

Benchtop to Bedside: Developments in Pulmonary MRI

Thesis submitted to the University of Nottingham for the degree of **Doctor of Philosophy**, **July 2025**.

Arthur Harrison 14252464

Supervised by

Thomas Meersmann Galina E. Pavlovskaya

School of Medicine

In memory of Derek Rimmer 1929 - 2013

My first inspiration for a life in science

Abstract

This work focuses on two distinct methods for pulmonary MRI: the development of contrast-based hyperpolarised noble gases production techniques and the implementation and application of free-breathing ¹H-based functional imaging at 0.5 T. An overview of the thesis and a list of published works are provided in Chapter 1, followed by four novel research chapters.

Established Spin Exchange Optical Pumping (SEOP) polarisers typically specialise in polarising a single noble gas isotope. Chapter 2 follows the construction and testing of a versatile preclinical polariser, able to produce large volumes of hyperpolarised ¹²⁹Xe, ⁸³Kr and ¹³¹Xe. The SEOP process necessitates the dilution of the hyperpolarised species and subsequent purification is required to maximise SNR; however, hyperpolarised krypton cannot be purified by established cryogenic-based methods. Therefore, Chapter 3 follows the engineering development of a novel purification apparatus for hyperpolarised ¹²⁹Xe and ⁸³Kr, using molecular hydrogen as a buffer gas, which is eliminated through controlled combustion with oxygen.

The established Phase Resolved Functional Lung (PREFUL) methodology is typically employed with 1.5 T MRI, limiting its deployment with scanners with an open architecture and in developing countries where low field systems are typical. In Chapter 4, the PREFUL algorithm is adapted in a custom-built pipeline for application with low field (0.5 T) MRI, accounting for a lower SNR and imaging rate. Conventional MRI scanners limit participants to lying horizontally, therefore these investigations do not account for changes to the lung due to gravity. In Chapter 5, the PREFUL pipeline is employed with Open MRI to probe changes in the distribution of lung ventilation and perfusion in different postures. This is the first MRI study to investigate the lungs in the inverted position.

Acknowledgements

This PhD has been an epic journey with many obstacles, from pandemics to invasions, that no-one could have anticipated. Even though it has taken slightly longer than expected, I have thoroughly enjoyed the process and I would encourage anyone with a love for their subject to embark on this journey for themselves. However, I could not have completed this work without the unwavering support from my family: Jacqui, Mark and Niamh. You are the pillars of my life; thank you for always being there when I needed you.

Thank you also to my supervisors, Thomas and Galina, for their guidance and the many interesting discussions we have had along the way. Additionally, thank you to my colleagues in the Sir Peter Mansfield Imaging Centre for their support and entertainment during many long days (and nights) in the office and the lab.

Finally, I would like to thank the University of Nottingham Mountaineering Club and the many friends I have made along the way for keeping me sane throughout these years with regular trips to the Lake District, Eryri, Scotland and beyond.

Contents

Abstra	act		ii
Ackno	owledgements		iii
List of	f Tables		viii
List of	f Figures		ix
Abbre	eviations		xvii
Chapte	er 1 This Thesis		1
1.1	Thesis Outline		. 1
1.2	Publications & Presentations		. 2
Chapte	ser 2 Development of a Prototype Pre-Clinica	al SEC)P
	Polariser		5
2.1	Introduction		. 9
2.2	Polarisation Theory		. 9
2.3	Polariser Development		. 28
2.4	Results		. 48
2.5	Discussion		. 51
2.6	Conclusion		. 52
2.7	References		. 53
Chapte	ter 3 Reactive Purification of Post-SEOP H	yper-	
	polarised Noble Gases using a Hydrogen	ı Buffe	er
	Gas		63
3.1	Introduction		. 67
3.2	Hyperpolarised Gas Purification		. 70
3.3	Hydrogen as a Buffer Gas		. 72

3.4	Aims
3.5	The Combustion system
3.6	Results
3.7	Discussion
3.8	Conclusion
3.9	References
Chapte	er 4 Development of PREFUL Post-Processing for
	Application at 0.5 T
4.1	Introduction
4.2	Aims
4.3	0.5 T Scanners
4.4	PREFUL Pipeline
4.5	Phase Assignment
4.6	Registration
4.7	Filtering
4.8	Reconstruction
4.9	Segmentation
4.10	Analysis
4.11	User Interface Implementation
4.12	Results
4.13	Discussion
4.14	Conclusions
4.15	References
Chapte	er 5 Gravitational Dependency of the Distribu-
	tion of Normal Lung Function 292
5.1	Introduction
5.2	Methodology and Materials
5.3	Results
5 4	Discussion

	5.5	Conclusions	473
	5.6	References	478
Ap	pen	dices 4	196
Ap	pen	dix A Neural Network Training 4	197
	A.1	Hyperparameters	498
Ap	pen	dix B Computational Hardware 4	199
	B.1	PREFUL Pipeline Application	500
	B.2	Neural Network Training	500
	В.3	Combustion System Operation	500

List of Tables

2.1	Required B_0 of each Resonant Circuit
2.2	Polariser Laser Modules
3.1	Controlled Valves of the Combustion System 91
4.1	T2* of Healthy Lung Parenchyma
4.2	Training Data Summary
4.3	Training Data Augmentation Limits
4.4	Accuracy of Trained Networks
4.5	Acquisition Parameters for the 0.5 T Upright Scanner 247
4.6	PREFUL Processing Parameters for the $0.5~\mathrm{T}$ Upright Scanner248
4.7	Acquisition Parameters for the 0.5 T Bruker Scanner \dots 253
4.8	PREFUL Processing Parameters for the 0.5 T Bruker Scanner254
4.9	Mean Values for FV and RVent in this study 265
4.10	Literature Values for Healthy RVent at Different Field Strengths 266
5.1	MRI Sequences of Protocol 1
5.2	PREFUL Processing Parameters
5.3	Resultant angle of Coronal Slices
5.4	B1 Impact to Functional Gradients
5.5	Coronal PREFUL Processing Summary
5.6	Sagittal PREFUL Processing Summary
5.7	Coronal Exhaled Signal Apical-to-Basal Variation 368
5.8	Sagittal Exhaled Signal Apical-to-Basal & Anterior-to-Posterior
	Variation 371

5.9	Sagittal Anterior and Posterior Exhaled Signal Apical-to-
	Basal Variation
5.10	Coronal Fractional Ventilation Apical-to-Basal Variation 378
5.11	Sagittal Fractional Ventilation Apical-to-Basal & Anterior-
	to-Posterior Variation
5.12	Sagittal Anterior and Posterior Fractional Ventilation Apical-
	to-Basal Variation
5.13	Coronal Normalised Perfusion Apical-to-Basal Variation 402
5.14	Sagittal Normalised Perfusion Apical-to-Basal & Anterior-
	to-Posterior Variation
5.15	Sagittal Anterior and Posterior Normalised Perfusion Apical-
	to-Basal Variation
5.16	Relative contribution to Tissue Density and Fractional ven-
	tilation
A 1	Training Hyperparameters

List of Figures

2.1	Boltzmann Polarisation at Low Temperatures	10
2.2	Optical Pumping of Rubidium D1 transition	17
2.3	Spin Transfer Mechanisms	21
2.4	The Polariser	30
2.5	Polariser Thermal Control System	31
2.6	Thermal Control Electronics	32
2.7	Polariser Magnetic Field Setup	33
2.8	Low Field NMR Apparatus	35
2.9	Polariser Optical Setup	38
2.10	Pneumatic manifolds of the Polariser Manifold	41
2.11	Polariser Flow Manifold Adaptations	44
2.12	Polariser Control Interface	46
2.13	Polariser Pre-Saturator Control Interface	47
2.14	Xenon-129 Results	49
2.15	Krypton-83 Results	50
2.16	Xenon-131 Results	51
3.1	Apparent Polarisation	70
3.2	Feasibility of using Hydrogen as a Buffer Gas for SEOP	73
3.3	Hydrogen-Oxygen-Nitrogen Flammability Diagram	75
3.4	Sequential Images Ignition Hydrogen in the Glass-Walled	
	Pt-Coated Vessel	77

3.5	Preliminary Investigations of the Hydrogen Combustion Re-
	sistance of Hyperpolarised Xenon-129 and Krypton-83 $$
3.6	Combustion System Hardware
3.7	Ignition Sources of the Combustion System
3.8	Combustion System Housing
3.9	Electronic Diagram of the Combustion System
3.10	Coil Spacing Optimisation
3.11	Field Magnitude Simulation
3.12	Field Characteristics Simulations
3.13	Pneumatic Diagram of the Combustion System 90
3.14	Hydrogen Delivery Cylinder
3.15	Piston-Based Oxygen Filling Station
3.16	Automated Oxygen Filling Station
3.17	Combustion System Control GUI
3.18	Post-Visualisation of Pressure Change GUI 98
3.19	Combustion System Safety Hardware
3.20	Single-Spark Automation Sequence Flowchart 106
3.21	Warm-up sequence Flow Chart
3.22	Batch-Mode Automation Sequence Flowchart
3.23	Batch-Mode Automation Sequence Flowchart, Continued 111
3.24	Re-compression Automation Sequence Flowchart
3.25	Pressure Curves of Batch-Mode Sequence
3.26	Hyperpolarised Results for Batch-Mode Sequence 116
3.27	Pressure Curves of Semi-Continuous-Mode Sequence 117
4.1	0.5 T Scanners used in this work
4.2	Overview of the PREFUL Pipeline
4.3	Diaphragm Position by Edge Detection
4.4	Comparison of Ventilation Markers

4.5	Frame Assignment to Inhalation or Exhalation 159
4.6	Respiratory Thresholding
4.7	Auto Respiratory Thresholding
4.8	Frequency Domain analysis to assist in Cardiac Phase As-
	signment
4.9	Demonstration of Cardiac Aliasing
4.10	Effects of Fitting Aliased Data
4.11	Cardiac Aliasing Quality Assurance
4.12	Feature-Based Registration based on Diaphragm Position 183
4.13	Typical Manual Segmentation Pipeline
4.14	Artificial Neural Networks
4.15	Computer Vision Classes
4.16	UNet Architecture
4.17	Model Training Progress
4.18	VGG-SegNet based Coronal 2D GRE Network Accuracy $$ 211
4.19	UNet based Coronal 2D GRE Network Accuracy 212
4.21	VGG-SegNet based Axial 2D GRE Network Accuracy 215
4.22	VGG-SegNet based Sagittal 2D GRE Network Accuracy $$ 216
4.23	$\operatorname{VGG-SegNet}$ based Coronal 2D HASTE Network Accuracy . 217
4.24	Ventilation Processing Pass vs Fail Assignment
4.25	Local Ventilation Analysis Methods
4.26	Fractional Ventilation Determination Methods
4.27	Ventilation Time-To-Peak Determination Methods 228
4.28	Jacobian Determinant Analysis Methods 230
4.29	Local Perfusion Analysis Methods
4.30	Perfusion Weighting Determination Methods
4.31	Perfusion Time-To-Peak Determination Methods
4.32	Process of Vessel Segmentation
4.33	Read Image-Series GUI

4.34	Ventilation Processing GUI
4.35	Ventilation Processing Quality Control GUI 241
4.36	Perfusion Processing GUI
4.37	Perfusion Processing Quality Control GUI
4.38	Display Processed Image-Series GUI
4.39	PREFUL Results from Coronal 2D GRE scans acquired us-
	ing the Upright Scanner
4.40	PREFUL Results from Axial 2D GRE scans acquired using
	the Upright Scanner
4.41	PREFUL Results from Sagittal 2D GRE scans acquired us-
	ing the Upright Scanner
4.42	PREFUL Results from Coronal 2D HASTE scans acquired
	using the Upright Scanner
4.43	PREFUL Results from Coronal 2D GRE scans acquired us-
	ing the Bruker Scanner #1
4.44	PREFUL Results from Coronal 2D GRE scans acquired us-
	ing the Bruker Scanner #2
4.45	PREFUL Results from Coronal 2D GRE scans acquired us-
	ing the Bruker Scanner #3
5.1	The Slinky Effect
5.2	Gravitational Plural Pressure
5.3	West's Zones
5.4	Ax et al. SPECT Investigation of Head-Down perfusion 316
5.5	Hopkins et al. ASL Investigation of Supine Perfusion 317
5.6	The 0.5T Open MRI Scanner (ASG Paramed, Italy) 322
5.7	Slice Positioning
5.8	Participant Postures
5.9	Breathing Regimes

5.10	Timeline of Protocol 2
5.11	Example Coronal Figure
5.12	Example Sagittal Figure
5.13	Lung Mask Processing
5.14	Periphery Segmentation Pipeline
5.15	Resultant angle of Coronal Slices
5.16	Demonstration of Signal Variation within a Uniform Spin-
	Density Phantom
5.17	B1-Transmit Variation in Phantom
5.18	B1-Receive Variation in Phantom
5.19	Sagittal Axes Interdependence
5.20	Posture Completion Summary
5.21	Example Coronal Exhaled Lung Reconstructions
5.22	Coronal Exhaled Signal Apical-to-Basal Distribution 367
5.23	Example Sagittal Exhaled Lung Reconstructions
5.24	Sagittal Exhaled Signal Apical-to-Basal & Anterior-to-Posterior
	Distribution
5.25	Sagittal Anterior and Posterior Exhaled Signal Apical-to-
	Basal Distribution
5.26	Example Coronal Fractional Ventilation Maps 374
5.27	Coronal Normal-Breathing Fractional Ventilation Apical-to-
	Basal Distribution
5.28	Coronal Deep-Breathing Fractional Ventilation Apical-to-Basal
	Distribution
5.29	Coronal Normal vs Deep Breathing Fractional Ventilation
	Apical-to-Basal Distribution
5.30	Example Sagittal Fractional Ventilation Maps
5.31	Sagittal Fractional Ventilation Apical-to-Basal & Anterior-
	to-Posterior Distribution

5.32	Sagittal Anterior and Posterior Fractional Ventilation Apical-
	to-Basal Distribution
5.33	Example Coronal Ventilation Time-To-Peak Maps 387
5.34	Coronal Ventilation Time-to-Peak Apical-to-Basal Distribu-
	tion
5.35	Coronal Normal vs Deep Breathing Ventilation Time-to-Peak
	Apical-to-Basal Distribution
5.36	Example Sagittal Ventilation Time-To-Peak Maps 389
5.37	Sagittal Ventilation Time-To-Peak Apical-to-Basal & Anterior-
	to-Posterior Distribution
5.38	Example Coronal Maps of the Jacobian Determinant 392
5.39	Coronal Apical-to-Basal Distribution of the Jacobian Deter-
	minant
5.40	Coronal Peripheral Apical-to-Basal Distribution of the Ja-
	cobian Determinant
5.41	Coronal Normal vs Deep Breathing Jacobian Determinant
	Apical-to-Basal Distribution
5.42	Example Sagittal Maps of the Jacobian Determinant $$ 396
5.43	Sagittal Apical-to-Basal & Anterior-to-Posterior Distribu-
	tion of the Jacobian Determinant
5.44	Example Coronal Perfusion Weighted Maps
5.45	Coronal Normalised Perfusion Apical-to-Basal Distribution . 400
5.46	Coronal Peripheral Normalised Perfusion Apical-to-Basal Dis-
	tribution
5.47	Coronal Normal vs Deep Breathing Normalised Perfusion
	Apical-to-Basal Distribution
5.48	Example Sagittal Perfusion Weighted Maps 404
5.49	Sagittal Normalised Perfusion Apical-to-Basal & Anterior-
	to-Posterior Distribution

5.50	Sagittal Anterior and Posterior Normalised Perfusion Apical-
	to-Basal Distribution
5.51	Example Coronal Perfusion Time-to-Peak Maps 411
5.52	Coronal Perfusion Time-to-Peak Apical-to-Basal Distribution 412
5.53	Coronal Normal vs Deep Breathing Perfusion Time-to-Peak
	Apical-to-Basal Distribution
5.54	Example Sagittal Perfusion Time-to-Peak Maps 413
5.55	Sagittal Perfusion Time-to-Peak Apical-to-Basal & Anterior-
	to-Posterior Distribution
5.56	The Weighted Slinky Effect
5.57	Prone Parenchyma Distribution
5.58	Model for Basal Fractional Ventilation Bias
5.59	Contributions to Perfusion in the Sitting Posture 456

Abbreviations

4DCT Four-Dimensional Computed Tomography.

ASL Arterial Spin Labelling.

bSSFP Balanced Steady-State Free Precession.

CNN Convolutional Neural Network.

COPD Chronic Obstructive Pulmonary Disease.

CT Computed Tomography.

DAQ Data Acquisition Board.

DCE Dynamic Contrast-Enhanced.

DECT Dual-Energy Computed Tomography.

DNP Dynamic Nuclear Polarisation.

ECG Electrocardiogram.

FD Fourier Decomposition.

FV Fractional Ventilation.

GRE Gradient Echo.

HASTE Half-Fourier Acquisition Single-shot Turbo spin-Echo.

HPV Hypoxic Pulmonary Vasoconstriction.

IoU Intersection over Union.

LFL Lower Flammability Limit.

MHD Magnetohydrodynamic.

MPD Matrix Pencil Decomposition.

MR Magnetic Resonance.

MRI Magnetic Resonance Imaging.

NMR Nuclear Magnetic Resonance.

nuFD Non-uniform Fourier-Decomposition.

OE Oxygen Enhanced.

PET Positron Emission Tomography.

PHIP Parahydrogen Induced Polarisation.

PREFUL Phase Resolved Functional Lung.

QTTP Perfusion Time-To-Peak.

SABRE Signal Amplification by Reversible Exchange.

SEOP Spin Exchange Optical Pumping.

SNR Signal to Noise Ratio.

SPECT Single Photon Emission Computed Tomography.

SPGRE Spoiled Gradient Echo.

UFL Upper Flammability Limit.

UTE Ultrashort Echo Time.

VDP Ventilation Defect Percentage.

VTTP Ventilation Time-To-Peak.

Chapter 1

This Thesis

1.1 Thesis Outline

Chapter 2 follows the development of a prototype pre-clinical polariser. this system is designed to be highly adaptable, capable of operating under the batch-mode and continuous-flow regimes, for the hyperpolarisation of xenon-129 and krypton-83.

Chapter 3 builds on the work of Chapter 2, to develop a novel apparatus for the reactive purification of post SEOP hyperpolarised noble gases, when hydrogen is employed as the sole buffer gas

Chapter 4 deviates from the previous chapters, exploring an alternative method for pulmonary MRI. This work seeks to adapt the Phase Resolved Functional Lung (PREFUL) post-processing method for application with 0.5 T MRI scanners.

Chapter 5 uses the home-built pipeline developed in Chapter 5, utilising its compatibility with open scanner architectures for a novel experiment

investigating the gravitational influences of regional lung function.

1.2 Publications & Presentations

1.2.1 First-Author Publications

Arthur Harrison, Thomas Meersmann, Galina E Pavlovskaya, Olga S. Pavlova, Alexander Makurenkov, Nikolay V Anisimov, Yury A. Pirogov Proof-of-principal of 1H PREFUL approaches at 0.5T: lung ventilation and perfusion maps of a healthy volunteer and a Covid-19 survivor, ISMRM Poster (2021)

Arthur Harrison, Max Filkins, Xinpei Wang, Stefano M. Collins, Sean P. Rigby, Chengbo Wang, Galina E. Pavlovskaya, Thomas Meersmann, A Device for the Oxidative Purification of Hyperpolarised Noble Gases after Spin Exchange Optical Pumping, EUROMAR Poster (2022)

Arthur Harrison, Max Filkins, Xinpei Wang, Stefano M. Collins, Sean P. Rigby, Chengbo Wang, Galina E. Pavlovskaya, Thomas Meersmann, Large Volume Purification of Hyperpolarised Noble Gases via Multi-Stage Combustion after Spin Exchange Optical Pumping, EUROMAR Poster/Flash Talk (2023)

Arthur Harrison, Thomas Meersmann, Galina E Pavlovskaya, Penny Gowland, Jan A Paul, Rashed Sobhan, Amanda Goodwin, Olivier Mougin, Investigating Gravitational Influence on Normal Lung Function Using PREFUL MRI on an Open Scanner, ISMRM Oral Presentation (2024)

Arthur Harrison, Thomas Meersmann, Galina E Pavlovskaya, Penny

Gowland, Jan A Paul, Rashed Sobhan, Amanda Goodwin, Olivier Mougin, Investigating Gravitational Influence on Normal Lung Function Using PREFUL MRI on an Open Scanner, BIC-ISMRM Oral Presentation (2024)

Arthur Harrison, Thomas Meersmann, Galina E Pavlovskaya, Penny Gowland, Jan A Paul, Rashed Sobhan, Olivier Mougin, Quantitative Assessment of Pulmonary Function Distribution in the Sitting, Supine, and Prone Postures Using PREFUL MRI with an Open Scanner, ISMRM Oral Presentation (2025)

1.2.2 Co-Author Publications

Max Filkins, **Arthur Harrison**, Guilhem J. Collier, Graham Norquay, Jim M. Wild, Sean P. Rigby, Galina E. Pavlovskaya, Thomas Meersmann, A standardized MRI phantom for dissolved phase 129Xe MRI, JMR (2024)

Stefano Marc Collins, Sean Rigby, Huw Williams, Galina Pavlovskaya, Thomas Meersmann, **Arthur Harrison**, Christophe Vallée, Thibaud Chevalier, Dina Lofficial, *Predicting the pore-structure transport relationship in disordered alumina using hyperpolarized xenon MRI and NMR cryodif-fusometry*, **CPM-9 International Workshop Poster Presentation** (2024)

Max Filkins, **Arthur Harrison**, Guilhem J. Collier, Graham Norquay, Jim M. Wild, Chengbo Wang, Sean P. Rigby, Galina E. Pavlovskaya, Thomas Meersmann, A standardized MRI phantom for dissolved phase 129Xe MRI,

BIC-ISMRM Poster Presentation (2024)

Z. Peggs, O. Mougin, A. Harrison, S.H. Needleman, N. Blockley, M Kim, G. Pavlovskaya, T. Meersmann, S. Francis, GJM. Parker, P. Gowland, R. Sobhan, Investigating the Effect of Gravity and Oxygen Signal Enhancement on Lung T2* with Upright 0.5T MR, ISMRM Poster Presentation (2025)

Max Filkins, **Arthur Harrison**, Guilhem J. Collier, Graham Norquay, Jim M. Wild, Chengbo Wang, Sean P. Rigby, Galina E. Pavlovskaya, Thomas Meersmann, A standardized MRI phantom for dissolved phase 129Xe MRI, **ISMRM Poster Presentation (2025)**

Laura Saunders, Guilhem Collier, Graham Norquay, Laurie Smith, Paul Hughes, Scarlett Strikland, Lotta Gustafsson, Thomas Newman, Megan Plowright, James Watson, Zoë Gabriel, Phillip Wade, James Meiring, James T Grist, Kher Lik Ng, **Arthur Harrison**, James Eaden, Jody Bray, Helen Marshall, David Capener, Martin Brook, Alberto Biancardi, James E Ball, Neil Stewart, Kevin M Johnson, Andy Swift, Smitha Rajaram, Lisa Watson, Paul Collini, Gary Mills, Rod Lawson, Jonathan Brooke, Phillip Molyneaux, Amanda Goodwin, Iain D Stewart, Ling-Pei Ho, Joseph Jacob, Thomas Meersman, Galina Pavlovskaya, Fergus Gleeson, Ian Hall, R Gisli Jenkins, A. A. Roger Thompson, and Jim Wild, 129Xe red blood cell chemical shift and T2* in patients hospitalised due to COVID-19 with and without residual lung abnormalities seen on CT, ISMRM Poster Presntation (2025)

Chapter 2

Development of a Prototype Pre-Clinical SEOP Polariser

Abstract

This chapter describes the design, construction, and testing of a highly adaptable, large-volume pre-clinical polariser for spin exchange optical pumping (SEOP) of noble gas isotopes. The polariser was developed to address the need for a versatile system capable of hyperpolarising multiple noble gas species, including the quadrupolar nucleus krypton-83, for which no large-quantity polariser (>0.5 L) previously existed in the literature.

The system incorporates five main components: a dual-zone forced-air heating system with PID temperature control, a double Helmholtz coil configuration, interchangeable laser mounting systems accommodating 60-180 W diode array lasers, adaptable pneumatic manifolds supporting both batch and continuous-flow operation modes, and integrated electronics for valve control. Furthermore, the system can be adapted into configurations for rubidium pre-saturation and cryogenic purification.

The polariser successfully demonstrated hyperpolarisation of xenon-129, krypton-83, and xenon-131 under various operating conditions. Polarisation measurements were conducted using a 9.4 T NMR spectrometer with thermal reference calibration. The modular design enables rapid reconfiguration for different experimental requirements, gas mixtures, and operating pressures, making it suitable for diverse pre-clinical research applications.

While the polariser successfully produced hyperpolarised noble gases, optimization of polarisation levels for each nucleus remains an ongoing objective requiring systematic characterization of operational parameters including temperature, pressure, laser power, and gas mixture composition.

Contents

2.1	Introduction	
2.2	Polarisation Theory	
	2.2.1 Hyperpolarisation Methods 11	
	2.2.2 SEOP Theory	
2.3	Polariser Development	
	2.3.1 Goals	
	2.3.2 Polariser Overview	
	2.3.3 Heating	
	2.3.4 Coils	
	2.3.5 Low field NMR	
	2.3.6 Laser and Optics	
	2.3.7 Gas Manifold	
	2.3.8 Electronics	
	2.3.9 Polarisation Determination 47	
2.4	Results	
	2.4.1 Xenon-129	
	2.4.2 Krypton-83	
	2.4.3 Xenon-131	
2.5	Discussion	
2.6	Conclusion	
2.7	References	

Declaration of Contributions

- **Arthur Harrison** Design, procurement, manufacture and testing of all components described in this chapter, excluding the double-Helmholtz coil pair and polariser frame. Operation of the polariser and optimization of hyperpolarisation parameters.
- Max Filkins Operation of the polariser and optimization of hyperpolarisation parameters. All NMR analysis.
- **Dr Xinpei Wang** Manufacture of the first iteration of this polariser, onto which the components described below were built.
- UoN Physics Electronics Workshop Built the mains-power electrical components of the polariser and interlock system.

2.1 Introduction

2.2 Polarisation Theory

The intensity of the signal observed from NMR active nuclei during MR experiments is proportional to the nuclear polarisation. All nuclei with a non-zero spin quantum number will possess a polarisation which is defined as the difference in the population between the spin states. Take for example a nucleus with a quantum spin number of I = 1/2, which has two spin sublevels m = +1/2 and m = -1/2; the polarisation is defined as:

$$P = \frac{N_{+} - N_{-}}{N_{+} + N_{-}} \tag{2.1}$$

where N_+ and N_- are the spin populations of the +1/2 and -1/2 states respectively. In an external magnetic field the energy of states diverge due the Zeeman effect. At equilibrium, the relative population of each of these sublevels is dependent on the internal energy of the ensemble. Therefore, this is known as the thermal polarisation. For a spin I=1/2, this thermal polarisation is defined by:

$$P_{Thermal} = \tanh\left(\frac{\gamma \hbar B_0}{2k_B T}\right) \tag{2.2}$$

where γ is the gyromagnetic ratio, B_0 is the applied magnetic field, and k_B is Boltzmann's constant. However, this description breaks down beyond spin I=1 due to the quantum mechanical difference in the probability of transitions between each spin state[1]. Instead, one must derive the spin polarisation using the magnetisation (M_z) and the maximum magnetisation

at zero kelvin $(M_z^{T\to 0K})$.

$$P = \frac{M_z}{M_z^{T \to 0K}} \tag{2.3}$$

Unfortunately, due to the small energy differences between spin state sublevels, spin polarisation drops very quickly with increasing temperature[1]. Figure 2.1 demonstrates this temperature dependence near absolute zero for the spin-active noble gas nuclei; the polarisation decays substantially within the range of just half a kelvin. By the time one reaches ambient temperatures, the polarisation of the noble gas nuclei is of the order of 10^{-4} .

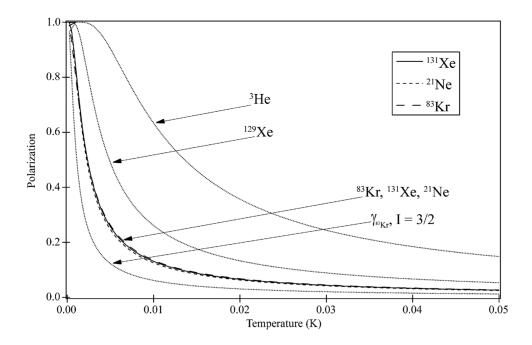


Figure 2.1: Temperature dependence of the thermal equilibrium polarisation of the spin active noble gas nuclei. Adapted from work by Stupic et al.[2] Including data simulated from a fictitious krypton isotope with the same gyromagnetic ratio but a spin I=3/2.

To attain the greatest SNR in MR experiments, it is desirable to conduct studies using nuclei with a high spin polarisation. This is particularly important for gas phase studies due to the low number density of atoms. However, it's often not feasible to cool samples to temperatures low enough to accomplish this, especially those involving living tissue. Instead, this can be achieved by artificially altering the spin state of the ensemble of nuclei, creating a hyperpolarised spin state, which is defined as a low nuclear spin temperature compared to the thermal temperature of the surroundings. This non-equilibrium state is temporary; the ensemble will 'relax' back to thermal polarisation given enough time. The temporal dependence of this relaxation depends on many factors and serves as the basis for many gas phase studies. Longitudinal relaxation times can vary greatly between the noble gas nuclei, from hours for helium-3 and xenon-129, to just seconds for the quadrupolar nuclei under identical conditions[1]. Therefore, noble gas hyperpolarisation must usually be conducted on-site when needed.

2.2.1 Hyperpolarisation Methods

The hyperpolarised spin state can be attained via a variety of techniques, each with its own strengths and weaknesses depending on the experimental requirements[3]. These techniques exploit phenomena such as enhanced Boltzmann polarization at low temperatures, selectively driven transitions within energy levels, and polarisation transfer via chemical reaction. The following section provides a brief description of several methods.

2.2.1.1 Brute Force Hyperpolarisation

Of the several different methods for creating the hyperpolarised state, the Brute Force [4, 5] method may be the easiest to grasp. The target nuclei are cooled to a very low temperature (< 3K) in the presence of a high

strength magnetic field (> 10T). Given sufficient time in this environment, the ensemble polarisation will stabilise to that denoted by Equation 2.2. This thermal polarisation is significantly increased compared to that at room temperature; however, notably the ensemble is not hyperpolarised in this state. By definition, hyperpolarisation is achieved via the rapid heating of the sample at a rate faster than the longitudinal relaxation time. This method can be applied to any nuclei with a non-zero quantum spin number, provided the longitudinal relaxation times allow enough time for sample heating. However, rapid decay of thermal equilibrium polarisation with temperature (Figure 2.1) necessitates very low temperatures which can be difficult to attain, significantly limiting the attainable polarisation levels. Furthermore, the time taken to stabilise at thermal equilibrium can be of the order of 10s to 100s of hours[5].

2.2.1.2 Dynamic Nuclear Polarisation (DNP)

DNP utilises the larger magnetic moments of electron spins, compared to nuclear spins, to amplify nuclear polarisation by orders of magnitude [6–10]. The theoretical limit of the attainable nuclear polarisation is dependent on the ratio of the gyromagnetic ratios of the electron and nucleus; for protons this is 660. The process typically involves introducing paramagnetic species (like stable radicals) to the sample, which, when kept at very low temperature and in a high magnetic field, will attain significant polarisation, similar to the brute force method described above. The transfer of this polarisation to the target nuclei is then facilitated by microwave radiation at specific frequencies, which, depending on the conditions, is achieved by one of several mechanisms: the Overhauser effect[11], the Solid effect[10], the Cross effect[12], and Thermal mixing[13]. Each mechanism varies in efficiency and is optimally applicable under different conditions,

such as the molecular mobility target nuclei; DNP can occur in liquids and solids. Modern 'dissolution' DNP approaches will then rapidly dissolve the hyperpolarised sample into a target molecule for investigation[14]. This polarisation method can be difficult to scale, requiring complex machinery which must be kept sterile for clinical applications. Furthermore, the fast longitudinal relaxation in the liquid state necessitates quick delivery to the patient immediately following production of the hyperpolarised species.

2.2.1.3 Parahydrogen Induced Polarisation (PHIP)

PHIP exploits the nuclear spin properties of parahydrogen (p-H₂), a nuclear spin isomer of molecular hydrogen [15–17]. At room temperature, molecular hydrogen can exist as two nuclear spin isomers, ortho- and para-, in an approximately 1:3 population ratio[18], although the p-H₂ state can be enriched at low temperatures in the presence of a catalyst. In this state, the two proton spins are in an antiparallel, singlet state, giving the molecule a total nuclear spin of zero and thus it will not produce an NMR signal. When parahydrogen undergoes a chemical reaction (generally hydrogenation), the H-H bond is broken and the hydrogen atoms are incorporated into a product molecule [19]. In this chemical environment, the enhanced polarisation of the incorporated hydrogen atoms becomes observable, creating amplified NMR signals for the host molecule. The attainable polarisation is dictated by factors such as the initial p-H₂ population and reaction rate. However, in this form, this method is limited to polarising unsaturated substrates that will react with molecular hydrogen. Modern methods of PHIP, such as Signal Amplification by Reversible Exchange (SABRE), circumvent this limitation, where polarisation transfer occurs through temporary association with a catalyst rather than chemical modification [20]. This expands the range of molecules that can be hyperpolarised using parahydrogenbased methods[21, 22]. The strength of this technique is that it doesn't require complex equipment and thus can be conducted at reduced cost compared to the other hyperpolarisation methods discussed. However, the necessity of a catalyst, usually heavy metal-based, can complicate in vivo experiments.

2.2.1.4 Spin Exchange Optical Pumping (SEOP)

SEOP employs alkali metal atoms as an intermediary for the transfer of spin angular momentum from incident laser photons to the nuclei of a gaseous species[23]. This two-stage process first involves the build-up of electron spin polarisation of an alkali metal vapour (optical pumping)[24], followed by transfer to the nucleus of the target species (spin exchange)[25]. This process is typically used for the hyperpolarisation of the noble gas isotope xenon-129, due to the relatively long lifetimes of the polarised spin state of this nucleus and because of its chemical inertness towards the highly reactive alkali metal vapour[23]. Attainable polarisation is dependent on a number of factors, such as laser power and alkali metal vapour density. Whilst enhancements of three to four orders of magnitude compared to Boltzmann polarisation are typical (at field strengths of 9.4 T), this process requires specialised equipment such as high-powered laser systems and bespoke glassware[1]. This method is employed in this work and is further explained in the following section.

2.2.2 SEOP Theory

Spin Exchange Optical Pumping is used to temporarily transfer noble gas nuclei into a non-equilibrium hyperpolarised spin state. The effective reduction in spin temperature enhances the magnetisation by many orders of magnitude compared to thermal equilibrium under standard conditions[7]. SEOP is a two-stage process whereby spin polarisation is passed from incident laser photons to the nuclear spin of a noble gas, via intermediary alkali metal atoms [23]. Circularly polarised laser photons are used to selectively pump the electron spin states of a saturated alkali metal vapour [26]. This electron polarisation is then transferred to noble gas nuclei via Fermi contact interactions either stochastically during binary collisions or coherently within temporary van der Waals molecules between the exchange pair [27]. This process often requires the mixing of a buffer gas (usually nitrogen) in order to quench the destructive radiation trapping originating from alkali metal fluorescence [23]. Of each of the alkali metals, SEOP has found most widespread success using rubidium, although caesium systems are available [28]. This is due to the low temperatures required to saturate the gas vapour, favourable cross sections during spin-exchange and spin-relaxation interactions and the required D1 resonance transition wavelength being well within the operating range of high intensity tunable laser systems [23].

2.2.2.1 Optical Pumping

The optical pumping process involves the selective excitation of the unpaired valence electrons of an alkali metal vapour, using resonant laser light. Figure 2.2 depicts a partial energy level diagram for a rubidium valence election, including the ground and lowest electron states. The incident laser light is tuned to the D1 transition of the rubidium vapour and left circularly polarised (photon spin $s_z = 1$) parallel to a static B_z field (helicity σ_+). Therefore, to maintain the total angular momentum of the system, induced excitation is limited to transitions from the $5S_{1/2}, m_j = -1/2$ state to the $5P_{1/2}, m_j = +1/2$ state (red solid arrow in Figure 2.2). From

this state, spontaneous decay will occur to the $5S_{1/2}, m_j = -1/2$ and $5S_{1/2}$, $m_j = +1/2$ sublevels with respective probabilities of 2/3 and 1/3[29]. However, due to collisions with other atoms or molecules in the gas phase, fluctuation spin-rotation interactions will result in the rapid equalisation of each of the $5P_{1/2}$ (L=1) sublevels. This is known as collisional mixing (Blue solid line in Figure 2.2), and results in an approximately equal probability of relaxation to each of the $5S_{1/2}$ sublevels (Blue dashed lines in Figure 2.2). Critically, in contrast to the p orbital, little mixing occurs between the $5S_{1/2}$, $m_j = -1/2$ and $5S_{1/2}$, $m_j = +1/2$ sublevels. This is due to the zero orbital angular momentum quantum number (L=0) associated with the spherically symmetric s orbital. Continuous irradiation will continue to pump the $5S_{1/2}$, $m_j = -1/2$ transition, continuing the optical pumping cycle. However, due to the $\Delta_m = +1$ selection rule, electrons in the $5S_{1/2}, m_j = +1/2$ sublevel will remain in this state for a significant period. Therefore, this process will result in the non-equilibrium build-up of the population of the $5S_{1/2}$, $m_j = +1/2$ sublevel, polarising the electrons.

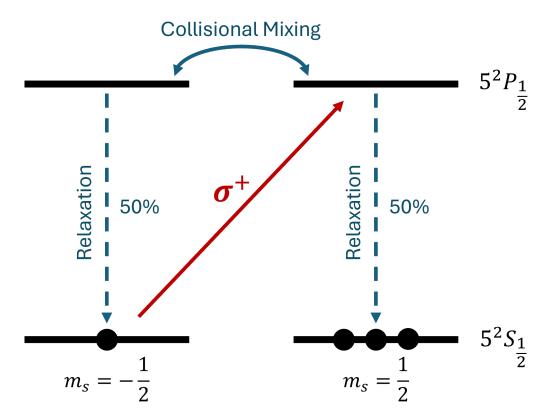


Figure 2.2: Partial energy level diagram of the valence electron of rubidium, including the ground state and the lowest excited state. Circularly polarised (σ^+) laser photons are used to excite the D₁ transition (red), continuously depleting the $5^2S_{m=-1/2}$ sublevel. Polarisation occurs due to the non-uniform spin population distribution created within the $5^2S_{1/2}$ energy level. Diagram not to scale.

As the election polarisation build up time is typically very fast, the alkali metal polarisation is normally expressed as the steady state polarisation, P_{AM} . This based on the relative rates of spin polarisation build up and destruction[29]:

$$P_{AM} = \frac{\gamma_{OP}}{\gamma_{OP} + \gamma_{SD}} \tag{2.4}$$

where γ_{OP} is the optical pumping rate and γ_{SD} is the spin destruction rate. The constructive term, γ_{OP} is caused by the incident laser light[29]:

$$\gamma_{OP} = \int \Phi(\lambda) \cdot \sigma(\lambda) d\lambda \tag{2.5}$$

where $\Phi(\lambda)$ is the laser flux, dependent on both laser power[30, 31] and linewidth,[32, 33] and $\sigma(\lambda)$ the absorption cross section, dependent on the alkali metal and the vapour density[1], and typically Lorentzian in shape[29]. Generally, SEOP using diode array lasers can benefit greatly from higher laser power[30, 31] and line narrowing is useful to ensure laser linewidth does not exceed the D_1 absorption line width[32, 33]. This absorption linewidth can be collisionally broadened through the introduction of high pressures of buffer gas[34], as discussed in Chapter 3.

However, the assumption of a uniform $\Phi(\lambda)$ in Equation 2.5 assumes an ideal scenario of uniform laser flux throughout the optically pumped medium. In real-world conditions, and assuming uniform alkali metal vapour density, absorption of the incident laser photons will result in a non-uniform optical pumping rate throughout the sample[1]. The alkali metal vapour absorbs incident laser light, acting as a translucent medium and casing an effective 'shadow' along the path of the laser. This magnitude of this phenomenon depends on both the density of the alkali metal vapour, the relative population of the $5^2S_{1/2}$, m=-1/2 sublevel and the geometry of the cell[29, 35]:

$$\frac{\mathrm{d}\Phi(\lambda,z)}{\mathrm{d}z} = -[AM]\sigma(\lambda)(1 - P_{AM}(z))\Phi(\lambda,z) \tag{2.6}$$

where z is the displacement through the cell along the laser path and [AM] is the density of the alkali metal vapour. In the extreme, all circularly polarised photons can be absorbed by the alkali metal vapour. If this condition occurs within the geometry of the cell, the resulting alkali metal

shadow will create a region of non-polarised (or 'dark') alkali metal. This can be highly destructive to the hyperpolarised state of the noble gas via the inverse of the spin exchange process described below.

The remaining term in equation 2.4, the spin destructive term γ_{SD} , has contributions from a number of sources:

$$\gamma_{SD} = \gamma_{trap} + \gamma_{BC} + \gamma_{vdW} \tag{2.7}$$

where γ_{trap} is due to radiation trapping, γ_{vdW} due to interactions with temporary van der Waals complexes, and γ_{BC} due to binary collisions of the electron spin polarised alkali metal atoms with additional species[1]. Radiation trapping stems from the reabsorption of energy lost due to relaxations between rubidium spin sublevels. This process is described further in Chapter 3. Depolarisation also occurs due to binary collisions between the electron polarised alkali metal and the other gas phase atoms. This depolarisation is dependent on the specific vapour density, [M], and spin destruction rate, κ_{sd} , between the atomic pair, summed for each additional gas phase species (i):

$$\gamma_{BC} = \sum_{i} \kappa_{sd}^{i}[M_{i}] \tag{2.8}$$

Literature values of for κ^i_{sd} for xenon-129 and krypton-83 and common species employed for SEOP can be found in work by Six et al.[36] As a very general rule, alkali metal collisions with larger atoms/molecules are the most destructive, with the magnitude for κ^i_{sd} following the order:

$$\kappa_{sd}^{Xe} > \kappa_{sd}^{Kr} >> \kappa_{sd}^{N_2} > \kappa_{sd}^{He} \tag{2.9}$$

As the magnitude of this destruction also depends on the number density of the corresponding gas phase atoms $[M_i]$, the higher the noble gas partial pressure, the lower the achievable alkali metal electron polarisation under otherwise identical conditions.

The final destructive term is due to temporary van der Waals complexes formed between the electron polarised alkali metal and other gas phase atoms. With sufficient lifetimes of these molecules, spin-rotation interactions can depolarise the atoms[37]. The form of γ_{vdW} is difficult to quantify; however, at lower pressures significant polarisation is lost within these complexes. As the pressure increases, so does the rate of formation of these molecules, although this is countered by their increasing rate of destruction[36]. Therefore, within the operating ranges of SEOP systems, the pressure dependence of this term is often considered negligible, especially compared to the contribution from binary collisions.

2.2.2.2 Spin Exchange

The second stage of SEOP is the transfer of the accumulated electron spin polarisation from the pumped alkali metal vapour to the nuclear spin of the noble gas. This takes place via Fermi-contact hyperfine interactions, inducing a 'flip-flop' exchange of polarisation between the species[27, 38]. This process is described by:

$$\alpha \overrightarrow{S} \cdot \overrightarrow{I} = \frac{\alpha}{2} [S_+ I_- + S_- I_+] + \alpha S_z I_z$$
 (2.10)

where I and S represent the nuclear and electron spins respectively and the term within the square brackets the 'flip-flop' exchange[38]. α is the coupling constant, proportional to the probability of finding the alkali metal valence electron at the nucleus of the noble gas:

$$\alpha = \frac{8\pi}{3} \gamma_S \gamma_I \hbar^2 \delta(r) \tag{2.11}$$

where γ_I and γ_S are the gyromagnetic ratios of the alkali metal electron and noble gas nucleus respectively and r the relative distance between the exchange pair[38]. Spin exchange can occur via two mechanisms: either during binary collisions between the alkali metal and noble gas atoms, or during temporary van der Waals molecules formed between the species. These pathways are demonstrated in Figure 2.3.

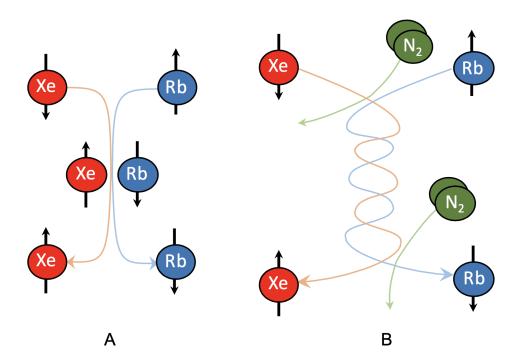


Figure 2.3: The mechanisms of spin transfer between the electron polarised rubidium atoms and noble gas nuclei, in this case xenon. Spin can be passed via binary collisions (A) or within temporary van der Waals complexes (B) often mediated by a third body, displayed here as nitrogen[23, 39].

Intuitively, binary collisions are two-body in nature, occurring at a rate that is dependent on both the specific alkali-metal—noble-gas cross-section and partial pressure of each species. Polarisation transfer within temporary van der Waals complexes is mediated by a third body, usually a nitrogen buffer gas. Alkali metal and noble gas atoms colliding in the presence of this third body can form weakly bound van der Waals molecules. This molecule will evolve freely, until it is dissociated by an additional third body. Polarisation transfer probability within this molecule is based on the lifetime of the complex, which is non-linearly dependent on gas pressure[25]. Both the rate of formation and destruction increase with increased pressure. However, the resulting reduced complex lifetime at higher pressures will suppress the ensemble polarisation transfer via this pathway. The relative occurrence of each pathway is also dependent on the noble gas nuclei. For the smaller noble gases, such as helium-3, polarisation transfer is normally dominated by binary collisions[23]. However, for heavier noble gases such as xenon-129 at sufficient partial pressure, spin transfer occurs mainly within van der Waals complexes[23].

The spin exchange rate, γ_{SE} , quantifies the transfer of spin polarisation from the pumped alkali metal elections to the noble gas nuclei:

$$\gamma_{SE} = [AM] \left(\langle \sigma v \rangle + \frac{\gamma_{AMNG}}{[NG]} \left(\frac{1}{1 + br} \right) \right)$$
(2.12)

where the first term within the brackets $\langle \sigma v \rangle$, is the velocity averaged binary spin exchange cross section and describes the spin transfer contribution from binary collisions[1]. The final term within the brackets quantifies spin transfer with the van der Waals dimers, characterised by the rate constant γ_{AMNG} , which is specific to each alkali metal to noble gas pair. The noble gas density is destructive to this term, as higher pressures will result in the breakup of these complexes. The buffer gas molecules within the mixture will also contribute to the dimer breakup, with these contributions quantified by $b = p_0(NG)/p_0(BG)$ the characteristic pres-

sure ratio and r = p(BG)/p(NG) the partial pressure ratio[1]. Literature values for each of these terms can be found in work by Six et al.[36] The dilution of xenon with nitrogen can be beneficial to the spin exchange of xenon, although detrimental to krypton due to the greater contribution to the dimer breakup[1]. Both spin transfer mechanisms are mediated by the number density of alkali metal vapour, [AM]. $\langle \sigma v \rangle$ and γ_{AMNG} are 2-3 orders of magnitude smaller for krypton-83 SEOP compared to that of xenon-129[36]. Therefore, to counter the reduced exchange rate, the alkali metal vapour pressure is increased by raising the temperature at which the SEOP is conducted[40].

Finally, using the pumped alkali metal polarisation, P_{AM} and the spin exchange rate γ_{SE} , one can quantify the time-dependent noble gas nuclear polarisation[34]:

$$P_{NG}(t_p) = \frac{\gamma_{SE}}{\gamma_{SE} + \Gamma} \cdot P_{AM} \left(1 - exp(-(\gamma_{SE} + \Gamma)t_p) \right)$$
 (2.13)

where Γ is the relaxation rate, or more generally the inverse of the noble gas longitudinal (T_1) relaxation time. Within the brackets is the time dependent term, reflecting the build-up in polarisation of the noble gas nuclei with time. For sufficiently long SEOP pump times (t_p) , this term vanishes; the noble gas has reached its steady state polarisation. If the spin exchange rate is much greater than the self-relaxation rate, the steady state noble gas polarisation will closely reflect that of the alkali metal electron polarisation. However, this is not the case for quadrupolar nuclei where Γ is high due to the fast quadrupolar relaxation[1].

2.2.2.3 Noble Gas Depolarisation

The above sections detail the mechanisms for the build up of noble gas nuclear spin polarisation during spin exchange optical pumping. However, the limiting factor to ensemble polarisation is characterised by the relaxation rate Γ . This term is omitted from Figure 2.2, although for rubidium, describes transitions between the $5S_{1/2}$ sublevels. High relaxation rates also serve to limit the lifetime of the hyperpolarised state after production as they cause the return to thermal spin equilibrium.

The relaxation rate is made up of a linear combination of relaxation pathways, which can be separated into two groups: intrinsic, due to interactions within the gas ensemble itself, and extrinsic, due to external factors:

$$\Gamma = \Gamma_{in} + \Gamma_{ex} \tag{2.14}$$

where Γ_{in} represents the intrinsic contributions and Γ_{ex} the extrinsic contributions. For hyperpolarised noble gas nuclei, the intrinsic contribution is made up of two primary components representing types of self-collisions of the atoms:

$$\Gamma_{in} = \Gamma_{BC} + \Gamma_{vdW} \tag{2.15}$$

where Γ_{BC} is due to binary collisions and Γ_{vdW} due to temporary van der Waals complexes of the gas phase species. The binary collisions term is dependent on the noble gas density:

$$\Gamma_{BC} = \gamma_{BC}^{NG}[NG] \tag{2.16}$$

where γ_{BC}^{NG} is the noble gas specific density dependence of the longitudinal relaxation and [NG] is in amagat. For xenon-129, this $\gamma_{BC}^{Xe129}=5.0\pm0.5\times10^{-6}~{\rm amagat^{-1}}s^{-1}$ [41] and krypton-83 $\gamma_{BC}^{Kr83}=1.6\pm0.1\times10^{-3}~{\rm amagat^{-1}}s^{-1}$ [42].

At low field, the contribution from temporary van der Waals molecules is independent of total gas density[43, 44]. Instead, this term is dependent on the relative concentration of noble gas and additional species, typically the buffer gas. For xenon-129, the contribution from this mechanism is dictated by:

$$\Gamma_{vdW}^{Xe129} = \frac{\gamma_{vdW}^{Xe}}{\left(1 + r\frac{[BG]}{[Xe]}\right)} \tag{2.17}$$

where γ_{vdW}^{Xe} is the Xe-Xe van der Waals relaxation term and $r = k_{BG}/k_{Xe}$, where k_{BG} and k_{Xe} represent the breakup rate coefficients for the buffer gas and xenon atoms, respectively. For xenon-129, $\gamma_{vdW}^{Xe129} = 6.72 \pm 0.1 \times 10^{-5} s^{-1}$ [43]. This term is complicated for the quadrupolar nuclei, due to the quadrupolar interactions within the noble gas dimer. For krypton-83, this contribution is described by [44]:

$$\Gamma_{vdW} = M_Q \frac{K}{\left(k_{kr} + k_{BG} \frac{[NG]}{[Kr]}\right)}$$
(2.18)

where K is the chemical equilibrium constant for krypton dimer formation, $K = [Kr_2]/[Kr]^2$, and M_Q represents the quadrupolar term in a linear molecule [45]:

$$M_Q \propto \frac{2I+3}{I^2(2I-1)} \left[\frac{eQq}{\hbar}\right]^2$$
 (2.19)

where eQq/\hbar is the quadrupolar coupling constant, dependent on electronic charge, e, the nuclear quadrupolar moment, Q, and the charge of the nucleus, q.

The relaxation rate due to extrinsic contributions, Γ_{ex} , during typical conditions for SEOP is made up of four terms:

$$\Gamma_{ex} = \Gamma_{\nabla B} + \Gamma_P + \Gamma_W \tag{2.20}$$

where $\Gamma_{\nabla B}$ is due to transit through magnetic field gradients, Γ_P due to interactions with paramagnetic species in the gas after SEOP, Γ_W due to interactions with the wall of the containment vessel, and for quadrupolar nuclei, Γ_Q due to electric quadrupolar couplings.

Molecular movement through magnetic field gradients (∇B) can be caused by diffusion or transit of the hyperpolarised species. The relaxation rate due to this effect is determined by[46]:

$$\Gamma_{\nabla B} \propto \frac{|\nabla B|^2}{B^2}$$
(2.21)

Furthermore, regions of very low or zero magnetic field will cause rapid depolarisation due to the degeneracy of the Zeeman energy levels.

Intermolecular magnetic dipole coupling with paramagnetic species in the gas phase can also cause depolarisation of the noble gas spin state. Typically for SEOP experiments this can stem from oxygen contamination in the gas mixture following SEOP, although this can also be caused by free radicals or any paramagnetic molecule in the gas phase. For oxygen-xenon-129 interactions, the resultant relaxation rate has been empirically

determined[47]:

$$\Gamma_{O_2}^{Xe129} = 0.388 \frac{p_{O_2}}{1.013 \text{ bar}} \frac{273 \text{ K}}{T} \left(\frac{300 \text{ K}}{T}\right)^{0.03}$$
(2.22)

where p_{O_2} is the oxygen partial pressure (in bar) and T is the gas temperature (in kelvin). For xenon in air ($p_{O_2} = 0.21$ bar) under standard conditions, $\Gamma_{O_2}^{Xe129} = 7.4 \times 10^{-2} s^{-1}$ (or $T_1 \approx 14$ s). The equivalent relaxation pathway for krypton-83 is approximately two orders of magnitude smaller than for xenon-129, due to the difference in the square of the gyromagnetic ratios[40]. However, unlike for the spin half noble gas nuclei[47, 48], the relaxation rate is difficult to empirically determine due to the magnetic quadrupole moment.

Interactions with the walls of the containment vessel can also cause relaxation of the hyperpolarised ensemble. The term Γ_W describes the spin destruction rate of this interaction, although omits contributions from magnetic quadrupolar moments. These surface relaxation mechanisms are complex, however can be summarised as[49, 50]:

$$\Gamma_W = \frac{1}{n} \frac{S}{V} \tag{2.23}$$

where S/V is the surface-to-volume ratio of the containment vessel and η the relaxivity of the cell walls, dependent on material, temperature and magnetic field strength. Particularly for xenon-129, relaxations of this kind stem from paramagnetic impurities within the glass surface. Surface coatings are often employed to create a barrier between the xenon-129 atoms and these impurities[51, 52], reducing the relaxation rate due to this mechanism.

Finally, Γ_W can be adapted for quadrupolar nuclei, introducing the term Γ_Q , which describes the relaxation due to surface interactions with the magnetic quadrupole moment of the nuclei. Therefore, this term is not present for the spin 1/2 nuclei, helium-3 and xenon-129. This relaxation mechanism stems from the interaction of the non-spherical charge distribution of the quadrupolar nucleus with the electric field gradients associated with the surface. For quadrupolar nuclei (I>1/2), this mechanism will likely dominate the relaxation rate and is sensitive factors such as surface chemistry, temperature and hydration[53, 54]. Indeed, this mechanism has been employed as a source of contrast using krypton-83[55, 56]. The relaxation rate due to this mechanism is described by[44, 57]:

$$\Gamma_Q \propto M_Q \tau_c$$
 (2.24)

where M_Q takes the same form as Equation 2.18, and represents the surface quadrupolar interactions and τ_c represents the correlation time, approximately equal to the average residency time at a particular adsorption site on the surface.

2.3 Polariser Development

2.3.1 Goals

Most modern polarisers are typically highly specialised devices, able to attain very high noble gas polarisation when operating with specific procedures and gas mixtures. Established 'large volume' polarisers are generally limited to the spin 1/2 nuclei, helium-3 and xenon-129, although previous

attempts have been made to polarise quadrupolar nuclei such as xenon-131 in these systems[58]. The aim of this project was to develop a highly adaptable, large volume pre-clinical polariser. This includes the capability to polarise multiple different noble gases with a range of buffer gases, operating under either batch or continuous-flow conditions. This includes the quadrupolar noble gas krypton-83, for which there exists no large quantity polariser (>0.5 L) in the literature; a key stepping stone on the path to future clinical application.

2.3.2 Polariser Overview

The polariser has five main components, the heating system, coils, laser and optics, gas manifolds, and electronics. Each of these components is described in the following sections, however the polariser as a whole is depicted in Figure 2.4.

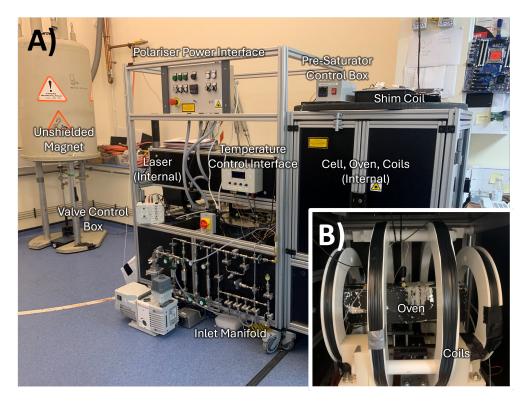


Figure 2.4: Labelled image of the polariser (A), with internal view of the Helmholtz coils and oven (B).

2.3.3 Heating

The temperature of the SEOP cell is controlled using a forced-air oven. The oven is manufactured primarily from aluminium, however the cell is supported by PEEK mounts due to the lower thermal conductivity of this material. This setup is depicted in Figure 2.5. The oven is fed by two heat guns, which can be controlled independently to create two heating zones. Three thermocouples (TC) are placed on the cell. Two are positioned immediately above the heat-gun inlets and used to control the output of each heat gun. The third thermocouple is placed on top of the cell to give insight into the in-cell conditions.

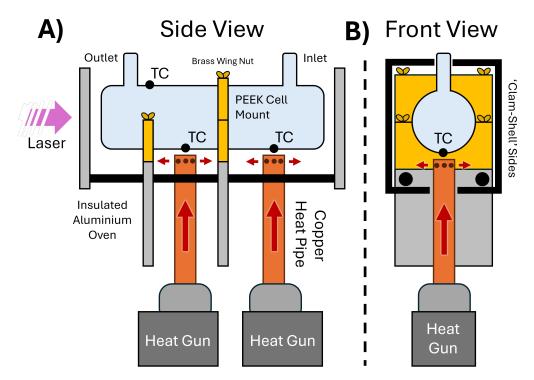


Figure 2.5: The thermal management system of the SEOP Cell, depicted from the side (A) and front (B) profile. Thermocouple placement is indicated by TC.

Each heat gun is controlled using manually tuned PID control software in an Arduino Mega 2560 (Arduino, Italy). The Arduino reads the temperature values from the thermocouples and sends an analogue control signal to each heat gun, which themselves receive 240 V ac power. The electronic diagram for this system is depicted in Figure 2.6.

