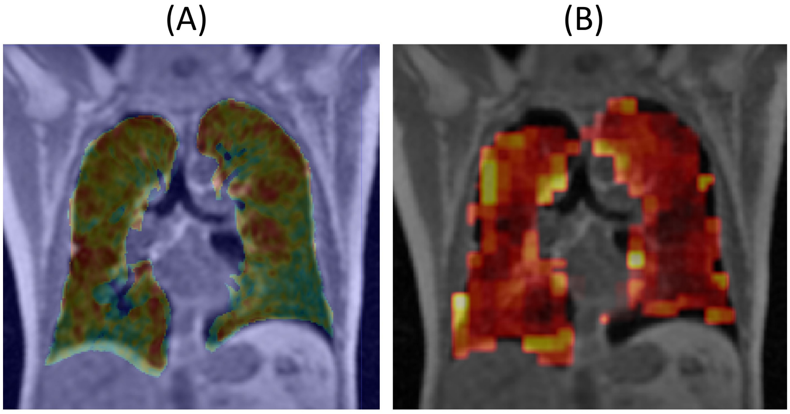

Figure 8.5.1: There were some limitations to the image registration: (A) mis-registration of the perfusion image to the ventilation anatomical image at the diaphragm. (B) Overlay of the registered but low-resolution α_{SW} map on the anatomical image.

As well as comparing the RBC oscillation maps to imaging data, we used a numerical model to simulate blood flow and phase lag. Patient-specific modelling of the pulmonary vasculature was achieved by using CT imaging and RHC measurements as the model inputs. The models demonstrated that negative values arise when ΔF was calculated based on t_{min} and t_{max} defined based on whole-lung blood flow values, in a parallel to the Two-Key oscillation mapping method. They also affirmed the phase heterogeneity we saw in the ϕ maps, and areas of increased t_{lag} roughly matched areas of increased ϕ and TTP. The computational model provided an estimate of the degree of distal remodelling in each patient, which offers insight to the extent and severity of their microvascular disease and may assist in the interpretation of the amplitude and phase maps. In both cases, the best agreement between the model-derived and RHC-derived PVR and mPAP values was found with grade 2 remodelling. This disease stage is characterised by medial hypertrophy with cellular intimal proliferation in the muscular pulmonary arteries of less than 300 μm in diameter [305]. This leads to occlusion of these small vessels, which offers an explanation for increased ϕ in areas of the lung which did not have sub-segmental blood clots. Patient CTEPH02 had RM = 49, and therefore was on the cusp of grade 3 remodelling, where at the earliest stages, fibrous tissue starts to form beneath the endothelium in the arterioles. This patient had higher average ϕ and CV_{ϕ} than patient CTEPH04, which may be demonstrative of the more advanced remodelling predicted by the computational model. Patient CTEPH02 also had a high symptom burden, with an EmPHasis-10 score of 40/50, where a higher score indicates worse quality of life [309]. Although modelling was only performed for two patients, by comparing the t_{lag} maps to the ϕ maps we have explored a potentially novel way to validate RBC oscillation phase mapping. Yet, there are several limitations to the computational model which may hinder its accuracy. To simplify the simulations, some assumptions were made, such as approximating the blood in the extra-acinar vessels as a Newtonian fluid and neglecting secondary reflections in the wave transmission model. Cardiac output remained constant in the simulations, although this may be reduced in disease, and this may have led to an overestimation of mPAP.

8.5.3 Future Work

The small number of lung disease patients included in this work (for the PCH, PCH-RLA and CTEPH groups) reduced the statistical power of the comparisons between groups and therefore increased the probability of a false negative. Future work will include evaluating our method in increased subject numbers, including additional PH subtypes. The whole-lung oscillation amplitude has previously been tested in a diagnostic algorithm, to distinguish between patients with different PH subtypes, patients with ILD and healthy people [209]. Where this algorithm struggled was in differentiating patients who had combined pre- and post-capillary disease; oscillation mapping may offer a solution to this by allowing regions of increased and reduced oscillation amplitude to be visualised and so it would be worthwhile to test our method in this group. It would also be very interesting to perform hyperpolarised ^{129}Xe MRI and oscillation phase mapping in a cohort of CTEPH patients pre- and post-PEA, to investigate post-surgery changes and, in particular, whether PEA results in a reduction in phase heterogeneity. This would help to investigate the role of ϕ, α_{SW}, CV_{ϕ} and CV_{SW} as biomarkers of microvascular disease and predicting response to PEA in CTEPH patients.

An advantage of RBC oscillation mapping is that the oscillation information is available ‘for free’ from dissolved-phase ^{129}Xe imaging; it is already encoded in k-space and so a separate image acquisition is not required. Therefore, RBC oscillation mapping can be easily applied to analyse retrospectively acquired patient data and we plan to employ this to the HRUK SSc/SSc-PAH study next (Table 4.3.1). A caveat to the applicability of our method is that the data must have been acquired using the updated of our dissolved-phase imaging sequence (as in [189]), because the previous version [35] had far fewer radial spokes and so the temporal resolution is too low for reliable oscillation mapping.

Initial attempts to validate our oscillation mapping method were made in this chapter, however more work is needed to fully substantiate our technique. Although we investigated inter- and intra-scan repeatability of the oscillation mapping metrics in healthy volunteers in the previous chapter, it would be worthwhile to evaluate this in patients with lung disease as well. Further comparison to DCE-MRI metrics, with increased subject numbers, is needed to better understand how oscillation amplitude relates to pulmonary perfusion. We only examined the PBV and TTP maps, because these were most clearly related to α_{SW} and ϕ, but it may be useful to compare these with the PBF and MTT maps too. Improving the image registration would allow for more in-depth regional comparisons to be made. Comparisons with metrics of the pulse wave from PREFUL MRI may also be insightful for both healthy subjects and those with pulmonary disease. Oscillation phase offers a parallel to pPTT, which has been found to be longer in CTEPH patients than in healthy people, but decreases following PEA [266]. The initial results presented here for the numerical lung perfusion model will be followed up with quantitative comparison of the simulated blood flow and phase lag maps with metrics from DCE-MRI, ^{129}Xe gas exchange MRI and oscillation phase mapping.

Another valuable aspect of validating novel methodology is cross-site validation. Our oscillation phase mapping code is freely available on GitHub (https://github.com/POLARIS-Sheffield/rbc-phase-map) and has been successfully implemented by Dr Peter Niedbalski from the University of Kansas Medical Center for data acquired with 1-point Dixon method ^{129}Xe gas exchange imaging. It is encouraging that our method seems to work well for data acquired at a different research facility, with a different MRI scanner, imaging sequence and parameters. We hope that future work will involve other sites implementing (and building on) our method and that this will allow us to evaluate cross-site reproducibility.

More work is needed to investigate the relationship between oscillation amplitude and phase and RHC measurements, such as PVR. RHC is the current gold standard for diagnosing PH, but is an invasive procedure which lacks regional sensitivity. A possible application of RBC oscillation mapping is the non-invasive detection of elevated PVR. Costelle et al. have proposed a physiological model based on an electric circuit analogue of the pulmonary capillaries, which can separate the competing affects of reduced capillary blood volume, which causes increased oscillation amplitude, and pre-capillary PH, which decreases oscillation amplitude [214]. They used Wang’s TL_{CO} prediction model [224], which was described in Chapter 5, to estimate capillary blood volume, V_C, from ^{129}Xe gas exchange MRI images. Correcting whole-lung oscillation amplitude for V_C accounts for reduced capillary blood volume in patients with gas exchange defects and allows for prediction of PVR. This is an intriguing application of the TL_{CO} model which links RBC oscillation amplitude to haemodynamic measurements. A limitation to this work is that the model is based on whole-lung measures of RBC:M, RBC:Gas and oscillation amplitude; there is potential for improved sensitivity by applying this model on a regional basis by using RBC oscillation mapping. V_C could also be estimated for each voxel by using the framework we presented in Chapter 5.

8.6 Conclusions

No conclusive evidence of microvascular damage was found for non-hospitalised long-Covid patients, which adds to the general paucity of diagnostic findings in these patients. Although no significant differences were found between the PCH-RLA patients and healthy volunteers, oscillation amplitude appeared elevated, which may reflect the fibrosis and ground-glass opacities identified by CT. In patients with CTEPH, amplitude and phase map heterogeneity were significantly greater than in healthy volunteers, despite similar α_{k0} values. Areas of increased phase difference approximately overlapped with regions of low gas transfer and perfusion and increased TTP, demonstrating the potential sensitivity of our method to identify lung damage at the microvascular level. By comparing our results to DCE-MRI blood volume maps, we may be able to differentiate between phase differences resulting from upstream thrombi, which impede downstream blood flow, and phase differences resulting from distal vasculopathy. This could be a useful tool for quantifying microvascular disease in CTEPH, which is important for predicting postoperative outcomes. Further work, including implementation of our method by another research group, is underway to continue the validation of our method and assessment of its utility as a biomarker of microvascular disease.

Chapter 9

Measuring Pulmonary Gas Exchange with Chemical Shift Saturation Recovery Spectroscopy and Imaging

9.1 Introduction

CSSR measures the temporal uptake of ^{129}Xe magnetisation into the alveolar membrane and RBCs [40]. The background literature on CSSR was reviewed in Section 3.3.1.1; here we briefly recap the main points of the CSSR technique and introduce the methodological advancements presented in this chapter.

Unlike the dissolved-phase ^{129}Xe spectroscopic imaging sequences used in previous chapters [35, 189], which provide a ‘snapshot’ of the dissolved ^{129}Xe signal at a single time-point, CSSR measures the time-resolved dissolved ^{129}Xe uptake. By selectively destroying the dissolved-phase ^{129}Xe magnetisation and waiting a fixed time before acquiring an FID, the dissolved-phase signal can be quantified at different uptake times. Fitting the uptake data to the analytical diffusion models described in Section 3.4 allows for the quantification of clinically relevant parameters such as the alveolar septal wall thickness (h). This is particularly beneficial for identifying lung fibrosis in ILD [39] and inflammation in COPD [314].

Another potentially interesting clinical application of the CSSR technique is in SSc, an autoimmune disease characterised by excessive collagen production in the connective tissue, which affects the skin and internal organs [315, 316]. The lungs are often involved in SSc, and pulmonary fibrosis and pulmonary hypertension are major causes of morbidity and mortality. ILD is estimated to affect 35% - 75% of SSc patients [317, 318] and is the leading cause of death for this patient group [12]. Early diagnosis of SSc-associated ILD (SSc-ILD) is therefore crucial, as is determining whether the fibrosis is stable or progressive [316]. High-resolution CT is the gold standard for identifying pulmonary fibrosis, however CT is not well-suited for repeated imaging and so monitoring disease progression and treatment response in SSc-ILD patients is challenging. The second most common cause of death in SSc patients is PAH, which has an estimated prevalence among SSc patients of between 5% and 15% [319]. Various other PH subtypes can also affect patients with SSc, such as left heart disease and chronic thromboembolism, and PH can further complicate SSc-ILD, resulting in worse outcomes than non-PH SSc-ILD patients [320]. The mechanisms behind pulmonary disease in SSc are not well-understood and there is a need for novel lung imaging techniques to assess the structural and functional changes associated with the disease. Stewart et al. performed CSSR in four SSc patients and found increased septal thickness in these patients when compared to healthy volunteers, despite these patients having little fibrosis and no PH [39].

To further assist the clinical application and utility of CSSR, it is important to be able to quantify the uncertainty on the model-derived parameters. The analytical models used to fit the CSSR data have a high degree of freedom which, coupled with noise in the CSSR uptake curves from low SNR at short delay times (t_{delay}) and signal variations from cardiogenic RBC oscillations [212], can lead to non-optimal curve fitting. However, in most CSSR literature, the uncertainty on the model-derived parameters is either not reported, or provided with no explanation of the error quantification method. Only one paper, by Stewart et al., described their method for deriving the errors on fit parameters [39]. They tested the model accuracy by re-fitting each data set with different weightings applied to specific data points and evaluated the effect on the sum of squares. Statistical methods such as bootstrapping [321] may be able to provide a more efficient way to estimate the uncertainty on the model parameters.

The standard CSSR sequence requires the use of multiple long temporal delays; typically this temporal information comes at the expense of spatial encoding. The CSSR parameters are therefore only available as whole-lung averages, reducing the sensitivity to early lung disease and to disease heterogeneity. Four imaging methods have been proposed for assessment of gas uptake dynamics in humans and small animals, which were described in detail in Section 3.3.1.1 and are summarised below:

1. Look-Locker CSSR, Kern et al. [155]: multiple excitations were acquired after one saturation in a Look-Locker-like scheme [156]. Spatial encoding was achieved using an undersampled stack-of-stars trajectory and low-rank plus sparse matrix decomposition.

2. SPLASH CSSR, Kern et al. [154]: regional CSSR data were acquired by using a multi-channel coil and the SPLASH method [152, 153].

3. CSI-CSSR, Stewart [158] and Ruppert at al. [159]: a CSI sequence was adapted such that CSSR saturation, delay and excitation occurred before acquiring each line of k-space.

4. Time-series spiral-IDEAL, Doganay et al. [99, 148, 149], Zanette et al. [150] and Friedlander et al. [151]: images were acquired using interleaved spiral-IDEAL MESI at different CSSR delay times.

CSI-CSSR is limited by its long scan time; Stewart's 1D CSI-CSSR method required a 15 s breath hold and in Ruppert's 2D implementation, each t_{delay} image required *n* breath holds to complete, where *n* is the matrix size in the phase-encoding direction. Hence, the 2D CSI-CSSR technique has only been tested in rats and is not easily applicable to humans. Kern's in vivo techniques for breath hold CSSR imaging use sparse and low-rank reconstruction and/or parallel imaging. These can be difficult to implement reliably, can lead to undersampling artefacts and require the use of multi-channel RF coils (not currently available at our site). Spiral MESI with IDEAL decomposition is advantageous for rapid image acquisition, but the requirement to loop the IDEAL imaging kernel for multiple spiral interleaves, TEs and CSSR delay times means that spiral-IDEAL CSSR imaging is still limited by breath hold constraints.

Mugler et al. proposed a method for simultaneously imaging the dissolved-phase and gas-phase ^{129}Xe components in 2010 [322]. Since then, more advanced methods such as the Dixon method and IDEAL have been developed for dissolved-phase ^{129}Xe imaging, which can further resolve the dissolved-phase signal into the RBC and M contributions. However, these methods are time-consuming and incorporating them into a CSSR imaging sequence requires undersampling and/or parallel imaging with multi-array coils to be feasible within a single breath hold. The Patz model does not require the dissolved-phase signal to be separated into the RBC and M contributions so methods for separating the RBC and M signals are not necessary. Therefore, Mugler's method offers a fast and straightforward way to image both the dissolved-phase and gas-phase signal in a CSSR-type experiment.

In this chapter, we explore different methods for CSSR spectroscopy and imaging. An improved CSSR analysis pipeline is presented, which uses bootstrapping to quantify the uncertainty on the Patz model parameters. We demonstrate its use in healthy participants and patients with SSc and/or PAH. In addition, we propose a CSSR imaging sequence based on the separation of gas and dissolved peaks by chemical shift dispersion and evaluate this sequence in four healthy volunteers. A Look-Locker CSSR spectroscopy sequence is investigated as a potential means to allow for an increased number of t_{delay} images in the regional CSSR sequence. This sequence is tested in four healthy volunteers and compared to 'standard' CSSR spectroscopy.

9.2 Theory

In Section 2.1.4, we introduced the concept of the chemical shift. To reiterate, the chemical shift is the difference in resonant frequency of nuclei in different chemical environments arising from electron cloud shielding of the magnetic field felt by the nuclei. If the nuclei are present in different chemical environments within the same imaging voxel, and if the frequency bandwidth per pixel is less than the chemical shift, the signal from the two species will be shifted along the readout direction. This is called the chemical shift artefact. By exploiting the ∼200 ppm chemical shift difference between the gas-phase and dissolved-phase ^{129}Xe, a careful choice of bandwidth can be used to image both resonances within the same FOV. To completely separate the dissolved-phase and gas-phase ^{129}Xe signals in the frequency encoding direction, the frequency bandwidth must be less than the chemical shift frequency difference:

$$BW < \frac{\gamma B_0 \Delta\chi_{GP-DP} \Delta x}{W_L} \tag{9.2.1}$$

where γ is the ^{129}Xe gyromagnetic ratio, B_0 is the field strength, $\Delta\chi_{GP-DP}$ is the chemical shift between the gas-phase and dissolved-phase ^{129}Xe spectral peaks, Δx is the image resolution and W_L is the width of the lungs (∼25 cm). If this condition is met, gradient echo imaging will result in side-by-side dissolved-phase and gas-phase ^{129}Xe images, as shown in Figure 9.2.1.

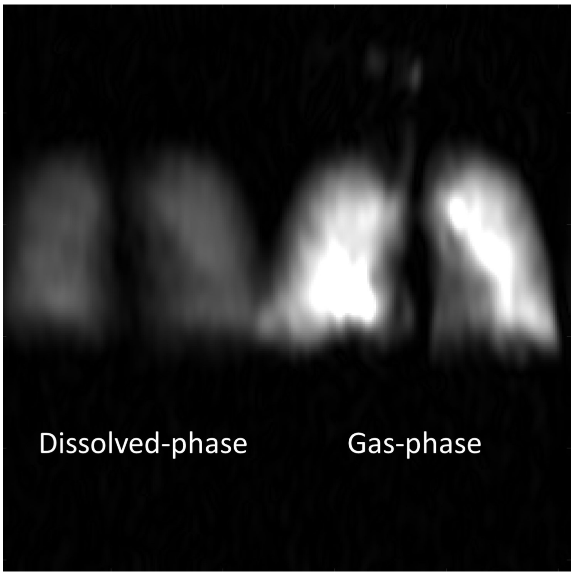

Figure 9.2.1: Combined dissolved-phase/gas-phase ^{129}Xe imaging of a healthy volunteer, acquired with a SPGR sequence: BW = 7 kHz, FOV = 50 cm, TR = 40 ms and dissolved-phase flip angle = 35°.

9.3 Methods

9.3.1 Subject Details

Nine healthy volunteers and six patients with SSc and/or PAH, who were enrolled on the Heart Research UK (HRUK) or PHoenix studies, were included in this work. Subject demographics are summarised in Table 9.3.1, along with RHC and PFT data where available.

9.3.2 CSSR Spectroscopy Pulse Sequences

Our implementation of the CSSR spectroscopy pulse sequence is shown in Figure 9.3.1A. At the start of each ‘block’, the signal from the dissolved-phase ^{129}Xe is destroyed using two frequency-selective 90° RF pulses (described in the following section) and spoiler gradients. After this, there is a fixed delay of length t_{delay}, over which gaseous ^{129}Xe diffuses from the alveolar airspaces, across the membrane barrier and into the capillary bloodstream. The dissolved-phase ^{129}Xe magnetisation is then excited using the frequency-selective RF pulse and an FID

Table 9.3.1: Demographic data for the healthy volunteers and SSc and/or PAH patients included in this chapter. PFT data were not available for the healthy volunteers or PAH patients.

	Healthy	SSc-PAH	PAH
n (female)	9 (4)	6 (5)	2 (1)
Age (years)	30 ± 5	60 ± 5	68 ± 8
mPAP (mm Hg)	—	38 ± 9	42 ± 4
PVR (WU)	—	5.5 ± 3.0	5.5 ± 1.8
TL_{CO} (% predicted)	—	60 ± 25	—
FEV_1 (% predicted)	—	89 ± 18	—

is acquired. The remaining transverse magnetization is destroyed using a spoiler gradient. This sequence block is repeated for each value of $t_{delay,n}$.

An accelerated CSSR spectroscopy pulse sequence is shown in Figure 9.3.1B. This uses a Look-Locker acquisition strategy [156] to acquire FIDs at several t_{delay} times simultaneously, as was first suggested by Kern et al. [155] and was described in Section 3.3.1.1. In this sequence, several excitation pulses are applied after the 90° dissolved-phase signal saturation pulses, with a small flip angle ($\sim$ 10°) such that the effect on the longitudinal magnetisation is small. Therefore, rather than sampling the magnetisation at increasing t_{delay} in a linear fashion and resetting with a saturation pulse each time, the Look-Locker approach allows the different t_{delay} times to run simultaneously. This results in significant time-saving when compared with the standard CSSR approach. To simulate the effect of the excitation pulses on the magnetisation evolution, the Look-Locker CSSR acquisition was simulated using a 2D finite difference method model, as in Kern et al. Simulation details can be found in Appendix A.

In the Look-Locker CSSR sequence, t_{delay} for the n^{th} excitation pulse is given by:

$$t_{delay,n} = td_i + (n-1)T_{min} + \sum_{1}^{(n-1)} x_n, \tag{9.3.1}$$

where td_i is the time between the last saturation pulse and the start of the first excitation pulse and x_n is an optional additional delay which can be added between two excitation pulses to achieve longer t_{delay} times. The minimum t_{delay} is limited by the minimum readout time (T_{min}) of the sequence. To sample smaller t_{delay} times and to increase the temporal resolution, the saturation-excitation blocks can be repeated with different values of td_i. Figure 9.3.1C shows the corresponding uptake data points ($F(t)$) from each excitation block.

Both CSSR sequences were programmed within the GE Healthcare EPIC pulse sequence programming environment and were based on the *fidall* pulse sequence database [323].

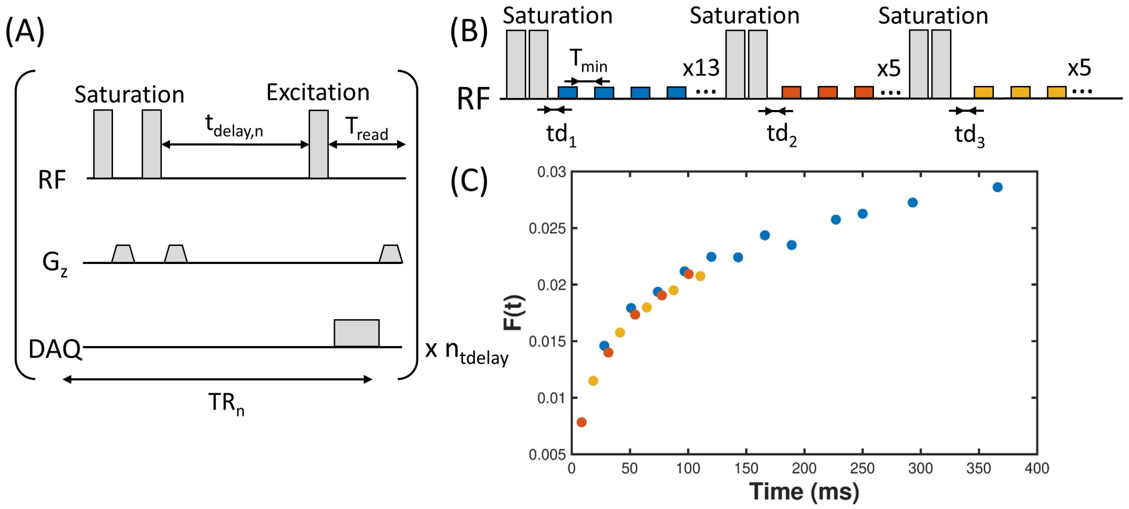

Figure 9.3.1: The pulse sequences used to perform CSSR spectroscopy. (A) Conventional CSSR: the dissolved-phase ^{129}Xe signal is depolarised with two 90° saturation RF pulses and spoiler gradients applied along the z direction (G_z) of width 1 ms, then is replenished by the transfer of ^{129}Xe from the gas-phase into the dissolved-phase over a time $t_{delay,n}$. After this, a third 90° pulse is used to excite the ^{129}Xe signal and acquire an FID. T_{read} is the fixed readout time and DAQ represents the position of the data acquisition window. The CSSR 'block' shown within the brackets is repeated for n_{tdelay} values of $t_{delay,n}$. (B) Look-Locker CSSR: multiple FIDs are acquired following each saturation, with $t_{delay,n}$ defined as the period between saturation and each subsequent small flip angle excitation (blue, red and yellow rectangles). For simplicity, the spoiler gradients and DAQ window are not shown here, but were applied as in (A). To increase the number of short t_{delay} times and the temporal resolution, three saturation blocks are used, each with a different offset (td_i) between the last saturation and first excitation pulse. (C) Example Look-Locker CSSR uptake curve, with the data points colour-coded to show which saturation block they were acquired after.

9.3.3 RF Pulse Design

A custom frequency-selective RF pulse was designed to both saturate and excite the dissolved-phase ^{129}Xe magnetisation, whilst minimising off-resonance excitation. This was done using tools from the GE Healthcare Multi-Nuclear Spectroscopy (MNS) research pack in MATLAB. The time and frequency profiles of this pulse are shown in Figure 9.3.2. The pulse width was 1.5 ms, chosen to be less than the T_2^* of the dissolved-phase ^{129}Xe (2.2 ms at 1.5 T [36]) to minimise signal decay. The pulse had a frequency pass-band of width 500 Hz, centred around 0 Hz. Thus, by setting the centre frequency to the frequency of dissolved-phase ^{129}Xe resonance, maximum RF power is delivered to the dissolved-phase resonance, whilst small side-bands at -3765 Hz to -3565 Hz and 3565 Hz to 3765 Hz deliver a small off-resonance excitation to the gas-phase resonance, which is needed for normalisation.

The nominal flip angles were 90° on-resonance and $< 1°$ off-resonance to minimise the excitation of the gas-phase. However ^{129}Xe's low γ and safety requirements associated with the specific absorbance rate (SAR) can limit the flip angles achievable for a given RF amplifier power. At high RF powers, any non-linearity of the RF amplifier in the high gain regime may also lead to pulse distortion and a delivered flip angle which is not equivalent to the desired flip angle [324, 98]. Careful design and testing of the RF pulse prior to in vivo application was therefore required. The true delivered on and off-resonance flip angles were investigated using a bag of hyperpolarised ^{129}Xe gas. To measure the on-resonance flip angle, 10 FIDs were acquired, with TR = 32 ms, flip angle = 90°, 512 spectral points and a bandwidth of 25 kHz. For the off-resonance flip angle measurement, 600 FIDs were acquired with the centre frequency off-set from the gas-phase frequency by 3650 Hz and TR = 100 ms. The on-resonance flip angle was also calibrated in vivo by acquiring 32 FIDs with TR = 75 ms, flip angle = 90° centred on the gas-phase resonance, 1024 spectral points and a bandwidth of 20 kHz. The transmit gain and centre frequency required to reach 90° were found using the Bloch-Siegert calibration sequence described in Section 4.2.1.

The FIDs were converted to spectra via Fourier transform, and the absolute amplitudes, S, in each case were fitted to a model of polarisation decay:

$$S = S_0 \cos^{n-1}\theta, \tag{9.3.2}$$

where S_0 is the initial signal amplitude, n is the number of RF pulses and θ is the measured flip angle.

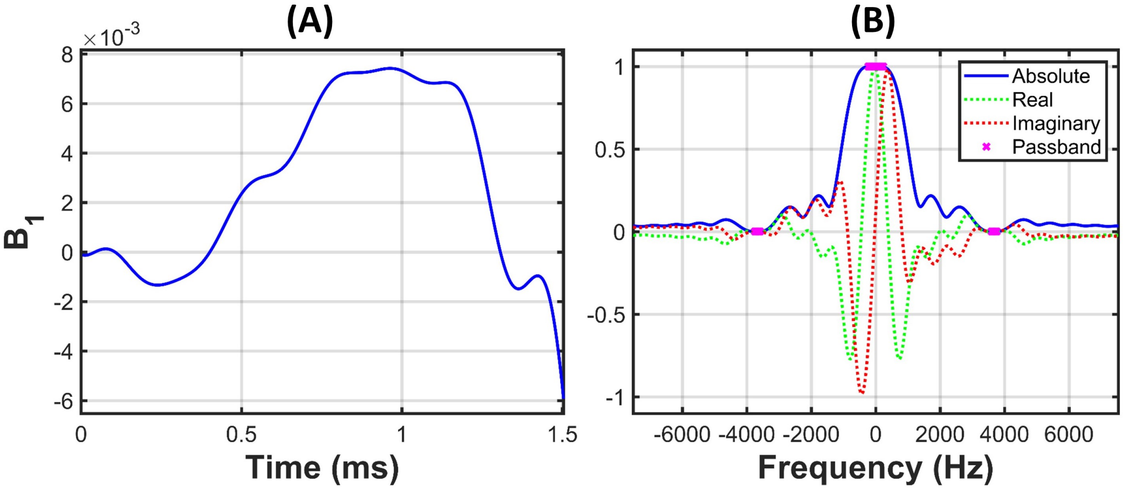

Figure 9.3.2: The frequency-selective CSSR RF pulse profile in (A) the time-domain and (B) the frequency-domain.

9.3.4 CSSR Spectroscopy Acquisition

All CSSR experiments were performed using a 1.5 T GE Artist clinical scanner and a transit-receive vest coil. For the healthy volunteers, a dose of 250 - 500 ml hyperpolarised ^{129}Xe was used, titrated with nitrogen to a total bag volume of 1 L (or 800 ml for volunteers who were $<$160 cm tall). For the patients, the ^{129}Xe dose was scaled according to their estimated lung volume. The total bag volume was 1/6 of the patient's total lung volume, as predicted by the GLI reference equations [227]. The ^{129}Xe concentration was then scaled based on the bag volume; for a bag volume of 1 L, 300 ml of ^{129}Xe was used.

The RF pulse described in the previous section was used for both saturation and excitation, with nominal flip angles of 90°/1° on the dissolved/gas-phase resonances respectively. A calibration scan was used to find the patient-specific transmit gain and gas-phase resonance frequency (as described in Section 4.2.1); the actual centre frequency of the scan was then found by adding 3650 Hz to the calibrated frequency so that the RF power was delivered to the centre of the dissolved-phase resonance (i.e. halfway between the RBC and M peaks, at a chemical shift of 205 ppm from the gas-phase resonance). 25 t_{delay} times were used for standard CSSR, from 6 - 605 ms, which were acquired in two sequential time sweeps:

t_{delay} = [6, 10, 15, 25, 45, 65, 95, 135, 175, 230, 280, 405, 605, 6, 12.5, 20, 35, 55, 75, 115, 155, 195, 255, 305, 505] ms

These values were chosen to accurately sample the gas exchange behaviour, whilst keeping the breath hold short (5 s). Denser sampling was used at the start of the sequence, where the gradient of the uptake curve is steeper and where the curve plateaus, and less dense sampling was used over the linear portion of the curve, where the t_{delay} times are longer and contribute more to the total scan time. Two time sweeps were used to attempt to average out any fluctuations in the dissolved-phase signal from cardiogenic oscillations in the RBC signal and bulk lung motion [325]. 512 spectral points were acquired over a bandwidth of 20 kHz, with T_{read} = 35.6 ms.