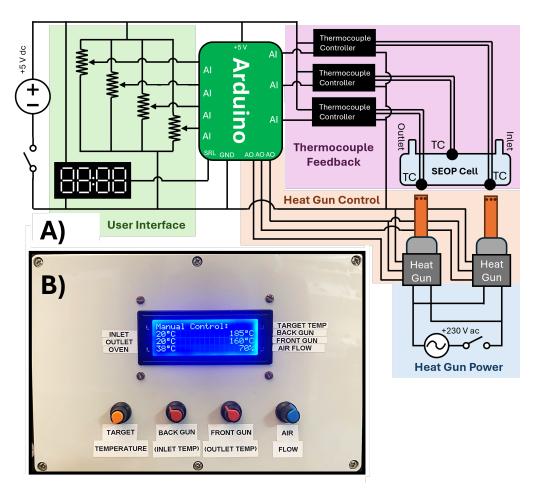


Figure 2.6: The interface used to control the temperate of the SEOP cell. Including the electronic diagram (A) and a picture of the physical interface (B).

2.3.4 Coils

A double Helmholtz coil pair is employed to create a uniform magnetic field within the optical pumping cell. This is depicted in Figure 2.7. However, the stray field from a nearby unshielded 9.4 T magnet (Oxford Instruments, UK) located near the polariser causes a significant magnetic field perpendicular to the laser path. This can permit unwanted transitions within alkali metal vapour, and reduce the attainable polarisation in proportion to the cosine of the angle between the resultant field and the laser path[1]. Therefore, a shim coil is installed on the roof of the polariser to reduce the

magnitude of this transverse stray field at the cell.

The coils are connected in series to a dedicated power supply. The direction of current within the coils can be changed using a custom-built 'switching box', although the use of this switching box is not part of standard operating procedure of the polariser. However, this is convenient if an experiment require the reversal of the magnetic field compared to the laser propagation direction, such as for polarisation build-up experiments.

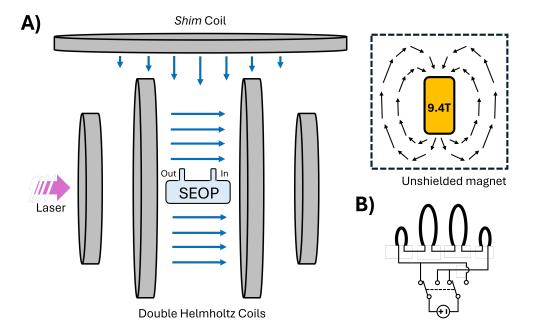


Figure 2.7: The double Helmholtz coil pair used for creating the static magnetic field within the SEOP cell, including a visual diagram (A), electronic diagram (B) and a representation of the stray field from the nearby unshielded magnet.

2.3.5 Low field NMR

All polarisation measurements quoted in this work were determined in reference to a thermally polarised sample of the same nuclei using a 9.4 T magnet (Bruker, Germany), described further in Section 2.3.9. However, whilst this method is accurate for measuring the polarisation of the deliv-

ered gas, this method requires the extraction of hyperpolarised gas from the cell, which can be highly inefficient for polarisation optimisation experiments. This will change the pumping conditions within the cell. Instead, the polarisation within the cell can be monitored in-situ using NMR and the B_0 field of the coils.

A 350 turn, 1 cm in diameter coil was used, described previously in work by Irwin et al.[59]. This coil is placed above the centre of the cell, secured using the PEEK mountings. A custom 'tuning box' was created to resonate at the natural frequency of either xenon-129 in a field of approximately 30 G, or krypton-83 in a field of approximately 70 G:

$$\omega_0 = \frac{1}{\sqrt{LC}} = \gamma B_0 \tag{2.25}$$

where L is the inductance of the coil, C is the capacitance of the capacitor and γ is the gyromagnetic ratio of the noble gas nuclei. The resonance of this circuit is manually selected using a switch to connect the desired capacitor. The resonant frequency of the noble gas nuclei is then matched to the resonance of the LC circuit using fine adjustments of B_0 . The custom switching tuning box is depicted in Figure 2.8.

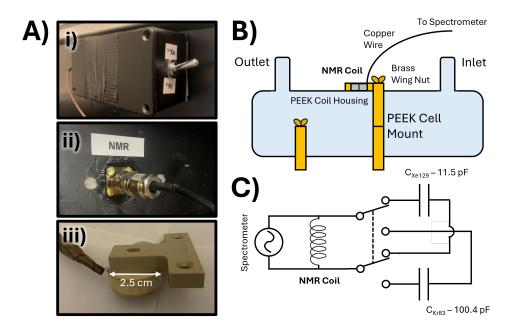


Figure 2.8: The switching-tuning box for on-cell NMR. Including pictures of the system (A) exterior (i) and interior (ii) to the polariser housing, as well as the coil attached to its mounting bracket (iii). A diagram of the coil mounting method (B) is presented, as well as the circuit diagram (C) of the system.

The resonant frequency of the xenon-129 circuit was 34.54 kHz, and that for the krypton-83 circuit was 11.68 kHz. The required B_0 field to detect each nucleus and hydrogen is detailed in Table 2.1.

Required B_0 Field			
Nuclei	Xenon-129 Circuit	Krypton-83 Circuit	
$^{1}\mathrm{H}$	8.11 G	2.74 G	
$^{83}{ m Kr}$	210.08 G	71.04 G	
^{129}Xe	29.12 G	9.85 G	

Table 2.1: Required B0 to tune the resonant frequency of ¹H, ⁸3Kr and ¹29Xe to each of the LC circuits.

Unfortunately, the stray field created by the external unshielded magnet (Figure 2.7) resulted in large B_0 inhomogeneity within the coil detection region. This resulted in a rapid decay of the FID and large linewidth of the

resulting spectral analysis. The use of the shim coil slightly reduced the FWHM of the xenon-129 peak. However, sufficient RF noise was detected from internal components of the system, primarily the heat guns and Peltier devices in the Optigrate lasers. Therefore, the SNR of the xenon signal was determined to be insufficient for reliable measurements. Furthermore, no signal was ever found for the krypton-83 nuclei, likely due to these factors and the reduced attainable polarisations of this nucleus[40].

2.3.6 Laser and Optics

The polariser has mountings capable of accepting three different narrowed laser models: either the 60 W or 95 W Shark lasers from Optigrate (USA), or the 180 W laser from QPC (USA). Details of these systems are presented in Table 2.2.

Laser Systems				
Manufacturer	Model	Power	Linewidth	
			(FWHM)	
OptiGrate	LC-795-65W	65 W	0.15 nm	
OptiGrate	LC-795-95W	95 W	0.20 nm	
QPC	6507-0003	180 W	0.25 nm	

Table 2.2: Details of the laser systems employed in the polariser.

The Shark lasers produce an approximately 90% linearly polarised beam. This is passed through a beam splitter to filter the non-linearly polarised light, which is discarded to the adjacent beam dump. The circular polarisation is then created using a quarter wave plate, positioned downstream in the optical path and oriented with its slow axis at 45 degrees to the linear polarisation. Finally, collimating optical elements are used to shape

the beam into a 25 mm diameter disk, the same size as the optical cells. This disk of light is used to illuminate the cell, a short distance away. The optical setup for the Shark lasers is displayed in Figure 2.9. Accurate alignment of each of these three components is required to achieve the greatest laser power illuminating the cell. This process is simplified for the QPC laser system, as all optical elements discussed above are integrated into the laser unit.

The optical system is designed for easy laser alignment. The three discrete components of the optical path, the laser, the optics and the SEOP cell are all movable, although the degrees of freedom are not consistent. The laser is positioned on rails and can move laterally relative to the optical path. The optics are mounted on a moveable jack and two fine-adjustment plates, for movement in all three axes relative to the optical path. Finally the SEOP cell mountings are fitted to a moveable jack, for vertical adjustment relative to the optical path.

The laser light transmitted through the cell is monitored using an optical spectrometer. This is used for monitoring the optical pumping system and troubleshooting errors. The intensity of the transmitted light yield information related to the rubidium vapour densities and electron polarisation. The remaining transmitted light, not used for spectral monitoring, is dumped into the back plate of the polariser. The heat deposited in this plate is radiated into the environment.

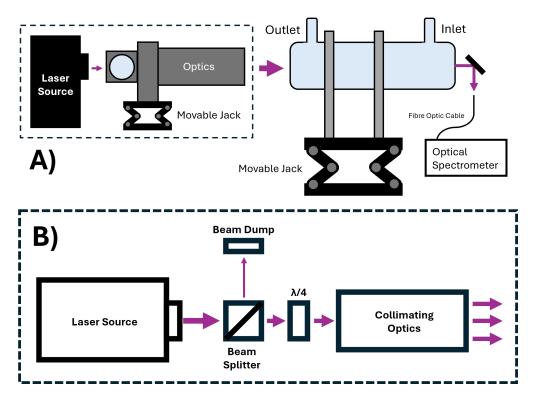


Figure 2.9: The optical setup of the polariser for the Shark laser system is illustrated, including a visual (A) and optical diagram (B) of the core equipment.

In addition to the laser systems, and not displayed in Figure 2.9, are the water-based thermal control systems for each laser.

2.3.7 Gas Manifold

Central to the polariser design are the pneumatic manifolds, for the containment and transport of gas within the system. The core manifold for all operational modes of the polariser is depicted in Figure 2.10. Gas mixtures for SEOP are prepared externally to the polariser, or ordered directly from commercial suppliers, and delivered to the system via regulated gas cylinders. These regulated cylinders are connected to one of three gas inlets, in the gas inlet manifold (Figure 2.10A). This gas then passes through one of four available gas purifiers, to remove potential contamination from

O₂ or H₂O. Each purifier is designated for a specific gas mix, due to the possibility of contamination via outgassing from the purifier when switching mixtures. A high (0.1 to 20 bar) pressure sensor is integrated into the manifold via a three-way plug valve. This sensor can monitor either the gas pressure from the delivery cylinders, or the gas pressure entering the SEOP cell, which may differ due to gas flow restriction form the purifiers or the metering valve.

Following delivery and purification in the inlet manifold, the gas mixtures are delivered to the SEOP cell (Figure 2.10B) via a controlled pneumatic valve. All controlled valves on the polariser are controlled via the interface box depicted in Figure 2.12. These valves are positioned close to the SEOP cell to reduce the volume of 'dead gas' within the transfer lines. Pneumatic valves are selected for this purpose to reduce exposure of the hyperpolarised gas to the magnetic field gradients typical of solenoid valves. However, both pneumatic valves in this system are controlled via high-pressure gas lines, which themselves are regulated by three-way solenoid valves positioned remote from the SEOP gas lines. This high-pressure manifold is not depicted in this work.

A typical SEOP cell for use in either batch-mode or continuous-flow experiments is depicted in 2.10. Gas is delivered into the SEOP cell via the inlet stem, positioned at the optical back of the cell. Following SEOP, the hyperpolarised gas is extracted via pressure differential from the outlet stem at the optical front of the cell. This transport line is regulated by the second controlled pneumatic valve. All controlled valves are set to normally-closed, to isolate the SEOP cell in the event of a power loss.

The final manifold within the polariser is the Post SEOP Utility Manifold (Figure 2.10C). This section is used for the set-up and cleaning routines of

the polariser; therefore, the hyperpolarised gas will typically pass through this system quickly. This manifold contains the final controlled valve, a solenoid valve for the vacuum outlet. This is used for cleaning the transfer line between hyperpolarised gas deliveries.

The final manifold depicted in Figure 2.10 is the Distribution Manifold (Figure 2.10D). This manifold is used to mediate the delivery of hyperpolarised gas to experiments. However, it is not integrated into the polariser itself and is instead positioned adjacent to the NMR/MRI spectrometer. During core operational modes, this manifold is connected directly to the polariser via a short length of transfer line. However, if hyperpolarised gas purification or compression is required, systems for these purposes are positioned between the polariser and this manifold. Furthermore, this manifold is always the final stage before hyperpolarised gas delivery to experiments. Finally, the pressure sensor on this manifold is used to monitor the hyperpolarised gas pressure during experiments.

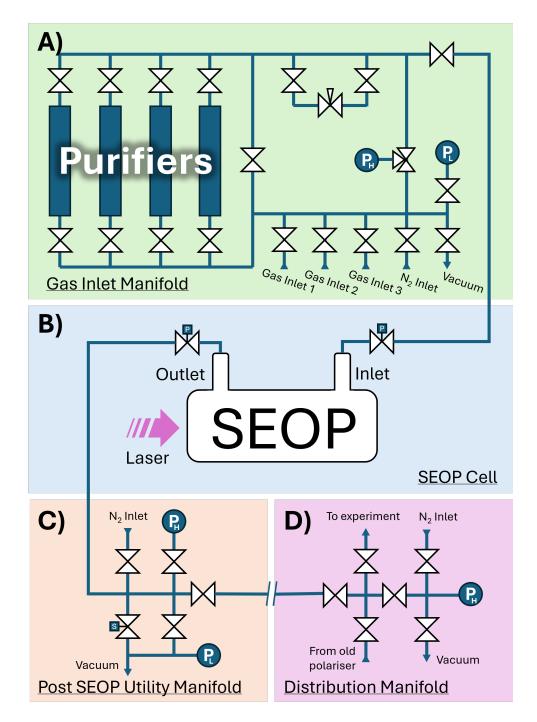


Figure 2.10: Pneumatic manifolds of the polariser, including the Gas Inlet Manifold (A), the SEOP Cell (B), the Post SEOP Utility Manifold (C) and the external Distribution Manifold (D).

The above description and Figure 2.10 concerns the gas manifolds for the core operation of the cell. However, this polariser can be easily adapted, either for additional functionality or for optimisation of specific operational

modes. Two pneumatic modifications frequently employed in our laboratory are depicted in Figure 2.11.

First, the SEOP cell can be substituted for the SEOP pre-saturator vessel (Figure 2.11A), which is optimised for continuous flow operation. The benefit of this cell is that the temperature dictating rubidium loading of the gas phase can be decoupled from the temperature of the SEOP chamber. Additionally, the lower volume of this chamber increases the rubidium surface to volume ratio, which improves vapour distributions. Furthermore, with the laser shield in position, the likelihood of a rubidium run-away is reduced, allowing for more precise control of the rubidium loading conditions. The temperature of the pre-saturator chamber is mediated via a heating cuff, controlled with the electrical diagram depicted in Figure 2.13. The cuff is thermally coupled to the chamber via a folded graphite sheet. This cell, with a rubidium pre-saturator, is less susceptible to polarisation reductions due to rubidium contamination, compared to the typical cell depicted in Figure 2.10B. The decoupling of the Rb saturation and SEOP temperatures means that the prior can be raised to account for a lower vapour pressure due to Rb contamination, without the forfeiture of the attainable polarisation associated with SEOP at higher temperatures.

This pre-saturator fitted cell also employs a rubidium trap, upstream of the pre-saturator chamber. Over long periods of gas flow, contaminated rubidium appears to 'creep' towards the source of the contamination. If this contamination stems from the gas mix, this means that eventually rubidium will make its way upstream to the inlet manifold. For batch-mode cells, the cell is normally considered poisoned and replaced before rubidium can reach the inlet stem. However, for the constant gas flows characteristic of continuous-flow mode, this effect can occur quickly, particularly when using highly contaminated gas mixes. Therefore, the rubidium trap consists of a

small glass vessel with the termination of the PFA inlet tube suspended at its centre. This suspension creates an effective break in continuous surface leading back to the inlet manifold and thus slows the rate of rubidium 'creep' upstream. Additionally, this trap serves as an additional purifier to the inlet gas stream, as any remaining impurities may react with this rubidium before reaching the pre-saturator chamber. Indeed, observation of these cells after significant use shows silver rubidium within the pre-saturator and brown deposits typical of rubidium contamination in the rubidium trap.

Figure 2.11B depicts a typical setup for the cryogenic purification of hyperpolarised xenon-129. The requirement for hyperpolarised gas purification is discussed further in Chapter 3. Typically, this apparatus is employed with the polariser operating under continuous-flow conditions and its operation follows largely standard operational procedures [30, 34, 60]. The flow-ratecontrolled hyperpolarised gas is passed through a cold finger, which itself is submerged in liquid nitrogen. In this system, the hyperpolarised xenon condenses to form a solid, while the now-redundant buffer gases are exhausted to the atmosphere. Xenon is accumulated in a high magnetic field until the requisite quantity of hyperpolarised medium has been collected. Then it is thawed using a vessel of room-temperature water. Typically, the gas flow, pressure, mixture composition and accumulation time are optimised to yield the greatest polarisation of xenon batch produced. Clinical investigations of the human lung commonly require 0.6 to 1.0 L of hyperpolarised xenon-129 and employ 3\% xenon mixtures [61]. However, the preclinical experiments for which this polariser was built typically require significantly reduced quantities of hyperpolarised gas (between 50 - 200 ml)[55]. Therefore, gas mixtures containing 1% xenon can be used without significant extension to the accumulation times [62] or changes to pumping conditions[36] which could forfeit polarisation. This optimisation can significantly increase the attainable polarisation during SEOP compared to 3% or 5% xenon mixtures. Finally, the concentrated hyperpolarised xenon is delivered to the experiment via the distribution manifolds. A Tedlar bag is sometimes introduced to the experimental setup to act as an 'overflow' volume. This prevents the experimental gas pressures from exceeding ambient pressure.

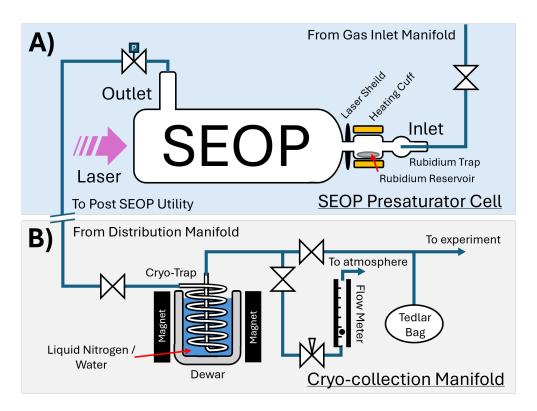


Figure 2.11: Typical substitutions or additions to the polariser pneumatics, including the Rb pre-saturator fitted flow-SEOP cell (A) and the cryogenic purification setup (B).

Two additional pneumatic substitutions are commonly made to the polariser, which are not depicted in Figure 2.11. The first is the apparatus for the purification of hyperpolarised species which utilise hydrogen as the sole buffer gas. This has been omitted from this section as Chapter 3 is devoted to its description. Secondly, a piston-based recompression system can be employed to increase the pressure of hyperpolarised gas delivery

to experiments. This system has been described previously in work by Hughes-Riley et al.[40]

2.3.8 Electronics

The pneumatic valves on the inlet and outlet of the cell are each controlled via high pressure (5 bar) lines, mediated by three-way solenoid valves. The power to these solenoid valves is controlled by the polariser control box, pictured in Figure 2.12. This control box permits either the manual or computer-based operation of the valves, determined using a manual switch. This control box can also control the vacuum line on the outlet manifold.

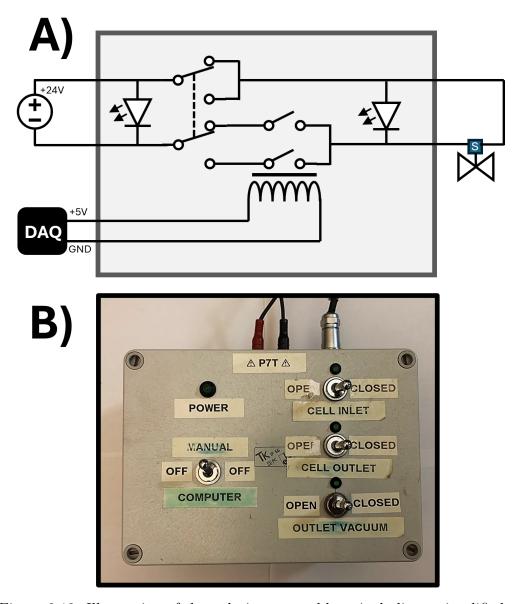


Figure 2.12: Illustration of the polariser control box, including a simplified circuit diagram (A) with only one solenoid valve depicted and picture of the physical interface (B).

The temperature of the heating cuff, depicted in Figure 2.11 is controlled using the pre-saturator control box. This apparatus reads the temperature of the thermocouple and mediates a 230 V AC current to the cuff, using an auto-tuned PID controller. This device is pictured in Figure 2.13.

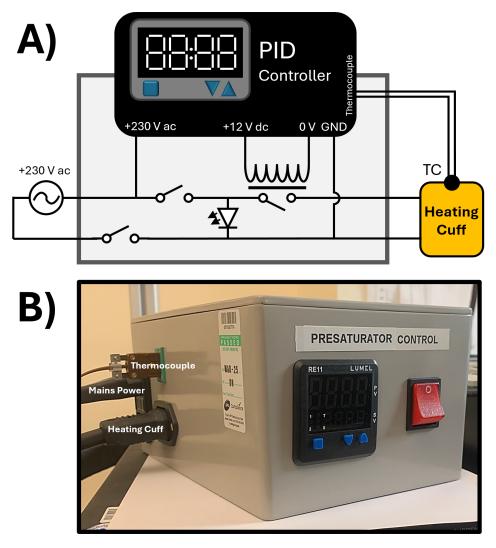


Figure 2.13: Illustration of the polariser control box, including a circuit diagram (A) and picture of the physical interface (B).

2.3.9 Polarisation Determination

All polarisation measurements were made using a 9.4 T Bruker Avance III spectrometer (Bruker, Germany). Gas was transferred from the SEOP cell into a glass vessel within the magnet, where is was subjected to a calibrated 90° RF pulse. The magnitude of the resultant signal is compared to a thermal reference sample to calculate the sample polarisation[1]:

$$P_0^{300K,9.4T} \frac{v_0^{B_0}}{v_0^{9.4T}} \frac{300K}{T} \tag{2.26}$$

In practice, differences in pressure, gas composition will change the magnitude of the detected signal, therefore the above equation can be adapted to account for these factors.

2.4 Results

The polariser was used to produce hyperpolarised xenon-129, krypton-83 and xenon-131. Due to the often changing dynamics in the SEOP cell, it is difficult to quantify the optimal polarisation attainable for each of these nuclei. Therefore the data presented is representative of a typical experimental day.

2.4.1 Xenon-129

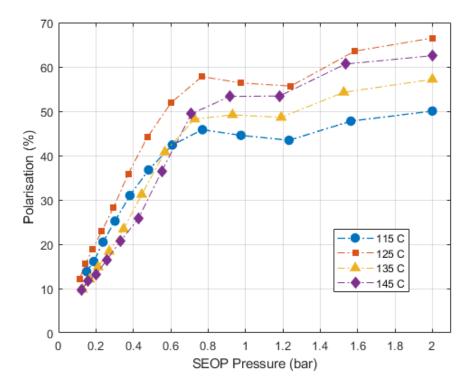


Figure 2.14: Polarisation vs SEOP pressure relationship of xenon-129 at a range of temperatures. Data was collected using the 65 W laser, with a $1\% \text{ Xe}^{129}$, $10\% \text{ N}_2$, 89% He mix in a 600 ml cell. Experiments were started at 2 bar with a pump time of 10 minutes. Gas batches were successively extracted from the cell in 5 minute intervals. The reduction in cell pressure and respective polarisation of the next batch yielded this plot.

2.4.2 Krypton-83

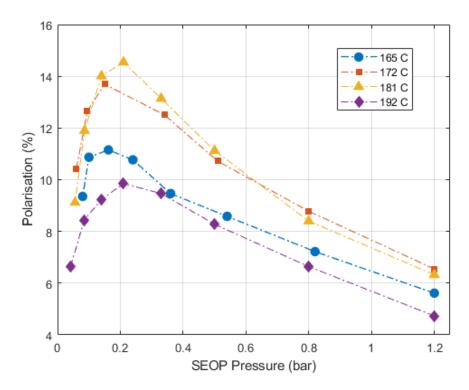


Figure 2.15: Polarisation vs SEOP pressure relationship of krypton-83 at a range of temperatures. Data was collected using the 65 W laser, with a 5% Kr (NA), 95% N_2 mix in a 600 ml cell. Experiments were started at 1.2 bar with a pump time of 8 minutes. Gas batches were successively extracted from the cell in 3 minute intervals. The reduction in cell pressure and respective polarisation of the next batch yielded this plot.

2.4.3 Xenon-131

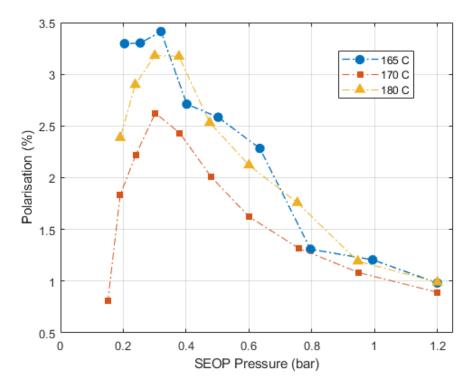


Figure 2.16: Polarisation vs SEOP pressure relationship of xenon-131 at a range of temperatures. Data was collected using the 65 W laser, with a 5% Xe (NA), 95% N_2 mix in a 300 ml cell. Experiments were started at 1.2 bar with a pump time of 8 minutes. Gas batches were successively extracted from the cell in 3 minute intervals. The reduction in cell pressure and respective polarisation of the next batch yielded this plot.

2.5 Discussion

The polariser has been demonstrated to produce hyperpolarised xenon-129, krypton-83 and xenon-131. However, significant work is required to determine optimal polarisation conditions and characterise the attainable polarisation of these nuclei using this system. This includes both batch mode operation of each nucleus, as well as continuous flow operation of xenon-129.

The strength of this polariser stems from its versatility: its ability to polarise multiple different nuclei with only minor changes to the experimental setup. However, this system is likely not able to match the magnitude of polarisation attainable using systems dedicated to specific nuclei or operational modes [63–68].

The operation of this polariser may be more complicated than established commercial systems as the operational parameters can vary significantly between nuclei. Hence stems the impetus to design simple intuitive features such as the gas delivery manifold (Figure 2.10) and thermo-control interface (Figure 2.6). However, standard operation between batch and continuous mode includes the swapping of the glass SEOP cell, which is typically beyond the realms of established commercial systems and requires trained personnel. Finally, the use of custom gas mixes, often produced in-house, can introduce contamination to the cell. This compounds the need for gas purification by the gas purifiers (Figure 2.10) and rubidium trap (Figure 2.11).

2.6 Conclusion

This work describes the components of a pre-clinical polariser, designed to hyperpolarise the noble gas isotopes xenon-129, krypton-83 and xenon-131. Key components include the dual heat-gun-fed oven, intuitive gas manifolds and interchangeable laser mountings. The primary benefit of this polariser as a pre-clinical research system is its versatility. This includes the capability to hyperpolarise these noble gas nuclei with different buffer gases including nitrogen, helium and hydrogen (discussed further in Chapter 3).

2.7 References

- Meersmann, T. & Brunner, E. Hyperpolarized xenon-129 magnetic resonance: concepts, production, techniques and applications eng (eds Meersmann, T. & Brunner, E.) Cambridge, 2015.
- Stupic, K. F., Cleveland, Z. I., Pavlovskaya, G. E. & Meersmann,
 T. Hyperpolarized 131Xe NMR spectroscopy. *Journal of Magnetic Resonance* 208, 58–69. doi:10.1016/j.jmr.2010.10.004 (2011).
- Nikolaou, P., Goodson, B. M. & Chekmenev, E. Y. NMR hyperpolarization techniques for biomedicine. *Chemistry A European Journal* 21, 3156–3166. doi:10.1002/CHEM.201405253, (Feb. 2015).
- Krjukov, E. V., O'Neill, J. D. & Owers-Bradley, J. R. Brute Force Polarization of 129Xe. Journal of Low Temperature Physics 140, 397–408. doi:10.1007/s10909-005-7323-4 (2005).
- Hirsch, M. L., Kalechofsky, N., Belzer, A., Rosay, M. & Kempf, J. G. Brute-Force Hyperpolarization for NMR and MRI. *Journal of the American Chemical Society* 137, 8428–8434. doi:10.1021/JACS. 5B01252, (July 2015).
- Abragam, A. & Goldman, M. Principles of dynamic nuclear polarisation. Reports on Progress in Physics 41, 395. doi:10.1088/0034-4885/41/3/002 (Mar. 1978).
- Ardenkjaer-Larsen, J. H. et al. Increase in signal-to-noise ratio of ¿ 10,000 times in liquid-state NMR. Proceedings of the National Academy of Sciences 100, 10158–10163. doi:10.1073/pnas.1733835100 (2003).

- Prisner, T. & Köckenberger, W. Dynamic nuclear polarization: New experimental and methodology approaches and applications in physics, chemistry, biology and medicine. *Applied Magnetic Resonance* 34, 213–218. doi:10.1007/S00723-008-0137-1/METRICS (Aug. 2008).
- Hausser, K. H. & Stehlik, D. Dynamic Nuclear Polarization in Liquids. Advances in Magnetic and Optical Resonance 3, 79–139. doi:10. 1016/B978-1-4832-3116-7.50010-2 (Jan. 1968).
- Borghini, M., de Boer, W. & Morimoto, K. Nuclear dynamic polarization by resolved solid-state effect and thermal mixing with an electron spin-spin interaction reservoir. *Physics Letters A* 48, 244–246. doi:10.1016/0375-9601(74)90487-3 (July 1974).
- 11. Overhauser, A. W. Polarization of Nuclei in Metals. *Physical Review* **92**, 411. doi:10.1103/PhysRev.92.411 (Oct. 1953).
- Jeffries, C. D. Dynamic Orientation of Nuclei by Forbidden Transitions in Paramagnetic Resonance. *Physical Review* 117, 1056. doi:10.1103/PhysRev.117.1056 (Feb. 1960).
- Andrew, E. R. Spin Temperature and Nuclear Magnetic Resonance in Solids. *Physics Bulletin* 22, 161. doi:10.1088/0031-9112/22/3/030 (Mar. 1971).
- Ardenkjaer-Larsen, J. H. On the present and future of dissolution-DNP. Journal of Magnetic Resonance 264, 3–12. doi:10.1016/J. JMR.2016.01.015 (Mar. 2016).
- 15. Bowers, C. R. & Weitekamp, D. P. Parahydrogen and Synthesis Allow Dramatically Enhanced Nuclear Alignment. *Journal of the American Chemical Society* **109**, 5541–5542. doi:10.1021/JA00252A049/ASSET/JA00252A049.FP.PNG{_}V03 (Sept. 1987).

- 16. Bowers, C. R. Sensitivity Enhancement Utilizing Parahydrogen. *eMa-qRes* **2007.** doi:10.1002/9780470034590.EMRSTM0489 (2007).
- 17. Pravica, M. G. & Weitekamp, D. P. Net NMR alignment by adiabatic transport of parahydrogen addition products to high magnetic field. *Chemical Physics Letters* **145**, 255–258. doi:10.1016/0009-2614(88)80002-2 (Apr. 1988).
- Green, R. A., Adams, R. W., Duckett, S. B., Mewis, R. E., Williamson,
 D. C. & Green, G. G. The theory and practice of hyperpolarization in magnetic resonance using parahydrogen. *Progress in Nuclear Magnetic Resonance Spectroscopy* 67, 1–48. doi:10.1016/J.PNMRS.
 2012.03.001 (Nov. 2012).
- 19. Eisenschmid, T. C., Kirss, R. U., Deutsch, P. P., Hommeltoft, S. I., Eisenberg, R., Bargon, J., Lawler, R. G. & Balch, A. L. Para Hydrogen Induced Polarization in Hydrogenation Reactions. *Journal of the American Chemical Society* 109, 8089–8091. doi:10.1021/JA00260A026/ASSET/JA00260A026.FP.PNG{_}V03 (Dec. 1987).
- Adams, R. W. et al. Reversible interactions with para-hydrogen enhance NMR sensitivity by polarization transfer. Science 323, 1708–1711. doi:10.1126/SCIENCE.1168877, (Mar. 2009).
- 21. Appleby, K. M., Mewis, R. E., Olaru, A. M., Green, G. G., Fairlamb, I. J. & Duckett, S. B. Investigating pyridazine and phthalazine exchange in a series of iridium complexes in order to define their role in the catalytic transfer of magnetisation from para-hydrogen. *Chemical Science* 6, 3981–3993. doi:10.1039/C5SC00756A (June 2015).
- 22. Zeng, H. et al. Optimization of SABRE for polarization of the tuberculosis drugs pyrazinamide and isoniazid. Journal of Magnetic Resonance 237, 73–78. doi:10.1016/J.JMR.2013.09.012 (Dec. 2013).

- Walker, T. G. & Happer, W. Spin-exchange optical pumping of noble-gas nuclei. Reviews of Modern Physics 69, 629–642. doi:10. 1103/revmodphys.69.629 (1997).
- 24. Happer, W. Optical Pumping. Reviews of Modern Physics 44, 169–249. doi:10.1103/revmodphys.44.169 (1972).
- 25. Happer, W., Miron, E., Schaefer, S., Schreiber, D., Van Wijngaarden, W. A. & Zeng, X. Polarization of the nuclear spins of noble-gas atoms by spin exchange with optically pumped alkali-metal atoms. *Physical Review A* 29, 3092. doi:10.1103/PhysRevA.29.3092 (June 1984).
- Kastler, A. Optical Methods of Atomic Orientation and of Magnetic Resonance*. Journal of the Optical Society of America 47, 460–465. doi:10.1364/JOSA.47.000460 (1957).
- Bouchiat, M. A., Carver, T. R. & Varnum, C. M. Nuclear Polarization in He3 Gas Induced by Optical Pumping and Dipolar Exchange.
 Physical Review Letters 5, 373–375. doi:10.1103/PhysRevLett.5.

 373 (1960).
- 28. Katz, O. & Firstenberg, O. Transverse optical pumping of spin states.

 *Communications Physics 2. doi:10.1038/s42005-019-0170-4

 (2019).
- 29. Wagshul, M. E. & Chupp, T. E. Optical pumping of high-density Rb with a broadband dye laser and GaAlAs diode laser arrays: Application to He 3 polarization. 40, 4447–4454. doi:10.1103/physreva. 40.4447 (1989).
- Ruset, I. C., Ketel, S. & Hersman, F. W. Optical Pumping System Design for Large Production of Hyperpolarized 129Xe. *Physical Review Letters* 96, 53002. doi:10.1103/PhysRevLett.96.053002 (2006).

- 31. Zook, A. L., Adhyaru, B. B. & Bowers, C. R. High capacity production of ¿65% spin polarized xenon-129 for NMR spectroscopy and imaging. eng. *Journal of magnetic resonance (San Diego, Calif. :* 1997) 159, 175–182. doi:10.1016/s1090-7807(02)00030-7 (2002).
- Zhu, H., Ruset, I. C. & Hersman, F. W. Spectrally narrowed externalcavity high-power stack of laser diode arrays. *Optics Letters* 30, 1342. doi:10.1364/ol.30.001342 (2005).
- 33. Parnell, S. R., Deppe, M. H., Parra-Robles, J. & Wild, J. M. Enhancement of X129e polarization by off-resonant spin exchange optical pumping. *Journal of Applied Physics* **108**, 64908. doi:10.1063/1.3478707 (2010).
- 34. Driehuys, B., Cates, G. D., Miron, E., Sauer, K., Walter, D. K. & Happer, W. High-volume production of laser-polarized 129Xe. *Applied Physics Letters* **69**, 1668–1670. doi:10.1063/1.117022 (Sept. 1996).
- 35. Appelt, S., Ünlü, T., Zilles, K., Shah, N. J., Baer-Lang, S. & Halling, H. Experimental studies of rubidium absolute polarization at high temperatures. *Applied Physics Letters* **75**, 427–429. doi:10.1063/1.124397 (July 1999).
- 36. Six, J. S., Hughes-Riley, T., Stupic, K. F., Pavlovskaya, G. E. & Meersmann, T. Pathway to Cryogen Free Production of Hyperpolarized Krypton-83 and Xenon-129. *PLoS ONE* 7, e49927. doi:10.1371/journal.pone.0049927 (2012).
- 37. Nelson, I. A. & Walker, T. G. Rb-Xe spin relaxation in dilute Xe mixtures. *Physical Review A* **65.** doi:10.1103/physreva.65.012712 (2001).

- 38. Brunner, E. Enhancement of surface and biological magnetic resonance using laser-polarized noble gases. *Concepts in Magnetic Resonance* 11, 313–335. doi:10.1002/(sici)1099-0534(1999)11: 5<313::Aid-cmr4>3.0.Co;2-c (1999).
- 39. Skinner, J. Optimisation of xenon-rich stopped-flow spin-exchange optical pumping for functional lung imaging (July 2017).
- 40. Hughes-Riley, T., Six, J. S., Lilburn, D. M. L., Stupic, K. F., Dorkes, A. C., Shaw, D. E., Pavlovskaya, G. E. & Meersmann, T. Cryogenics free production of hyperpolarized 129Xe and 83Kr for biomedical MRI applications. *Journal of Magnetic Resonance* 237, 23–33. doi:10.1016/j.jmr.2013.09.008 (2013).
- Hunt, E. R. & Carr, H. Y. Nuclear Magnetic Resonance of Xe129 in Natural Xenon. *Physical Review* 130, 2302. doi:10.1103/PhysRev. 130.2302 (June 1963).
- Cleveland, Z. I. & Meersmann, T. Binary-collision-induced longitudinal relaxation in gas-phase 83Kr. *Journal of Chemical Physics* 129. doi:10.1063/1.3029663/187814 (Dec. 2008).
- Chann, B., Nelson, I. A., Anderson, L. W., Driehuys, B. & Walker,
 T. G. 129Xe-Xe Molecular Spin Relaxation. *Physical Review Letters* 88, 4. doi:10.1103/PHYSREVLETT.88.113201, (2002).
- Cleveland, Z. I. & Meersmann, T. Density-independent contributions to longitudinal relaxation in 83Kr. ChemPhysChem 9, 1375–1379. doi:10.1002/CPHC.200800134 (July 2008).
- Gordon, R. G. Kinetic Theory of Nuclear Spin Relaxation in Gases.
 The Journal of Chemical Physics 44, 228–234. doi:10.1063/1.
 1726451 (Jan. 1966).

- 46. Cates, G. D., Schaefer, S. R. & Happer, W. Relaxation of spins due to field inhomogeneities in gaseous samples at low magnetic fields and low pressures. *Physical Review A* 37, 2877. doi:10.1103/PhysRevA. 37.2877 (Apr. 1988).
- 47. Jameson, C. J., Jameson, A. K. & Hwang, J. K. Nuclear spin relaxation by intermolecular magnetic dipole coupling in the gas phase. 129Xe in oxygen. *The Journal of Chemical Physics* **89**, 4074–4081. doi:10.1063/1.454842 (Oct. 1988).
- 48. Saam, B., Happer, W. & Middleton, H. Nuclear relaxation of He3 in the presence of O2. *Physical Review A* 52, 862–865. doi:10.1103/PHYSREVA.52.862 (1995).
- 49. Driehuys, B., Cates, G. D. & Happer, W. Surface Relaxation Mechanisms of Laser-Polarized 129Xe. *Physical Review Letters* **74**, 4943. doi:10.1103/PhysRevLett.74.4943 (June 1995).
- Fitzsimmons, W. A., Tankersley, L. L. & Walters, G. K. Nature of Surface-Induced Nuclear-Spin Relaxation of Gaseous He3. *Physical Review* 179, 156. doi:10.1103/PhysRev.179.156 (Mar. 1969).
- 51. Breeze, S. R., Lang, S., Moudrakovski, I., Ratcliffe, C. I., Ripmeester, J. A., Simard, B. & Santyr, G. Coatings for optical pumping cells and extending the lifetime of hyperpolarized xenon. *Journal of Applied Physics* 86, 4040–4042. doi:10.1063/1.371325 (Oct. 1999).
- 52. Breeze, S. R., Lang, S., Moudrakovski, I., Ratcliffe, C. I., Ripmeester, J. A., Santyr, G., Simard, B. & Zuger, I. Coatings for optical pumping cells and short-term storage of hyperpolarized xenon. *Journal of Applied Physics* 87, 8013–8017. doi:10.1063/1.373489 (June 2000).
- 53. Cleveland, Z. I., Stupic, K. F., Pavlovskaya, G. E., Repine, J. E., Wooten, J. B. & Meersmann, T. Hyperpolarized83Kr and129Xe NMR

- Relaxation Measurements of Hydrated Surfaces: Implications for Materials Science and Pulmonary Diagnostics. *Journal of the American Chemical Society* **129**, 1784–1792. doi:10.1021/ja065994t (2007).
- 54. Stupic, K. F., Cleveland, Z. I., Pavlovskaya, G. E. & Meersmann, T. Quadrupolar relaxation of hyperpolarized krypton-83 as a probe for surfaces. *Solid State Nuclear Magnetic Resonance* **29**, 79–84. doi:10. 1016/j.ssnmr.2005.08.008 (2006).
- 55. Meersmann, T. & Pavlovskaya, G. E. Surface Quadrupolar Relaxation (SQUARE) Contrast in Pulmonary MRI With Hyperpolarized Kr-83 English (eds Albert, M. S. & Hane, F. T.) 293–304 (Academic Press Ltd-Elsevier Science Ltd, London, 2017).
- 56. Stupic, K. F., Elkins, N. D., Pavlovskaya, G. E., Repine, J. E. & Meersmann, T. Effects of pulmonary inhalation on hyperpolarized krypton-83 magnetic resonanceT1relaxation. *Physics in Medicine and Biology* 56, 3731–3748. doi:10.1088/0031-9155/56/13/001 (2011).
- 57. Abragam, A. & Hebel, L. C. The Principles of Nuclear Magnetism.

 American Journal of Physics 29, 860–861. doi:10.1119/1.1937646

 (Dec. 1961).
- 58. Molway, M. J. et al. Dramatic improvement in the "Bulk" hyperpolarization of 131Xe via spin exchange optical pumping probed using in situ low-field NMR. Journal of Magnetic Resonance **354**, 107521. doi:10.1016/J.JMR.2023.107521 (Sept. 2023).
- 59. Irwin, R. Development of Hyperpolarised Xenon 129 Production and MRI for a Clinical Setting PhD thesis (University of Nottingham, 2020).

- 60. Norquay, G., Parnell, S. R., Xu, X., Parra-Robles, J. & Wild, J. M. Optimized production of hyperpolarized 129Xe at 2 bars for in vivo lung magnetic resonance imaging. *Journal of Applied Physics* 113. doi:10.1063/1.4776763/818813 (Jan. 2013).
- 61. Stewart, N. J., Norquay, G., Griffiths, P. D. & Wild, J. M. Feasibility of human lung ventilation imaging using highly polarized naturally abundant xenon and optimized three-dimensional steady-state free precession. *Magnetic resonance in medicine* **74**, 346–352. doi:10.1002/MRM.25732 (Aug. 2015).
- 62. Kuzma, N. N., Patton, B., Raman, K. & Happer, W. Fast Nuclear Spin Relaxation in Hyperpolarized Solid 129Xe. *Physical Review Letters* 88, 147602. doi:10.1103/PhysRevLett.88.147602 (Mar. 2002).
- 63. Nikolaou, P. et al. Near-unity nuclear polarization with an open-source 129Xe hyperpolarizer for NMR and MRI. Proceedings of the National Academy of Sciences 110, 14150–14155. doi:10.1073/PNAS.1306586110 (Aug. 2013).
- 64. Birchall, J. R. et al. XeUS: A second-generation automated opensource batch-mode clinical-scale hyperpolarizer. *Journal of Magnetic Resonance* **319**, 106813. doi:10.1016/J.JMR.2020.106813 (Oct. 2020).
- 65. Birchall, J. R., Chowdhury, M. R. H., Nikolaou, P., Chekmenev, Y. A., Shcherbakov, A., Barlow, M. J., Goodson, B. M. & Chekmenev, E. Y. Pilot Quality-Assurance Study of a Third-Generation Batch-Mode Clinical-Scale Automated Xenon-129 Hyperpolarizer. *Molecules* 27. doi:10.3390/MOLECULES27041327, (Feb. 2022).
- 66. Hersman, F. W. et al. Large Production System for Hyperpolarized 129Xe for Human Lung Imaging Studies. Academic Radiology 15, 683-692. doi:10.1016/j.acra.2007.09.020 (June 2008).

- 67. Norquay, G., Collier, G. J., Rao, M., Stewart, N. J. & Wild, J. M. Xe129-Rb Spin-Exchange Optical Pumping with High Photon Efficiency. *Physical Review Letters* **121**, 153201. doi:10.1103/PhysRevLett. 121.153201 (Oct. 2018).
- 68. Norquay, G., Collier, G. J., Rodgers, O. I., Gill, A. B., Screaton, N. J. & Wild, J. Standalone portable xenon-129 hyperpolariser for multicentre clinical magnetic resonance imaging of the lungs. *The British Journal of Radiology* 95, 20210872. doi:10.1259/BJR.20210872 (2022).

Chapter 3

Reactive Purification of

Post-SEOP Hyperpolarised

Noble Gases using a Hydrogen

Buffer Gas

Abstract

Spin-Exchange Optical Pumping (SEOP) requires buffer gases to enhance optical pumping efficiency, but these dilute the final hyperpolarised gas mixture, resulting in reduced signal intensity for MRI applications. While conventional phase-change separation methods exist for xenon purification, they are limited to batch processing and can result in significant polarization losses.

This work explores the use of molecular hydrogen as a buffer gas, which can be eliminated through reaction with oxygen to produce water vapour that is subsequently removed via condensation. The combustion system was designed and constructed, featuring a 4 L glass-walled reaction vessel with integrated safety systems, automated control software, and dual ignition mechanisms (platinum catalyst and high-voltage electrodes) and pneumatic recompression capabilities for gas delivery.

Two operational modes were developed and tested: batch mode for rapid purification of single gas deliveries, and semi-continuous mode for processing multiple smaller deliveries. Preliminary results demonstrate that the hyperpolarised spin states of xenon-129 and krypton-83 survive the combustion process at low magnetic fields with minimal polarization loss due to the purification process. Batch mode experiments achieved noble gas concentrations of $71\pm2\%$ for xenon and $81\pm2\%$ for krypton.

This reactive purification approach offers advantages over traditional purification methods, including cryogen-free operation and applicability to both xenon and krypton. The work establishes the foundation for a new paradigm in hyperpolarised gas purification that could facilitate wider clinical deployment of hyperpolarised noble gas MRI.

Contents

3.1	Introd	Introduction	
	3.1.1	Buffer Gases for SEOP 67	
	3.1.2	Apparent Polarisation 69	
3.2	2 Hyper	polarised Gas Purification	
3.5	3 Hydro	Hydrogen as a Buffer Gas	
3.4	4 Aims		
3.5	5 The C	The Combustion system	
	3.5.1	Hardware	
	3.5.2	Polariser Control	
	3.5.3	Control Software	
	3.5.4	Safety	
	3.5.5	Operational Modes	
	3.5.6	Gas Mix estimation	
3.6	6 Result	s	
	3.6.1	Batch Mode	
	3.6.2	Semi-Continuous Mode	
3.7	7 Discus	Discussion	
3.8	8 Conclu	Conclusion	
3.9) Refere	References	

Declaration of Contributions

- **Arthur Harrison** Design, procurement, manufacture and testing of all components and processed described in this chapter, unless otherwise stated.
- Max Filkins Assisted in testing of components and processes described in this chapter and the manufacture of pneumatic components.
- Xinpei Wang Early testing of the combustion system core pressure vessel and developed early iterations of the pneumatic manifold (not discussed in this work).
- Thomas Meersmann Design of the core combustion pressure vessel.
- Alan Dokes (Workshop) Manufacture of the core combustion pressure vessel.
- Stuart Salter (Workshop) Assisted in the manufacture of the system housing, including the operation of all heavy power tools.

3.1 Introduction

3.1.1 Buffer Gases for SEOP

In order to obtain the highest noble gas polarisations, buffer gases are commonly introduced to SEOP gas mixtures. These gases serve two primary purposes, to prevent depolarisation via radiation trapping and to increase the efficiency of laser photon absorption by the alkali metal vapour[1, 2].

3.1.1.1 Radiation trapping

Radiation trapping stems from the reabsorption of the photons produced due to the optical pumping process; due to the relaxation of the alkali metal electron to the ground state sublevels[3]. These non-circularly polarised photons can cause unwanted transitions within the alkali metal electron energy levels, which themselves can spontaneously emit photons due to relaxation. Therefore, if left unchecked, this phenomenon will lead to the rapid depolarisation of the alkali metal due to the depletion of the $S_{1/2}$, m = +1/2 state. A quenching agent is required within the gas mixture to quench this process, by absorbing the free energy. Without the introduction of a sufficient partial pressure of this gas, rubidium fluorescence can be observed as a violet glow[4].

Typically, molecular nitrogen is introduced to the SEOP gas mixture as a quenching agent to prevent significant depolarisation due to radiation trapping. Its utility owes to its triple covalent bond, which not only grants a high quenching cross section[1, 5], but can also absorb lots of energy into its vibrational energy levels[6]. Furthermore, this molecule has a low rubidium

spin destruction cross section $(\kappa_{sd}^{N_2})$ and will not react chemically with the alkali metal vapour. The effectiveness of this mechanism is dependent on the partial pressure of nitrogen, $[N_2]$. This is described by the branching ratio ω_{γ} , the probability that radiation trapping occurs [2, 5, 7]:

$$\omega_{\gamma} = \frac{3}{3 + 835[N_2]} \tag{3.1}$$

where $[N_2]$ is in units of amagat. At least 5 kPa of nitrogen is generally required within an SEOP gas mixture to effectively suppress the depolarisation associated with radiation trapping. [1, 2] Although this increases with the rubidium vapour density. Therefore, greater quantities are required for krypton SEOP due to the associated higher temperatures [6].

3.1.1.2 Pressure broadening the D1 Transition

Buffer gases are employed to increase optical pumping efficiency using diode array lasers. Pure rubidium vapour requires radiation of 794.7 nm to stimulate the D1 transition.[3] However, laser systems possess a broadband spectral output. Therefore, without significant line narrowing, a portion of the laser power will be wasted, reducing the efficiency of the optical pumping process. The absorption profile linewidth can be collisionally broadened via the introduction of a buffer gas.[8] This allows the alkali metal vapour to absorb a greater range of wavelengths to better suit the spectral profile of the laser. Molecular nitrogen can be used for this purpose; however, the greatest efficiency can be attained using helium as it has the lowest rubidium spin destruction cross section ($\kappa s d^{He}$). This allows for the introduction of the many hundreds of kPa required for line broadening, without significant penalty to the attainable rubidium polarisation. To date, gas

mixtures comprising $\rm Xe/N_2$ or $\rm Xe/N_2/He$ have found the most widespread success.

3.1.2 Apparent Polarisation

The previous sections in this work have discussed at length the need for a buffer gas within the SEOP gas mixture. However, the addition of these buffer gases will inevitably reduce the concentration of hyperpolarised noble gas within the resultant mix. Without further processing of the hyperpolarised gas mix, this will reduce the observable signal intensities within NM-R/MRI applications, due to the reduced number density of the polarised species. This phenomenon is described by the apparent polarisation, P_{app} , the effective polarisation observed during experimentation when accounting for the dilution of the hyperpolarised nuclei[9]:

$$P_{app} = P_{NG} \times \frac{[NG]}{\sum_{i} [M_i]} \tag{3.2}$$

where P_{NG} is the polarisation of the noble gas, [NG] is the concentration of the noble gas, i represents each species in the gas phase (excluding the alkali metal vapour), and $[M_i]$ represents the concentration of each of these constituent parts. For example, if a fictitious noble gas ensemble is polarised to 100%, but this is achieved with a 1% noble gas concentration, then the resulting apparent polarisation would be 1%. Considering a more scientifically relevant scenario: it is possible to polarise 5% xenon mixtures to polarisations of 60%[9], however this will result in an apparent polarisation of only 3%.

Given this effect of reducing the hyperpolarised noble gas concentration,

one may assume that SEOP should be conducted with maximal concentrations of this target nuclei. After all, greater partial pressures of this gas will contribute to the pressure broadening of the D1 transition, and only a small concentration of nitrogen is required to quench the radiation trapping. However, increasing the concentration of the target nuclei will reduce the attainable polarisation of the ensemble (see Chapter 2). Therefore, one must balance these effects in order to attain the optimal observable signals for MR experiments. This is demonstrated in work by Hughes-Riley et al.[10], Figure 3.1.

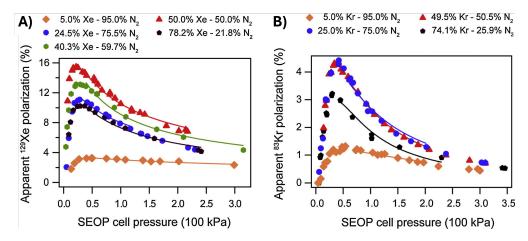


Figure 3.1: Demonstration of the effect of SEOP cell pressure on the apparent polarisation of xenon-129 (A) and krypton-83 (B) for a range of noble gas concentrations[10].

3.2 Hyperpolarised Gas Purification

To attain the greatest signal in NMR experiments, it is beneficial to purify the hyperpolarised gas species, thereby eliminating the penalty from gas dilution (at full concentration). However, to be useful, the noble gas hyperpolarised spin state must survive the purification process.

3.2.0.1 Phase Change Separation

Following SEOP, hyperpolarised xenon-129 is typically purified via a phase separation method. For Xe/N₂ or Xe/N₂/He gas mixes, this process utilises the xenon's significantly higher boiling point compared to that of the other components. The post-SEOP gases are passed through a cryogen trap, generally a high-surface-to-volume glass apparatus submerged in a cryogen bath. The hyperpolarised gas mixture is continuously flowed through this apparatus, accumulating condensed xenon. Upon collection of a sufficient quantity of xenon, the hyperpolarised medium is thawed to return it to a pure gaseous state. This separation method has been employed to produce concentrated hyperpolarised xenon with apparent polarisations of 50%.[8] However, as it involves the storage of solid xenon, this process is inherently batch mode. It is also difficult to automate and thus less suited to widespread clinical applications [6]. Although the longitudinal relaxation time of solid xenon-129 within a magnetic field is reasonably long[11] polarisation is lost during the accumulation phase. Within simple setups, often up to 1/3 of the polarisation is forfeited. [6, 11, 12] Furthermore, this separation method is not feasible for krypton SEOP, due to fast quadrupolar relaxations in the condensed state.[13, 14]

Conventional post-SEOP cryogen traps employ a freeze-thaw method for gas separation where concentrated hyperpolarised xenon is accumulated via condensation. Imai et al.[15] have presented an alternative approach to this process, where instead the buffer gas is condensed, leaving just the hyperpolarised xenon in the gas stream. This method circumvents many of the disadvantages associated with conventional cryogen traps and is capable of continuously supplying concentrated hyperpolarised xenon gas. Isobutene was employed as the buffer gas for this method, proving

effective as a quenching agent due to its double bond. This compound has a high boiling point of 266.2 K, over 100 K greater than that of xenon. Therefore, by passing the post-SEOP gas mixture through a tubular spiral cryogen trap, liquid isobutene is accumulated, leaving a purified gas flow of hyperpolarised xenon. The greatest apparent polarisation obtained using this method was 11.9%.

3.3 Hydrogen as a Buffer Gas

Previously, molecular hydrogen has been largely overlooked as a quenching agent for SEOP. Small molecules without a double or triple bond generally possess very low IR quenching cross sections, making them unsuitable as a buffer gas. This is evident for methane; as SEOP using this molecule as a quenching agent yields very low spin polarisations [16, 17]. Hydrogen molecules appear to be a rare exception to this rule, although they still possess a quenching cross section an order of magnitude lower than that of nitrogen[5]. Hydrogen's ability to quench the destructive rubidium fluorescence has been previously investigated by Rogers et al.[17] The D2 fluorescence was used to give quantitative insight into the radiation quenching efficiency of nitrogen, hydrogen and methane at a range of pressures. As evident in Figure 3.2A, above partial pressures of 40-60 kPa, hydrogen acts as a remarkably efficient quenching agent. Little fluorescence was detected within this regime, even within the high rubidium density environments associated with krypton SEOP, with efficiency comparable to the well-established quenching agent, nitrogen. However, below 40 kPa considerable radiation was detected, signifying that the hydrogen failed to prevent significant radiation trapping. This is in agreement with previous measurements of the quenching cross section. For effective suppression of radiation

trapping comparable to nitrogen, approximately ten times the quantity of hydrogen is required. Rogers et al.[17] conducted SEOP experiments using highly concentrated hydrogen as the sole buffer gas, attaining polarisations similar to those expected using pure nitrogen, albeit not at lower pressures. Polarisation curves comparing the effectiveness of each buffer gas are presented in Figure 3.2B. The highest spin polarisations attained using hydrogen as the sole buffer gas were 63% and 29% for xenon-129 and krypton-83 respectively, which was comparable to that of nitrogen under similar conditions. Therefore, the utilisation of hydrogen as a buffer gas during SEOP has the potential to attain spin polarisations comparable to nitrogen.

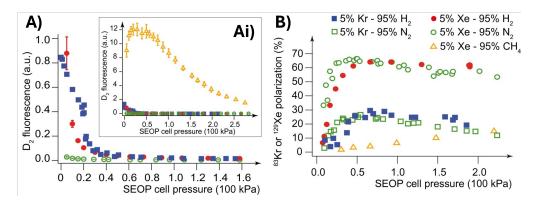


Figure 3.2: Previous work by Rodgers et al.[17] demonstrating the feasibility of molecular hydrogen as a buffer gas for SEOP compared to nitrogen and methane. The D_2 florescence (A) was measured for each buffer gas in a 95% concentration. The polarisation attained (B) from SEOP using 5% noble gas, 95% buffer gas mixtures.

Whilst molecular nitrogen is inert to the rubidium vapour used for SEOP, hydrogen will react. Optical pumping of rubidium in the presence of hydrogen can lead to on-resonance formation of rubidium hydrides (RbH)[18]. Indeed, this phenomenon was observed as white depositions on the cell walls during SEOP, although to a lesser extent on the laser entry window[17]. However, after approximately 30 hours of SEOP, no significant reduction in the attained polarisation was observed, suggesting that depo-

sitions of this kind have little effect on the efficiency of the optical pumping process. Work by Rohrbaugh et al.[19] demonstrated that T1 reductions of xenon-129 under these conditions (field strength, cell profile and temperature) would have a negligible effect on the SEOP process. Previously, relaxation times of the quadrupolar isotope xenon-131 have been shown to increase with RbH surface build-up[20]. Therefore, depositions of this kind may reduce krypton depolarisation due to surface quadrupolar interactions, yielding higher polarisations[17].

3.3.0.1 Reactive Separation

The use of molecular hydrogen as a buffer gas has one great advantage compared to nitrogen: it can be eliminated via an oxidative reaction process following SEOP. Hydrogen reacts exothermically with oxygen to produce water vapour and yields approximately 286 kJ per mole of hydrogen:

$$2H_2 + O_2 \to 2H_2O$$
 (3.3)

Hydrogen reacts readily with oxygen at concentrations of between 4% and 94% at 1 atmosphere pressure[21], forming a colourless flame. At lower pressures, the flammability range is narrowed: raising the Lower Flammability Limit (LFL) and lowering the Upper Flammability Limit (UFL). Below approximately 0.076 atmospheres, no hydrogen-oxygen mixture will combust[22]. The addition of a third species also reduces flammability, as depicted in the flammability diagram for hydrogen-oxygen-nitrogen mixtures[21] in Figure 3.3. The auto-ignition temperature of hydrogen-oxygen mixtures is approximately 839 K[22] and this temperature is largely independent of pressure[23]. However, combustion can be initiated at much

lower temperatures via an external ignition source such as a spark or a flame. The resulting exothermic reaction is self sustaining as long as sufficient hydrogen and oxidiser are present. The reaction of hydrogen-oxygen mixtures is comparatively easy to initiate through external ignition, requiring a very low energy input of approximately 19 μ J[24]. This reaction can also be catalysed using a metal catalyst, particularly platinum, palladium and rhodium[25]. The heat generated by this adsorptive reaction process can ignite hydrogen and oxygen mixtures when the auto-ignition temperature is reached.