All participants underwent CSSR spectroscopy. A subset of four healthy volunteers were also scanned with the accelerated Look-Locker CSSR sequence (Figure 9.3.1B), using the same ^{129}Xe dose as was used for standard CSSR. The bandwidth also remained the same, but the number of spectral points was decreased to 256 in order to reduce T_{min}, which was 23 ms. As shown in Figure 9.3.1B, 23 excitation pulses were used and three saturation pulse blocks. The first ten excitation pulses used T_{min}, as did the excitation pulses after the second and third saturation. An additional delay was added to the final three excitations after the first saturation; $x_{11,12,13}$ = (15, 50, 50) ms. The values of $td_{1,2,3}$ were 5 ms, 8.5 ms and 18.5 ms respectively. The resulting t_{delay} values for Look-Locker CSSR are listed below, which culminated in a total required breath hold time of <1 s.

t_{delay} = [5, 28, 51, 74, 97, 120, 143, 166, 189, 212, 250, 308, 331, 8.5, 31.5, 54.5, 77.5, 100.5, 18.5, 41.5, 64.5, 87.5, 110.5] ms

Two different flip angles were needed for the accelerated CSSR sequence: 90° for dissolved-phase saturation (as in CSSR spectroscopy) and 10° for excitation. This was achieved by manually scaling the RF pulse amplifier power. The RF pulse described in Section 9.3.3 was used for both saturation and excitation, and the nominal flip angle was set to 90° on the scanner user interface as before. This was done so that all SAR and scaling calculations were performed using the highest flip angle in the sequence, to ensure patient safety. The 90° pulse used the maximum RF amplifier power and the 10° pulse was set to be one ninth of this value, i.e. amplifier linearity was assumed.

9.3.5 CSSR Spectroscopy Analysis

The CSSR spectroscopy analysis pipeline is shown in Figure 9.3.3.

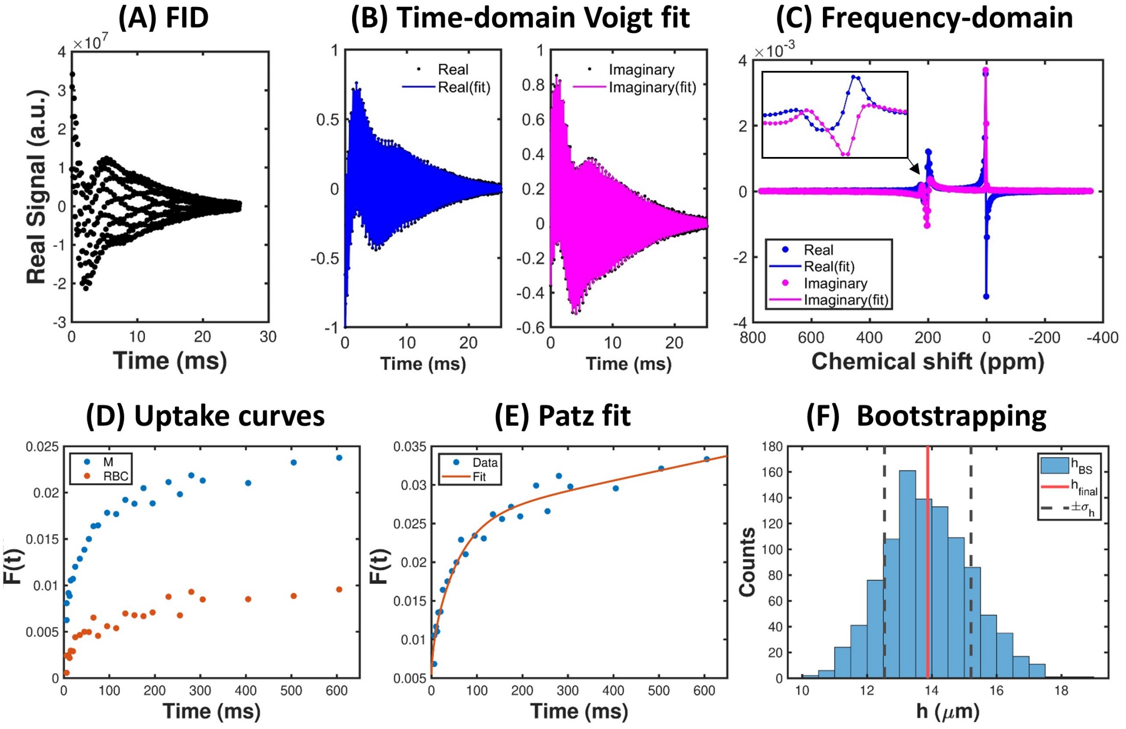

Figure 9.3.3: CSSR spectroscopy analysis pipeline: (A) the FIDs acquired at each t_{delay} were (B) fit to a Voigt lineshape. (C) The corresponding frequency domain fit. The inset shows the fit to the dissolved-phase peaks in more detail. (D) The RBC and M peak intensities from the Voigt fit were normalised by the gas peak intensity, scaled for flip angle differences and plotted against time as uptake curves. (E) The combined (RBC + M) dissolved-phase uptake curves were fitted to the Patz model. (F) The final parameters values and uncertainties were found from the distribution of values over 1000 fitted synthetic datasets.

At each t_{delay}, the acquired FID (Figure 9.3.3A) was normalised by its maximum absolute value and fitted to a time-domain Voigt function to quantify the gas, RBC and M signal intensities (Figure 9.3.3A). The Voigt lineshape ($V(t)$) is the product of a Lorentian lineshape ($L(t)$) and a Gaussian lineshape ($G(t)$), defined by:

$$Re[V(t)] = \sum_{n=1}^{3} A_n \exp(-\pi\Gamma_n^L t)\exp\left[-\left(\frac{\pi\Gamma_n^G t}{2\sqrt{\ln 2}}\right)^2\right]\cos[(\omega_{0,n}-\omega_t)t+\phi_n] \tag{9.3.3}$$

$$Im[V(t)] = \sum_{n=1}^{3} A_n \exp(-\pi\Gamma_n^L t)\exp\left[-\left(\frac{\pi\Gamma_n^G t}{2\sqrt{\ln 2}}\right)^2\right]\sin[(\omega_{0,n}-\omega_t)t+\phi_n] \tag{9.3.4}$$

where $n = 1, 2, 3$ represents the Gas, M and RBC resonances, A = FID amplitude/spectral peak intensity, $\omega_{0,n} - \omega_t$ is the frequency difference between the ^{129}Xe resonance frequencies and the RF transmit/receive frequency and ϕ is the peak phase. Γ^L and Γ^G are the Lorentzian and Gaussian linewidths, from which the proportion of Gaussian broadening, F_G, can be calculated:

$$F_G = \frac{\Gamma^G}{\Gamma^G + \Gamma^L} \text{x} 100\% \tag{9.3.5}$$

Figure 9.3.3C shows the corresponding frequency-domain fit. The CSSR uptake curves ($F(t)$) for the M and RBC signals were given by $A_M(t)/A_{Gas}(t)$ and $A_{RBC}(t)/A_{Gas}(t)$, respectively. Example $F(t)$ curves, plotted against t_{delay}, are shown in Figure 9.3.3. $F_M(t)$ and $F_{RBC}(t)$ were summed and corrected for flip angle differences between the gas and dissolved-phase, by multiplying $F(t)$ by:

$$c = \frac{\sin(0.9)}{\sin(90)} \tag{9.3.6}$$

For the accelerated CSSR sequence data, an additional correction factor was required to correct for the systematic bias resulting from the 10° excitation RF pulse in the Look-Locker acquisition scheme:

$$F_{LL}(t) = \frac{1}{\cos^{\mathrm{m}}(10)} \frac{\sin(\theta)}{\sin(10)} \frac{A_M(t) + A_{RBC}(t)}{A_{Gas}}, \tag{9.3.7}$$

where m is the index of the excitation pulse with respect to the preceding saturation pulse and A_{Gas} is the amplitude of the gas signal from the preceding saturation. The off-resonance flip angle was not well-characterised when the on-resonance flip angle of the RF pulse was scaled from 90° to 10°, because the RF pulse induced signal decay was minimal and the fit to Equation 9.3.2 unreliable (the fitted flip angle was dependent on the initial parameter input). The flip angle scaling was therefore found manually, by finding the value of θ which gave the best agreement between $F_{LL}(t)$ and $F(t)$ from standard CSSR , on a case-by-case basis.

The corrected $F(t)$ data were then fitted to the Patz model (Figure 9.3.3E). The Patz model [142], described in Section 3.4, is based on solving the 1D diffusion equation in a slab geometry with plug flow and is given by the following equation:

$$F(t) = F_0 + \frac{\lambda h}{2} \frac{S}{V} \left(\frac{\tau - t}{\tau} \right) f(q) + \lambda h \frac{S}{V} \left[\frac{t}{\tau} + \frac{8h^2}{D\pi^4 \tau} g(q) \right], \tag{9.3.8}$$

where $q = \frac{Dt}{h^2}$ and $f(q)$ and $g(q)$ were defined in Equations 3.4.4 and 3.4.5. All parameters were defined in Section 3.4. As a reminder: F_0 is an offset term, λ is the xenon Ostwald solubility, D is the dissolved ^{129}Xe diffusion coefficient, h is the septal wall thickness, S/V is the alveolar surface area to volume ratio and τ

is the ^{129}Xe capillary transit time. In this work, λ = 0.1 and D = 3.3 x 10^{-6} cm^{2}s^{-1} were used [93], although the true values of these parameters are not well known.

To increase the robustness of the Patz model fit and to reduce the sensitivity to the initial inputs to the fitting function, a multi-step fitting protocol was implemented, using a bounded fitting function (`fminsearchbnd`). To reduce the number of free parameters per fit, S/V and F_0 were found first by fitting the data points from the first 100 ms to the short-time Butler approximation [142]:

$$F(t) = F_0 + \lambda \frac{S}{V} \sqrt{\frac{4Dt}{\pi}}, \tag{9.3.9}$$

The initial inputs for S/V and F_0 in the fitting functions were found by running fits with a range of input values and choosing the values at the fit with the lowest residual sum of squares (RSS). All data points were then fitted to Equation 9.3.8, using the fitted values for S/V and F_0. The initial inputs for h and τ in the fit were also found by minimising the RSS over a range of input values.

Bootstrapping was used to obtain the final parameters and their uncertainties (Figure 9.3.3F). The residuals, R of the full Patz model fit were calculated and used to simulate 1000 synthetic data sets:

$$D_{synth} = f(t, P) + R_{synth}, \tag{9.3.10}$$

where $f(t, P)$ is the Patz model with parameters $P = (F_0, S/V, h, \tau)$ and R_{synth} are residuals chosen at random (with replacement) from R. Each synthetic data set was then fit to the Patz model and the fit parameters stored in an array, P_{boot}. According to Central Limit Theorem, P_{boot} is normally distributed and hence the final parameter values and uncertainties were given by the mean and standard deviation of P_{boot}.

9.3.6 CSSR Imaging

CSSR imaging was implemented using a low-bandwidth 2D gradient echo sequence, such that dissolved-phase and gas-phase ^{129}Xe images appeared within the same FOV due to the chemical shift artefact, as described in Section 9.2.

The 2D gradient echo pulse sequence used for CSSR imaging is shown in Figure 9.3.4. CSSR with seven t_{delay} values (6 ms, 15 ms, 45 ms, 95 ms, 175 ms, 280 ms and 405 ms) was repeated for every phase-encoding step. This choice of t_{delay} values was made based on retrospectively undersampling the CSSR spectroscopy uptake curve for one healthy volunteer dataset, fitting to the Patz model and comparing the agreement with the parameters from the fully-sampled uptake curve. Using Equation 9.2.1, the bandwidth required for complete separation of the gas-phase and dissolved-phase components in the frequency-encoding direction was found to be 7 kHz, with a pixel size of 1.25 cm along the frequency-encoding direction and assuming a lung width of 25 cm. A fully-sampled Cartesian k-space trajectory was used, with centric ordering. The gradient waveforms were designed using tools from the MNS research pack in MATLAB. The sequence parameters were: 1 coronal projection covering the whole lung, FOV = 50 cm x 50 cm, matrix size = 40 x 20, readout direction = right to left, TR/TE = 20 ms/3.1 ms and dissolved-phase/gas-phase flip angle = 90°/0.9°. The ^{129}Xe dose was 1 L and the breath hold time was 24 s.

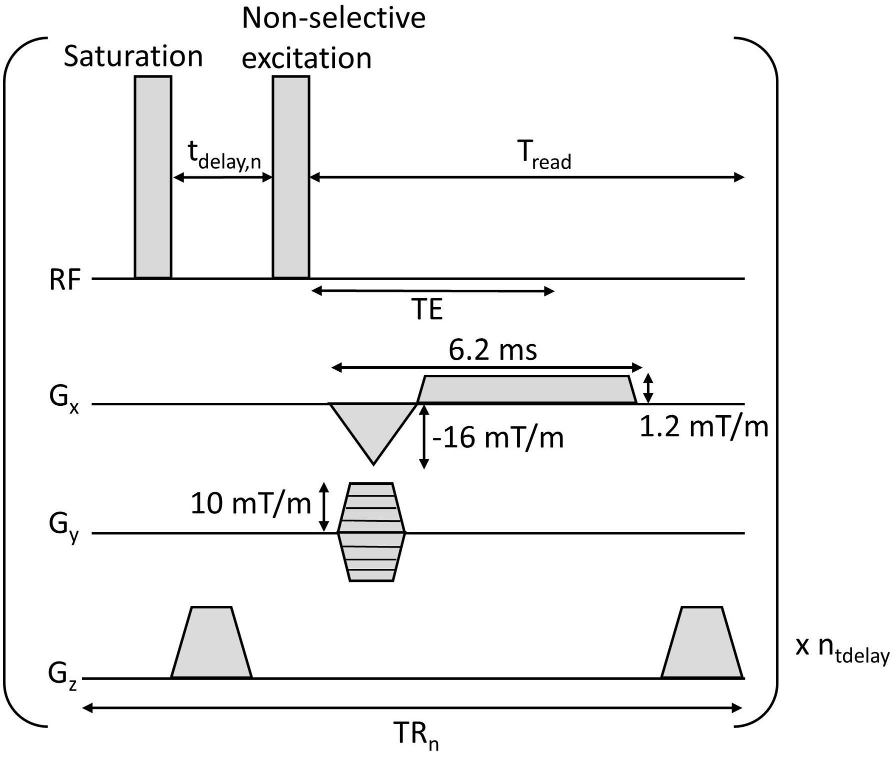

Figure 9.3.4: Pulse sequence diagram for the CSSR imaging sequence. Each CSSR 'block' was repeated for every phase-encoding step; i.e. each t_{delay} was repeated 20 times. To reduce the number of 90° RF pulses, only one saturation RF pulse was used per CSSR block. Centric Cartesian phase encoding was used, with a maximum phase-encoding gradient strength of 10 mTm^{-1}. Spoiler gradients were used to destroy the transverse magnetisation both after the saturation pulse and at the end of each TR.

9.3.7 CSSR Image Analysis

For each t_{delay}, the k-space data were reconstructed using the inverse Fourier transform. The gas-phase signal was masked via thresholding; this mask was shifted to the right to mask the dissolved-phase signal. The masked dissolved-phase image ($S_{DP}(x,y,t)$) was normalised by the masked gas-phase image ($S_{GP}(x,y,t)$) and corrected for dissolved-phase/gas-phase flip angle differences and the T_2^* decay that occurred over the readout time:

$$F_{image}(x,y,t) = c\frac{S_{DP}(x,y,t)}{S_{GP}(x,y,t)}\exp\left(\frac{TE}{T_2^*}\right) \tag{9.3.11}$$

where c is the flip angle scaling factor given in Equation 9.3.6, TE = 3.1 ms (measured from the end of the excitation pulse to the centre of the read-out gradient) and T_2^* = 2.2 ms for the dissolved-phase signal at 1.5 T [36]. Each pixel of the normalised image was then fit to the Patz model, using the fitting procedure described in Section 9.3.5, to obtain parameter and error maps for h, S/V and τ.

9.3.8 Statistical Analysis

Normality of variables was tested for with Shapiro-Wilk tests. Differences between the CSSR parameters between the healthy volunteers and patients were tested for using Student's t-test. Correlations between variables were evaluated using Pearson correlation coefficient for normally distributed variables and Spearman's correlation coefficient for non-normal variables. A significance level of $p<0.05$ was used for all tests.

9.4 Results

9.4.1 RF Pulse Design

Figure 9.4.1 shows the signal decay when the frequency-selective RF pulse was repeatedly applied to a bag of hyperpolarised ^{129}Xe gas. Using Equation 9.3.2, the on-resonance flip angle was found to be 87° and the off-resonance flip angle was ∼0.7° (assuming negligible T_1 decay). Very good saturation performance was also found in vivo. As shown in Figure 9.4.2A, one RF pulse was sufficient to almost entirely destroy the gaseous ^{129}Xe magnetisation when applied on-resonance. The calibrated on-resonance flip angle was 89.7°. To ensure complete dissolved-phase signal saturation, two saturation pulses were used per CSSR block, in line with previous work [325]. However, when we compared the results of CSSR spectroscopy results using one and two saturation pulses, there was no difference in the uptake curves, as demonstrated in Figure 9.4.3.

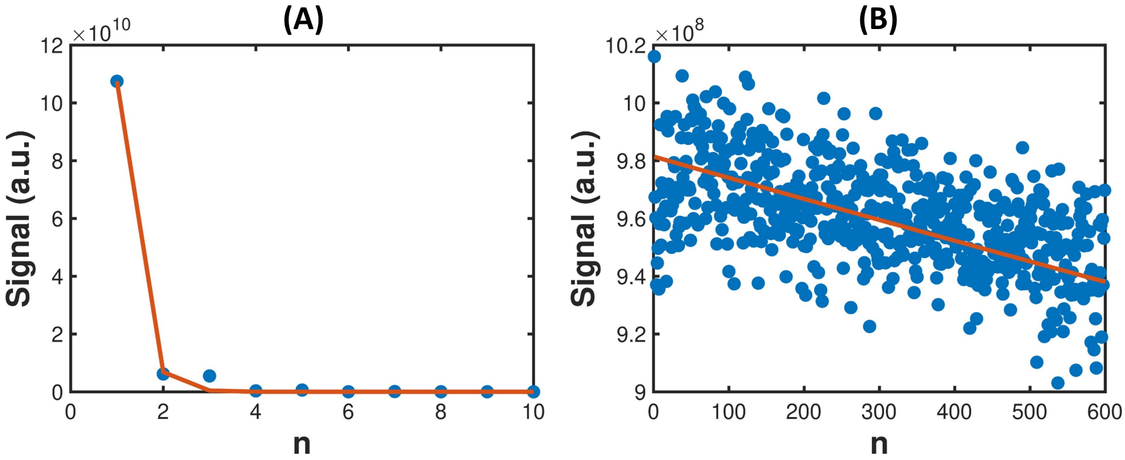

Figure 9.4.1: Gaseous ^{129}Xe signal decay as a function of RF pulse number acquired from a bag of hyperpolarised ^{129}Xe gas when the RF power was centred (A) on the gas-phase resonance and (B) 3650 Hz downstream of the gas-phase resonance.

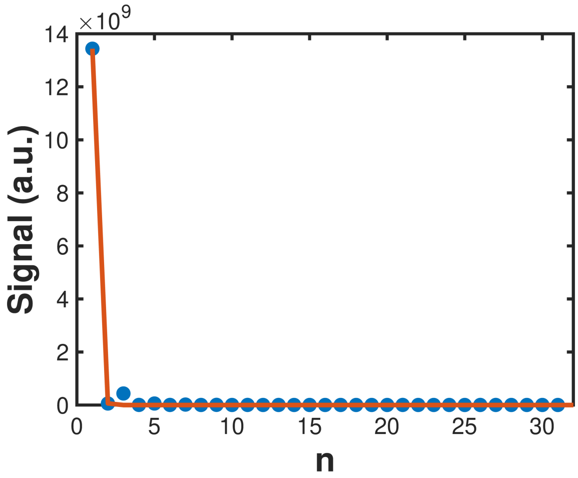

Figure 9.4.2: On-resonance ^{129}Xe signal decay as a function of RF pulse number acquired in a healthy volunteer.

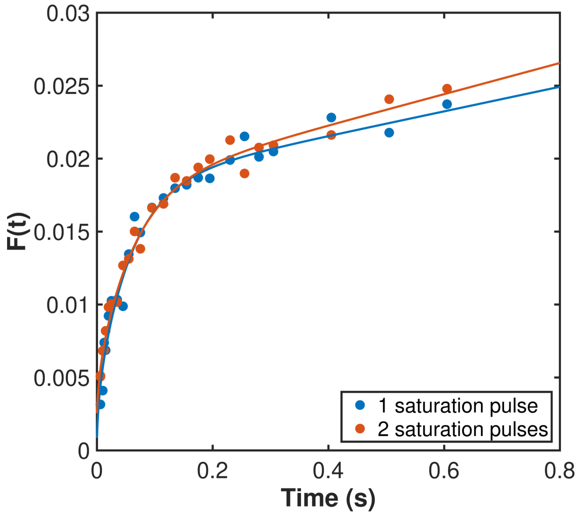

Figure 9.4.3: CSSR uptake curves acquired in the same healthy volunteer using one saturation pulse per CSSR block (blue) and two saturation pulses per block (red).

9.4.2 CSSR Spectroscopy

Subject demographics and CSSR results are shown in Table 9.4.1. Due to the small number of non-SSc PAH patients, these patients were combined with the SSc-PAH patients in order to make comparisons with the healthy volunteer group. Boxplots comparing the values of h, S/V, τ, and their associated fitting errors between the two groups are shown in Figure 9.4.4. h was significantly higher (p = 0.01) in the SSc-PAH and PAH patients (h = 17.2 ± 2.2 μm) than in the healthy volunteers (14.5 ± 0.9 μm), but no significant differences were found between the groups for S/V or τ. The mean percentage parameter error was higher for τ (28%) than for h (10%) or S/V (8%) and in general, errors were greater for the patients than for the healthy volunteers, although not significantly. Overall, Gaussian broadening was found to contribute around 40 - 60% to the gas, membrane and RBC resonance lineshapes. No differences in the proportion of Gaussian broadening were found between the healthy volunteers and the SSc/PAH patients.

Table 9.4.1: Subject demographics and Patz model metrics from CSSR spectroscopy.

	Healthy	All Patients	SSc-PAH	PAH
n (female)	9 (4)	8 (6)	6 (5)	2 (1)
Age	30 ± 6	65 ± 7	64 ± 7	68 ± 8
h (μm)	14.5 ± 0.9	17.2 ± 2.2	16.6 ± 2.2	18.7 ± 1.5
h error (μm)	1.2 ± 0.2	2.0 ± 0.9	1.9 ± 0.9	2.2 ± 1.5
S/V (cm^{-1})	278 ± 52	287 ± 115	281 ± 123	304 ± 128
S/V error (cm^{-1})	22 ± 4	19 ± 6	19 ± 8	19 ± 1
τ (s)	1.91 ± 0.59	1.85 ± 0.66	1.68 ± 0.45	2.34 ± 1.17
τ error (s)	0.44 (0.25 - 1.11)	0.56 ± 0.23	0.49 ± 0.16	0.78 ± 0.35
$F_{G,Gas}$	61 ± 5	58 ± 6	56 ± 7	64 ± 4
$F_{G,M}$	37 ± 4	38 ± 5	40 ± 4	30 ± 2
$F_{G,RBC}$	62 ± 7	66 ± 5	66 ± 5	66 ± 9

Figure 9.4.5 shows example CSSR spectra at t_{delay} = 230 ms for a healthy participant and three patients and their corresponding uptake curves. The amplitude of the RBC peak was reduced in all patients; this reduction in signal is particularly apparent in patient SSc-PAH1 (Figure 9.4.5Aii). Both PAH patients and some of the SSc-PAH patients also had an increased relative amplitude of the M peak, reflecting the alveolar septal thickening measured by h. In these subjects, the $F(t)$ uptake curves had a greater amplitude and saturated at a later time than the healthy volunteers, as shown in Figure 9.4.5.

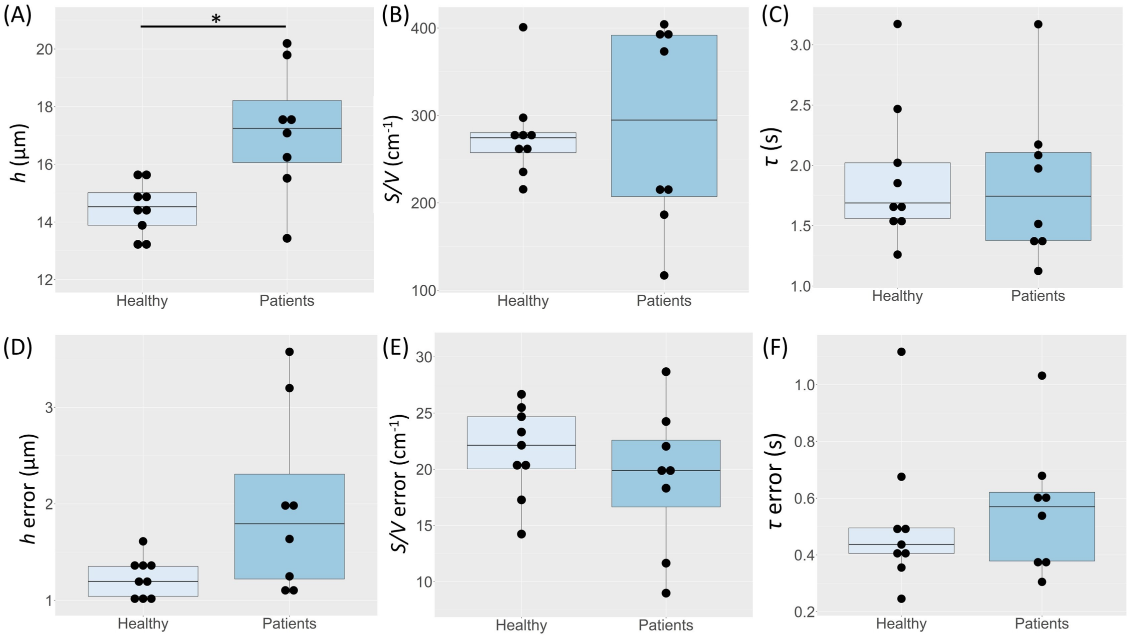

Figure 9.4.4: Comparison of the CSSR parameters derived from the Patz model (A-C) and their errors (D - F) between healthy volunteers and patients with SSc and/or PAH.

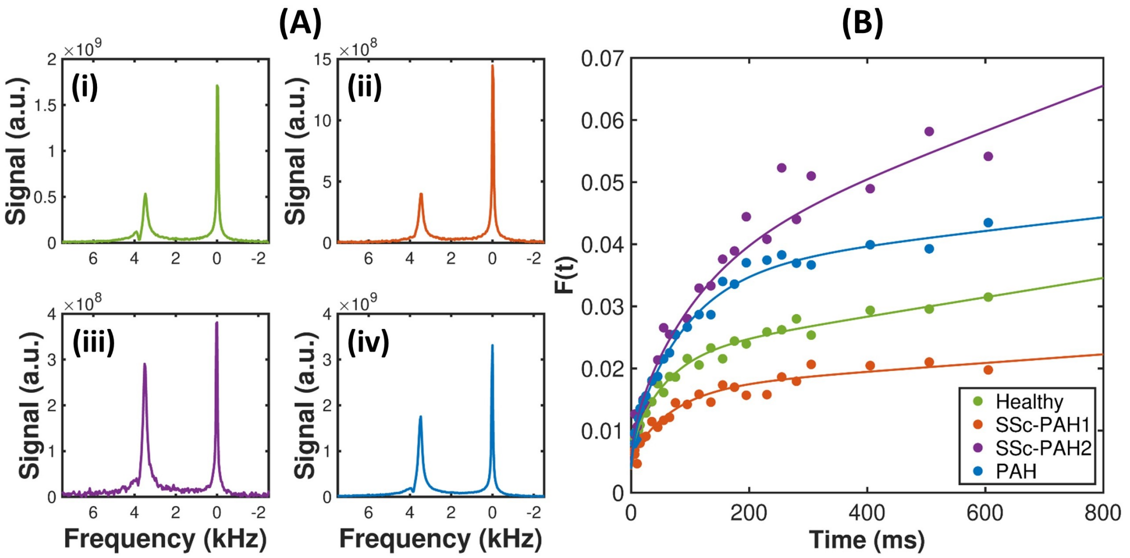

Figure 9.4.5: (A) CSSR spectra at t_{delay} = 230ms for (i) a healthy volunteer (HV6), (ii) an SSc-PAH patient (SSc-PAH1), (iii) a second SSc-PAH patient (SSc-PAH2) and (iv) a PAH patient (PAH2) and (B) their uptake curves.

9.4.3 CSSR Imaging

Gas-phase and dissolved-phase ^{129}Xe images acquired simultaneously with the 2D gradient echo sequence at each t_{delay} time are shown for one healthy volunteer in Figure 9.4.6. The SNR decreased with increasing t_{delay}, as shown in Figure 9.4.6H. The final image (t_{delay} = 405 ms) was discarded for two of the CSSR imaging cases (HV6 and HV7) due to low SNR.

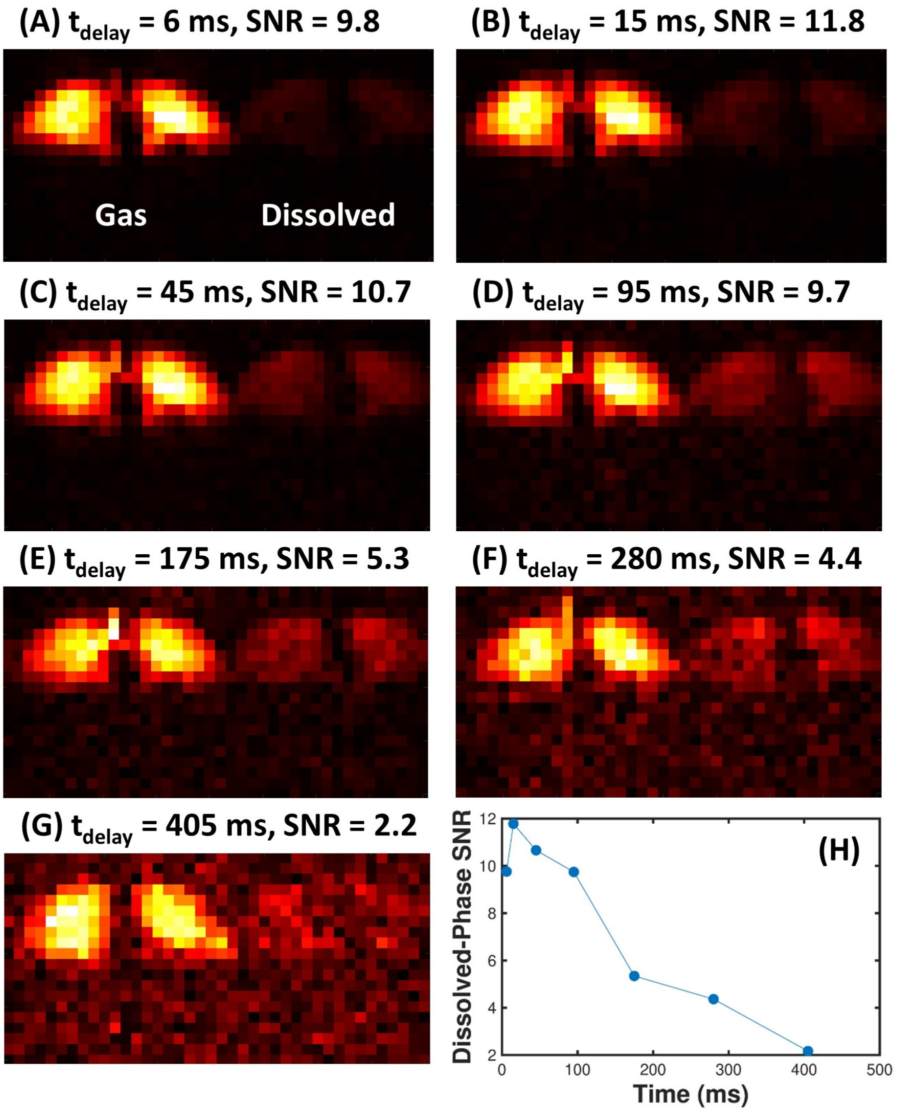

Figure 9.4.6: (A-G) Gas-phase and dissolved-phase ^{129}Xe images acquired with the CSSR imaging sequence acquired at each t_{delay} for a healthy volunteer and (H) SNR of the dissolved-phase image plotted against t_{delay}.

The choice of t_{delay} values for CSSR imaging was made based on retrospectively removing $F(t)$ points from the uptake curve of one of the healthy volunteers (HV6) and comparing the resulting Patz model fit with the fully-sampled Patz fit. Table 9.4.2 shows the results of this process. With the seven delay times used in the CSSR imaging sequence there was very good agreement between the parameters. However, when the data point at t_{delay} = 405 ms was removed, as was the case for the first implementation of the CSSR imaging sequence due to the low SNR, there was worse agreement. In particular, τ was underestimated by ∼60%. If instead, six t_{delay} times were used but the data point at 280 ms was removed rather than the 405 ms data point, better agreement was achieved and the fitted uptake curve appeared more similar to the uptake curve with seven t_{delay} values (Figure 9.4.7).

It was not possible to acquire more than six usable t_{delay} images because of the SNR decrease towards the end of the sequence. For this reason, the t_{delay} = 280 ms images were not acquired for HV8 or HV9, reducing the breath hold time to 17 s. Figure 9.4.8 shows the combined gas-phase and dissolved-phase images for HV8 and the normalised dissolved-phase signal at each each t_{delay} time. The SNR for the t_{delay} = 405 ms image was sufficient (SNR = 6.0) for inclusion in the Patz model fitting.

Table 9.4.2: Patz model fit parameters for the fully-sampled and retrospectively undersampled CSSR uptake curves for one healthy participant. For the CSSR uptake curve with n_{tdelay} = 6, two combinations of t_{delay} times were tested, with either 280 ms or 405 ms as the final t_{delay}.

n_{tdelay}	h (μm)	S/V (cm^{-1})	τ (s)
25	14.3 ± 1.2	274 ± 21	1.26 ± 0.25
7	14.4 ± 1.4	284 ± 22	1.30 ± 0.38
6 (280 ms)	13.0 ± 1.7	285 ± 26	0.72 ± 0.21
6 (405 ms)	13.9 ± 1.4	283 ± 21	1.16 ± 0.30

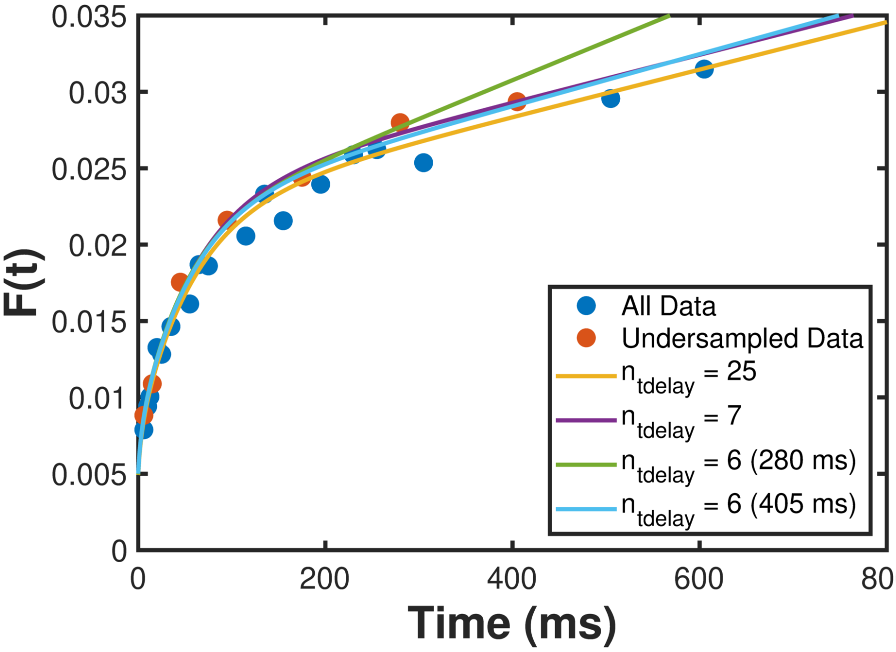

Figure 9.4.7: CSSR uptake curve and Patz model fits found from undersampling the CSSR data. n_{tdelay} is the number of data points included in the Patz model fit.

The parameter and error maps for one of the healthy volunteers (HV8) are shown in Figure 9.4.9. The h and S/V maps were relatively uniform for all volunteers, with mean CV of 0.27 ± 0.07 and 0.17 ± 0.06, respectively. The error maps for h and S/V exhibited increased heterogeneity when compared with the parameter maps (CV = 0.49 ± 0.12 for h error and 0.59 ± 0.21 for S/V error). The τ

and τ error maps were much less uniform (CV = 1.01 ± 0.29 and 1.09 ± 0.28, respectively).

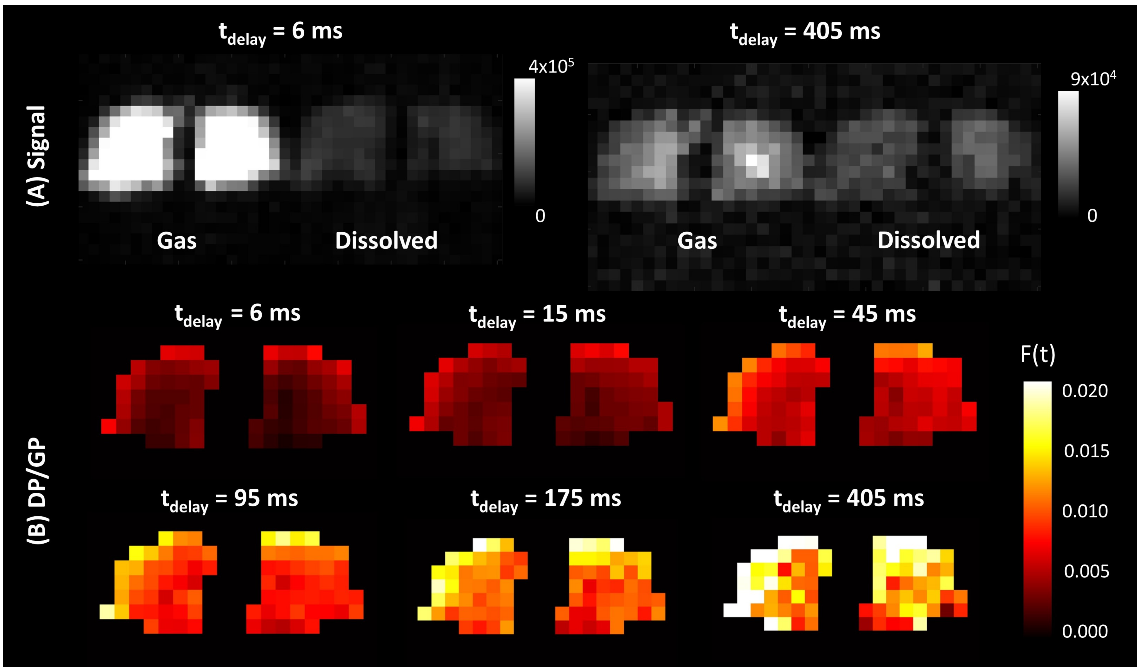

Figure 9.4.8: (A) ^{129}Xe gas-phase (GP) and dissolved-phase (DP) images for HV8, acquired simultaneously with a 2D gradient echo sequence, for t_{delay} = 6 ms and t_{delay} = 405 ms. (B) Normalised dissolved-phase images at each t_{delay}, showing the uptake of ^{129}Xe to the alveolar membrane and RBCs over time.

Table 9.4.3 compares the Patz model parameters from spectroscopy and imaging for each of the healthy volunteers who underwent both scans. As shown by the Bland-Altman plot in Figure 9.4.10, CSSR imaging tended to overestimate h but underestimate S/V and τ. The absolute fitting errors were larger for imaging than spectroscopy for h and S/V, but smaller for τ because the values of τ were also reduced. The percentage error for τ was larger for imaging (34% on average across the four healthy volunteers) than for spectroscopy (23%).

In three out of four of the healthy volunteers who underwent CSSR imaging, the uptake curves did not show good agreement with those from CSSR spectroscopy (Figure 9.4.11). With the exception of HV6, the uptake curves from CSSR imaging diverged from those from spectroscopy as t_{delay} increased. However, running the CSSR imaging sequence with the phase and frequency encoding gradients set to zero in one healthy volunteer resulted in an uptake curve that showed good agreement with CSSR spectroscopy, but not imaging (Figure 9.4.12).

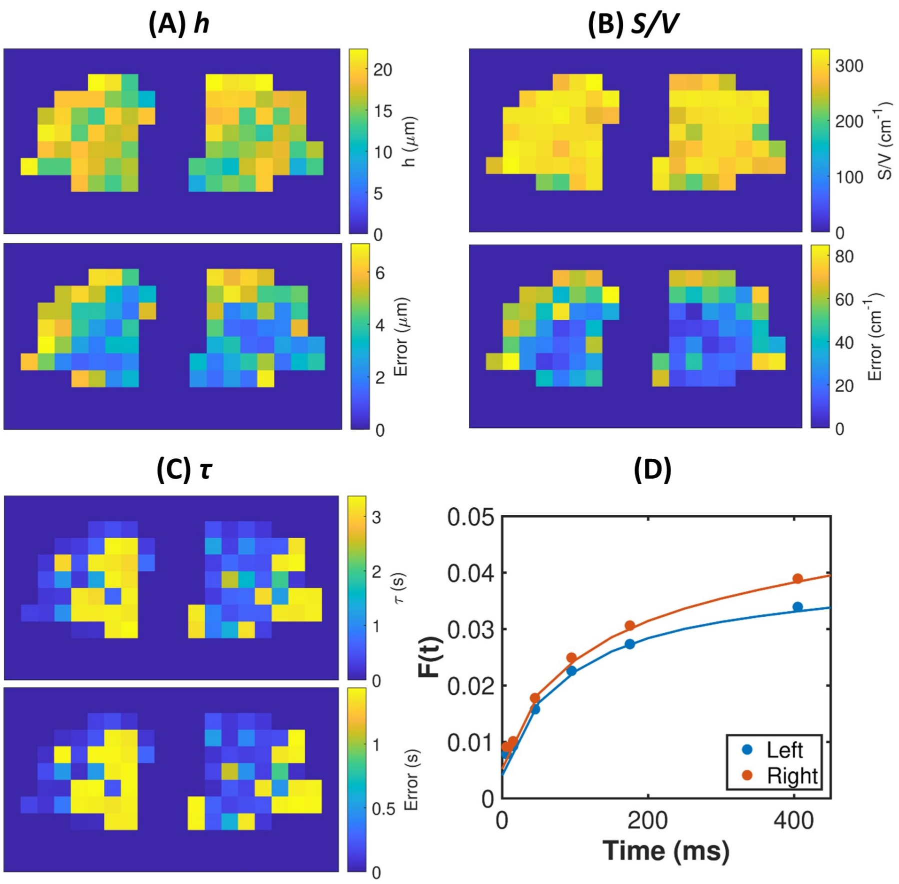

Figure 9.4.9: Parameter and error maps for (A) *h*, (B) *S/V* and (C) from the regional CSSR Patz model fit in a healthy participant (HV8). (D) Gas uptake curves derived from the average normalised DP signal in the left and right lungs.

Table 9.4.3: Comparison of the Patz model parameters between CSSR spectroscopy and imaging. For imaging, the average values of the parameter and uncertainty maps are given as mean ± standard deviation for normally distributed maps and median (range) for non-normally distributed maps.

		h (μm)	S/V (cm^{-1})	τ (s)
HV6	Spectroscopy	14.3 ± 1.2	274 ± 21	1.26 ± 0.25
	Imaging	14.0 ± 4.8	254 (77 - 329)	0.68 (0.04 - 3.29)
	Imaging Error	3.2 (0.3 - 6.9)	40 ± 20	0.26 (0.01 - 1.47)
HV7	Spectroscopy	15.7 ± 1.0	401 ± 20	2.02 ± 0.50
	Imaging	17.7 (1.1 - 23.2)	304 (180 - 331)	0.23 (0.01 - 3.29)
	Imaging Error	5.3 (0.1 - 7.8)	57 ± 22	0.08 (0.00 - 1.47)
HV8	Spectroscopy	14.5 ± 1.4	280 ± 23	1.69 ± 0.44
	Imaging	16.8 ± 3.2	301 (200 - 328)	0.87 (0.14 - 3.37)
	Imaging Error	3.7 (1.3 - 7.1)	34 (4 - 85)	0.27 (0.04 - 1.44)
HV9	Spectroscopy	14.7 ± 1.0	235 ± 14	1.62 ± 0.36
	Imaging	15.0 (0.8 - 21.7)	241 (97 - 328)	1.01 (0.01 - 3.34)
	Imaging Error	2.5 (0.1 - 7.2)	18 (1 - 107)	0.31 (0.00 - 1.5)

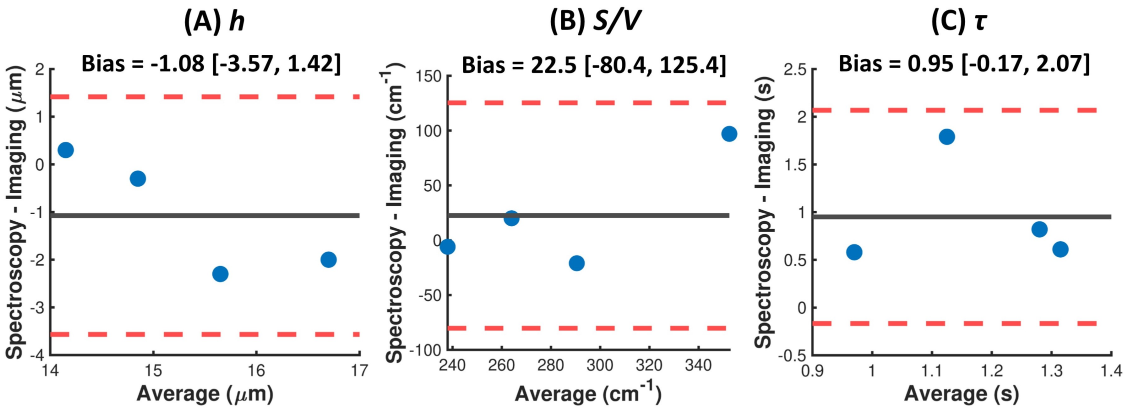

Figure 9.4.10: Bland-Altman plots for (A) h, (B) S/V and (C) τ, from CSSR spectroscopy and imaging. The mean bias (black line) and limits of agreement (red lines) are given above each plot.

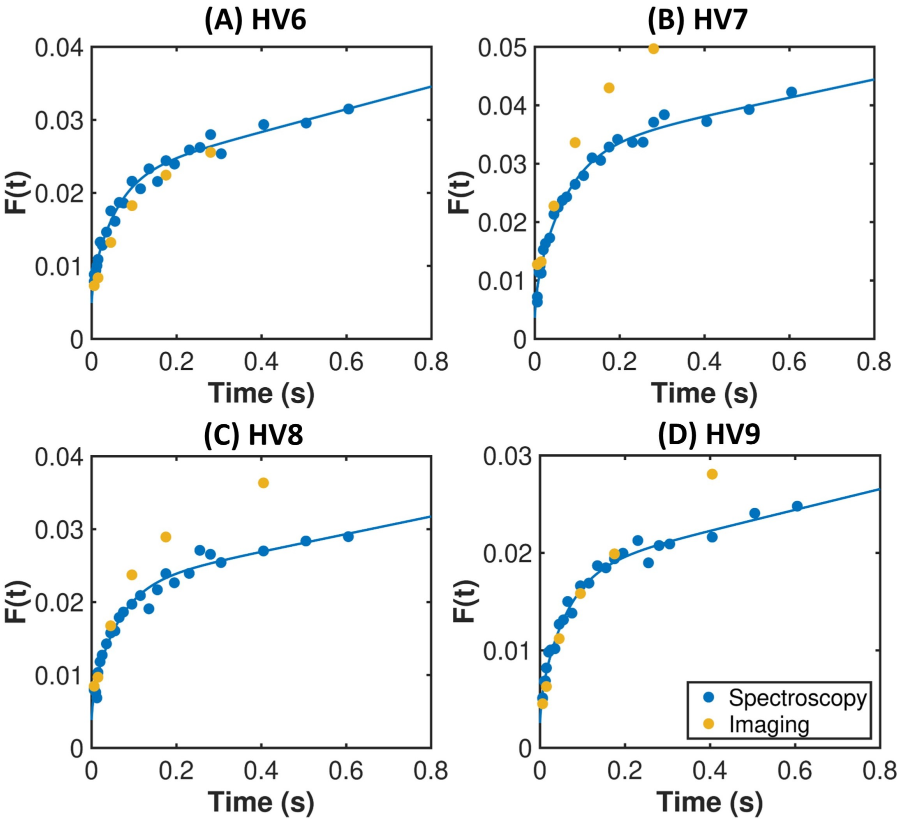

Figure 9.4.11: CSSR uptake curves from spectroscopy and imaging (whole-lung average) for each of the four healthy volunteers who underwent CSSR imaging.

9.4.4 Accelerated CSSR Spectroscopy

A solution to some of the problems of the CSSR imaging sequence, namely the large number of 90° pulses and the limited number of t_{delay} images achievable, may be found by using a Look-Locker acquisition strategy. Before attempting to implement a Look-Locker CSSR imaging sequence, an accelerated spectroscopy sequence was investigated using modelling and was tested in four healthy volunteers. The scan time for this sequence was decreased seven-fold compared to

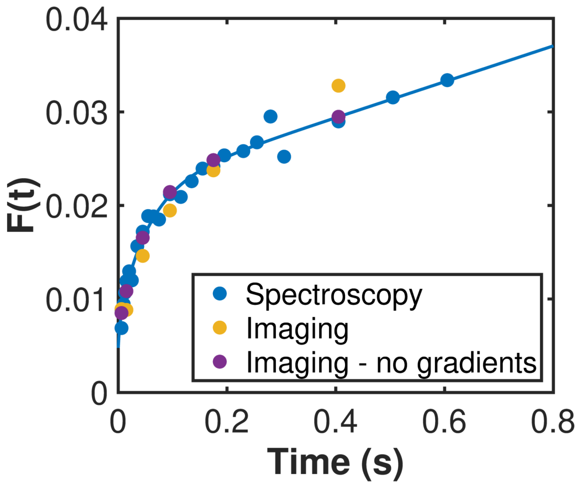

Figure 9.4.12: CSSR uptake curves obtained from a healthy volunteer using the CSSR spectroscopy sequence (blue), CSSR imaging (yellow) and the CSSR imaging sequence with the frequency and phase encoding gradients nulled (purple).

the standard CSSR spectroscopy sequence (from 5 s to 0.7 s).

Simulations using a 2D finite difference model demonstrated a small systematic bias of 3% in the magnetisation, resulting from the Look-Locker acquisition scheme (Figure 9.4.13). This was corrected for in the experimental uptake curves by dividing the m^{th} data point after each saturation by $\cos^{m}(10)$.

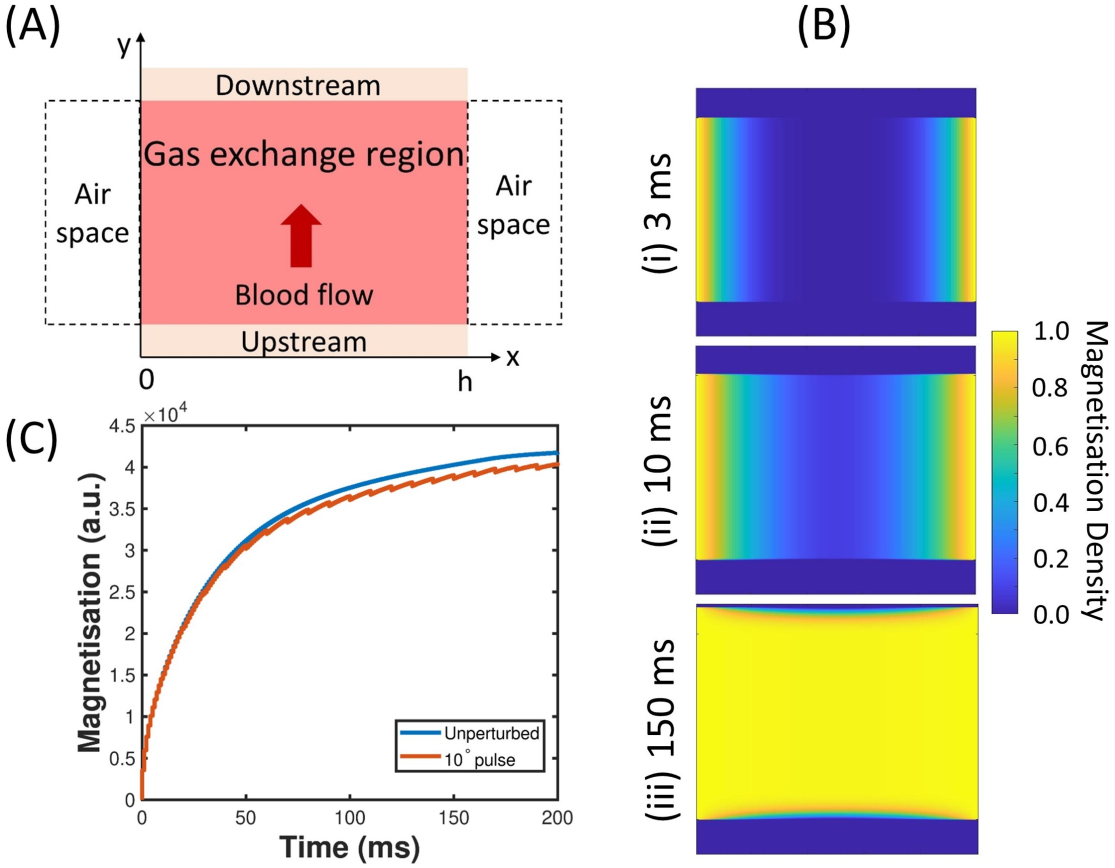

Figure 9.4.13: Finite difference modelling of ^{129}Xe gas transfer from the alveolar airspace into the bloodstream under the influence of small flip angle excitation pulses. (A) The 2D slab geometry used for the model: ^{129}Xe diffuses from the airspace and into the gas exchange region of the capillary, where diffusion takes place. Upstream blood flows through the gas exchange region and onwards downstream, removing ^{129}Xe from this region. (B) ^{129}Xe magnetisation density within the 2D model at three time points: (i) at 3 ms, ^{129}Xe has started to diffuse into the capillary, but all magnetisation density is contained within the gas exchange region, (ii) at 10 ms, fresh blood containing no dissolved ^{129}Xe magnetisation has begun to flow into the gas exchange region from upstream and, downstream, some ^{129}Xe has started to flow out of the gas exchange region, (iii) after 150 ms, the gas exchange region has become saturated with ^{129}Xe and more ^{129}Xe has been transported out of the gas exchange region via the flow of blood. (C) Simulated magnetisation evolution from the finite difference model with and without small flip angle excitation pulses.

Figure 9.4.14A shows the spectra at 250 ms for a healthy volunteer (HV4) from the Look-Locker CSSR spectroscopy sequence, demonstrating minimal excitation of the gas-phase peak and good SNR, despite the low flip angle excitation. Good agreement was found between the Look-Locker CSSR and standard CSSR uptake curves using an off-resonance flip angle of ∼0.035° in the flip angle scaling (Figure 9.4.14B).

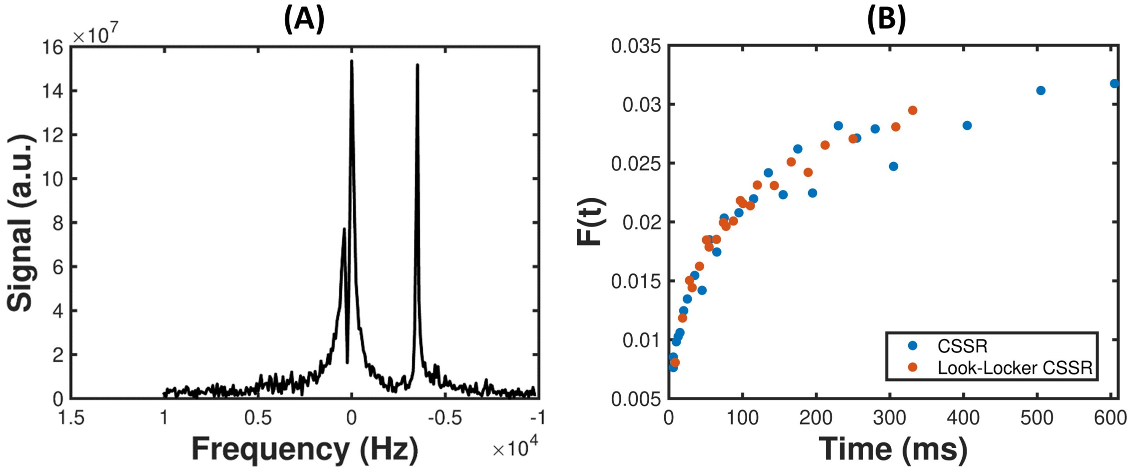

Figure 9.4.14: (A) CSSR spectra at 250 ms, acquired with the accelerated Look-Locker CSSR sequence for a healthy volunteer and (B), their uptake curves from standard CSSR and Look-Locker CSSR. The Look-Locker CSSR uptake curve was manually scaled to overlay with the standard CSSR uptake curve.

Table 9.4.4 compares the Patz model parameters from the standard CSSR and accelerated CSSR sequences for each subject. There was reasonable agreement between *h*, *S/V* and τ between the two sequences, with Bland-Altman mean biases of -0.70 µm, 21.5 cm^{-1} and 0.38 s respectively. As shown in the Bland-Altman plots in Figure 9.4.15, the residuals were evenly distributed, although *h* tended to be overestimated and *S/V* and τ tended to be underestimated for Look-Locker CSSR.

Table 9.4.4: Comparison of the Patz model parameters between the standard CSSR spectroscopy sequence and the accelerated Look-Locker sequence for four healthy volunteers.

	CSSR			Look-Locker CSSR		
Subject ID	*h* (µm)	*S/V* (cm^{-1})	τ (s)	*h* (µm)	*S/V* (cm^{-1})	τ (s)
HV2	13.2 ± 1.0	266 ± 22	2.47 ± 0.68	13.6 ± 1.6	229 ± 24	1.06 ± 0.33
HV3	13.3 ± 1.4	257 ± 27	1.56 ± 0.41	14.3 ± 2.3	236 ± 31	1.09 ± 0.40
HV4	15.0 ± 1.6	279 ± 25	1.52 ± 0.41	14.3 ± 1.3	283 ± 22	1.07 ± 0.31
HV6	14.3 ± 1.2	274 ± 21	1.26 ± 0.25	16.4 ± 1.5	242 ± 18	2.06 ± 0.85

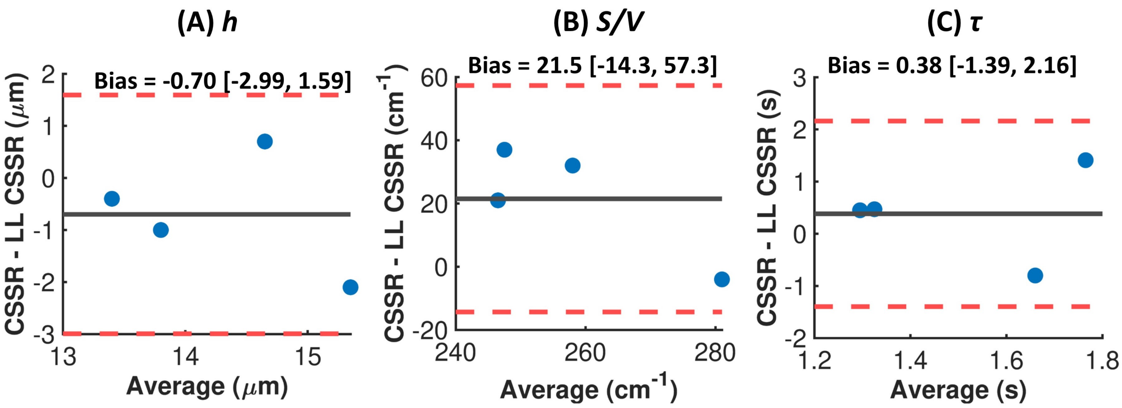

Figure 9.4.15: Bland-Altman plots for (A) h, (B) S/V and (C) τ, from standard CSSR and Look-Locker CSSR spectroscopy. The mean bias (black line) and limits of agreement (red lines) are given above each plot.