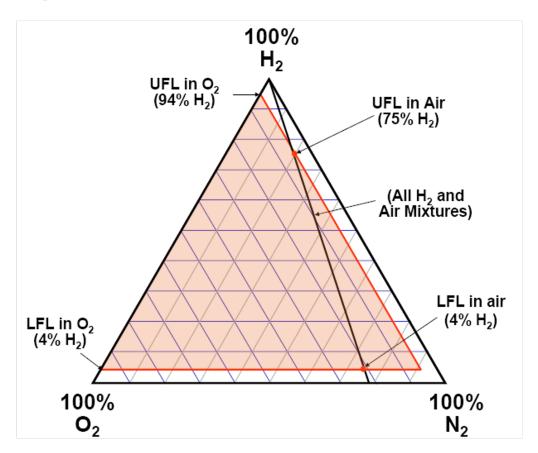


Figure 3.3: Tertiary flammability diagram of hydrogen-oxygen-nitrogen mixtures. Shaded orange regions represent the flammability range, between the LFL and UFL. The black line represents the oxygen-nitrogen mix of air. Adapted from work by Ring et al.[21]

Following the work by Rogers et al[17], demonstrating the feasibility of using molecular hydrogen as a quenching agent for SEOP, a novel and

potentially revolutionary mechanism for the purification of hyperpolarised gases has been introduced. Using hydrogen as the sole component of a buffer gas mixture, permits the reactive purification of the hyperpolarised species. The buffer gas reacts with molecular oxygen to produce water vapour, which can be subsequently removed via condensation at ambient temperatures. Thus the hyperpolarised species is purified by this two-stage process involving reaction followed by phase separation. However, fundamental to this process is the conservation of the noble gas hyperpolarised state during the oxidation process. Significant destruction of spin polarisation during the harsh conditions associated with combustion would render this purification method redundant.

Previous work has demonstrated the combustion resistance of the xenon-129 hyperpolarised state[26]. Although insufficient as a sole quenching agent,[5, 16] methane was introduced to a standard SEOP gas mixture (5% Xe, 85% N₂, 10% CH₄). Employed for continuous flow optical pumping, the hyperpolarised gas mixture was used to fuel a small flame within an upright 9.4T small bore scanner. Significant signal was attained from the post combustion gases, validating the survival of the hyperpolarised spin state within these conditions[26]. The majority of signal loss was attributed to the presence of excess oxygen, destructive to the polarisation due to oxygen's paramagnetic characteristics. Therefore, it was determined that the low spin temperature of the hyperpolarised state is resistant to the high thermal temperatures of the combustion process.

Preliminary investigations of the hyperpolarised spin state lifetime in the presence of the oxidation of hydrogen have been conducted by Rogers et al.[17] Batch mode SEOP was performed using molecular hydrogen as the sole buffer gas and the post SEOP gases were combined with a known volume of oxygen in a small glass walled combustion vessel. A catalyst powder

(5% Pt, Al₂O₃) coating the walls of the vessel was employed to stimulate the oxidation process. Approximately 100 ms after the introduction of the gas to this environment, localised reaction sites were observed on the vessel walls, followed by the primary combustion event 20 ms later. Within 200 ms, the oxidation process appeared largely complete, corresponding to a significant pressure reduction within the vessel. Figure 3.5 depicts a typical combustion process within this vessel.

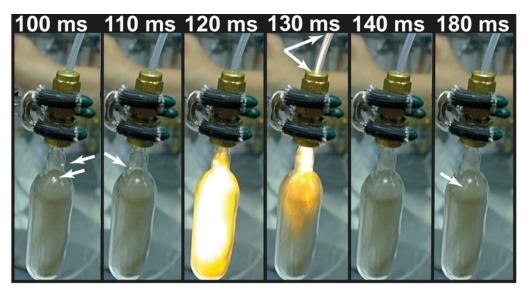


Figure 3.4: Sequential images from the preliminary investigations of the hydrogen combustion resistance of xenon-129 and krypton-83. The glass-walled reaction vessel, coated with the powder Pt-based catalyst, is depicted at sequential stages of the combustion process, including initial localised reaction sites at 100 ms, followed by the rapid ignition of the full vessel at 120 ms after gas delivery. Adapted from work by Rogers et al.[17]

During the combustion process, the vessel was suspended within the bore of a 9.4 T superconducting magnet (Bruker, USA). Using a series of low flip angle pulses, the NMR spectrometer was used to continuously probe the noble gas nuclear spin polarisation during the reaction process. The evolution of the observed signal from xenon-129 and krypton-83 during combustion is presented in Figure 3.5, coupled with the corresponding pressure variation within the reaction vessel. This demonstrates that these combustion events, in the presence of a high magnetic field, had little effect on the

hyperpolarised state of the xenon-129 and krypton-83 nuclei. While the total gas pressure within the vessel dropped significantly, there was negligible impact on the observed signal from the hyperpolarised species, beyond contributions from known depolarisation pathways. For example, the post combustion xenon mixture experienced a greater rate of depolarisation compared to the control experiment due to the presence of excess paramagnetic oxygen after the combustion[17]. Indeed, the effect was amplified for the experiment utilising a greater partial pressure of oxygen. Surprisingly, the krypton experienced a reduced depolarisation rate following the combustion event; a consequence of the competitive co-absorption of water molecules on the glass-walled combustion vessel [17]. Previous work has demonstrated that this effect can reduce krypton quadrupolar surface interactions[27]. Additionally, due to its lower gyromagnetic ratio and smaller electron cloud compared to xenon, krypton polarisation is largely insensitive to excess paramagnetic oxygen[10]. These factors contribute to the extended lifetime of the krypton hyperpolarised state due to the hydrogen oxidation process.

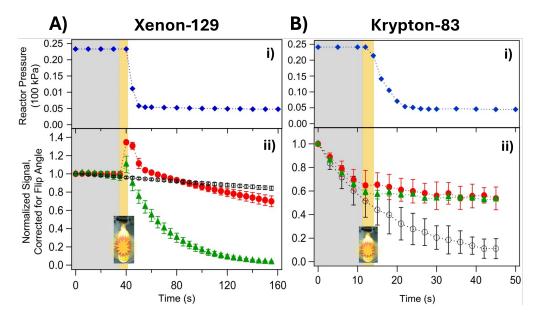


Figure 3.5: Adapted results from work by Rodgers et al.[17], demonstrating the combustion resistance of xenon-129 (A) and krypton-83 (B) hyperpolarised spin state. A typical pressure profile (i) for each combustion event is plotted alongside the flip-angle-corrected NMR signal (ii) from the hyperpolarised species. Three experiments are presented utilising a small (red circles) or a large (green circles) excess of oxygen, or a control experiment with no oxygen (black open circles). The increase in xenon signal associated with the combustion event was due to the addition of extra hyperpolarised gas from the inlet lines. This phenomenon was not observed for krypton due to the fast quadrupolar relaxation within the transport tubes [28].

3.4 Aims

Rodgers et al.[17] previously demonstrated the combustion resistance of the xenon-129 and krypton-83 hyperpolarised states. However, these experiments were completed with a small quantity of gas, within a high field environment (9.4 T), for a near-instantaneous reaction. However, these conditions may not be feasible for the widespread deployment of this technique. Therefore, the combustion resistance of the xenon-129 and krypton-83 hyperpolarised state, should be tested at low-field (<100G), larger gas quantities and for prolonged reactions. Furthermore, an extraction mech-

anism is required for the purified gases to be useful for studies.

Therefore, the aim of this work is to develop a prototype device capable of purifying large quantities of hyperpolarised noble gases at low field, as well as recompressing the resulting mix for delivery to experiments. These processes must take as little time as safely possible in order to reduce the signal lost due to depolarisation. For xenon gas mixtures, mixing with oxygen must be reduced as much as possible due to the fast relaxation associated with contact with paramagnetic molecules [10].

3.5 The Combustion system

3.5.1 Hardware

The combustion system is centred around a vertically-orientated, glass-walled cylinder, approximately 5.65 L in volume. This vessel is 50 cm in length with an inner diameter of 12 cm and a wall thickness of 5.2 mm, and is expected to sustain up to 5 bar of internal pressure, although this hasn't been tested. Each end of the combustion vessel is covered with acrylic end caps, sealed with O-rings and held in place by calibrated springs. Figure 3.6 presents a diagram of the combustion system, coupled with an image of the device.

The combustion system houses a piston which is used to pneumatically recompress the post-oxidation gases. This piston splits the glass vessel into two regions: the reaction vessel above the piston and the compression volume below. The SEOP gas mixture is stored within the combustion chamber during the oxidation and recompression process. Following the completion of the combustion process, the compression chamber is pres-

surised relative to the reaction vessel, extending the piston and compressing the gases above. The piston is later retracted via the use of a vacuum pump. At 15 cm in length, this piston reduces the usable volume of the reaction vessel to approximately 4 L when fully retracted.

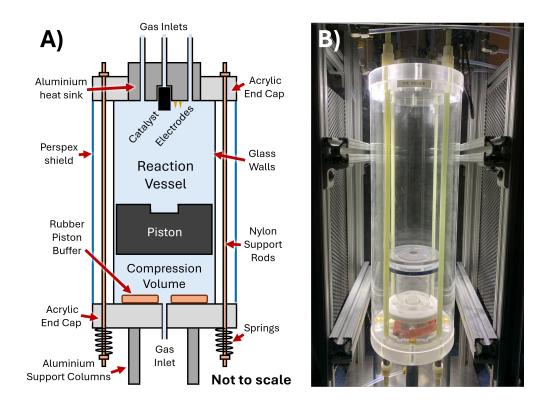


Figure 3.6

3.5.1.1 Ignition Sources

The system incorporates two methods for initiating a hydrogen-oxygen reaction: a Pt-coated catalyst and a pair of high-voltage electrodes. These features are depicted in Figure 3.7. Each of the ignition mechanisms is housed in an aluminium heat sink, located at the top of the reaction vessel. This heat sink is in direct contact with the catalyst, which can heat significantly during the reaction process. Therefore, this housing must be able to endure the high temperatures associated with the reaction process, shielding the rest of the components.

The custom-made (Johnson MattheyTM, UK) catalyst comprises an aluminium oxide (Al_2O_3) foam structure, washed with a 1% wt. covering of platinum. The catalyst can facilitate hydrogen-oxygen reactions on its surface at room temperature. These reactions will heat up the structure and thus the surrounding gas, until the self-ignition point of the hydrogen-oxygen mixture is reached. The catalyst is cylindrical in shape, with a length of 4 cm and a diameter of 2 cm. It is mounted within an oxygen inlet to the reaction vessel, and a hole of radius 1 mm has been bored through the centre of its long axis to allow for easier gas flow through the catalyst. It protrudes about 1 cm into the combustion chamber (the piston has an indent to account for this when fully extended). This ensures that the gas mixture within the chamber can reach the catalyst, even whilst the inlet oxygen is passing through the system.

The electrodes allow for a greater level of control of the ignition process compared to the catalyst. Whereas the catalyst initiates reaction analogous to a glow-plug, the electrodes are analogous to a spark plug. The electrodes are positioned with their tips approximately 2 mm apart. A high voltage is applied across this apparatus, which can create a plasma between the electrodes. This plasma can ignite the hydrogen-oxygen gas mix. The electronic circuit diagram for the electrodes is incorporated into Figure 3.9. A single AA battery, combined with a step-up transformer, is used to create an inter-electrode voltage of approximately 15,000 V.

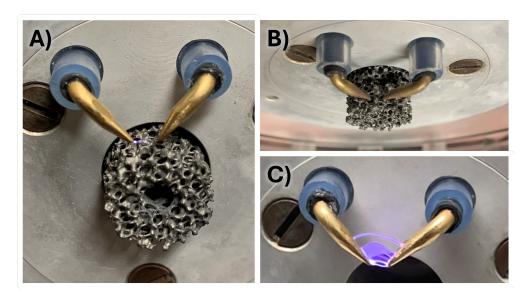


Figure 3.7: The ignition sources of the combustion system, viewed from a below (A) and side (B) profile. Additionally, the plasma created by the activation of the electrodes in air (C) is presented.

3.5.1.2 System Housing

The core system is placed within an aluminium housing. This housing serves multiple purposes, primarily for protection of the operators as detailed in Section 3.5.4. Additionally, this housing is used to mount the pneumatic and electrical control systems, and equipped with wheels to permit the easy transportation of the system. The housing is depicted in Figure 3.8. The sides of the system are removable for system maintenance.

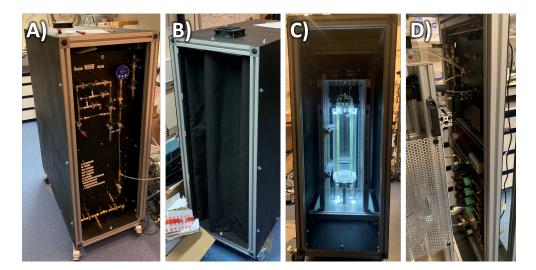


Figure 3.8: The housing of the combustion system, pictured from: the front (A) which holds the manual control pneumatics, the back both with (B) and without (C) the curtain, permitting access to the system, and the reverse of the front panel (D) with the side panel removed, housing the controlled pneumatic and electronic systems.

3.5.1.3 Electronics

The electronics diagram for the system is depicted in Figure 3.9. The system interfaces with a control computer via a Data Acquisition Board (DAQ) (National Instruments, USA). This board controls the solenoid valves using relays and receives analogue voltage signals from the pressure transducers.

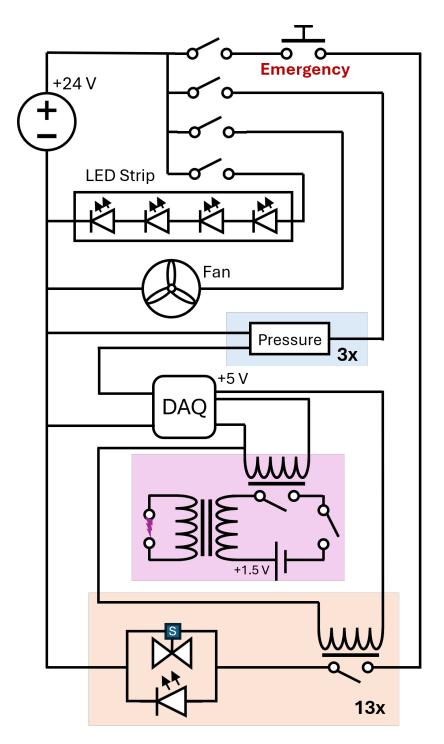


Figure 3.9: The simplified electronic diagram of the combustion system. The circuit for a single solenoid valve is depicted in the **orange box**. The circuit for the high voltage electrodes is depicted in the **magenta box**.

3.5.1.4 Magnetic Field

A static magnetic field is created within the chamber using a four circular coils. These four coils are connected in series, are each approximately 30 cm in diameter and have 112 turns. Magnetic field gradients can cause depolarisation of the noble gas nuclei (Equation 2.21). Therefore, the position of these coils was optimised to yield the lowest field gradients within the reaction vessel. This optimisation was achieved by calculating the resultant magnetic field from the four coils in 3D. Modelling each coil as 15 discreet segments the field is calculated using the Biot-Savart law:

$$d\vec{B} = \frac{\mu_0 I}{4\pi} \frac{d\vec{l} \times \hat{r}}{r^2} \tag{3.4}$$

where μ_0 is the permeability of free space, I is the current in the coils, r is the distance from the coil element, $d\vec{l}$ is the vector of the coil element and \hat{r} is the vector from this element to the point in space where the field is to be calculated. Both the average and maximum field gradients within the reaction vessel region were calculated for respective spacings of the internal and external coils. The resulting 2D optimisation plots are presented in Figure 3.10.

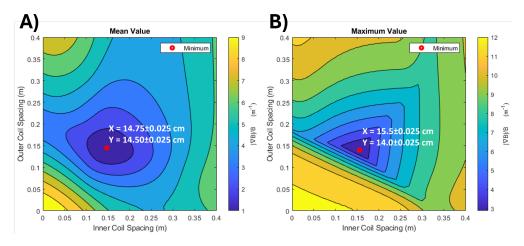


Figure 3.10: Plots to determine the optimal internal and external spacing of the double-coil pair for the lowest mean (A) and maximum (B) relative field gradient within the reaction vessel.

Based on the spacing optimisation simulations, the coils were positioned with an internal spacing of 15.5 cm and an outer spacing of 14 cm. This is depicted in Figure 3.11, with the resultant average field within the reaction vessel calculated for varying currents.

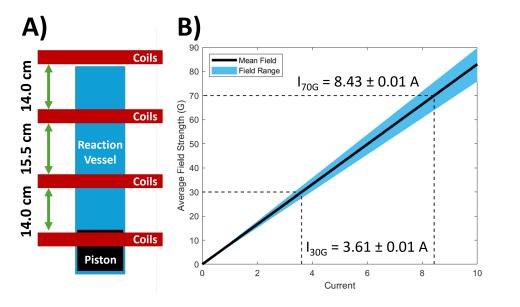


Figure 3.11: Diagram of the optimum coil spacing about the reaction vessel (A) and plot of the resultant mean field within this vessel (B) for different currents.

Using the same field calculations as used in Figure 3.10, the field char-

acteristic for the coil spacing used in the system is presented in Figure 3.12.

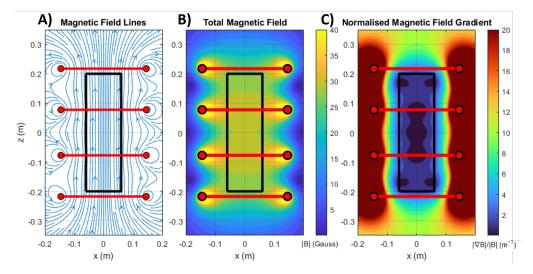


Figure 3.12: Simulations of the static field created by the double coil pair, showing field lines (A), field magnitude (B) and relative field gradient (C). All simulation presented use a current of 3.61 A.

3.5.1.5 Pneumatics

PFA tubing with Swagelok connections is used to transport the gases. Brass connections are used in any areas in contact with hyperpolarised gases, to reduce the depolarisation during gas transfer. Gas flow is controlled by a number of pneumatic valves, each connected to a central control computer. Figure 3.13 describes the complete pipework for the combustion system. The pressures in the reaction vessel and compression volume are monitored independently.

There are four gas inlet lines into the system, three into the reaction vessel and one into the compression volume. Two of the inlets above the piston (including through the centrally housed catalyst) are used to deliver oxygen. These holes are 1 mm (through the catalyst) and 2 mm in radius, designed to vary the rate of oxygen delivery. The oxygen flow through the

catalyst is controlled using a metering valve, calibrated to 0.86 mbar/s into the reaction vessel with the piston fully retracted (approx 3.44 ml/s). The final inlet into the reaction vessel is 1 mm in diameter and is of dual use: first as an inlet for the post-SEOP mixture, then later as an outlet for the dilute hyperpolarised noble gas. The gas line into the compression volume is 2 mm in diameter, and is capable of drawing a vacuum in this volume or pressurising it with atmospheric or compressed air.

Several manual valves are present in the system. These are present in areas where automation is not required such as the set-up of the device. Additionally, these valves are placed as redundancy, should the control software fail. This is the case for the control of the piston and the vacuuming of the reaction vessel.

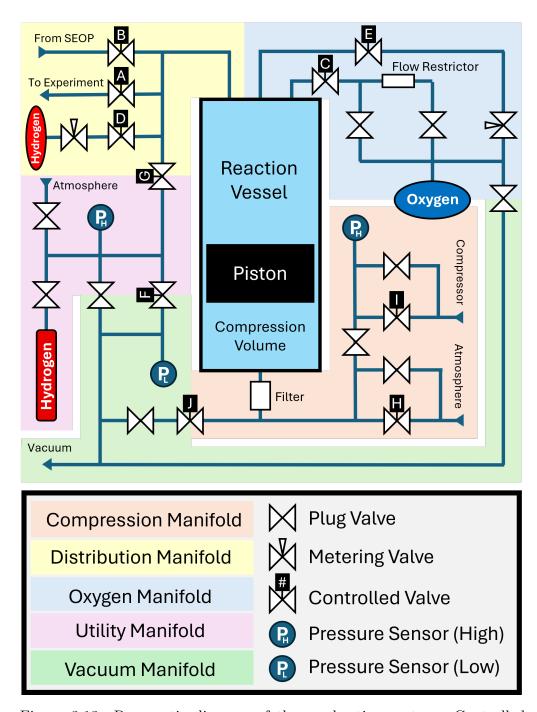


Figure 3.13: Pneumatic diagram of the combustion system. Controlled valve labels refer to descriptions in Table 3.1. All Pneumatic valves are controlled via high pressure gas line, mediated by three-way solenoid valves (not depicted).

The controlled valves in the stem are a combination of pneumatic and solenoid valves, dependent on whether they will come into contact with hyperpolarised gas. These valves are detailed in Table 3.1.

Controlled Valves of the Combustion System.			
#	Name	Type	Normal State
A	Delivery Outlet	Pneumatic	Closed
В	SEOP Inlet	Pneumatic	Closed
С	Fast Oxygen Inlet	Pneumatic	Closed
D	Hydrogen Inlet	Pneumatic	Closed
Е	Slow Oxygen Inlet	Pneumatic	Closed
F	Top Vacuum	Pneumatic	Closed
G	Utility Line	Pneumatic	Open
Н	Bottom Atmosphere	Solenoid	Closed
Ι	Bottom Compressor	Solenoid	Closed
J	Bottom Vacuum	Pneumatic	Open

Table 3.1: Details of the controlled valves within the pneumatic manifold of the combustion system.

3.5.1.6 Gas delivery

Post SEOP Gases

The combustion system receives post SEOP gases directly from the polariser and serves as an intermediary between the outlet manifold and the polariser distribution manifold in Figure 2.10. For experiments that monitor the polarisation loss due to the combustion process, and where direct SEOP polarisation validation is required, researchers fitted an additional gas line to bypass the combustion system.

Hydrogen

The delivery of pure hydrogen occurs via two methods. For initial volumes prior to the combustion process, delivery is achieved using a small transfer cylinder, pressurised to 5 to 6 bar. Multiple deliveries from this apparatus are performed as required, with each yielding approximately 50-60 mbar of gas pressure in the reaction vessel under standard conditions. For subsequent deliveries of hydrogen, generally to feed the combustion process, an additional hydrogen supply system is attached directly to the distribution manifold, as seen in Figure 3.13. This feeds the reaction vessel via an intermediary metering valve, calibrated to deliver hydrogen at twice the flow rate of oxygen delivered by the oxygen metering valve.



Figure 3.14: Image of the hydrogen transfer cylinder, including plug valve and female quick connect.

Oxygen

Oxygen is delivered to the reaction vessel via the oxygen manifold, depicted in Figure 3.13. Through this manifold, oxygen can flow into the chamber following three distinct paths. The first two paths constitute the fast oxygen delivery lines. These paths both enter the reaction vessel via the same inlet and are controlled automatically via the same pneumatic valve. However, prior to this valve, two parallel paths permit the selection of two flow rates, either with no specific restriction of the gas flow, or with a flow restrictor (a length of 1/8" OD tube). Manual operation of two plug valves determines the selection of these paths before the automated sequence commences. The final path provides a metered flow of oxygen, using the oxygen metering valve. This route enters the reaction vessel through the catalyst and is controlled via its own pneumatic valve.

Rather than connecting an oxygen cylinder directly to the system, oxygen was delivered via an intermediary vessel. This was attached to the oxygen manifold via a quick-connect. Originally, a simple party balloon was employed for this purpose, attached to the quick connect using elastic bands. However, this apparatus broke frequently and introduced contamination via the powder coating on the interior surfaces of the balloons. The balloon was therefore replaced with a Tedlar bag (similar to those used for hyperpolarised xenon-129 delivery to humans[29]), connected via a JACO connector to the quick connect. This apparatus yielded a significantly improved operational lifetime, as well as eliminating new powder contamination of the pneumatic manifold. However, this device provided a non-linear delivery profile for the oxygen during near-total delivery from the vessel. This occurred because, when nearly depleted, the walls of the Tedlar bag would partially seal the gas outlet and restrict the flow of oxygen. Therefore, in the final iteration of this design, foam was introduced into the Tedlar bag to prevent it from self-sealing at the end of the delivery.

Originally, the oxygen transfer vessel was filled using a pneumatic pistonbased apparatus, depicted in Figure 3.15. Whilst simple to operate, it was difficult to repeatably fill the transfer vessel to the same volume of gas.

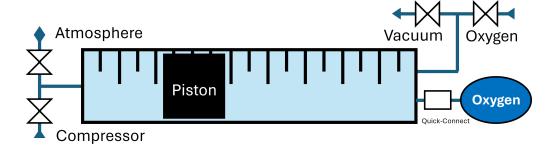


Figure 3.15: Design of the piston-based Oxygen Filling station.

As evident in section 3.5.6, accurate delivery of oxygen was of paramount importance for generating the greatest purity of hyperpolarised gas. There-

fore, an automated oxygen filling apparatus was developed. This utilised a slow oxygen flow rate through a calibrated metering valve, to produce a user-selected quantity of gas by flowing for a predetermined duration. The apparatus was controlled via an Arduino and is depicted in Figure 3.16.

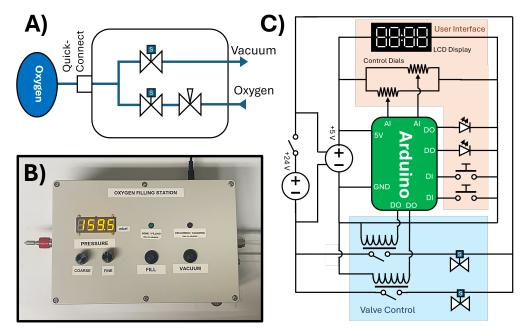


Figure 3.16: Apparatus for filling the oxygen delivery vessel, including the pneumatic diagram (A), an image of the interface (B) and the electronics diagram (C). Arduino serial connection to the LCD display is omitted for simplicity.

While the oxygen filling device represented a significant improvement over the previous apparatus, partly because it was fully automated, it was not without its faults. Primarily, changes in experimental conditions such as room temperature and pressure would slightly change the oxygen flow rate, reducing the accuracy of the device. The flow rate variations were generally repeatable within a few hours, however the next iteration may benefit from a calibrated mass-flow metre to control the oxygen delivery, rather than a simple timing device. Finally, the device is calibrated to produce a specific volume of gas to generate a user-determined pressure within the reaction vessel. Described by the ideal gas law, this is also dependent on the internal temperature of the reaction vessel. The delivery timing calculations are based on an assumed vessel temperature of 20°C, however future iterations of this device should account for temperature variations in the reaction vessel.

3.5.2 Polariser Control

The combustion system is also capable of controlling the polariser described in Chapter 2. This is achieved using the same methods as the solenoid valves, where then DAQ sends a 5 V signal to the relays in the polariser control box (Figure 2.12). This functionally is useful for the automated purification of hyperpolarised gas, especially when multiple deliveries are required from the polariser such as during the semi-continuous flow mode.

3.5.3 Control Software

The combustion system is typically controlled using an external computer (Computational Hardware detailed in Appendix B). The control software (coded in MATLAB 2024a) allows the user to control the gas flow in and out of the combustion system, whilst displaying the real-time pressure changes within the chamber. The GUI used for this is depicted in Figure 3.17.

The state of each pneumatic valve is controlled by passing a digital signal through the DAQ board output to the respective solenoid valves. This control system has a maximum update frequency of 5000 Hz, although the response time of the pneumatic valves will be significantly slower due to the non-instantaneous gas flow used to operate the valves. The valves can be controlled either manually in the software, or by the pre-encoded automated delivery sequences. A similar method is used to apply the potential across

the electrodes, although this is terminated after 0.05 s.

Within the control software, a data acquisition object seeks pressure information (in the combustion vessel and recompression volume) from the DAQ at a rate of 0.1 Hz. Each data dump collects 100 pressure data points and associated time points. These are averaged to ten groups, before plotting on the pressure plot in the GUI. Therefore, the temporal resolution of the plotted pressure data is 0.01 s. The display window is 10 s wide, plotting the previous 8 s of pressure data. This object is also used to run the automated sequences, described in Section 3.5.3.2.

An additional data acquisition object seeks temperature information from a thermocouple placed on the glass exterior of the reaction vessel.

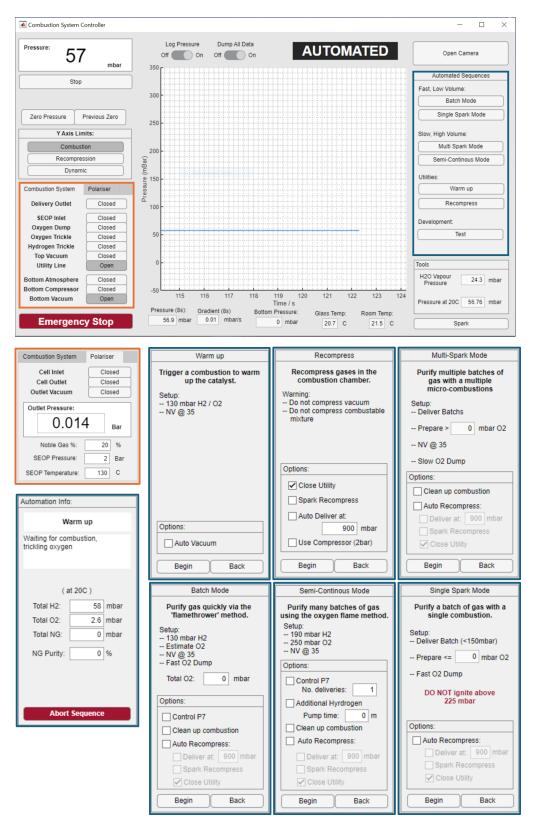


Figure 3.17: Control GUI for the Combustion System, including alternative panels highlighted in the **orange** and **blue** boxes.

Pressure and temperature data is stored in a FIFO buffer, then in a binary

file to prevent memory overflow for experiments of significant duration. this data can be later saved if desired for future analysis. The position state of each valve for each time-point is also saved in this way.

3.5.3.1 Display Pressure Curves

The saved pressure and temperature data can be visualised using the display interface depicted in Figure 3.18. This interface can overlay up to six datasets, each time-shifted independently depending on the associated valve positions.

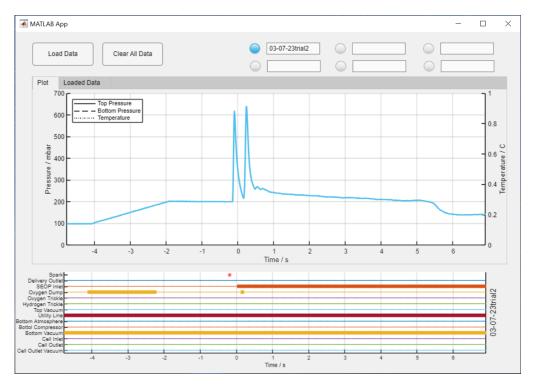


Figure 3.18: Display interface for visualising the pressure and temperature variation during experiences, with associated valve states.

3.5.3.2 Automation

Due to the fast reaction rates typically associated with hydrogen-oxygen ignition, it is necessary to automate the purification process. This automa-

tion is achieved within the control software, using nested switch statements. Both the sequence and an stage within said sequence are assigned to variables. During each iteration of the data acquisition object, the software look up code to execute based on these parameters. Therefore, the refresh rate of the automated sequences is 10 Hz.

An example nested switch statement is depicted below, for two typical steps of automation. To initiate the automated sequence, the 'Sequence' variable is set to the desired sequence and the 'Step' variable set to one. In the example code, the software will check the pressure within the reaction vessel ten times a second, until it surpasses 100 mbar. A this point, a timer is started for step two, and the 'Step' variable updated as such. The code will then continuously check how much time as passed, until this parameter reaches 10 s. This is the end of this automation sequence, so the Step variable is set to 0, which in turn will set the set the Sequence variable to None.

This method of iteratively updating the auto sequence is necessary to permit the continued update of the pressure plot in Figure 3.17. However, it means the sequence can only update at a rate of 10 Hz, thus fast ignition evens could be missed. For this reason, the automation programming is capable of looking back at the last 8 s of pressure data in the FIFO buffer.

3.5.4 Safety

Operation of the combustion system comes with its own risks, both to the operators and the equipment itself. Primarily, these stem from the use of the highly flammable hydrogen gas and the risk of vessel rupture due to over-pressure. Therefore, the system was designed to incorporate a number of safety features, both active and passive. These measures aim both to reduce the likelihood of an accident, or mitigate the effects if one occurs. Aside from a significant hydrogen leak, the greatest risk to the operators and equipment associated with the combustion system would likely be due to the rupture of the central reaction vessel. Rupture of the vessel may launch projectiles towards the operators and the equipment, particularly the glass of the vessel walls. This is most likely to occur due to the rapid pressure increase associated with the oxidation of a large quantity of hydrogen gas. A quick back-of-the-envelope calculation indicates that a stoichiometric mixture of hydrogen and oxygen will increase in pressure 10-fold during complete combustion, reaching a temperature of nearly 3000 K. Measures taken to reduce the risk of this happening fall under three categories: hardware, software and operation.

3.5.4.1 Hardware

The safety precautions built into the system are depicted in Figure 3.19. The glass walled reaction vessel is expected to sustain internal pressures of up to 4 atmospheres pressure differential. However, two physical mechanisms are in place to ensure this condition is never reached. First, a pressure release valve calibrated to a 1 bar pressure differential is mounted to the pipework beneath the recompression chamber. This item's primary purpose is to prevent the over-pressure of the vessel during the recompression sequence. Secondly, the acrylic end caps are secured to the reaction vessel using four springs in parallel. This feature allows the reaction vessel to open, should pressure within surpass a calibrated limit: the springs are compressed to a calibrated length to permit venting at a 0.4 atmosphere over-pressure. This feature should reduce the risk of a catastrophic failure of the reaction vessel due to over-pressure, save for two factors:

• The opening created in the pressure vessel is small compared to the volume of the vessel; therefore, it may not be able to vent an adequate quantity of gas in a short-time period.

• The rate of vessel opening is dependent on the mass of the acrylic end caps, as dictated by Newton's second law. Therefore, for very fast increases in pressure, they may not open sufficiently rapidly to prevent the failure of the reaction vessel.

Should the springs be insufficient to release the end caps, the nylon rods supporting them are designed to be sacrificial. In the event of significant over-pressure, they will break under the tensile stress, fully releasing the contents of the reaction vessel. In this eventuality, the steel end plates should prevent the end caps from causing injury.

The system also includes a number of safety measures to mitigate the risk to the operators and equipment should the internal reaction vessel fail catastrophically. This is in the form of a hierarchical shield system, comprised of a primary, secondary and tertiary shield. The primary shield is made up of a Perspex cylinder and capped by the same acrylic end caps as the reaction vessel. This Perspex shield contains a number of blow-out holes, directed away from the opening of the system housing. The second shield comprises three sides of an aluminium grating, with clear Perspex on the fourth side to allow viewing of the system by the internal camera. This shield is capped with steel plates. The final shield is an aluminium frame, with one side covered by a thick curtain to allow the operator to access the system. This layer serves a dual purpose, both as a shield and for mounting electronic and pneumatic control systems. Finally, at the base of the compression chamber, a rubber buffer is sometimes installed. This is to protect the piston and bottom end cap from a sudden impact, although should be removed for accurate internal gas mix calculations.

The risk of a hydrogen fire is mitigated by storing the pure hydrogen separately from the combustion system, thus minimising the quantity of combustible gas available to the system at any one time. However, this is not the case for the hydrogen/noble gas mix, which is connected via the SEOP polariser (Chapter 2). A hand-held combustible-gas meter is used to monitor for hydrogen leaks, although in future this capability should be integrated into the system itself.

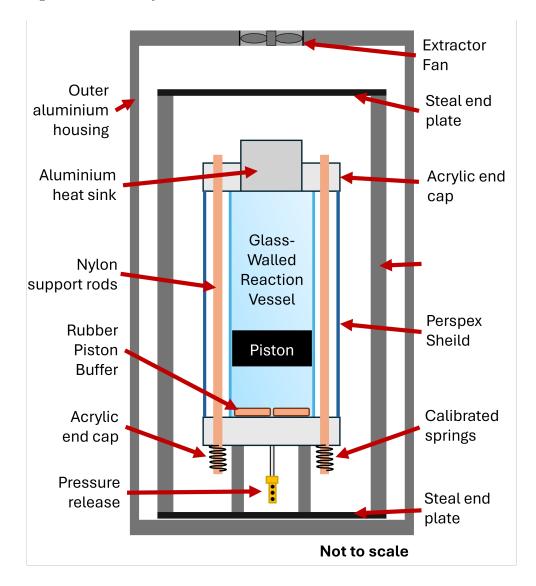


Figure 3.19: Diagram of physical hardware designed to protect the operators and equipment. Including active and passive measures to for preventing and mitigating risks associated with the combustible medium.

3.5.4.2 Software

In addition to the hardware-based measures, safety controls are integrated into the control software. These include limiting the activation of the 'spark plug' to low pressure mixtures, below 300 mbar. This should reduce the risk of a catastrophic failure of the reaction vessel due to the rapid increase in pressure associated with combustion, although it does not abate the risk of a catalytically induced reaction.

Additionally, the system will not permit the activation of the recompression sequence while the reaction vessel is under vacuum conditions. This is to protect the piston and internal components of the reaction vessel, as pressurising the compression volume under vacuum conditions would cause the piston to impact the top acrylic cap and aluminium catalyst housing.

3.5.4.3 Operation

To protect the central reaction vessel, all pre-recompression operation within the chamber is limited to low-pressure gas mixes (< 0.5 bar). This ensures that even with an erroneous combustion, the pressure will not spike to above the rupture pressure of the vessel. Furthermore, when the combustion is expected, internal pressures are maintained below 0.25 bar.

Emergency-stop buttons are integrated into the system, both physically on the outside of the housing and virtually in the control software. When pressed, the emergency-stop buttons cease all automated sequences, close all valves in the system, and ensure the piston is retracted to the bottom of the chamber. The control software is also able to operate this emergency stop protocol, should the reaction vessel exhibit unexpected internal conditions.

Finally, core functionalities such as control of the piston position and vacuum of the reaction vessel can be manually bypassed without the control software. The default position of both the top utility and bottom vacuum pneumatic valves is set to 'open' (Table 3.1), therefore the system can be controlled manually in the event of power loss or control software malfunction.

3.5.5 Operational Modes

Multiple operational modes are programmed into the control software, each optimised for different gas quantities or noble bas nuclei.

3.5.5.1 Single Spark

The simplest method of purifying hyperpolarised gas mixtures is to deliver a stoichiometric mixture of hydrogen and oxygen and activate the electrodes to ignite the mixture. This sequence is named 'Single Spark' and is described in Figure 3.20.

Single-Spark Automated Sequence

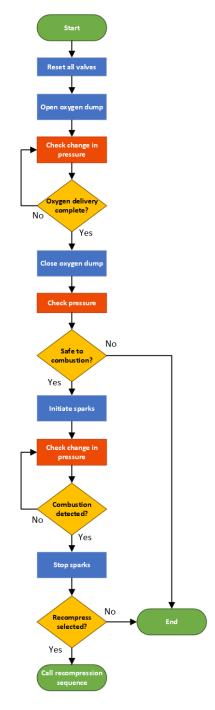


Figure 3.20: Single-Spark Automation Sequence Flowchart

However, while simple to implement, this sequence is limited in the quantity of gas it can produce due to structural integrity of the reaction vessel. If a stoichiometric is expected to expand to ten times its starting pressure suing a reaction, and the vessel can sustain up to 5 bar, than the max-

imum pressure of gas that ignited is 500 mbar. In reality, to keep safe working bounds, this is limited to 250 mbar. Therefore, for a 5% noble gas mix, the maximum quantity of hyperpolarised gas that can be produced is approximately 30 ml.

An alternative method to purify a larger quantity of gas, is to continually react gas mixtures as they enter the chamber, or between multiple deliveries of gas.

3.5.5.2 Catalyst Warm-up Sequence

For these sequences, the reaction vessel must be primed, this can be achieved using a combustion event initiated using the electrodes or the catalyst.

Warm-Up Automated Sequence

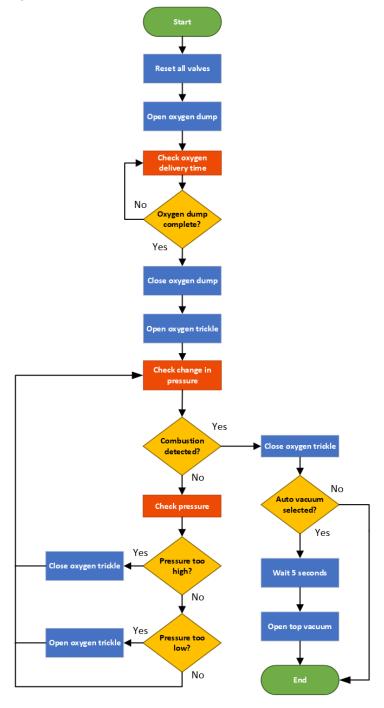


Figure 3.21: Catalyst Warm-up sequence Flow Chart

3.5.5.3 Batch Mode

Batch mode is designed to purify a full delivery of gas from the polariser very quickly. This mode utilises a primed reaction vessel with excess oxygen. Delivery of a post SEOP gas mixture into this regime results in the immediate reaction of the hydrogen. This is because the temperature of the gas is above the self-ignition point of a hydrogen-oxygen mix. Extra oxygen can be delivered gas required to sustain the reaction, therefore the quantity of gas purified by this mode is limited by the delivery from the polariser. A clean up combustion is typically required to attain the greatest noble gas concentrations

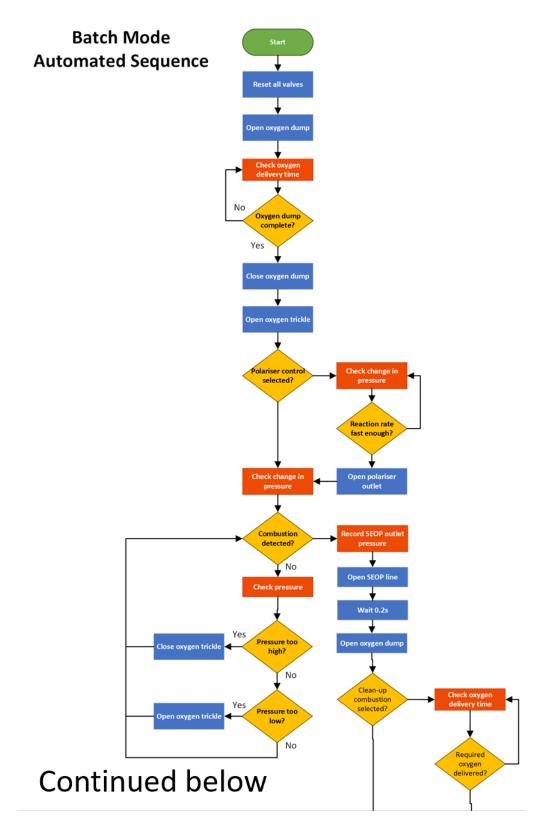


Figure 3.22: Batch-Mode Automation Sequence Flowchart

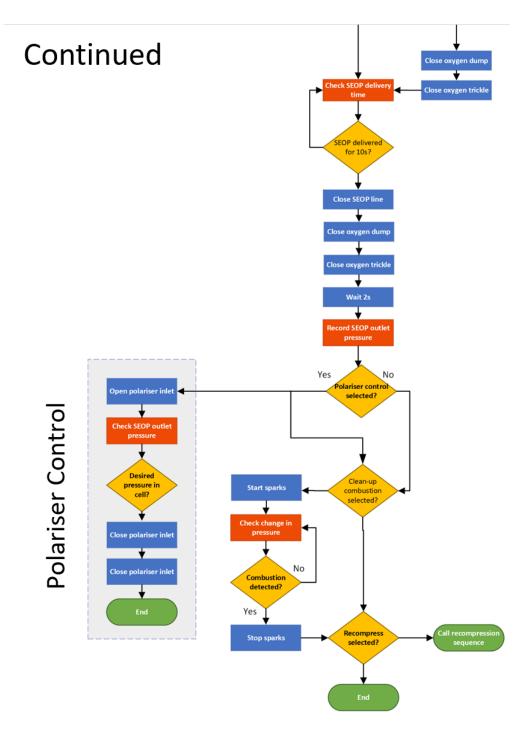


Figure 3.23: Batch-Mode Automation Sequence Flowchart, Continued

The hyperpolarised spin state resistance to the batch mode process is demonstrated in Figure 3.26.

3.5.5.4 Semi-Continuous Mode

The semi-continuous mode sequence is designed to gradually purify many deliveries of hyperpolarised gas mix slowly. This method requires a catalyst based warm-up step. This method utilises an inverted oxygen flame in the catalyst, within a hydrogen environment. This hydrogen reservoir can be topped up using gas from the polariser, or using the hydrogen delivery metering valve calibrated to double the flow rate of the oxygen delivery. SEOP deliveries can extinguish the flame, however it can be relit using the electrodes. This mode is not suitable for the quadrupolar nuclei due to the fast quadrupolar relaxation.

3.5.5.5 Re-compression

The recompression sequence is automated using the pipeline in Figure 3.24. This permits fast recompression using both atmospheric and compressed air. The delivery lines can also open at a predetermined point during recompression for optimal signal recovery.

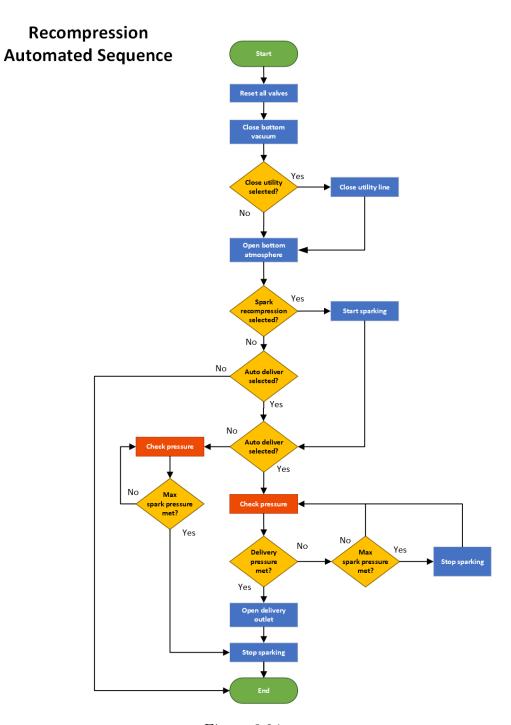


Figure 3.24

3.5.6 Gas Mix estimation

It is difficult to track the reaction using changes in pressure due to the large variations of pressure inside the reaction vessel that are not necessarily detected by the thermocouple. However, the purity of the resulting mix is calculated based on the known gas input into the system. This is based upon the measured priming hydrogen and oxygen quantities, the gas delivered from the polariser (accounting for cell temperature), the oxygen quantity in the delivery Tedlar bag, or the total time the oxygen metering valve was open. Water vapour partial pressure is estimated based the measured temperature of the glass, using the Buck equation[30]:

$$P = 0.61121 \exp\left((18.678 - \frac{T}{234.5})(\frac{T}{257.14 + T})\right)$$
(3.5)

3.6 Results

3.6.1 Batch Mode

A pressure curve for a batch-mode sequence is presented in Figure 3.25.

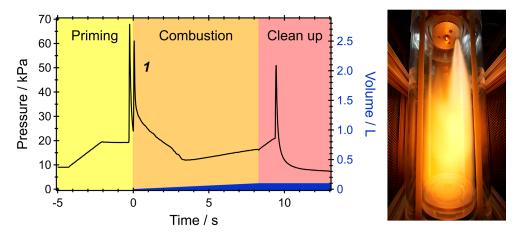


Figure 3.25: Pressure within the reaction vessel (A) during a batch-mode sequence and associated image of the SEOP gas delivery (B). This mode concentrates a single delivery of noble gas (100 ml demonstrated) upon delivery to the combustion chamber. The vessel is primed via spark-triggered combustion with excess Oxygen before SEOP gas is delivered, combusting on entry. A final spark-ignited 'clean-up' combustion increases the noble gas purity. Blue volume denotes the quantity of noble gas purified using a 10% noble gas / 90% $\rm H_2$ mixture. Pressure spikes on delivery of SEOP gas are not representative of pressure within the vessel.

The resistance of the hyperpolarised spin state to the batch mode process is demonstrated in Figure 3.26. In this figure, the relative signal enhancement (RSE) is employed as a unitless metric to capture the many changing dynamics of the reaction process. These include contributions from the change in polarisation and gas purity. However, as it is a relative factor, it omits the magnitude of the polarisation. RSE is calculated based on the apparent polarisation (AP), at several time points (t):

$$RSE = \frac{AP(t)}{AP(t=0)} \tag{3.6}$$

where AP(t=0) is the apparent polarisation prior to combustion. Therefore, the RSE metric is effectively the relative apparent polarisation. In Figure 3.26, each experiment used a 20% noble gas mix; therefore, the maximum RSE, representing 100% gas purity and no loss to polarisation, is five. For the 71% and 81% respective xenon-129 and krypton-83 purities

depicted in this figure, the RSEs with no penalty to the polarisation would be 3.55 and 4.05. By back-calculating the polarisation losses due to relaxation in the chamber (dashed red line), the hyperpolarised state of both xenon-129 and krypton-83 appears largely unaffected by the combustion process.

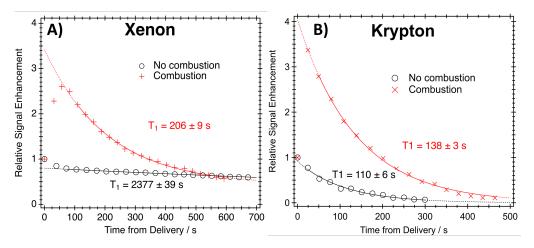


Figure 3.26: Pressure corrected signal enhancement and decay after purification via batch mode (without clean up step) (red +'s) of each noble gas, compared to a non-combusted delivery (black circles). Un-combusted T1 measurements were collected at a chamber pressure of 14 ± 1 kPa, with 20% Xe/Kr and 80% H₂ mixtures. Post combustion T1 measurements were collected at 4.3 ± 0.5 kPa, $71\pm2\%$ Xe, $10\pm2\%$ O₂, $19\pm2\%$ H₂, plus water vapour and 3.5 ± 0.5 kPa, $81\pm2\%$ Kr, $12.5\pm2\%$ O₂, $6.5\pm2\%$ H₂, plus water vapour. Data corrected for an assumed water vapour pressure of 2.2kPa at 19° C.

3.6.2 Semi-Continuous Mode

A pressure curve for an example semi-continuous automated experiment with six gas deliveries for the hyperpolariser is demonstrated in Figure 3.27.

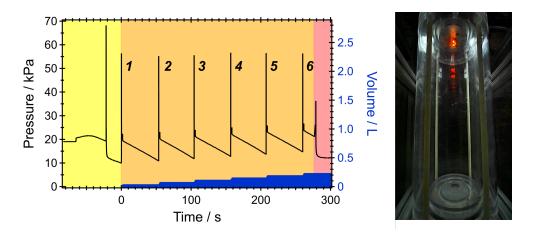


Figure 3.27: Pressure within the reaction vessel (A) during semi-continuous-mode sequence and associated image of the SEOP gas delivery (B). Semi-continuous mode concentrates multiple deliveries of noble gas via gradual reaction. The vessel is primed via catalytic-triggered combustion to initiate an Oxygen-fed flame at the catalyst. Multiple batches (6 pictured) of SEOP gas can then be delivered into the vessel to attain the desired noble gas volume (220 ml demonstrated), whilst the buffer gas combusts via the flame. A 'clean-up' combustion finalises the purification. Blue volume denotes the quantity of noble gas purified using a 10% noble gas / 90% H₂ mixture. Pressure spikes on delivery of SEOP gas are not representative of pressure within the vessel.

Further experimental work is required to investigate the combustion resistance of xenon-129 during this operational mode.

3.7 Discussion

As shown in Figure 3.26, the hyperpolarised spin states of xenon-129 and krypton-83 appear to sustain during the combustion process of batch-mode operation. This suggest that these spin states are resistant to the combustion process at low field and during extended reactions. Future work should focus on further development of this system and further testing of the hyperpolarised spin resistance under other operational modes.

This system has a number of benefits over established cryogenic separation

techniques. Primarily, this is the only method currently able to purify the quadrupolar nuclei, krypton-83 for hyperpolarised studies. This represents a fundamental stepping stone in the path for future clinical applications of this nucleus for respiratory MRI. Additionally, the consumable overhead of this method is lower compared to established systems. This reactive separation technique is cryogen-free; therefore, it does not require a source of liquid nitrogen. This resource can require significant infrastructure to supply and store and is accompanied by safety risks when handling. Furthermore, this technique foregoes helium as the buffer gas, which is non-renewable. Instead, this system employs a pure hydrogen buffer gas, which can be produced relatively simply via electrolysis of water. Finally, the fully automated reaction sequence eliminates the need for specialist operator training and reduces susceptibility to user errors during operation.

While this method has a reduced consumable overhead, the equipment requirements of this prototype are significantly more complex than typical cryogenic separation hardware. In comparison to the cryogenic separation apparatus described in Figure 2.11, this system has a larger footprint and comprises many more components, each of which could be a point of failure. Furthermore, although this system alleviates safety concerns from cryogens, the system uses significant quantities of hydrogen. While every effort has been made to reduce the risk of a thermal explosion and mitigate the effects should it occur, the hazard remains.

This system introduces a number of extrinsic depolarisation mechanisms, although the relative contribution from each of these mechanisms has not been investigated. Interactions of the hyperpolarised nuclei with the internal surfaces of the reaction vessel (wall relaxation) likely contributed to the depolarisation of the ensemble. In subsequent iterations, the hyperpolarised xenon-129 lifetime in the system could be extended by coating the

walls with a siliconising agent; however, this was omitted in this system for adaptability between xenon-129 and krypton-83. The relaxation rate due to this mechanism likely increased during recompression, particularly for krypton-83, due to the increased surface-to-volume ratio. Additionally, the low molar mass of hydrogen-based gas mixtures will result in greater diffusion within the reaction vessel. Therefore, depolarisation due to magnetic field gradients is likely increased for hydrogen gas mixtures in the system compared to mixtures utilising pure nitrogen as the buffer gas. However, no experiments were completed to quantify this effect.

Finally, for xenon-129, significant depolarisation is observed following combustion due to excess oxygen. Even with a purely stoichiometric mixture, trace oxygen appears to remain following an ignition event. This quantity is increased under excess oxygen conditions; therefore, precise gas metering is of paramount importance. Relaxation due to this mechanism is accelerated with the recompression of below-ambient gas mixtures, due to the increased partial pressure of oxygen. One could reduce the contribution from this depolarisation mechanism by using excess hydrogen in the reaction process. However, this may be problematic following delivery of the gas, especially if it is to be inhaled by a patient. As per Figure 3.3, hydrogen gas in air exceeding 4% by volume poses a risk of ignition.

This system is limited by the volume of gas produced during SEOP and the pressure limits imposed on the system for safety. To purify the quantities of gas typically required for clinical lung studies (0.6 to 1.0 L[29]) using the current iteration of this system would require a fast throughput from the SEOP polariser, or the storage of hyperpolarised gas. However, adapting the SEOP gas flow-rate may incur a significant penalty to the initial polarisation of the hyperpolarised gas. Furthermore, although the hyperpolarised state of each nucleus is demonstrated to be resistant to the purification pro-

cess, wall relaxation in the reaction vessel and the persistence of oxygen in the gas mixture following purification reduce the feasibility of storing the purified hyperpolarised gas. Therefore, to minimise polarisation losses, the resultant gas mixture should be consumed soon after processing. Alternatively, non-purified post-SEOP gases could be stored before delivery to the reaction vessel. This would alleviate depolarisation contributions due to contact with oxygen. However, this would require the storage of significant quantities of hydrogen: to produce 1 L of xenon from a typical 3% mix would require the storage of over 30 L of hydrogen. Without further development, this constitutes a significant hazard.

Future versions of this technology could permit the storage and ignition of larger quantities of gas, prior to delivery to purification. As discussed above, this would prevent depolarisation due to prolonged contact with oxygen and may be facilitated by an all-aluminium combustion vessel. Alternatively, many small quantities of gas, from either a storage reservoir or from the cell itself, could be purified in rapid succession. This may limit depolarisation by shortening the time for which the polarised species is in contact with oxygen and surfaces.

3.8 Conclusion

This work describes the development and operation of a prototype apparatus to demonstrate the feasibility of the purification of large quantities (100 ml to 220 ml) of hyperpolarised xenon-129 and krypton-83 via reactive separation. Instead of the nitrogen-helium buffer gas employed with conventional SEOP methods, this technique employs a hydrogen buffer gas, which is removed via reaction with oxygen following SEOP. At present,

this cryogen-free technique is the only method for purifying hyperpolarised krypton-83, a key milestone for the future clinical applications of this hyperpolarised nucleus.

The system is capable of purifying hyperpolarised gas via two automated operational sequences, batch and semi-continuous modes, providing rapid purification of a single delivery or prolonged accumulation of multiple deliveries of gas, respectively. The hyperpolarised spin states of both xenon-129 and krypton-83 are demonstrated to be resistant to the reaction process during batch mode operation, with negligible depolarisation observed from the reaction process itself. Instead, observed depolarisation is thought to stem from wall relaxation and the interaction with residual oxygen remaining in the gas phase following purification. The system is also capable of recompressing the resultant below-ambient-pressure purified gas; however, contributions to the depolarisation of the hyperpolarised gas are exacerbated by this process.

The system described in this work served its purpose as a prototype device to demonstrate the feasibility of this technology to purify a significant quantity of hyperpolarised xenon-129 and krypton-83 at low field. However, in its current state, it is not applicable to clinical studies due to residual oxygen and hydrogen in the gas phase, which cause depolarisation and pose a safety risk to the patient. Future iterations of this device could focus on either purifying a large quantity of hyperpolarised gas via a single ignition event, or numerous small quantities in rapid succession.

3.9 References

- Happer, W. Optical Pumping. Reviews of Modern Physics 44, 169–249. doi:10.1103/revmodphys.44.169 (1972).
- 2. Wagshul, M. E. & Chupp, T. E. Optical pumping of high-density Rb with a broadband dye laser and GaAlAs diode laser arrays: Application to He 3 polarization. **40**, 4447–4454. doi:10.1103/physreva. 40.4447 (1989).
- Walker, T. G. & Happer, W. Spin-exchange optical pumping of noble-gas nuclei. Reviews of Modern Physics 69, 629–642. doi:10. 1103/revmodphys.69.629 (1997).
- 4. Carter, T. S. Absorption and fluorescence of rubidium vapour. German. *Physikalische Zeitschrift* **11**, 632–633 (1910).
- 5. Hrycyshyn, E. S. & Krause, L. Inelastic collisions between excited alkali atoms and molecules. VII. Sensitized fluorescence and quenching in mixtures of rubidium with H2, HD, D2, N2, CH4, CD4, C2H4, and C2H6. Canadian Journal of Physics 48, 2761 (1970).
- 6. Meersmann, T. & Brunner, E. Hyperpolarized xenon-129 magnetic resonance: concepts, production, techniques and applications eng (eds Meersmann, T. & Brunner, E.) Cambridge, 2015.
- 7. Brunner, E. Enhancement of surface and biological magnetic resonance using laser-polarized noble gases. *Concepts in Magnetic Resonance* 11, 313–335. doi:10.1002/(sici)1099-0534(1999)11: 5<313::Aid-cmr4>3.0.Co;2-c (1999).
- 8. Driehuys, B., Cates, G. D., Miron, E., Sauer, K., Walter, D. K. & Happer, W. High-volume production of laser-polarized 129Xe. Ap-

- plied Physics Letters **69**, 1668–1670. doi:10.1063/1.117022 (Sept. 1996).
- Six, J. S., Hughes-Riley, T., Stupic, K. F., Pavlovskaya, G. E. & Meersmann, T. Pathway to Cryogen Free Production of Hyperpolarized Krypton-83 and Xenon-129. *PLoS ONE* 7, e49927. doi:10.1371/journal.pone.0049927 (2012).
- Hughes-Riley, T., Six, J. S., Lilburn, D. M. L., Stupic, K. F., Dorkes,
 A. C., Shaw, D. E., Pavlovskaya, G. E. & Meersmann, T. Cryogenics free production of hyperpolarized 129Xe and 83Kr for biomedical MRI applications. *Journal of Magnetic Resonance* 237, 23–33. doi:10.1016/j.jmr.2013.09.008 (2013).
- Kuzma, N. N., Patton, B., Raman, K. & Happer, W. Fast Nuclear Spin Relaxation in Hyperpolarized Solid 129Xe. *Physical Review Letters* 88, 147602. doi:10.1103/PhysRevLett.88.147602 (Mar. 2002).
- 12. Gatzke, M., Cates, G. D., Driehuys, B., Fox, D., Happer, W. & Saam, B. Extraordinarily slow nuclear spin relaxation in frozen laser-polarizedXe129. *Physical Review Letters* **70**, 690–693. doi:10.1103/physrevlett.70.690 (1993).
- Cowgill, D. F. & Norberg, R. E. Spin-Lattice Relaxation and Chemical Shift of Kr83in Solid and Liquid Krypton. *Physical Review B* 8, 4966–4974. doi:10.1103/physrevb.8.4966 (1973).
- Cowgill, D. F. & Norberg, R. E. Pulsed NMR studies of self-diffusion and defect structure in liquid and solid krypton. 13, 2773–2781. doi:10.1103/physrevb.13.2773 (1976).
- 15. Imai, H., Yoshimura, H., Kimura, A. & Fujiwara, H. Continuous flow production of concentrated hyperpolarized xenon gas from a dilute

- xenon gas mixture by buffer gas condensation. Scientific Reports 7. doi:10.1038/s41598-017-07695-7 (2017).
- Parkhomenko, A. I. & Shalagin, A. M. An alkali metal vapor laser amplifier. Journal of Experimental and Theoretical Physics 119, 24– 35. doi:10.1134/s106377611407005x (2014).
- 17. Rogers, N. J. et al. Molecular hydrogen and catalytic combustion in the production of hyperpolarized 83Kr and 129Xe MRI contrast agents. Proceedings of the National Academy of Sciences 113, 3164–3168. doi:10.1073/pnas.1600379113 (2016).
- 18. Tam, A., Moe, G. & Happer, W. Particle Formation by Resonant Laser Light in Alkali-Metal Vapor. *Physical Review Letters* **35**, 1630–1633. doi:10.1103/physrevlett.35.1630 (1975).
- Rohrbaugh, S., Wang, H. T. J., Singh, J., Tobias, W. A. & Cates,
 G. D. Magnetic decoupling of 129Xe nuclear spin relaxation due to
 wall collisions with RbH and RbD coatings. *Physical Review A* 86,
 43413. doi:10.1103/PhysRevA.86.043413 (2012).
- Wu, Z., Happer, W., Kitano, M. & Daniels, J. Experimental studies of wall interactions of adsorbed spin-polarizedXe131nuclei. *Physical Review A* 42, 2774–2784. doi:10.1103/physreva.42.2774 (1990).
- 21. Ring, R. W., Stott, J. E. & Hales, C. Modeling the Risk of Fire/Explosion Due to Oxidizer/Fuel Leaks in the Ares I Interstage 2008.
- 22. NASA. Glenn Safety Manual Chapter 6: Hydrogen (NASA, Cleveland, Nov. 2019).
- Cain, T. M. Autoignition of hydrogen at high pressure. Combustion and Flame 111, 124–132. doi:10.1016/S0010-2180(97)00005-9 (Oct. 1997).

- 24. Ono, R., Nifuku, M., Fujiwara, S., Horiguchi, S. & Oda, T. Minimum ignition energy of hydrogen—air mixture: Effects of humidity and spark duration. *Journal of Electrostatics* **65**, 87–93. doi:10.1016/J. ELSTAT.2006.07.004 (Feb. 2007).
- 25. Jennings, T. J., Voge, H. H. & Armstrong, W. E. Catalysts for initiating the hydrogen-oxygen reaction at 78 °K. *Journal of Catalysis* **24**, 493–501. doi:10.1016/0021-9517(72)90133-9 (Mar. 1972).
- Stupic, K. F., Six, J. S., Olsen, M. D., Pavlovskaya, G. E. & Meersmann, T. Combustion resistance of the 129Xe hyperpolarized nuclear spin state. *Phys. Chem. Chem. Phys.* 15, 94–97. doi:10.1039/c2cp43382f (2013).
- 27. Cleveland, Z. I., Stupic, K. F., Pavlovskaya, G. E., Repine, J. E., Wooten, J. B. & Meersmann, T. Hyperpolarized83Kr and129Xe NMR Relaxation Measurements of Hydrated Surfaces: Implications for Materials Science and Pulmonary Diagnostics. *Journal of the American Chemical Society* 129, 1784–1792. doi:10.1021/ja065994t (2007).
- 28. Stupic, K. F., Cleveland, Z. I., Pavlovskaya, G. E. & Meersmann, T. Quadrupolar relaxation of hyperpolarized krypton-83 as a probe for surfaces. *Solid State Nuclear Magnetic Resonance* **29**, 79–84. doi:10. 1016/j.ssnmr.2005.08.008 (2006).
- 29. Stewart, N. J., Norquay, G., Griffiths, P. D. & Wild, J. M. Feasibility of human lung ventilation imaging using highly polarized naturally abundant xenon and optimized three-dimensional steady-state free precession. *Magnetic resonance in medicine* **74**, 346–352. doi:10.1002/MRM.25732 (Aug. 2015).

30. Buck, A. L., Buck & L., A. New Equations for Computing Vapor Pressure and Enhancement Factor. JApMe **20**, 1527–1532. doi:10. 1175/1520-0450(1981)020(1981).

Chapter 4

Development of PREFUL

Post-Processing for

Application at 0.5 T

Abstract

This chapter presents the development of a home-built Phase Resolved Functional Lung (PREFUL) post-processing pipeline optimized for application with 0.5 T scanners. While PREFUL methodology has been successfully demonstrated at 1.5 T, low-field implementation presents unique challenges. These can manifest as reduced SNR, imaging rates and image resolution compared to studies at 1.5 T.

The pipeline was developed using datasets acquired from two 0.5 T systems: an Open MRI system (ASG Paramed, Italy) and a conventional bore system (Bruker, Germany). Key innovations include: automated respiratory phase assignment using 2D lung area measurements; capability to reconstruct an undersampled cardiac cycle; a dual-stage batch-wise registration pipeline combining feature-based and intensity-based methods; and analysis methods for extracting ventilation and perfusion dynamics where SNR is limited. Furthermore, five convolutional neural networks (CNNs) were trained for automated lung segmentation across different acquisition planes and sequences.

Successful ventilation-perfusion processing was demonstrated for 2D gradient echo sequences in coronal and sagittal planes, with mixed results in the axial plane and for coronal 2D HASTE sequences. However, validation against established imaging modalities for assessment of pathological conditions is required for clinical deployment.

This work demonstrates the feasibility of PREFUL methodology at 0.5 T, potentially enabling functional lung imaging in resource-limited settings and facilitating research in postural lung dynamics through open-bore scanner compatibility.

Contents

4.1	Introd	Introduction		
	4.1.1	Functional Lung Imaging		
	4.1.2	Fourier Decomposition		
	4.1.3	Phase Resolved Functional Lung [MRI] 141		
4.2	Aims			
4.3	0.5 T	Scanners		
4.4	PREF	TUL Pipeline		
4.5	Phase	Assignment		
	4.5.1	Ventilation		
	4.5.2	Perfusion		
4.6	Regist	ration		
	4.6.1	Background		
	4.6.2	Lung MR Registration		
	4.6.3	Typical Errors of MR Pulmonary Registration . 180		
	4.6.4	Dual-Stage Registration Pipeline 182		
	4.6.5	Batch-wise Registration		
	4.6.6	Pipeline Integration		
4.7	Filteri	ng		
	4.7.1	Digital Filtering		
	4.7.2	Filter Design		
	4.7.3	Filter Application		
4.8	Recon	struction		
	4.8.1	Algorithm		
	4.8.2	Observations		

4.9	Segmentation		
	4.9.1	Manual	
	4.9.2	Automated	
4.10	Analys	is	
	4.10.1	Ventilation	
	4.10.2	Perfusion	
	4.10.3	Vessel Segmentation	
4.11	User Ir	nterface Implementation	
	4.11.1	Automation	
4.12	Results	245	
	4.12.1	0.5 T Upright Scanner	
	4.12.2	0.5 T Bruker Scanner	
4.13	Discuss	sion	
	4.13.1	Quality of Functional Maps	
	4.13.2	Application with Pathology 260	
	4.13.3	Processing Functional Dynamics 262	
	4.13.4	Comparison to Literature	
	4.13.5	Drawbacks of PREFUL	
	4.13.6	Future Work	
4.14	Conclu	sions	
4.15	Refere	nces	

Declaration of Contributions

- Arthur Harrison All programming of the home-built PREFUL pipeline.

 Labelling of some of the Coronal 2D GRE images, and all images for the other acquisition types. Training and hyperparameter optimisation of SegNet based Coronal 2D GRE network. Training of the Axial and Sagittal 2D GRE networks and the Coronal 2D HASTE network.
- **Dr Olivier Mougin** Optimisation and acquisition of all scans on the 0.5 T Upright scanner (ASG Paramed, Italy) in University of Nottingham..
- **Team at Moscow State University, Russia** Optimisation and acquisition of all scans on the 0.5 T Bruker (USA) scanner in Moscow State University.
- Matilda Chalk Most image labelling for Coronal 2D GRE neural network. Training and hyperparameter optimisation for the UNet-based Coronal 2D GRE network. These training paraments were also used for the Axial and Sagittal 2D GRE networks and the Coronal 2D HASTE network.

4.1 Introduction

4.1.1 Functional Lung Imaging

Currently, X-ray-based methods such as Computed Tomography (CT) dominate the field of structural lung imaging[1]. CT scans have found widespread clinical application as they can provide excellent spatial resolution within short scan times[2]. While this modality is capable of investigating lung function, using methods such as Dual-Energy Computed Tomography (DECT) or Four-Dimensional Computed Tomography (4DCT), these techniques are not widespread in the clinical setting. For functional lung assessment, SPECT scanning serves as a well-established standard for evaluating regional lung ventilation and perfusion[3, 4]. However, these methods necessitate that the patient be exposed to ionising radiation. Radiation dose constraints therefore serve to limit assessment frequency, particularly in paediatric patients due to the greater radiosensitivity of developing tissues and their generally longer life expectancy[5, 6].

4.1.1.1 Proton-Based Lung MRI

Proton-MRI offers an alternative radiation-free approach for pulmonary imaging. This highly adaptable technique is capable of characterising different tissue types, where x-ray-based modalities may have difficulty. These techniques can gain functional contrast by tracking variations in spin density, such in Fourier Decomposition (FD)[7], or by magnetically-tagging blood entering the lung, such as in Arterial Spin Labelling (ASL)[8], or by introducing an exogenous agent to change the relaxation dynamics of protons in the lung tissue, such as in Dynamic Contrast-Enhanced (DCE)[9]

or Oxygen Enhanced (OE)[10] MRI.

However, pulmonary characteristics can complicate the MR acquisition process. Firstly, lung regions have a very low tissue density, approximately 0.1 g/cm³ in healthy individuals[11]. Therefore, even under identical conditions, the SNR within lung tissue is significantly reduced compared to surrounding tissue. In addition, the numerous air-tissue interfaces within the lungs create highly inhomogeneous magnetic field gradients, due to the susceptibility differences between the tissue (diamagnetic) and oxygen (paramagnetic) in the void space. This results in rapid spin de-phasing during gradient echo imaging, characterised by the very short T2* of lung parenchyma (Table 4.1). Therefore, MR investigations of the lungs can be highly challenging, requiring pulse sequences with very short echo times.

T2* of Healthy Lung Parenchyma				
B0 Field Strength	T2*			
0.55 T	$8.2 \pm 2.4 \text{ ms}[12]$			
1.5 T	$2.11 \pm 0.27 \text{ ms}[13]$			
3.0 T	$0.74 \pm 0.10 \text{ ms}[13]$			

Table 4.1: Literature values for the T2* of Healthy lung parenchyma for a range of field strengths

4.1.1.2 Low Field Pulmonary Imaging

As demonstrated in Table 4.1, the T2* value of lung parenchyma is highly dependent on magnetic field strength. Therefore, lung MRI methods could benefit greatly from reduced B_0 field strengths (≤ 0.5 T). This would not only reduce susceptibility artifacts, but also preserve signal from the lung tissue that would otherwise be lost due to rapid de-phasing at higher field. This may enable a greater detection of subtle anomalies and improved over-

all image quality in the organ. However, disregarding associated changes in relaxation mechanisms, the SNR of MR imaging is dependent on B_0 field strength:

$$SNR \propto \sqrt{B_0} \tag{4.1}$$

SNR can be increased by increasing voxel sizes and slice thickness in 2D studies, although this may beget the inability to resolve smaller pulmonary features due to partial volume effects[11]. Alternatively, while all field strengths typically benefit from averaging, this is especially true for low-field studies.

4.1.1.3 Respiratory Gating

MR signal averaging of the pulmonary regions is complicated by the high degree of movement of this organ during to normal respiration. One of a number of techniques can be employed to reduce this effect, known as methods for respiratory gating. The three most common methods employed for signal averaging within the lung are detailed below:

Breath Hold While not strictly a respiratory gating method, the simplest technique for reducing the effect of respiratory motion is to ask the participant to hold their breath at specific stages of the respiratory cycle. The scanning sequence is initiated for the duration of the breath-hold, when movement of the lungs and abdominal organs is assumed to be minimal. Commonly, this method utilises a deep inspiration or exhaled breath-hold; however, repeatability between sessions can be difficult. This method is limited to acquisition times of 15-30 seconds or as long as the participant can hold their breath.

However, this period is often reduced in patients with respiratory conditions. Furthermore, this method is limited to investigations of compliant participants, thus excluding studies of very young children.

Prospective gating In contrast to breath-hold imaging, this method permits the patient to breathe normally during the scanning procedure. The patient's breathing is then monitored in real-time and used to initiate the scanning sequence at specific points within the respiratory cycle. Commonly, the gating window is positioned around the exhaled state, where respiratory motion is minimal. However, this method typically extends the imaging procedure as scans are only acquired during selected portions of the respiratory cycle.

Retrospective gating Similar to prospective gating, for this scanning procedure the patient is permitted to breathe normally. However, unlike prospective gating, information is continuously acquired for a predetermined period, irrespective of the breathing cycle. This data is then retrospectively binned, based on its specific point within the respiratory cycle, which is either measured during the acquisition or determined from the data itself (self-gating). This technique captures the full respiratory cycle, permitting four-dimensional analysis and extraction of dynamic functional information. However, this method typically requires complex post-processing which can introduce uncertainty to the result. Furthermore, although possible using x-ray-based modalities, radiation doses are typically higher due to the continuous acquisition.

Alternatively, instead of 'freezing' the lungs in place by respiratory gating, the complete dynamic free-breathing time-series can be analysed to yield functional information. With appropriate registration - deforming each image such that the parenchyma appears stationary - the signal variation of each voxel correlated with the respiratory and cardiac cycles can be used to attain insight into local ventilation and perfusion. This analysis method is known as Fourier Decomposition (FD) [7] and has served as the basis for a plethora of functional pulmonary MRI studies across a range of field strengths.

4.1.2 Fourier Decomposition

Fourier Decomposition (FD)[7] is a frequency domain analysis technique used for generating pulmonary functional information from a proton-based image time-series of the lungs. The core principle of FD is that the noise-free time varying signal (S(t)) within each stationary lung voxel can be separated into three constituent parts:

$$S(t) = S_0 + A_{vent}\cos(2\pi f_{resp}t + \phi_{resp}) + A_{perf}\cos(2\pi f_{card}t + \phi_{card}) \quad (4.2)$$

where S_0 is the time independent signal intensity associated with the voxel, A_{vent} and A_{perf} are the amplitudes of signal variation due to ventilation and perfusion respectively, which vary depending on their respective frequencies: f_{resp} and f_{card} , with phase offsets ϕ_{resp} and ϕ_{card} with respect to time, t. The first component, S_0 , is the only time-independent part of Equation 4.2, and can provide a qualitative structural representation of the lung. The two cosine terms relate to the signal variation due to the respiratory and cardiac cycles. These terms are assumed to be independent and thus can be isolated via filtering in the frequency domain.

Typically, this technique involves capturing many MR images of free-breathing lungs in quick succession, at a sampling frequency sufficient to satisfy the Nyquist criterion for both the respiration and particularly the cardiac cycle. Each voxel will exhibit a time varying signal at frequencies correlated to each of these cycles, although breathing irregularities and motion artifacts can complicate the analysis. The amplitude of these variations can relate to the change in air and blood associated with the voxel.

For perfusion (Q), it is intuitive to imagine that if a voxel possess a greater blood flow, the variation in the number of spins within that voxel is increased. Therefore, to a first approximation, and disregarding associated changes in relaxation mechanics and blood oxygenation effects (BOLD), the signal attained using MRI should correlate with perfusion:

$$Q \approx A_{perf} \tag{4.3}$$

For ventilation, the relationship is slightly more complicated. Previous work by Zapke et al.[14] demonstrated that for specific pulse sequences (particularly T1-weighted sequences), the ventilation (V) within each voxel can be inversely related to the signal amplitude in that voxel due to air dilution effects on proton density:

$$V \approx \frac{1}{A_{vent}} \tag{4.4}$$

Therefore, the time varying signal of each voxel can be decomposed using a standard Fourier transform, with the intensity of resultant respiratory and cardiac peaks containing voxel-wise information relating to local ventilation and perfusion.

This technique has been validated with with helium-3 MRI[15] and employed to investigate asthma[16] and COPD[17], where a high correlation in VDP was observed between the modalities. Furthermore, ventilation measurements exhibit strong correlations with those derived using fluorine-19 MRI[18, 19]. FD has also been used to investigate pulmonary perfusion in both cystic fibrosis young patients[20] and patients preparing for operations to remove lung cancer[21].

FD exhibits a high degree of reproducibly for investigating ventilation and perfusion in healthy volunteers[19]. This technique has been validated with helium-3 MRI[15] and SPECT/CT[22] and employed to investigate asthma[16] and COPD[17], where high correlations in VDP were observed between the modalities. Furthermore, ventilation measurements exhibit strong correlations with those derived using fluorine-19 MRI[18, 19]. FD has also been used to investigate pulmonary perfusion in both young patients with cystic fibrosis[20] and patients preparing for lung cancer resection operations[21].

With appropriate registration, this method is relatively easy to implement, and permits direct comparison of ventilation and perfusion information for each voxel. However, this method requires uniform temporal sampling, with accuracy dependent on the consistency of patient breathing patterns. Appropriate frequency bands must be selected for the respiratory and cardiac components, although spectral leakage artifacts can be introduced with a non-integer number of cycles sampled. Furthermore, phase information is typically disregarded. Therefore, a number of new techniques have been developed based on the principle of FD, which try to address some of the limitations of this methodology. These include Non-uniform Fourier-Decomposition (nuFD)[23], Matrix Pencil Decomposition (MPD)[24], and Phase Resolved Functional Lung (PREFUL)[25] MRI.

4.1.2.1 Non-uniform Fourier Decomposition

Non-uniform Fourier-Decomposition (nuFD) is an advancement of FD, to account for irregular sampling and time varying respiratory and cardiac characteristics [23]. This includes variations in the rate and depth of patient breathing, which is typical of pathology. This technique takes such signals and, whilst maintaining their sequence, artificially temporally redistributes the sampled points to generate a signal varying at a uniform frequency. This signal undergoes a non-uniform Fourier transform [26, 27], or is interpolated to simulate uniform temporal sampling and regular Fourier transform applied [28]. When implemented correctly, this has the effect of narrowing the spectral peak associated with the respiratory and cardiac cycles compared to standard FD methods. Then like for FD, intensities of these peaks are then used to attain functional information. However, complex phase information is also captured from the spectral frequency bins, to generate 'signal delay' information relative to a reference voxel. While this is possible for standard FD methods, phase errors are common due to the spectral peak overlaps associated with variations in the respiratory and cardiac cycles[23].

4.1.2.2 Matrix Pencil Decomposition

Instead of direct analysis on the frequency domain, MPD employs spectral analysis based on the Matrix Pencil approach linearised least-squares fitting[24]. Following filtering to isolate contributions from the respiratory and cardiac cycles, the variation of the MRI signal of each registered voxel is modelled as the sum of exponentially damped sinusoids[29]. For the respiratory cycle, and without the noise term, this takes the form:

$$S_r(n) = \sum_{p=1}^u a_{r,p} z_{r,p}^n \tag{4.5}$$

where u is the number of detected respiratory frequency components, $a_{r,p}$ is the complex respiratory amplitude $(a_{r,p} \equiv |a_{r,p}|e^{i\phi_{r,p}})$ and $z_{r,p}^n$ the simulated damped exponentials with signal poles:

$$z_{r,p}^n \equiv e^{(i\omega_{r,p} - R_{r,p})\Delta t} \tag{4.6}$$

where for each pole at time, t, $\omega_{r,p}$ is the frequency and $R_{r,p}$ is the damping factor. The signal poles are used to form a matrix:

$$\mathbf{Z_r} = \begin{bmatrix} z_{r,1}^1 & z_{r,2}^1 & \dots & z_{r,u}^1 \\ z_{r,1}^2 & z_{r,2}^2 & \dots & z_{r,u}^2 \\ \vdots & \vdots & \ddots & \vdots \\ z_{r,1}^N & z_{r,2}^N & \dots & z_{r,u}^N \end{bmatrix}$$
(4.7)

where N is the total number of signal points, or the number of acquired MRI images in the time-series. Generally, for a low-pass-filtered respiratory signal, one pole will possess $\omega \approx 0$ and thus represent the signal offset analogous to S_0 in Equation 4.2. The voxel-wise complex amplitudes can be obtained using the Moore-Penrose pseudo-inverse of the matrix and the local signal time-course $(S_n(x, y))$:

$$\mathbf{a}_r(x,y) = \mathbf{Z}_r^+ \mathbf{s}(x,y) \tag{4.8}$$

where \mathbf{Z}^+ is the Moore-Penrose pseudo-inverse of matrix Z. The voxel-wise summation of each of these complex amplitudes yields the total respiratory

amplitude, A_r :

$$A_r(x,y) = \sum_{p=1}^{u} a_{r,p}(x,y)$$
 (4.9)

A similar method is employed for the cardiac contributions and this technique can also yield phase information to calculate blood arrival times [24].

Although computationally expensive, this automated method can provide improved functional analysis compared to typical FD methods, especially for participants with irregular breathing patterns[24]. This technique has so far been used to quantify ventilation and perfusion abnormities in patients[30] and in children with cystic fibrosis and congenital diaphragmatic hernia[31].

4.1.3 Phase Resolved Functional Lung [MRI]

In contrast to the above described techniques, PREFUL performs the majority of analysis in the spatial domain. Therefore, this technique draws significant parallels to retrospective gating, and is typically self-gated to reconstruct a single ventilation or perfusion cycle from the many-image dynamic time series[25]. Combining the registered images in this way greatly improves the typically poor SNR associated with imaging the lungs with proton-MRI and allows for regional function investigation. The PREFUL pipeline includes three core steps:

Phase Assignment The goal of PREFUL is to increase lung SNR via averaging many images together. However, you can't simply combine any frame with another due to the motion of the lungs between acquisitions. Furthermore, to retain functional information, only frames of a similar ventilation or cardiac phase can be combined. Therefore, the first stage of PREFUL is to determine each frame's phase with respect to the ventilation and perfusion cycle, or conduct self-gating.

Image Registration An image series for PREFUL can take many minutes to acquire and subjects are expected to breathe throughout. However, this means the lungs are constantly in motion. Lung tissue moves within the field of view, and thus it occupies different voxels at different points during the ventilation cycle. Prior to averaging, each image must be deformed such that the lung parenchyma appears stationary throughout the time series. Consequently, the second stage of PREFUL is the registration of all frames to the same point of the ventilation cycle.

Interpolation The penultimate-stage of PREFUL is to combine the registered images to create a single breath or heartbeat. Sorting each frame by its previously assigned phase would result in sporadically sampled ventilation and cardiac cycles. Therefore, all frames are interpolated via weighted mean to predetermined equidistant phase points.

Analysis Finally, changes throughout the reconstructed ventilation and perfusion cycle can be used to gain information based on the air and blood-flow dynamics. Strategies for this analysis and the associated biomarkers for pathology are currently an active area of development [25, 32–34].

This process can yield both amplitude and phase offset information to attain local ventilation, perfusion and respective time-to-peak metrics[25, 32]. Contemporary studies also employ flow-loop analysis ventilation analysis[33] as well as blood arrival time for perfusion[34, 35]. However, this

process can be difficult to automate, particularly during respiratory and cardiac binning.

The validity of PREFUL has been extensively explored across multiple repeatability and validation studies with other imaging modalities[36]. Ventilation results have been validated against hyperpolarised xenon-129[33, 37] and fluorine-19[38] MRI, with high correlations in defect detection between the modalities. The validity of derived perfusion results has also been explored, with significant overlap of defect regions determined using SPECT[39]. The repeatability of this technique has been demonstrated with several multi-site, multi-vendor patient studies[32, 40–42].

PREFUL has been employed to investigate ventilation and perfusion abnormalities stemming from a wide range pulmonary pathologies, including asthma[43, 44], COPD[32], pulmonary embolism[41], post-acute COVID-19[45] and chronic thromboembolic pulmonary hypertension[46]. This technique has been used to track responsiveness to COPD treatment[47], as well as to predict the probability of graft loss following double lung transplantation, with a greater accuracy than established spirometry methods[48]. The feasibility of using this method to investigate infants and neonates has also been demonstrated[49, 50].

In the above studies, PREFUL is typically applied to image-series captured at 1.5 T for both 2D[25, 32] and 3D (ventilation only) studies[36, 42, 51]. However, investigations have also utilised this processing method on image-series captured at 0.55 T[52] and 3 T[53].

PREFUL studies typically employ a SPGRE sequence to acquire the dynamic image series, due to its accessibility and robustness at 1.5 T and 3 T[54]. Alternatively, bSSFP sequences are typically used for FD methods[23] and have been employed for PREFUL at low field[52]. Due to the

high T2/T1 ratio of blood[14], these can yield improved SNR in the lung compared to SPGRE techniques[54], however, are susceptible to banding artifacts[55]. Additionally, UTE sequences have been employed for FD based methods when scanning at 3 T[56], due to the short T2* associated with this field strength[13].

4.2 Aims

Typically, the PREFUL algorithm has been employed in conjunction with data collected at 1.5 T, due to the compromise between attainable SNR and T2* inhomogeneity. The aim of this work was to develop a home-built processing pipeline based on the PREFUL algorithm, optimised for application with data collected at 0.5 T. Previously this has been achieved in conjunction with a high-performance prototype 0.55 T whole-body Free.Max MRI system (Siemens Healthcare, Germany) using a bSSFP sequence[52]. However, this methodology has not yet been successfully applied to older 0.5 T MRI systems, or low-field systems with an open architecture.

The successful application of PREFUL with these low-field devices contributes to the democratisation of MRI; demonstrating the feasibility of this modality's deployment in low- and middle-income countries. Furthermore, application of this methodology to open-MRI systems permits the investigation of lung dynamics in a range of postures beyond the horizontal.

The custom-PREFUL pipeline was developed solely in MATLAB 2024b (MathWorks, USA), based on the original work by Voskrebenzev et al.[25]. It was applied in conjunction with off-the-shelf 2D GRE and HASTE sequences, to generate functional maps describing both ventilation and perfusion dynamics. While three-dimensional acquisitions were not explored

in this work, the PREFUL pipeline was developed to process single-slice 2D acquisitions in the coronal, sagittal and axial imaging planes.

4.3 0.5 T Scanners

The custom-PREFUL algorithm was developed using image-series acquired on two 0.5 T MRI systems each pictured in Figure 4.1.

- 0.5 T Tomikon S50 (Bruker, USA) Located at Moscow State University, Russia, this system is equipped with a superconducting magnet (Magnex, UK) with a bore diameter of 60 cm, a 2 kW RF transmitter LPPA 2120 (Dressler, Germany) and a S630 gradient system with a maximum power of 16.7 mT/m with a rise time of 0.5 ms. The built-in 1H whole body coil was used in all experiments. Scan management and initial data processing were performed using XWinNMR v.1 and ParaVision v.1. This system is referred to as the 'Bruker Scanner' henceforth...
- 0.5 T MROpen Evo (ASG Paramed, Italy) Located at the University of Nottingham, UK, this system is equipped with a superconducting magnet, with an open bore width of 56 cm, and a gradient system with maximum power of 20 mT/m with a rise time of 0.6 ms. All experiments were conducted using the commercial 4-channel receive body coil. Scan management and initial data-processing (including B0 field correction and zero filling) were conducted using the inbuilt MR-GUI Pro software. This system is referred to as the 'Upright Scanner' henceforth.

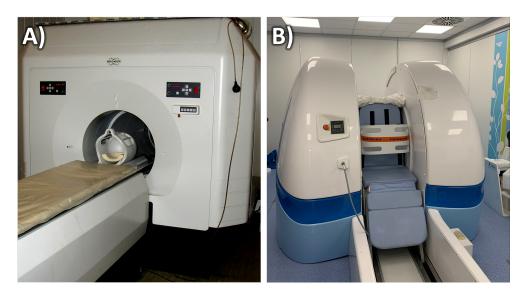


Figure 4.1: Depictions of the 0.5 T scanners utilised in this work: the Tomikon S50 Bruker system (A) and the MROpen ASG Paramed system (B).

4.4 PREFUL Pipeline

An overview of the PREFUL pipeline employed in this work is presented in Figure 4.2.

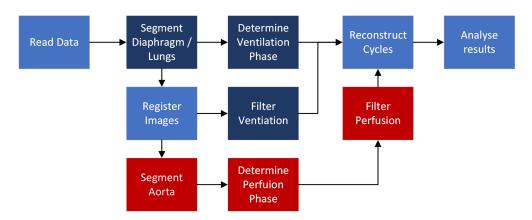


Figure 4.2: Overview of the PREFUL Pipeline used in this work. Blue cells correspond to general processing, Dark blue cells to ventilation specific processing and Red cells to perfusion processing.

4.5 Phase Assignment

A note on 'phase' terminology:

When discussing the reconstruction of the ventilation or perfusion cycles, it is easy to become lost in the various references to phase. This confusion may relate to the phase of the respiratory and cardiac cycle, or the corresponding ventilation and perfusion phase of the reconstructed frames or the phase reconstruction of the MR images. Importantly, phase in this context does not correspond to the phase associated with the reconstruction of MRI images. At the outset, we will coin several definitions to assist with future understanding:

Cycle Phase The phase of each frame within the breath or heartbeat cycle of the person, which is by definition a continuous cycle from 0 to 2π , shall be referred to as respiratory or cardiac phase henceforth.

Reconstructed Phase The artificial phase assigned to each frame with respect to the reconstructed image series, which exists within a single cycle of 0 to 2π and is the phase alluded to when discussing 'Phase Assignment', shall be referred to as ventilation or perfusion phase henceforth.

Lung State Defined as the condition of the lung at a certain time, with respect to the absolute ventilation and perfusion. The core aim of the PREFUL algorithm is to combine frames of the same lung state. As both the respiratory and cardiac cycles vary semi-independently, it is uncommon to find two frames with the same ventilation and cardiac state. This work will attempt to decouple these cycles to permit their consideration independently.

Ventilation reconstruction uses the respiratory phase as the determining factor for reconstructed phase. Perfusion reconstruction uses the cardiac phase.

4.5.1 Ventilation

Methods for determining the ventilation and perfusion phase for each slice can be divided into two camps: Physical methods, using external equipment to measure the respiratory cycle[57]. Or software methods, using the acquired image series itself to determine these parameters.

4.5.1.0.1 Physical Methods

Respiratory bellows A common method to track the respiratory cycle is to measure changes in thoracic circumference. As a subject inhales, the intercostal muscles contract, raising the rib cage and in turn increasing the size of the chest. This can be measured using a device called respiratory bellows[58, 59]: an elastic belt placed around the chest connected to a piezo-electric transducer to measure displacement, generating an electrical signal proportional to the respiratory state. Connected to the scanner, this method has found widespread application as gating for MRI studies of the lungs[60, 61]. Images are only acquired within a designated respiratory range to reduce artifacts from diaphragmatic motion. This method is very reliable and less affected by additional motion[62].

Optical tracking (Camera or Laser) A less common method for tracking the respiratory cycle is via optical tracking of the subject, either tracking the rise and fall of their chest or shoulders. For open scanners, motion tracking using optical cameras can be employed [63]. For conventional scanners, laser range finders combined with mirrors can be used to detect the thoracic movement [64]. This method requires no physical changes to the setup for each subject.

No physical gating methods were explored as part of this work.

4.5.1.0.2 Software Methods

Diaphragm Position Established PREFUL methodology determines the ventilation phase of each frame using the relative position of the diaphragm: it is normally nearest to the apex when exhaled and moves away with inhalation. Diaphragm position is a favourable indicator for this purpose for several reasons:

- Physiological differences between the tissue of the diaphragm and abdominal organs, and the lung parenchyma can provide good MRI contrast between these features. This increases the SNR of the boundary and makes it easier to determine the position of the diaphragm-lung interface.
- Movement of this boundary is normally larger in magnitude than movement of the rest the thoracic cavity during respiration. Therefore, tracking this feature should provide the greatest sensitivity during the respiratory cycle.
- The diaphragm is present in all coronal slices of healthy lungs.

 Therefore, ventilation phase can be determined for each slice in multi-slice acquisitions.

However, this marker is not without its drawbacks. If the diaphragm position is not measured relative to another physiological feature, then it is highly sensitive to movement of the subject within the FOV. For example, if a sitting subject were to slouch during a scan, then their diaphragm may appear artificially lower in the FOV for the rest of the sequence, irrespective of respiratory phase. Additionally, the diaphragm is not the sole mechanism by which the lung expands. During inhalation, the intercostal and accessory muscles contract, lifting the ribs upward and outward. This contribution may result in a decoupling of the ventilation state from the diaphragm position, especially during deep breathing when this mechanism is amplified. Furthermore, this method is often focused on a single lung. Any variation in diaphragm movement between both lungs may create errors during reconstruction.

- 2D Lung Area An alternative marker for determining lung ventilation phase is the two-dimensional area of the thoracic cavity for each frame. This indicator is not measured relative to the FOV, and thus vertical movement of the subject will not skew the results, assuming the complete lung remains in frame. Furthermore, as this method utilises changes of the full lung, it accounts for movements beyond the diaphragm position, albeit those in plane with the captured image. However, manual masking of each frame is time consuming and generally not feasible for the many-image acquisition necessary for PREFUL. Therefore, this process requires automatic segmentation of the lung regions for widespread application.
- 3D Lung Volume The natural progression from 2D area is to mask the complete 3D volume of the thoracic cavity. This shares all the same benefits as 2D, whilst accounting for expansion in every plane. This method would also require automatic segmentation for widespread application. However, PREFUL requires fast acquisition of frames

to resolve the cardiac cycle; therefore, this method is currently only demonstrated on higher field strength scanners. This method was not explored in this work.

Zero Frequency Intensity (Navigator pulse) An alternative method to determine the respiratory state is to track the intensity changes of the centre frequency of the MRI image's frequency domain. This is the simplest of the software-based methods, requiring little processing or user input. It is often used as a prospective self-gating method in thoracic MRI studies.

4.5.1.1 Determining the Respiratory Marker

4.5.1.1.1 Diaphragm Position - Intensity Based

The simplest method for determining the position of the diaphragm is to use the average intensity of an ROI encompassing the lung-diaphragm interface. Generally, the diaphragm has a higher signal intensity than the lung parenchyma. Therefore, when the subject is fully exhaled and the muscle is at its most cranial position, the average intensity of the ROI is greatest. Conversely, when the subject is fully inhaled, a greater portion of the ROI is lung parenchyma, resulting in a lower average intensity. Therefore, this signal variation can be used to track the respiratory cycle. This method is easy to implement; however, it relies on the assumption that the relative signal intensity between the muscle from lung tissue remains constant. This is often not the case for several reasons:

• The cardiac cycle contributes to observed signal intensities, thus, any result must be low-pass filtered.

- The characteristics of the diaphragm may change during expansion and contraction resulting in variation of its observed signal intensity.
- As more of the lung enters the ROI, more lung vessels contribute to the mean intensity of this region. These features usually exhibit an increased signal intensity compared to the surrounding lung parenchyma, therefore can change the ROI's average intensity. This is observed in Figure 4.4Aii as a small increase in the measured signal each time the lung reaches maximum inhalation.

Consequently, as this method is directly dependent on intensities, it is highly susceptible to noise. The filtered signal of the diaphragm position, determined by signal intensity, is presented in Figure 4.4Aii.

4.5.1.1.2 Diaphragm Position - Edge Detection

Instead of relative intensity, edge detection relies simply on existence of a difference in intensity between the lung and the diaphragm. Consequently, it is significantly less susceptible to relative variation in signal intensities.

A single ROI is selected over the centre of the diaphragm, encompassing its full range of motion during the scans, from just above its position at full exhalation to just below its position at maximum inhalation during the acquisition. Selecting the width of this ROI is a trade-off between accuracy and SNR. A wide slice will result in a greater SNR simply by including more voxels in the calculation. However, a healthy diaphragm is naturally curved; therefore, its position will vary within the ROI and by a greater magnitude with increasing segment width. This will result in an averaging of the diaphragm position and effectively reducing its apparent range of motion. For the greatest accuracy, this ROI placement is done

manually and adjusted for optimal performance. However, this process can be assisted or automated using the output of the neural network detailed in section 4.9.2. If applied successfully, this automated method will position the ROI to encompass the centre of the diaphragm-lung boundary of the right lung for each frame in the image sequence.

This segment, designated by the ROI, is extracted for each frame and used to determine the diaphragm position. First, each segment is subjected to vertical edge detection using the kernel (k) displayed, adapted to suit the width of the segment:

$$k = \begin{pmatrix} -1 & -1 & -1 \\ -1 & -1 & -1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

This kernel is oriented such as to amplify vertical changes in voxel intensity. This processing creates a one-dimensional vector corresponding to the average vertical change in signal intensity across the segment, as shown in Figure 4.3. A low to high signal boundary, characteristic of the lung-diaphragm interface, will result in a positive result. A high to low signal boundary, typical of the diaphragm-stomach interface for example, will result in a negative result. A centre of mass calculation is taken of the three points about the maxima, to estimate the sub-pixel location of the lung-diaphragm boundary. Omitting this step would result in a discretised position determination, where the diaphragm position is assigned to the nearest pixel. This would reduce the accuracy of the algorithm, a problem exacerbated for lower resolution acquisitions where the diaphragm may only move a few pixels.

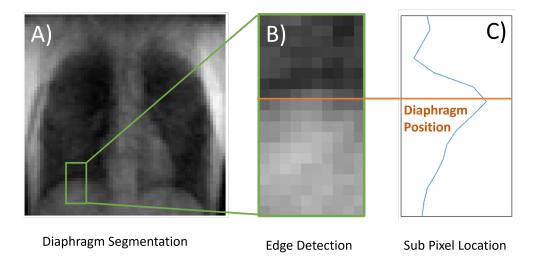


Figure 4.3: Method of diaphragm detection using edge detection. A ROI is selected encompassing the diaphragm in each image. One-dimensional edge detection is then applied to this region, followed by a centre of mass calculation about the most maximum point.

This method was found to be visually far more robust than simply using the average intensity of the ROI. This, therefore, was the desired procedure for determining diaphragm position going forward and was implemented into the PREFUL pipeline. The filtered signal of the diaphragm position, determined by edge detection, is presented in Figure 4.4Bii.

4.5.1.1.3 2D Lung Area

As a healthy individual inhales, the lungs will expand in volume. When imaged with a 2D slice, this will be observed as the 2D area of the lung increasing. Likewise, exhalation will be observed as a reduction in the 2D lung area. Therefore, if the lung is segmented for each image, this property can be used as a marker for the respiration phase.

Similar to detecting the diaphragm position by edge detection, this method determines the position in the ventilation cycle by image features, rather than signal intensity. Therefore, it is less susceptible to contributions from the cardiac cycle. Additionally, this method circumvents pitfalls of using the diaphragm position, such as participant movement in the FOV and only considering a single lung. However, this method requires segmentation of the lungs in each frame of the acquisition. This process is discussed in Section 4.9.2; however, if done manually, it would be highly time consuming. Furthermore, variation in segmentation between frames will lead to errors in the determined respiratory phase. Therefore, for widespread application, this process should employ an automatic segmentation method. This can be provided by a convolutional neural network, also discussed in that section. The number of voxels included in each mask can be used as a respiratory marker. The filtered signal of the 2D lung area is presented in Figure 4.4Cii.

This method was visually found to be very robust, although it was occasionally exposed to errors associated with the automatic segmentation process. Furthermore, as discussed in Section 4.9.2, this segmentation method will only work when trained on similar data, such as field strength, acquisition properties and differences due to healthy vs pathological tissue. However, this method was also implemented in the PREFUL pipeline.

4.5.1.1.4 Zero-Frequency (Navigator Pulse)

The final software based method investigated as a marker for the respiratory phase was the intensity of the zero-frequency peak in k-space. This peak is related to the average intensity of the image. At maximal inhalation, the average intensity of the image as a whole is expected to be reduced as the low intensity lung parenchyma takes up an increased portion of the FOV. Likewise, at maximal exhalation, the low intensity lung parenchyma will be displaced by the higher intensity abdominal organs, thus the average

intensity of the image will increase. Therefore, this property can be used as a marker for the respiratory phase.

In this work, this metric as simulated retrospectively by applying a 2D Fourier transform to each image in the image sequence. The average intensity of the central 3x3 pixel grid in k-space, filtered along the time sequence, is presented in Figure 4.4Dii. The plot created by this method shares similarities to that of the diaphragm position, determined by signal intensity. However, this variation appears slightly noisier. Therefore, this method may not be accurate enough to be used as the ventilation marker, when alternatives are available. An additional observation is the increased value of the first images in the series. This is likely because the spin ensemble has not reached a steady state in the 2D GRE sequence used to acquire this data.

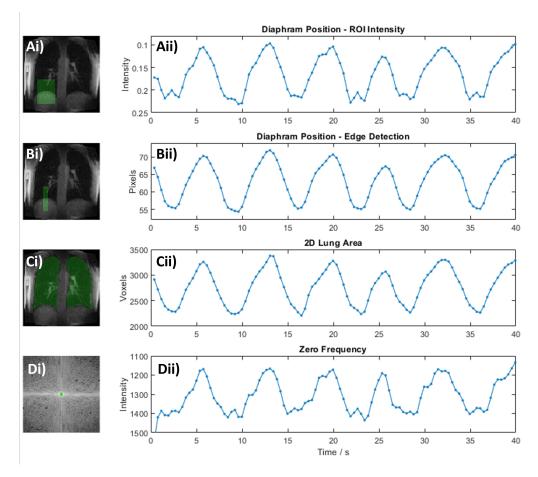


Figure 4.4: Demonstration of discussed methods of respiratory phase tracking, including diaphragm position, detected by signal intensity (A) and edge detection (B), 2D lung area (C) and zero-frequency intensity (D). Each is presented with an example image of the process (i), with the effective ROI used for analysis overlaid in green. These are depicted alongside the resulting variation in the marker (ii) during 40 seconds of deep breathing respiration. Note, the size of the green overlay for the zero frequency has been exaggerated to be visible in this plot.

Intensity-based methods for determining the respiratory marker are more sensitive to contributions from the cardiac cycle, compared to the region-based method. However, a marker determined using all of the above methods should be filtered to remove contributions from other sources of noise. This is achieved using the ventilation filtering method described in section 4.7.

4.5.1.2 Ventilation Phase Assignment

When a marker for the respiratory cycle has been assigned to each frame, the associated ventilation phase can be determined. However, the respiratory cycle is prone to significant variation in amplitude and thus frames of the same cycle phase may not necessarily share the same lung state. For example, the lung at 50% of the respiratory cycle will be significantly larger when deep breathing compared to tidal breathing, even though these states share the same cycle phase. Therefore, the respiratory phase is not generally equivalent to the assigned ventilation phase. Instead, the ventilation phase is calculated for each frame directly from the assigned marker [25]. This marker for the respiratory cycle is assigned the value (R), and is described for each frame using the equation [25]:

$$R_i(t) = A\cos(2\pi f t) + K \tag{4.10}$$

where A is the amplitude range, f is the respiratory frequency, K the offset and i the associated frame number. Therefore, ft represents the artificial ventilation phase Φ^V . Frames corresponding to inhalation or exhalation may share the same respiratory marker, whilst differing in phase. Therefore, before Equation 4.10 can be applied, the frames must be separated into two groups depending on if they correspond to inhalation (R_{in}) or exhalation (R_{ex}) . This is achieved for each frame using the relative value of the respiratory markers of the frames prior to and preceding it. When 2D lung area is employed as the respiratory marker and one assumes that the cycle is symmetrically sinusoidal, the following rules applies:

$$R_{i-1} > R_{i+1}$$
 then $R_i \in \text{Exhalation}$

$$R_{i-1} < R_{i+1}$$
 then $R_i \in Inhalation$

where *i* denotes the image number in the acquisition. This rule must be amended for the scenario where both markers are equal, in which case the frame type is equal to that of the previous frame. The application of this rule is shown in Figure 4.5. Note, this rule applies to respiratory markers where inhalation results in a increasing marker value, such as for 2D lung area. For markers where the opposite is true, such as intensity based approaches, then the inverse rule applies. Furthermore, as this rule can not be applied to the first of last images in the acquisition sequence, these are omitted.

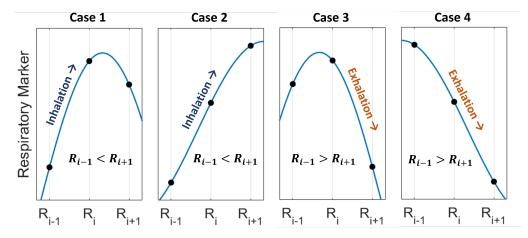


Figure 4.5: Demonstration of the assignment rules to determine if each frame belongs to the inhalation or exhalation regime. The **blue line** represents an ideal continuous variation in the respiratory marker, where the **black dots** represent temporally adjacent frames on this curve. Cases 1 and 2 both satisfy the rule: $R_{i-1} < R_{i+1}$, therefore belong to the inhalation regime. Cases 3 and 4 each satisfy the rule $R_{i-1} > R_{i+1}$, therefore belong to the exhalation regime. An additional 5th case is omitted, for the rare occurrence that: $R_{i-1} = R_{i+1}$.

With the successful assignment of images in the acquisition cycle to either the inhalation or exhalation regime, Equation 4.10 can be adapted for either condition to calculate the artificial ventilation phase (Φ^V) for each frame[25]:

$$\Phi_{Inhalation}^{V}(R) = \frac{\cos^{-1}\left(\frac{R_{in}(t) - K}{A}\right)}{2\pi}$$
(4.11)

and

$$\Phi_{Exhalation}^{V}(R) = 1 - \frac{\cos^{-1}\left(\frac{R_{in}(t) - K}{A}\right)}{2\pi}$$
(4.12)

The amplitude (A) is calculated based on half of the difference between the maximum and minimum values of the respiratory markers associated with the image series. The offset (K) is the midpoint between these points. Using these equations, an artificial ventilation phase is assigned to each image. This process is summarised in Figure 4.6A, with the variation in respiratory marker (in this case lung volume) presented in Figure 4.6Ai, with each frame designated to either inhalation or exhalation. The corresponding visualisation of the ventilation phase assigned to each image is presented in Figure 4.6Aiii. This process uses a direct calculation to assign the respiratory phase, therefore there is no variation from the artificial sinusoidal distribution. The images located exactly at the artificial 0 and π phase points are those from which the amplitude in Equations 4.11 and 4.12 was determined.

4.5.1.2.1 Thresholding

When asked to free-breathe, the subject will often exhibit a large variation in the amplitude of the respiratory cycle. For example, a healthy individual may take many tidal breaths, interspersed with occasional deeper breaths. As the ventilation phase is determined directly from the maximum and minimum values of the respiratory marker, these larger breaths contribute

to the full range of the amplitude (Equations 4.11 and 4.12). This often results in the undersampling of much of the resulting artificial ventilation cycle as seen in Figure 4.6Aiii, particularly at the points of maximum exhalation and inhalation, due to the reduced number of respiratory cycles of this amplitude. This inequality can be quantitatively measured using the number of images contributing to each frame, which is dependent on the number of frames and the interpolation kernel bandwidth, described in Section 4.8. The resulting distribution for the respiratory cycle presented in Figure 4.6A is plotted in Figure 4.6Aiv. This sampling inequality may result in a significantly reduced SNR of the ventilation cycle corresponding to the points of maximal inhalation and exhalation, compared to the frames in-between.

One method to reduce this effect is to threshold the respiratory marker. This will omit images with the greatest marker amplitude from the phase determination calculation and ensure a more temporally uniform distribution of images contributes to the reconstructed sequence. This process can be completed manually, by moving the position of each threshold and visually inspecting the resulting reconstruction quality. When the respiratory marker is thresholded, all images outside of the bounds are discounted from the reconstruction process (Section 4.8), with the amplitude and offset terms in Equations 4.11 and 4.12 calculated based on half of the difference between the bounds and the midpoint between them, respectively. The information presented to the user to assist with the manual positioning of the thresholds is presented in Figure 4.6, with Figure 4.6A representing a pre-thresholded and Figure 4.6B representing a thresholded state. The reconstruction quality metrics (Figure 4.6iv) are dependent on the number of frames and interpolation bandwidth chosen in the reconstruction step (Section 4.8), each of which are kept constant for both parts of Figure 4.6.

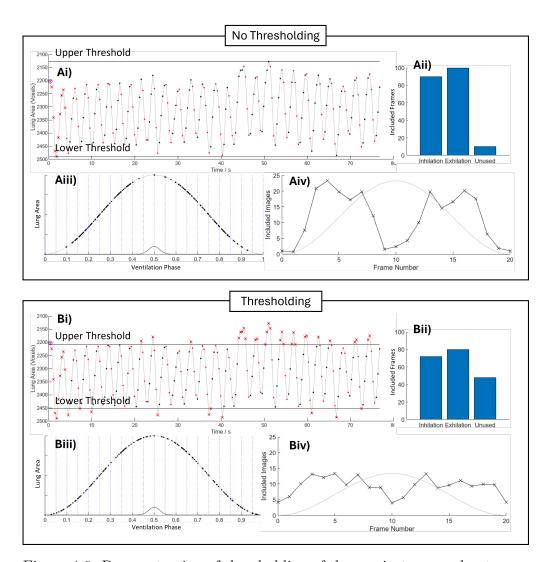


Figure 4.6: Demonstration of thresholding of the respiratory marker to create a more temporally uniform distribution of frames along the artificially reconstructed ventilation cycle. Plots are presented for the un-thresholded (A) and thresholded (B) case. This includes the variation in respiratory marker with time (i), where images determined to correspond to exhalation are marked with black dots, and inhalation with red dots. Images are discounted either due to being outside of the thresholded bounds, or at the start of the imaging sequence and are denoted as red x's. The total number of images in each of these groups is used to populate a bar chart (ii). The temporal distribution of frames along the artificial ventilation cycle is plotted visually (iii) onto a sinusoidal curve, where each dot represents an image and each vertical dashed line represents a reconstructed frame, and quantitatively (iv), as the number of images averaged to create each reconstructed frame using the interpolation detailed in section 4.8.

This thresholding process can also be automated. The aim of the thresholding step is to increase the relative number of images used to reconstruct the most inhaled and exhaled frames; this is achieved by moving the threshold bounds towards the mean respiratory metric value; however, this will also result in a reduced number of frames included in the reconstruction of the full cycle. These properties can be quantified and used to calculate a Quality metric (Q) for each position of the thresholding bounds. To determine the position of the upper threshold (representing the inhaled state), this metric is calculated using:

$$Q(b) = \frac{\sum_{k=\frac{n}{2}-1}^{\frac{n}{2}+1} b_k}{\sum_{k=1}^{n} b_k}$$
(4.13)

where (k) denotes each frame, (n) denotes the total number of frames and (b) the total (weighted) number of images used to reconstruct each frame, calculated using the denominator of Equation 4.19. The numerator of this equation represents the number of images used to reconstruct the three frames corresponding to the most inhaled state. This is adapted for the exhaled state when calculating the optimal position of the lower bound. The denominator represents the total number of images used to construct the full cycle. Each threshold (upper and lower) is stepped independently between the respective maximum and the central respiratory marker position and the quality metric calculated for each of these steps. The resulting variation of this metric for a typical acquisition is plotted in Figure 4.6. The point corresponding to the maximum quality metric for each bound is then determined to be the optimal threshold position. Figure 4.6 depicts this process for the data presented in Figure 4.6 and the resulting threshold position is presented in Figure 4.6B.

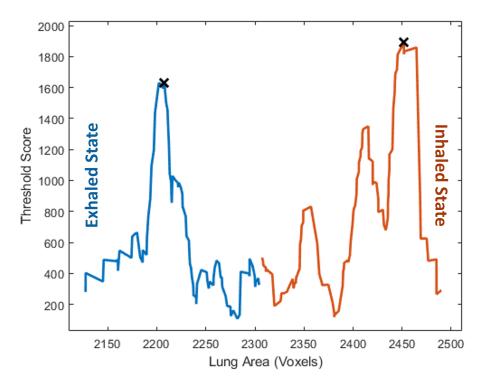


Figure 4.7: Variation of the calculated Quality metric with movement of the respiratory thresholds, in this case using 2D lung area as the respiratory marker. Movement of the inhaled threshold is depicted in **orange**, and exhaled in **blue**. The determined most optimal positions of each of these bounds are denoted as **black x's**.

This process of ventilation phase assignment assumes that images of the same respiratory marker (and of consistent inhalation/exhalation regime) will share the same lung state. In other words, a voxel will possess the same parenchymal density across images with the same assigned ventilation phase. However, due to the high degree of variation in the amplitude of the respiratory cycle, and the known hysteresis in the pressure-volume curve [65] associated with this process, this assumption is likely false. For example, the lung at maximal inhalation during tidal breathing may share a similar respiratory marker value with a lung that is partially inhaled during deep breathing. Images acquired in these different physiological states may therefore exhibit significantly different physiological properties. This is not accounted for in this work, and further investigation is required

into the magnitude of this potential error. Furthermore, as previously discussed, there may be a decoupling of respiratory marker and lung state depending on the choice of marker. This would create further errors in the determination of ventilation phase.

4.5.2 Perfusion

4.5.2.0.1 Physical Methods

Electrocardiogram (ECG) ECG is widely used for cardiac gating in imaging studies of the heart [66, 67]. This method employs electrical signals from the heart, detected using electrodes attached to the skin, to determine cardiac phase. Using this information, the position in the cardiac cycle can be determined for either prospective or retrospective gating. However, placing the electrodes on the patient can be time consuming, and considerations must be made for Magnetohydrodynamic (MHD) effects at high field [68, 69].

Peripheral Pulse Detection An alternative method for indirectly determining the cardiac phase is by detecting peripheral pulse waves[70]. This is usually achieved by placing a detector on the patient's finger or toe. However, although this method doesn't suffer from the same MHD effects at high field, it is less preferred than ECG due to the long and unpredictable delay between myocardial contraction and arrival of the pulse wave to the detector[71].

4.5.2.0.2 Software Methods

Heart Analysis If the MRI slice position captures the heart, analysis of that organ could be used to generate a marker for the cardiac phase. Analogous to diaphragm position as a marker for ventilation phase, this technique would rely on structural image information, rather than absolute voxel intensity. And it could therefore be susceptible to less noise. At the simplest level, this could be done by tracking the position of the edge of the heart; it will expand and contract as the heart beats. However, a potentially more robust method would be to train a neural network to recognise the state of the heart in the context of the complete image stack[72, 73]. Therefore, to be able to determine each frame's point in the cardiac cycle. This technique has not been explored in this work.

ROI Intensity Although simple to implement, this method is not without its drawbacks. Similar to the ROI intensity method for determining the respiratory cycle, this marker is highly susceptible to contributions beyond that of the cardiac cycle. These contributions come primarily from the respiratory cycle, as the tissue within the ROI can move around as the subject breathes. Therefore, the time-series of this marker must be filtered to remove these contributions. Errors of this nature may also be reduced by determining this marker using previously registered frames. Furthermore, this method is susceptible to intensity variations due to noise from the MRI acquisition, which is amplified if the ROI is smaller. Finally, any movement of the subject can cause this method to break down. Changes in intensity due to the tissue moving within the slice should be somewhat countered during filtering. However, should the tissue leave the imaging plane

entirely, this method may no longer work. Modern PREFUL methodology uses an iterative-based approach for optimising this ROI[39, 74]. The standard deviation of each voxel intensity along the time series is used to determine those which exhibit the greatest variation. These are assumed to correspond to the central lung vessels and thus those with the greatest variation are included in the ROI. The number of voxels included is steadily increased, until the quality of the subsequent perfusion phase assignment ceases to improve.

4.5.2.1 Determining the Cardiac Marker

4.5.2.1.1 Aorta Intensity

To create a marker for the cardiac cycle, an ROI is placed manually over the aorta. This process is assisted using the automated segmentation network detailed in section 4.9.2. The region between the segmented lungs is assumed to contain the aorta, therefore this region is isolated using morphological operations. However, this process often benefits from manual segmentation of these vessels. Manual refinement is assisted through the visualisation of the greatest frequency components for each voxel, as shown in Figure 4.8, which gives good contrast between the aorta and the surrounding tissue. This allows the user to see which part of the anatomy has the greatest signal variation within the expected range for the cardiac cycle, and thus include this in the ROI. This approach both increases the SNR by omitting weaker contributions from the cardiac cycle and addresses cases where the aorta may not be immediately obvious in some MRI data.

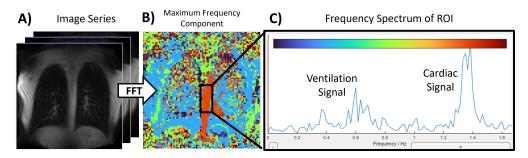


Figure 4.8: Demonstration of the frequency analysis method used to assist in the ROI selection for cardiac phase assignment. Including the input image series (A), the associated voxel-wise maximum component analysis of frequency space (B) and the frequency spectrum of the displayed ROI (C).

The marker for the cardiac cycle varies at a greater frequency than that of the respiratory cycle. Therefore, for the imaging rates typical at 0.5 T, it is more useful to visualise the variation in the perfusion signal in the frequency domain, rather than a time-series. Generally, the contribution from the cardiac cycle should form a distinctive peak, with an SNR ideally greater than two, although this is not always the case, as variation of the heart rate throughout the acquisition sequence will result in a broadening of this peak. At this point, this plot should be used for quality assurance, to iterate the ROI to resolve a greater SNR of the cardiac cycle.

This cardiac phase determination method is based on ROI intensity; therefore, it may experience contributions from the respiratory cycle. To reduce this effect, the variation in the cardiac marker is filtered using the using the cardiac filtering method described in Section 4.7.