9.5 Discussion

In this chapter, we demonstrated several methods for quantifying temporal gas exchange in the lungs using ^{129}Xe spectroscopy. We developed an improved analysis pipeline for reconstructing and fitting the CSSR uptake data. In most previous literature, a Lorentzian lineshape has been used to fit the CSSR spectral peaks in the frequency-domain. However, the spectral resolution in the frequency-domain is low for high bandwidths. This can lead to inaccurate fitting, especially for the gas-phase resonance, which typically has a width which is narrower than the spectral resolution. Fitting the CSSR data in the time-domain is therefore advantageous. Additionally, the Voigt line function is easier to implement in the time-domain, where it is a simple multiplication of the Lorentzian and Gaussian line functions, in comparison to the frequency-domain, where it is a convolution of the two functions. Improved fitting of the ^{129}Xe spectral peaks can be achieved with a Voigt lineshape because the lineshapes are not purely Lorentzian. There is known to be substantial Gaussian broadening to the ^{129}Xe peaks from the changes in magnetic susceptibility at the boundaries between the different chemical environments within the lung [137]. We found a Gaussian contribution of 40% and 60% to the ^{129}Xe M and gas linewidths, respectively, which is in line with the values of ∼45% and ∼75% found by Norquay et al. [326]. We also found ∼65% Gaussian broadening of the RBC peak, which was assumed to be purely Lorentzian in [326]. Including a Gaussian component within the spectral fitting function should improve its accuracy compared to using a triple Lorentzian fit. However this does also increase the number of degrees of freedom in the fitting function.

In the CSSR literature, the errors on the parameters derived from diffusion models are often not quoted. However, knowledge of the uncertainties on the values of h, S/V and τ is important for the clinical translation of the technique, especially for quantifying fibrosis and/or inflammation in patients with lung disease. Without this information, we cannot know how much confidence can be placed on the CSSR metrics. Additionally, the CSSR uptake curves are often noisy and can be affected by cardiogenic oscillations in the RBC signal [212]. We observed oscillations in several CSSR uptake curves. In some cases, for example the healthy volunteer shown in Figure 9.4.14, these oscillations were only present in the RBC signal (A_{RBC}/A_{Gas}) and not the membrane signal (A_M/A_{Gas}), so are likely to be related to RBC signal variation over the cardiac cycle. However, in other cases, for example in Figure 9.3.3, the oscillations are more apparent in the membrane signal. This may be a result of a systematic offset between the first and second sweep of t_{delay} values, possibly associated with decay of the gas-phase signal over the duration of the sequence. Note: the ordering of the t_{delay} values was given in Section 9.3.4.

In this work, bootstrapping was used to estimate the uncertainty on the Patz model-derived CSSR parameters. The errors on h and S/V were found to be $\leq$10% over all subjects, whereas τ had a greater average error of $\sim$30%. Sensitivity analysis has previously shown that the Patz model is less sensitive to variations in τ than it is to h and S/V [327]. A larger range of τ values may yield similar fitting accuracy, increasing the uncertainty on this parameter.

Table 9.5.1 summarises the values of h, S/V and τ from the literature. The values of h found for the healthy volunteers were similar to micro-CT measurements of the alveolar septal wall thickness [328], but were higher than previous estimates from CSSR. S/V was also higher than in previous work, but in line with CT-based measurements [329]. Discrepancies between our results and the CSSR literature may be due in part to differences in experimental methodology and analysis procedure. There is not currently a standardised sequence for CSSR and so results may be dependent on the acquisition parameters used, such as the choice of delay times. Additionally, the diffusion coefficient and Ostwald solubility for dissolved ^{129}Xe used in the Patz and MOXE CSSR fitting models are not well-known. Different values have been used in the literature; for example, we used λ = 0.1, but Chang et al. used λ = 0.2 [330]. We used a value of 3.3 x 10^{-10}m^{2}s^{-1} for the ^{129}Xe diffusion coefficient, but this value originates from a single animal study [93] which derived this value using an assumed septal thickness and has not yet been validated either in humans or in vitro. The diffusion coefficient is assumed to be the same in the pulmonary tissue and blood, however ^{129}Xe diffuses faster in liquid solution than organic polymers by up to three orders of magnitude [331, 332], so this assumption might not be correct. The development of physiologically realistic dissolved-phase ^{129}Xe phantoms by Filkins et al. [333] and Willmering et al. [334] may facilitate better characterisation of the dissolved ^{129}Xe diffusion coefficients.

Factors such as lung inflation and participant age also affect CSSR metrics [141]. At TLC, h and S/V are reduced due to alveolar distension and associated stretching of the membrane layer [335]. h has been found to be increased in older people due to age-related septal wall thickening [39].

Although the values we found for h were larger than other CSSR estimates, this parameter was still sensitive to differences between the healthy volunteers and lung disease patients, for whom h was significantly larger. Our results are similar to Stewart et al. [39], who also found septal thickening in patients with SSc using CSSR spectroscopy. 80% of SSc patients develop lung fibrosis [336], and so the increase in h could be attributed to fibrosis of the alveolar membrane. CSSR has not previously been used to assess gas exchange in patients with

Table 9.5.1: Lung parameters found from the Patz model fits and their comparison with literature values. In Ref. [330] the parameters were found using the MOXE. S/V and τ were not provided in Ref. [38]. For Ref. [147], the healthy volunteers were separated into two groups: 'healthy young' (HY) and age-matched controls (AMC), and τ was not reported.

Work	n	Age (years)	Lung inflation	h (μm)	S/V (cm^{-1})	τ (s)
This work	9	30 ± 6	FRC + bag	14.5 ± 0.9	278 ± 52	1.91 ± 0.59
Patz, 2011 [142]	3	Unknown	'near FRC'	12.5 ± 1.5	118 ± 58	1.5
Chang, 2014 [330]	10	34 ± 18	FRC + bag	9.2 ± 6.5	210 ± 50	1.3 ± 0.3
Qing, 2014 [38]	13	26 ± 9	TLC	7.1 ± 0.9	-	-
Stewart, 2015 [39]	10	45 ± 16	FRC + bag	9.7 ± 1.5	128 ± 36	2.5 ± 0.7
Stewart, 2017 [141]	9	59 ± 8	FRC + bag	11.6 ± 1.0	194 ± 70	-
Xie, 2019 [147]	9	HY: 26 ± 2, AMC: 66 ± 7	FRC + bag	HY: 8.8 ± 0.9, AMC: 10.6 ± 2.3	HY: 196 ± 18, AMC: 133 ± 18	-

PAH. These preliminary results revealed increased septal thickness in both PAH patients (19.8 ± 3.2 μm and 17.7 ± 1.1 μm), which may be due to chronic inflammation of the small vessels [337]. There were also differences in the CSSR spectra between the healthy volunteers and patients. Notably, the RBC peak was substantially reduced or not visible for the SSc-PAH and PAH patients, probably because of reduced capillary perfusion as a result of pre-capillary PH. The ability to detect these changes with CSSR could be highly beneficial for future studies, in particular those which aim to measure longitudinal changes or treatment response, because ^{129}Xe is non-invasive and does not use ionising radiation. However, in both this work and the work of Stewart et al., the patients were older than the healthy volunteers. Therefore, further investigation with age-matched healthy controls is required to assess the impacts of lung disease on the alveolar septal thickness.

Lung disease is often spatially heterogeneous, so developing methods for mapping h using regional CSSR imaging is clinically valuable. Here, we presented initial results for a low-bandwidth gradient echo CSSR imaging sequence, which simultaneously imaged the gas-phase and dissolved-phase ^{129}Xe signal by exploiting the chemical shift between the two resonances. This allowed for CSSR images to be acquired at multiple delay times within a reasonable breath hold time, without the need for undersampling or multi-array coils. By fitting each pixel of the normalised dissolved-phase signal to the Patz model using the fitting procedure outlined in Section 9.3.5, maps were generated of not just the Patz model parameters but also their uncertainties. Reasonable agreement was found between spectroscopy and imaging parameters, however there were differences between the spectroscopy uptake curves and the whole-lung mean uptake from CSSR imaging. Initially, we thought that the increased $F(t)$ values for CSSR imaging could be a result of incomplete dissolved-phase signal saturation, since the imaging sequence only used one RF saturation pulse per block instead of the two used for spectroscopy. Although we confirmed that one saturation pulse was able to achieve sufficient saturation, B_1 field inhomogeneities may mean that the delivered flip angle was reduced in some parts of the lung. However, if this was the case, we would expect to see higher signal at all t_{delay} times, whereas at short t_{delay} times there was good agreement between the imaging and spectroscopy uptake curves. Running the CSSR imaging sequence with the frequency and phase-encoding gradients switched off resulted in an uptake curve which showed very good agreement with CSSR spectroscopy, suggesting that the differences are related to the spatial encoding gradients. To explore the reason for the discrepancy between CSSR imaging and spectroscopy further, a future experiment could be to perform CSSR imaging with the phase-encoding and temporal-encoding interleave ordering reversed. In other words, instead of repeating each CSSR block with the same t_{delay} at each phase-encoding step before moving onto the next t_{delay}, each line of k-space would be repeated for each t_{delay} before moving onto the next line. This would mean that the acquisition of each t_{delay} image would be spread out over the entire sequence duration. Another potential limitation related to the phase-encoding is that we used centric phase-encode ordering to maximise the SNR, however this can lead to blurring. With sequential encoding, the SNR may be lower, but initial data points which were acquired before a steady-state was reached are less important.

Our CSSR imaging method was limited by the lack of longer t_{delay} images, which constrains the fit accuracy in the linear section of the $F(t)$ curves, leading to an underestimation of τ. It was not possible to acquire more than six usable t_{delay} images because the SNR decreased towards the end of the sequence due to gas T_1 decay and the XTC effect [79]. During the delay time between the saturation and excitation pulses, there is rapid exchange between the gas-phase and depolarised dissolved-phase ^{129}Xe, which replenishes the dissolved signal, but also depletes the gas polarisation. To produce each t_{delay} image, the CSSR block (containing two 90° pulses separated by t_{delay}) was repeated for each of the 20 phase-encoding steps. Hence, acquiring six t_{delay} images involves saturating the dissolved-phase signal 120 times and applying 240 90° RF pulses. Attempting to acquire a seventh t_{delay} image was unsuccessful and the final images did not satisfy the Rose SNR criterion [269]. The choice of t_{delay} times was also limited by the breath hold time. With seven t_{delay} times the required breath hold time was 24 s, which was difficult for healthy volunteers and would be very challenging for patients with lung disease. With six t_{delay} images, the breath hold was a more manageable 17 s.

The need to repeat the CSSR block at each phase-encoding step within a reasonable breath hold time placed a limitation on the achievable spatial resolution. Similar to the work of Ruppert et al. [159], we used 2D coronal projection imaging, so a more detailed regional analysis of the Patz model parameter maps was not possible. Kern et al. used 3D radial spatial encoding in their regional CSSR sequence, allowing them to identify regional trends in the alveolar septal thickness [155]. *h* was observed to increase in the anterior to posterior direction in both healthy participants and those with COPD, although the trend was reduced in the COPD patients. This is likely due to the increased pulmonary blood in the posterior lung due to gravitational effects related to the subjects' supine position. Another limitation to our work is that the use of a combined dissolved-phase and gas-phase imaging sequence means that we cannot separate the dissolved-phase signal into the M and RBC components. Although this is not necessary for fitting the signals to the Patz model, it means that the MOXE cannot be used. We also lose potentially interesting spectral information by only acquiring the gas-phase and dissolved-phase ^{129}Xe signal magnitudes. Although it is currently too slow to be practical in humans, CSI-CSSR has the advantage of acquiring ^{129}Xe spectra for every imaging voxel, allowing the mapping of the RBC, M and gas ratios, frequencies and T_2^*, which may be affected by lung disease.

One possible solution to some of the issues with our sequence may be found by using the Look-Locker acquisition approach published by Kern et al. [155]. An advantage of this method is that the number of 90° RF pulses is greatly reduced, lessening the depolarisation of the gas-phase signal. Additionally, the ability to run several t_{delay} times simultaneously means that more t_{delay} images could be acquired within a breath hold. To confirm the feasibility of the Look-Locker acquisition strategy and to demonstrate the reproducibility of the method presented in Ref. [155], 2D finite difference method simulations were performed. A small percentage difference of 3% was found between the magnetisation from the Look-Locker acquisition and the unperturbed magnetisation. We tested a Look-Locker CSSR spectroscopy sequence in four healthy volunteers and found moderate agreement with standard CSSR spectroscopy, although this result may be slightly biased because the Look-Locker uptake curves were manually scaled to match those acquired with standard CSSR spectroscopy. This was done instead of using a scaling factor based on the dissolved-phase/gas-phase flip angles because the very small gas-phase flip angle was difficult to measure accurately. The Look-Locker CSSR results were therefore dependent on the standard CSSR results, however the scaling is only important in the estimation of S/V; it would be possible to extract *h* from the Look-Locker uptake curves without this scaling. Future work will involve developing a Look-Locker CSSR imaging sequence, similar to that of Kern et al. [155], but using 2D gradient echo simultaneous gas-phase and dissolved-phase imaging as in our existing CSSR imaging sequence.

An alternative method for combining spectral and spatial measurements of dissolved ^{129}Xe uptake is to perform time-resolved MESI with rapid spiral k-space encoding [149, 150]. Development of such a sequence and its comparison against the gradient echo CSSR imaging sequence presented in this chapter would be an interesting continuation of this project.

It would be valuable to investigate the repeatability of the CSSR imaging maps. *h* determined from CSSR spectroscopy has previously been demonstrated to have good repeatability, with a similar coefficient of variation between scans to that of PFT [141]. Kern et al. found good reproducibility of *h* and S/V derived from the M signal using the SPLASH regional CSSR technique, although τ was not reproducible [154]. The effect of factors such as lung inflation and B_0 field strength on CSSR spectroscopy and mapping also requires evaluation. Once our CSSR imaging sequence has been fully validated against CSSR spectroscopy, its reproducibility established and the effects of external factors understood, the goal will then be to test this sequence in patients with lung disease. In particular, it would be interesting to see whether the h maps are able to identify regional fibrosis in patients with ILD. Ruppert et al. found increased h values in the irradiated right lung of a rat when compared to the non-irradiated left lung, demonstrating the sensitivity of this parameter to regional lung damage related to pneumonitis and capillary bed destruction [159].

Another possible area of future work is the improvement of the gas exchange models used to fit the CSSR uptake data. The Patz model and MOXE are analytical models which solve the diffusion equation over simplified geometries. More realistic geometries, which include a non-uniform h, can be modelled using numerical modelling approaches such as the finite element method (FEM). This could improve the accuracy of modelled-derived parameters especially for ILD, which is characterised by a spatially heterogeneous tissue thickness. Moreover, the analytical models are limited by the need for many interrelated fit parameters, so numerical modelling could act as a validity check for these parameters. Stewart et al. developed three FEM gas exchange models; a 3D cylindrical model and two image-derived models, which used 2D histology images and micro-CT images respectively [338]. These models, and our initial work to build on them by developing a combined gas-phase and dissolved-phase FEM model, are described in Appendix B.

9.6 Conclusions

Non-invasive measurements of clinically relevant parameters, such as alveolar septal thickness, can be made using CSSR spectroscopy. We have reported an error quantification method for estimating the uncertainty on these parameters, which should aid their clinical interpretation. CSSR was capable of measuring septal thickening in patients with SSc and/or PAH, although further comparison with age-matched healthy controls is required to separate age-related changes from pathologic changes. CSSR imaging with a low-bandwidth gradient echo sequence allowed for regional measurement of alveolar septal thickness, which could help quantify fibrosis in heterogeneous lung disease. Further comparison of this method against CSSR spectroscopy is needed in order to validate our approach. Future work will involve further optimisation of the CSSR imaging sequence, for example by implementing a Look-Locker acquisition strategy to increase the number of t_{delay} images which can be acquired in a breath hold, and testing the sequence in patients with lung disease.

Chapter 10

Conclusion

The focus of this thesis has been the development and validation of hyperpolarised ^{129}Xe MRI/S methods for measuring pulmonary gas transfer and haemodynamics. In this chapter, the main findings from each chapter of original research are summarised, along with some proposed future work (more details on the future work were provided in each experimental chapter). This chapter concludes with a discussion of some of the emerging themes from this PhD project and potential future research directions.

10.1 Summary

In **Chapter 5**, the physiological interpretation of ^{129}Xe gas exchange metrics was investigated by comparing them to standard clinical measures of gas exchange. Three methods for modelling TL_{CO} from ^{129}Xe MRI metrics were developed and compared. We applied Wang et al.'s physiology-based model [339] to a large cohort of asthma and COPD patients and found different linear coefficient values. The prediction of TL_{CO} was improved by using a random forest regression machine learning model, which used patient age, sex and VV and RBC:Gas from ^{129}Xe MRI as features. A framework was developed to use the random forest model to create maps of predicted TL_{CO}, K_{CO} and V_A, providing additional regional information and aiding the comprehensibility of ^{129}Xe gas exchange MRI. The models showed good generalisability to a separate group of post-COVID-19 patients and healthy volunteers, however more work is needed to confirm their validity in data acquired from different patient groups and research sites. Future work for this project will include further assessment of our models with different patient groups, in particular those with PVD, and with data acquired at 3 T to further evaluate their generalisability.

In **Chapter 6**, a CS reconstruction for non-Cartesian dissolved ^{129}Xe MESI was implemented, with the purpose of reducing scan time. Through the comparison of retrospectively analysed patient and healthy volunteer data, we demonstrated good preservation of the dissolved ^{129}Xe signals and signal ratios using CS, compared with conventional gridding reconstruction. CS reconstruction allows for the scan time to be halved (7 s instead of 14 s) by acquiring fewer radial spokes, which will be more comfortable for patients with breathlessness. Although a prospectively undersampled sequence has not yet replaced the fully-sampled sequence in our research protocols, the CS reconstruction is useful for cases of patient breath hold failure or significant motion in the second half of the scan. The CS reconstruction only requires the first 7 s of data, which means that images are salvageable in such cases. Preliminary experiments with NA dissolved ^{129}Xe imaging in healthy participants showed that reasonable image SNR could be achieved using CS reconstruction, despite the lower signal strength in comparison to enriched ^{129}Xe. A possible continuation of this project would be the further improvement of image quality and SNR by incorporating prior knowledge of the T_1 and RF pulse-induced decay of the gas-phase signal into the CS reconstruction, as well as further validation in different disease populations.

Over the last few years, the amplitude of cardiogenic oscillations of the ^{129}Xe RBC signal has emerged as a potential biomarker of pulmonary microvascular function. In **Chapter 7**, we have contributed to the growing body of research on RBC oscillations by proposing a novel sliding window keyhole reconstruction technique for mapping the oscillation amplitude and phase. This method can be applied retrospectively to dissolved-phase ^{129}Xe imaging data, without the need for additional acquisitions. The work presented in this chapter builds upon previous methodology: whereas previous methods assumed that the RBC oscillations remained in phase over the lung, we showed that the phase is in fact spatially dependent. Regional trends in the RBC oscillation amplitude and phase were examined, as well as their inter- and intra-scan repeatability, in order to establish these metrics in healthy volunteers before evaluation in patients with lung disease in the next chapter.

Chapter 8 builds upon the previous chapter by applying the RBC oscillation mapping method to post-COVID-19 and CTEPH patients. No significant differences were found between hospitalised COVID-19 patients, non-hospitalised long-Covid patients and healthy volunteers, although the oscillation amplitude and phase were higher in hospitalised COVID-19 patients who had residual lung abnormalities. Overall, the results in COVID-19 patients did not provide evidence of microvascular disease in this cohort. In four CTEPH patients, the whole-lung oscillation amplitude and phase were similar to those of the healthy volunteers, but the heterogeneity of both maps were significantly higher. Preliminary attempts to validate our methodology against DCE-MRI, the gold standard for measuring perfusion, and ^{129}Xe gas exchange MRI found that areas of increased phase difference approximately corresponded to regions of low gas transfer and perfusion and increased TTP. These initial results demonstrate the potential of RBC oscillation amplitude and phase mapping for identifying microvascular damage, but evaluation in increased patient numbers, including different PH subtypes, is required to corroborate this.

In **Chapter 9**, methods for CSSR spectroscopy, imaging and analysis were presented. An efficient method for estimating the uncertainty on Patz model parameters was developed using bootstrapping. This is important for determining how much confidence can be placed on the CSSR-derived metrics. A significantly higher alveolar septal thickness was found in patients with SSc and/or PAH when compared to healthy volunteers, using CSSR spectroscopy. However, the volunteers were not age-matched and an age correction was not applied. Additionally, the number of participants included in this chapter was small; the acquisition of more CSSR data in patients with SSc and/or PAH is underway. A combined dissolved-phase/gas-phase gradient echo sequence that exploits the chemical shift artefact was developed for regional CSSR imaging. This allowed for spatial mapping of the alveolar septal thickness, which will be a useful tool for quantifying regional fibrosis and emphysema. More work is needed to optimise and validate this approach against CSSR spectroscopy and Dixon/MESI-type sequences before it can be used for lung disease patients. Other techniques for CSSR imaging may be explored in the future, such as a Look-Locker type acquisition strategy. A ^{129}Xe multi-channel array coil has become available at our site [340], which may be useful for increasing spatial sensitivity in a CSSR imaging sequence. Another possible avenue of future research is the investigation of numerical gas exchange models which could help the interpretation of CSSR data. Preliminary work towards a 3D geometric model of ^{129}Xe diffusion in the alveolar airspaces, membrane and capillaries is described in Appendix B. This model could allow for simultaneous prediction of ^{129}Xe diffusion coefficients and CSSR data.

10.2 Conclusions and Outlook

Hyperpolarised ^{129}Xe functional lung MRI holds huge clinical potential, providing sensitivity to early lung disease and non-ionising measurement of disease progression and treatment response. In particular, dissolved-phase ^{129}Xe MRI is uniquely poised to probe regional gas exchange and can provide a wealth of information which is otherwise unavailable from other imaging modalities or lung function tests. However, despite its many advantages, there are several barriers to the widespread adoption of this technique and it has mostly been limited as a research tool. Since 2015, our group has had UK Medicine and Healthcare Regulatory Authority approval to perform routine clinical lung imaging using hyperpolarised gas. In the USA, Polarean Imaging were recently granted FDA approval for their hyperpolarised ^{129}Xe contrast agent, which should help to increase access to ^{129}Xe lung MRI. One of the most significant hurdles to the expansion of ^{129}Xe MRI is the high cost associated with the required infrastructure, particularly for hyperpolarisation. To make hyperpolarised ^{129}Xe MRI more accessible, a priority needs to be reducing operational costs.

CS reconstruction can be used to decrease scan time for dissolved-phase ^{129}Xe MRI and possibly reduce the required dose, by allowing the use of either lower cost NA Xe or a lower volume of enriched ^{129}Xe, both of which would make ^{129}Xe MRI cheaper. Sequences which combine multiple ^{129}Xe measurements into one scan, such as single-breath ventilation and gas exchange imaging [174] or combined CSSR and DW-MRI sequences [146, 147], are enticing and may be facilitated by the use of CS. Low field (<0.5 T) MRI is associated with reduced material and operational costs and has additional advantages specific to lung imaging, namely, improved ^{1}H SNR and longer ^{129}Xe T_2^* [207]. Our group is starting to investigate ^{129}Xe MRI at 0.5 T. However, performing dissolved-phase ^{129}Xe imaging at low field strength will be challenging because the ^{129}Xe resonant frequency differences are field strength dependent. The frequency shift between the M and

RBC resonances is 200 Hz at 3 T, but is reduced to only 20 Hz at 0.3 T, which may prohibit their separation [207]. Another aspect to consider when trying to increase the accessibility of ^{129}Xe MRI is the dissemination and interpretation of results by clinicians, physiologists and patients. Involving these stakeholders in the process is important and a project is currently underway as part of the MAGNIFY study, led by our research group, to co-produce a toolkit to support the communication of ^{129}Xe MRI results [341].

A prominent theme throughout this thesis is the progression from whole-lung metrics towards regional mapping. Lung diseases are often heterogeneous in nature, so spatial encoding of lung function measurements is important for detecting pathologic changes. In chapters 5, 7, 8 and 9 we described efforts to create regional maps of metrics that have generally only been measured as global averages. The benefits of this were especially apparent in Chapter 8, where we found that the coefficient of variation of the RBC oscillation amplitude and phase maps were more sensitive to differences between healthy people and those with CTEPH than the average map values.

The purpose of this thesis was the methodological development and initial validation of dissolved-phase ^{129}Xe MRI/S techniques, however in future, more emphasis needs to be placed on demonstrating their impact. The role of ^{129}Xe gas exchange MRI in driving treatment decisions and assisting in the diagnosis and management of patients needs more substantiation to propel clinical translation. For example, the CSSR-derived alveolar septal thickness holds potential as a measure of fibrosis, and could conceivably act as a non-invasive biomarker for treatment response to anti-fibrotic drugs. RBC oscillation amplitude and phase mapping have shown potential for measuring pulmonary microvascular function and a worthwhile continuation of this project would be to investigate post-surgery changes in CTEPH patients undergoing PEA, to ascertain whether this method can help predict patient outcomes.

In conclusion, ^{129}Xe MRI is a versatile and sensitive method for measuring pulmonary gas exchange function, and the work presented in this thesis will hopefully contribute towards the wider clinical translation of this technique.

Bibliography

[1] World Health Organization. The top 10 causes of death. https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death, 2024.

[2] British Lung Foundation. The battle for breath - the impact of lung disease in the uk. https://www.asthmaandlung.org.uk/battle-breath-report, 2016.

[3] British Lung Foundation. Estimating the economic burden of respiratory illness in the UK. https://www.asthmaandlung.org.uk/economic-burden-report, 2017.

[4] Forum of International Respiratory Societies. The global impact of respiratory disease. Third Edition. European Respiratory Society. firsnet.org/images/publications/FIRS_Master_09202021.pdf, 2021.

[5] S.A. Quaderi and J.R. Hurst. The unmet global burden of COPD. *Global Health, Epidemiology and Genomics*, 3:e4, 2018.

[6] H. Hamedani, K. Ma, D. DiBardino, R. Baron, S. Kadlecek, F. Amzajerdian, K. Ruppert, I. Duncan, and R.R. Rizi. Simultaneous imaging of ventilation and gas exchange with hyperpolarized Xe-129 MRI for monitoring patients with endobronchial valve interventions. *Am J Respir Crit Care Med*, 205(9):E48–E50, 2022.

[7] P.J. Martinez-Pitre, B.R. Sabbula, and M. Cascella. Restrictive lung disease. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing; 2025 Jan-. Available from: https://www.ncbi.nlm.nih.gov/books/NBK560880/, 2023.

[8] K.M. Antoniou, G.A. Margaritopoulos, S. Tomassetti, F. Bonella, U. Costabel, and V. Poletti. Interstitial lung disease. *European Respiratory Review*, 23(131):40–54, 2014.

[9] L. Lonzetti, M. Zanon, G.S. Pacini, S. Altmayer, D.M. de Oliveira, A.S. Rubin, F.F. Gazzoni, M.C. Barros, and B. Hochhegger. Magnetic resonance imaging of interstitial lung diseases: A state-of-the-art review. *Respiratory Medicine*, 155:79–85, 2019.

[10] K.C. Meyer and C.A. Decker. Role of pirfenidone in the management of pulmonary fibrosis. *Ther Clin Risk Manag*, 13:427–437, 2017.

[11] J. Fukihara and Y. Kondoh. COVID-19 and interstitial lung diseases: A multifaceted look at the relationship between the two diseases. *Respir Investig*, 61:601–617, 2023.

[12] A.J. Tyndall, B. Bannert, M. Vonk, P. Airò, F. Cozzi, P.E. Carreira, D.F. Bancel, Y. Allanore, U. Müller-Ladner, O. Distler, F. Iannone, R. Pellerito, M. Pileckyte, I. Miniati, L. Ananieva, A.B. Gurman, N. Damjanov, A. Mueller, G. Valentini, G. Riemekasten, M. Tikly, L. Hummers, M.J.S. Henriques, P. Caramaschi, A. Scheja, B. Rozman, E. Ton, G. Kumánovics, B. Coleiro, E. Feierl, G. Szucs, C. Alberto Von Mühlen, V. Riccieri, S. Novak, C. Chizzolini, A. Kotulska, C. Denton, P.C. Coelho, I. Kötter, I. Simsek, P. García de la Pena Lefebvre, E. Hachulla, J.R. Seibold, S. Rednic, J. tork, J. Morovic-Vergles, and U.A. Walker. Causes and risk factors for death in systemic sclerosis: a study from the EULAR scleroderma trials and research (EUSTAR) database. *Annals of the Rheumatic Diseases*, 69(10):1809–1815, 2010.

[13] N.D. Weatherley, J.A. Eaden, N.J. Stewart, B.J. Bartholmai, A.J. Swift, S.M. Bianchi, and J.M. Wild. Experimental and quantitative imaging techniques in interstitial lung disease. *Thorax*, 74(6):611–619, 2019.

[14] W. Seeger, Y. Adir, J. A. Barbera, H. Champion, J. G. Coghlan, V. Cottin, T. De Marco, N. Galie, S. Ghio, S. Gibbs, F. J. Martinez, M. J. Semigran, G. Simonneau, A. U. Wells, and J. L. Vachiery. Pulmonary hypertension in chronic lung diseases. *J Am Coll Cardiol*, 62(25 Suppl):D109–16, 2013.

[15] M. K. Han, V. V. McLaughlin, G. J. Criner, and F. J. Martinez. Pulmonary diseases and the heart. *Circulation*, 116(25):2992–3005, 2007.

[16] S.H. Oldroyd, G. Manek, and A. Bhardwaj. Pulmonary hypertension. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing; 2025 Jan-. Available from: https://www.ncbi.nlm.nih.gov/books/NBK482463/, 2024.

[17] T. Dahhan, S.S. Kaushik, M. He, J.G. Mammarappallil, V.F. Tapson, H.P. McAdams, T.A. Sporn, B. Driehuys, and S. Rajagopal. Abnormalities in hyperpolarized 129Xe magnetic resonance imaging and spectroscopy in two patients with pulmonary vascular disease. *Pulm Circ*, 6(1):126–31, 2016.

[18] J.M. Wild, H. Marshall, M. Bock, L.R. Schad, P.M. Jakob, M. Puderbach, F. Molinari, E.J.R. Van Beek, and J. Biederer. MRI of the lung (1/3): methods. *Insights into Imaging*, 3:345–353, 2012.

[19] O. Togao, R. Tsuji, Y. Ohno, I. Dimitrov, and M. Takahashi. Ultrashort echo time (UTE) MRI of the lung: Assessment of tissue density in the lung parenchyma. *Magn Reson Med*, 64(5):1491–1498, 2010.

[20] K.M. Johnson, S.B. Fain, M.L. Schiebler, and S. Nagle. Optimized 3D ultrashort echo time pulmonary MRI. *Magn Reson Med*, 70(5):1241–1250, 2013.

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

[22] J.P. Mugler, B. Driehuys, J.R. Brookeman, G.D. Cates, S.S. Berr, R.G. Bryant, T.M. Daniel, E.E. deLange, J.H. Downs, C.J. Erickson, W. Happer, D.P. Hinton, N.F. Kassel, T. Maier, C.D. Phillips, B.T. Saam, K.L. Sauer, and M.E. Wagshul. Mr imaging and spectroscopy using hyperpolarized Xe-129 gas: Preliminary human results. *Magn Reson Med*, 37(6):809–815, 1997.

[23] M. Kirby, S. Svenningsen, A. Owrangi, A. Wheatley, A. Farag, A. Ouriadov, G.E. Santyr, R. Etemad-Rezai, H.O. Coxson, D.G. McCormack, and G. Parraga. Hyperpolarized 3He and 129Xe MR imaging in healthy volunteers and patients with chronic obstructive pulmonary disease. *Radiology*, 265(2):600–610, 2012. PMID: 22952383.

[24] L. Ebner, M. He, R.S. Virgincar, T. Heacock, S.S. Kaushik, M.S. Freemann, H.P. McAdams, M. Kraft, and B. Driehuys. Hyperpolarized 129xenon magnetic resonance imaging to quantify regional ventilation differences in mild to moderate asthma: A prospective comparison between semiautomated ventilation defect percentage calculation and pulmonary function tests. *Investigative radiology*, 52(2):120127, February 2017.

[25] G. Santyr, N. Kanhere, F. Morgado, J.H. Rayment, F. Ratjen, and M.J. Couch. Hyperpolarized gas magnetic resonance imaging of pediatric cystic fibrosis lung disease. *Academic Radiology*, 26(3):344–354, 2019. Special Issue: Pulmonary Functional Imaging.

[26] H.-F. Chan, N.J. Stewart, J. Parra-Robles, G.J. Collier, and J.M. Wild. Whole lung morphometry with 3D multiple b-value hyperpolarized gas MRI and compressed sensing. *Magn Reson Med*, 77(5):1916–1925, 2017.

[27] H.-F. Chan, N.J. Stewart, G. Norquay, G.J. Collier, and J.M. Wild. 3D diffusion-weighted 129Xe MRI for whole lung morphometry. *Magn Reson Med*, 79(6):2986–2995, 2018.

[28] A.L. Sukstanskii and D.A. Yablonskiy. Lung morphometry with hyperpolarized 129Xe: Theoretical background. *Magn Reson Med*, 67:856866, 2012.

[29] J. Parra-Robles, H. Marshall, R.A. Hartley, C.E. Brightling, and J.M. Wild. Quantification of lung microstructure in asthma using a 3He fractional diffusion approach. In *Proc. 22nd Intl. Soc. Mag. Reson. Med. 3529*, volume 22, page 3529, 2014.

[30] S.S. Kaushik, Z.I. Cleveland, G.P. Cofer, G. Metz, D. Beaver, J. Nouls, M. Kraft, W. Auffermann, J. Wolber, H.P. McAdams, and B. Driehuys. Diffusion-weighted hyperpolarized 129Xe MRI in healthy volunteers and subjects with chronic obstructive pulmonary disease. *Magn Reson Med*, 65(4):1154–1165, 2011.

[31] Z.I. Cleveland, R.S. Virgincar, Y. Qi, S.H. Robertson, S. Degan, and B. Driehuys. 3D MRI of impaired hyperpolarized Xe-129 uptake in a rat model of pulmonary fibrosis. *NMR in Biomedicine*, 27(12):1502–1514, 2014.

[32] S.S. Kaushik, S.H. Robertson, M.S. Freeman, M. He, K.T. Kelly, J.E. Roos, C.R. Rackley, W.M. Foster, H.P. McAdams, and B. Driehuys. Single-breath clinical imaging of hyperpolarized Xe-129 in the airspaces, barrier, and red blood cells using an interleaved 3D radial 1-point Dixon acquisition. *Magn Reson Med*, 75(4):1434–1443, 2016.

[33] J.M. Wang, S.H. Robertson, Z.Y. Wang, M. He, R.S. Virgincar, G.M. Schrank, R.M. Smigla, T.G. O'Riordan, J. Sundy, L. Ebner, C.R. Rackley, P. McAdams, and B. Driehuys. Using hyperpolarized Xe-129 MRI to quantify regional gas transfer in idiopathic pulmonary fibrosis. *Thorax*, 73(1):21–28, 2017.

[34] D. Mummy, Z. Wang, E. Bier, R.M. Tighe, B. Driehuys, and J. Mammarappallil. Hyperpolarized 129Xe magnetic resonance imaging measures of gas exchange in non-specific interstitial pneumonia. *Am J Respir Crit Care Med*, 201, 2020.

[35] G.J. Collier, J.A. Eaden, P.J.C. Hughes, S.M. Bianchi, N.J. Stewart, N.D. Weatherley, G. Norquay, R.F. Schulte, and J.M. Wild. Dissolved Xe-129 lung MRI with four-echo 3D radial spectroscopic imaging: Quantification of regional gas transfer in idiopathic pulmonary fibrosis. *Magn Reson Med*, 85(5):2622–2633, 2021.

[36] K. Qing, K. Ruppert, Y. Jiang, J.F. Mata, G.W. Miller, Y.M. Shim, C. Wang, I.C. Ruset, F.W. Hersman, T.A. Altes, and J.P. Mugler. Regional mapping of gas uptake by blood and tissue in the human lung using hyperpolarized xenon-129 MRI. *J Magn Reson Imaging*, 39(2):346–359, 2014.

[37] Z. Wang, E.A. Bier, A. Swaminathan, K. Parikh, J. Nouls, M. He, J.G. Mammarappallil, S. Luo, B. Driehuys, and S. Rajagopal. Diverse cardiopulmonary diseases are associated with distinct xenon magnetic resonance imaging signatures. *Eur Respir J*, 54(6), 2019.

[38] K. Qing, J.P. Mugler, T.A. Altes, Y. Jiang, J.F. Mata, G.W. Miller, I.C. Ruset, F.W. Hersman, and K. Ruppert. Assessment of lung function in asthma and copd using hyperpolarized Xe-129 chemical shift saturation recovery spectroscopy and dissolved-phase MRI. *Nmr in Biomedicine*, 27(12):1490–1501, 2014.

[39] N.J. Stewart, G. Leung, G. Norquay, H. Marshall, J. Parra-Robles, P.S. Murphy, R.F. Schulte, C. Elliot, R. Condliffe, P.D. Griffiths, D.G. Kiely, M.K. Whyte, J. Wolber, and J.M. Wild. Experimental validation of the hyperpolarized Xe-129 chemical shift saturation recovery technique in healthy volunteers and subjects with interstitial lung disease. *Magn Reson Med*, 74(1):196–207, 2015.

[40] K. Ruppert, J.R. Brookeman, K.D. Hagspiel, B. Driehuys, and J.P. Mugler. NMR of hyperpolarized Xe-129 in the canine chest: spectral dynamics during a breath-hold. *Nmr in Biomedicine*, 13(4):220–228, 2000.

[41] S. Patz, I. Muradian, M.I. Hrovat, I.C. Ruset, G. Topulos, S.D. Covrig, E. Frederick, H. Hatabu, F.W. Hersman, and J.P. Butler. Human pulmonary imaging and spectroscopy with hyperpolarized 129Xe at 0.2T. *Academic Radiology*, 15(6):713–727, 2008.

[42] R.W. Brown, Y.C.N. Norman Cheng, E. Mark Haacke, M.R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging : Physical Principles and Sequence Design*. John Wiley and Sons, New Jersey, 2014.

[43] D.W. McRobbie, E.A. Moore, M.J. Graves, and M.R. Prince. *MRI From Picture to Proton*. Cambridge University Press, Cambridge, 2017.

[44] I.I. Rabi, J.R. Zacharias, S. Millman, and P. Kusch. A new method of measuring nuclear magnetic moment. *Phys Rev*, 53:318, 1938.

[45] F. Bloch, W.W. Hansen, and M. Packard. Nuclear induction. *Phys Rev*, 69:127, 1946.

[46] F. Bloch, W.W. Hansen, and M. Packard. The nuclear induction experiment. *Phys Rev*, 70:474–485, 1946.

[47] E.M. Purcell, H.C. Torrey, and R.V. Pound. Resonance absorption by nuclear moments in a solid. *Phys Rev*, 69:37–38, 1946.

[48] R.R. Ernst and W.A. Anderson. Application of fourier transform spectroscopy to magnetic resonance. *Review of Scientific Instruments*, 37(1):93–102, 01 1966.

[49] R. Damadian. Tumor detection by nuclear magnetic resonance. *Science*, 171(3976):1151–1153, 1971.

[50] P.C. Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *nature*, 242(5394):190–191, 1973.

[51] D.I. Hoult and R.E. Richards. The signal-to-noise ratio of the nuclear magnetic resonance experiment. *J Magn Reson*, 24(1):71–85, 1976.

[52] G. Biro. k-space explorer. https://github.com/birogeri/kspace-explorer.

[53] S. Posse, R. Otazo, S.R. Dager, and J. Alger. MR spectroscopic imaging: Principles and recent advances. *J Magn Reson Imaging*, 37(6):1301–1325, 2013.

[54] F. Wiesinger, E. Weidl, M.I. Menzel, M.A. Janich, O. Khegai, S.J. Glaser, A. Haase, M. Schwaiger, and R.F. Schulte. IDEAL spiral CSI for dynamic metabolic MR imaging of hyperpolarized [1-13C]pyruvate. *Magn Reson Med*, 68(1):8–16, 2012.

[55] J.M. Pauly. Non-Cartesian reconstruction. https://mri-q.com/uploads/3/4/5/7/34572113/pauly-non-cartesian_recon.pdf.

[56] J.I. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski. Selection of a convolution function for Fourier inversion using gridding (computerised tomography application). *IEEE Transactions on Medical Imaging*, 10(3):473–478, 1991.

[57] P.J. Beatty, D.G. Nishimura, and J.M. Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Transactions on Medical Imaging*, 24(6):799–808, 2005.

[58] W.Q. Malik, H.A. Khan, D.J. Edwards, and C.J. Stevens. A gridding algorithm for efficient density compensation of arbitrarily sampled fourier-domain data. In *IEEE/Sarnoff Symposium on Advances in Wired and Wireless Communication, 2005.*, pages 125–128, 2005.

[59] O.N. Jaspan, R. Fleysher, and M.L. Lipton. Compressed sensing MRI: a review of the clinical literature. *British Journal of Radiology*, 88(1056):20150487, 10 2015.

[60] M. Lustig, D. Donoho, and J.M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med*, 58(6):1182–95, 2007.

[61] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trendső in Machine Learning*, 3(1):1–122, 2010.

[62] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

[63] H. Marshall, N.J. Stewart, H.F. Chan, M.D. Rao, G. Norquay, and J.M. Wild. In vivo methods and applications of xenon-129 magnetic resonance. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 122:42–62, 2021.

[64] Y. Shukla, A. Wheatley, M. Kirby, S. Svenningsen, A. Farag, G. E. Santyr, N. A. Paterson, D. G. McCormack, and G. Parraga. Hyperpolarized 129Xe magnetic resonance imaging: tolerability in healthy volunteers and subjects with pulmonary disease. *Acad Radiol*, 19(8):941–51, 2012.

[65] A.M. Oros and N.J. Shah. Hyperpolarized xenon in NMR and MRI. *Physics in Medicine and Biology*, 49(20):R105–R153, 2004.

[66] S. Fain, M.L. Schiebler, D.G. McCormack, and G. Parraga. Imaging of lung function using hyperpolarized helium-3 magnetic resonance imaging: Review of current and emerging translational methods and applications. *J Magn Reson Imaging*, 32(6):1398–1408, 2010.

[67] D. Shea and D. Morgan. The Helium-3 shortage: Supply, demand, and options for congress. *Congressional Research Service*, page R41419, 01 2011.

[68] R.K. Harris. Encyclopedia of nuclear magnetic resonance. *in-Chief DM Grant and RK Harris*, 5:3301, 1996.

[69] W.H. Isbister, P.F. Schofield, and H.B. Torrance. Measurement of the solubility of xenon-133 in blood and human brain. *Phys Med Biol*, 10:243–50, 1965.

[70] M.S. Albert, V.D. Schepkin, and T.F. Budinger. Measurement of 129Xe T1 in blood to explore the feasibility of hyperpolarized 129Xe MRI. *Journal of Computer Assisted Tomography*, 19(6):975–978, 1995.

[71] J. Ladefoged and A.M. Andersen. Solubility of xenon-133 at 37degC in water, saline, olive oil, liquid paraffin, solutions of albumin, and blood. *Physics in Medicine and Biology*, 12(3):353–358, 1967.

[72] J. Wolber, A. Cherubini, A. S. K. Dzik-Jurasz, M. O. Leach, and A. Bifone. Spin-lattice relaxation of laser-polarized xenon in human blood. *Proc Natl Acad Sci U S A*, 96(7):3664–3669, 1999.

[73] R.Y. Chen, F.C. Fan, S. Kim, K.M. Jan, S. Usami, and S. Chien. Tissue-blood partition coefficient for xenon: temperature and hematocrit dependence. *J Appl Physiol Respir Environ Exerc Physiol*, 49(2):178–83, 1980.

[74] S.-Y. Yeh and R.E. Peterson. Solubility of carbon dioxide, krypton, and xenon in aqueous solution. *Journal of Pharmaceutical Sciences*, 53(7):822–824, 1964.

[75] K.W. Miller, N.V. Reo, A.J. Schoot Uiterkamp, D.P. Stengle, T.R. Stengle, and K.L. Williamson. Xenon NMR: chemical shifts of a general anesthetic in common solvents, proteins, and membranes. *Proc Natl Acad Sci U S A*, 78(8):4946–9, 1981.

[76] J. Wolber, I.J. Rowland, M.O. Leach, and A. Bifone. Perfluorocarbon emulsions as intravenous delivery media for hyperpolarized xenon. *Magn Reson Med*, 41(3):442–9, 1999.

[77] K. Sakai, A.M. Bilek, E. Oteiza, R.L. Walsworth, D. Balamore, F.A. Jolesz, and M.S. Albert. Temporal dynamics of hyperpolarized 129Xe resonances in living rats. *J Magn Reson B*, 111(3):300–4, 1996.

[78] J. Wolber, D.J. McIntyre, L.M. Rodrigues, P. Carnochan, J.R. Griffiths, M.O. Leach, and A. Bifone. In vivo hyperpolarized 129Xe NMR spectroscopy in tumors. *Magn Reson Med*, 46(3):586–91, 2001.

[79] K. Ruppert, J.R. Brookeman, K.D. Hagspiel, and J.P. Mugler. Probing lung physiology with xenon polarization transfer contrast (XTC). *Magn Reson Med*, 44(3):349–357, 2000.

[80] A. Bifone, Y.Q. Song, R. Seydoux, R.E. Taylor, B.M. Goodson, T. Pietrass, T.F. Budinger, G. Navon, and A. Pines. NMR of laser-polarized xenon in human blood. *Proc Natl Acad Sci U S A*, 93(23):12932–6, 1996.

[81] J.P. Mugler and T.A. Altes. Hyperpolarized 129Xe MRI of the human lung. *J Magn Reson Imaging*, 37(2):313–331, 2013.

[82] G. Norquay, G. Leung, N.J. Stewart, G.M. Tozer, J. Wolber, and J.M. Wild. Relaxation and exchange dynamics of hyperpolarized 129Xe in human blood. *Magn Reson Med*, 74(2):303–311, 2015.

[83] W. Kilian, F. Seifert, and H. Rinneberg. Dynamic NMR spectroscopy of hyperpolarized 129Xe in human brain analyzed by an uptake model. *Magn Reson Med*, 51(4):843–847, 2004.

[84] M. Rao, N.J. Stewart, G. Norquay, P.D. Griffiths, and J.M. Wild. High resolution spectroscopy and chemical shift imaging of hyperpolarized 129 Xe dissolved in the human brain in vivo at 1.5 tesla. *Magn Reson Med*, 75(6):2227–2234, 2016.

[85] M.R. Rao, N.J. Stewart, P.D. Griffiths, G. Norquay, and J.M. Wild. Imaging human brain perfusion with inhaled hyperpolarized 129Xe MR imaging. *Radiology*, 286(2):659–665, 2018.

[86] J. ChaconCaldera, A. Maunder, M. Rao, G. Norquay, O.I. Rodgers, M. Clemence, C. Puddu, L.R. Schad, and J.M. Wild. Dissolved hyperpolarized xenon129 MRI in human kidneys. *Magn Reson Med*, 83(1):262–270, 2020.

[87] X. Xu, G. Norquay, S.R. Parnell, M.H. Deppe, S. Ajraoui, R. Hashoian, H. Marshall, P.D. Griffiths, J. Parra-Robles, and J.M. Wild. Hyperpolarized 129Xe gas lung MRISNR and T2* comparisons at 1.5 T and 3 T. *Magn Reson Med*, 68(6):1900–1904, 2012.

[88] Z.Y. Wang, M. He, E. Bier, L. Rankine, G. Schrank, S. Rajagopal, Y.C. Huang, C. Kelsey, S. Womack, J. Mammarappallil, and B. Driehuys. Hyperpolarized Xe-129 gas transfer MRI: the transition from 1.5T to 3T. *Magn Reson Med*, 80(6):2374–2383, 2018.

[89] J. Wolber, A. Cherubini, M.O. Leach, and A. Bifone. Hyperpolarized 129Xe NMR as a probe for blood oxygenation. *Magn Reson Med*, 43(4):491–496, 2000.

[90] G. Norquay, G. Leung, N.J. Stewart, J. Wolber, and J.M. Wild. (129) Xe chemical shift in human blood and pulmonary blood oxygenation measurement in humans using hyperpolarized (129) Xe NMR. *Magn Reson Med*, 77(4):1399–1408, 2017.

[91] R.H. Acosta, L. Agulles-Pedrós, S. Komin, D. Sebastiani, H.W. Spiess, and P. Blümler. Diffusion in binary gas mixtures studied by NMR of hyperpolarized gases and molecular dynamics simulations. *Physical Chemistry Chemical Physics*, 8(36):4182–4188, 2006.

[92] X. J. Chen, H. E. Möller, M. S. Chawla, G. P. Cofer, B. Driehuys, L. W. Hedlund, and G. A. Johnson. Spatially resolved measurements of hyperpolarized gas properties in the lung in vivo. part I: Diffusion coefficient. *Magn Reson Med*, 42(4):721–728, 1999.

[93] K. Ruppert, J.F. Mata, J.R. Brookeman, K.D. Hagspiel, and J.P. Mugler III. Exploring lung function with hyperpolarized 129Xe nuclear magnetic resonance. *Magn Reson Med*, 51(4):676–687, 2004.

[94] A. Comment, S. Jannin, J.-N. Hyacinthe, P. Miéville, R. Sarkar, P. Ahuja, P.R. Vasos, X. Montet, F. Lazeyras, J.-P. Vallée, P. Hautle, J.A. Konter, B. van den Brandt, J.-P. Ansermet, R. Gruetter, and G. Bodenhausen. Hyperpolarizing gases via dynamic nuclear polarization and sublimation. *Phys. Rev. Lett.*, 105:018104, Jul 2010.

[95] T.G. Walker and W. Happer. Spin-exchange optical pumping of noble-gas nuclei. *Rev. Mod. Phys.*, 69:629–642, Apr 1997.

[96] J.W. Plummer, K. Emami, A. Dummer, J.C. Woods, L.L. Walkup, and Z.I. Cleveland. A semi-empirical model to optimize continuous-flow hyperpolarized 129Xe production under practical cryogenic-accumulation conditions. *J Magn Reson*, 320:106845, 2020.

[97] N.J. Stewart, G. Norquay, P.D. Griffiths, and J.M. Wild. Feasibility of human lung ventilation imaging using highly polarized naturally abundant xenon and optimized three-dimensional steady-state free precession. *Magn Reson Med*, 74(2):346–52, 2015.

[98] G. Leung, G. Norquay, R.F. Schulte, and J.M. Wild. Radiofrequency pulse design for the selective excitation of dissolved Xe-129. *Magn Reson Med*, 73(1):21–30, 2015.

[99] O. Doganay, T. Wade, E. Hegarty, C. McKenzie, R.F. Schulte, and G.E. Santyr. Hyperpolarized Xe-129 imaging of the rat lung using spiral ideal. *Magn Reson Med*, 76(2):566–576, 2016.

[100] N. De Zanche, N. Chhina, K. Teh, C. Randell, K.P. Pruessmann, and J.M. Wild. Asymmetric quadrature split birdcage coil for hyperpolarized 3He lung MRI at 1.5T. *Magn Reson Med*, 60:431–438, 2008.

[101] I. Dregely, I.C. Ruset, G. Wiggins, A. Mareyam, 3rd J.P. Mugler, T.A. Altes, C. Meyer, K. Ruppert, L.L. Wald, and F.W. Hersman. 32-channel phased-array receive with asymmetric birdcage transmit coil for hyperpolarized xenon-129 lung imaging. *Magn Reson Med*, 70(2):576–83, 2013.

[102] JB West. *Respiratory Physiology*. Lippencott Williams and Wilkins, Baltimore, MD, 8th edition, 2008.

[103] K. Lamb, D. Theodore, and B.S. Bhutta. Spirometry. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing; 2025 Jan-. Available from: https://www.ncbi.nlm.nih.gov/books/NBK560526/, 2017.

[104] E. Weibel. What makes a good lung? *Swiss Med Wkly*, 139:375386, 2009.

[105] E.R. Weibel. *Morphometry of the human lung*. Springer, 1963.

[106] Matthias Ochs, Jens R. Nyengaard, Anja Jung, Lars Knudsen, Marion Voigt, Thorsten Wahlers, Joachim Richter, and Hans Jørgen G. Gundersen. The number of alveoli in the human lung. *American Journal of Respiratory and Critical Care Medicine*, 169(1):120–124, 2004. PMID: 14512270.

[107] W.S. Blakemore, R.E. Forster, J.W. Morton, and C.M. Ogilvie. A standardized breath holding technique for the clinical measurement of the diffusing capacity of the lung for carbon monoxide. *J Clin Invest*, 36(1 Part 1):1–17, 1957.

[108] M.J. Hegewald. Diffusing capacity. *Clin Rev Allergy Immunol*, 37(3):159–66, 2009.

[109] P. Modi and M. Cascella. Diffusing capacity of the lungs for carbon monoxide. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing. Available from:: https://www.ncbi.nlm.nih.gov/books/NBK556149/, Jan 2024 2023.

[110] E.J. Heckman and G.T. OConnor. Pulmonary function tests for diagnosing lung disease. *JAMA*, 313(22):2278–2279, 2015.

[111] N.S. Goh, R.K. Hoyles, C.P. Denton, D.M. Hansell, E.A. Renzoni, T.M. Maher, A.G. Nicholson, and A.U. Wells. Short-term pulmonary function trends are predictive of mortality in interstitial lung disease associated with systemic sclerosis. *Arthritis & Rheumatology*, 69(8):1670–1678, 2017.

[112] K.K. Brown, K.R. Flaherty, V. Cottin, G. Raghu, Y. Inoue, A. Azuma, J.T. Huggins, L. Richeldi, S. Stowasser, W. Stansen, R. Schlenker-Herceg, T.M. Maher, and A.U. Wells. Lung function outcomes in the INPULSIS (R) trials of nintedanib in idiopathic pulmonary fibrosis. *Respiratory Medicine*, 146:42–48, 2019.

[113] M.C. McCormack. Facing the noise: Addressing the endemic variability in D-LCO testing. *Respiratory Care*, 57(1):17–23, 2012.

[114] G. Barisione and V. Brusasco. Pulmonary function tests in systemic sclerosis-associated interstitial lung disease: new directions and future prospects. *Current Opinion in Physiology*, 22, 2021.

[115] G.S. Zavorsky and I. van der Lee. Can the measurement of pulmonary diffusing capacity for nitric oxide replace the measurement of pulmonary diffusing capacity for carbon monoxide? *Respir Physiol Neurobiol*, 241:9–16, 2017.

[116] J.M.B. Hughes and A.T. Dinh-Xuan. The DLNO/DLCO ratio: Physiological significance and clinical implications. *Respiratory Physiology & Neurobiology*, 241:17–22, 2017.

[117] B. Degano, T. Soumagne, T. Delaye, P. Berger, T. Perez, A. Guillien, J.-L. Pellegrin, D. Launay, N. Magy-Bertrand, C. Agard, K.P. Tiev, T. Hua-Huy, C. Tardiff, V. Diaz, A. Chambellan, and A.T. Dinh-Xuan. Combined measurement of carbon monoxide and nitric oxide lung transfer does not

improve the identification of pulmonary hypertension in systemic sclerosis. *European Respiratory Journal*, 50(4):1701008, 2017.

[118] S.R. Hopkins. Ventilation/perfusion relationships and gas exchange: Measurement approaches. *Compr Physiol*, 10(3):1155–1205, 2020.

[119] Y. Jernudd-Wilhelmsson, Y. Hornblad, and G. Hedenstierna. Ventilation-perfusion relationships in interstitial lung disease. *Eur J Respir Dis*, 68(1):39–49, 1986.

[120] A.G. Agustí, J. Roca, J. Gea, P.D. Wagner, A. Xaubet, and R. Rodriguez-Roisin. Mechanisms of gas-exchange impairment in idiopathic pulmonary fibrosis. *Am Rev Respir Dis*, 143(2):219–25, 1991.

[121] P.D. Wagner. The multiple inert gas elimination technique (MIGET). *Intensive Care Med*, 34(6):994–1001, 2008.

[122] A. Sanchez-Crespo, J. Petersson, S. Nyren, M. Mure, R.W. Glenny, J.O. Thorell, H. Jacobsson, S.G. Lindahl, and S.A. Larsson. A novel quantitative dual-isotope method for simultaneous ventilation and perfusion lung spet. *Eur J Nucl Med Mol Imaging*, 29(7):863–75, 2002.

[123] J. Petersson, A. Sanchez-Crespo, M. Rohdin, S. Montmerle, S. Nyren, H. Jacobsson, S.A. Larsson, S.G. Lindahl, D. Linnarsson, R.W. Glenny, and M. Mure. Physiological evaluation of a new quantitative SPECT method measuring regional ventilation and perfusion. *J Appl Physiol (1985)*, 96(3):1127–36, 2004.

[124] G. Musch and J.G. Venegas. Positron emission tomography imaging of regional pulmonary perfusion and ventilation. *Proc Am Thorac Soc*, 2(6):522–7, 508–9, 2005.

[125] G. Musch and J.G. Venegas. Positron emission tomography imaging of regional lung function. *Minerva Anestesiol*, 72(6):363–7, 2006.

[126] E.A. Hoffman and D. Chon. Computed tomography studies of lung ventilation and perfusion. *Proc Am Thorac Soc*, 2(6):492–8, 506, 2005.

[127] C. Do, J. Deaguero, A. Brearley, X. Trejo, T. Howard, G.P. Escobar, and B. Wagner. Gadolinium-based contrast agent use, their safety, and practice evolution. *Kidney360*, 1(6):561–568, 2020.

[128] N.J. Stewart, L.J. Smith, H.-F. Chan, J.A. Eaden, S. Rajaram, A.J. Swift, N.D. Weatherley, A. Biancardi, G.J. Collier, D. Hughes, G. Klafkowski, C.S. Johns, N. West, K. Ugonna, S.M. Bianchi, R. Lawson, I. Sabroe, H. Marshall, and J.M. Wild. Lung MRI with hyperpolarised gases: current & future clinical perspectives. *The British Journal of Radiology*, 95(1132):20210207, 2022.

[129] M.S. Albert, V.D. Schepkin, and T.F. Budinger. Measurement of 129Xe T1 in blood to explore the feasibility of hyperpolarized 129Xe MRI. *J Comput Assist Tomogr*, 19(6):975–8, 1995.

[130] M.E. Wagshul, T.M. Button, H.F. Li, Z. Liang, C.S. Springer, K. Zhong, and A. Wishnia. In vivo MR imaging and spectroscopy using hyperpolarized129Xe. *Magn Reson Med*, 36(2):183–191, 1996.

[131] K. Sakai, A.M. Bilek, E. Oteiza, R.L. Walsworth, D. Balamore, F.A. Jolesz, and M.S. Albert. Temporal dynamics of hyperpolarized 129Xe resonances in living rats. *J Magn Reson B*, 111(3):300–4, 1996.

[132] S.D. Swanson, M.S. Rosen, B.W. Agranoff, K.P. Coulter, R.C. Welsh, and T.E. Chupp. Brain MRI with laser-polarized 129Xe. *Magn Reson Med*, 38(5):695–8, 1997.