4.5.2.2 Perfusion Phase Assignment

Unlike for the respiratory cycle, the cardiac cycle is less prone to variations in amplitude. Therefore, the determined cardiac phase for each frame should adequately correspond to the reconstructed perfusion phase.

The signal variation is assumed to be sinusoidal, thus permitting phase determination via fitting to a sine curve. However, the heart rate can vary throughout the acquisition sequence and thus fitting the whole time-series at once will result in a compounding error in the determined phase. This is demonstrated in the work by Voskrebenzev et al.[25] Instead, the fitting is done piecewise on smaller sections of the time-series. The number of parts is user selected, with changes in this parameter having a profound effect on the fit quality. It is not always immediately clear why the number of fitting parts has such an effect on the total fit quality, although as a rule of thumb:

- A greater number of parts results in a lower number of points per fit,
 which reduces the quality of each fit but responds to changes in heart
 rate well. This approach is good for acquisitions where the heart rate
 varies a lot but is more sensitive to noise. Generally, it results in
 slower processing times.
- A lower number of parts results in a greater number of points per fit.
 This approach increases the quality of each fit as it is less sensitive to noise. However, it does not respond well to changes in heart rate.
 This method is good for acquisitions where the heart rate doesn't change much. Generally, this method results in faster processing times.

In general, one should aim for a lower number of parts and increase as necessary to best fit the data. When the optimum number of parts has been selected, the average number of cardiac markers per fit is equal to the total number of included frames divided by the number of parts. However, these parts are not of constant size throughout the time series, as the fitting algorithm sizes them such that they start and end at local maxima

of the cardiac marker[25]. This may improve the accuracy of the fitted parameters. Each part can then be fitted to the equation:

$$C_n(t) = A_n \cos(2\pi f_n t + \varphi_n) + K_n \tag{4.14}$$

where C is the cardiac marker, n denotes the fitting part, A denotes the amplitude, f denotes the frequency, φ denotes the phase offset and K denotes the marker offset. Therefore, the fitted perfusion phase (Ψ^P) of each frame can be calculated according to:

$$\Psi^P = 2\pi f_n t + \varphi_n \tag{4.15}$$

To assist with the fitting algorithm, the variation in the cardiac marker within each part is pre-processed. As mentioned previously, the data is filtered to reduce the contribution from both the cardiac cycle and movement of the subject. This filtering is commonly achieved via the application of a high pass filter above the contribution from the respiratory cycle (details in Section 4.7). For acquisitions proving difficult to fit, this can also be done with a band pass filter, however, one must be cautious against over-filtering the data. If this happens, the fit may seem exemplary, but at penalty to the validity of the final result. All filtering is done using a zero-phase infinite impulse response filter. Following filtering, each part is normalised independently via two steps:

- 1. The mean marker value of each part is subtracted.
- 2. Each part is scaled such that its average RMS value is equal to $2/\pi$.

Each of the normalisation steps assumes there are sufficient cycles within

each part to reduce the effect of outliers and partially sampled cycles. They have the effect of suppressing the value K such that it can be omitted as a fitting parameter, as well as setting the amplitude to approximately 1 to assist the fitting algorithm.

When the data is prepared, each part (n) is fitted to a sine curve with fitting parameters: amplitude (A), frequency (f) and phase offset (φ) . To further assist with fitting, starting values and limits are set for the parameters. The amplitude has been normalised; therefore, its starting value is set to one although this is not bounded. The phase offset is started at π , but not limited. The frequency start point is selected based on the greatest frequency contribution from the cardiac cycle. This start point is determined by the greatest frequency contribution to the variation of the cardiac marker within the pass band of the high-pass filter. This can also be manually selected if desired. The fitted frequency parameter is limited to ± 0.25 Hz (15 bpm) of the start frequency. This range is denoted as (Δf_{fit}) . Each part is fitted using a non-linear least squares regression [75–77]. This algorithm returns the R² value to be used as a fit quality measure. Any part with an R² lower than 0.7 is omitted from future processing. This value can be adapted in the user interface. Then, using Equation 4.15 with the fitting results for each part, the cardiac phase (and therefore perfusion phase) can be calculated for every image. Finally, this phase is normalised to between 0 and 2π .

The fitting process is iterative. When completed manually, the user should gradually increase the number of fitting parts until a maximal number of images are assigned a phase with sufficient R^2 to be included in the perfusion reconstruction. This process is accelerated by multithreading the fitting of each part independently. Furthermore, the iterative process can be automated, with the number of fitted parts increased sequentially

to a predetermined maximum. The number of parts for which the greatest number images exhibit sufficient R^2 , is used for the subsequent processing. Furthermore, if a number of parts returns all images of sufficient R^2 , then the iteration is aborted and this value is used.

4.5.2.2.1 Cardiac Aliasing

It is generally accepted that to adequately sample a time-varying sinusoidal signal, one must do so with at least double the sampling frequency of said signal. This is referred to as satisfying the Nyquist condition. The resting heart-rate of a healthy individual is expected to be between 60 and 100 bpm (1 to 1.67 Hz). This would require a sampling frequency of at least 3.3 Hz to satisfy the Nyquist frequency for the full range of expected rates. Therefore, established PREFUL methodology will often acquire images at a rate between 3-5 Hz to adequately sample variation due to the cardiac cycle. However, the lower sensitivity and longer T2* of lung parenchyma when investigated at lower field (Table 4.1) often results in longer imaging times. If the cardiac cycle is undersampled, then this variation will appear as an aliased signal in the frequency domain of the cardiac marker. This is demonstrated in Figure 4.9, where the aliased signal folds linearly back onto the lower frequencies, effectively mirrored at the Nyquist frequency. In isolation, this signal is indistinguishable from a lower frequency contribution. However, if the expected cardiac frequency is roughly known, for example by external measurement, then the correct cardiac frequency can be determined from this spectrum. This can subsequently be used to constrain the fitting algorithm to the appropriate frequency. However, this method comes with its own caveats:

• The heart rate must be measured either before, after or ideally during

the scan to select the correct frequency for fitting.

• If the acquisition period encompasses the full time between images, the apparent amplitude of the cardiac signal will decrease as the cardiac frequency increases. This is due to the averaging of a larger portion of the cycle and is demonstrated for a simple model system in Figure 4.9C. Therefore, it is unfavourable to scan at a rate such that the signal aliases further than halfway back through the frequency domain, or:

$$f_{cardiac} > \frac{3}{2} f_{Nyquist}$$
 (4.16)

where $f_{cardiac}$ is the cardiac frequency and $f_{Nyquist}$ is the Nyquist frequency determined by the sampling rate.

- When sampling at a lower rate, the cardiac frequency may alias near
 or onto the respiratory frequency. This would make the signals difficult or impossible to decouple via filtering.
- When the cardiac frequency is close to the Nyquist frequency, it can be hard to determine if the signal is aliased or not. Moreover, with sufficient variation of the cardiac frequency, the signal may only be aliased for some of the acquisition. Additionally, cardiac signal may vary, fluctuating above and below the Nyquist frequency at different times during the same acquisition. Furthermore, for cardiac frequencies near the Nyquist frequency, the restriction bounds on the fitted frequency may encompass both the real and aliased peak positions. This may create errors in the phase determination step as the fitting algorithm could fit to either frequency. To prevent this, the following rule should apply:

$$|f_{cardiac} - f_{Nyquist}| > \frac{\Delta f_{fit}}{2}$$
 (4.17)

Therefore, it may be possible to recover perfusion information from an undersampled time-series. Using the Equation 4.16, the minimum sampling frequency required to sample the full resting heart-rate range is approximately 2.2 Hz. However, this could result in a reduction of approximately 40% of the apparent amplitude of the cardiac signal.

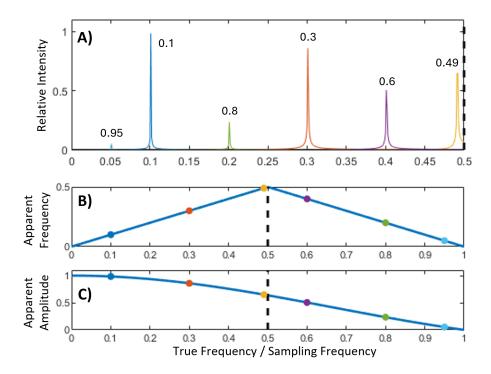


Figure 4.9: Demonstration of the effect of aliasing on the frequency domain of the cardiac marker. The effect was modelled using a constant sampling frequency to sample artificial sinusoidal functions with varying frequency. The resultant frequency domain of the signal is presented in **A**, for 6 example frequencies. The x axis is scaled relative to the sampling frequency. This is broken down to demonstrate the effect of changing sinusoidal frequency on the apparent frequency **B**, and the apparent amplitude **C**. In each case, the Nyquist frequency is indicated as a **black dashed line**. The amplitude variation is based on the assumption that signal is acquired uniformly throughout the acquisition window, analogous to the shutter-speed of an optical camera. However, MRI acquisitions do not function this way.

If the cardiac frequency is determined incorrectly, for example by assuming the signal is aliased when it is not, then there may be errors in the generated perfusion maps. These errors are explored in Figure 4.10. For the perfusion weighted maps, the resultant errors are negligible, as this map is based on the amplitude of signal variation due to the cardiac cycle. However, this error can have a significant impact on the QTTP maps, as these are based on the phase of signal variation with respect to the cardiac cycle. This effect can be simulated using an undersampled sinusoidal signal. Using Equation 4.15, the determined erroneous calculated phase for the time-points appears to have been reflected about π and shifted π degrees out of phase. This distribution is plotted in Figure 4.10Biii. Therefore, while the perfusion weighted maps may only suffer from a reduced SNR due to undersampling, the QTTP maps present a significant possibility for errors. Therefore, great care should be taken in the assignment of cardiac frequency during the PREFUL pipeline.

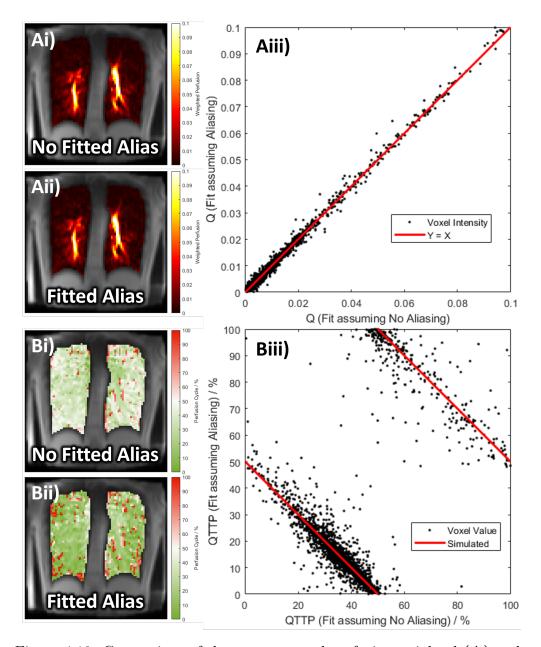


Figure 4.10: Comparison of the reconstructed perfusion weighted (A) and QTTP (B) maps of a typical healthy participant, when the fitting algorithm is constrained to an aliased and non-aliased frequency domain peak. The reconstructed maps are displayed for fitting to the non-aliased (i) and aliased (ii) peak. These are compared on a voxel-wise basis (iii) to display the potential errors associated with incorrect frequency assignment.

Finally, considering the distribution of QTTP values in Figure 4.10Biii, there may be a physiological mechanism for determining if the correct frequency was chosen based on the cardiac peak in the frequency spectrum. To aid in visualisation, histograms of the variation of QTTP values are

presented in Figure 4.11. Due to the reconstruction method employed, the peak aorta signal occurs at 25% of the cardiac cycle. Then physiologically, one would expect healthy lung parenchyma to reach peak perfusion shortly afterwards. Highlighted in Figure 4.11, for this acquisition, this occurs when the non-aliased peak has been used for fitting. Therefore, this may be the correctly processed result. Further investigation with additional methods of cardiac frequency determination are required to confirm this hypothesis, which may be affected by pathology. At this time, this quality control step has not been implemented in the PREFUL pipeline.

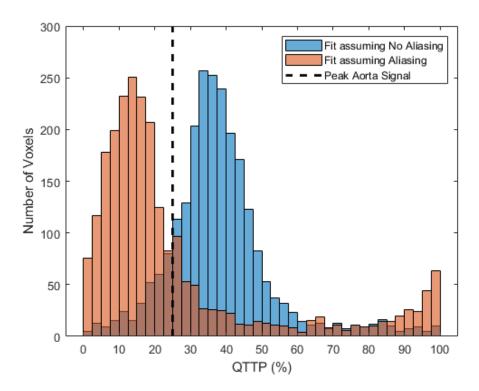


Figure 4.11: Distribution of reconstructed QTTP values when fitted to the cardiac signal assuming no aliasing (blue) and aliasing (orange). The time of peak aorta signal is depicted as a black dashed line.

4.6 Registration

4.6.1 Background

Core to FD based methods is the accurate registration of the lungs between varying respiratory states, such that the parenchyma appears stationary throughout the cycle[7]. Image registration can be thought of as an optimisation problem: for registration between two images, one must determine the optimal transformation which yields the lowest difference between the registered and target image. Biomedical registration methods can therefore be characterised based on the method for calculating this difference. Typically, techniques will either use the intensity of the voxels, or the similarity between a predetermined feature map. These methods are therefore termed intensity-based registration and feature-based registration, respectively.

The above example concerns registration between two images, where one serves as the target while the other is transformed. This constitutes another characterisation of image registration methods, pairwise or groupwise. Pairwise registration involves just two images, where one is transformed to match the other. Groupwise registration can involve multiple images, where all images are transformed to a calculated intermediary position between them.

The transformation function can include multiple components. These transformations can be applied globally to the entire image, or locally to specific regions. The transformation types with increasing complexity are:

Translation Voxel shift.

Rigid Voxel shift and rotation.

Similarity Voxel shift, rotation and scale.

Affine Voxel shift, rotation, scale and sheer.

Additional to these methods, is deformable registration (non-rigid registration). In this technique, local transformations are applied specifically to certain voxels.

4.6.2 Lung MR Registration

Lung registration techniques for images acquired with MRI typically use non-rigid intensity-based techniques. A number of toolkits are available for this purpose, including the Advanced Normalisation Toolkit (ANTs)[78], Elastix[79], NiftyReg[80] and the commercial software fMRLung 3.0 (Siemens, USA). For implementation solely in MATLAB, two further options are available: the inbuilt MATLAB Demons-based registration function[81, 82] and the Forsberg[83] toolkit.

4.6.2.1 GOREG

Established PREFUL methodology typically employs a two-stage approach to pulmonary image registration. Instead of registering all frames directly to an intermediary position of the respiratory cycle, Group-Oriented Registration (GOREG) registers each image to predetermined frames positioned at equidistant phase points of the respiratory cycle[84]. These registered frames are then averaged to form discrete groups. Intergroup pairwise registration then occurs, starting with the groups at start and end expiration points to the adjacent group in the respiration cycle, in the direction of the centre respiratory position. This process is repeated for all groups, to

calculate the stepwise transformation fields towards the centre respiratory position. Finally, the resulting deformation fields are used to register each frame individually to this position. All registration steps in GOREG employ non-rigid intensity-based pairwise registration. This method has been employed for 2D[84] and 3D[85] datasets.

4.6.3 Typical Errors of MR Pulmonary Registration

Beyond incorrect assignment of registration parameters, errors in pulmonary MR registration at low field appear to stem from three sources:

1. Low SNR of lung parenchyma, particularly at 0.5 T.

This can result in few real features internal to the lungs for the intensity-based algorithm to use, which would be typical of images captured using other modalities such as CT. Therefore, these algorithms may try to perform registration based on the noise within the void regions. Ultimately this would result in inaccurate results, although this may not be immediately obvious to the operator due to the aforementioned lack of visual features in these regions. For each successive pairwise registration, these errors can compound from noise associated with the frame being deformed and the target frame. The prevalence of these errors can be further exacerbated with the low SNRs associated with imaging at low field. The GOREG pipeline averages the frames together to create groups. These groups will exhibit an increased SNR compared to the input images. However, this process precedes the pairwise registration of each frame to its respective group's target frame and can therefore be sensitive to the errors discussed above.

2. Failure of non-rigid registration over large displacements.

Non-rigid deformable registration algorithms can fail when large voxel displacements are required between lung states, such as between fully exhaled and fully inhaled lungs. This may be corrected with the optimisation of parameters associated with registration, although acquisition specific optimisation may hamper universal application of these tools. The GOREG pipeline accounts for this source of error by repeatedly applying non-rigid registration only between frames of a similar respiratory state. Therefore, large voxel displacements are typically not required. However, this process can create compounding errors if they occur early in the stepwise process.

3. Varying signal between frames of the same respiratory phase.

Intensity-based registration algorithms utilise the same signal variations for registration as are employed for functional analysis. They assume that variations in signal intensity, used for respiratory registration, stem purely from the respiratory cycle. However, regions of parenchyma may exhibit significant variation in MR signal intensity throughout the cardiac cycle. Pairwise intensity-based registration algorithms, which have no context of the dynamics of the lung, may attempt to 'correct' for this variation. Thus, creating errors in the resulting deformation maps between images captured at differing points of the cardiac cycle. This is visually most evident in the central lung vessels, which can vary in intensity significantly throughout this cycle. Often it is possible to distinguish the cardiac pulse wave transiting the largest vessels. These dynamics may be erroneously registered by a pairwise intensity-based registration algorithm, which will result in these vessels moving around in the reconstructed perfusion cycle. One might assume cardiac contributions can be removed prior to registration via filtering. However, due to the movement of the lung, the registration step is a prerequisite of filtering.

This work presents a registration pipeline that employs a number of techniques to account for these sources of error. This includes two key methods implementing dual-stage registration, and batch-wise registration.

4.6.4 Dual-Stage Registration Pipeline

This work utilises a dual-stage approach to lung registration to overcome the second error typical of pulmonary MR registration. In the first stage, the bulk voxel displacement is crudely performed using feature-based registration. This is followed by a non-rigid intensity-based registration algorithm for displacements on the fine scale. This method partially accounts for the first and third error associated with pulmonary registration techniques, as feature-based registration algorithms can be less sensitive to low SNR and cardiac signal variations than intensity-based techniques.

4.6.4.1 Feature Based Registration Methods

This work employs two methods of region-based registration, each based on features previously determined in this PREFUL pipeline: the diaphragm position and 2D lung region. The selection of the feature-based method for registration is dependent on the accuracy of the determined features. For example, in low-quality datasets where accurate segmentation of each frame was unsuccessful, the diaphragm position method can be used.

4.6.4.1.1 Linear Deformation based on Diaphragm Position

The simplest of these methods is a linear deformation, based on the position of the diaphragm within each image. This technique is therefore only possible in imaging planes where the diaphragm is visible. This method assumes that most of the movement in the lung during respiration is along the apical-to-basal axis, thus respiratory motion can largely be nullified by linearly deforming each image such that the diaphragm appears stationary throughout the image series. This transformation is applied to a region determined by an 'anchor point' specific to the image series and the image-specific sub-pixel diaphragm position. The anchor point is usually placed just above the hilum, either manually or based on the automated segmentation algorithm. The process of this linear deformation registration step is summarised in Figure 4.12.

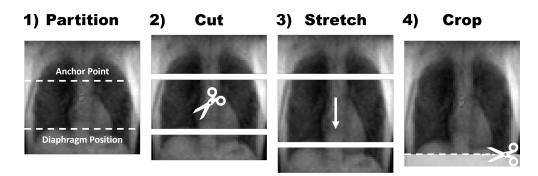


Figure 4.12: Demonstration of the feature-based registration method using the diaphragm position. The process is depicted for an example frame, at the start of inspiration position. Therefore, the linear deformation resembles a stretch. The final stage constitutes a crop of the image to maintain image matrix size. For the case of linear compression, this region is padded with zeros.

4.6.4.1.2 Affine Transinformation based on 2D Lung Region

Alternatively, the feature-based registration can employ the segmented mask of each image, generated using the trained CNN. An affine transformation is applied to each lung independently so as to register the mask of each image to a target mask. The resulting deformation field is then applied to the associated image.

4.6.4.2 Intensity based Registration Methods

Two intensity-based registration algorithms are integrated into this PRE-FUL pipeline: the MATLAB demons-based function[81, 82] and the Forsberg algorithm. Each of these functions applies pairwise non-rigid intensity-based registration[83].

4.6.5 Batch-wise Registration

A batch-wise registration algorithm was developed to account for the first and third errors typically associated with pulmonary MR image registration. In contrast to GOREG, averaging occurs prior to the non-rigid registration step; thus, this algorithm is not applied to individual images.

Using the interpolation method detailed in Section 4.8, each frame is interpolated based on its assigned respiratory phase, to create 20 equidistant respiratory phase 'batches'. Depending on the number of frames contributing to each batch, this process should increase the SNR of each batch compared to the input data. Furthermore, this method can reduce the impact of the cardiac cycle on the registration algorithm. If the interpolated frames are assumed to be distributed evenly throughout the cardiac cycle, then the time varying signal from these contributions should be minimised during averaging. This assumption becomes more valid the more images are included in each interpolated batch.

Following construction of the respiratory batches, non-rigid intensity-based registration is applied in a pairwise fashion from each batch to a target batch, usually at the intermediary respiratory position. The resulting deformation fields for each transformation are then stored to determine the transformation required of each frame individually. However, the slight difference of each constituent frame from their respective averaged batch will result in errors in the registration. If the displacement of lung tissue between the batches can be approximated as linear, then the deformation field associated with an individual frame can be estimated. Therefore, these parameters are back-interpolated from the deformation fields associated with the respiratory batches, based on the relative respiratory phase of each frame and batch independently. Finally, these estimated deformation fields are applied to each frame in the series.

An additional benefit of this process is that it requires fewer applications of the non-rigid registration algorithm, which can be computationally expensive to perform. As it is only applied between the 20 reconstructed batches, it doesn't scale with the total number of frames in the acquired image series. This is beneficial for large sets or where computational power is limited.

4.6.6 Pipeline Integration

Both the dual-stage and batch-mode registration algorithms are integrated into this pipeline. The user can toggle each of these processes and select which feature-based and/or intensity-based algorithms to employ. When the dual-stage and batch mode processes are deployed in tandem, the feature-based registration is implemented before the formation of the batches. This is because the feature-based deformation methods are less sensitive to

the low SNR and cardiac signal variations and thus can be applied to each individually.

These methods require further exploratory work to investigate the validity of the assumptions used and the accuracy of the resulting registration.

4.7 Filtering

Key to the PREFUL pipeline is the filtering of the time-dependent data to decouple contributions from the cardiac and respiratory cycles. Primarily this applies to each voxel in the time series; it also applies to the markers for the cardiac and respiratory cycle. In this work, the same filtering method is applied to both.

4.7.1 Digital Filtering

The time-varying data are filtered computationally using a digital filter. Unlike analogue filtering where input data are filtered with physical hardware, digital filters are applied computationally. Digital filters can be divided into two groups, characterised by the length of their impulse response [86]:

- Infinite Impulse Response (IIR) filters, whilst being less computationally intensive than their counterpart, produce a non-linear phase response. This results in a non-linear shifting of the frequency components of the signal.
- Finite Impulse Response (FIR) filters can be designed to introduce no phase distortion. All frequencies are shifted by the same

4.7. FILTERING

amount and hence any phase relationships are maintained. However,

the implementation of these filters is generally more memory and

computationally expensive.

An FIR filter type was selected for this work for its favourable phase char-

acteristics. However, this should be less important later due to the zero-

phase application of the filter. Furthermore, no significant computation

time difference between the two methods was observed in this work.

4.7.2Filter Design

Four types of filters are encoded in this PREFUL pipeline: Low-Pass, High-

Pass, Band-Pass, and Band-Stop. Each of these can be selected by the user

in each processing tab. However, in general the low-pass filter is used for

respiratory processing and the high-pass filter for cardiac processing.

With the type of filter selected, the filter is created using inbuilt MAT-

LAB tools. The FIR filter selected for processing is generated using the

'kaiserwin' algorithm[87], chosen for its linear magnitude response within

the pass-band. The user-editable properties of the filters are the pass-band

frequency and stop-band frequencies. As the relative magnitude of these

values is important, and will change for each filter type, the user-interface

will assist the user in selecting them correctly. In general, they should be

placed at a frequency in between the contributions from the respiratory

and cardiac cycles. Further to this, the following parameters are fixed in

the back-end:

• Stop band attenuation: 60 dB

• Pass band ripple: 0 dB

187

• Minimum Order: Any

• Scale Passband: True

4.7.3 Filter Application

The filter is applied using inbuilt MATLAB functions. To preserve the temporal location of features in the filtered waveform, the data are zero-phase filtered. This is done by processing the data in both the forward and reverse directions[88]. The result has zero phase distortion as opposed to being shifted as would happen if it were filtered in just one direction. Application of this filter is done along the x-dimension for the cycle markers and in the z-dimension for each voxel. The processing time of this function will increase with the number of voxels in the image, resulting in a cost function of:

$$O(N_w \times N_h)$$

where N_w is the number of voxels in width and N_h is the number of voxels in height. Therefore, this process is multi-threaded to reduce computation time.

4.8 Reconstruction

The final step of the reconstruction pipeline is to combine all the included images into a single ventilation and perfusion cycle. Sorting each of the frames by their assigned phase will result in a sporadically sampled cycle. To simulate a constant sampling rate, the images are interpolated onto an equidistant time grid. Furthermore, this process increases SNR by signal averaging. The number of interpolated frames is selected by the user and each of their phases calculated so as to space them equally throughout a single cycle.

4.8.1 Algorithm

To estimate each voxel value of the interpolated images, non-parametric kernel regression is employed[89, 90]. This can be written as:

$$E(Y|X) = m(X) \tag{4.18}$$

where the conditional expectation value E of Y (interpolated frames) relative to the value X (input frames) can be estimated using the function m. A locally weighted mean, based on the Nadaraya-Watson estimator is used to approximate m for each interpolated image (h):

$$m_h(x) = \frac{\sum_{i=1}^n K_h(dx_{h,i})y_i}{\sum_{i=1}^n K_h(dx_{h,i})}$$
(4.19)

where i denotes the input frames, y the input voxel intensity and K the weighting kernel, dependent on the relative phase x. In this work, the weighting kernel is calculated using a Gaussian:

$$K_h(x - x_i) = exp\left(\frac{dx_{h,i}^2}{2c^2}\right) \tag{4.20}$$

where c, the standard deviation controlling the kernel bandwidth, is chosen by the user. For every interpolated frame, the contribution from each input frame is calculated based on the difference in phase value:

$$dx_{h,i} = |x_h - x_i| (4.21)$$

This calculation also wraps to account for the repeating nature of the cycle. This process is applied to every voxel and is multithreaded to reduce processing times. This process can also be used to calculate the true phase of each of the interpolated images by combining the phase of the input frames instead of the voxel value. Furthermore, as well as returning the estimated voxel value, this process returns the weighted number of included frames for each process as the denominator in Equation 4.20. This results in a 'frame quality' metric, as utilised previously in Equation 4.13.

4.8.2 Observations

Example plots for the frame quality of a reconstructed ventilation and cardiac cycle are displayed in Figure 4.6iv. The frame quality metric is based on the weighted number of images combined to create each frame. For the perfusion cycle, these values usually appear approximately uniform for all frames within the reconstructed sequence. However, this is not the case for the ventilation cycle. The contribution to images at the limits of the ventilation cycle (fully inhaled or exhaled) is reduced compared to those in the intermediary. Primarily, this is because of the variation in respiratory amplitude during the acquisition. If all points are included in the ventilation phase assignment, then there is reduced sampling at the greatest amplitude. This effect is more pronounced at the inhaled limit due to greater variation in respiratory amplitude. Thresholding of the ventilation marker reduces this effect; however, it is not alleviated completely. Errors

likely persist due to two primary reasons:

- During thresholding, a middle ground must be chosen. Such as to omit the ends of the greatest cycles whilst still retaining sufficient data points and thus variation in the lung state. However, this threshold often must surpass the amplitude of the lowest peaks. Therefore, phase points near the limits of the thresholding still suffer reduced sampling.
- Secondly, the algorithm projects the respiratory marker onto a sinusoidal curve, when in practice post-thresholding this approximation is flawed. A complete marker cycle better fits this assumption; however, a thresholded cycle is not complete. Therefore, at their new artificial limits, the markers do not exhibit sinusoidal properties and are changing at a much greater rate. Furthermore, the inverse applies for cycles where the threshold surpasses their amplitude. They are projected onto the same sinusoidal relationship and therefore their greatest amplitude points contribute to the over sampling of the intermediary phase points.

Two options present themselves to reduce this effect:

- Ensure there is little to no variation in the respiratory amplitude during the scan. This is often unfeasible even for healthy patients.
 A breathing restriction could be put in place; either by restricting airflow or expansion of the thoracic cavity. However, this could be uncomfortable to the patient and may change the lung characteristics.
- Only include respiratory cycles of nearly identical amplitude. However, this would require very long scan times to collect sufficient acceptable data points.

Additionally, the phase points associated with inhalation and exhalation may not experience the same amount of sampling. This is signified by the differing amplitude of included frames in the quality plot. This can occur when the subject exhales at a faster rate than they inhale, or vice versa, resulting in fewer phase points within this regime.

4.9 Segmentation

To display the only anatomically useful regions of the functional maps, the lungs must be segmented as a region of interest. This can be achieved either manually or by automatic segmentation. Each approach was implemented in this pipeline using the processes detailed below. The process of segmenting the lungs may be more straightforward than other anatomical regions of the body, due to the high contrast between the lung parenchyma and walls of the thoracic cavity, typical of proton pulmonary MRI.

4.9.1 Manual

The simplest method for segmenting the lungs is by manually outlining the thoracic cavity. In Matlab, this process can be assisted using the Image Segmenter App, part of the Image Processing Toolbox. However, to achieve accurate results, this can be a meticulous process, made especially difficult with the use of a mouse. Therefore, this process can be prone to human error, which can be abated by implementing several semi-automated processing steps, detailed in Figure 4.13. The first step is a rough manual segmentation of the interior of the thoracic cavity. This is followed by an iterative region-growing algorithm[91, 92] (known as snakes or active contours) to expand the user defined mask to the walls of the thoracic cavity.

The number of iterations used is selected by the user, depending on the size of the image. The user may need to adapt this mask with the 'paint brush' tool if the algorithm expands beyond the thoracic cavity, for example by expanding into a low signal region such as the stomach, which constitutes a low signal region adjacent to the thoracic cavity. Finally, a morphological erode operation is applied to the mask to remove contributions from the wall of the thoracic cavity, typically by one voxel.

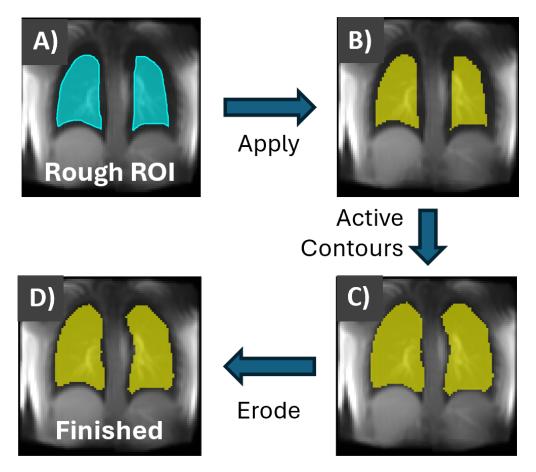


Figure 4.13: Typical process for the assisted manual segmentation of the lung. Depicting the manually-drawn rough interior of the thoracic cavity (A) and resulting mask (B). The active contours algorithm is applied to expand the mask to the walls of the thoracic cavity (C). Finally, the mask is eroded to reduce contributions from the wall of the thoracic cavity (D).

Individual images captured during the acquisition sequence may exhibit a low SNR. This could reduce the accuracy of the resultant segmentation. Instead, a reference image is generated for each dataset to facilitate manual segmentation. This is achieved using the mean value of each voxel for all images in the time-series, each registered to an intermediary lung state.

This process is highly reliable, as significant errors in segmentation will be immediately apparent to the operator. Therefore, this process is most useful for lower quality data sets with reduced SNR. However, by nature this process requires user input and thus can be time consuming. This is exacerbated for processes requiring the segregation of each acquired image, and is thus not feasible for the many-hundred-image sequences typical of PREFUL. In such cases, it is useful to automate this process.

4.9.2 Automated

The PREFUL algorithm benefits greatly from automation of the segmentation process, especially with translation to the clinical setting. Automating this process reduces the requirement for user input, accelerating this postprocessing and increasing repeatability. Furthermore, approaches requiring the segmentation of multiple images, such as using 2D lung area to determine respiratory phase, are not feasible for manual completion by the user.

Under ideal conditions, the lungs may appear to be a simple organ to implement automated segmentation. The low parenchymal signal intensities associated with pulmonary MRI typically result in significant contrast between these regions and the surrounding tissue. Therefore, an automated thresholding function may suffice to segment the void space. However, variation between participants combined with signal drop-off in the peripheries of some low-field scanners often renders this method insufficient. An alternative technique is to employ artificial intelligence to segment these regions. This method is explored below.

4.9.2.1 AI for Image Segmentation

Artificial intelligence is defined as the simulation of human cognitive function. Hailed as the fourth industrial revolution [93] after steam, electricity and computers, AI is finding more and more application in the world around us. To attempt to summarise in this text the latest developments would be futile. Given its rapid rate of development, the information would be obsolete within the year. However, to name just a few advances, progress has included development in large language models [94], driverless cars [95], and assistive medical diagnosis technologies [96].

Many AI innovations can be characterised as two camps, machine learning and deep learning. Machine learning is defined as the ability of an algorithm to analyse data, learn from this data, and then apply this learning to make informed decisions. Whilst powerful, these algorithms still require human intervention during the learning process, defining features important to analysis, and tweaking the algorithm if the output is not as required. For example, consider a program to recognise the difference between a cat or a dog; the algorithm could be programmed to include features such as the tail length and ear shape.

Deep learning is a subset of machine learning and can be considered its evolution. Now the system can determine for itself which features are important. It does this using an algorithm with a logical structure like that of a human brain would make decisions. This is called an artificial neural network. Take the previous cat/dog example as before, a deep learning algorithm may decide for itself that in fact nose shape and total mass are the best indicators of species. However, it can often settle upon features far more abstract, and therefore when these models are trained, can act as somewhat of a 'black box'. That is to say, it will return a result often

without the user having a full understanding of how the decision was made.

Therefore, although requiring much less human intervention than machine learning, training a deep learning algorithm requires much more data. Furthermore, the quality of this network will only be as good as the data it learns from. For example, if you train a deep network on purely ginger cats and black dogs, and then show it a black cat, it may reasonably distinguish it as a canine.

4.9.2.1.1 Artificial Neural Networks

Artificial neural networks, as their name suggests, were developed to simulate the connections in the human brain[97]. Therefore, a brief description of their biology is useful for comparison. By the simplest model, a neuron system can be thought of as having four components, each with its own function:

Cell body (Soma) Containing the nucleus and controlling the function of the neuron.

Dendrites Branch-like structures extending from the cell body responsible for communication and receiving signals from other neurons.

Axon Carries electrical impulses along the length of the neuron and ends in the axon terminal, where signals are passed to other neurons.

Synapse The chemical junction between the axon terminal of one neuron and the dendrites of another. A chemical reaction here effectively decides how information is transmitted between neurons, called the action potential.

The human brain has approximately 86 billion of these neurons. Connections within this network form the basis of our ability to process information and form complex thought.

The 'neurons' in an artificial neural network consists of three components, each analogous to their biological counterparts:

Node Where computation is performed, based on the cell body.

Weights Similar to the axon, in charge of signal transfer and connecting to other nodes.

Activation Like a synapse, dictates the transfer of information from one node to another.

An artificial neural network can be thought of as a structure of layers, made up of three groups: the input, the output, and the hidden layer(s). Named as such, the input and output layers accept and return the resultant information respectively. There is generally only a single instance of each. However, a neural network typically has many hidden layers, and it is within these layers that the complex calculation takes place. These layers are made up of artificial neurons as described previously, each connected to every neuron in adjacent layers, but not to neurons in the same layer.

The input layer nodes are created to reflect each of the variables input into the network. For example, that could be tail length and total mass for the dog/cat differentiator network. For classification problems requiring a discrete result, the output layer nodes are initialised such to reflect the desired outputs: cat or dog in this case. Information is passed through these layers such that a probability is produced from the output layer. Importantly, without alteration this network is only capable of returning

the probability of the input belonging to a cat or a dog. Should it be given the properties of a rabbit for example, it will still return the likelihood of the creature belonging to the feline or canine groups.

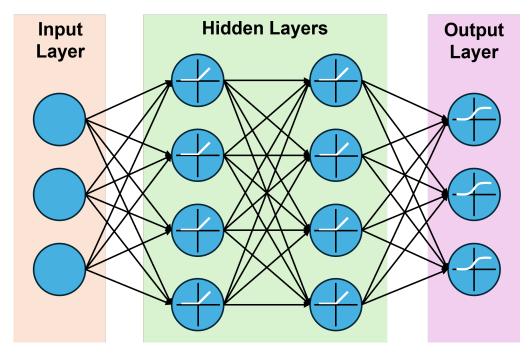


Figure 4.14: Diagram of the core components of an artificial neural network, including nodes (blue circles) and weights (black arrows), within the input, hidden and output layers. Activation functions are visualised within some of the nodes, including rectified linear units and sigmoid functions within the hidden and output layers, respectively. Note: modern artificial neural networks will typically contain many more layers in the hidden section than the two displayed here.

Each node accepts a weighted sum from all its inputs and then applies an activation function, mimicking the action potential of its biological counterpart. This function can take multiple forms in the hidden layers, such as a Rectified Linear Unit (Depicted in the nodes of the hidden layers in Figure 4.14), which suppresses negative values and allows the function to act non-linearly. In the output layer a sigmoid activation function is often used (depicted in the nodes of the output layer in Figure 4.14) such that the network generates a probability.

A typical Neural network will have many variables or trainable parame-

ters, often into the tens of millions. The act of training the network is effectively an optimisation problem of a non-linear function; finding a combination of parameters that result in the lowest error from an expected result. The parameters can start randomly assigned and data is submitted to the network. The result from the output layer is then compared to the ground truth and the difference is used to determine the error using a loss function. This loss is then used to adjust the training parameters via a process called backpropagation[98]. Hyperparameters are used to control how much these trainable parameters are adjusted in response to the associated loss. Repeated many times, this process should increase the network accuracy, effectively 'training' it. This often requires significant computing capabilities to avoid long training times. However, with the development of GPUs and cloud computing infrastructure in the last decade, this has become more accessible.

An example of a common artificial neural network is a convolutional neural network, specifically designed for image processing and object detection, such as that required for recognising road signs in driverless cars. This network type is discussed further in the next section. Recurrent neural networks share a similar architecture to convolutional neural networks but have built-in feedback loops that permit the algorithm to 'remember' previous data. These are used when the input data is sequential, such as in language processors such as AI assistants (Siri (Apple, USA), Alexa (Amazon, USA) etc.), or translators (Google Translate (Google, USA) etc.).

4.9.2.2 Convolutional Neural Networks (CNN)

Widely used for object detection and classification, convolutional neural networks[99] are inspired by the function of the visual cortex, where sig-

nals from the retina are processed as if by a hierarchical system of cells. For example, when reading text some cells detect shapes/patterns to form individual letters. This information is then passed to further cells with greater context to form syllables. These are then processed by other cells to form words and so on to form sentences before the information is sent to higher centres in the brain. Inspired by this understanding, the first CNN was generated, comprised of a hierarchical mathematical model using convolution and down-sampling[100]. This works like the model of the visual cortex, processing only part of an image at a time to extract features. Then combining these features with additional context to attain insight into the entire system. In a CNN, each of the input nodes can represent a pixel of the input image. Convolutional neural networks generally comprise of six layer types:

Convolution layer A layer where features are extracted via convolution with a filter. These filters can perform operations such as blurring, sharpening or edge detection.

Activation layer Applies an activation function discussed previously.

Pooling layer Reduces the dimensionality of an image often by taking the maximum value or averaging a region. Spatial pooling can also be referred to as downsampling.

Batch Normalisation layer Controls input distortion by normalising and scaling to make the network more robust to initial parameter values and covariance shift.

Dropout layer A layer that randomly ignores the activation layer to prevent over-fitting.

Fully Connected layer An output layer used for object classification.

These networks are commonly used for computer vision applications. Four typical classes of computer vision are defined as:

Image Classification Globally classifying the main subject of an image.

Classification and Localisation Classifying and locating the main subject of an image.

Object Detection Classifying and locating multiple subjects of an image.

Semantic Segmentation Applying a classification to each individual pixel of an image.

The application of each of these classes is demonstrated in Figure 4.15. For the image segmentation task in this work, a network capable of semantic segmentation is required.

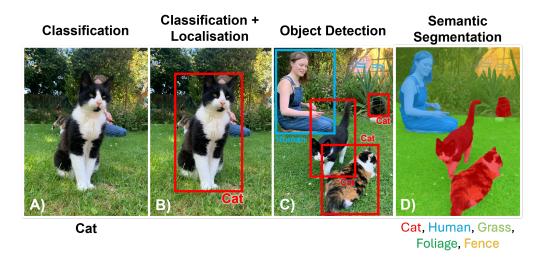


Figure 4.15: Demonstration of the application of each discussed computer vision classification technique. Classification and segmentation were applied through manual evaluation for these examples.

4.9.2.2.1 CNN Architectures

Designing a CNN architecture from scratch is beyond the scope of this work. Typically, models are trained using pre-existing architectures, which

may be adapted slightly to suit local requirements. Below are descriptions of two models that are relevant to this work:

UNet UNet was designed specifically for biomedical applications[101]. Variations of this model[102] have found widespread application in biomedical image segmentation[103, 104]. This architecture is comprised of two parts: contraction and expansion. Displayed graphically in Figure 4.16, these resemble the letter 'U', for which this network was named. The contraction segment consists of several sequential blocks of convolution and max pooling layers, which reduce the spatial information whilst increasing feature information. The expansion segment combines feature and spatial information and can be thought of as a decoder. Information is passed between the contraction and expansion layers to guide localisation.

SegNet SegNet[105] is similar to UNet in architecture, although it is less memory intensive. This is because SegNet passes a reduced quantity of information from the contraction layers to the expansion path. Although not specifically designed for biomedical applications, the strength of SegNet in this case is the readily available pre-trained models in Matlab. This includes the VGG-16 network[106], capable of recognising 1000 different objects. However, these objects are not necessarily related to biomedical imaging.

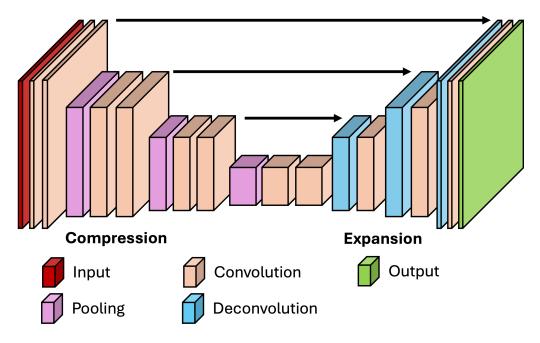


Figure 4.16: Simplified visualisation of the UNet architecture, including 3 pooling layers. Information is passed between successive neighbouring layers and also between contraction and expansion blocks, with both pathways indicated by the **black arrows**.

4.9.2.2.2 Network Training

With the network architecture selected, model parameters must be optimised before the system can produce meaningful predictions. This process typically requires a significant quantity of training data, depending on the task complexity and model architecture. The training data comprises input data for which the expected output is known. For semantic segmentation tasks, this is generally in the form of numerous human-labelled images. However, human labelling can introduce inconsistencies in the training dataset and is very resource-intensive. These data are typically split into three groups: the training, validation, and test sets. The first two sets are used for model training and hyperparameter tuning, respectively. The final group is used to determine performance metrics such as Intersection over Union (IoU)[107], pixel accuracy, and mean average precision for the

resulting model and is not used during the training process.

When training a model from scratch, the weights and biases are typically initialized at random. Inputting the training data to this system will yield erroneous results; however, this is the first stage of model training. These results are compared to the expected output, and a loss metric is generated based on the loss function. This loss is then used to adjust the model parameters via a process called backpropagation[98], before more training data is passed through the model and the process repeated. Effectively, this training process is an iterative optimisation problem, which aims to determine a set of model parameters for which the loss is minimised. Training is deemed complete when convergence criteria are met, such as validation loss plateauing, reaching maximum epochs, or achieving satisfactory performance metrics. It is possible to over-fit a trained network, where the model learns training-specific patterns that fail to generalize to unseen data, often detected by divergence between training and validation loss.

4.9.2.2.3 Transfer Learning

An alternative method for model training is transfer learning[108], which can significantly reduce training times and data requirements. Instead of initialising model parameters at random, this method utilises a pre-trained model trained on large-scale datasets to determine the starting values for model weights and biases. The existing model, often trained on related but distinct tasks, may have captured generalisable feature representations which are transferable between tasks. Common approaches include fine-tuning all parameters with reduced learning rates, or freezing early layers while retraining only the final classification layers. This can improve trained model performance, especially in scenarios where training data or

computational resources are limited. However, transfer learning effectiveness depends on the similarity between source and target tasks.

4.9.2.3 Training Lung Segmentation Networks

The aim of this section was to train a number of artificial neural networks capable of the semantic segmentation of the lung, from MR images captured using the 0.5 T scanners described in this work. This includes four scanning modalities: 2D GRE scans in the coronal, axial and sagittal planes, as well as 2D HASTE scans in the coronal plane. These networks were based on pre-existing architectures and trained using inbuilt MATLAB tools. Models trained to an acceptable accuracy were integrated into the home-built PREFUL pipeline. Scanner protocols for all scans used for model training are described in Tables 4.5 and 4.7.

4.9.2.3.1 Image Labelling

Prior to manual segmentation, images were pre-processed to meet the requirements of the network architecture. This included changing the image size to 128 by 128 voxels and scaling the image values to between 0 and 255 (8-bit). A benefit of PREFUL for generating labelled data for training a neural network is the large number of images captured for each scan. Therefore, many hundreds of training images can be created from only a handful of acquisitions.

The aim of the first network trained was to segment coronal lung images captured using a standard 2D GRE sequence. Ground truth labels for training the coronal 2D GRE network were manually generated using the Medical Image Labeller application, part of MATLAB's medical imaging

toolbox. In total, 3552 training images were generated from MR images of 15 healthy individuals across each of the acquisition types. This encompassed participants in the postures: supine, prone, sitting and upside down (although scans were rotated or flipped as required to ensure the base of the lung remained at the bottom of the image and left lung on the right). All voxels within all images were assigned one of three classes: Left Lung, Right Lung and Background.

Left Lung For voxels within the left lung.

Right Lung For voxels within the right lung.

Background For all other voxels in the image, regardless of anatomy.

Generally, the segmentation of each of the 2D GRE coronal images was achieved using a semi-automated process. The rough edges of the lungs were painted onto each image by the user, ensuring a complete loop was made within each of the lungs. Then a custom function was employed to snap this rough outline to the edges of the thoracic cavity. This function followed the steps:

- 1. Fill holes.
- 2. Active contours, 5 iterations.
- 3. Close function (dilate then erode) using a disk shaped kernel of diameter 3 pixels.

Finally, any remaining segmentation errors were manually corrected.

For the image labelling of the later networks, the previously trained 2D coronal network was applied to assist with the initial rough outline of the

lung. Any errors in the resulting labels were corrected and the above function was applied if necessary. The number of training images and the number of healthy participants used for model training is summarised in Table 4.2.

Image labelling was performed by two people, Matilda Chalk and Arthur Harrison, with details of their respective contributions in Table 4.2. No inter-observer repeatability studies were employed in this work.

For the coronal and axial scan, a 200-image sequence could be fully labelled in approximately three hours, equating to just over one minute per frame. Sagittal scans were generally slightly faster to label as only one lung is in frame. Haste scans typically took slightly longer to label, due to the lower quality of the images. This lower SNR often resulted in the failure of the semi-automated edge-finding algorithm, requiring the labelling process to be fully manual.

Labelled Training Data							
Acquisition	Plane	Images	Participants	Scanner(s)	Labeller(s)		
2D GRE	Coronal	3552	15	0.5 T Open	90% MC		
				0.5 T Bruker	10% AH		
2D GRE	Axial	135	10	0.5 T Open	100% AH		
2D GRE	Sagittal	200	8	0.5 T Open	100% AH		
2D HASTE	Coronal	150	13	0.5 T Open	100% AH		

Table 4.2: Number of training images manually labelled and number of healthy participants for each acquisition type. MC represents Matilda Chalk, AH represents Arthur Harrison.

An image augmentation function was employed to adapt the training data each time it was passed through the training algorithm. This was to increase both the noise and variation in the training data to prevent overfitting. The augmentation function applied a number of random transformations to each labelled training image, within the bounds specified in Table 4.3.

Augmentation Limits				
Transformation Type	Limits			
Rotation (Degrees)	-10 to 10			
X Scale	0.8 to 1.2			
Y Scale	0.8 to 1.2			
X Translation (Voxels)	-10 to 10			
Y Translation (Voxels)	-10 to 10			

Table 4.3: Limiting bounds to random transformations applied to each training imaged by the augmenter function.

4.9.2.3.2 Network Generation

Two networks were generated using inbuilt MATLAB protocols for the semantic segmentation of the Coronal 2D GRE images. These were based on the UNet and SegNet architectures, each with three pooling layers and adapted to include a Dice classification output layer. The UNet model was trained from scratch, however the SegNet network was initialised with the pre-trained VGG-16 model for transfer learning. Each model was created to classify the three classes designated by the training data.

Networks for the Axial 2D GRE and Coronal 2D HASTE images were trained via transfer learning, with initialised weights and biases from the prior trained Coronal 2D GRE SegNet network. As the output labels were the same across each of these acquisitions, all weights and biases were taken from the Coronal 2D GRE network. For the Sagittal 2D GRE network, the pre-trained Coronal 2D GRE SegNet network was also used for

transfer learning, however the classification layers were replaced to reflect the single lung within the image. These replacement classification layers were generated using inbuilt MATLAB protocols.

4.9.2.3.3 Network Training

All training data was partitioned into training, validation and test sets with the respective ratios 75:20:5. All model training and hyperparameter optimisation was completed within MATLAB's Experiment Manager environment, part of the Deep Learning Toolbox. Training was completed using the computer hardware detailed in Appendix B. Hyperparameter optimisation is beyond the scope of this thesis. The hyperparameters used to train each of the models are listed in Appendix A.1. The training progress plots the Coronal 2D GRE UNet network is depicted in Figure 4.17. Training times for each network was generally between 10 and 20 minutes.

The Axial 2D GRE, Sagittal 2D GRE, and Coronal 2D HASTE networks had significantly less training data compared to the Coronal 2D GRE network (Table 4.2). To address this limitation, transfer learning using the pre-trained Coronal 2D GRE model was employed. This approach utilised the extensive labelled dataset from the Coronal 2D GRE network to compensate for the limited training data from the other MR acquisition modalities.

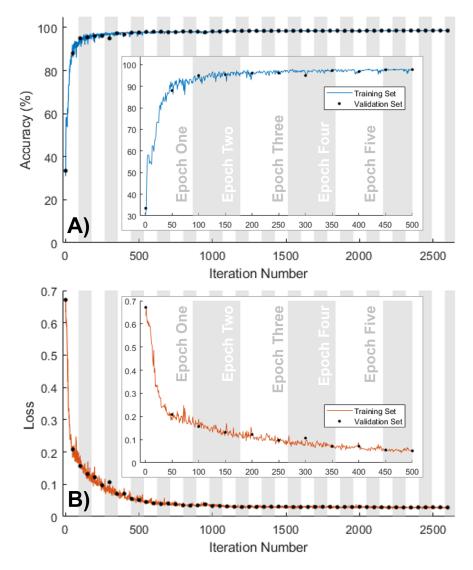


Figure 4.17: Training progress plots for the 2D GRE SegNet network, including the change in accuracy and loss due to model training.

4.9.2.3.4 Training Results

Initially, the UNet and pre-initialised VGG-16 SegNet networks were trained on the Coronal 2D GRE images. The accuracy of the trained networks was contrasted using Intersection over Union (IoU) score[107] of each network's respective test set in Table 4.4. Additionally, typical segmentation results for each of the networks are presented in Figures 4.18 and 4.19, accompanied by heat maps summarising the types of error within the network.

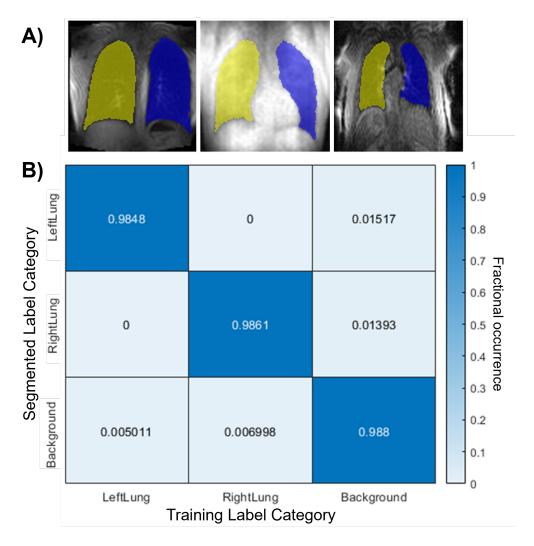


Figure 4.18: Summary of the accuracy of the trained Coronal 2D GRE network, based on the pre-initialised VGG-16 SegNet architecture. This figure includes examples of typical semantic segmentations of three typical pulmonary MR images from the test set (A). In these segmentations the segmented right lung is depicted in **yellow**, the left in **blue**, while the background is omitted. Additionally, the network accuracy heat map, also based on the test set, is presented (B), detailing the relative frequency of each voxel segmentation category.

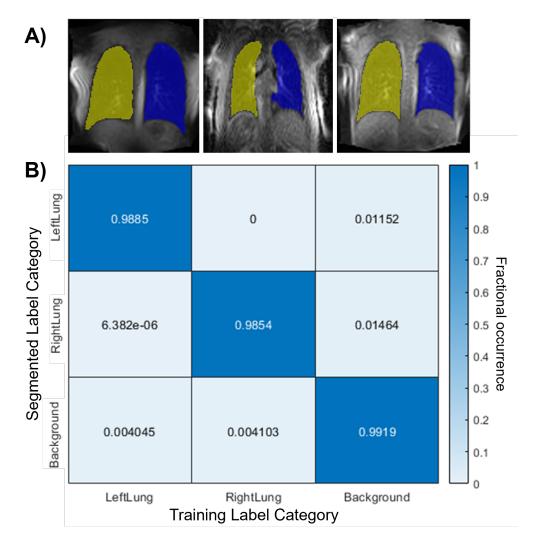


Figure 4.19: Summary of the accuracy of the trained Coronal 2D GRE network, based on the UNet architecture. This figure includes examples of typical semantic segmentations of three typical pulmonary MR images from the test set (A). In these segmentations the segmented right lung is depicted in **yellow**, the left in **blue**, while the background is omitted. Additionally, the network accuracy heat map, also based on the test set, is presented (B), detailing the relative frequency of each voxel segmentation category.

The accuracy of each of the 2D Coronal GRE trained networks was evaluated on the Axial 2D GRE and Coronal 2D HASTE test data. The 2D sagittal GRE test data was omitted from this process due to the difference in the output labels. A summary of the efficacy of each trained network is presented in Figure 4.20.

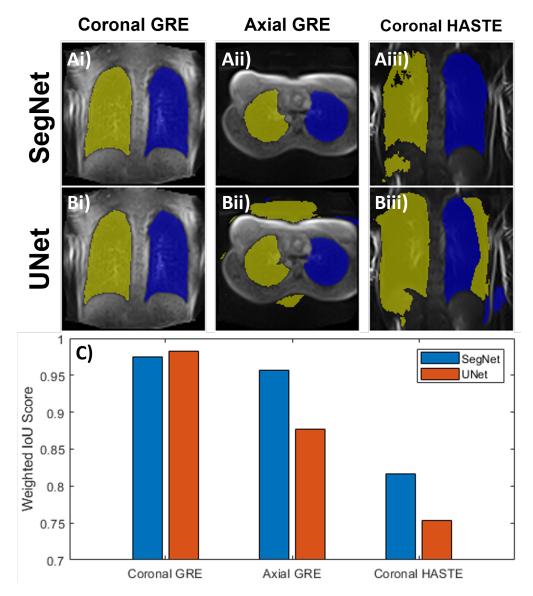


Figure 4.20: Summary of the application of the trained Coronal 2D GRE networks on the additional MR image data types with the same voxel classes. The figure includes the segmentation of Coronal GRE (i), Axial GRE (ii) and Coronal HASTE (iii) MR images, by the trained SegNet (A) and the UNet (B) - based networks. This is accompanied by a comparison of the Weighted IoU score (C) for the application of each network on each acquisition type.

Both the UNet and pre-trained SegNet network architectures appear to be able to segment the coronal 2D GRE datasets well, yielding IoU scores of over 98% (Table 4.4). However, it appears that the UNet based network is slightly more accurate. However, when applied to pulmonary MR images acquired in different planes and modalities, the SegNet based network per-

formed better. Therefore, this network may be more robust at segmenting the lungs. The UNet network, while more efficient at the task for which it is trained, may be more specialised to this data type.

This observation may be explained by the method used to train each of the networks. The SegNet based network was trained via transfer learning, using the VGG-16 network which is capable of recognising 1000 everyday objects. Therefore, when trained on the lung images, this model may have a more fundamental recognition of the shape of the pulmonary organs, rather than just pattern recognition specific to the limited training data. Consequently, this model was subsequently used for transfer learning of the additional networks. The results of this training, for the Axial 2D GRE, Sagittal 2D GRE and Coronal 2D HASTE data, are presented in Figures 4.21, 4.22 and 4.23, respectively.

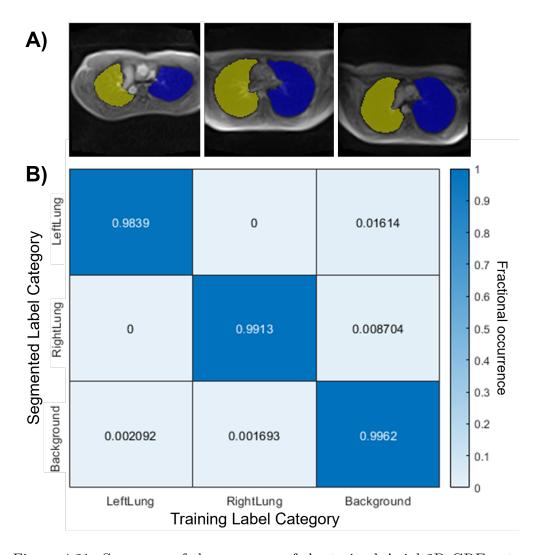


Figure 4.21: Summary of the accuracy of the trained Axial 2D GRE network, based on the pre-trained Coronal 2D GRE SegNet architecture. This figure includes examples of typical semantic segmentations of three typical pulmonary MR images from the test set (A). In these segmentations the segmented right lung is depicted in **yellow**, the left in **blue**, while the background is omitted. Additionally, the network accuracy heat map, also based on the test set, is presented (B), detailing the relative frequency of each voxel segmentation category.