[133] 3rd J.P. Mugler, B. Driehuys, J.R. Brookeman, G.D. Cates, S.S. Berr, R.G. Bryant, T.M. Daniel, E.E. de Lange, 3rd J.H. Downs, C.J. Erickson, W. Happer, D.P. Hinton, N.F. Kassel, T. Maier, C.D. Phillips, B.T. Saam, K.L. Sauer, and M.E. Wagshul. Mr imaging and spectroscopy using hyperpolarized 129Xe gas: preliminary human results. *Magn Reson Med*, 37(6):809–15, 1997.

[134] S.S. Kaushik, M.S. Freeman, S.W. Yoon, M.G. Liljeroth, J.V. Stiles, J.E. Roos, W. Foster, C.R. Rackley, H.P. McAdams, and B. Driehuys. Measuring diffusion limitation with a perfusion-limited gashyperpolarized 129Xe gas-transfer spectroscopy in patients with idiopathic pulmonary fibrosis. *J Appl Physiol*, 117:577 – 585, 2014.

[135] N.J. Stewart, N.S. Higano, S. Mukthapuram, M.M. Willmering, W. Loew, M. West, A. Arnsperger, R. Pratt, M.R. Rao, R.F. Schulte, J.M. Wild, and J.C. Woods. Initial feasibility and challenges of hyperpolarized Xe MRI in neonates with bronchopulmonary dysplasia. *Magn Reson Med*, 90(6):2420–2431, 2023.

[136] M. Kocher Wulfeck, D.G. Mummy, S. Zhang, D. Kadi, S. Goins, A. Swaminathan, R.M. Tighe, B. Driehuys, and J.G. Mammarappallil. Hyperpolarized 129Xe MRI and spectroscopy of gas exchange abnormalities in chronic hypersensitivity pneumonitis. *Respiratory Medicine*, 234:107827, 2024.

[137] E.A. Bier, S.H. Robertson, G.M. Schrank, C. Rackley, J.G. Mammarappallil, S. Rajagopal, H.P. McAdams, and B. Driehuys. A protocol for quantifying cardiogenic oscillations in dynamic (129) Xe gas exchange spectroscopy: The effects of idiopathic pulmonary fibrosis. *NMR Biomed*, 32(1):e4029, 2019.

[138] J. Mata, S. Guan, K. Qing, N. Tustison, Y. Shim, J.P. Mugler, T. Altes, J. Huaromo, and B. Mehrad. Evaluation of regional lung function in pulmonary fibrosis with xenon-129 MRI. *Tomography*, 7(3):452–465, 2021.

[139] S. Guan, N. Tustison, K. Qing, Y.M. Shim, J. Mugler, T. Altes, D. Albon, D. Froh, B. Mehrad, J. Patrie, A. Ropp, B. Miller, J. Nehrbas, and J. Mata. 3D single-breath chemical shift imaging hyperpolarized Xe-129 MRI of healthy, CF, IPF, and COPD subjects. *Tomography*, 8(5):2574–2587, 2022.

[140] J. Wolber, A. Cherubini, D. Santoro, G.S. Payne, M.O. Leach, and A. Bifone. Linewidths of hyperpolarized 129Xe NMR spectra in human blood at 1.5 T. In *Proc 8th Intl Soc Magn Reson Med*, volume 8, page 970, 2000.

[141] N.J. Stewart, F.C. Horn, G. Norquay, G.J. Collier, D.P. Yates, R. Lawson, H. Marshall, and J.M. Wild. Reproducibility of quantitative indices of lung function and microstructure from 129Xe chemical shift saturation recovery (CSSR) MR spectroscopy. *Magn Reson Med*, 77(6):2107–2113, 2017.

[142] S. Patz, I. Muradyan, M.I. Hrovat, M. Dabaghyan, G.R. Washko, H. Hatabu, and J.P. Butler. Diffusion of hyperpolarized 129Xe in the lung: a simplified model of 129Xe septal uptake and experimental results. *New Journal of Physics*, 13(015009), 2011.

[143] K. Ruppert, K. Qing, J.T. Patrie, T.A. Altes, and J.P. Mugler III. Using hyperpolarized xenon-129 MRI to quantify early-stage lung disease in smokers. *Academic Radiology*, 26:355–366, 2019.

[144] A.L. Kern, D.-H. Park, J. Fuge, J.M. Hohlfeld, F. Wacker, M.M. Hoeper, K.M. Olsson, and J. Vogel-Claussen. Loss of pulmonary capillaries in idiopathic pulmonary arterial hypertension with low diffusion capacity is accompanied by early diffuse emphysema detected by 129Xe MRI, 2024. https://-doi.org/10.1007/s00330-024-11209-1.

[145] M. Zhang, H. Li, H. Li, X. Zhao, Q. Zhou, Q. Rao, Y. Han, Y. Lan, H. Deng, X. Sun, X. Lou, C. Ye, and X. Zhou. Quantitative evaluation of lung injury caused by PM2.5 using hyperpolarized gas magnetic resonance. *Magn Reson Med*, 84(2):569–578, 2020.

[146] J.P. Zhong, H.T. Zhang, W.W. Ruan, J.S. Xie, H.D. Li, H. Deng, Y.Q. Han, X.P. Sun, C.H. Ye, and X. Zhou. Simultaneous assessment of both lung morphometry and gas exchange function within a single breath-hold by hyperpolarized Xe-129 MRI. *NMR in Biomedicine*, 30(8), 2017.

[147] J. Xie, H. Li, H. Zhang, X. Zhao, L. Shi, M. Zhang, S. Xiao, H. Deng, K. Wang, H. Yang, X. Sun, G. Wu, C. Ye, and X. Zhou. Single breath-hold measurement of pulmonary gas exchange and diffusion in humans with hyperpolarized (129) Xe MR. *NMR Biomed*, 32(5):e4068, 2019.

[148] O. Doganay, E. Stirrat, C. McKenzie, R.F. Schulte, and G.E. Santyr. Quantification of regional early stage gas exchange changes using hyperpolarized Xe-129 MRI in a rat model of radiation-induced lung injury. *Medical Physics*, 43(5), 2016.

[149] O. Doganay, M. Chen, T. Matin, M. Rigolli, J.A. Phillips, A. McIntyre, and F.V. Gleeson. Magnetic resonance imaging of the time course of hyperpolarized Xe-129 gas exchange in the human lungs and heart. *European Radiology*, 29(5):2283–2292, 2019.

[150] B. Zanette and G. Santyr. Accelerated interleaved spiral-IDEAL imaging of hyperpolarized Xe-129 for parametric gas exchange mapping in humans. *Magn Reson Med*, 82(3):1113–1119, 2019.

[151] Y. Friedlander, B. Zanette, A. Lindenmaier, D. Li, S. Kadlecek, G. Santyr, and A. Kassner. Hyperpolarized Xe-129 MRI of the rat brain with chemical shift saturation recovery and spiral-IDEAL readout. *Magn Reson Med*, 2021.

[152] L. An, S. Warach, and J. Shen. Spectral localization by imaging using multielement receiver coils. *Magn Reson Med*, 66(1):1–10, 2011.

[153] L. An and J. Shen. Image-guided spatial localization of heterogeneous compartments for magnetic resonance. *Medical Physics*, 42(9):5278–5286, 2015.

[154] A.L. Kern, M. Gutberlet, K. Qing, A. Voskrebenzev, F. Klime, T.F. Kaireit, C. Czerner, H. Biller, F. Wacker, K. Ruppert, J.M. Hohlfeld, and J. Vogel-Claussen. Regional investigation of lung function and microstructure parameters by localized 129 Xe chemical shift saturation recovery and dissolved-phase imaging: A reproducibility study. *Magn Reson Med*, 81(1):13–24, 2019.

[155] A.L. Kern, M. Gutberlet, A. Voskrebenzev, F. Klimes, A. Rotarmel, F. Wacker, J.M. Hohlfeld, and J. Vogel-Claussen. Mapping of regional lung microstructural parameters using hyperpolarized Xe-129 dissolved-phase MRI in healthy volunteers and patients with chronic obstructive pulmonary disease. *Magn Reson Med*, 81(4):2360–2373, 2019.

[156] D.C. Look and D.R. Locker. Time saving in measurement of NMR and EPR relaxation times. *Rev Sci Instrum.*, 41:250 – 251, 1970.

[157] R. Otazo, E. Candes, and D.K. Sodickson. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. *Magn Reson Med*, 73(3):1125–1136, 2015.

[158] N.J. Stewart. *Hyperpolarised 129Xe Magnetic Resonance Imaging Techniques for Assessment of Human Lung Function*. Phd thesis, University of Sheffield, Sheffield, UK, August 2016.

[159] K. Ruppert, L. Loza, H. Hamedani, M. Ismail, J. Chen, I.F. Duncan, H. Profka, S. Kadlecek, and R.R. Rizi. Regional variations in hyperpolarized 129Xe lung MRI: Insights from CSI-CSSR and CSSR in healthy and irradiated rat models. *Magn Reson Med*, 93(3):902–915, 2025.

[160] Z.I. Cleveland, G.P. Cofer, G. Metz, D. Beaver, J. Nouls, S.S. Kaushik, M. Kraft, J. Wolber, K.T. Kelly, H.P. McAdams, and B. Driehuys. Hyperpolarized Xe-129 MR imaging of alveolar gas uptake in humans. *Plos One*, 5(8), 2010.

[161] J.P. Mugler, T.A. Altes, I.C. Ruset, I.M. Dregely, J.F. Mata, G.W. Miller, S. Ketel, J. Ketel, F.W. Hersman, and K. Ruppert. Simultaneous magnetic resonance imaging of ventilation distribution and gas uptake in the human lung using hyperpolarized xenon-129. *Proc Natl Acad Sci U S A*, 107(50):21707–21712, 2010.

[162] S.D. Swanson, M.S. Rosen, K.P. Coulter, R.C. Welsh, and T.E. Chupp. Distribution and dynamics of laser-polarized 129Xe magnetization in vivo. *Magn Reson Med*, 42:1137–1145, 1999.

[163] W.T. Dixon. Simple proton spectroscopic imaging. *Radiology*, 153(1):189–194, 1984.

[164] P.J. Niedbalski, C.S. Hall, M. Castro, R.L. Eddy, J.H. Rayment, S. Svenningsen, G. Parraga, B. Zanette, G.E. Santyr, R.P. Thomen, N.J. Stewart, G.J. Collier, H.F. Chan, J.M. Wild, S.B. Fain, G.W. Miller, J.F. Mata, J.P. Mugler, B. Driehuys, M.M. Willmering, Z.I. Cleveland, and J.C. Woods. Protocols for multi-site trials using hyperpolarized Xe-129 MRI for imaging of ventilation, alveolar-airspace size, and gas exchange: A position paper from the Xe-129 MRI clinical trials consortium. *Magn Reson Med*, 2021.

[165] A.D. Hahn, J. Kammerman, M. Evans, W. Zha, R.V. Cadman, K. Meyer, N. Sandbo, and S.B. Fain. Repeatability of regional pulmonary functional metrics of hyperpolarized Xe-129 dissolved-phase MRI. *J Magn Reson Imaging*, 50(4):1182–1190, 2019.

[166] A.M. Matheson, M.J. McIntosh, H.K. Kooner, J. Lee, V. Desaigoudar, E. Bier, B. Driehuys, S. Svenningsen, G.E. Santyr, M. Kirby, M.S. Albert, Y. Shepelytskyi, V. Grynko, A. Ouriadov, M. Abdelrazek, I. Dhaliwal, J.M. Nicholson, and G. Parraga. Persistent 129Xe MRI pulmonary and CT vascular abnormalities in symptomatic individuals with post-acute COVID-19 syndrome. *Radiology*, 305(2):466–476, 2022.

[167] R.L. Eddy, D. Mummy, S. Zhang, H. Dai, A. Bechtel, A. Schmidt, B. Frizzell, F.V. Gerayeli, J.A. Leipsic, J.M. Leung, B. Driehuys, L.G. Que, M. Castro, D.D. Sin, and P.J. Niedbalski. Cluster analysis to identify long COVID phenotypes using 129Xe magnetic resonance imaging: a multicentre evaluation. *European Respiratory Journal*, 63(3), 2024.

[168] Z.Y. Wang, S.H. Robertson, J. Wang, M. He, R.S. Virgincar, G.M. Schrank, E.A. Bier, S. Rajagopal, Y.C. Huang, T.G. O'Riordan, C.R. Rackley, H.P. McAdams, and B. Driehuys. Quantitative analysis of hyperpolarized Xe-129 gas transfer MRI. *Medical Physics*, 44(6):2415–2428, 2017.

[169] D.G. Mummy, E.A. Bier, Z.Y. Wang, J. Korzekwinski, L. Morrison, C. Barkauskas, H.P. McAdams, R.M. Tighe, B. Driehuys, and J.G. Mammarappallil. Hyperpolarized Xe-129 MRI and spectroscopy of gas-exchange abnormalities in nonspecific interstitial pneumonia. *Radiology*, 301(1):211–220, 2021.

[170] A.D. Hahn, K.J. Carey, G.P. Barton, L.A. Torres, J. Kammerman, R.V. Cadman, K.E. Lee, M.L. Schiebler, N. Sandbo, and S.B. Fain. Hyperpolarized 129Xe MR spectroscopy in the lung shows 1-year reduced function in idiopathic pulmonary fibrosis. *Radiology*, 305(3):688–696, 2022. PMID: 35880982.

[171] M.M. Willmering, L.L. Walkup, P.J. Niedbalski, H. Wang, Z.Y. Wang, E.B. Hysinger, K.C. Myers, C.T. Towe, B. Driehuys, Z.I. Cleveland, and J.C. Woods. Pediatric Xe-129 gas-transfer MRI-feasibility and applicability. *J Magn Reson Imaging*, 2022.

[172] M.A. Anderson, B.I. Masokano, J.W. Plummer, M.M. Willmering, K. Krivchenia, C.T. Towe, and L.L. Walkup. Xe ventilation and gas-exchange MRI abnormalities in children with neuroendocrine cell hyperplasia of infancy (NEHI). *Pediatric Pulmonology*, 60(1):e27410, 2025.

[173] L.J. Rankine, Z. Wang, C.R. Kelsey, E. Bier, B. Driehuys, L.B. Marks, and S.K. Das. Hyperpolarized 129Xe magnetic resonance imaging for functional avoidance treatment planning in thoracic radiation therapy: A comparison of ventilation- and gas exchange-guided treatment plans. *Int J Radiat Oncol Biol Phys*, 111(4):1044–1057, 2021.

[174] P.J. Niedbalski, M.M. Willmering, R.P. Thomen, J.P. Mugler III, J. Choi, C. Hall, and M. Castro. A single-breath-hold protocol for hyperpolarized 129Xe ventilation and gas exchange imaging. *NMR in Biomedicine*, 36(8):e4923, 2023.

[175] M.M. Willmering, Z.I. Cleveland, L.L. Walkup, and J.C. Woods. Removal of off-resonance xenon gas artifacts in pulmonary gas-transfer MRI. *Magn Reson Med*, 86:907–915, 2021.

[176] A.D. Hahn, J. Kammerman, and S.B. Fain. Removal of hyperpolarized Xe-129 gas-phase contamination in spectroscopic imaging of the lungs. *Magn Reson Med*, 80(6):2586–2597, 2018.

[177] G.H. Glover and E. Schneider. Three-point Dixon technique for true water/fat decomposition with B0 inhomogeneity correction. *Magn Reson Med*, 18(2):371–83, 1991.

[178] I. Muradian, S. Patz, J.P. Butler, G.P. Topulos, M.I. Hrovat, S. Covrig, S. Ketel, I. Russet, and F.W. Hersman. Hyperpolarized 129Xe human pulmonary gas exchange with 3-point Dixon technique. In *Proc. 14th Intl. Soc. Mag. Reson. Med.*, page 1297, Seattle, Washington, USA, 2006.

[179] S.B. Reeder, A.R. Pineda, Z. Wen, A. Shimakawa, H. Yu, J.H. Brittain, G.E. Gold, C.H. Beaulieu, and N.J. Pelc. Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): application with fast spin-echo imaging. *Magn Reson Med*, 54(3):636–44, 2005.

[180] S.B. Reeder, J.H. Brittain, T.M. Grist, and Y.-F. Yen. Least-squares chemical shift separation for 13C metabolic imaging. *J Magn Reson Imaging*, 26(4):1145–1152, 2007.

[181] B. Zanette, E. Stirrat, S. Jelveh, A. Hope, and G. Santyr. Physiological gas exchange mapping of hyperpolarized Xe-129 using spiral-IDEAL and MOXE in a model of regional radiation-induced lung injury. *Medical Physics*, 45(2):803–816, 2018.

[182] J. Tsao and Y. Jiang. Hierarchical IDEAL: robust waterfat separation at high field by multiresolution field map estimation. In *Proc. 16th Intl. Soc. Mag. Reson. Med.*, page 653, 2008.

[183] J. Tsao and Y. Jiang. Hierarchical IDEAL: Fast, robust, and multiresolution separation of multiple chemical species from multiple echo times. *Magn Reson Med*, 70(1):155–159, 2013.

[184] K. Qing, N.J. Tustison, J.P. Mugler, J.F. Mata, Z. Lin, L. Zhao, D. Wang, X. Feng, J.Y. Shin, S.J. Callahan, M.P. Bergman, K. Ruppert, T.A. Altes, J.M. Cassani, and Y.M. Shim. Probing changes in lung physiology in COPD using CT, perfusion MRI, and hyperpolarized xenon-129 MRI. *Academic Radiology*, 26(3):326–334, 2019.

[185] L. Myc, K. Qing, M. He, N. Tustison, Z. Lin, A.W. Manichaikul, J. Patrie, J. Cassani, R.N. Nunoo-Asare, Y. Huang, Z. Obaida, S. Tafti, A.M. Ropp, G.W. Miller, J. Mata, T. Altes, J. Mugler, and Y.M. Shim. Characterisation of gas exchange in COPD with dissolved-phase hyperpolarised xenon-129 MRI. *Thorax*, 76(2):178–181, 2020.

[186] M. He, K. Qing, N.J. Tustison, Z. Beaulac, T.W. King, T.B. Huff, M. Paige, K. Earasi, R. Nunoo-Asare, S. Struchen, M. Burdick, Z. Zhang, A. Ropp, G.W. Miller, J.T. Patrie, J.F. Mata, J.P. Mugler 3rd, and Y.M. Shim. Characterizing gas exchange physiology in healthy young electronic-cigarette users with hyperpolarized 129Xe MRI: A pilot study. *Int J Chron Obstruct Pulmon Dis*, 16:3183–3187, 2021.

[187] K. Qing, T.A. Altes, J.P. Mugler III, N.J. Tustison, J. Mata, K. Ruppert, P. Komlosi, X. Feng, K. Nie, L. Zhao, Z. Wang, W.F. Hersman, I.C. Ruset, B. Liu, Y.M. Shim, and W.G. Teague. Pulmonary MRI with hyperpolarized xenon-129 demonstrates novel alterations in gas transfer across the air-blood barrier in asthma. *Medical Physics*, 51(4):2413–2423, 2024.

[188] J. Kammerman, A.D. Hahn, R.V. Cadman, A. Malkus, D. Mummy, and S.B. Fain. Transverse relaxation rates of pulmonary dissolved-phase hyperpolarized Xe-129 as a biomarker of lung injury in idiopathic pulmonary fibrosis. *Magn Reson Med*, 84(4):1857–1867, 2020.

[189] G.J. Collier, L.J. Smith, L.C. Saunders, A.J. Swift, H. Marshall, N.J. Stewart, G. Norquay, P.J.C. Hughes, A.A.R. Thomspson, and J.M. Wild. Age, sex, and lung volume dependence of dissolved xenon-129 mri gas exchange metrics. *Magn Reson Med*, 92(4):1471–1483, 2024.

[190] L.C. Saunders, G.J. Collier, H.-F. Chan, P.J.C. Hughes, L.J. Smith, J. Watson, J. Meiring, Z. Gabriel, T. Newman, M. Plowright, P. Wade, J.A. Eaden, S. Thomas, S. Strickland, L. Gustafsson, J. Bray, H. Marshall, D.J. Capener, L. Armstrong, J. Rodgers, M. Brook, A.M. Biancardi, M.R. Rao, G. Norquay, O. Rodgers, R. Munro, J.E. Ball, N.J. Stewart, A. Lawrie, G. Jenkins, J. Grist, F. Gleeson, R.F. Schulte, K.M. Johnson, F.J. Wilson, A. Cahn, A.J. Swift, S. Rajaram, G.H. Mills, L. Watson, P.J. Collini, R. Lawson, A.A.R.

Thompson, and J.M. Wild. Longitudinal lung function assessment of patients hospitalized with COVID-19 using 1H and 129Xe lung MRI. *Chest*, 164(3):700–716, 2023.

[191] H. Marshall, L. Smith, A. Biancardi, G. Collier, H.-F. Chan, P. Hughes, J. Astley, R. Munro, S. Rajaram, D. Capener, J. Bray, I. Smith, L. Armstrong, L. Hardaker, T. Fihn-Wikander, R. Hughes, and J. Wild. A comparison of 129Xe MRI and advanced lung function testing in patients with asthma and /or COPD: The NOVELTY ADPro substudy. *European Respiratory Journal*, 58(suppl 65):PA1872, 2021.

[192] N.D. Weatherley, N.J. Stewart, H.-F. Chan, M. Austin, L.J. Smith, G. Collier, M. Rao, H. Marshall, G. Norquay, S.A. Renshaw, S.M. Bianchi, and J.M. Wild. Hyperpolarised xenon magnetic resonance spectroscopy for the longitudinal assessment of changes in gas diffusion in IPF. *Thorax*, 74(5):500–502, 2019.

[193] J.T. Grist, M. Chen, G.J. Collier, B. Raman, G. Abueid, A. McIntyre, V. Matthews, E. Fraser, L.P. Ho, J.M. Wild, and F. Gleeson. Hyperpolarized Xe-129 MRI abnormalities in dyspneic patients 3 months after COVID-19 pneumonia: Preliminary results. *Radiology*, 301(1):1353–E360, 2021.

[194] J.T. Grist, G.J. Collier, H. Walters, M. Kim, M. Chen, G.A. Eid, A. Laws, V. Matthews, K. Jacob, S. Cross, A. Eves, M. Durrant, A. McIntyre, R. Thompson, R.F. Schulte, B. Raman, P.A. Robbins, J.M. Wild, E. Fraser, and F. Gleeson. Lung abnormalities detected with hyperpolarized 129Xe MRI in patients with long COVID. *Radiology*, 305(3):709–717, 2022.

[195] K. Ruppert, J.F. Mata, J.R. Brookeman, K.D. Hagspiel, and J.P. Mugler. Exploring lung function with hyperpolarized Xe-129 nuclear magnetic resonance. *Magn Reson Med*, 51(4):676–687, 2004.

[196] K. Ruppert, J.F. Mata, H.T.J. Wang, W.A. Tobias, G.D. Cates, J.R. Brookeman, K.D. Hagspiel, and J.P. Mugler. XTC MRI: Sensitivity improvement through parameter optimization. *Magn Reson Med*, 57(6):1099–1109, 2007.

[197] I. Dregely, J.P. Mugler, I.C. Ruset, T.A. Altes, J.F. Mata, G.W. Miller, J. Ketel, S. Ketel, J. Distelbrink, F.W. Hersman, and K. Ruppert. Hyperpolarized xenon-129 gas-exchange imaging of lung microstructure: First case studies in subjects with obstructive lung disease. *J Magn Reson Imaging*, 33(5):1052–1062, 2011.

[198] I. Dregely, I.C. Ruset, J.F. Mata, J. Ketel, S. Ketel, J. Distelbrink, T.A. Altes, J.P. Mugler, G.W. Miller, F.W. Hersman, and K. Ruppert. Multiple-exchange-time xenon polarization transfer contrast (MXTC) MRI: Initial results in animals and healthy volunteers. *Magn Reson Med*, 67(4):943–953, 2012.

[199] I. Muradyan, J.P. Butler, M. Dabaghyan, M. Hrovat, I. Dregely, I. Ruset, G.P. Topulos, E. Frederick, H. Hatabu, W.F. Hersman, and S. Patz. Single-breath xenon polarization transfer contrast (SB-XTC): Implementation and

initial results in healthy humans. *J Magn Reson Imaging*, 37(2):457–470, 2013.

[200] F. Amzajerdian, K. Ruppert, H. Hamedani, R. Baron, Y. Xin, L. Loza, T. Achekzai, I.F. Duncan, Y.W. Qian, M. Pourfathi, S. Kadlecek, and R.R. Rizi. Measuring pulmonary gas exchange using compartment-selective xenon-polarization transfer contrast (XTC) MRI. *Magn Reson Med*, 85(5):2709–2722, 2021.

[201] F. Amzajerdian, H. Hamedani, R. Baron, L. Loza, I. Duncan, K. Ruppert, S. Kadlecek, and R. Rizi. Simultaneous quantification of hyperpolarized xenon-129 ventilation and gas exchange with multi-breath xenon-polarization transfer contrast (XTC) MRI. *Magn Reson Med*, 90(6):2334–2347, 2023.

[202] S. Mänsson, J. Wolber, B. Driehuys, P. Wollmer, and K. Golman. Characterization of diffusing capacity and perfusion of the rat lung in a lipopolysaccaride disease model using hyperpolarized Xe-129. *Magn Reson Med*, 50(6):1170–1179, 2003.

[203] B. Driehuys, G.P. Cofer, J. Pollaro, J.B. Mackel, L.W. Hedlund, and G.A. Johnson. Imaging alveolar-capillary gas transfer using hyperpolarized 129Xe MRI. *Proc Natl Acad Sci U S A*, 103(48):18278–18283, 2006.

[204] Y.V. Chang. MOXE: A model of gas exchange for hyperpolarized 129 Xe magnetic resonance of the lung. *Magn Reson Med*, 69(3):884–890, 2013.

[205] A.L. Kern, K.M. Olsson, T.F. Till Kaireit, F.K. Wacker, J.M. Hohlfeld, M.M. Hoeper, and J. Vogel-Claussen. Assessment of interstitial membrane permeability using a generalized model of 129Xe septal uptake in the lung. In *Proc. 28th Intl. Soc. Mag. Reson. Med.*, page 2371, 2020.

[206] Y. Taylor, G.J. Parker, M. Kim, and P. Jimenez-Royo. A novel compartmental model of 129Xe gas exchange using blind estimation of blood concentration. In *Proc. 31st Intl. Soc. Mag. Reson. Med. 4496*, volume 31, page 4496, 2023.

[207] S. Perron and A. Ouriadov. Hyperpolarized 129Xe MRI at low field: Current status and future directions. *J Magn Reson*, 348:107387, 2023.

[208] A.K. Venkatesh, K.S. Hong, L. Kubatina, Y. Sun, R.V. Mulkern, F.A. Jolesz, and M.S. Albert. Direct observation of the transport of 129Xe from the lung-gas to the tissue and the blood. In *Proc. 9th Intl. Soc. Mag. Reson. Med*, volume 9, page 954, 2001.

[209] E. A. Bier, F. Alenezi, J. Lu, Z. Wang, J. G. Mammarappallil, B. O'Sullivan-Murphy, A. Erkanli, B. Driehuys, and S. Rajagopal. Noninvasive diagnosis of pulmonary hypertension with hyperpolarised 129Xe magnetic resonance imaging and spectroscopy. *ERJ Open Res*, 8(2), 2022.

[210] G. Norquay, G.J. Collier, and J.M. Wild. Temporal correlation of alveolar-capillary 129xe signal dynamics with the cardiac cycle. In *Proc. 29th Intl. Soc. Mag. Reson. Med. 3564*, volume 29, 2021.

[211] J. Alastruey, P.H. Charlton, V. Bikia, B. Paliakaite, B. Hametner, R.M. Bruno, M.P. Mulder, S. Vennin, S. Piskin, A.W. Khir, A. Guala, C.C. Mayer, J. Mynard, A.D. Hughes, P. Segers, and B.E. Westerhof. Arterial pulse wave modeling and analysis for vascular-age studies: a review from VascAgeNet. *Am J Physiol Heart Circ Physiol*, 325(1):H1–h29, 2023.

[212] K. Ruppert, T. A. Altes, J. F. Mata, I. C. Ruset, F. W. Hersman, and J. P. Mugler III. Detecting pulmonary capillary blood pulsations using hyperpolarized xenon-129 chemical shift saturation recovery (CSSR) MR spectroscopy. *Magn Reson Med*, 75(4):1771–1780, 2016.

[213] P.J. Niedbalski, E.A. Bier, Z. Wang, M.M. Willmering, B. Driehuys, and Z.I. Cleveland. Mapping cardiopulmonary dynamics within the microvasculature of the lungs using dissolved 129Xe MRI. *J Appl Physiol*, 129(2):218–229, 2020.

[214] A. Costelle, J. Lu, S. Leewiwatwong, B.E. Westerhof, D. Mummy, S. Rajagopal, and B. Driehuys. Combining hyperpolarized 129Xe MR imaging and spectroscopy to noninvasively estimate pulmonary vascular resistance. *J Appl Physiol*, 138(3):623–633, 2025. PMID: 39873409.

[215] J. Lu, F. Alenezi, E. Bier, S. Leewiwatwong, D. Mummy, S. Kabir, S. Rajagopal, S. Robertson, P. J. Niedbalski, and B. Driehuys. Optimized quantitative mapping of cardiopulmonary oscillations using hyperpolarized 129Xe gas exchange MRI: Digital phantoms and clinical evaluation in CTEPH. *Magn Reson Med*, 91(4):1541–1555, 2024.

[216] R. Lethmate, F.T.A.W. Wajer, Y. Crémillieux, D. van Ormondt, and D. Graveron-Demilly. Dynamic MR-imaging with radial scanning, a post-acquisition keyhole approach. *EURASIP Journal on Advances in Signal Processing*, 2003(5):794878, 2003.

[217] G. Norquay, G.J. Collier, M. Rao, N.J. Stewart, and J.M. Wild. ^{129}Xe-rb spin-exchange optical pumping with high photon efficiency. *Phys. Rev. Lett.*, 121:153201, Oct 2018.

[218] L.J. Smith, G.J. Collier, H. Marshall, P.J.C. Hughes, A.M. Biancardi, M. Wildman, I. Aldag, N. West, A. Horsley, and J.M. Wild. Patterns of regional lung physiology in cystic fibrosis using ventilation magnetic resonance imaging and multiple-breath washout. *European Respiratory Journal*, 52(5):1800821, 2018.

[219] R.F. Schulte, L. Sacolick, M.H. Deppe, M.A. Janich, M. Schwaiger, J.M. Wild, and et al. Transmit gain calibration for nonproton MR using the Bloch-Siegert shift. *NMR in Biomedicine*, 24:1068–1072, 2011.

[220] F. Bloch and A. Siegert. Magnetic resonance for nonrotating fields. *Phys. Rev.*, 57:522–527, Mar 1940.

[221] F. Wiesinger, E. Weidl, M.I. Menzel, M.A. Janich, O. Khegai, S.J. Glaser, A. Haase, M. Schwaiger, and R.F. Schulte. IDEAL spiral CSI for dynamic

metabolic MR imaging of hyperpolarized [1-13C]pyruvate. *Magn Reson Med*, 68(1):8–16, 2012.

[222] P.J. Niedbalski, J.L. Lu, C.S. Hall, M. Castro, J.P. Mugler, Y.M. Shim, and B. Driehuys. Utilizing flip angle/TR equivalence to reduce breath hold duration in hyperpolarized Xe-129 1-point Dixon gas exchange imaging. *Magn Reson Med*, 2021. Preprint. 10.1002/mrm.29040.

[223] K. Ruppert, F. Amzajerdian, H. Hamedani, Y. Xin, L. Loza, T. Achekzai, I.F. Duncan, H. Profka, S. Siddiqui, M. Pourfathi, F. Sertic, M.F. Cereda, S. Kadlecek, and R.R. Rizi. Assessment of flip angle-TR equivalence for standardized dissolved-phase imaging of the lung with hyperpolarized 129Xe MRI. *Magn Reson Med*, 81(3):1784–1794, 2019.

[224] Z.Y. Wang, L. Rankine, E.A. Bier, D. Mummy, J.L. Lu, A. Church, R.M. Tighe, A. Swaminathan, Y.C.T. Huang, L.G. Que, J.G. Mammarappallil, S. Rajagopal, and B. Driehuys. Using hyperpolarized Xe-129 gas-exchange MRI to model the regional airspace, membrane, and capillary contributions to diffusing capacity. *J Appl Physiol*, 130(5):1398–1409, 2021.

[225] B. Aguilaniu, J. Maitre, S. Glénet, A. Gegout-Petit, and H. Guénard. European reference equations for CO and NO lung transfer. *European Respiratory Journal*, 31(5):1091–1097, 2008.

[226] P.H. Quanjer, S. Stanojevic, T.J. Cole, X. Baur, G.L. Hall, B.H. Culver, P.L. Enright, J.L. Hankinson, M. S. M. Ip, J. Zheng, J. Stocks, and t. E. G. L. F. Initiative. Multi-ethnic reference values for spirometry for the 395-yr age range: the global lung function 2012 equations. *European Respiratory Journal*, 40(6):1324–1343, 2012.

[227] S. Stanojevic, B.L. Graham, B.G. Cooper, B.R. Thompson, K.W. Carter, R.W. Francis, and G.L. Hall. Official ERS technical standards: Global lung function initiative reference values for the carbon monoxide transfer factor for caucasians. *European Respiratory Journal*, 50(3):1700010, 2017.

[228] M. Munkholm, L. Jacob J.L. Marott, L. Bjerre-Kristensen, F. Madsen, O.F. Pedersen, P. Lange, B.G. Nordestgaard, and J. Mortensen. Reference equations for pulmonary diffusing capacity of carbon monoxide and nitric oxide in adult caucasians. *European Respiratory Journal*, 52(1), 2018.

[229] J.W. Plummer, M.M. Willmering, Z.I. Cleveland, C. Towe, J.C. Woods, and L.L. Walkup. Childhood to adulthood: Accounting for age dependence in healthy-reference distributions in 129Xe gas-exchange MRI. *Magn Reson Med*, 89(3):1117–1133, 2023.

[230] D. Mummy, S. Zhang, A. Bechtel, J. Lu, J. Mammarappallil, S. Leewiwatwong, A. Costelle, A. Swaminathan, and B. Driehuys. Functional gas exchange measures on (129)Xe MRI and spectroscopy are associated with age, sex, and BMI in healthy subjects. *Frontiers in Medicine (Lausanne)*, 11:1342499, 2024.

[231] J.B. West. *Respiratory Physiology*, chapter 3. Lippencott Williams and Wilkins, Baltimore, MD, 8th edition, 2008.

[232] N. Woodhouse, J.M. Wild, M.N.J. Paley, S. Fichele, Z. Said, A.J. Swift, and E.J.R. van Beek. Combined helium-3/proton magnetic resonance imaging measurement of ventilated lung volumes in smokers compared to never-smokers. *J Magn Reson Imaging*, 21(4):365–369, 2005.

[233] H. K. Reddel, M. Gerhardsson de Verdier, A. Agustí, G. Anderson, R. Beasley, E. H. Bel, C. Janson, B. Make, R. J. Martin, I. Pavord, D. Price, C. Keen, A. Gardev, S. Rennard, A. Sveréus, A. T. Bansal, L. Brannman, N. Karlsson, J. Nuevo, F. Nyberg, S. S. Young, and J. Vestbo. Prospective observational study in patients with obstructive lung disease: NOVELTY design. *ERJ Open Research*, 5(1):00036–2018, 2019.

[234] H. Marshall, J.M. Wild, L.J. Smith, L. Hardaker, T. Fihn-Wikander, H. Müllerová, and R. Hughes. Functional imaging in asthma and COPD: design of the NOVELTY ADPro substudy. *ERJ Open Research*, pages 00344–2022, 2023.

[235] K. Ng, L.C. Saunders, G. Collier, J. Grist, S. Strickland, L. Gustafsson, L.J. Smith, P.J.C. Hughes, N. Kainth, R.I. Evans, A. Laws, S.R. Thomas, J.L. Rodgers, S.A. Jones, T.L. Newman, M.L. Plowright, L. Dryhurst-Pearce, K. Jacob, A. McIntyre, D. Capener, J. Bray, M. Durrant, K. Yeung, H. Walters, V. Matthews, L. Watson, A. Elbehairy, G. Vuddamalay, G.A. Eid, M. Viktorie, M. Cox, V. Harris, M. Lachut, N. Mulvey, W. Hickes, N. Rahman, A. Horsley, H.E. Davies, R.A.A. Thompson, J.M. Wild, F.V. Gleeson, and E. Fraser. The EXPLAIN (hyperpolarised xenon magnetic resonance pulmonary imaging in patients with Long-COVID) study: The primary outcome. *Am J Respir Crit Care Med*, 209:A6478, 2024.

[236] B.L. Graham, V. Brusasco, F. Burgos, B.G. Cooper, R. Jensen, A. Kendrick, N.R. MacIntyre, B.R. Thompson, and J. Wanger. 2017 ERS/ATS standards for single-breath carbon monoxide uptake in the lung. *European Respiratory Journal*, 49(1):1600016, 2017.

[237] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[238] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[239] F.C. Horn, M.H. Deppe, H. Marshall, J. Parra-Robles, and J.M. Wild. Quantification of regional fractional ventilation in human subjects by measurement of hyperpolarized 3He washout with 2D and 3D MRI. *J Appl Physiol*, 116(2):129–139, 2014. PMID: 24311749.

[240] J.B. Hittner, K. May, and N.C. Silver. A monte carlo evaluation of tests for comparing dependent correlations. *J Gen Psychol*, 130(2):149–168, 2003.

[241] G.Y. Zhou. Toward using confidence intervals to compare correlations. *Psychol Methods*, 12(4):399–413, 2007.

[242] B. Diedenhofen and J. Musch. cocor: A comprehensive solution for the statistical comparison of correlations. *PLoS ONE*, 10(4):e0121945, 2015.

[243] R.H. Sansores, R.T. Abboud, C. Kennell, and N. Haynes. The effect of menstruation on the pulmonary carbon monoxide diffusing capacity. *Am J Respir Crit Care Med*, 152(1):381–4, 1995.

[244] J.M.B. Hughes, A.f.R. Technology, and Physiology. *Physiology and Practice of Pulmonary Function*. Association for Respiratory Technology and Physiology, 2009. (Great Britain).

[245] W.J. Garrison, K. Qing, M. He, L. Zhao, N.J. Tustison, J.T. Patrie, J.F. Mata, Y.M. Shim, A.M. Ropp, T.A. Altes, 3rd J.P. Mugler, and G.W. Miller. Lung volume dependence and repeatability of hyperpolarized (129)Xe MRI gas uptake metrics in healthy volunteers and participants with COPD. *Radiology Cardiothoracic Imaging*, 5(3):e220096, 2023.

[246] A. Bechtel, E. Bier, J. Lu, A. Church, J. Korzekwinski, D. Mummy, and B. Driehuys. Establishing a hemoglobin correction for 129Xe gas exchange MRI. In *Proc. 31st Intl. Soc. Mag. Reson. Med.*, page 1184, London, UK, 2022.

[247] G.J. Collier, P.J.C. Hughes, F.C. Horn, H.F. Chan, B. Tahir, G. Norquay, N.J. Stewart, and J.M. Wild. Single breath-held acquisition of coregistered 3D (129) Xe lung ventilation and anatomical proton images of the human lung with compressed sensing. *Magn Reson Med*, 82(1):342–347, 2019.

[248] S. Perron, D.G. McCormack, G. Parraga, and A. Ouriadov. Undersampled diffusion-weighted 129Xe MRI morphometry of airspace enlargement: Feasibility in chronic obstructive pulmonary disease. *Diagnostics*, 13(8):1477, 2023.

[249] Q. Zhou, H. Li, Q. Rao, M. Zhang, X. Zhao, L. Shen, Y. Fang, H. Li, X. Liu, S. Xiao, L. Shi, Y. Han, C. Ye, and X. Zhou. Assessment of pulmonary morphometry using hyperpolarized 129Xe diffusion-weighted MRI with variable-sampling-ratio compressed sensing patterns. *Medical Physics*, 50(2):867–878, 2023.

[250] J.W. Plummer, R. Hussain, A.S. Bdaiwi, S.A. Soderlund, X. Hoyos, J.M. Lanier, W.J. Garrison, J. Parra-Robles, M.M. Willmering, P.J. Niedbalski, Z.I. Cleveland, and L.L. Walkup. A decay-modeled compressed sensing reconstruction approach for non-cartesian hyperpolarized Xe MRI. *Magnetic Resonance in Medicine*, 92(4):1363–1375, 2024.

[251] M. Blumenthal, C. Holme, V. Roeloffs, S. Rosenzweig, P. Schaten, N. Scholand, J. Tamir, X. Wang, and M. Uecker. mrirecon/bart: version 0.8.00. https://doi.org/10.5281/zenodo.7110562, 2022.

[252] J.L. Olesen, A. Ianus, L. Østergaard, N. Shemesh, and S.N. Jespersen. Tensor denoising of multidimensional MRI data. *Magn Reson Med*, 89(3):1160–1172, 2023.

[253] S.A. Soderlund, A.S. Bdaiwi, J.W. Plummer, J.C. Woods, L.L. Walkup, and Z.I. Cleveland. Improved diffusion-weighted hyperpolarized 129Xe lung MRI with patch-based higher-order, singular value decomposition denoising. *Academic Radiology*, 31:5289–5299, 2024. doi: 10.1016/j.acra.2024.06.029.

[254] N.J. Stewart, J. de Arcos, A.M. Biancardi, G.J. Collier, L.J. Smith, G. Norquay, H. Marshall, A.C.S. Brau, R.M. Lebel, and J.M. Wild. Improving xenon-129 lung ventilation image snr with deep-learning based image reconstruction. *Magnetic Resonance in Medicine*, 92(6):2546–2559, 2024.

[255] B. Kang, W. Lee, H. Seo, H.-Y. Heo, and H. Park. Self-supervised learning for denoising of multidimensional mri data. *Magnetic Resonance in Medicine*, 92(5):1980–1994, 2024.

[256] C. Lazarus, P. Weiss, A. Vignaud, and P. Ciuciu. An empirical study of the maximum degree of undersampling in compressed sensing for T(2)(*)-weighted MRI. *Magnetic Resonance Imaging*, 53:112–122, 2018.

[257] J.P. Mugler III and J.R. Brookeman. Signal-to-noise considerations for parallel imaging with hyperpolarized gases. In *Proc. 13th Intl. Soc. Mag. Reson. Med. 108*, volume 13, page 108, 2005.

[258] S. Ajraoui, J. Parra-Robles, and J.M. Wild. Incorporation of prior knowledge in compressed sensing for faster acquisition of hyperpolarized gas images. *Magn Reson Med*, 69(2):360–369, 2013.

[259] A. Bechtel, J. Lu, D. Mummy, E. Bier, S. Leewiwatwong, J. Mugler III, S. Kabir, A. Church, and B. Driehuys. Establishing a hemoglobin adjustment for Xe gas exchange MRI and MRS. *Magn Reson Med*, 90(4):1555–1568, 2023.

[260] N.D. Weatherley, J.A. Eaden, P.J.C. Hughes, M. Austin, L. Smith, J. Bray, H. Marshall, S. Renshaw, S.M. Bianchi, and J.M. Wild. Quantification of pulmonary perfusion in idiopathic pulmonary fibrosis with first pass dynamic contrast-enhanced perfusion MRI. *Thorax*, 76(2):144–151, 2021.

[261] K. Hueper, J. Vogel-Claussen, M.A. Parikh, J.H. Austin, D.A. Bluemke, J. Carr, J. Choi, T.A. Goldstein, A.S. Gomes, E.A. Hoffman, S.M. Kawut, J. Lima, E.D. Michos, W.S. Post, M.J. Po, M.R. Prince, K. Liu, D. Rabinowitz, J. Skrok, B.M. Smith, K. Watson, Y. Yin, A.M. Zambeli-Ljepovic, and R.G. Barr. Pulmonary microvascular blood flow in mild chronic obstructive pulmonary disease and emphysema. the MESA COPD study. *Am J Respir Crit Care Med*, 192(5):570–80, 2015.

[262] I. Vlahos, M.C. Jacobsen, M.C. Godoy, K. Stefanidis, and R.R. Layman. Dual-energy CT in pulmonary vascular disease. *The British Journal of Radiology*, 95(1129):20210699, 2022.

[263] C.S. Johns, D.G. Kiely, and A.J. Swift. Novel imaging techniques in pulmonary hypertension. *Current Opinion in Cardiology*, 33(6):587–593, 2018.

[264] C. Ozer Gokaslan, S. Gokaslan, E. Demirel, and S. Sengul Ayan. Use of phase-contrast mri to measure aortic stiffness in young-onset hypertension: a pilot study. *Cardiology in the Young*, 33(2):266270, 2023.

[265] H.B. Grotenhuis, J.J.M. Westenberg, P. Steendijk, R.J. van der Geest, J. Ottenkamp, J.J. Bax, J.W. Jukema, and A. de Roos. Validation and reproducibility of aortic pulse wave velocity as assessed with velocity-encoded MRI. *J Magn Reson Imaging*, 30(3):521–526, 2009.

[266] G.H. Pöhler, F. Löffler, F. Klime, L. Behrendt, A. Voskrebenzev, C.C. González, M. Westhoff-Bleck, F. Wacker, and J. Vogel-Claussen. Validation of phase-resolved functional lung (PREFUL) magnetic resonance imaging pulse wave transit time compared to echocardiography in chronic obstructive pulmonary disease. *J Magn Reson Imaging*, 56(2):605–615, 2022.

[267] G.H. Pöhler, F. Klimes, A. Voskrebenzev, L. Behrendt, C. Czerner, M. Gutberlet, S. Cebotari, F. Ius, C. Fegbeutel, C. Schoenfeld, T.F. Kaireit, E.F. Hauck, K.M. Olsson, M.M. Hoeper, F. Wacker, and J. Vogel-Claussen. Chronic thromboembolic pulmonary hypertension perioperative monitoring using phase-resolved functional lung (PREFUL)-MRI. *J Magn Reson Imaging*, 52(2):610–619, 2020.

[268] N.R. Zwart, K.O. Johnson, and J.G. Pipe. Efficient sample density estimation by combining gridding and an optimized kernel. *Magn Reson Med*, 67(3):701–710, 2012.

[269] A. Rose. *Vision: Human and Electronic*. Plenum Press, 1973.

[270] N.B. Karatzas and J. Lee Gde. Propagation of blood flow pulse in the normal human pulmonary arterial system. analysis of the pulsatile capillary flow. *Circulation Research*, 25(1):11–21, 1969.

[271] W.A. Littler, R. Bojorges-Bueno, and J. Banks. Cardiovascular dynamics in women during the menstrual cycle and oral contraceptive therapy. *Thorax*, 29(5):567–70, 1974.

[272] A.C. Henderson, R.C. Sá, R.J. Theilmann, R.B. Buxton, G.K. Prisk, and S.R. Hopkins. The gravitational distribution of ventilation-perfusion ratio is more uniform in prone than supine posture in the normal human lung. *J Appl Physiol*, 115(3):313–324, 2013.

[273] K. Wasserman, J. Butler, and A.V. Kessel. Factors affecting the pulmonary capillary blood flow pulse in man. *J Appl Physiol*, 21(3):890–900, 1966.

[274] M.L. Costa, J. Lu, M.M. Willmering, A.S. Bdaiwi, J.W. Plummer, L. Walkup, B. Driehuys, and Z.I. Cleveland. Sex and age-dependence of 129Xe MRI cardiopulmonary oscillations in healthy volunteers. In *B80-1. METHODOLOGICAL ADVANCEMENTS IN PULMONARY IMAGING*, pages A4506–A4506. American Thoracic Society, 2024.

[275] W.G. Murphy. The sex difference in haemoglobin levels in adults mechanisms, causes, and consequences. *Blood Reviews*, 28(2):41–47, 2014.

[276] A. Costelle, J. Lu, A. Bechtel, E. Bier, S. Kabir, J. Mammarrappallil, H. Dai, D. Mummy, and B. Driehuys. Quantifying cardiogenic oscillations of hyperpolarized xe gas exchange mr spectra in a healthy reference cohort. In *Proc. 31st Intl. Soc. Mag. Reson. Med. 0857*, volume 31, page 0857, 2023.

[277] J. Lu, S. Leewiwatwong, K. Du, S. Zhang, Y-C. Huang, D. Mummy, and B. Driehuys. Repeatability of cardiopulmonary oscillations measured with imaged with 129Xe MRI. In *Proc. 32nd Intl. Soc. Mag. Reson. Med. 3192*, volume 32, page 3192, 2024.

[278] D.S. Feuer, E.M. Handberg, B. Mehrad, J. Wei, C.N. Bairey Merz, C.J. Pepine, and E.C. Keeley. Microvascular dysfunction as a systemic disease: A review of the evidence. *The American Journal of Medicine*, 135(9):1059–1068, 2022.

[279] N.D. Weatherley, J.A. Eaden, P.J.C. Hughes, M. Austin, L. Smith, J. Bray, H. Marshall, S. Renshaw, S.M. Bianchi, and J.M. Wild. Quantification of pulmonary perfusion in idiopathic pulmonary fibrosis with first pass dynamic contrast-enhanced perfusion MRI. *Thorax*, 76(2):144–151, 2021.

[280] J.D. McFadyen, H. Stevens, and K. Peter. The emerging threat of (micro)thrombosis in COVID-19 and its therapeutic implications. *Circulation Research*, 127(4):571–587, 2020.

[281] M. Ackermann, S.E. Verleden, M. Kuehnel, A. Haverich, T. Welte, F. Laenger, A. Vanstapel, C. Werlein, H. Stark, A. Tzankov, W.W. Li, V.W. Li, S.J. Mentzer, and D. Jonigk. Pulmonary vascular endothelialitis, thrombosis, and angiogenesis in Covid-19. *New England Journal of Medicine*, 383(2):120–128, 2020.

[282] J.M. Wild, F.V. Gleeson, S. Svenningsen, J.T. Grist, L.C. Saunders, G.J. Collier, M. Sharma, S. Tcherner, A. Mozaffaripour, A.M. Matheson, and G. Parraga. Review of hyperpolarized pulmonary functional Xe MR for long-COVID. *J Magn Reson Imaging*, 59(4):1120–1134, 2024.

[283] E.W. Ely, L.M. Brown, and H.V. Fineberg. Long covid defined. *New England Journal of Medicine*, 391(18):1746–1753, 2024.

[284] I.Y. Zhou, M. Mascia, G.A. Alba, M. Magaletta, L.C. Ginns, P. Caravan, and S.B. Montesi. Dynamic contrast-enhanced mri demonstrates pulmonary microvascular abnormalities months after SARS-CoV-2 infection. *Am J Respir Crit Care Med*, 207(12):1636–1639, 2023.

[285] L. Leber, A. Beaudet, and A. Muller. Epidemiology of pulmonary arterial hypertension and chronic thromboembolic pulmonary hypertension: identification of the most accurate estimates from a systematic literature review. *Pulmonary Circulation*, 11(1):2045894020977300, 2021.

[286] B.R. Sabbula, A. Sankari, and J. Akella. Chronic thromboembolic pulmonary hypertension. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing. Available from: https://www.ncbi.nlm.nih.gov/books/NBK549836/, March 2024.

[287] D. Jenkins. Pulmonary endarterectomy: the potentially curative treatment for patients with chronic thromboembolic pulmonary hypertension. *European Respiratory Review*, 24(136):263–271, 2015.

[288] J. Pepke-Zaba, H.-A. Ghofrani, and M.M. Hoeper. Medical management of chronic thromboembolic pulmonary hypertension. *European Respiratory Review*, 26(143):160107, 2017.

[289] J.E. Cannon, L.Su, D.G. Kiely, K. Page, M. Toshner, E. Swietlik, C. Treacy, A. Ponnaberanam, R. Condliffe, K. Sheares, D. Taboada, J. Dunning, S. Tsui, C. Ng, D. Gopalan, N. Screaton, C. Elliot, S. Gibbs, L. Howard, P. Corris, J. Lordan, M. Johnson, A. Peacock, R. Mackenzie-Ross, B. Schreiber, G. Coghlan, K. Dimopoulos, S.J. Wort, S. Gaine, S. Moledina, D.P. Jenkins, and J. Pepke-Zaba. Dynamic risk stratification of patient long-term outcome after pulmonary endarterectomy. *Circulation*, 133(18):1761–1771, 2016.

[290] F. Dardi, A. Manes, D. Guarino, S.M. Suarez, A. Loforte, M. Rotunno, D. Pacini, N. Gali, and M. Palazzini. Long-term outcomes after pulmonary endarterectomy. *Ann Cardiothorac Surg*, 11(2):172–174, 2022.

[291] Y. Gordon, S. Partovi, M. Müller-Eschner, E. Amarteifio, T. Bäuerle, M.-A. Weber, H.-U. Kauczor, and F. Rengier. Dynamic contrast-enhanced magnetic resonance imaging: fundamentals and application to the evaluation of the peripheral perfusion. *Cardiovascular Diagnosis and Therapy*, 4(2), 2014.

[292] D.L. Levin, Q. Chen, M. Zhang, R.R. Edelman, and H. Hatabu. Evaluation of regional pulmonary perfusion using ultrafast magnetic resonance imaging. *Magn Reson Med*, 46(1):166–171, 2001.

[293] M. Ingrisch and S. Sourbron. Tracer-kinetic modeling of dynamic contrast-enhanced MRI and CT: a primer. *J Pharmacokinet Pharmacodyn*, 40(3):281–300, 2013.

[294] B.S. Ebrahimi, M.H. Tawhai, H. Kumar, K.S. Burrowes, E.A. Hoffman, M.L. Wilsher, D. Milne, and A.R. Clark. A computational model of contributors to pulmonary hypertensive disease: impacts of whole lung and focal disease distributions. *Pulm Circ*, 11(4):20458940211056527, 2021.

[295] B.S. Ebrahimi. *A structure-based multi-scale model evaluation of hemodynamics and the impact of remodelling in pulmonary hypertension*. Phd thesis, University of Auckland, Auckland, New Zealand, May 2022.

[296] B.S. Ebrahimi, H. Kumar, M.H. Tawhai, K.S. Burrowes, E.A. Hoffman, and A.R. Clark. Simulating multi-scale pulmonary vascular function by coupling

computational fluid dynamics with an anatomic network model. *Frontiers in Network Physiology*, 2, 2022.

[297] K.S. Burrowes, P.J. Hunter, and M.H. Tawhai. Anatomically based finite element models of the human pulmonary arterial and venous trees including supernumerary vessels. *J Appl Physiol*, 99(2):731–738, 2005.

[298] A.R. Clark, M.H. Tawhai, E.A. Hoffman, and K.S. Burrowes. The interdependent contributions of gravitational and structural features to perfusion distribution in a multiscale model of the pulmonary circulation. *J Appl Physiol*, 110(4):943–955, 2011.

[299] A.R. Clark and M.H. Tawhai. Temporal and spatial heterogeneity in pulmonary perfusion: A mathematical model to predict interactions between macro- and micro-vessels in health and disease. *The ANZIAM Journal*, 59(4):562580, 2018.

[300] B. Haefeli-Bleuer and E.R. Weibel. Morphometry of the human pulmonary acinus. *The Anatomical Record*, 220(4):401–414, 1988.

[301] Y.C. Fung and S.S. Sobin. Theory of sheet flow in lung alveoli. *J Appl Physiol*, 26(4):472–488, 1969. PMID: 5775333.

[302] Y.C. Fung and S.S. Sobin. Elasticity of the pulmonary alveolar sheet. *Circulation Research*, 30(4):451–469, 1972.

[303] Y.C. Fung. Theoretical pulmonary microvascular impedance. *Ann Biomed Eng*, 1:221245, 1972.

[304] B. Duan and M. Zamir. Viscous damping in onedimensional wave transmission. *The Journal of the Acoustical Society of America*, 92(6):3358–3363, 12 1992.

[305] D. Heath and J.E. Edwards. The pathology of hypertensive pulmonary vascular disease. *Circulation*, 18(4):533–547, 1958.

[306] N.J. Tustison, P.A. Cook, A.J. Holbrook, H.J. Johnson, J. Muschelli, G.A. Devenyi, J.T. Duda, S.R. Das, N.C. Cullen, D.L. Gillen, M.A. Yassa, J.R. Stone, J.C. Gee, and B.B. Avants. The ANTsX ecosystem for quantitative biological and medical imaging. *Scientific Reports*, 11(1):9068, 2021.

[307] A. Klein, J. Andersson, B.A. Ardekani, J. Ashburner, B. Avants, M.C. Chiang, G.E. Christensen, D.L. Collins, J. Gee, P. Hellier, J.H. Song, M. Jenkinson, C. Lepage, D. Rueckert, P. Thompson, T. Vercauteren, R.P. Woods, J.J. Mann, and R.V. Parsey. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage*, 46(3):786–802, 2009.

[308] M. Schiwek, S.M.F. Triphan, J. Biederer, O. Weinheimer, M. Eichinger, C.F. Vogelmeier, R.A. Jörres, H.U. Kauczor, C.P. HeuSSel, P. Konietzke, O. von Stackelberg, F. Risse, B.J. Jobst, and M.O. Wielpütz; COSYCONET study

group. Quantification of pulmonary perfusion abnormalities using DCE-MRI in COPD: comparison with quantitative CT and pulmonary function. *Eur Radiol*, 32(3):1879–1890, 2022.

[309] Pulmonary Hypertension Association UK. EmPHasis-10 questionnaire. https://www.phauk.org/pha-uk-resources/emphasis-10-questionnaire/, 2013.

[310] K.L. Ng, L. Saunders, G. Collier, J. Grist, I. Dunstan, R. Evans, M. Lachut, J. Ablott, S. Strickland, L. Gustafsson, L. Smith, S. Thomas, J. Rodgers, G. Vuddamalay, N. Kainth, A. Laws, E. Hedley, S. Jones, P. Hughes, T. Newman, M. Plowright, L. Dryhurst-Pearce, A. Elbehairy, K. Jacob, A. McIntyre, D. Capener, J. Bray, M. Durrant, K. Yeung, H. Walters, L. Watson, B. Johnson, O. Rodgers, R. Munro, V. Madhusudhan Stisova, M. Cox, D. Jakymelen, V. Harris, V. Matthews, G. Abu-Eid, N. Mulvey, W. Hickes, D. Parramon, N. Rahman, A. Horsley, H. Davies, J. Wild, F. Gleeson, E. Fraser, A.A.R. Thompson, and S. Shapiro. S82 persistently raised serum amyloid A in never-hospitalised long-COVID patients without association with lung or coagulation abnormalities: the EXPLAIN study (hyperpolarised xenon magnetic resonance pulmonary imaging in patients with long-COVID). *Thorax*, 79(Suppl 2):A59–A60, 2024.

[311] H. Li, X. Zhao, Y. Wang, X. Lou, S. Chen, H. Deng, L. Shi, J. Xie, D. Tang, J. Zhao, L.-S. Bouchard, L. Xia, and X. Zhou. Damaged lung gas exchange function of discharged COVID-19 patients detected by hyperpolarized 129Xe MRI. *Science Advances*, 7(1):eabc8180, 2021.

[312] C. Gerges, M. Gerges, R. Friewald, P. Fesler, P. Dorfmüller, S. Sharma, K. Karlocai, N. Skoro-Sajer, J. Jakowitsch, B. Moser, S. Taghavi, W. Klepetko, and I.M. Lang. Microvascular disease in chronic thromboembolic pulmonary hypertension. *Circulation*, 141(5):376–386, 2020.

[313] T. Heimann, M. Eichinger, G. Bauman, A. Bischoff, M. Puderbach, and H.-P. Meinzer. Automated scoring of regional lung perfusion in children from contrast enhanced 3D MRI. *Progress in Biomedical Optics and Imaging - Proceedings of SPIE*, 8315:28–, 02 2012.

[314] K. Qing, J.P. Mugler III, T.A. Altes, Y. Jiang, J.F. Mata, G.W. Miller, I.C. Ruset, F.W. Hersman, and K. Ruppert. Assessment of lung function in asthma and copd using hyperpolarized 129Xe chemical shift saturation recovery spectroscopy and dissolved-phase MRI. *NMR in Biomedicine*, 27(12):1490–1501, 2014.

[315] M.B. Bolster and R.M. Silver. Lung disease in systemic sclerosis (scleroderma). *Baillière's Clinical Rheumatology*, 7(1):79–97, 1993. Lung Disease in Rheumatic Disorders.

[316] M. Martín-López and P.E. Carreira. The impact of progressive pulmonary fibrosis in systemic sclerosis-associated interstitial lung disease. *J Clin Med*, 12(20), 2023.