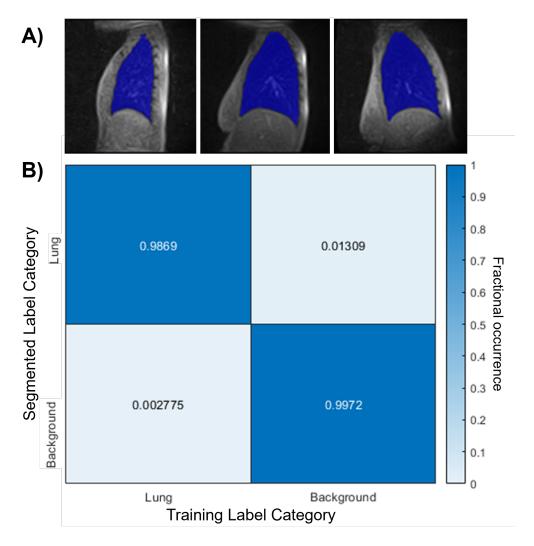


Figure 4.22: Summary of the accuracy of the trained Sagittal 2D GRE network, based on the pre-trained Coronal 2D GRE SegNet architecture. This figure includes examples of typical semantic segmentations of three typical pulmonary MR images from the test set (A). In these segmentations the segmented lung is depicted in **blue** and the background is omitted. Additionally, the network accuracy heat map, also based on the test set, is presented (B), detailing the relative frequency of each voxel segmentation category.

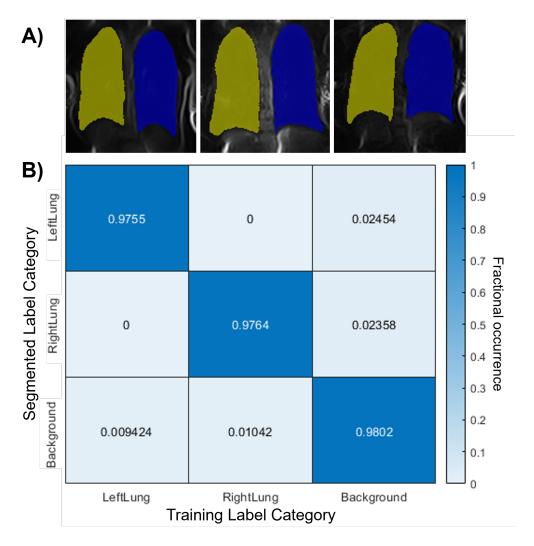


Figure 4.23: Summary of the accuracy of the trained Coronal 2D HASTE network, based on the pre-trained Coronal 2D GRE SegNet architecture. This figure includes examples of typical semantic segmentations of three typical pulmonary MR images from the test set (A). In these segmentations the segmented right lung is depicted in **yellow**, the left in **blue**, while the background is omitted. Additionally, the network accuracy heat map, also based on the test set, is presented (B), detailing the relative frequency of each voxel segmentation category.

Finally, the accuracy of each model was evaluated using the test set associated with each network. These results are presented in Table 4.4.

Trained Network Accuracy					
Trained Network	Architecture	Weighted IoU Score			
Coronal 2D GRE	SegNet	0.974			
Coronal 2D GRE	UNet	0.982			
Axial 2D GRE	SegNet	0.991			
Sagittal 2D GRE	SegNet	0.991			
Coronal 2D HASTE	SegNet	0.958			

Table 4.4: Comparison of the accuracy of each Trained Networks using the weighted IoU score[107] determined using each network's test data set.

4.9.2.3.5 Pipeline Integration

Each of the SegNet based networks were integrated into the home-built PREFUL pipeline. These networks serve to automatically segment the thoracic regions, either for ROI designation or for ventilation phase assignment by 2D lung area. However, following the neural network based segmentation an additional processing step is applied to the segmented lungs to reduce noise and remove potential errors. This processing followed three steps:

- 1. Fill holes All enclosed spaces within each of the segmented lungs are removed.
- 2. Remove islands All but the largest island by 2D area for each lung segmentation are removed.
- 3. Smooth edges A morphological close operation with a disk shaped kernel of diameter of 3 voxels is applied to each segmented lung.

Usually these operations have a minor effect on the segmented lung mask.

However, they may reduce the error in segmentation if the networks experience data that varies from their training data. For example, in Figure 4.20Aiii and 4.20Bii, this processing would have significantly reduced the segmentation error.

4.9.2.4 Network Evaluation

This section follows the training of five CNNs, each trained to segment 2D data acquired using specific sequences or specific imaging planes, acquired at 0.5 T. Each of these networks appears to be highly capable at segmenting healthy lungs when imaged under these specific conditions. Therefore, each of the Segnet-based CNNs was integrated into the PREFUL pipeline for lung segmentation and respiratory phase determination.

However, these networks were trained on data with specific acquisition parameters. Varying these parameters can significantly change the contrast within each MR image, even within the same pulse sequence. Furthermore, these networks have been trained on data acquired of relatively few individuals. Therefore, in future, these networks should be tested on a larger range of datasets with a greater variance in acquisition parameters and a wider range of individuals. These networks were trained via transfer learning from the VGG-16 network with the hypothesis that the resulting networks will be highly adaptable to new data types. However, future retraining may be necessary, potentially using these networks as the basis for transfer learning.

The above discussion is limited to healthy volunteers. No analysis has been performed on the capability of these CNNs to segment diseased lungs. If pathology were to alter the appearance of the lungs in their MR images then the networks may exhibit a reduced performance. This represents a

significant limitation to this work. If the networks become too inaccurate, then they will no longer be useful for determining the respiratory phase of each acquired image. Therefore, future work should explore the performance of these networks for segmenting diseased lungs, with a range of pathology types, and perform retraining if necessary.

Further exploratory work may use the labelled training data to conduct transfer learning using a pretrained UNet network, an architecture more commonly used for medical image segmentation tasks[103, 104]. Additionally, the training data could be compiled into a single set, used to train a single network capable of segmenting the lungs from any of the available image series employed in this study.

4.10 Analysis

With successful PREFUL reconstruction, two lung cycles are generated representing an artificial ventilation and perfusion cycle. The final step is to glean useful voxel-wise information from these cycles relating to the local ventilation and perfusion. The original PREFUL paper [47] presents two primary methods for this: using the relative intensities of voxels in the cycle (fractional ventilation and perfusion weighted) and the time each voxel takes to reach its maximum value or time-to-peak (VTTP and QTTP).

4.10.1 Ventilation

This PREFUL pipeline is not always successful in processing image series.

This could be due to low-quality input data, or the incorrect selection of processing parameters. Furthermore, a data-set may exhibit successful ven-

tilation processing, whilst the perfusion fails, or vice versa. Therefore, the success of each of these processing types is determined separately. For ventilation, the correlation between the average fractional ventilation within the lung for each reconstructed frame of the ventilation cycle frame, and the associated respiratory markers is used to automatically categorise the processing output as either a pass or a fail. Figure 4.24 displays the relationship of these metrics for two example processing outputs. This assumes that with lung expansion, the air-fraction within the lung increases. Therefore, the pass threshold for the correlation coefficient of these metrics is set to ≥ 0.5 . This threshold was selected based on the data available during development of this PREFUL pipeline, which only contains healthy volunteers. Therefore, further work is required to validate the reliability of this threshold.

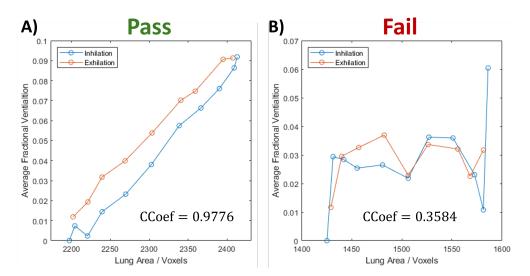


Figure 4.24: Demonstration of the process for determining if ventilation processing has been successful. Displayed for an example case of successful processing (A) and unsuccessful processing (B), based on the correlation of the respiratory marker and average fractional ventilation within the lung for each case. Error bars omitted.

4.10.1.1 Fractional Ventilation

The voxel-wise fractional ventilation is determined using the method presented by Zapke et al.[14] The signal (S) of each voxel is expected to vary throughout the ventilation cycle. One can define S_{ex} to be the signal at greatest exhalation and S_{in} to be the signal at greatest inhalation. With successful registration, the total amount of lung parenchyma within each voxel through the ventilation cycle should remain constant, however the air fraction increases as one draws a breath. Therefore, the average density of the lung tissue will decrease. This should increase again as the subject exhales. This results in lower signals observed with inspiration compared to exhalation. This variation, with associated errors, is presented for a typical lung voxel in Figure 4.25. Using the signal from each voxel with time (S(t)), one can calculate the time dependent fractional ventilation[14]:

$$FV(t) = \frac{S_{ex} - S(t)}{S_{ex} - S_{noise}}$$

$$\tag{4.22}$$

where S_{noise} is equivalent to the background noise, usually omitted for PREFUL[25]. This equation can be altered to give the fractional ventilation at greatest inhalation FV[14]:

$$FV = \frac{S_{ex} - S_{in}}{S_{ex}} \tag{4.23}$$

The values of S_{ex} and S_{in} can be estimated from the reconstructed ventilation cycle using the signal intensity at the time points representing greatest exhalation and inhalation (0% and 50% of the ventilation cycle). This technique is demonstrated in Figure 4.25. A typical fractional ventilation map, generated using this method, is presented in Figure 4.26Ci. The voxel-wise

error can be calculated using the standard error of the time-point used in the calculation. However, this methodology presents a number of drawbacks:

- It assumes that the point of maximum ventilation (thus minium signal) occurs at the point of maximum inhalation. If this property occurs earlier or later in some regions, then the resultant map of the fractional ventilation may not accurately reflect the air within these zones. In the extreme, then the local ventilation appears our of phase with the respiratory cycle, this method can yield negative fractional ventilation values.
- It is based on the intensity of only two points in the phase cycle. Furthermore, as demonstrated in Figure 4.6iv, the corresponding reconstructed frames will often have used the fewest images for reconstruction. This is reflected as the larger error bars at the limits (0% and 50%) of the ventilation cycle in Figure 4.25. Therefore, this method is highly sensitive to errors, especially at the lower sensitivities associated with scanning at low-field.

These limitations may affect the reliability of pulmonary ventilation assessment using this method. Therefore, an alternative method was developed to reduce the effect of the above limitations. The variation in the signal intensity within each voxel is assumed to be sinusoidal with respect to the ventilation phase. If this assumption is correct, then this signal can be fitted to a sine curve using the equation:

$$S = A_{Vent}\sin(2\pi ft + \phi_{Vent}) + C \tag{4.24}$$

where A_{Vent} is the ventilation amplitude, f the frequency, t the position in the ventilation cycle of the reconstructed frame, ϕ_{Vent} the phase offset of the fitted ventilation cycle and C the magnitude offset. This fitting is based on an artificially reconstructed cycle with a known frequency of 1 Hz. Therefore, the fitted variables are A, ϕ and C. The weighted fitting, based on error of the reconstructed voxels, is achieved via non-linear least-squares regression[75–77]. Depicted in Figure 4.25, these values can be used to estimate the voxel intensity at inspiration and exhalation:

$$S_{ex} = C + A_{Vent} \tag{4.25}$$

$$S_{in} = C - A_{Vent} \tag{4.26}$$

Substituting these into Equation 4.23 gives:

$$FV = \frac{2A_{Vent}}{C + A_{Vent}} \tag{4.27}$$

This method uses the weighted variation in voxel intensity of all frames of the reconstructed image sequence, therefore, the resultant maps of the fractional ventilation may be less susceptible to noise. This fitting method is demonstrated in Figure 4.25. An example map of the fractional ventilation, generated using this method, is presented in Figure 4.26Ai. Each value of the fractional ventilation is accompanied by a map of the associated error of each voxel, calculated using the standard error of each fitting parameter.

An additional method explored for determining the fractional ventilation

of each voxel, was to simply use the greatest variation in signal along the reconstructed sequence. Therefore, for Equation 4.23, S_{in} is determined using the value S_{-}^{V} , and S_{ex} is determined using S_{+}^{V} . This method is demonstrated for a typical voxel in Figure 4.25, with a reconstructed fractional ventilation map using this method presented in Figure 4.26Bi. Like the first method discussed, this method is sensitive to noise as it only uses two points for ventilation determination. However, unlike the first method, this method doesn't produce negative fractional ventilation values.

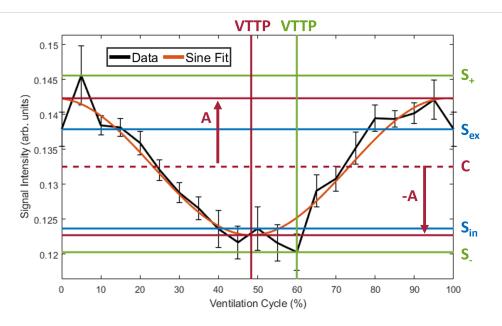


Figure 4.25: Demonstration of methods used to determine local fractional ventilation, for the three methods discussed in this work. The local signal variation with respect to the reconstructed ventilation cycle is plotted as a black line, with error bars representing the standard error calculated during the reconstruction process. A sinusoidal curve is also fitted to the data (Orange line). S_{in} , S_{ex} and VTTP values are determined from the ventilation cycle limits (Blue lines), the maximal change in signal (Green lines), and using the sinusoidal fitting (Red lines).

Example Fractional ventilation maps generated using each of the three methods discussed are demonstrated in Figure 4.26 for a free-breathing healthy volunteer. Each map is accompanied by a histogram to visualise the spread of the data. The image series used to reconstruct these maps

were acquired on the 0.5 T Upright scanner, with acquisition/processing parameters detailed in Tables 4.5 and 4.6.

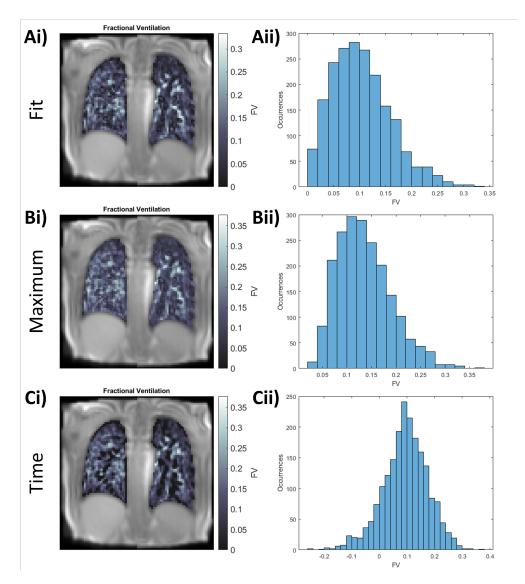


Figure 4.26: Comparison of analysis methods used to determine the voxel-wise fractional ventilation generated from the PREFUL processing. Maps of the fractional ventilation (i) and associated histograms (ii) are displayed for the sinusoidal fitting (A), maximum difference (B) and temporal referencing (C) methods. Underlying data was captured using a Coronal 2D GRE acquisition on the 0.5 T Upright scanner, whilst the participant was in the supine posture, breathing freely.

This PREFUL pipeline utilises the sinusoidal fitting method for calculating the local fractional ventilation. Voxel-wise errors are calculated based on the standard error of the weighted fit parameters.

4.10.1.2 Ventilation Time-To-Peak

VTTP may provide insight into the air filling dynamics of the lung. This metric is referenced to the start of the ventilation cycle at 0%. Two methods present themselves to determine the voxel-wise VTTP, demonstrated in Figure 4.25. The first is to use the point within the reconstructed ventilation cycle for which the voxel exhibits the lowest signal and thus assumed maximum air fraction[25]. Whilst simple to implement, this method only uses the signal of a single reconstructed frame and therefore is susceptible to noise. Furthermore, this method creates a discretised VTTP result, dependent on the number of reconstructed ventilation frames selected during processing. This is demonstrated in Figure 4.27Bii.

Alternatively, the fitted phase term (ϕ_{Vent}) from Equation 4.24 can be used to determine VTTP using all of the weighted points of the reconstructed ventilation cycle. This term can then be used to calculate the fractional VTTP using the equation:

$$VTTP_{\%} = 100\% \times \frac{\left|\frac{3\pi}{2} - \phi_{Vent}\right|_{0}^{2\pi}}{2\pi}$$
 (4.28)

This method uses all frames of the reconstructed ventilation cycle, theretofore may be less susceptible to noise. A comparison of each analysis method for a typical dataset is provided in Figure 4.27.

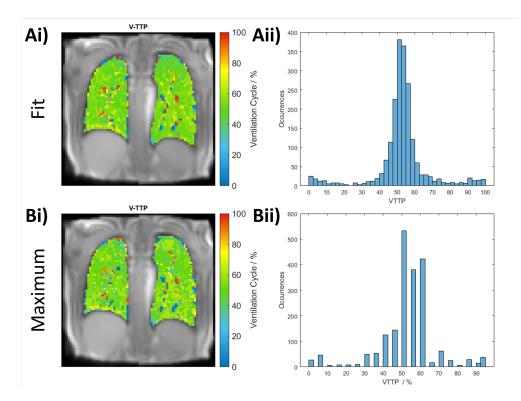


Figure 4.27: Comparison of analysis methods used to determine the voxel-wise VTTP generated from the PREFUL processing. Maps of the VTTP (i) and associated histograms (ii) are displayed for the sinusoidal fitting (A) and maximum difference (B) methods. Underlying data was captured using a Coronal 2D GRE acquisition on the 0.5 T Upright scanner, whilst the participant was in the supine posture, breathing freely.

This PREFUL pipeline utilises the sinusoidal fitting method for calculating the local VTTP. Voxel-wise errors are calculated based on the standard error of the weighted fit parameters.

4.10.1.3 Jacobian Determinant

The Jacobian determinant is an alternative method for investigating the regional expansion of the lung tissue, which may yield results related to the local ventilation. This analysis method is generally used for CT studies[109–111], but can also be used in conjunction with lung MRI[112]. Instead of the signal-based calculations used for the fractional ventilation, the Jacobian determinant is a region-based approach. This method uses

the results of the non-rigid, intensity based registration step described previously. Using the forward displacement map, the 2D Jacobian can be calculated for each voxel (v)[113]:

$$J_{v}(x,y) = \begin{pmatrix} \frac{\partial v_{x}}{\partial x} & \frac{\partial v_{x}}{\partial y} \\ \frac{\partial v_{y}}{\partial x} & \frac{\partial v_{y}}{\partial y} \end{pmatrix}$$
(4.29)

from which, the determinant is calculated, yielding a single number quantification of local expansion. Values greater than 1 represent areas of expansion, values lower than 1 represent areas of contraction. The Jacobian determinant can be calculated using the displacement field generated via the successful registration between the reconstructed maximally exhaled and inhaled states. This may yield information related to local maximum ventilation associated with the image series. However, similar to the fractional ventilation calculations, this processing is based on only two lung states. Furthermore, these lung states typically exhibit the highest errors in voxel value compared to their counterparts. Therefore, this method may be highly sensitive to errors.

An alternative method for calculating the local maximum expansion, may use the temporal variation in voxel-wise Jacobian determinant. Applying registration from the maximally exhaled state to each subsequent reconstructed state, permits the plotting of an expansion vs respiratory cycle relationship. Assuming this relationship is sinusoidal, one can fit a sine curve to this distribution, similar to Equation 4.24. Then the amplitude may be a more reliable representation of local expansion. Additionally, the phase of the associated fitted curve, could be used to gain temporal dynamic information, the Jacobian [Determinant] Time-To-Peak (JTTP), which is analogous to VTTP. Example maps of the Jacobian determinant

and JTTP using this processing method are presented in Figure 4.28. However, further testing is required to test the validity of these results.

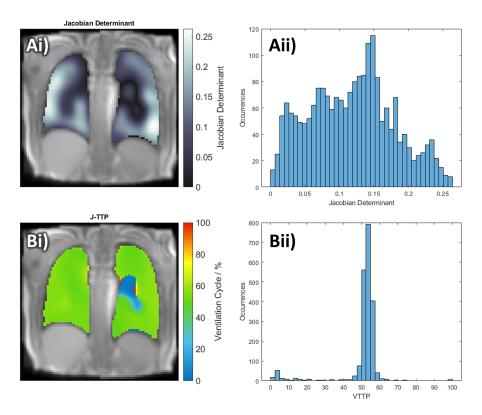


Figure 4.28: Comparison of analysis methods used to determine the voxel-wise Jacobian Determinant and JTTP generated from the PREFUL processing. A Jacobian Determinant map (Ai) and associated JTTP Map (Bi) is presented for a typical dataset, with accompanying histograms (ii). Underlying data was captured using a Coronal 2D GRE acquisition on the 0.5 T Upright scanner, whilst the participant was in the supine posture, breathing freely.

This PREFUL pipeline utilises the sinusoidal fitting method for calculating the local Jacobian determinant and JTTP. Voxel-wise errors are calculated based on the standard error of the fitted parameters.

4.10.2 Perfusion

To determine if perfusion processing has been successful, the number of frames for which the perfusion phase was successfully assigned is used. This threshold is arbitrarily set to $\geq 50\%$; this value was chosen based on the datasets available for development of this pipeline. Therefore, this threshold may need to be optimised when investigating pathology. Furthermore, this threshold does not take into account the size of the underlying dataset.

4.10.2.1 Perfusion Weighted

The local weighted perfusion is calculated based on the variation of signal intensity of each voxel of the reconstructed perfusion cycle. This can be achieved using simply the difference between the greatest voxel intensity and lowest voxel intensity: S_{+}^{Q} and S_{-}^{Q} . This method is demonstrated in Figure 4.29. The magnitude of the perfusion weighted (Q_w) variation can then be defined as:

$$Q_w = S_+^Q - S_-^Q (4.30)$$

However, like for the fractional ventilation determination, this method is reliant only on two reconstructed frames and therefore is susceptible to noise. This analysis can be adapted to consider the weighted contribution from all reconstructed frames by fitting a sinusoidal curve to the variation of each voxel through the perfusion cycle. For a high-pass filtered signal, the fitted equation takes the form:

$$S = A_{Perf} \sin(2\pi f t + \phi_{Perf}) \tag{4.31}$$

where A_{Perf} is the perfusion amplitude, f the frequency, t the phase point corresponding to the reconstructed frame, and ϕ_{perf} the phase offset of the fitted perfusion cycle. The magnitude of perfusion weighted variation is

then defined as:

$$Q_w = 2A_{Perf} (4.32)$$

This sinusoidal weighted fitting method is demonstrated in Figure 4.29, however, if the signal variation within the reconstructed perfusion cycle is not sinusoidal in shape, then this analysis method can be unreliable.

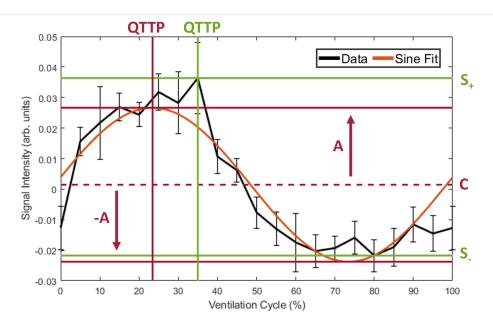


Figure 4.29: Demonstration of methods used to determine local weighted perfusion values, for the two methods discussed in this work. The local signal variation with respect to the reconstructed perfusion cycle is plotted as a black line, with error bars representing the standard error calculated during the reconstruction process. A sinusoidal curve is also fitted to the data (Orange line). S_+ , S_- and QTTP values are determined from the maximal change in signal (Green lines), and using the sinusoidal fitting (Red lines).

Example perfusion weighted maps generated using the methods discussed are demonstrated for a typical acquisition from a free-breathing healthy volunteer. Each map is accompanied by a histogram to visualise the spread of the data. The image series used to reconstruct these maps was acquired on the 0.5 T Upright scanner, with acquisition/processing parameters detailed

in Table 4.5 and 4.6.

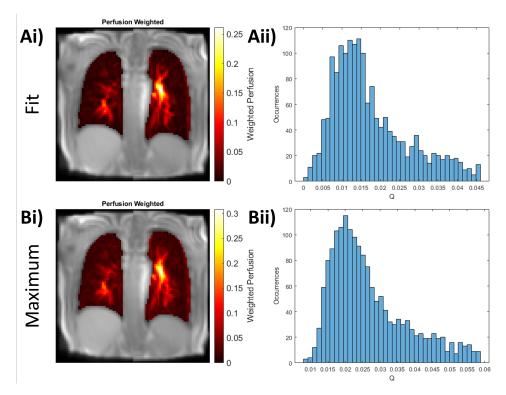


Figure 4.30: Comparison of analysis methods used to determine the voxel-wise perfusion generated from the PREFUL processing. Maps of the perfusion weighting (i) and associated histograms (ii) are displayed for the sinusoidal fitting (A) and maximum difference (B) methods. Underlying data was captured using a Coronal 2D GRE acquisition on the 0.5 T Upright scanner, whilst the participant was in the supine posture, breathing freely.

In the above examples, the maximal change method appears to have a greater SNR, however this work uses the sinusoidal fitting method for perfusion weighting determination. Errors of the fitted parameters are calculated based on the standard error of the weighted fit parameters.

4.10.2.2 Perfusion Time-To-Peak

The QTTP may provide insight into the perfusion filling dynamics of the lung. This metric is referenced to the peak value of the cardiac marker, in this case the intensity of the aorta. For this PREFUL reconstruction, this

is located at 25% of the perfusion cycle. Similar to VTTP determination, the simplest method to calculate the voxel-wise QTTP is to isolate the frame of greatest signal within the reconstructed perfusion cycle. However, this method is susceptible to errors due to the small amount of included data. Furthermore, it creates a discretised result, demonstrated in Figure 4.31. Alternatively, the fitted phase term from Equation 4.31 can be used to calculate the fractional QTTP of each voxel:

$$QTTP_{\%} = 100\% \times \frac{\left|\frac{\pi}{2} - \phi_{Perf}\right|_{0}^{2\pi}}{2\pi}$$
 (4.33)

This method uses all frames of the reconstructed perfusion cycle, theretofore may be less susceptible to noise. A comparison of each analysis method for a typical dataset is provided in Figure 4.31.

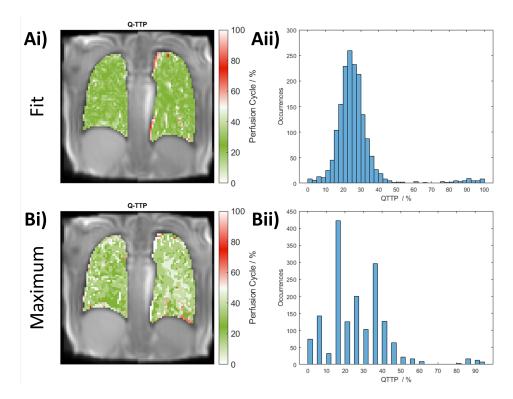


Figure 4.31: Comparison of analysis methods used to determine the voxel-wise QTTP generated from the PREFUL processing. Maps of the QTTP (i) and associated histograms (ii) are displayed for the sinusoidal fitting (A) and maximum difference (B) methods. Underlying data was captured using a Coronal 2D GRE acquisition on the 0.5 T Upright scanner, whilst the participant was in the supine posture, breathing freely.

This PREFUL pipeline utilises the sinusoidal fitting method for calculating the local QTTP. Voxel-wise errors are calculated based on the standard error of the weighted fit parameters.

4.10.3 Vessel Segmentation

While primary vessels can appear appealing to view in perfusion-weighted images, the signal variation in voxels associated with the primary vessels does not represent local pulmonary blood-flow[114]. Therefore, these should be removed from both ventilation and perfusion maps for accurate analysis. This work utilises a home-built k-means clustering based

approach to segment the primary vessels based on their elevated signal intensity compared to surrounding tissue. Typically, this is conducted on the perfusion weighted maps, as these provide the greatest vessel-to-parenchyma contrast. However, for acquisitions where perfusion processing was unsuccessful, it can be applied to the average of all registered images. A k-means clustering segmentation algorithm[115] is applied to the data. Starting with two clusters within the lung, the algorithm is applied with a successively greater number of clusters (up to 20). For each segmentation, the cluster encompassing the voxels with the lowest average intensity is assumed to represent the lung parenchyma. All other voxels are assigned to the vessels. The optimum number of clusters is automatically chosen based on a user selected threshold, such that greatest proportion of the lung mask is assigned to vessels, whilst not surpassing this threshold. This process is depicted in Figure 4.32.

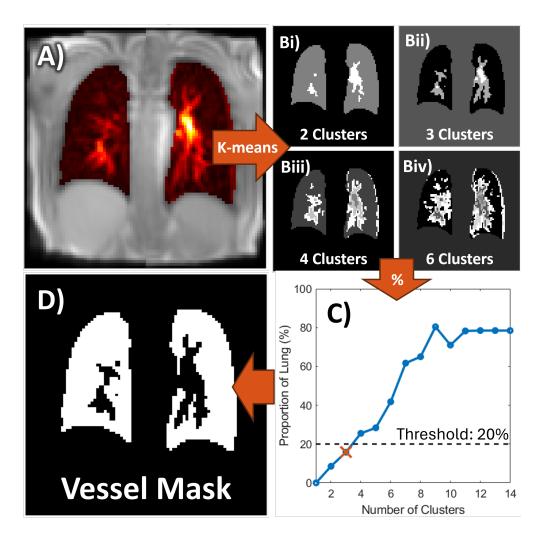


Figure 4.32: Process for segmenting the primary vessels following PREFUL processing. The perfusion weighted map (A) is segmented multiple times using a k-means clustering based approach. The resulting segmentations (B) are displayed for 2 (i), 3 (ii), 4 (iii) and 6 (iv) internal clusters. The proportion of lung taken up by the vessels is calculated for each segmentation, and plotted against the number of clusters (C). The segmentation with the greatest number of clusters, with a vessel proportion below the user-defined threshold, is used for the final vessel segmentation mask (D).

However, this process is only capable of removing the resolvable primary vessels. Smaller vessels cannot be distinguished from lung parenchyma. Therefore, while the blood in these vessels is not yet contributing to pulmonary gas exchange, it will contribute to the perfusion weighted maps. This problem exists at all field strengths, although it is exacerbated in this work due to the large voxel sizes required for imaging at low-field. In fu-

ture, this process may benefit from using shape information rather than simply intensity[116].

4.11 User Interface Implementation

The PREFUL algorithm is very complex, with many steps and many variables within each of those steps. Although many of these parameters can be estimated or standardised, there is always the possibility of the model breaking down when presented with lower quality input data. Furthermore, without adequate quality assurance, errors may go unnoticed, unless the output result is anomalous. The best results are often achieved with time-consuming optimisation of the processing parameters specific to each dataset. This is particularly important for processing data collected on low-field scanners, as the images can suffer from lower SNR compared those captured at 1.5 T[25].

A user-friendly interface was developed using MATALB's App Designer Feature, to enable manual utilisation and optimisation of the PREFUL algorithm. Five discrete stages facilitate the application of the pipeline:

1. Importing MRI data, Figure 4.33.

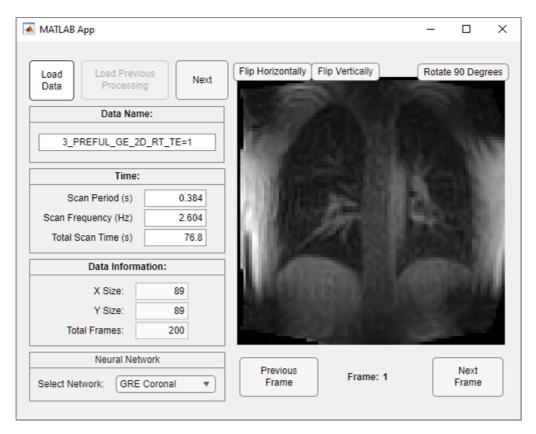


Figure 4.33: The interface used for reading the MR image series and initialising acquisition parameters. If necessary, the data can be transformed via rotation or reflection to position the lungs vertical within frame. The segmentation network is also selected using this interface.

2. Ventilation Processing, Figure 4.34.

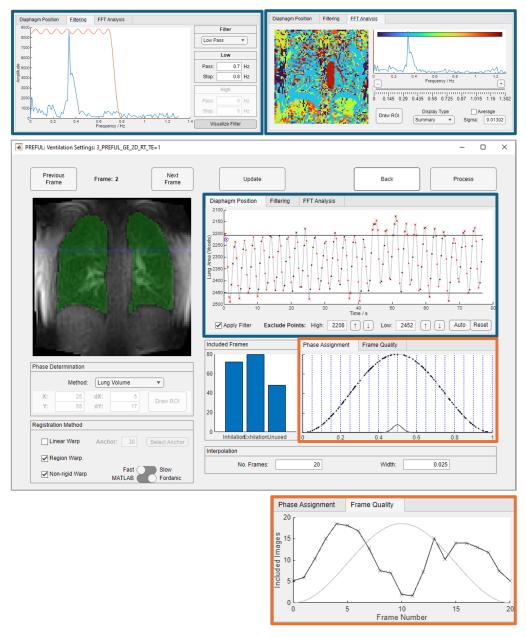


Figure 4.34: The interface for processing the ventilation and registration components of the PREFUL pipeline. Alternative windows are displayed in the **blue** and **red** boxes. Typical optimisations include selecting and thresholding the respiratory marker (and segmenting an ROI over the diaphragm if necessary), selecting the registration method, and determining ventilation cycle interpolation characteristics.

3. Ventilation Quality Assurance, Figure 4.35.

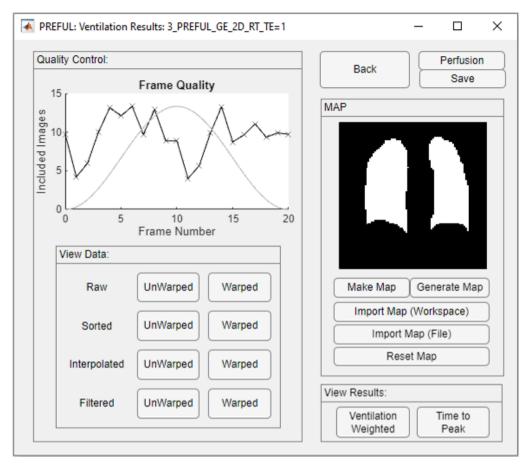


Figure 4.35: The interface for checking the quality of the ventilation processing and creating the lung segmentation mask. Processed image series can be viewed using the Display Interface, Figure 4.38.

4. Perfusion Processing, Figure 4.36.

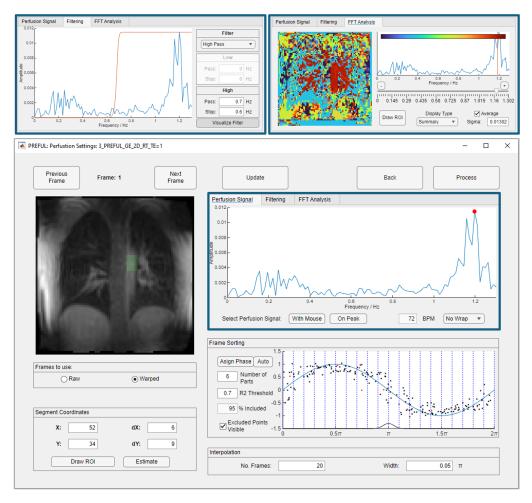


Figure 4.36: The interface for processing the perfusion components of the PREFUL pipeline. Alternative windows are displayed in the **blue** and **red** boxes. Typical optimisations include selecting the aorta ROI, by-part fitting of the signal variation due to the cardiac cycle, and determining perfusion cycle interpolation characteristics.

5. Perfusion Quality assurance, Figure 4.37.

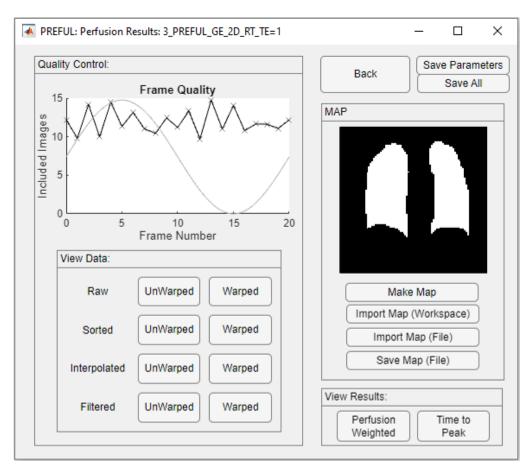


Figure 4.37: The interface for checking the quality of the perfusion processing and adapting the lung segmentation mask if necessary. Processed image series can be viewed using the Display Interface, Figure 4.38.

Additional to this, an interface was developed to visualise the dynamic image sequences. This is displayed in Figure 4.38 and primarily used for quality control.



Figure 4.38: The interface used to display the dynamic PREFUL imageseries, generally for quality control purposes. Additional functionality includes the writing of the displayed image-series to either the MATLAB workspace, a number of JPEG images, or an animated gif file.

Each interface is coded as its own independent MATLAB app which can be easily navigated back and forth through the PREFUL process, while maintaining the selection of processing parameters. This permits fast optimisation of processing parameters.

4.11.1 Automation

Following the selecting of the input data and appropriate trained CNN, almost all stages of this PREFUL pipeline provide some level of automation. These automation methods have been described throughout this text, including optimisation of the respiratory thresholding and by-part fitting of the sinusoidal signal variation due to the cardiac cycle. Furthermore, all ROI selection is automated using the segmentation network, including selection of diaphragm, the linear registration anchor and the region between the lungs containing the aorta. Therefore, it is possible to apply this algorithm with minimal operator input. With full automation, functional maps are typically generated within 1 to 2 minutes (for an 89×89 voxel, 200-image series dataset, using the computational hardware described in Appendix B). However, the greatest SNR in the resultant functional maps can generally be attained via manual optimisation of processing parameters, particularly for low SNR input data. In this case, this process can take a trained operator between 5 and 10 minutes to complete per dataset.

4.12 Results

To demonstrate the capability of this home-built pipeline to execute the PREFUL algorithm, it was applied to seven dynamic MR image-series, captured of apparently healthy individuals. This includes four datasets from the 0.5 T Upright scanner and three from the 0.5 T Bruker scanner, each with variations the in respective acquisition parameters.

4.12.1 0.5 T Upright Scanner

Four datasets were acquired on the Upright scanner for post-processing with the PREFUL pipeline. The first three of these were collected using a 2D GRE sequence, applied in each of the orthogonal imaging planes: coronal, sagittal and axial. An additional dataset was collected using a HASTE sequence in the coronal plane. Acquisition parameters for each dataset are detailed in Table 4.5 with corresponding PREFUL processing parameters in Table 4.6.

Acquisition Parameters				
Figure	4.39	4.41	4.40	4.42
Scan Type	2D GRE	2D GRE	2D GRE	2D
				HASTE
Imaging Plane	Coronal	Sagittal	Axial	Coronal
No. Images	200	200	200	200
Acquisition Time	0.385	0.385	0.385	0.405
				(0.81)**
TR (ms)	5.7	5.7	5.7	405
TE (ms)	1.2	1.2	1.2	108
Flip angle (°)	10	10	10	90
				(120)***
Matrix Size (Voxels)	64 x 64 (89			
	x 89)*	x 89)*	x 89)*	x 89)*
FOV (cm ²)	32 x 32	32 x 32	32 x 32	32 x 32
Slice Thickness (mm)	20	20	20	20

Table 4.5: Acquisition parameters for the data acquired on the 0.5 T Upright Scanner. * Data zero-filled to 89 x 89 voxels. ** Data acquired for 0.405s, followed by a delay of the same length for spin de-phasing. *** Excitation pulse: 90° , refocusing pulse: 120° , echo spacing of 8 ms for an echo train of 46.

PREFUL Processing Parameters				
Figure	4.39	4.41	4.40	4.42
Scan Type	2D GRE	2D GRE	2D GRE	2D HASTE
Imaging Plane	Coronal	Sagittal	Axial	Coronal
Registration Method	Region,	Region,	Region,	Region, Fors-
	Forsberg	Forsberg	Forsberg	berg
Respiratory marker	2D Area	2D Area	2D Area	2D Area
Respiratory filtering	Low Pass	Low Pass	Low Pass	Low Pass
method (Hz)	Pass: 0.4	Pass: 0.4	Pass: 0.4	Pass: 0.3
	Stop: 0.5	Stop: 0.5	Stop: 0.5	Stop: 0.4
No. interp. ventilation frames	20	20	20	20
Ventilation interp. width	0.05π	0.05π	0.05π	0.05π
Cardiac Aliasing	No	Yes	No	Yes
Cardiac filtering	High Pass	High Pass	High Pass	Band Pass
method (Hz)	Pass: 0.5	Pass: 0.5	Pass: 0.5	Pass: 0.15, 0.4
	Stop: 0.4	Stop: 0.4	Stop: 0.4	Stop: 0.2, 0.3
Cardiac Fitting R2 Thresh.	0.7	0.7	0.7	0.7
No. interp. perfusion frames	20	20	20	20
Perfusion interp. width	0.05π	0.05π	0.05π	0.05π

Table 4.6: PREFUL processing parameters for the data acquired on the 0.5 T Upright Scanner.

The first results presented in this work were acquired of a participant in the supine posture in the coronal plane using a 2D GRE sequence. Both ventilation and perfusion processing were successful, with a respiratory correlation coefficient of 0.9776 and 79.95% of the images included in the perfusion processing. The functional maps generated from this dataset are presented in Figure 4.39.

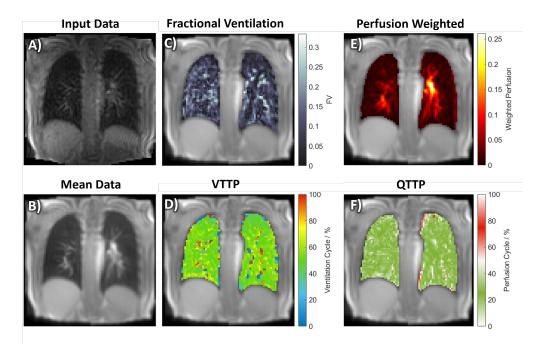


Figure 4.39: Typical results from PREFUL processing of Coronal 2D GRE image data, acquired on the 0.5 T Upright Scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D), perfusion weighted (E) and QTTP (F). Data acquired whilst the participant was supine and breathing freely. Mean FV = 0.1 \pm 0.06.

The next results presented were acquired of a participant in the supine posture in the axial plane using a 2D GRE sequence. Both ventilation and perfusion processing were successful, with a respiratory correlation coefficient of 0.963 and 95.79% of the images included in the perfusion processing. The functional maps generated from this dataset are presented in Figure 4.40.

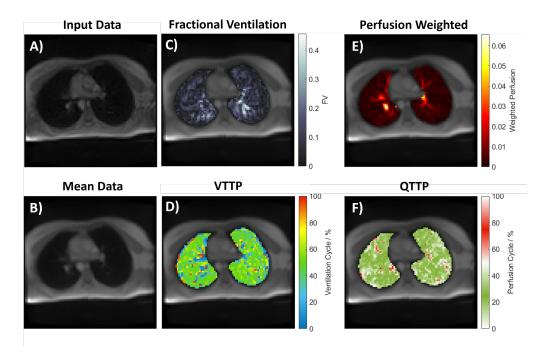


Figure 4.40: Typical results from PREFUL processing of Axial 2D GRE image data, acquired on the 0.5 T Upright Scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D), perfusion weighted (E) and QTTP (F). Data acquired whilst the participant was supine and breathing deeply. Mean FV = 0.09 \pm 0.06.

Image-series were also acquired of a participant in the supine posture in the sagittal plane using a 2D GRE sequence. Both ventilation and perfusion processing were successful, with a respiratory correlation coefficient of 0.9702 and 81.58% of the images included in the perfusion processing. The functional maps generated from this dataset are presented in Figure 4.41:

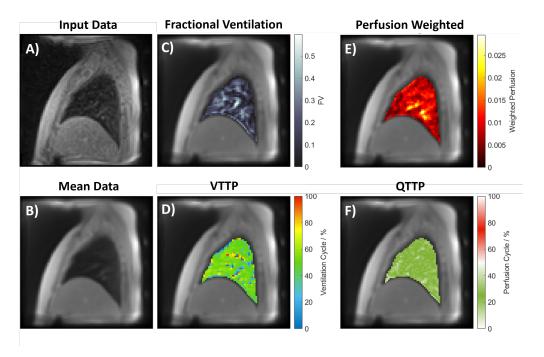


Figure 4.41: Typical results from PREFUL processing of Sagittal 2D GRE image data, acquired on the 0.5 T Upright Scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D), perfusion weighted (E) and QTTP (F). Data acquired whilst the participant was supine and breathing freely. Mean FV = 0.17 \pm 0.01.

The final results presented from the 0.5 T upright scanner were acquired of a participant in the sitting posture, used a 2D HASTE sequence in the coronal plane. Both ventilation and perfusion processing were successful, with a respiratory correlation coefficient of 0.5552 and 85.79% of the images included in the perfusion processing. The functional maps generated from this dataset are presented in Figure 4.42.

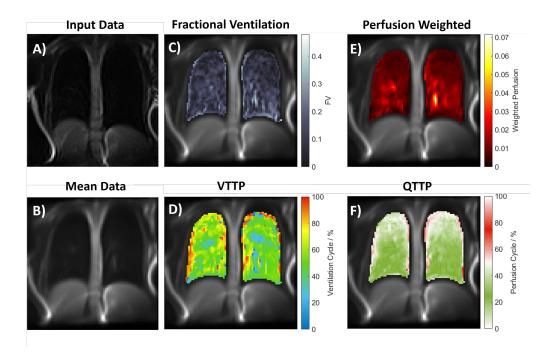


Figure 4.42: Typical results from PREFUL processing of Coronal 2D HASTE image data, acquired on the 0.5 T Upright Scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D), perfusion weighted (E) and QTTP (F). Data acquired whilst the participant was sitting and breathing freely. Mean FV = 0.11 ± 0.06 .

4.12.2 0.5 T Bruker Scanner

All datasets collected using the 0.5 T Bruker scanner were acquired in the coronal plane using a 2D GRE sequence with varying acquisition parameters. Acquisition parameters for each dataset are detailed in Table 4.7 with corresponding PREFUL processing parameters in Table 4.8.

Acquisition Parameters			
Figure	4.43	4.44	4.45
Scan Type	2D GRE	2D GRE	2D GRE
Imaging Plane	Coronal	Coronal	Coronal
No. Images	200	200	200
Acquisition Time	0.341	0.845	0.341
TR (ms)	3.21	9.6	3.21
TE (ms)	2.14	2.7	2.14
BW (Hz)	200	52.6	200
FOV (cm ²)	22 x 25	27.5 25	22 x 25
Matrix Size (Voxels)	44 x 50	88 x 80	128 x 128
Slice Thickness (mm)	NA*	15	10

Table 4.7: Acquisition parameters for the data acquired on the 0.5 T Bruker Scanner. * Non-slice-selective acquisition.

PREFUL Processing Parameters			
Figure	4.43	4.44	4.45
Scan Type	2D GRE	2D GRE	2D GRE
Imaging Plane	Coronal	Coronal	Coronal
Registration Method	Region,	Region,	Linear,
	Forsberg	Forsberg	Forsberg
Respiratory marker	2D Area	2D Area	Diaphragm
			Position
Respiratory filtering method	Low Pass	Low Pass	Low Pass
(Hz)	Pass: 0.45	Pass: 0.3	Pass:
	Stop: 0.55	Stop: 0.35	Stop:
No. interp. ventilation frames	20	20	20
Ventilation interp. width	0.05π	0.05π	0.05π
Cardiac Aliasing	No	Yes	No
Cardiac filtering method	High Pass	High Pass	High Pass
(Hz)	Pass: 0.8	Pass: 0.35	Pass:
	Stop: 0.7	Stop: 0.3	Stop:
Cardiac Fitting R2 Thresh.	0.7	0.7	0.7
No. interp. perfusion frames	20	20	20
Perfusion interp. width	0.05π	0.05π	0.05π

Table 4.8: PREFUL processing parameters for the data acquired on the 0.5 T Bruker Scanner.

The first results from presented in this work the 0.5 T Bruker scanner were acquired using a non-slice-selective 2D GRE sequence. Both ventilation and perfusion processing were successful, with a respiratory correlation coefficient of 0.7045 and 88.95% of the images included in the perfusion processing. The functional maps generated from this dataset are presented

in Figure 4.43.

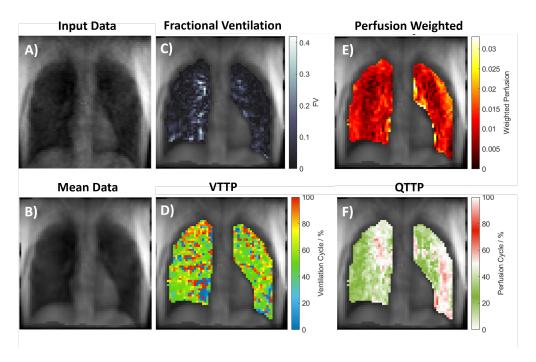


Figure 4.43: Results from PREFUL processing of Coronal 2D GRE nonslice selective image data acquired using the Bruker 0.5 T scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D), perfusion weighted (E) and QTTP (F). Mean $FV = 0.13 \pm 0.08$.

The next results presented in this work were acquired using a slice-selective 2D GRE sequence, acquired at a higher resolution than the previous result, although at a reduced imaging frequency. Therefore, the cardiac cycle was not adequately sampled and only the ventilation processing was successful, with a respiratory correlation coefficient of 0.7673. The functional maps generated from this dataset are presented in Figure 4.44.

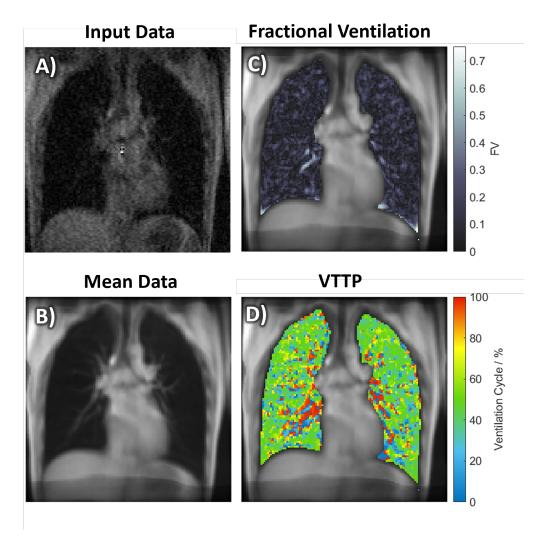


Figure 4.44: Results from PREFUL processing of Coronal 2D GRE nonslice selective image data acquired using the Bruker 0.5 T scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D). Perfusion processing was not successful for this dataset. Mean FV = 0.13 ± 0.08 .

The final results presented in this work were acquired using a slice-selective 2D GRE sequences capture at a higher frequency than the previous result. The input image series exhibited a low SNR, therefore only the ventilation processing was successful with a respiratory correlation coefficient of 0.9478. Furthermore, due to the low SNR of the input data, it was not possible to accurately segment each image in the series. Therefore, the diaphragm position was used as the respiratory marker and the linear warp method was used for the feature-based registration step. The functional maps generated

from this dataset are presented in Figure 4.45.

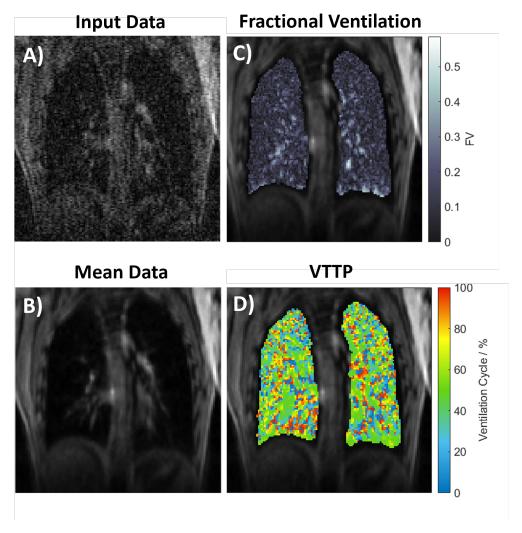


Figure 4.45: Results from PREFUL processing of Coronal 2D GRE non-slice selective image data acquired using the Bruker 0.5 T scanner. Including a frame of the input data (A), the average registered frame (B), fractional ventilation (C), VTTP (D). Perfusion processing was not successful for this dataset. Mean FV = 0.06 ± 0.06 .

4.13 Discussion

4.13.1 Quality of Functional Maps

4.13.1.1 2D Gradient Echo Experiments

When applied with GRE sequences this PREFUL algorithm can create visually pleasing results. This includes for both ventilation and perfusion on the 0.5 T Upright scanner (Figures 4.39, 4.40 and 4.39). Applied to the 0.5 T Bruker data (Figures 4.43, 4.44 and 4.45), ventilation was generally deemed successful, whereas perfusion processing often failed. This may be due to the apparent lower SNR associated with these datasets.

While processing may have been successful for these datasets, the results have not been validated against other modalities. Therefore, it is not possible to determine if the results' accuracy reflect regional ventilation and perfusion in the lung. This pipeline has been employed to study regional variation in lung function with respect to participant posture, yielding results that reflected established theory (Chapter 5). However, in future this pipeline must be validated compared to established methods of lung function investigation.

4.13.1.1.1 Rib artifacts in non-slice selective acquisitions

The non-slice selective acquisition (Figure 4.43) exhibited artefacts that resembled lateral bands across the image. These horizontal bands of delayed VTTP are most easily viewed in Figure 4.43 and represent out-of-phase regions. When using just the linear region-based registration method, these

appeared to 'bounce' up and down in the registered sequence. These artefacts were therefore determined to stem from the rib cage. The rib cage remains approximately stationary within this projection throughout the respiratory cycle, although it is subject to registration along with the rest of the lung. Therefore, although appearing to move vertically in the reconstructed ventilation cycle, this artefact will likely cause significant errors in the results as the difference in signal intensity of the ribs to intercostal muscles is likely greater than the signal change associated with respiration.

4.13.1.1.2 Through plane movement in axial scans

The axial acquisitions using the 0.5 T upright scanner appeared to yield good results. However, these scans are particularly sensitive to throughplane movement, more so than the coronal and sagittal planes due to the large displacement of the diaphragm. In an effort to avoid regions of large vertical displacement of lung parenchyma, this imaging plane was placed high in the chest (see Figure 5.7). However, no investigation was undertaken to check the magnitude of the parenchymal movement in this region. Therefore, until these investigations have taken place, results in the axial plane may be unreliable for functional lung assessment.

4.13.1.2 2D HASTE Experiments

This algorithm was also applied to data acquired using a 2D HASTE sequence. Although the presented dataset was processed successfully for ventilation and perfusion, this was an outlier. In Chapter 5, the HASTE scans were removed from the protocol due to their unreliability for PREFUL processing. This may be due to a combination of factors. HASTE scans are typically T2* weighted[117] and may therefore be more sensitive to spin

density variations due to blood rather than tissue density tissue due to the comparatively long T2 of blood[12]. This is reflected in the maximum component analysis used for guiding the ROI over the aorta (Figure 4.8) as when plotted this way the signal in the whole lung appears to vary the most due to cardiac variations. However, the low imaging rate of the HASTE scans employed in this work (Table 4.5) means it is difficult to adequately sample the cardiac cycle, even when accounting for signal aliasing. Furthermore, it is common for the signal cardiac peak to alias to frequencies similar to the respiratory cycle. If this occurs, then it is difficult to isolate each of these components via filtering. For the HASTE results presented in this work, perfusion filtering was performed using a band-pass filter (Table 4.6). Typically, this means that either the ventilation processing will fail due to significant contributions from the cardiac cycle, or the cardiac processing will fail due to inadequate temporal sampling. Although this acquisition sequence has been employed previously for PREFUL at 1.5T[41], it appears that the HASTE acquisitions employed at 0.5 T in this work are not suitable for PREFUL processing. Further optimisation may yield useful results at this field strength, however this was not completed in this work.

4.13.2 Application with Pathology

The development of this pipeline was conducted using datasets captured of only healthy volunteers. Furthermore, no attempt has been made in this work to characterise pathology, or assign regional ventilation or perfusion defects. Therefore, before deployment of this pipeline to investigate lung defects a number of investigations must take place to ensure the validity of the results. Initially, a number of assumptions have been made in this

work based on the dynamics of healthy lungs. Primarily, these are:

- Diaphragm position and 2D lung area are appropriate indicators of respiratory state and the state of adjacent frames can be used to determine if frames correspond to the inhalation or exhalation regime.
- The variation in signal due to the cardiac cycle is assumed to be sinusoidal for cardiac phase assignment.
- The variations in signal in the reconstructed ventilation and perfusion cycle can each be fitted to a sinusoidal curve for voxel-wise functions analysis.
- The displacement between neighbouring respiratory states can be modelled as linear for the batch-wise registration method.

In addition to these assumptions, the accuracy of the trained CNNs must be validated for segmenting diseased lungs. Following the validation of these processing assumptions, this pipeline must be employed for investigations of pathology, and the resulting functional maps analysed to determine if this method is sensitive to the changes associated with changes in lung dynamics. Furthermore, it must be determined if sufficient contrast exists between healthy and diseased tissue for meaningful assignment of ventilation and perfusion defects. This process will require validation using additional imaging modalities.

4.13.3 Processing Functional Dynamics

4.13.3.1 Perfusion Dynamics

During the perfusion processing, the signal variation due to the cardiac cycle is assumed to be sinusoidal. This assumption is made twice, at the cardiac phase assignment step and and during the final analysis. This assumption may provide good approximation of perfusion dynamics, especially for fitting low SNR data, however it may not be correct. Indeed, in Figure 4.29, the signal variation appears as an initial pulse followed by a period with little variation. This is similar to to the variation in pulmonary artery pressure due to a single heartbeat. Further iterations of this work should explore different fitting models for perfusion.

This perfusion analysis method is based purely on signal intensity, without any data specific normalisation step. Therefore, it is difficult to compare the magnitude of these results between individuals. It is possible to normalise the perfusion weighted results[32]:

$$Q_N = \frac{Q_w}{S_{Blood}} \tag{4.34}$$

where Q_N is the normalised perfusion, and S_{Blood} is the signal from a voxel containing only blood. This value can be used to calculate the Quantified Perfusion $(Q_{Quant})[32]$:

$$Q_{Quant} = \frac{Q_N}{2t_{perf}} \tag{4.35}$$

where t_{perf} is the time between two heart beats. This metric represents a

fully quantitive lung perfusion biomarker[118]. Typically, the signal from a voxel sampling the aorta is used for $S_{Blood}[32]$. However, beyond effects due to differences in relaxation mechanics of blood compared to lung parenchyma, the larger voxel sizes associated with scanning at low field in this study were determined to complicate this method. Often the aorta was subject to partial volume effects. Therefore, no further normalisation process was undertaken for perfusion in this work, and comparison between participants during experimental studies should use relative change within the perfusion weighted map.

4.13.3.2 Ventilation Dynamics

The calculated values of the fractional ventilation are dependent on the target frame chosen during the registration step. This is eloquently demonstrated by Klimes et al.[119]. Instead, regional ventilation can provide a better representation of local ventilation, independent of the lung state selected for registration. Fractional ventilation is equivalent to regional ventilation, only when the target frame is the point of maximal inspiration. Regional ventilation is calculated using[119]:

$$RVent = \frac{S_{mid}}{S_{in}} - \frac{S_{mid}}{S_{ex}} \tag{4.36}$$

where S_{mid} is voxel-wise average respiratory signal. Applying this equation to the sinusoidal fitting method employed in this work, $S_{mid} = C$ in equation 4.24. Therefore, adapting Equation 4.36 or use with the sinusoidal fitting method:

$$RVent = \frac{2CA_{Vent}}{C^2 + A_{Vent}^2} \tag{4.37}$$

Therefore, future iterations of this pipeline should introduce calculations of the regional ventilation. However, even with this alteration, the magnitude of the observed ventilation is highly dependent on subject breathing patterns during the acquisition, and substituent thresholding of the respiratory marker. Therefore, ventilation investigations using PREFUL typically used flow-loop metrics as marker for heath and disease. Future iterations of this work should explore the use of this metric in conjunction with low-field imaging.

For some lung regions, highlighted in the VTTP maps, the lung ventilation appeared to be 180 degrees out of phase with the respiratory cycle. That is, as the lung expanded, the MR signal in that region increased, thus the air content was assumed to decrease. This paradoxical result may be indicative of regions of competitive ventilation or gas trapping [65, 120]. However, this phenomenon was observed for apparently healthy volunteers, therefore, it may instead be due to errors in registration step. For example, if a vessel is registered incorrectly, those voxel to which it appears to move to upon full participant inhalation, will exhibit this behaviour.

Finally, the magnitude of the fractional ventilation is dependent on the breathing patterns of the participant and the selection of the positioning thresholding. This is a problem inherent to this free breathing technique and contributes to the use of flow loop analysis for ventilation assessment instead of the fractional ventilation.

4.13.4 Comparison to Literature

PREFUL is typically employed in conjunction with data collected at 1.5 T[25, 32, 33, 51, 53]. However, in this study, no repeatability experi-

ments have been conducted at this field strength. Therefore, to validate the ventilation results, the functional parameters acquired in this study were compared to those from healthy participants in the literature.

As discussed above, modern PREFUL investigations typically use regional ventilation as a parameter to explore ventilation. Therefore, to compare with these studies, this value was calculated (Equation 4.36) and the mean taken for each dataset presented in this chapter. The median values for the fractional ventilation and regional ventilation for each scanner are presented in Table 4.9.

Mean Values of Helathy FV and RVent			
Scanner	FV	RVent	
0.5 T Open (N=4)	0.11 (0.10-0.14)	0.12 (0.11-0.16)	
0.5 T Bruker (N=3)	0.13 (0.08-0.13)	0.14 (0.08-0.14)	
All (N=7)	0.11 (0.09-0.13)	0.12 (0.10-0.14)	

Table 4.9: Mean values of FV and RVent for the healthy participants investigated in this chapter. Quoted results represent the median values with 25th and 75th percentiles in parentheses.

Values for the regional ventilation from literature are presented in Table 4.10 for healthy volunteers investigated at a range of field strengths in the supine posture.

Literature Values for Healthy RVent			
0.55 T	1.5 T	3.0 T	
0.13 (0.09-0.16)[52]	0.14 (0.10-0.16)[53]	0.18 (0.14-0.27)[53]	
	0.14 (0.11-0.17)[51]		
	0.11 (0.09-0.21)[33]		
	0.18 (0.14-0.24)[32]		
	0.17 (0.13-0.25)[32]		

Table 4.10: Literature values for the average RVent within the lung for healthy volunteers investigated using PREFUL in 2D at different field strengths. Quoted results represent the median values with 25th and 75th percentiles in parentheses.

The values of the regional ventilation acquired in this study appear to be similar to those acquired at different field strengths. However, no statistical analysis has been performed. Furthermore, the values acquired in this study represent only a small number of participants, so a wider repeatability study should be performed in the future.

It was not possible to compare the perfusion dynamics acquired in this work to values from the literature, as no data normalisation was performed in this chapter. This is due to the partial volume effects of the blood filled voxels as discussed above. In the future, this normalisation step should be implemented and repeatability studies performed.

4.13.5 Drawbacks of PREFUL

PREFUL is a highly complex algorithm. The high number of optimisable parameters associated with image acquisition and post-processing (which are often interdependent) can lead to a high degree of variability in results. Therefore, standardisation of these processes in conjunction with

with healthy and diseased lungs may be necessary.

The PREFUL algorithm derives local ventilation and perfusion information directly from variations in voxel intensity with the respiratory and cardiac cycles, respectively[25]. This method assumes all signal variation is due to variations in air and blood flow that contribute to gas exchange. However, this assumption is likely not valid[114]. No distinction is made between areas of anatomical dead-space or anatomical shunt, such as the bronchioles or pulmonary arteries. Furthermore, no distinction is given to the initially inspired air occupying areas of anatomical dead-space, which had been expelled during the previous expiration. These factors may be accounted for with appropriate models of these functions; however, this is was not explored this work.

All signal-based FD methods assume consistent proportionality of signal intensity to spin density throughout each image and the image-sequence as a whole. Therefore, these methods typically omit influences from T1, T2/T2* variations, through-plane motion of parenchyma (for 2D acquisitions) and intra-voxel motion of excited spins[74].

- T2* is known to vary based on parenchyma density. Varying signal influences can therefore manifest regionally[13, 121–124] and depending on the respiratory phase[122, 123].
- Motion of the lung parenchyma orthogonal to the 2D imaging plane during breathing can result in differing features of the lung imaged at different points of the respiratory cycle. Whilst accurate registration can account for in-plane lung motion, these algorithms cannot account for through-plane motion of pulmonary features. This phenomenon is expected to occur to some degree for all imaging planes, however with the greatest magnitude for axial acquisitions, due to the

relatively high degree of movement of the diaphragm during normal respiration. This error can be reduced using 3D acquisitions[51].

• Inflow and outflow of excited spins to a voxel can cause signal attenuation if the rate of this is a similar timescale as the echo time in gradient echo experiments. Blood flow varies significantly within the lung[65]; measured in the main pulmonary artery up to approximately 100 cm/s at peak systole[125] and up to 1 cm/s in the pulmonary capillaries[126]. For the approximately 2 ms TE typical of the GRE sequences employed in this work (Tables 4.5 and 4.7), this would result in approximately 2 mm and 0.02 mm of spin displacement between excitation and echo, respectively. This is smaller than the 5 mm × 5 mm typically employed in this work, therefore this effect is assumed to have a negligible impact on imaging lung features downstream of the main pulmonary artery.

4.13.6 Future Work

Further iterations of this pipeline may include the automation of the remaining manual processing steps, such as improved segmentation of the aorta for cardiac phase assignment. An additional CNN could be employed for this purpose. This pipeline is only designed to process 2D single slice data, therefore it could be adapted to process 2D multi-slice or 3D acquisitions. The processing methods employed in this work must be validated both for accuracy and repeatability. This could include repeated participant investigations, multicentre studies and validation with additional imaging modalities. The capability of this pipeline for investigating diseased lungs can also be validated in the same way.

This work has explored PREFUL using data sets captured with 2D GRE

and HASTE acquisitions. Previous deployment of the PREFUL algorithm at 0.55 T used a bSSFP sequence[52], therefore this could be explored in future work. Additionally, the capability of this pipeline for processing data captured at 1.5 T could also be explored. Finally, this work has investigated ventilation dynamics using the FV, VTTP and Jacobian determinant, and perfusion dynamics using the perfusion weighted and QTTP maps. Further work could investigate if additional metrics typical of PREFUL are possible with the SNR limitations observed at this field strength, including flow loop analysis and pulmonary pulse wave transit time[32].

4.14 Conclusions

This work demonstrates the application of PREFUL methodology[25] for functional lung imaging at 0.5 T. The home-built pipeline was employed with two different 0.5 T systems with varying hardware configurations, indicating the robustness and generalizability of the developed methods. Demonstration of this algorithm at field strengths below 1.5 T, typical of PREFUL acquisitions[25, 32], contributes to the democratisation of this advanced pulmonary MRI technique. Furthermore, the compatibility with open-bore scanners, which are typically low field, enables investigation of lung function in various postural conditions, opening new research avenues in gravitational physiology and postural effects on pulmonary function (Chapter 5).

This PREFUL pipeline possesses several adaptations from the original algorithm [25] to address the challenges associated with low field imaging, including reduced SNR, dynamic image acquisition rates and image resolution, compared to equivalent studies at 1.5 T. These adaptations are

integrated into several stages of the PREFUL processing:

• Phase Assignment

For accurate assignment of the respiratory state of each frame within low resolution datasets, this algorithm is capable of determining the sub-pixel diaphragm position using an edge detection method. Alternatively, lung volume can be used for determining the respiratory state; this method may be less sensitive to alterations in participant breathing regimes associated with different postures, a capability available when scanning using open scanners. For perfusion, the low imaging rates in this study often meant the cardiac cycle was undersampled. This pipeline permitted the determination of the cardiac cycle via sinusoidal fitting of a cardiac cycle that doesn't satisfy the Nyquist condition.

• Image Registration

A dual-stage registration approach was implemented to reduce the prevalence of errors from non-rigid registration algorithms when large voxel displacements are required. This is achieved by using feature-based registration for bulk registration, followed by non-rigid intensity-based registration for fine-scale registration. Additionally, a batchwise method was implemented whereby only reconstructed frames are subjected to intensity-based non-rigid registration. This reduced errors associated with the low SNR of lung parenchyma, as well as signal variations due to the cardiac cycle.

• Functional Analysis

Functional analysis methods were implemented that utilised signal variation throughout the full reconstructed ventilation and perfusion cycles. This method is less sensitive to noise than the original twopoint analysis methods[25] and returns non-discreet time-to-peak results.

Additionally, five CNNs were trained, achieving >95% accuracy for lung segmentation tasks across multiple acquisition types and imaging planes, enabling automated respiratory phase determination and assisting with ROI selection of the diaphragm and aorta. While the pipeline is capable of processing data with minimal user-input, the algorithm is integrated with an intuitive interface for optimisation of processing parameters and quality assurance of each processing step.

Future investigations should focus on validation of this algorithm against established modalities, as well as repeatability and reproducibility studies. Furthermore, this pipeline was developed and tested exclusively on healthy volunteers, thus it requires validation in pathological conditions to test its capability for functional defect assignment. The successful completion of these validation studies would position this technology for clinical translation and broader research adoption.

4.15 References

- Tian, Y. & Nayak, K. S. New clinical opportunities of low-field MRI: heart, lung, body, and musculoskeletal. *Magnetic Resonance Materials in Physics, Biology and Medicine 2023 37:1* 37, 1–14. doi:10.1007/S10334-023-01123-W (Oct. 2023).
- Candemir, S. & Antani, S. A review on lung boundary detection in chest X-rays. *International Journal of Computer Assisted Radiology* and Surgery 14, 563. doi:10.1007/S11548-019-01917-1 (Apr. 2019).
- Roach, P. J., Gradinscak, D. J., Schembri, G. P., Bailey, E. A., Willowson, K. P. & Bailey, D. L. SPECT/CT in V/Q Scanning. Seminars in Nuclear Medicine 40, 455–466. doi:10.1053/J.SEMNUCLMED. 2010.07.005 (Nov. 2010).
- Israel, O. et al. Two decades of SPECT/CT the coming of age of a technology: An updated review of literature evidence. European Journal of Nuclear Medicine and Molecular Imaging 2019 46:10 46, 1990–2012. doi:10.1007/S00259-019-04404-6 (July 2019).
- Sulieman, A., Almogren, K. S. & Tamam, N. Radiation exposure estimation in pediatric patients during computed tomography imaging procedures. *Radiation Physics and Chemistry* 204, 110653. doi:10.1016/J.RADPHYSCHEM.2022.110653 (Mar. 2023).
- Mathews, J. D. et al. Cancer risk in 680 000 people exposed to computed tomography scans in childhood or adolescence: data linkage study of 11 million Australians. BMJ 346. doi:10.1136/BMJ.F2360 (May 2013).

- 7. Bauman, G. et al. Non-contrast-enhanced perfusion and ventilation assessment of the human lung by means of fourier decomposition in proton MRI. Magnetic Resonance in Medicine 62, 656–664. doi:10. 1002/mrm.22031 (2009).
- 8. Mai, V. M. & Berr, S. S. MR Perfusion Imaging of Pulmonary Parenchyma Using Pulsed Arterial Spin Labeling Techniques: FAIRER and FAIR. J. Magn. Reson. Imaging 9, 483–487. doi:10.1002/(SICI)1522-2586(199903)9:3<483::AID-JMRI18>3.0.CO;2-(1999).
- Hatabu, H., Tadamura, E., Levin, D. L., Chen, Q., Li, W., Kim, D., Prasad, P. V. & Edelman, R. R. Quantitative assessment of pulmonary perfusion with dynamic contrast- enhanced MRI. Magnetic Resonance in Medicine 42, 1033–1038. doi:10.1002/(SICI)1522-2594(199912)42:6<1033::AID-MRM7>3.0.CO;2-7, (1999).
- Edelman, R. R., Hatabu, H., Tadamura, E., Li, W. & Prasad, P. V. Noninvasive assessment of regional ventilation in the human lung using oxygen-enhanced magnetic resonance imaging. *Nature medicine* 2, 1236–1239. doi:10.1038/NM1196-1236 (Nov. 1996).
- 11. Wild, J. M. et al. MRI of the lung (1/3): methods. Insights into imaging 3, 345–353. doi:10.1007/S13244-012-0176-X (Aug. 2012).
- 12. Li, B., Lee, N. G., Cui, S. X. & Nayak, K. S. Lung parenchyma transverse relaxation rates at 0.55 T. *Magnetic Resonance in Medicine* **89**, 1522. doi:10.1002/MRM.29541 (Apr. 2022).
- 13. Yu, J., Xue, Y. & Song, H. K. Comparison of Lung T2* During Free-Breathing at 1.5T and 3.0T with Ultrashort Echo Time (UTE) Imaging. Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine 66, 248. doi:10.1002/MRM.22829 (2011).

- 14. Zapke, M. et al. Magnetic resonance lung function A breakthrough for lung imaging and functional assessment? A phantom study and clinical trial. Respiratory Research 7. doi:10.1186/1465-9921-7-106, (Aug. 2006).
- 15. Bauman, G. et al. Lung ventilation- and perfusion-weighted Fourier decomposition magnetic resonance imaging: In vivo validation with hyperpolarized 3He and dynamic contrast-enhanced MRI. Magnetic Resonance in Medicine 69, 229–237. doi:10.1002/MRM.24236, (2013).
- 16. Capaldi, D. P., Sheikh, K., Eddy, R. L., Guo, F., Svenningsen, S., Nair, P., McCormack, D. G. & Parraga, G. Free-breathing Functional Pulmonary MRI: Response to Bronchodilator and Bronchoprovocation in Severe Asthma. *Academic Radiology* 24, 1268–1276. doi:10.1016/J.ACRA.2017.04.012 (Oct. 2017).
- Capaldi, D. P. et al. Free-breathing Pulmonary 1H and Hyperpolarized 3He MRI: Comparison in COPD and Bronchiectasis. Academic Radiology 22, 320–329. doi:10.1016/J.ACRA.2014.10.003 (Mar. 2015).
- 18. Kaireit, T. F., Gutberlet, M., Voskrebenzev, A., Freise, J., Welte, T., Hohlfeld, J. M., Wacker, F. & Vogel-Claussen, J. Comparison of quantitative regional ventilation-weighted fourier decomposition MRI with dynamic fluorinated gas washout MRI and lung function testing in COPD patients. *Journal of Magnetic Resonance Imaging* 47, 1534–1541. doi:10.1002/JMRI.25902; CTYPE: STRING: JOURNAL (June 2018).
- 19. Lederlin, M., Bauman, G., Eichinger, M., Dinkel, J., Brault, M., Biederer, J. & Puderbach, M. Functional MRI using Fourier decomposition of lung signal: Reproducibility of ventilation- and perfusion-

- weighted imaging in healthy volunteers. European Journal of Radiology 82, 1015–1022. doi:10.1016/J.EJRAD.2012.12.003 (June 2013).
- 20. Bauman, G., Puderbach, M., Heimann, T., Kopp-Schneider, A., Fritzsching, E., Mall, M. A. & Eichinger, M. Validation of Fourier decomposition MRI with dynamiccontrast-enhanced MRI using visual and automated scoring ofpulmonary perfusion in young cystic fibrosis patients. European Journal of Radiology 82, 2371–2377. doi:10.1016/j.ejrad.2013.08.018 (2013).
- 21. Sommer, G., Bauman, G., Koenigkam-Santos, M., Draenkow, C., Heussel, C. P., Kauczor, H. U., Schlemmer, H. P. & Puderbach, M. Non-contrast-enhanced preoperative assessment of lung perfusion inpatients with non-small-cell lung cancer using Fourier decompositionmagnetic resonance imaging. *European Journal of Radiology* 82. doi:10.1016/j.ejrad.2013.06.030 (2013).
- Bauman, G. et al. Pulmonary functional imaging: Qualitative comparison of fourier decomposition MR imaging with SPECT/CT in porcine lung. Radiology 260, 551–559. doi:10.1148/RADIOL.11102313, (Aug. 2011).
- 23. Bondesson, D., Schneider, M. J., Gaass, T., Kühn, B., Bauman, G., Dietrich, O. & Dinkel, J. Nonuniform Fourier-decomposition MRI for ventilation- and perfusion-weighted imaging of the lung. *Magnetic Resonance in Medicine* 82, 1312–1321. doi:10.1002/MRM.27803, (Oct. 2019).
- 24. Bauman, G. & Bieri, O. Matrix pencil decomposition of time-resolved proton MRI for robust and improved assessment of pulmonary ventilation and perfusion. *Magnetic Resonance in Medicine* **77**, 336–342. doi:10.1002/MRM.26096, (Jan. 2017).

- 25. Voskrebenzev, A., Gutberlet, M., Klimeš, F., Kaireit, T. F., Schönfeld, C., Rotärmel, A., Wacker, F. & Vogel-Claussen, J. Feasibility of quantitative regional ventilation and perfusion mapping with phase-resolved functional lung (PREFUL) MRI in healthy volunteers and COPD, CTEPH, and CF patients. Magnetic Resonance in Medicine 79, 2306–2314. doi:10.1002/mrm.26893 (2018).
- Liu, Q. H. & Nguyen, N. An accurate algorithm for nonuniform fast Fourier transforms (NUFFT's). *IEEE Microwave and Guided Wave Letters* 8, 18–20. doi:10.1109/75.650975 (Jan. 1998).
- 27. Greengard, L. & Lee, J. Y. Accelerating the nonuniform fast Fourier transform. SIAM Review 46, 443–454. doi:10.1137/S003614450343200X; JOURNAL: JOURNAL: SIREAD; PAGE: STRING: ARTICLE/CHAPTER (Aug. 2004).
- 28. Zhang, K. & Kang, J. U. Graphics processing unit accelerated non-uniform fast Fourier transform for ultrahigh-speed, real-time Fourier-domain OCT. Optics Express, Vol. 18, Issue 22, pp. 23472-23487 18, 23472-23487. doi:10.1364/0E.18.023472 (Oct. 2010).
- Lin, Y. Y., Hodgkinson, P., Ernst, M. & Pines, A. A Novel Detection–Estimation Scheme for Noisy NMR Signals: Applications to Delayed Acquisition Data. *Journal of Magnetic Resonance* 128, 30–41. doi:10.1006/JMRE.1997.1215 (Sept. 1997).
- 30. Doellinger, F. et al. Contrast agent-free functional magnetic resonance imaging with matrix pencil decomposition to quantify abnormalities in lung perfusion and ventilation in patients with cystic fibrosis. Frontiers in Medicine 11. doi:10.3389/FMED.2024.1349466, (2024).
- 31. Streibel, C. et al. Long-term pulmonary outcome of children with congenital diaphragmatic hernia: functional lung MRI using matrix-

- pencil decomposition enables side-specific assessment of lung function. European Radiology 34, 3773–3785. doi:10.1007/S00330-023-10395-8, (June 2024).
- 32. Pöhler, G. H., Klimeš, F., Behrendt, L., Voskrebenzev, A., Gonzalez, C. C., Wacker, F., Hohlfeld, J. M. & Vogel-Claussen, J. Repeatability of Phase-Resolved Functional Lung (PREFUL)-MRI Ventilation and Perfusion Parameters in Healthy Subjects and COPD Patients.

 Journal of Magnetic Resonance Imaging 53, 915–927. doi:10.1002/JMRI.27385, (Mar. 2021).
- 33. Kaireit, T. F. et al. Flow Volume Loop and Regional Ventilation Assessment Using Phase-Resolved Functional Lung (PREFUL) MRI: Comparison With 129Xenon Ventilation MRI and Lung Function Testing. Journal of Magnetic Resonance Imaging 53, 1092–1105. doi:10.1002/JMRI.27452 (Apr. 2021).
- 34. Pöhler, G. H. et al. Validation of Phase-Resolved Functional Lung (PREFUL) Magnetic Resonance Imaging Pulse Wave Transit Time Compared to Echocardiography in Chronic Obstructive Pulmonary Disease. Journal of Magnetic Resonance Imaging 56, 605–615. doi:10.1002/JMRI.28016, (Aug. 2022).
- 35. Wernz, M. M. et al. Feasibility, Repeatability, and Correlation to Lung Function of Phase-Resolved Functional Lung (PREFUL) MRIderived Pulmonary Artery Pulse Wave Velocity Measurements. Journal of Magnetic Resonance Imaging 60, 2216–2228. doi:10.1002/JMRI.29337, (Nov. 2024).
- Klimeš, F., Voskrebenzev, A., Wacker, F. & Vogel-Claussen, J. Three-Dimensional Phase Resolved Functional Lung Magnetic Resonance Imaging. *Journal of Visualized Experiments* 2024. doi:10.3791/66385, (June 2024).

- 37. Marshall, H. et al. 129Xe and Free-Breathing 1H Ventilation MRI in Patients With Cystic Fibrosis: A Dual-Center Study. Journal of Magnetic Resonance Imaging 57, 1908–1921. doi:10.1002/JMRI. 28470, (June 2023).
- 38. Klimeš, F. et al. Comparison of Free-Breathing 3D Phase-Resolved Functional Lung (PREFUL) MRI With Dynamic 19F Ventilation MRI in Patients With Obstructive Lung Disease and Healthy Volunteers. Journal of Magnetic Resonance Imaging 60, 1416–1431. doi:10.1002/JMRI.29221, (Oct. 2024).
- 39. Behrendt, L. et al. Validation of Automated Perfusion-Weighted Phase-Resolved Functional Lung (PREFUL)-MRI in Patients With Pulmonary Diseases. Journal of Magnetic Resonance Imaging 52, 103–114. doi:10.1002/JMRI.27027 (July 2020).
- 40. Behrendt, L. et al. A dual center and dual vendor comparison study of automated perfusion-weighted phase-resolved functional lung magnetic resonance imaging with dynamic contrast-enhanced magnetic resonance imaging in patients with cystic fibrosis. Pulmonary Circulation 12. doi:10.1002/PUL2.12054, (Apr. 2022).
- 41. Moher Alsady, T. et al. Multicenter Standardization of Phase-Resolved Functional Lung MRI in Patients With Suspected Chronic Thromboembolic Pulmonary Hypertension. Journal of Magnetic Resonance Imaging 59, 1953–1964. doi:10.1002/JMRI.28995, (June 2024).
- 42. Klimeš, F. et al. Dual-Center Repeatability, Consistency, and Comparison of 3D Phase-Resolved Functional (PREFUL) Ventilation MRI at 3 and 1.5 T in Healthy Volunteers. NMR in Biomedicine 38, e70101. doi:10.1002/NBM.70101; CTYPE: STRING: JOURNAL (Sept. 2025).

- 43. Friedlander, Y. et al. Phase-Resolved Functional Lung (PREFUL) MRI to Quantify Ventilation: Feasibility and Physiological Relevance in Severe Asthma. Academic Radiology 31, 3416–3426. doi:10.1016/ J.ACRA.2024.01.039 (Aug. 2024).
- 44. Foo, C. T., Langton, D., Donovan, G. M., Thompson, B. R., Noble, P. B. & Thien, F. Assessment of ventilation heterogeneity in severe asthma using phase-resolved functional lung magnetic resonance imaging. *Physiological Reports* 13. doi:10.14814/PHY2.70423, (June 2025).
- 45. Ouyang, T., Tang, Y., Klimes, F., Vogel-Claussen, J., Voskrebenzev, A. & Yang, Q. Phase-resolved Functional Lung (PREFUL) MRI May Reveal Distinct Pulmonary Perfusion Defects in Postacute COVID-19 Syndrome: Sex, Hospitalization, and Dyspnea Heterogeneity. *Journal of Magnetic Resonance Imaging* **61**, 851–862. doi:10.1002/JMRI.29458, (Feb. 2025).
- 46. Pöhler, G. H. et al. Chronic Thromboembolic Pulmonary Hypertension Perioperative Monitoring Using Phase-Resolved Functional Lung (PREFUL)-MRI. Journal of Magnetic Resonance Imaging 52, 610–619. doi:10.1002/JMRI.27097, (Aug. 2020).
- 47. Voskrebenzev, A. et al. PREFUL MRI Depicts Dual Bronchodilator Changes in COPD: A Retrospective Analysis of a Randomized Controlled Trial. Radiology: Cardiothoracic Imaging 4. doi:10.1148/RYCT.210147, (Apr. 2022).
- 48. Vogel-Claussen, J. et al. Phase-resolved Functional Lung (PREFUL) MRI-derived Ventilation and Perfusion Parameters Predict Future Lung Transplant Loss. Radiology 307. doi:10.1148/RADIOL.221958, (May 2023).

- 49. Zanette, B. et al. Clinical Feasibility of Structural and Functional MRI in Free-Breathing Neonates and Infants. Journal of Magnetic Resonance Imaging 55, 1696–1707. doi:10.1002/JMRI.28165, (June 2022).
- 50. Dyke, J. P., Voskrebenzev, A., Blatt, L. K., Vogel-Claussen, J., Grimm, R., Worgall, S., Perlman, J. M. & Kovanlikaya, A. Assessment of lung ventilation of premature infants with bronchopulmonary dysplasia at 1.5 Tesla using phase-resolved functional lung magnetic resonance imaging. *Pediatric Radiology* 53, 1076–1084. doi:10.1007/S00247-023-05598-6, (May 2023).
- 51. Klimeš, F. et al. 3D phase-resolved functional lung ventilation MR imaging in healthy volunteers and patients with chronic pulmonary disease. Magnetic resonance in medicine 85, 912–925. doi:10.1002/MRM.28482 (Feb. 2021).
- 52. Lévy, S. et al. Free-Breathing Low-Field MRI of the Lungs Detects Functional Alterations Associated With Persistent Symptoms After COVID-19 Infection. *Investigative Radiology* 57, 742–751. doi:10. 1097/RLI.000000000000000892, (Nov. 2022).
- 53. Glandorf, J. et al. Comparison of phase-resolved functional lung (PREFUL) MRI derived perfusion and ventilation parameters at 1.5T and 3T in healthy volunteers. PLOS ONE 15, e0244638. doi:10. 1371/JOURNAL.PONE.0244638 (Dec. 2020).
- 54. Hahn, J. J., Voskrebenzev, A., Behrendt, L., Klimeš, F., Pöhler, G. H., Wacker, F. & Vogel-Claussen, J. Sequence comparison of spoiled gradient echo and balanced steady-state free precession for pulmonary free-breathing proton MRI in patients and healthy volunteers: Correspondence, repeatability, and validation with dynamic

- contrast-enhanced MRI. NMR in Biomedicine 37. doi:10.1002/NBM.5209, (Nov. 2024).
- 55. Graves, M. J. & Mitchell, D. G. Body MRI artifacts in clinical practice: A physicist's and radiologist's perspective. *Journal of Magnetic Resonance Imaging* 38, 269–287. doi:10.1002/JMRI.24288, (Aug. 2013).
- 56. Peggs, Z. J., Brooke, J. P., Bolton, C. E., Hall, I. P., Francis, S. T. & Gowland, P. A. Free-Breathing Functional Pulmonary Proton MRI: A Novel Approach Using Voxel-Wise Lung Ventilation (VOLVE) Assessment in Healthy Volunteers and Patients With Chronic Obstructive Pulmonary Disease. *Journal of magnetic resonance imaging:*JMRI 61. doi:10.1002/JMRI.29444 (Feb. 2025).
- 57. Lewis, C. E., Prato, F. S., Drost, D. J. & Nicholson, R. L. Comparison of respiratory triggering and gating techiques for the removal of respiratory artifacts in MR imaging. *Radiology* 160, 803–810. doi:10.1148/RADIOLOGY.160.3.3737921, (1986).
- 58. Ehman, R. L., McNamara, M. T., Pallack, M., Hricak, H. & Higgins, C. B. Magnetic resonance imaging with respiratory gating: techniques and advantages. *AJR. American journal of roentgenology* **143**, 1175–1182. doi:10.2214/AJR.143.6.1175 (1984).
- 59. Santelli, C., Nezafat, R., Goddu, B., Manning, W. J., Smink, J., Kozerke, S. & Peters, D. C. Respiratory bellows revisited for motion compensation: Preliminary experience for cardiovascular MR. Magnetic Resonance in Medicine 65, 1097–1102. doi:10.1002/mrm.22687 (2011).
- 60. François, C. J. et al. Renal arteries: isotropic, high-spatial-resolution, unenhanced MR angiography with three-dimensional radial phase

- contrast. Radiology **258**, 254–260. doi:10.1148/RADIOL.10100443 (Jan. 2011).
- Johnson, K. M., Fain, S. B., Schiebler, M. L. & Nagle, S. Optimized 3D ultrashort echo time pulmonary MRI. Magnetic resonance in medicine 70, 1241–1250. doi:10.1002/MRM.24570 (2013).
- 62. Yoon, J.-W., Noh, Y.-S., Kwon, Y.-S., Kim, W.-K. & Yoon, H.-R. Improvement of Dynamic Respiration Monitoring Through Sensor Fusion of Accelerometer and Gyro-sensor. *J Electr Eng Technol* 8, 742. doi:10.5370/JEET.2013.8.?.742 (2013).
- 63. Bortolotti, L., Gowland, P., Clennell, I., Mougin, O., Glover, P. & Bowtell, R. Using Body Motion Tracking Data at Various Patient Position to Implement Respiration Gating in a 0.5 T Paramed ASG Upright Scanner in ISMRM Workshop on Motion Detection & Correction (2022), 4.
- 64. Krug, J. W., Odenbach, R., Boese, A. & Friebe, M. Contactless respiratory monitoring system for magnetic resonance imaging applications using a laser range sensor. *Current Directions in Biomedical Engineering* 2, 719–722. doi:10.1515/cdbme-2016-0156 (2016).
- 65. West, J. B. (B. West's respiratory physiology: the essentials eng (ed Luks, A.) Philadelphia, 2021.
- 66. Lanzer, P., Charles Barta, I., Elias Botvinick, M. H., D Wiesendanger, H. U., Gunnar Modin, D. I. & Higgins, C. B. ECG-Synchronized Cardiac MR imaging: Method and Evaiuatio&. *Radiology* (1985).
- 67. Chia, J. M., Fischer, S. E., Wickline, S. A. & Lorenz, C. H. Performance of QRS Detection for Cardiac Magnetic Resonance Imaging With a Novel Vectorcardiographic Triggering Method. *JOURNAL OF MAGNETIC RESONANCE IMAGING* 12, 678–688 (2000).

- 68. Chakeres, D. W., Kangarlu, A., Boudoulas, H. & Young, D. C. Effect of Static Magnetic Field Exposure of Up to 8 Tesla on Sequential Human Vital Sign Measurements. *JOURNAL OF MAGNETIC RESONANCE IMAGING* 18, 346–352. doi:10.1002/jmri.10367 (2003).
- Jekic, M., Ding, Y., Dzwonczyk, R., Burns, P., Raman, S. V. & Simonetti, O. P. Magnetic Field Threshold for Accurate Electrocardiography in the MRI Environment. *Magn Reson Med* 64, 1586–1591. doi:10.1002/mrm.22419 (2010).
- Lenz, G. W., Haacke, E. M. & White, R. D. Retrospective cardiac gating: A review of technical aspects and future directions. *Magnetic Resonance Imaging* 7, 445–455. doi:10.1016/0730-725X(89) 90399-8 (Sept. 1989).
- 71. Bert, R. J. et al. ECG Gating Is More Precise Than Peripheral Pulse Gating When Quantifying Spinal CSF Pulsations Using Phase Contrast Cine MRI. Academic Radiology 27, 552–562. doi:10.1016/J. ACRA.2019.06.015 (Apr. 2020).
- 72. Ogawa, R., Kido, T., Shiraishi, Y., Yagi, Y., Yoon, S. S., Wetzl, J., Schmidt, M. & Kido, T. Neural network–based fully automated cardiac resting phase detection algorithm compared with manual detection in patients. *Acta Radiologica Open* 11, 20584601221137772. doi:10.1177/20584601221137772 (Oct. 2022).
- 73. Wood, G. et al. Automated detection of cardiac rest period for trigger delay calculation for image-based navigator coronary magnetic resonance angiography. Journal of Cardiovascular Magnetic Resonance 25. doi:10.1186/s12968-023-00962-9 (2023).
- 74. Voskrebenzev, A., Klimeš, F., Wacker, F. & Vogel-Claussen, J. Phase-Resolved Functional Lung MRI for Pulmonary Ventilation and Per-

- fusion (V/Q) Assessment. Journal of visualized experiments: JoVE. doi:10.3791/66380, (Aug. 2024).
- 75. Seber, G. A. F. & Wild, C. J. Nonlinear Regression. Wiley Series in Probability and Statistics doi:10.1002/0471725315 (Feb. 1989).
- 76. Dumouchel, W. & O'brien, F. Integrating a Robust Option into a Multiple Regression Computing Environment, 41–48. doi:10.1007/978-1-4613-9154-8{_}3 (1991).
- 77. Holland, P. W. & Welsch, R. E. Robust regression using iteratively reweighted least-squares. *Communications in Statistics Theory and Methods* **6**, 813–827. doi:10.1080/03610927708827533 (Jan. 1977).
- 78. Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A. & Gee, J. C. A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage* **54**, 2033–2044. doi:10.1016/j.neuroimage.2010.09.025 (Feb. 2011).
- 79. Klein, S., Staring, M., Murphy, K., Viergever, M. A. & Pluim, J. P. Elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging* 29, 196–205. doi:10.1109/ TMI.2009.2035616, (Jan. 2010).
- 80. Schneider, C. A., Rasband, W. S. & Eliceiri, K. W. NIH Image to ImageJ: 25 years of image analysis. *Nature Methods* **9**, 671–675. doi:10.1038/NMETH.2089, (July 2012).
- 81. Thirion, J. P. Image matching as a diffusion process: An analogy with Maxwell's demons. *Medical Image Analysis* **2**, 243–260. doi:10. 1016/S1361-8415(98)80022-4 (1998).
- 82. Vercauteren, T., Pennec, X., Perchant, A. & Ayache, N. Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* 45. doi:10.1016/J.NEUROIMAGE.2008.10.040, (2009).

- 83. Forsberg, D. Fordanic / Image-registration 2015.
- 84. Voskrebenzev, A., Gutberlet, M., Kaireit, T. F., Wacker, F. & Vogel-Claussen, J. Low-pass imaging of dynamic acquisitions (LIDA) with a group-oriented registration (GOREG) for proton MR imaging of lung ventilation. *Magnetic Resonance in Medicine* 78, 1496–1505. doi:10.1002/MRM.26526, (Oct. 2017).
- 85. Klimeš, F., Voskrebenzev, A., Gutberlet, M., Grimm, R., Wacker, F. & Vogel-Claussen, J. Evaluation of image registration algorithms for 3D phase-resolved functional lung ventilation magnetic resonance imaging in healthy volunteers and chronic obstructive pulmonary disease patients. *NMR in Biomedicine* **36.** doi:10.1002/NBM.4860, (Mar. 2023).
- 86. Oppenheim, A. V. Discrete-time signal processing eng (ed Schafer,R. W.) Upper Saddle River, N.J, 2010.
- 87. Kaiser, J. F. & Schafer, R. W. On the use of the I0-Sinh Window for Spectrum Analysis. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **28**, 105–107. doi:10.1109/TASSP.1980.1163349 (1980).
- 88. Gustafsson, F. Determining the initial states in forward-backward filtering. *IEEE Transactions on Signal Processing* **44**, 988–992. doi:10. 1109/78.492552 (1996).
- 89. Nadaraya, E. A. On Estimating Regression. https://doi.org/10.1137/1109020
 9, 141–142. doi:10.1137/1109020 (July 2006).
- 90. Watson, G. Smooth Regression Analysis. Sankhy: The Indian Journal of Statistics 26, 359–372 (1964).

- 91. Chan, T. F. & Vese, L. A. Active contours without edges. *IEEE Transactions on Image Processing* **10**, 266–277. doi:10.1109/83.902291 (Feb. 2001).
- 92. Whitaker, R. T. A Level-Set Approach to 3D Reconstruction from Range Data. *International Journal of Computer Vision* **29**, 203–231. doi:10.1023/A:1008036829907/METRICS (1998).
- 93. Bai, C., Dallasega, P., Orzes, G. & Sarkis, J. Industry 4.0 technologies assessment: A sustainability perspective. *International Journal of Production Economics* **229**, 107776. doi:10.1016/J.IJPE.2020. 107776 (Nov. 2020).
- 94. Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X. & Gao, J. Large Language Models: A Survey (Feb. 2024).
- 95. Parekh, D., Poddar, N., Rajpurkar, A., Chahal, M., Kumar, N., Joshi, G. P. & Cho, W. A Review on Autonomous Vehicles: Progress, Methods and Challenges. *Electronics* 2022, Vol. 11, Page 2162 11, 2162. doi:10.3390/ELECTRONICS11142162 (July 2022).
- Alowais, S. A. et al. Revolutionizing healthcare: the role of artificial intelligence in clinical practice. BMC Medical Education 2023 23:1
 1–15. doi:10.1186/S12909-023-04698-Z (Sept. 2023).
- 97. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics* 5, 115–133. doi:10.1007/BF02478259/METRICS (Dec. 1943).
- 98. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536. doi:10. 1038/323533A0; KWRD=SCIENCE (1986).

- 99. Cun, L., Henderson, J., Le Cun, Y., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. & Jackel, L. D. Handwritten Digit Recognition with a Back-Propagation Network. Advances in Neural Information Processing Systems 2 (1989).
- 100. Fukushima. Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements. *IEEE Transactions on Systems Science and Cybernetics* **5**, 322–333. doi:10.1109/TSSC.1969.300225 (1969).
- 101. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 9351, 234–241. doi:10.1007/978-3-319-24574-4{_}28 (2015).
- 102. Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J. & Maier-Hein, K. H. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* 18, 203–211. doi:10.1038/S41592-020-01008-Z;SUBJMETA=114,1564,308,575,631,692; KWRD=IMAGE+PROCESSING, TRANSLATIONAL+RESEARCH (Feb. 2021).
- 103. Chen, X. et al. Recent advances and clinical applications of deep learning in medical image analysis. Medical Image Analysis 79, 102444. doi:10.1016/J.MEDIA.2022.102444 (July 2022).
- 104. Zhou, S. K. et al. A Review of Deep Learning in Medical Imaging: Imaging Traits, Technology Trends, Case Studies with Progress Highlights, and Future Promises. Proceedings of the IEEE 109, 820–838. doi:10.1109/JPROC.2021.3054390 (May 2021).
- 105. Badrinarayanan, V., Kendall, A. & Cipolla, R. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation.

- IEEE Transactions on Pattern Analysis and Machine Intelligence 39, 2481–2495. doi:10.1109/TPAMI.2016.2644615 (Nov. 2015).
- 106. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (Sept. 2014).
- 107. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I. & Savarese, S. Generalized Intersection over Union: A Metric and A Loss for Bounding Box Regression. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 2019-June, 658–666. doi:10.1109/CVPR.2019.00075 (Feb. 2019).
- 108. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H. & He, Q. A Comprehensive Survey on Transfer Learning. *Proceedings of the IEEE* 109, 43–76. doi:10.1109/JPROC.2020.3004555 (Nov. 2019).
- 109. Bhatt, S. P., Bodduluri, S., Hoffman, E. A., Newell, J. D., Sieren, J. C., Dransfield, M. T. & Reinhardt, J. M. Computed tomography measure of lung at risk and lung function decline in chronic obstructive pulmonary disease. American Journal of Respiratory and Critical Care Medicine 196, 569–576. doi:10.1164/RCCM.201701-00500C/SUPPL{_}FILE/DISCLOSURES.PDF (Sept. 2017).
- 110. Parr, D. G. Quantifying the lung at risk in chronic obstructive pulmonary disease does emphysema beget emphysema? *American Journal of Respiratory and Critical Care Medicine* **196**, 535–536. doi:10. 1164/RCCM.201705-0962ED/SUPPL{_}FILE/DISCLOSURES.PDF{\\$}AUTHORDISCLOSURES (Sept. 2017).
- 111. Li, F. *et al.* Latent traits of lung tissue patterns in former smokers derived by dual channel deep learning in computed tomography images.

- Scientific Reports 11, 1-15. doi:10.1038/S41598-021-84547-5; SUBJMETA=166,1785,4037,639,692,699,985; KWRD=BIOMEDICAL+ENGINEERING, CHRONIC+OBSTRUCTIVE+PULMONARY+DISEASE (Dec. 2021).
- 112. Tan, F., Zhu, X., Chan, M., Deveshwar, N., Willmering, M. M., Lustig, M. & Larson, P. E. Z. Pulmonary Ventilation Analysis Using 1H Ultra-Short Echo Time (UTE) Lung MRI: A Reproducibility Study. bioRxiv, 2023.10.22.563196. doi:10.1101/2023.10.22.563196 (Oct. 2023).
- 113. Liu, S., Leiva, V., Zhuang, D., Ma, T. & Figueroa-Zúñiga, J. I. Matrix differential calculus with applications in the multivariate linear model and its diagnostics. *Journal of Multivariate Analysis* 188. doi:10.1016/j.jmva.2021.104849 (Mar. 2022).
- Levin, D. L., Schiebler, M. L. & Hopkins, S. R. Physiology for the pulmonary functional imager. *European journal of radiology* 86, 308–312. doi:10.1016/j.ejrad.2016.09.027 (2017).
- 115. Arthur, D. & Vassilvitskii, S. k-means++: The Advantages of Careful Seeding. doi:10.5555/1283383.1283494.
- 116. Bayat, M., Fatemi, M. & Alizad, A. Background Removal and Vessel Filtering of Noncontrast Ultrasound Images of Microvasculature. *IEEE Transactions on Biomedical Engineering* 66, 831–842. doi:10. 1109/TBME.2018.2858205, (Mar. 2019).
- 117. Chavhan, G. B., Babyn, P. S., Thomas, B., Shroff, M. M. & Mark Haacke, E. Principles, Techniques, and Applications of T2*-based MR Imaging and Its Special Applications 1. Radiographics 29, 1433. doi:10.1148/RG.295095034 (Sept. 2009).

- 118. Kjørstad, Å., Corteville, D. M., Fischer, A., Henzler, T., Schmid-Bindert, G., Zöllner, F. G. & Schad, L. R. Quantitative lung perfusion evaluation using Fourier decomposition perfusion MRI. *Magnetic resonance in medicine* **72**, 558–562. doi:10.1002/MRM.24930 (2014).
- 119. Klimeš, F. et al. Free-breathing quantification of regional ventilation derived by phase-resolved functional lung (PREFUL) MRI. NMR in Biomedicine 32. doi:10.1002/NBM.4088 (June 2019).
- 120. Cotes, J. E., Maynard, R. L., Pearce, S. J., Nemery, B. B., Wagner, P. D. & Cooper, B. G. Lung Function 7th edition. eng (Wiley, Newark, 2020).
- 121. Stock, K. W., Chen, Q., Hatabu, H. & Edelman, R. R. Magnetic resonance T2/* measurements of the normal human lung in vivo with ultra-short echo times. *Magnetic Resonance Imaging* 17, 997–1000. doi:10.1016/S0730-725X(99)00047-8, (1999).
- 122. Hatabu, H., Alsop, D. C., Listerud, J., Bonnet, M. & Gefter, W. B. T2* and proton density measurement of normal human lung parenchyma using submillisecond echo time gradient echo magnetic resonance imaging. European Journal of Radiology 29, 245–252. doi:10.1016/S0720-048X(98)00169-7 (1999).
- 123. Theilmann, R. J., Arai, T. J., Samiee, A., Dubowitz, D. J., Hopkins, S. R., Buxton, R. B. & Prisk, G. K. Quantitative MRI measurement of lung density must account for the change in T*2 with lung inflation. *Journal of Magnetic Resonance Imaging* 30, 527–534. doi:10.1002/JMRI.21866, (Sept. 2009).
- 124. Higano, N. S. *et al.* Quantification of neonatal lung parenchymal density via ultrashort echo time MRI with comparison to CT. *Journal*

- of Magnetic Resonance Imaging 46, 992–1000. doi:10.1002/JMRI. 25643, (Oct. 2017).
- 125. Ley, S. et al. MRI measurement of the hemodynamics of the pulmonary and systemic arterial circulation: Influence of breathing maneuvers. American Journal of Roentgenology 187, 439–444. doi:10. 2214/AJR.04.1738, (Aug. 2006).
- 126. Merrikh, A. A. & Lage, J. L. Effect of blood flow on gas transport in a pulmonary capillary. *Journal of Biomechanical Engineering* 127, 432–439. doi:10.1115/1.1894322, (June 2005).

Chapter 5

Gravitational Dependency of the Distribution of Normal Lung Function

Abstract

Understanding the factors influencing regional pulmonary functional is of great importance for improving clinical care. Established theory attributes systematic variation of ventilation and perfusion to gravitational[1, 2] and physiological mechanisms[3]. However, due to the spatial constraints common to modern scanners, few modern imaging studies have probed the maximal difference in functional distribution in the human lung between the gravitationally opposite upright and inverted positions[4].

Twenty-three healthy volunteers (aged 21±2 years) underwent functional pulmonary imaging in the sitting, head-down, supine, and prone postures. This variation in participant positioning was facilitated using a 0.5 T Open MRI scanner with functional information generated using a home-built PREFUL algorithm[5]. Quantitative analysis was performed to examine positional differences in functional parameters across all four postures.

Asymmetric distributions of tissue density between gravitationally opposite postures suggest additional physiological contributions from the structure of the primary vessels and early bronchi, and the shape of the thoracic cavity. Tissue compression appeared to be a primary determinant of the distribution of ventilation in the lung[6]. However, a consistent bias towards the basal regions, irrespective of gravitational orientation, suggested additional mechanical contributions based on the changing shape of the thoracic cavity.

Regional perfusion largely aligned with predictions based on established gravitational mechanisms[1, 2, 7], however, these models cannot fully account for the small variations observed within isogravitational planes. Therefore, additional physiological factors may contribute to these observations.

Contents

5.1	Introduction	
	5.1.1	Clinical Implications
	5.1.2	Historical Perspective on Regional Lung Function 299
	5.1.3	Contributions to Normal Regional Lung Function 302
	5.1.4	The Current Study
	5.1.5	Aim of this Study
5.2	Metho	dology and Materials
	5.2.1	Equipment
	5.2.2	Independent Variables
	5.2.3	Participants
	5.2.4	Study Visits
	5.2.5	Data Analysis
5.3	Result	s
	5.3.1	General Results
	5.3.2	Variation in Exhaled Signal Intensity 365
	5.3.3	Variation in Ventilation Dynamics 374
	5.3.4	Variation in Perfusion Dynamics 398
5.4	Discussion	
	5.4.1	Tissue Compression
	5.4.2	Ventilation Dynamics
	5.4.3	Perfusion Dynamics
	5.4.4	Ventilation and Perfusion Matching 468
	5.4.5	Sources of Error
	5.4.6	Future Work

5.5	Conclusions
5.6	References

Declaration of Contributions

- Arthur Harrison Created and maintained the ethics. Designed the protocol. Recruited and screened participants for Protocol 1. Present for all scanning days; coordinating and positioning participants during Protocol 1 and running the scanner for Protocol 2. Processed and analysed all data presented in this chapter.
- **Dr Olivier Mougin** Created and optimised all scans and operated the scanner for Protocol one.
- **Dr Rashed Sobhan** Helped run the scanner, helped design the T1, T2* and B0 field map protocol. Processed T1 and T2* data, the results of which are not discussed in this chapter.
- **BMedSci Students** Recruiting and screened participants for and helping run Protocol 2.

5.1 Introduction

The human respiratory system is a complex and dynamic structure, with the lungs playing a central role as a gas exchanger between the atmosphere and our bloodstream[8]. Sufficient supply of air (ventilation) and blood (perfusion) throughout the organ is critical; however, these do not distribute evenly[3]. Therefore, the mechanisms behind regional ventilation and perfusion distribution and how they correspond to each other, is of central importance to our understanding of pulmonary gas exchange. Insight into this relationship has evolved in complexity over time, progressing from the assumption of functional uniformity, to recognising gravitational effects, and now appreciating structural and mechanical contributions. Our understanding of this topic continues to advance with improvements in imaging technologies.

5.1.1 Clinical Implications

Expanding our knowledge on the normal distribution of ventilation and perfusion in the lungs, as well as the mechanisms governing them, has far-reaching implications for clinical care[9]. This knowledge serves as the foundation for the interpretation of diagnostic imaging and the assessment of pulmonary diseases. Additionally, it enables clinicians to differentiate between regional functional changes typical of a healthy lung and alterations due to pathological changes[10, 11]. Indeed, many pulmonary diseases are characterised by changes in lung ventilation and perfusion, for example:

• **Asthma:** Regional bronchoconstriction narrows airways, increasing their resistance and altering the ventilation in areas of the lung which

they serve [12].

- Chronic Obstructive Pulmonary Disease (COPD): A collective term encompassing regional airway constriction and alveolar emphysema, resulting in uneven ventilation distribution with both areas of increased and reduced air flow[13].
- Pulmonary embolism: Blockage in the pulmonary arteries caused by a blood clot, altering regional perfusion via the restriction of flow to blood vessels[14].

Furthermore, knowledge of functional heterogeneity typical of healthy lungs can help with the treatment of lung disease. This knowledge allows for tailoring personalised medicine and drug delivery approaches, targeting treatment to where it is required and mitigating damage or side effects in healthy, high-functioning parts of the lungs[15, 16]. It can also guide methods in clinical care settings, such as strategies for mechanical ventilation[17] and surgical and post-operative care[18]. Finally, in addition drug delivery methods, understanding aerosol deposition is important for the understanding of particle deposition of contaminants and pollutants for occupational and public health[19]. Deepening our knowledge of functional lung physiology has wide-reaching implications for intervention, diagnosis and treatment for the improvement of patient outcomes. Further development in this field and application in clinical settings will allow for more precise and effective respiratory care[20, 21].

5.1.2 Historical Perspective on Regional Lung Function

Today, a multitude of studies aim to understand pathophysiological features and their response to treatments in diseased lungs. Core to that research our understanding of the physiological features and mechanisms determining the function of lungs in their healthy state. Until the 1950s, pulmonary functional distribution was assumed to be uniform throughout the organ [3]. However, studies on regional gas exchange in animal models and humans during this decade revealed heterogeneous characteristics [22, 23. These investigations were the first to demonstrate greater ventilation and perfusion at the base of the lungs compared to the apex. Modalities based on the detection of radioactive gases, were employed to investigate the vertical distribution of lung function. External scintillator counters were used to detect inert gases such as the metastable isotope ${}^{133}Xe^m$ that were either inhaled or injected intravenously. These demonstrated a trend of greater blood flow [24–27] and ventilation [28–30] at the base of the lungs compared to the apex. It was therefore theorised that the observed vertical distribution of lung ventilation and perfusion could be attributed to gravitational forces. These early experimental results contributed to the development of two models to explain these phenomena: West's zonal model for perfusion[2] and the later-termed 'Slinky effect'[1] for ventilation[31]. These findings, paralleled the first theories regarding the relationship between ventilation and perfusion in the lungs. These were determined highly correlated, as despite each's individual variability, gas exchanged remained highly efficient. With both ventilation and perfusion increasing down the lung, this matching was attributed to the shared influence of gravity [32]. Indeed, this perspective has been presented in textbooks and taught to

generations of medical students[3, 6].

Limitations of imaging technology prior to the 1970s necessitated the summation of lung function within isogravitational planes. However, as methodologies with greater spatial resolution became available, physiologists observed regional changes in lung perfusion that could not be attributed to gravitational effects alone. Studies on humans [33, 34] and animal models[35] observed variations in blood flow within isogravitational planes that could not be accounted for by hydrostatic or West's zonal conditions. These studies were the first to suggest that the shape and structure of the pulmonary vascular tree may contribute to blood flow distribution within the lungs[3]. Canine models of the 1980s were the first to cast doubt on the solely gravitational dependence of ventilation distribution. Non-uniformity of ventilation was observed within isogravitational planes[36], not consistent with the classic model for regional ventilation. Therefore, it was theorised that a complex interplay between the shape of the thoracic cavity [37] and the relative motion of the lung lobes [38] contribute to variations in regional ventilation. These observations marked a shift in the understanding of pulmonary functional distribution to the modern perspective, expanding from simple gravitational models to more nuanced explanations accounting for the complex structural and mechanical properties of the lungs and surrounding tissues.

Many investigations have utilised animal models for studies of lung function distribution. Animal models were practically useful where the technology of the time did not permit equivalent investigation in humans. One such technology is microsphere methods[39]: an early method for reconstructing three-dimensional maps of ventilation and blood flow. However, this method is based on the ex vivo analysis of the lung and thus typically requires the subject to be euthanised. These method confirmed the

observed heterogeneity of lung function within isogravitational planes, observing this phenomenon both in the ventilation [40] and perfusion [41] distribution of goats. It was also used to demonstrate that the distributions of ventilation [42] and perfusion [43] are spatially correlated to themselves - with high flow regions neighbouring high flow regions, and low flow regions neighbouring low flow regions. These observations gave weight to the theory that the geometry of the airways and vascular trees may contribute to functional distribution. Microspheres were used to demonstrate the tight matching of ventilation and perfusion under normal conditions in large laboratory animals [40, 44]. As this method can take a snapshot of lung function at roughly the time of administration, it was also employed for studies where the gravitational environment of lungs was altered. Perfusion distribution was investigated in primates suspended upside down[45] and a study on pigs aimed to remove the effect of gravity entirely, utilising brief periods of weightlessness on parabolic flights [46]. Each found gravity to be an important, but not predominant determinant of pulmonary blood flow distribution within these animals.

Approaching the modern day, a plethora of scanning methodologies is available for lung function investigations. Positron Emission Tomography (PET) methods have been used to observe vertical heterogeneity in the ventilation and perfusion of humans[47, 48] and dogs[49]. Computed Tomography (CT) with inhaled xenon has demonstrated ventilation distributions in animals[50, 51] while Electron Beam CT was employed to visualise perfusion heterogeneity in humans[52]. Single Photon Emission Computed Tomography (SPECT) has been used to investigate functional distribution, contributing some of the first observations of hilar to peripheral gradients of blood flow[53]. Although, the magnitude of this effect has since been disputed as an artifact of the methodology[54–56]. The redis-

tribution of lung function was investigated in a variety of postures using SPECT[57]. Subject received a radioactive nuclides dose both orally and intravenously, whilst assuming different postures. These were then detected whilst the subject laid down to negate the effects of parenchyma redistribution [7, 58]. This included perfusion distribution in the head-down posture and noted a non-opposite distribution to the upright lung[4]. Indirect measures of functional distribution, based on the profile of exhaled gas, have been used to investigate the effects of weightlessness on the distribution of lung function during parabolic flights and in space. Observing a significant, but not total reduction in vertical heterogeneity[59, 60] whilst sustaining similar although slightly diminished quality of ventilation and perfusion matching[61]. Magnetic resonance imaging (MRI) has been utilised to investigate the distribution of pulmonary ventilation and perfusion, as well as the relationships between them [62]. Studies have confirmed the functional redistribution between the supine and prone postures [63, 64]. Expanding in scope with contrast-enhanced methods, utilising agents such as oxygen to change the relaxation mechanics of hydrogen [65–67], or directly imaged contrast agents such as fluorine [68, 69] and the hyperpolarised noble gases helium-3[70–73] and xenon-129[74–76]. These studies have afforded great insight into the distribution of lung function and relationships of gas exchange in the lungs[71].

5.1.3 Contributions to Normal Regional Lung Function

The regionality of healthy lung function is thought to stem from a combination of both gravitational and physiological mechanisms; however, the relative contribution of each of these is still an active are of discussion [77].

For simplicity this explanation will focus on the mechanisms controlling functional distribution on the large to medium scales, omitting mechanisms below the length scales of the acinus. This explanation separates description of lung function distribution into ventilation and perfusion. Although, that is not to suggest they are unrelated. Furthermore, the following description generally concerns functional distribution in healthy lungs. Pathophysiology can have a significant effect on ventilation and perfusion distributions beyond the scope of this description [8].

5.1.3.1 Pulmonary Ventilation

5.1.3.1.1 Gravitational Contributions to Regional Ventilation

The lungs are a highly deformable and compressible organ and weigh approximately a kilogram for a normal adult[6]. Therefore, this organ will compress under its own weight due to the force of gravity. The resulting vertical gradient in the deformation of lung parenchyma contributes to a non-uniform distribution of lung ventilation; increasing from the least dependent (the top) to the most dependent (the bottom) of the organ[6]. To aid in the explanation of this phenomenon, it is useful to imagine the lung as a compliant, edge-wound spring, i.e. a SlinkyTM[1]. This effect is demonstrated in Figure 5.1.

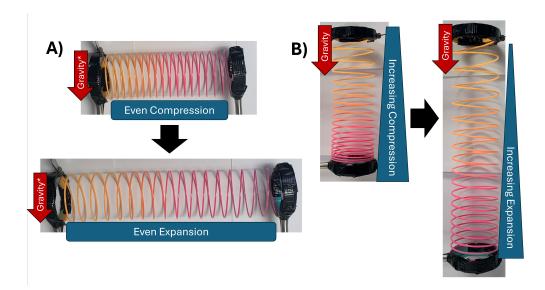


Figure 5.1: Demonstration of the Slinky Effect[1] using a circularly wound spring. When positioned perpendicular to the force of gravity (A), each loop is positioned at equidistant intervals. This is analogous to a uniform distribution of lung parenchyma along the isogravitational plane. When the spring is extended, all coils expand uniformly. However, when positioned parallel to the gravitational vector (B), the gravitationally dependent loops are spaced closer than those in the superior position. This is analogous to the non-uniform compression of parenchyma under its own weight. When the spring is expanded, the more compressed coils expand to a greater degree.

Applying this analogy to the lung, it is supported in the thoracic cavity by negative pleural surface pressure[6]. This pressure increases down the lung; in an upright posture a typical apical to basal distance of 30cm results in a pressure change of approximately 0.75kPa due to gravity[78]. Therefore, for a lung with otherwise uniform compliance, alveoli at different vertical heights will populate different points on the pressure-volume curve (Figure 5.2). Similar to gravitational changes in the spring (Figure 5.1), changing pleural surface pressure results in changes in the size of the alveoli: experiencing more compression the lower they are in the lung. The pressure-volume relationship is non-linear; these more compressed alveoli occupy a steeper part of the curve and are therefore more compliant[79]. A change in pressure here results in a greater change in volume compared to

the less compressed alveoli on the shallower part of the curve. Therefore, the same inspiration will result in a greater expansion of alveoli as one moves further down the lung. Thus, in an upright posture, lung ventilation increases apex to base. This effect remains consistent for a horizontal lung, although with a reduced magnitude to the smaller ventral to dorsal distance. Reproduced from [78].

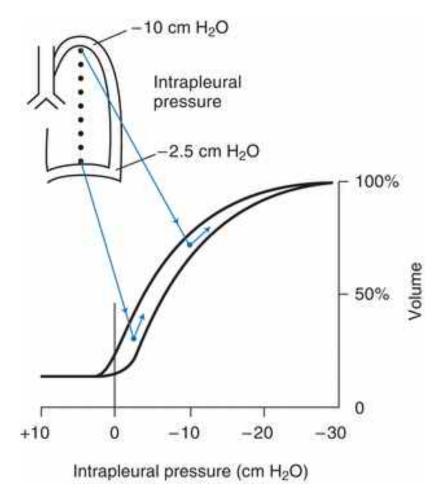


Figure 5.2: Hydrostatic forces generate a vertical gradient of pleural pressure within the thoracic cavity. This causes alveolar units at different vertical heights to occupy different parts of the pressure-volume curve. In the upright lung, the basal regions will experience a less negative pleural pressure resulting in a relative compression of parenchyma compared to regions near the apex. These compressed regions occupy a more compliant (steeper) part of the pressure-volume curve and therefore expand to a greater degree than those near the apex during inspiration. Reproduced from [6].