[317] A.-M. Hoffmann-Vold, H. Fretheim, A.-K. Halse, M. Seip, H. Bitter, M. Wallenius, T. Garen, A. Salberg, C. Brunborg, O. Midtvedt, M.B. Lund, T.M. Aaløkken, and O. Molberg. Tracking impact of interstitial lung disease in systemic sclerosis in a complete nationwide cohort. *Am J Respir Crit Care Med*, 200(10):1258–1266, 2019.

[318] U.A. Walker, A. Tyndall, L. Czirják, C. Denton, D. Farge-Bancel, O. Kowal-Bielecka, U. Müller-Ladner, C. Bocelli-Tyndall, and M. Matucci-Cerinic. Clinical risk assessment of organ manifestations in systemic sclerosis: a report from the EULAR scleroderma trials and research group database. *Annals of the Rheumatic Diseases*, 66(6):754–763, 2007.

[319] M. Naranjo and P.M. Hassoun. Systemic sclerosis-associated pulmonary hypertension: Spectrum and impact. *Diagnostics (Basel)*, 11(5), 2021.

[320] M. Hinchcliff, A. Fischer, E. Schiopu, and V.D. Steen. Pulmonary hypertension assessment and recognition of outcomes in scleroderma (PHAROS): Baseline characteristics and description of study population. *The Journal of Rheumatology*, 38(10):2172–2179, 2011.

[321] A. Luz. The bootstrap or why you should care about uncertainty. https://datalesdatales.medium.com/the-bootstrap-or-why-you-should-care-about-uncertainty-c3fe1a5f36b5, 2018.

[322] J.P. Mugler 3rd, T.A. Altes, I.C. Ruset, I.M. Dregely, J.F. Mata, G.W. Miller, S. Ketel, J. Ketel, F.W. Hersman, and K. Ruppert. Simultaneous magnetic resonance imaging of ventilation distribution and gas uptake in the human lung using hyperpolarized xenon-129. *Proc Natl Acad Sci U S A*, 107(50):21707–21712, 2010.

[323] R.F. Schulte. *MNS Research Pack fidall PSD Manual*, 2021.

[324] S.S. Kaushik, M.S. Freeman, Z.I. Cleveland, J. Davies, J. Stiles, R.S. Virgincar, S.H. Robertson, M. He, K.T. Kelly, W.M. Foster, H.P. McAdams, and B. Driehuys. Probing the regional distribution of pulmonary gas exchange through single-breath gas- and dissolved-phase 129Xe MR imaging. *J Appl Physiol*, 115(6):850–860, 2013.

[325] K. Ruppert, L. Loza, F. Amzajerdian, H. Hamedani, R. Baron, S. Kadlecek, and R. Rizi. Quantitative measure of lung structure and function obtained from hyperpolarized xenon spectroscopy. *J Vis Exp.*, 10(101), 2023.

[326] G. Norquay, G.J. Collier, and J.M. Wild. Regional quantification of 129Xe gaussian broadening in the lungs with time-domain voigt fitting of 3D CSI. In *Proc. 32nd Intl. Soc. Mag. Reson. Med. 4756*, volume 32, page 4756, 2024.

[327] Y. Taylor, F.J. Wilson, M. Kim, and G.J.M. Parker. Sensitivity analysis of models of gas exchange for lung hyperpolarised 129Xe MR. *NMR in Biomedicine*, 37(12):e5239, 2024.

[328] D.M. Vasilescu, A.B. Phillion, D. Kinose, S.E. Verleden, B.M. Vanaudenaerde, G.M. Verleden, D. Van Raemdonck, C.S. Stevenson, C.J. Hague, M.K. Han, J.D. Cooper, T.L. Hackett, and J.C. Hogg. Comprehensive stereological assessment of the human lung using multiresolution computed tomography. *J Appl Physiol*, 128(6):1604–1616, 2020.

[329] H.O. Coxson, R.M. Rogers, K.P. Whittall, Y. D'Yachkova, P.D. Pare, F.C. Sciurba, and J.C. Hogg. A quantification of the lung surface area in emphysema using computed tomography. *Am J Respir Crit Care Med*, 159(3):851–856, 1999.

[330] Y.V. Chang, J.D. Quirk, I.C. Ruset, J.J. Atkinson, F.W. Hersman, and J.C. Woods. Quantification of human lung structure and physiology using hyperpolarized 129Xe. *Magn Reson Med*, 71(1):339–344, 2014.

[331] J. Wolber, S.J. Doran, M.O. Leach, and A. Bifone. Measuring diffusion of xenon in solution with hyperpolarized 129Xe NMR. *Chemical Physics Letters*, 296(3):391–396, 1998.

[332] F. Junker and W.S. Veeman. Xenon self-diffusion in organic polymers by pulsed field gradient NMR spectroscopy. *Macromolecules*, 31(20):7010–7013, 1998.

[333] M. Filkins, A. Harrison, G.J. Collier, G. Norquay, J.M. Wild, S.P. Rigby, G.E. Pavlovskaya, and T. Meersmann. A standardized MRI phantom for dissolved phase 129 Xe MRI. *JOURNAL OF MAGNETIC RESONANCE OPEN*, 21, DEC 2024.

[334] M.M. Willmering, B.J. Albert, J.W. Plummer, J. Greer, L.L. Walkup, D.M. Lindquist, and Z.I. Cleveland. A thermally polarized, dissolved-phase 129Xe phantom for quality-control and multisite comparisons of gas-exchange imaging. *J Magn Reson*, 371:107829, 2025.

[335] A.J. Hajari, D.A. Yablonskiy, A.L. Sukstanskii, J.D. Quirk, M.S. Conradi, and J.C. Woods. Morphometric changes in the human pulmonary acinus during inflation. *J Appl Physiol*, 112(6):937–943, 2012.

[336] D. Khanna, D.P. Tashkin, C.P. Denton, E.A. Renzoni, S.R. Desai, and J. Varga. Etiology, risk factors, and biomarkers in systemic sclerosis with interstitial lung disease. *Am J Respir Crit Care Med*, 201(6), 2020.

[337] A. Huertas, L. Tu, M. Humbert, and C. Guignabert. Chronic inflammation within the vascular wall in pulmonary arterial hypertension: more than a spectator. *Cardiovascular Research*, 116(5):885–893, 11 2019.

[338] N.J. Stewart, J. Parra-Robles, and J.M. Wild. Finite element modeling of 129Xe diffusive gas exchange NMR in the human alveoli. *J Magn Reson*, 271:21–33, 2016.

[339] Z. Wang, A. Swaminathan, E. Bier, R.M. Tighe, Y.T. Huang, S. Rajagopal, and B. Driehuys. Using 129Xenon MR gas exchange MRI to measure the membrane and capillary components of DLCO and KCO. *Am J Respir Crit Care Med*, 201, 2020.

[340] C. Puddu, M. Rao, X. Xu, M.H. Deppe, G.J. Collier, A. Maunder, H.-F. Chan, N. De Zanche, F. Robb, and J.M. Wild. An asymmetrical whole-body birdcage RF coil without RF shield for hyperpolarized 129Xe lung MR imaging at 1.5 T. *Magn Reson Med*, 86(6):3373–3381, 2021.

[341] MAGNIFY SRC: Pulmonary magnetic resonance imaging for cystic fibrosis. Online: https://www.cysticfibrosis.org.uk/the-work-we-do/research/cf-research-topics/understanding-and-treating-symptoms/magnify-src-pulmonary-magnetic-resonance Accessed 17/03/25.

[342] MIT. Finite difference methods. Online: http://web.mit.edu/-course/16/16.90/BackUp/www/pdfs/Chapter13.pdf.

[343] B. Mauroy. Following red blood cells in a pulmonary capillary. *ESAIM: Proc.*, 23:48–65, 2008.

[344] D. Rubenstein, W. Yin, and M.D. Frame. *Biofluid Mechanics: An Introduction to Fluid Mechanics, Macrocirculation, and Microcirculation*. Academic Press, 2011.

Appendix A

2D Numerical Gas Exchange Simulation

^{129}Xe gas exchange in the alveoli, including the transport of ^{129}Xe magnetisation out of the gas exchange region via blood flow, was simulated using a 2D finite difference method (FDM) model. The FDM approximates the spatial derivatives of a differential equation by dividing the space and time domains into discrete intervals in a mesh and using the difference between the values at neighbouring nodes [342].

The model geometry is shown in Figure A.0.1A, where 'upstream' denotes blood flowing into the gas exchange region with no dissolved ^{129}Xe and 'downstream' contains blood with dissolved ^{129}Xe outside of the gas exchange region. The width of the slab (h) was 10 μm and the length of the gas exchange region was 0.5 mm, which was estimated from the capillary transit time, $\tau \sim 1$ s and blood flow velocity, which was assumed to be a linear (plug) flow of v $\sim$ 0.5 mm s^{-1} [343, 344]. The upstream and downstream capillary lengths were 0.1 mm, chosen so that the simulated downstream magnetisation would not reach the end of the model, avoiding end effects.

To minimise the computational power required, the 2D geometry was divided into 500 horizontal lines so that the finite difference equation only had to be solved once (in 1D). The 1D diffusion equation (Equation 3.4.1) was solved in Matlab using the FDM approximation:

$$\frac{\partial M_i}{\partial t} - D\delta_x^2 M_i = 0, \tag{A.0.1}$$

where M_i is the magnetisation at each node, δ_x is the forwards difference operator and D is the ^{129}Xe diffusion coefficient, taken to be 3.3 x10^{-6} cm^{2}s^{-1}.

The solution, represented in Figure A.0.1B, was found for an arbitrary time period chosen to be longer than the total time of the 2D model. The horizontal lines were assumed to be independent of one another, such that if there was no blood flow, each one would be equivalent to the 1D solution. The total magnetisation was given by a Y x X x J matrix, where Y was the total number of horizontal lines

(500), X was the number of x points (100) and J was the number of time points (20,000). Figure A.0.1C demonstrates how an individual line in the 2D geometry is merely the solution of the 1D simulation.

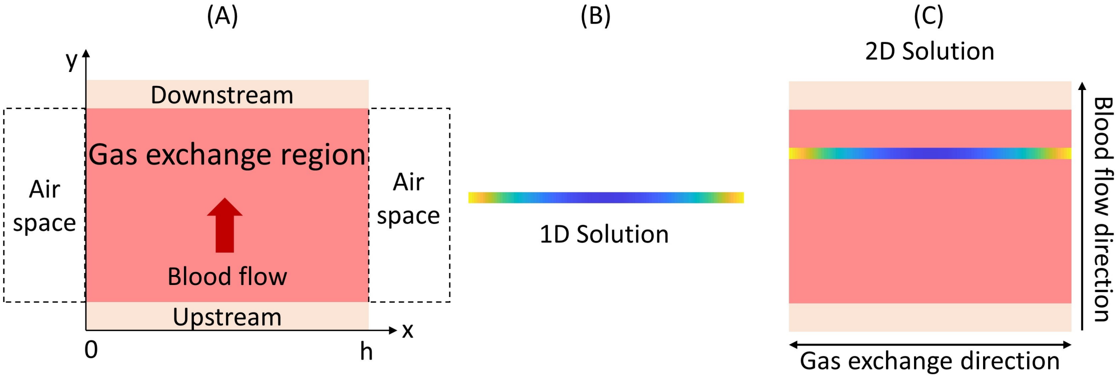

Figure A.0.1: (a) The geometry used for the 2D FDM model. ^{129}Xe enters from the airspace into the gas exchange region, where diffusion takes place. Blood flows from upstream through the gas exchange region and onwards downstream, taking ^{129}Xe out of this region. To simulate this, the geometry was divided into horizontal lines. The diffusion equation was solved for a single line (b) and then the 2D solution (c) is made up of copies of this solution at different times.

To include blood flow, each line was shifted forwards according to the capillary transit time. The index of these 'shift times' was found from $\frac{\Delta y}{v\Delta t}$, where Δy and Δt are the y and time steps, which were 1.4 µm and 10 µs respectively. At each shift time the magnetisation density of a horizontal line was set to the value of the 1D solution at the next time such that there was an overall flow of velocity v. This meant that the flow of fresh blood into the gas exchange region and the magnetisation leaving the gas exchange region downstream were captured in the simulation.

Appendix B

3D Numerical Gas Exchange Simulation

B.1 Introduction

In this thesis, CSSR uptake data were fit to the Patz model. However, this is an analytical model which uses a simple 1D model of an alveolar septum with a uniform thickness. Gas exchange modelling may be improved by using numerical methods such as the finite element method (FEM) to simulate more realistic geometries. FEM analysis involves splitting the domain where the problem is defined into small 2D elements in a mesh and solving the problem in each of these elements to reach an approximate solution, as illustrated in Figure B.1.1.

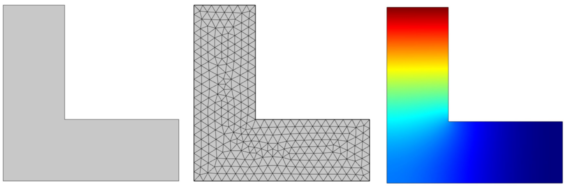

Figure B.1.1: The FEM involves dividing a geometry into small discrete parts, then approximating the unknown function for each element. This creates a system of equations which are recombined to form a global system of equations which give the final solution.

FEM modelling of ^{129}Xe gas exchange in the lungs has been demonstrated for several different geometries by Stewart et al. [338]. The first of their models used a cylinder to represent a pulmonary capillary, surrounded by a tissue shell. The signal from dissolved-phase ^{129}Xe transported out of the gas exchange region by the blood was taken into account by splitting the cylinder in half and removing the tissue compartment for one of the halves, as shown in Figure B.1.2.

Two image-based models were also presented in Ref. [338], which used 2D histology images and 3D micro-CT images from ex vivo lung tissue biopsies respectively. The 2D histology model was segmented and extruded to make a pseudo-3D model and the same principle as the cylinder used to deal with the blood flow. The micro-CT model used 3D volumetric images, which offers the most realistic geometry, but a significantly lower spatial resolution than the histology-based

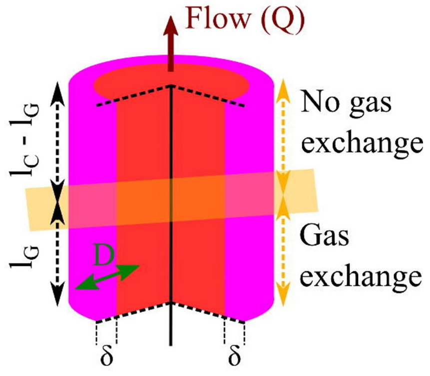

Figure B.1.2: The cylindrical geometry from Ref. [338], where l_G and l_C are the lengths of the gas exchange region and total capillary length, respectively. D and Q represent diffusion and perfusion. Copyright 2016 Elsevier, reproduced with permission from [338].

model. A limitation of this model is that the images did not include capillaries from outside of the gas exchange region, so the effects of blood flow could not easily be included.

A differential diffusion equation was solved over these geometries with Dirichlet boundary conditions imposed at the air-tissue barrier, representing a constant inflow of gaseous ^{129}Xe to the dissolved compartments. This is equivalent to assuming an infinite reservoir of gas-phase ^{129}Xe, which is justified because only a small (∼2%) proportion of the gaseous ^{129}Xe diffuses into the membrane and capillary blood. However, by simulating the gas-phase magnetisation in addition to the dissolved-phase magnetisation, more accurate solutions may be reached. The entire journey of ^{129}Xe from inhalation, diffusion and uptake could be simulated, including the Brownian motion of the gas-phase ^{129}Xe.

In this section, we describe the progress made so far towards the development of a combined gas-phase/dissolved-phase FEM gas exchange model.

B.2 Methods

FEM modelling was performed using COMSOL Multiphysics (5.6, Burlington, MA), which was accessed via visualisation nodes on the University of Sheffield ShARC high performance computer cluster. To represent both the gaseous and dissolved-phase ^{129}Xe magnetisation, the alveolus was modelled as a sphere of radius R_α = 0.15 mm surrounded by a thin membrane layer of width h = 10 μm (Figure B.2.1). In the preliminary simulation, this membrane was not further separated into tissue and capillary compartments and so no blood flow was included. In the following simulation, the membrane was divided into a tissue layer of width 1 μm and a capillary layer of width 10 μm. To account for the flow of blood containing dissolved ^{129}Xe out of the gas exchange region and into the bloodstream, an outflow tube of radius 35 μm was included to simulate a pulmonary venule. The length of the venule was chosen to be greater than the product of the blood flow velocity and the total simulation time, such that the magnetisation in the tube continued to contribute to the total magnetisation for the entire simulation. To reduce computation time, the axial symmetry of the spherical geometry was exploited and simulations were run for a 2D slice of the sphere which was then rotated around the $r = 0$ axis to obtain the 3D solution.

The geometries were meshed using triangular elements. The minimum element size was 0.006 μm in the tissue and capillary layers and the maximum element size was 0.1 μm in the tissue layer and 1 μm in the capillary layer. These values were chosen to ensure that the mesh consisted of at least ten elements in the radial direction for each layer. For the alveolar airspace, predefined mesh size settings were used (minimum element size = 0.226 μm, maximum element size = 67 μm).

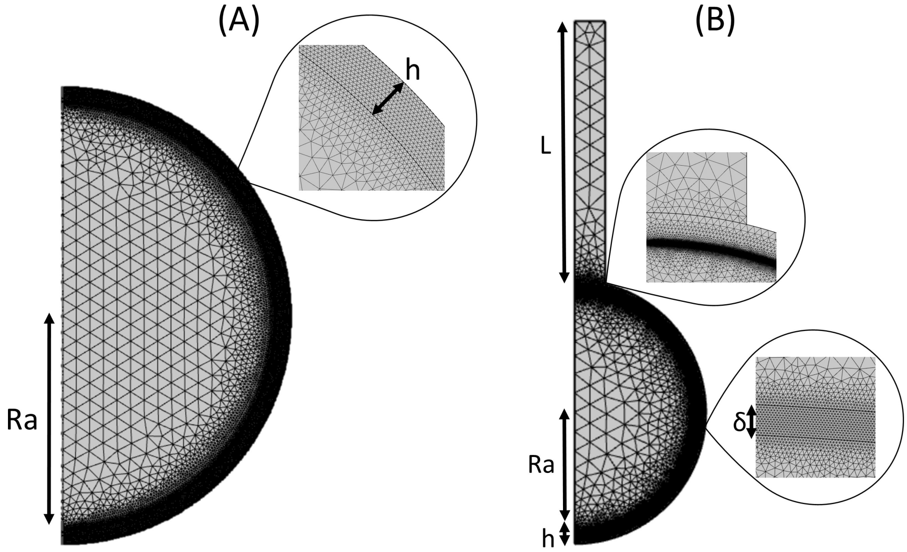

Figure B.2.1: Geometries of the CSSR FEM models: (A) a spherical alveolus was considered, which was surrounded by a membrane layer of width h. (B) To include blood flow effects in the alveolus model, the membrane layer was divided into a tissue layer of width δ and a capillary layer of width $h - \delta$. The contribution from ^{129}Xe magnetisation outside of the gas exchange region was incorporated by the addition of an outflow tube, representing a small blood vessel.

Gas exchange was simulated for an ideal CSSR experiment, assuming perfect dissolved-phase saturation such that the dissolved-phase magnetisation is initially zero. Magnetisation evolution was modelled using the partial differential equation for diffusion:

$$\frac{\partial \mathbf{M}}{\partial t} = \nabla \cdot (D \nabla \mathbf{M}) \tag{B.2.1}$$

This equation was solved within the gas-phase compartment and dissolved-phase compartment. The large difference in diffusion coefficients between the air space and the dissolved-phase and the low Ostwald solubility lead to a discontinuity at the tissue-gas boundary. This amounts to a ‘stiff-spring’ boundary condition [338], given by

$$(-D\nabla M_2)\cdot\hat{n} = v_{ss}(K_{21}M_1 - M_2), \tag{B.2.2}$$

where D is the dissolved ^{129}Xe diffusion coefficient, taken to be 3.3 x 10^{-6} m^2 s^{-1}, M_i is the magnetisation in region i, $\hat{n}$ is the unit vector perpendicular to the boundary, v_{ss} is the stiff-spring velocity and K_{21} is the ratio of Ostwald solubilities between the two regions. v_{ss} is a non-physical velocity which must be large enough that a continuous ^{129}Xe flux can cross the boundary and that $(K_{21}M_1 - M_2)$ can approach zero [338]. The value used throughout this work was 1000 ms^{-1}. Zero-flux boundary conditions were enforced at all other boundaries. The initial conditions used for the simulation were $\mathbf{M}$ = 1 mol m^{-3} within the alveolar airspace and $\mathbf{M}$ = 0 in the dissolved-phase compartment.

To model blood flow, the dissolved-phase was separated into tissue and capillary compartments and in the capillary, the convection-diffusion equation (Equation B.2.3) was used to include a blood flow term. There is a linear increase after ~0.1 s in the ^{129}Xe uptake curve from blood containing dissolved ^{129}Xe flowing out of the gas exchange region.

$$\frac{\partial \mathbf{M}}{\partial t} = \nabla\cdot(D\nabla\mathbf{M}) - \nabla\cdot(v\mathbf{M}) \tag{B.2.3}$$

The time-dependent magnetisation with and without blood flow was found by integrating the magnetisation density in each compartment over its volume.

B.3 Results

Figure B.3.1 shows the diffusion of magnetisation from the alveolar airspace and into the membrane at different time points, without blood flow. With a membrane thickness of h = 10 µm, the dissolved-phase magnetisation saturated at ~0.4 s, which is later than expected by the Patz model. With the inclusion of blood flow in the model, the transport of ^{129}Xe magnetisation into the pulmonary capillary and out of the gas exchange region can be observed (Figure B.3.2B).

There was good agreement between the FEM and Patz model when the Patz model septal width was equal to twice the numerical model value. This is due to geometry differences; in the single alveolus model the dissolved-phase is sourced with gas on one side, whereas in the Patz model, the dissolved-phase is supplied with gas from both sides of the septal slab. Figure B.3.3A shows the dissolved-phase magnetisation over time for the FEM sphere model with no blood flow, along with the Patz model result. Fitting the simulated dissolved-phase magnetisation from the FEM model with blood flow to the Patz model resulted in good agreement and a fitted septal width of 9.6 µm (Figure B.3.3B).

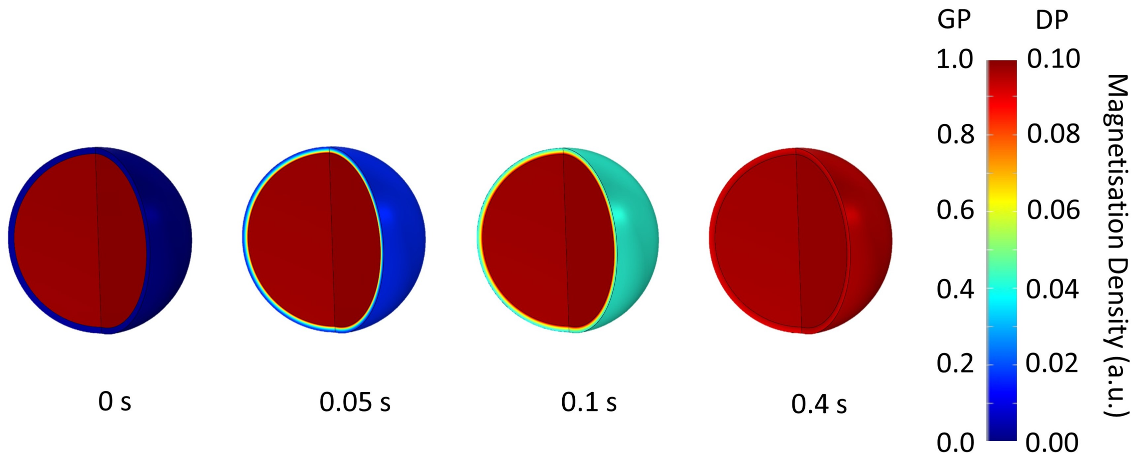

Figure B.3.1: The diffusion of ^{129}Xe from the alveolar airspace into the dissolved-phase in the alveolar sphere model, with no blood flow, shown at four time points. GP = gas-phase, DP = dissolved-phase.

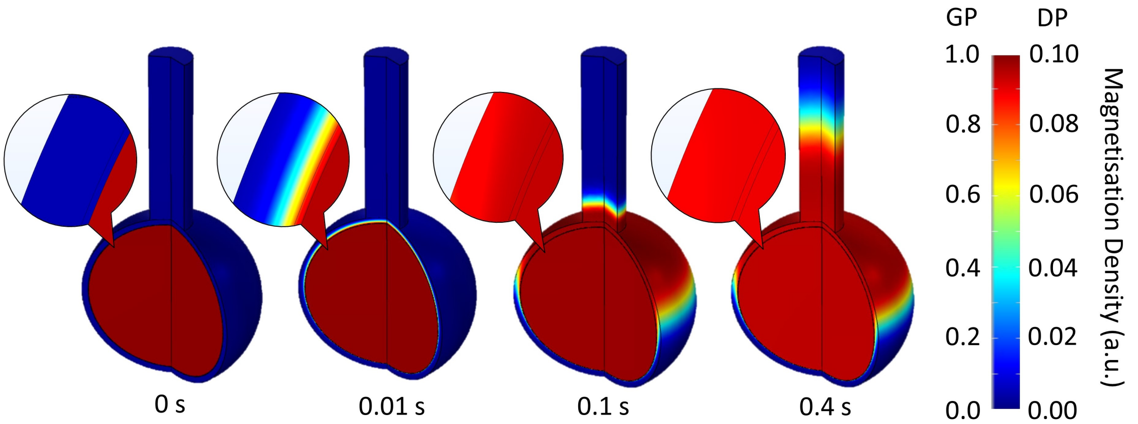

Figure B.3.2: The diffusion of ^{129}Xe from the alveolar airspace into the dissolved-phase in the alveolar sphere model, with blood flow, shown at four time points. This model separated the dissolved-phase compartment into tissue and capillary, as shown by the zoomed in images.

B.4 Discussion

We have started to develop a numerical gas exchange model which captures both the gaseous and dissolved-phase ^{129}Xe magnetisation, something that has not been done previously. Instead of Dirichlet boundary conditions, a discontinuous boundary was applied at the gas-dissolved interface, governed by the Ostwald solubility and the differences in diffusion coefficients between the two media. Comparison of our modelling results with the Patz model showed that the dissolved-phase thickness in the alveoli model was equivalent to double this thickness in the Patz model, due to geometry differences.

Although in its early stages, this model will eventually allow the journey of ^{129}Xe to be temporally mapped from the alveolar airspace into the lung tissue and capillaries. Future work will involve using the numerical model to simulate both the

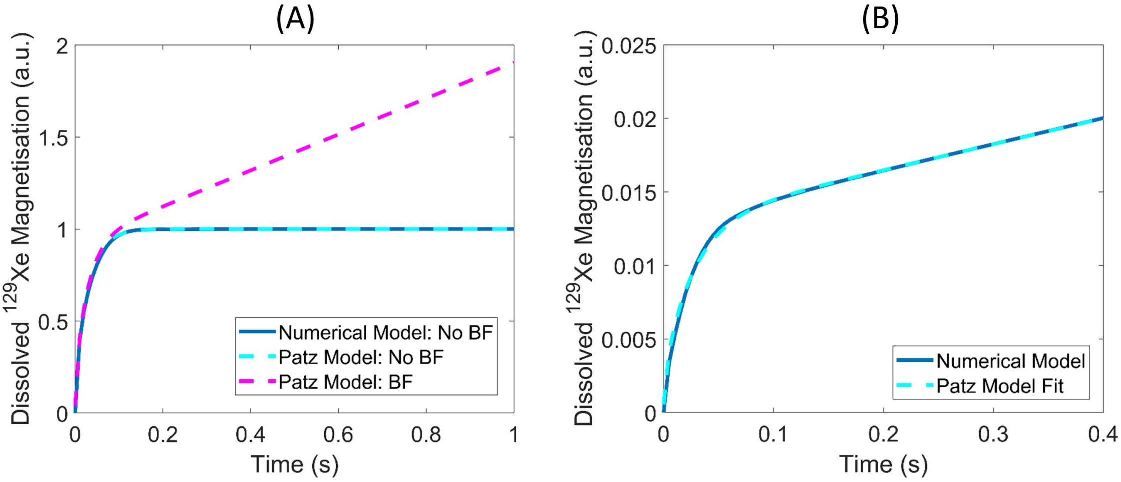

Figure B.3.3: (A) The dissolved-phase magnetisation simulated with the sphere model for no blood flow (h = 5 μm) and the Patz model results for blood flow and no blood flow (h = 10 μm). (B) The dissolved-phase magnetisation simulated with the sphere model with blood flow (h = 5 μm) and the Patz model fit to the simulated data.

Brownian diffusion of the gaseous ^{129}Xe and the diffusion across the tissue barrier into the bloodstream using the Bloch-Torrey equation. This can then be compared to experimental results from healthy subjects and to results from analytical modelling. For diseased lungs, the model may be able to offer insight to the underlying disease mechanisms by changing parameters such as the tissue thickness, alveolar radius and perfusion.

There are some limitations to our FEM model. Blood flow was modelled as a linear flow, but pulmonary capillary blood flow is known to be pulsatile. In healthy people, the ^{129}Xe RBC signal oscillates with an amplitude of ∼15% (see Chapter 7) and this is increased/decreased in lungs with post/pre-capillary disease. A pulsatile blood flow could be implemented into the FEM model by using an input waveform which is periodic and oscillates at the heart beat frequency. Using a 2D axially symmetric geometry reduced computation time in comparison to a fully 3D model, but required the assumption that 2D diffusion can be scaled up to 3D diffusion. Since diffusion is a random process, there may be some differences between diffusion in the 2D slice and diffusion in the 3D sphere. However, any differences would occur on very small time scales and the overall behaviour is likely to be the same. Although the model presented in this section uses a more realistic geometry than the 1D analytical models, it is still a simplified representation which may limit its usefulness in providing information about pulmonary disease. More realistic geometries which incorporate a heterogeneous tissue thickness can be modelled using micro-CT lung images, as was done by Stewart et al. [338].

B.5 Conclusions

Initial steps have been made towards a full 3D alveolar airspace model, which could allow simultaneous prediction of the ^{129}Xe diffusion coefficients and microstructural parameters from CSSR spectroscopy.