5.1.3.1.2 Physiological Contributions to Regional Ventilation

The above description of the dependence of ventilation on gravitational forces relies on the assumed uniformity of mechanical properties within the lung[6]. However, this is often not the case: the lung can possess regional differences in resistance, compliance or both[6]. The resistance dictates the difficulty for air to navigate the airways and reach the alveoli, whilst the compliance represents the readiness of the alveoli to expand. Therefore, the lung can be thought of as a series of parallel units, each comprised of a small tube representing the resistance, feeding a container with a specific compliance [80]. Multiple factors can affect these properties, with the most relevant being:

- Airway dilation can change the resistance of airways and thus the ventilation of the regions of which the serve[9].
- Non-uniform distributions of surfactant can alter regional alveolar surface tension and consequently their compliance, resulting in heterogeneous ventilation distributions[6, 81, 82]

Although, the above examples are not thought to be typical of a healthy lung[6].

The 'time constant of a lung unit' is the product of the regional resistance and compliance and describes the rate of change of ventilation given a sudden change in pleural pressure [80]. Inequalities in time constant across lung unit can create a phenomenon termed 'Pendelluft' [83], where upon inspiration, air that initially occupied fast-filling units can transfer to slower-filling units. This reduces the lung's effective tidal volume and can pose a challenge when measuring regional ventilation [9].

A significant contributor to the non-uniformity of regional ventilation in healthy lungs is the structure of the airway tree[3]. Inequalities at each bifurcation can lead to heterogeneous airway resistance for different regions of the lung. This is primarily a temporal phenomenon, causing inequality in the rate of expansion of different parts of the lung[9]. Alveoli served with shorter, wider airways will exhibit less resistance those served with longer and narrower airways. Reduced resistance means inspired air can reach the alveoli more easily and thus will expand more readily with inspiration. The extent to which this determines the spatial distribution of ventilation depends on the duration of inspiration. Given enough time, all alveoli with common pleural surface pressure will achieve the same expansion ratio. However, if the inspiration is brief, then regions with high resistance may not have sufficient time to inflate, leading to heterogeneity in ventilation. As neighbouring lung regions share a common history of branching in the airway tree, and consequently possessing similar resistances. Therefore, ventilation heterogeneity due to this mechanism is often spatially correlated with high flow regions neighbouring high flow regions and vice versa[42]. In addition to changes in regional resistance, the uneven contraction of respiratory muscles can create regional heterogeneity of ventilation[9]. For example, increased contraction of the intercostal and accessory muscles, associated with distortion of the rib cage, can increase apical ventilation [84]. This can be associated with forced inspiration from residual volume. Furthermore, heightened contraction of the diaphragm can increase air flow to the base of the lungs. This is observed in many postures, but particularly when lying on one's side[85].

5.1.3.2 Pulmonary Perfusion

5.1.3.2.1 Gravitational Contributions to Regional Perfusion

Pulmonary circulation operates at much lower pressures compared to systemic circulation, at approximately 25/10 mmHg[86]. This low perfusion pressure means that hydrostatic pressure differences (approximately 23 mmHg in an upright lung with a basal to apical distance of 30cm) within the lung play a significant role in the distribution of blow flow. This mechanism is described by the long-established Zonal model for perfusion, developed by West et al.[2] This model describes blood flow as a product of the relative magnitude of the arterial, venous and alveolar pressures. Both arterial and venous pressures increase vertically down the lungs by the same magnitude due to the gravitational hydrostatic forces. Although, this change is relatively greater for the arterial pressure due to the greater contribution from the heart. Whereas, the alveolar pressure remains equal throughout the lung, assuming patent airways. West's zonal model sees the lung divided into three vertically stacked zones. The lower-most zone, zone three, sees the greatest perfusion, and its mechanism is the most intuitive. Here, blood flow depends on the relative magnitude of the arterial and venous pressures, each greater than the alveolar pressure due to hydrostatic effects. Moving up the lung, both the arterial and venous pressures are reduced, and zone two begins when the venous pressure falls below that of the alveolar pressure. Here, the alveolar pressure can be though of 'squeezing' the capillaries, restricting pulmonary blood flow[87]. Therefore, perfusion in this zone is dependent on the relative magnitude of the arterial and alveolar pressure. In zone one, the uppermost portion of the lung, alveolar pressure surpasses both arterial and venous pressures, causing capillary collapse and preventing blood flow. While zone one doesn't exist in a normal lung, it can occur in situations of high alveolar pressure or low pulmonary pressure. Finally, an additional zone is present in the most dependent part of the lung. Sometimes called zone four, this zone sees a reduction in blood flow down the lung due to the high interstitial pressure compressing alveolar blood vessels[27]. Importantly, these zones do not represent fixed anatomical regions but rather physiological conditions that can shift based on a multitude of factors such as body position, lung volume, cardiovascular status, atmospheric pressure and external (G) forces.

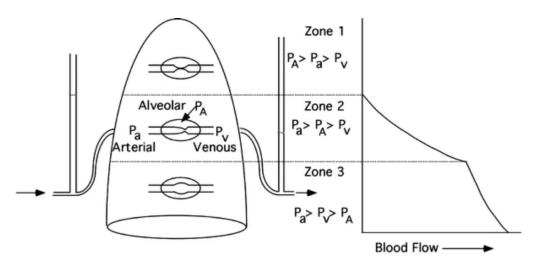


Figure 5.3: Depiction of West's zones of pulmonary perfusion[2], based on the interplay of arterial (P_a) , venous (P_v) , and alveolar (P_A) pressures. Estimated pulmonary blood flow is depicted for zones 1, 2, and 3 within the lung. However, zone 4 is omitted from this diagram. Reproduced from [2].

Further to gravitational effects due to hydrostatic forces, the lung will also see regional perfusion heterogeneity due to tissue compression[88]. The same mechanism as described for ventilation, known as the 'Slinky Effect'[1]. The higher tissue and thus alveolar density in the most dependent region of the lungs results in a higher density of alveolar capillaries. However, while measured blood flow per volume is increased due to this effect, the blood flow per alveolus does not increase. This distinction must be

made for any investigation of perfusion distribution[58].

5.1.3.2.2 Physiological Contributions to Regional Perfusion

In addition to systematic gravitational effects, physiological mechanisms contribute to the distribution of blood flow in the lung. However, their contribution to the perfusion distribution in the lungs remains an area of scientific debate. One such mechanism is the geometry of the vascular tree; taking the form of a dichotomously branching fractal pattern[6]. Like a tree, it is similar across all length scales such that the early branching will resemble the late branching. This is thought to be due to its evolutionary benefits [89]. During early development, a recursive fractal network requires fewer instructions to form a complete vascular tree, instead of having specific DNA for every vascular branching[90]. Additionally, a fractal network also creates the largest possible surface area within a given volume [6]. Finally, a fractal distribution minimises the amount of material needed to create the vascular network, and the amount of blood to fill it [55]. This has the effect of minimising the work required from the heart to perfuse the pulmonary regions [91]. However, this structure inherently creates heterogeneity of the blood flow distribution due to asymmetry of branching resistances. Thus, each of the generations of branching may not perfectly split blood flow 50:50. For example, a two percent heterogeneity (49:51) of blood flow compounded over 17 generations will result in a two-fold maximal difference in perfusion between lung regions. This mechanism is genetically determined and remains consistent during ageing, resulting in spatially correlated blood flow distributions, as neighbouring lung regions will share vascular history[92].

5.1.3.3 Ventilation and Perfusion Matching

The cornerstone of efficient pulmonary gas exchange is the relationship between ventilation and perfusion within each lung unit. This relationship is often denoted by the ventilation-perfusion ratio, V/Q. The lungs are most efficient with a V/Q close to 1. This is due to the similarity between the oxygen content in the atmosphere (21%) and the oxygen capacity of blood (20%) with a normal haemoglobin level. Thus, the lungs require a similar total quantity of air and blood of approximately 5-6 L per minute[6]. It is clinically useful to describe V/Q variation in the lungs with the highly simplified three-compartment model. These compartments are termed:

- Shunt An area with perfusion but no ventilation, hence a V/Q of zero. Perfusion to this area is wasted.
- Dead-space An area with ventilation but no perfusion, hence V/Q
 of infinity. Ventilation to this area is wasted.
- Ideal Collective area encompassing all remaining ventilation and perfusion. This is the only compartment that can contribute to gas exchange.

In a healthy lung, only the conducting airways represent true anatomical dead space, and there are rarely any true shunt compartments[9]. Instead, in this model, any wasted ventilation or perfusion due to V/Q mismatch is assigned to the dead-space or shunt compartments respectively. Although simple, this model allows clinicians to build up a quantitative description of the lungs. Using clinical measures of gas exchange, the distribution of ventilation and perfusion can be assigned to each of the virtual compartments. This provides a useful description of the lungs in health and disease.

It is often more useful to adapt this relationship to the alveolar ventilation to perfusion ratio, Va/Q, to only account spaces contributing to gas exchange. This ratio determines the concentration of oxygen and carbon dioxide in any lung unit, given the concentration of each molecule in the inhaled gas and blood entering from the pulmonary artery. As discussed previously, both ventilation and perfusion vary significantly within the lung. With such variability, one might assume there would be significant variation in Va/Q. However, this is not the case. In healthy lungs ventilation and perfusion are highly correlated to each other, resulting in Va/Q ratios close to 1 and thus highly efficient gas exchange. Measures of the regional matching of ventilation and perfusion have resulted in an average correlation coefficient of above 0.7. The exact mechanism behind this correlation remains a topic of scientific debate, although it is thought to be a combination of shared gravitational and physiological effects. Intuitively, any heterogeneity of ventilation and perfusion due to gravitational effects will be caused by a common force. Therefore, functional non-uniformity due to the 'Slinky effect' [1] and West's Zonal model [2] will apply along the same vector. A simple explanation for the physiological matching is that the lung is constructed in such a way to match regional ventilation and perfusion. This passive mechanism is a result of similarities between the airway and vascular trees, each following the same path. Therefore, any inequalities in blood flow and airway resistance may be shared at each bifurcation. Whilst this results in a heterogeneous functional distribution, correlation between regional ventilation and perfusion is maintained. However, whilst elegant in its simplicity, this explanation omits the contribution to ventilation distribution due to differences in regional compliance; often a greater determinant to ventilation distribution than airway resistance. This could be explained via a concept termed symmorphosis [93], where structural design matches functional demand. This may result in a correlation between

regional compliance and airway geometry. Indeed, correlation has been noted in the diameter of airway branches and the volume of peripheral lung served by said branch.

The above explanations for ventilation and perfusion matching are based on passive processes. However, the distribution of perfusion in the lungs can be actively altered via a mechanism called Hypoxic Pulmonary Vasoconstriction (HPV)[94]. This mechanism serves to redistribute blood away from areas with low ventilation (Shunt), to match ventilation and blood flow and thus optimise gas exchange. This is achieved by contraction of the pulmonary vascular smooth muscle due to the local tension of oxygen[95]. This mechanism is generally not thought to be present to a significant degree in healthy lungs; however, one's susceptibility to it is an inherent genetic characteristic of the individual[96]. This mechanism may be activated with lung disease, or with the reduction in oxygen partial pressure typical of ascent to high altitude[9].

5.1.4 The Current Study

This work aims to build upon previous studies and provide more human data regarding normal regional lung function. However, decoupling the relative contributions of physiological and gravitational mechanisms is not a straightforward endeavour. There is currently no method capable of imaging the lungs without the presence of gravity. Studies must instead vary the relative direction of gravity experienced by the lungs and use the resulting change in functional distribution to infer the contribution from gravity. This variation in the relative direction of gravity is achieved by positioning the participant in a range of postures to alter the angle of the lungs to the vertical. While the effects of gravity on human pulmonary function have

been studied extensively using modern imaging modalities, these studies are generally limited to the subject lying horizontally. This is due to the spatial restrictions inherent to the design of modern imaging systems. As a result, few investigations have probed the distribution of lung function in the vertical lung. In this position, one would expect a maximal contribution from gravitational effects due to the larger apical-to-basal distance compared to the anterior-to-posterior. Humans spend most of their awake time in the upright position, either whilst standing or sitting down. Therefore, an understanding of functional mechanisms in this posture is of great importance. Additionally, although not a common posture for most people, imaging the lung whilst in the opposite orientation, or upside down (or head-down), could provide valuable insight into lung function distribution by maximally changing the relative direction of gravity.

Few studies have sought to compare functional redistribution in human lungs between the upright and head-down posture. Glenny et al. used primates as a bipedal animal model for humans, employing fluorescent microspheres to map the distribution of pulmonary blood flow in four postures: upright, head-down, supine, and prone [45]. They found the greatest perfusion heterogeneity due to gravity in the upright posture, followed by head-down and supine, with almost no gradient whilst prone. Combined with previous studies that observed persistent functional inequity in microgravity [61], this work deduced gravity to be a secondary determinant of blood flow distribution. This challenged the traditional gravitational model of perfusion, suggesting that mechanisms due to vascular structure and anatomy are likely greater contributors [3]. Peterson et al. have conducted a series of studies investigating the redistribution of lung function distribution using SPECT [4, 7, 58]. This method is based on the detection of radiotracers that, when administered intravenously or orally via

aerosol, remain in place in the lung for a period. The first study aimed to investigate the contribution from redistribution of lung parenchyma to observed blood flow distribution in the supine and prone posture [58]. It found that this mechanism has a major influence on the observed distribution of perfusion. Following this work, the researchers set out to investigate the distribution of ventilation and perfusion in the upright lung[7]. They applied radiotracers while the subjects were seated upright to be investigated whilst supine. They found a much greater influence from gravity on the distribution of function in the upright posture compared to the horizontal lung. Both ventilation and, with a greater magnitude, perfusion increased towards the most dependent lung regions, resulting in a decreasing ventilation-perfusion ratio. Finally, the group set out to investigate the maximal influence of posture on blood flow distribution, mapping perfusion in the upright and head-down posture [4] The results of this study are summarised in Figure 5.4. Using a tilt table, radiotracers were administered intravenously while the subjects were positioned upright and upside down, to be detected whilst supine. This work observed a significant apical redistribution of blood flow when the subject was in the head-down posture compared to upright, although not to the magnitude required to observe an inverse distribution. Therefore, they concluded that whilst gravity plays a role in the distribution of regional blood flow, it is a secondary contribution behind that of lung structure.

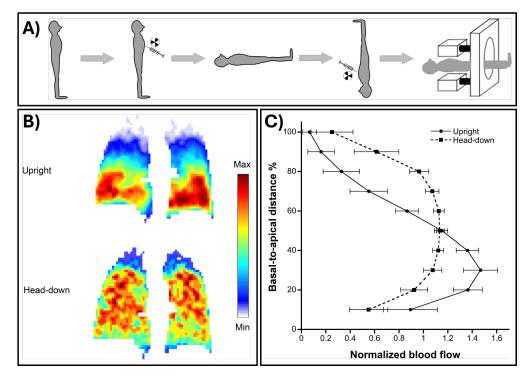


Figure 5.4: Summary of results from Ax et al. [4], investigating the apical-to-basal distribution of perfusion in the upright and head-down lung using SPECT. Radionuclide particles were administered intravenously to participants in the upright and head-down postures (A). These were investigated using SPECT (B) whilst the participants were supine. Averaging these maps over ten isogravitational planes provided a comparison of blood flow between the gravitationally opposite postures in the human lung (C).

Numerous previous studies have employed MRI to investigate the distribution of pulmonary ventilation [65–69, 71, 72, 76] and perfusion [1, 62–64, 97, 98], however, these have been generally been limited to patients in a horizontal postures. This is due to the spatial constraints associated with the horizontal cylindrical bore, common to most scanners. Therefore, studies of this nature have constrained participants to lying prone, supine or in a lateral recumbent position. These studies have reported significant apical-to-basal gradients in pulmonary function in the supine posture [62, 63, 71], with generally [98] more uniform distributions when prone [62, 63, 71, 97].

Hopkins and colleagues have published many studies using proton MRI to investigate functional distributions in the lung, employing arterial spin

labelling[98–100] to investigate pulmonary perfusion and oxygen enhanced MRI[66] for specific ventilation in the supine and prone posture. Furthermore, their observations of vertical gradients of parenchymal density in supine humans led to the terming of the Slinky effect[1]. Figure 5.5 summarises the results of this study, including observations of the anterior-to-posterior gradients in lung perfusion and tissue density in the supine lung. These were found to be highly related, leading to the conclusion that anterior-to-posterior perfusion variations in the supine posture are primarily dependent on tissue density/composition.

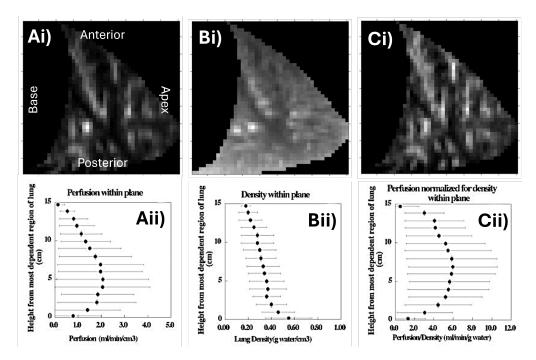


Figure 5.5: Summary of results from Hopkins et al.[1], investigating the contribution from tissue density to the distribution of perfusion in the supine lung. Sagittal maps of perfusion (Ai) and proton density (Bi) were created using arterial spin labelling and FLASH proton density, respectively. These were combined to create a tissue-density-normalised map (Ci). Averaging these maps along isogravitational planes created distributions of the perfusion (Aii), lung density (Bii), and tissue density normalised perfusion (Cii).

Spatial constraints typical of conventional horizontal cylindrical-bore MRI scanners can be alleviated with the use of Open MRI. These scanners are

designed to allow a far greater degree of subject positioning; however, this is generally at the expense of field strength and homogeneity. This technology has been employed for studies that would be difficult with conventional MRI, including musculoskeletal system[101, 102], gastric emptying[103] and respiratory kinematics investigations[104]. These investigations have positioned participants in a number of postures, including supine, prone, sitting upright, and standing[101, 102, 104]. Furthermore, open MRI has been used to investigate the stomach whilst participants adopted a head-down posture[103]. However, due to difficulties associated with lung imaging with proton MRI, no studies have utilised the opportunities associated with open MRI for pulmonary investigations. This limitation includes investigating the distribution of lung function in the upright lung. Therefore, this study aims to investigate the maximal change in pulmonary function, by placing participants in the upright, supine, prone and head-down postures.

Functional information will be derived for each using proton-based Phase Resolved Functional Lung (PREFUL) methodology[5] during free breathing. A full description of this method can be found in Chapter 4.

5.1.4.1 Health Risks of a Head-Down Posture

Assuming an inverted posture is not without its risks, although it is common in activities such as yoga and gymnastics, or with the use of inversion tables. These activities should only take place with due care and consideration. Common side effects of an inverted posture can include an elevated blood pressure in the head and eyes, dizziness, headaches and nausea[105]. Long term inversion of healthy individuals can lead to death, with the risk to the subject increasing significantly if they have underlying health conditions[106–108]. This can be via several mechanisms:

- Asphyxiation The weight of the organs resting on the lungs when upside down makes it harder to breathe and will eventually impede gas exchange sufficiently to be fatal. This is called positional asphyxia. This risk is greater if the subject is overweight or obese due to the greater weight on the lungs[109].
- Heart Failure Our circulatory system has evolved to work with gravity. Therefore, when upside down, the heart must work harder to pump the blood around your body. Blood must travel a greater vertical height to reach your feet, compared to your head when upright. Therefore, a greater systolic pressure is required to overcome the increased hydrostatic pressure difference. The risk of heart failure is increased if the subject has a cardiovascular condition[106–108].
- Brain haemorrhage The muscles in your legs assist with the return of blood to the heart when upright. However, your head does not have such a mechanism for blood return when upside down. With sustained inversion, this can result in blood pooling in the brain. Combined with the increased blood pressure of the inverted posture, this can lead to ruptured blood vessels and brain haemorrhages. The risk of a brain haemorrhage is elevated if the subject has high blood pressure [106–108].

Recoded incidents describe fatalities due to long periods of inversion of healthy individuals occurring after approximately 20 hours[106–108]. However, this time can reduce significantly with pathology to just tens of minutes[109]. Studies on this process are understandably hard to conduct. Animal models have been used to investigate this phenomenon[108], but studies on humans are largely retrospective following fatal accidents[106]. This is especially true for individuals who are obese, whom generally exhibit

reduced lung function anyway[110].

Therefore, for the safety of the subject, any inversion period should be short, and participation should exclude those who are elderly, pregnant, or have health conditions such as cardiovascular problems, hypertension, or are overweight.

5.1.5 Aim of this Study

The aim of this study was to investigate the effect of changing the direction of the force due to gravity experienced by the lungs. This was achieved by scanning the participants in different positions, thus orientating their lungs at different angles to Earth's gravitational force. The maximal change in the distribution of ventilation and perfusion, as well as their variation within isogravitational planes, was used to infer the relative contributions of gravitational and physiological factors on normal lung function.

In the sitting and head-down postures, the gravitational vector is parallel to the long axis of the lungs: towards the base and the apex, respectively. When the lungs are positioned horizontally, the gravitational vector is perpendicular to this axis; towards the posterior when supine and anterior when prone. Furthermore, functional variation was also investigated in a number of isogravitational planes: along the apical-to-basal axis in both supine (coronal and sagittal planes) and prone positions (sagittal plane), and along the anterior-to-posterior direction in the sitting position (sagittal plane).

5.2 Methodology and Materials

5.2.1 Equipment

5.2.1.1 Scanner

All MRI experiments for this study were carried out using a 0.5 T Open MRI Scanner (ASG Paramed, Italy) located in the University of Nottingham Medical school. This scanner is equipped with a four-channel receive body coil, and a gradient system with a maximum power of 20 mT/m and rise time 0.6 ms.

Illustrated in Figure 5.6, this scanner comprises two parallel cylindrical magnetic poles separated by a gap of width 56cm. The static magnetic field (B0) is oriented horizontally between these poles, with a spherical region of magnetic field homogeneity of approximately 20cm in diameter at the isocentre. The space between these magnetic poles remains accessible, facilitating a diverse range of experimental configurations. The system incorporates a movable, articulated support platform with a central hinge, enabling variable subject positioning. In the horizontal configuration resembling a bed, the system permits conventional supine imaging protocols. However, when the central hinge is actuated to 90 degrees, the support mechanism provides a seated configuration for upright studies. Furthermore, removing the support mechanism entirely enables studies of the subject whilst standing. These examples represent a subset of the potential experimental orientations enabled by this open-bore architecture. Indeed, this study will employ atypical human postures to explore maximal changes in pulmonary functional distribution.

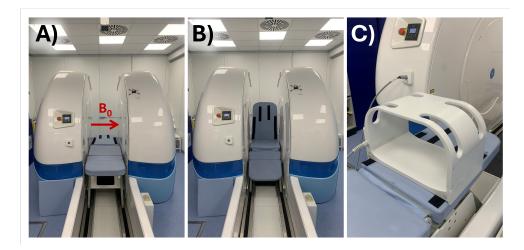


Figure 5.6: The 0.5T open MRI scanner (ASG Paramed, Italy) employed in this study is pictured with the central support mechanism articulated to resemble a bed (A) and a chair (B). The commercial body coil (ASG Paramed, Italy) used for all experiments in this study (C); resting horizontally, exterior to the scanner. The B₀ field direction is labelled in red.

5.2.1.2 Scanning Protocols

Scout Scans TR/TE=8/84.3ms, FOV=320x320x100 mm (15mm slices, 5mm gap, 5 slices), matrix size=192x128x5. 5x3 images acquired in 3 different orientations (axial, coronal, sagittal).

2D Gradient Echo (GRE) TR/TE = 5.7/1.2ms, FOV = 320×320 mm, matrix size = 64×64, cartesian sampling, slice thickness = 20mm (single slice: either axial, sagittal or coronal), flip angle = 10 degrees, TR per slice 365ms.

Half-Fourier Acquisition Single-shot Turbo spin-Echo (HASTE)

TR/TE = 405/108ms, FOV = 320x320 mm, slice thickness = 20mm, acquired matrix size of 64x64, refocusing pulse of 120 degrees (excitation=90), echo spacing of 8ms for an echo train of 46.

Henceforth, all scanning protocols will be referenced by the above names.

5.2.1.3 Slice Positioning

All slices were positioned by eye, with the rough location of each plane depicted in Figure 5.7. The slice in each imaging plane had its own positioning protocol:

Coronal Slice Positioned posterior to the heart while still obtaining a partial volume of the aorta for the perfusion processing. A general rule for this positioning is to place the slice such that it captures approximately half of the aorta, above the heart. The slice may also be tilted slightly from the vertical to match the angle of the spine to capture more of the lungs and to compensate if the participant is slouching. For each subject and each posture, the positioning of this slice was checked by collecting 10 2D GRE images to ensure the heart did not enter the frame.

Sagittal Slice Positioned over approximately the centre of the right lung. This lung was chosen to avoid cardiac interference associated with the left lung. Furthermore, due to the curvature of the diaphragm, the length of the lung is greatest in the dorsal region. This resulted in a greater likelihood that the base of the lung in this area could leave the FOV during a respiratory cycle, compared to investigations in the coronal plane where the slice is positioned more anterior to this region. Therefore, deep breathing experiments were omitted in this plane due to the greater associated expansion of the lung compared to normal breathing.

Axial Slice Positioned just above the heart while still capturing the top of the aorta for perfusion processing. This area was selected to observe the greatest lateral expansion in the lung area in later ventilation processing. This expansion was due to the activation of the intercostal muscles in this region. Furthermore, participants were only investigated whilst breathing deeply in this plane to maximise this effect. 10 frames were captured using the 2D GRE sequence to check that changes in the lung 2D area could be observed during the deep breathing respiratory cycle.

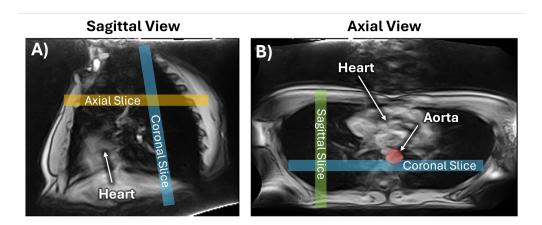


Figure 5.7: Diagram of the slice positioning used in this study. Rough slice positions are depicted as coloured overlays on two sample MRI images in the sagittal (A) and axial (B) planes, captured using the Scout sequence of a participant in the sitting posture. The heart is labelled in both images, and the aorta is labelled and highlighted in **red** in the axial image. The axial slice, shown in **yellow**, is located just above the heart. The sagittal slice (**green**) is located over approximately the centre of the left lung. The coronal slice (**blue**) is located behind the heart, angled as required to avoid the organ while still sampling the diaphragm. This slice should also partially sample the aorta. The slice overlays in this figure are not to scale. The highlighted axial and sagittal slice positions do not represent the position depicted in the axial and sagittal views.

5.2.2 Independent Variables

5.2.2.1 Participant Postures

Each subject was investigated while assuming each of at least three postures, sitting, head-down and supine. Some participants were also investigated in the prone position. These were chosen such as to maximally vary the relative direction of gravity experienced by the lungs. Each posture required a different setup and subject positioning, depicted in Figure 5.8:

- Supine The hinge at the centre of the movable bed was set to 180 degrees, resulting in a horizontal flat surface. The subject was positioned lying on their back on this bed and the coil placed around their torso. This posture resulted in the lungs assuming a horizontal orientation with the dorsal lung the most dependent. The bed was then moved into position such that the lungs occupied the isocentre of the magnet.
- Prone The bed setup was identical to that for the supine posture; however, instead of lying on their back, the participants were instructed to lie on their front, with their arms above their head. This resulted in the lungs being oriented horizontally with the ventral regions being the most dependent. The bed was then moved into position such that the lungs occupied the isocentre of the magnet.
- Sitting The central hinge was set to 80 degrees to create a seat. The participant was directed to sit on the resulting seat and again the coil was fitted around their abdomen. This posture resulted in the lungs assuming a upright vertical orientation. A cushion was placed in front of the subject, within the coil. This was to restrict forward movement; however, it was soft enough so as not to restrict breathing. The seat was then moved into position such that the lungs occupied the isocentre of the magnet.
- **Head-Down** Unlike the previous three postures for which the participant was positioned in the desired posture outside and later moved into the isocentre of the magnet, this posture required the participant

to be positioned in situ. Like the setup for the sitting posture, the central hinge was set to 80 degrees to create a platform that resembled a chair. The coil was placed at a height determined by eye, based on the height of the subject and such that their shoulders would be level with its bottom. If necessary, the subject could be later raised or lowered with the use of foam pads on the base of the chair and initially a pillow was placed on the base of the chair for the comfort of the participant. With the coil in place, the chair was then moved to the isocentre of the magnet. Then, in the space behind the bed, two large foam blocks were placed to create a platform level with the top of the chair. To enter the head down posture, the participants were instructed to climb upon this platform and lower themselves headfirst through the coil. This manoeuvre was assisted by a researcher acting as a 'spotter', positioned in front of the chair. When in the correct posture, the participants were upside-down facing the back of the chair and supporting themselves with their forearms above their head and their waist on the chair top. This posture resulted in the lungs assuming a near-vertical upside-down orientation. Importantly, the participants were positioned such that their torso was not pressed up against the back of the chair, which could have restricted breathing or created motion during respiration due to the expansion of the chest. For all scans whilst the participant was in this position, a researcher remained in the scanning room to monitor their well-being and be positioned to quickly release them from the posture if necessary. To exit this posture, the subjects were asked to climb back out reverse of the way they entered. However, if this were not possible for them to exit this way or in an emergency, they could be released by removing the front of the coil and moving forward from the top of the chair.

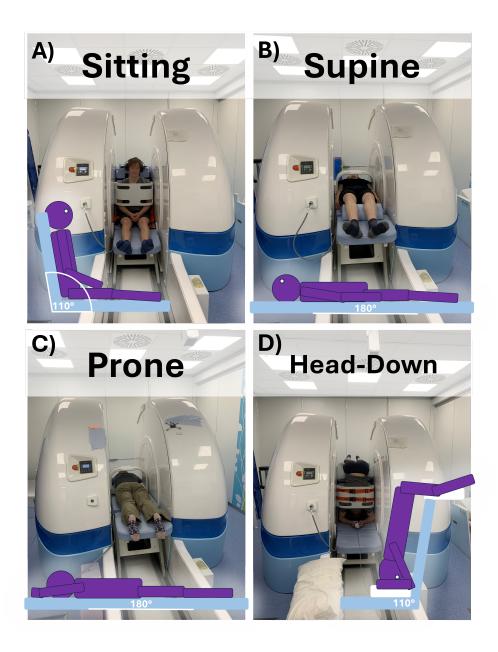


Figure 5.8: These images show the participant postures employed to vary the relative direction of the gravitational vector to the lungs: sitting (A), supine (B), prone (C) and head-down posture (D). Each image is accompanied by a respective diagram of the participant posture.

Each posture was chosen to maximally vary the orientation of the lungs to the vertical; however, in practice this positioning method would not result in the participants' lungs perfectly assuming the desired orientation. While a good approximation can be made for the lungs in horizontal orientation $(90 \pm 3 \text{ degrees from vertical})$, this deviation is greater for the upright and

head-down positions. For both postures, the bed was set to 80 degrees; therefore, in the upright posture the orientation of the lungs would typically be 10 ± 10 degrees from vertical, estimated from participant scans. The potential for deviation was greater in the head-down posture, due to the nature of the posture and because subjects were asked to assume a comfortable position for the scans. This permitted a high degree of deviation from the desired orientation as participants would position themselves to slightly horizontally. This is evident in Figure 5.15C. Therefore, the orientation of the lungs in this posture would typically be 20 ± 20 degrees to the vertical, estimated from participant scans.

5.2.2.2 Breathing Exercises

During the study, the subjects were asked to practice several different breathing exercises. This was to investigate how the lungs and associated functional distribution change in response to different inspiration volumes in each posture. These breathing exercises are depicted in Figure 5.9, including:

- Normal ("free") breathing Regular breaths over the tidal volume as the subject would conduct normally. As it was assumed the subject would be doing this without direction, they would only be instructed to free breathe if they were currently undertaking a different breathing exercise. This is because, paradoxically, being told to 'breathe normally' can make a person concentrate on each breath and in doing so may this result in abnormal breathing patterns.
- Deep breathing Larger and generally slower breaths than for normal breathing. Participants were not coached on this breathing exercise; however, these breaths were expected to result in an increased

tidal volume. This breathing exercise achieves a greater activation of the diaphragm and intercostal muscles.

The inspiration volumes were not measured externally for any of these exercises.

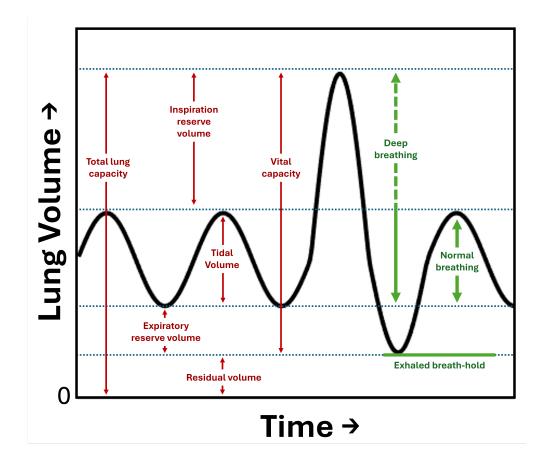


Figure 5.9: Depiction the breathing regimes employed by this study. The formal definition of lung volumes are shown in **red**, whilst the breathing exercises used in this study are indicated in **green**.

5.2.3 Participants

21 healthy volunteers were recruited for the study (5 male, 16 female), with ethics approval from the University of Nottingham Medical School Ethics Committee, no. FMHS 236-0323.

5.2.3.1 Eligibility

Further to the general eligibility requirements for having an MRI, the participants were selected based on:

- Age 18-35 years
- BMI above 18.5, below 30 kg m^{-2}
- No known heart conditions
- Blood pressure under 140/90 mmHg
- Not currently taking medications (not including contraceptive pill)
- No joint disorders

Furthermore, the FOV in the scanner is limited in size. Therefore, to ensure their entire lung regions fit within frame, particularly while deep breathing, an additional criterion was introduced:

• Under 180 cm in height

However, this requirement has an inadvertent impact on recruitment for this study: a gender imbalance. The average height of a 25- to 35-year-old man in England in 2021 was $177.6(\pm 6.9)$ cm[111], just below the inclusion criterion. Therefore, a significant proportion of men were ineligible for this study. If one only considers this criterion, 36.4% of this male demographic were ineligible. This had far less of an effect on the recruitment of women, who have an average height of $164.4(\pm 6.6)$ cm (for 25 to 35-year-olds in England in 2021[111]). Therefore, only 1% of this demographic were ineligible due to this criterion. As a result of this, the gender balance of

this study was heavily skewed towards women. Out of the 21 participants investigated, 16 were women and 5 were men.

5.2.3.2 Recruitment

Participants were recruited via word of mouth and posters within the SP-MIC. Initial communication was conducted via email. Each subject was sent the participant information sheet, study consent form and standard MRI safety questionnaire. This was completed in good time to enable them to make an informed decision regarding their participation in the study. During this period, they were able to ask any questions they desired. For those interested, this interaction also involved establishing if they met the general inclusion criteria. If they met the criteria and were still happy to participate, they were invited to attend an in-person screening visit.

5.2.4 Study Visits

A study in two parts:

This study was conducted in two cohorts. Approximately halfway through the study, the protocol was amended to remove the HASTE scans as these were deemed too unreliable for PREFUL post processing in their current form.

An additional posture was introduced – Prone, to maximally reverse the relative direction of gravity experienced by the supine lung. Furthermore, to measure posture dependent anterior to posterior functional distributions, sagittal scans using the 2D GRE sequence were introduced. However, this was omitted for the head-down posture as there was not sufficient time

within the 15 minutes permitted for safety.

These changes resulted in two protocols, to be referred to as Protocol 1 and Protocol 2. Of the 21 participants, 10 took part in Protocol 1 and 11 took part in Protocol 2.

5.2.4.1 Screening Visit

The screening visits were completed at the SPMIC or a convenient location on the University of Nottingham Park Campus. The purpose of the visit was to ensure they met all the study criteria, complete the requisite forms, and allow them to ask any further questions they may have. During this visit, the participants were given a paper copy of the Participant Information Sheet to keep. They also completed the Consent Form, MRI Safety Form and Tattoo Form, if required. The screening section of the Case Report Form was also completed by the study investigator present.

As part of this visit, the blood pressure of each participant was checked to ensure it was in the acceptable range. This was done using a portable OMRON blood pressure monitor. Their blood pressure was checked three times and an average taken to avoid 'white coat syndrome', where the stress of being checked artificially increases blood pressure. If the blood pressure was over 140/90 mmHg, it would be checked again ten minutes later. If the pressure remained outside of inclusion criteria, then the participant would be removed from the study. Their doctor would also be informed of the incidental finding.

At the end of the screening visit, the eligible participants were given instructions on how to find the SPMIC satellite department in the medical school and given a map of the building if necessary. They were also instructed to wear MRI-appropriate clothing, should they not wish to change into scrubs. Finally, participants were advised not to eat a significant meal within the hour prior to the scanning visit, to avoid discomfort due to pressure on the lower oesophageal sphincter while assuming the head-down posture.

5.2.4.2 Scanning Visit

The scanning visits were scheduled over email and were always within two months of the screening visit. Upon arrival at the screening visit, participants underwent a final screening to confirm their eligibility for their study. This was done to ensure no details had changed since the screening visit: for example, if their blood pressure had risen above acceptance criteria. Blood pressure was checked via the same method as per the screening visit and the 'Study visit' section of the CRF completed. Following the final screening, the subjects were given the opportunity to change into scrubs, if required.

Details of the scans conducted for each participant in each posture are detailed in Section 5.2.1.2. For both protocols, the participants were first scanned in the sitting posture, followed by head-down and supine. Participants completing protocol 2 then took part in an additional posture prone. This order was chosen so as to familiarise the participants with the scanning protocols before undertaking the head-down posture and reduce setup time between scans for rearranging the scanner setup. For example, a similar setup was used for the sitting and head-down scans, as well as for the prone and supine scans.

It was often deemed more efficient to scan multiple participants during the same session. In this case, each participant was scanned in each posture sequentially. For example, if there were two subjects in the same session, each subject would be scanned in the first posture before moving onto the next. Generally, three participants were scanned in a session for protocol 1 and two participants scanned in a session for protocol 2, each resulting in a total session length of approximately three hours. The protocol descriptions that follow concern a study session with only one participant.

Descriptions for each protocol for each of the three postures are presented in the sections below. Each protocol ended upon completion of the experiments in the final posture: supine for protocol 1 and prone for protocol 2. Upon completion of this position, the subject was removed from the scanner and permitted to change back into their regular clothes, if necessary. This marked the end of the scanning session, and the participant was released to go about their day. There were no repeat visits.

5.2.4.2.1 Scanning Protocol

A timeline for the protocol is presented in Figure 5.10, with the experiments conducted for each posture summarised in Table 5.1.

1. Sitting - For both protocols, the first scanning posture was sitting upright (Figure 5.8A). This order was chosen to familiarise each participant with the scanning environment and procedure while in a comfortable position. Furthermore, this position required a similar scanner bed set up to the head-down posture, reducing down-time between scans for the bed to be reset. The hinge at the centre of the support platform was set to 110 degrees as per Figure 5.8A. The coil was positioned such that the top was approximately level with the subject's shoulders while sitting. If multiple subjects were scanned in the same session, the coil remained in this position throughout

each of their sitting scans. If there was significant variation between the height of the participants being scanned, they could be raised or lowered to fit the coil with foam pads beneath them.

The first experiments were the scout scans. These were conducted to ensure the subject was positioned correctly in the field of view and for slice positioning (Figure 5.7). Correct subject positioning was determined by the location of the lungs within frame. The whole lung needed to be visible, while leaving room below the lung to accommodate its expansion due to deep breathing. If the participant was positioned satisfactorily, then the resultant scout images were used to position a coronal slice over the lungs. The position of this 2D slice was confirmed by collecting ten 2D GRE images while the participant was deep breathing, ensuring the diaphragm remained visible for all frames. If required, the subject could be raised or lowered via the addition or removal of foam pads beneath them, before repeating the above process. When positioned satisfactorily, this slice position was used for all single slice coronal acquisitions of this posture.

When the slice position was confirmed, the acquisition was ready to begin. The first scans were the coronal 200-image 2D GRE normal breathing experiments. Immediately following the completion of this sequence was the 200-image coronal HASTE normal breathing experiment. As they should already be breathing normally, no instruction was needed for the participant. Indeed, their only indication that the next experiment had started was the change in the acoustic noise from the scanner. The next scans were the coronal deep breathing experiments, 200-image 2D GRE and 200-image HASTE sequences. The subject was asked to deep-breathe, and upon confirmation they had complied, the scans were initiated. The scan parameters remained the

same as for the corresponding normal breathing experiments. There was no pause nor instruction given to the participant between the transition from 2D GRE to HASTE experiments.

Following the completion of these scans, the sitting section of Protocol 1 was complete. The participant was removed from the scanner and returned to a waiting room. There would then follow at least a 10-minute wait period before the participant continued with the next posture, as per the protocol.

2. **Head-down** - The second part of protocol two was the scans of the participant in the head-down posture (Figure 5.8D). After familiarisation with the protocol during the sitting scans, these scans had to adhere to a strict 15-minute time limit for the safety and comfort of the participant. This covered the full time the participant was inverted, including positioning, scanning sequences and down time between scans. Therefore, efficiency was imperative.

The participants were first asked to practice entering and exiting the head-down posture. This was to familiarise themselves with each procedure and because generally they reported entering a more comfortable position on the second attempt. This practice attempt was not included in the 15-minute scan time limit as the participant was permitted some recovery while upright before entering the process. When they reported feeling ready to proceed, the participant was instructed to enter the head-down posture once again. At this point, the 15-minute time limit began. Before checking the position with the scanner, their lateral position could be corrected by eye as they might not have positioned themselves vertically in the coronal plane. This could happen when they did not position their forearms directly below their hips and therefore their body assumed a slight lateral

angle. With permission from the participant, this could be quickly corrected by slightly sliding their hips left or right, so their body appeared to be vertical in this plane. At this time, the angle of their body in the sagittal plane was not corrected as this position was chosen by the participant to be comfortable. When in position, the researcher, remaining in the room to monitor the participant, signalled to the scanner operator to begin the Scout sequence. This sequence lasted approximately 40s and was used for final positioning quality control. As achieving perfect participant positioning could be time-consuming, a lower quality of positioning was acceptable for this posture compared to the others. Although it remained imperative that the full lung was captured in frame. If necessary, they could be raised or lowered with the addition or removal of pillows/thin foam blocks below their forearms. They could also be asked to assume a more vertical posture in the sagittal plane to quickly raise them. A single additional scout sequence was permitted for positioning quality control during the 15-minute time limit. If their position remained unsatisfactory, the participants were asked to exit the scanner and permitted some down time while upright to reset the time limit. When reporting feeling ready to continue once again, the positioning protocol was repeated from the point prior to completion of the practice positioning. With satisfactory positioning determined by the scout images, the slice was positioned with the same criteria as detailed in Figure 5.7. Ten coronal 2D GRE images were captured while the participant was deep breathing to ensure their full lungs remained in frame. If necessary, repositioning could occur via repeating the procedure above.

With the satisfactory positioning of the participant, the full scan-

ning procedure was ready to begin. If still breathing deeply from the positioning scans, the subject was asked to breathe normally and the coronal 200-image 2D GRE sequence was initiated. The 200-image HASTE sequence was initiated immediately upon its conclusion. Combined, these sequences took approximately three minutes. When completed, the researcher communicated with the participant to confirm they were happy to continue. If not, they were removed from the scanning posture and the following scans omitted. If happy to proceed, the participant was instructed to begin breathing deeply and upon their compliance, the scanner operator signalled to begin the next scans. The additional coronal 200-image 2D GRE and 200-image HASTE sequences were initiated in sequence for a further three minutes of scanning time. These were the final scans for protocol 1 in the head-down posture.

With the completion of the final scans, the head-down posture imaging was complete, and the participant was asked to exit the posture but remain sitting upright on top of the foam blocks behind the chair. The researcher within the room communicated with the participant to assess their condition and if they were experiencing any side effects. When they no longer felt any effects, they were assisted down from the blocks and to a waiting area where they could recover from the head-down posture for at least ten minutes before moving on to the next scanning posture.

3. Supine - The final posture for Protocol 1 was supine. For this posture, the hinge at the centre of the support mechanism was set to 180 degrees, or horizontal, to resemble a bed (Figure 5.8B). Positioning was far easier in this posture than for sitting and head-down and no time limit was imposed. The coil was positioned over the participant

such that their shoulders were level with its end and then the bed moved to position it at isocentre of the magnet. The Scout sequence was used to confirm the position of the participant was correct. If not, the participants were asked to slide themselves up or down the bed as required, and this was checked again with an additional scout sequence. Coronal slice positioning followed the same criteria as the sitting and head-down postures and was checked with 10 images captured with the 2D GRE sequence while the participant was deep breathing.

Once satisfactory positioning was achieved, the normal breathing coronal 200-image 2D GRE sequence was initiated, followed immediately by the 200-image HASTE sequence. Upon completion of these scans, the participant was asked to breathe deeply and the same sequences were repeated. These were the final scans for the supine position and concluded the experiments for Protocol 1. Upon completion of these scans, the participant was removed from the scanner and taken to the waiting area.

4. **Prone** - The final posture for Protocol 2 was the newly introduced prone position (Figure 5.8C). The participant was placed on their front with the coil enclosing their torso, level with their shoulders. Positioning was first confirmed by eye and then with the Scout acquisition. If adjustment was needed the participant was asked to slide themselves up or down the scanner bed before repeating the scout acquisitions. This posture was introduced to complement the supine posture for investigations in the sagittal plane. For this reason, all non-scout coronal-experiments were omitted. Therefore, following satisfactory participant and slice positioning, detailed in Figures 5.8C and 5.7 respectively, the only scans in this posture were the 200-image

sagittal normal breathing 2D GRE experiments.

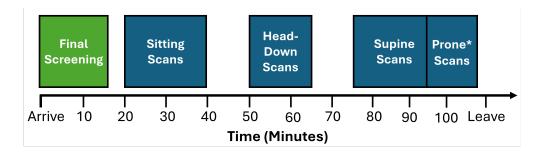


Figure 5.10: Timeline for Protocols 1 & 2, comprising a final screening, followed by scans in four postures. The first three postures were separated by a ten-minute break for participant recovery and reconfiguring the participant support mechanism. However, there was no break between the supine and prone postures as participants were instructed simply to roll over in the scanner. * denotes only for protocol 2.

Protocol 1								
Sequence	Plane	Breathing	Sitting	Head-Down	Supine	Prone**	Time (m)	
Scout	All	Normal	√	√	✓	√	0:31	
2D GRE	Coronal	Normal	✓	✓	✓	X	1:17	
HASTE*	Coronal	Normal	√	✓	✓	X	2:53	
2D GRE	Coronal	Deep	√	√	✓	X	1:17	
HASTE*	Coronal	Deep	√	✓	✓	X	2:53	
2D GRE	Sagittal	Normal	√	X	✓	✓	1:17	

Table 5.1: Timeline of MRI Sequences completed for each posture in Protocol 1. * Only in Protocol 1, ** Only in Protocol 2.

5.2.5 Data Analysis

The scanner acquires each image as a 64 x 64 pixel matrix and applies zerofilling and filtering before the Fourier transform. A warp is then applied to each image to counter the known B0 magnetic field inhomogeneities, which are calculated based on the size and position of the slice. The effect of this is sometimes visible in the periphery or corners of the image (particularly in Figure 5.7). Finally, the scanner interpolates the image onto an 89 x 89 voxel grid and outputs each image as its own DICOM file.

5.2.5.1 PREFUL

Functional maps were generated by applying the PREFUL algorithm[5] to the resultant image series. This was done using a home-built pipeline, applied in MATLAB 2024b. A full description of this process can be found in Chapter 4. Processing parameters used for this study are detailed in Table 5.2.

PREFUL Processing Parameters					
Registration Method	Region, Forsberg				
Respiratory marker	2D Lung Area				
Respiratory filtering method	Low Pass				
(Hz)	Pass: 0.4				
	Stop: 0.5				
No. interp. ventilation frames	20				
Ventilation interp. width	0.05π				
Cardiac marker	Aorta Intensity				
Cardiac filtering method	High Pass				
(Hz)	Pass: 0.5				
	Stop: 0.4				
Cardiac Fitting R2 Thresh.	0.7				
No. interp. perfusion frames	20				
Perfusion interp. width	0.05π				

Table 5.2: PREFUL processing parameters.

5.2.5.2 Regional Analysis of Functional Maps

5.2.5.2.1 Lung ROI Sectioning

The PREFUL processing detailed above resulted in the creation of five 2D functional maps for each acquisition: FV, VTTP, weighted perfusion, QTTP and Jacobian determinant. Maps of the exhaled signal intensity were also created. To compare the distributions of these functions for each posture, it was useful to reduce the dimensionality of the data. This was achieved by projecting the variation within each map onto two orthogonal axes: apical-to-basal for the coronal acquisitions, and apical-to-basal and

anterior-to-posterior for the sagittal acquisitions. Functional distributions along these axes were assumed to be independent of each other, although the validity of this assumption is discussed later.

All acquisitions were transformed as necessary (by rotation or reflection) to standardise the lung orientation between postures. In the coronal plane, this resulted in a lung orientation equivalent to observing the lungs as viewed from in front of a sitting individual, with the apex at the top and the right lung to the left of the image. In the sagittal plane, the lungs were transformed to orient the apex and anterior surface to the top and left of the image, respectively. Examples of each standardised orientation are depicted in Figures 5.11A and 5.12Ai.

The projection of each functional result onto each axis was achieved via an averaged zonal method, similar to previous studies[4, 7, 58]. Instead of plotting each voxel's value individually, the mean value of multiple distinct zones was used to visualise the posture-dependent distributions. The number of zones was chosen to balance resolution and SNR, due to the width and number of voxels combined by each zone. This method significantly improved the SNR of the results compared to plotting the means of individual lines of pixels. Along the apical-to-basal axis, each lung was sectioned into several vertically stacked rows of equal height. This was achieved by isolating the pixels nearest to the apex and to the base, then distributing each row's boundary at equidistant intervals between them, rounded to the nearest pixel. Typical examples of the resulting sectioned zones for each imaging plane, sectioned into five zones, are depicted in Figures 5.11A and 5.12Ai. Each zone was assigned a fractional position along the apical-to-basal axis, corresponding to the vertical position of its centre. This fractional approach was chosen in preference to an absolute distance to permit easy combination of participant data and comparison between

postures. However, this method is sensitive to the size of the lung at the intermediary respiratory state. For each participant, the functional values and their associated errors for all voxels within each zone were used to calculate the weighted mean and standard error of the zone. These values were then combined using a weighted mean across all participants, yielding a single functional distribution for each posture and breathing regime. These values were used to populate a graph of the variation along the apical-to-basal axis. Examples of such plots, with artificial data points, are presented in Figures 5.11B and 5.12Aii. Figure 5.11 represents the standard format employed for the presentation of coronal results in this chapter: a plot of the functional distribution with respect to the associated axis, presented next to an example mask to aid in the visualisation of the data. The apical-to-basal axis is orientated vertically, matching the orientation of the lung depicted alongside. However, this means that, contrary to standard convention, in these plots the dependent variable is plotted on the horizontal axis.

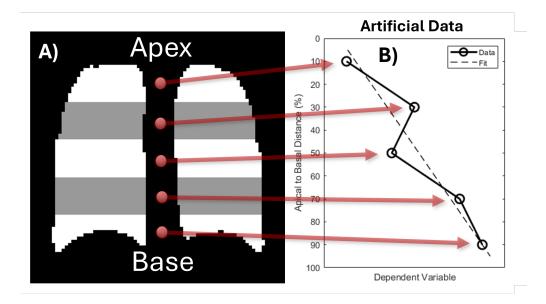


Figure 5.11: This figure depicts the standard format for plotting the functional variation along the apical-to-basal axis of coronal acquisitions in this work. An example five-zone sectioned mask (A) is presented, with an accompanying artificial apical-to-basal distribution (B). The **red** arrows indicate which lung sections correspond to their respective data points. The **dashed line** represents the linear fit of the artificial data.

The investigation along the anterior-to-posterior axis follows a similar method to that of the apical-to-basal axis. However, instead of vertically stacked rows, the lung is sectioned into several horizontally adjacent columns of equal width. The columns are distributed based on the location of the most anterior and most posterior pixels. A five-zone example sectioned mask is depicted in Figure 5.12Bi. The mean value of each zone was calculated, which in turn were averaged over each participant. The resulting functional distributions were plotted against the anterior-to-posterior distance, oriented horizontally to reflect the lateral adjacency of the zones. Figure 5.12 shows the standard format used in this chapter for presenting sagittal results.

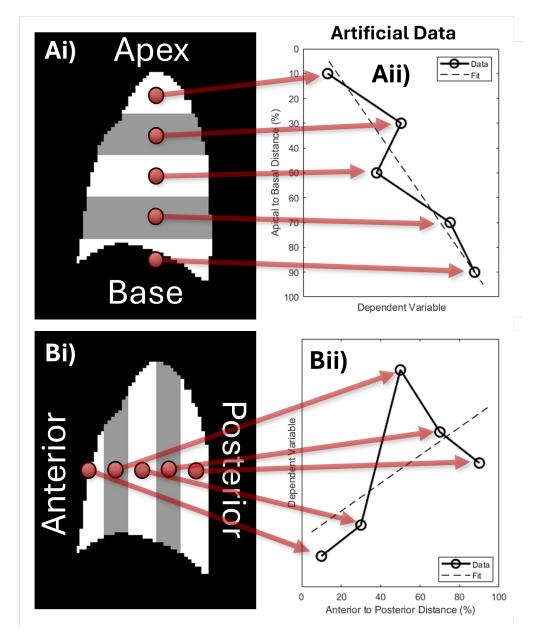


Figure 5.12: This figure depicts the standard format for plotting the functional variation along the apical-to-basal and anterior-to-posterior axes of sagittal acquisitions in this work. An example five-zone sectioned mask is presented; it is sectioned along the apical-to-basal (Ai) and anterior-to-posterior axes (Bi). The masks are plotted adjacent to their corresponding artificial apical-to-basal (Aii) and anterior-to-posterior (Bii) distributions. The red arrows indicate which lung sections correspond to their respective data points. The dashed line represents the linear fit of the artificial data.

For all investigations in the coronal plane in this study, the lung was sectioned into 10 zones in order to attain adequate SNR. The number of zones was reduced to five for investigations in the sagittal plane, to reflect the

lower number or participants investigated.

5.2.5.2.2 Mask Adaptation

The above method for sectioning each lung is based on the associated mask. For the FV, the VTTP, the perfusion, and the QTTP, this is based on the mask of the intermediary lung state produced by the PREFUL processing. However, as the exhaled signal intensity and map of the Jacobian determinant is based on the exhaled lung, these are sectioned based on the mask of the maximally exhaled state.

Figures 5.11 and 5.12 depict the sectioning process based on a complete full mask of the lungs. However, this was a simplification for explanatory purposes and in reality each mask is adapted slightly before applying the sectioning algorithm. Figure 5.13 depicts this adaptation process applied to all masks. Starting with the base masks (Figure 5.13Ai and 5.13Bi), the major pulmonary vessels are removed. These do not represent the state of the local lung parenchyma[56] and thus they can skew the results of the central lung regions. The vessel removal is achieved using the k-means clustering method described in Chapter 4. Which is applied to either the perfusion-weighted data for the FV, VTTP, normalised perfusion and QTTP, or to the signal intensity of the most exhaled state for the exhaled signal intensity and map of the Jacobian determinant. Example lung masks with the vessels highlighted are presented in Figures 5.13Aii and 5.13Bii.

In addition to the removal of the major pulmonary vessels, the corners of the mask were also removed. This serves two purposes: firstly, these are the areas are likely to experience registration errors. They are surrounded by tissue not associated with the lung parenchyma on multiple sides. Generally, this surrounding tissue returns a significantly greater signal in MR images than that of the lung. Therefore, any slight errors in the registration can create artefacts in the functional maps. Secondly, these areas can skew the segmentation algorithm, which is based on the most peripheral pixels. These pixels are often those at the corners, especially at the base of the lung. Furthermore, these corner regions may extend significantly beyond the rest of the lung, particularly in the basal-posterior region. This will skew the placement of the lung segments and create a basal region with a reduced number of voxels, evident in Figure 5.12Ai. The corner removal step is achieved by the application of a morphological 'Open' operation using a 2D structuring element in the form of a disk with radius 5 voxels. This radius was chosen by eye to best suit the resolution of the acquired data, corresponding to approximately 1.8 cm. Example maps with the corners highlighted are presented in Figures 5.13Aiii and 5.13Biii. These features are determined for each acquisition and subsequently removed from the base mask.

Removal of both the major pulmonary vessels and lung corners from the base mask results in the mask used by the sectioning algorithm, here, called the 'analysis mask'. Examples of these are depicted in Figures 5.13Aiv and 5.13Biv for the coronal and sagittal planes, respectively.

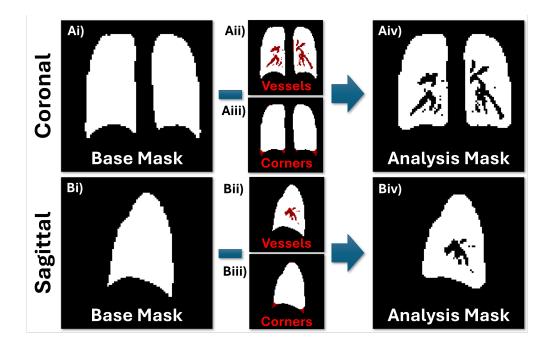


Figure 5.13: Mask preparation prior to segmentation for an example mask in the coronal (A) and sagittal (B) planes. Starting from the generated lung mask (i), this figure details the removal of the vessels (ii) and corners (iii) to create the mask used for functional distribution analysis (iv).

In the perfusion-weighted and Jacobian maps, a significant central to peripheral gradient was observed, likely due to the central lung vessels. Although, this gradient was not investigated, it skewed the variation of each function when investigated along each axis. Therefore, it was beneficial to only investigate variation within parenchyma located at an equal distance to the centre of the lung. In this case, the lung peripheries were segmented for this purpose. The four-step method for automatically segmenting the peripheral segments of a coronal mask is depicted in Figure 5.14 for the coronal plane. Starting with the base mask (Figure 5.14A), a 'Close' morphological operation is applied using a 2D structuring element in the shape of a diamond with equivalent dimensions as that of the acquisition. This results in Figure 5.14B, where the area between the two lungs has been filled. The second step is the inversion of this mask, generating Figure 5.14C, where the non-lung areas are now highlighted, except between the

lungs. A final morphological operation is applied to this mask, a 'Dilation' using a 2D structuring element resembling a disk, with a radius chosen to reflect the desired width of the peripheral segment. The width of this segment was selected to balance SNR through the number of voxels included in each zone, while avoiding contributions from central regions. For this work, a width of 7 voxels (approximately 2.5 cm) was chosen to match the resolution of the data. This expands the non-lung areas, extending them into those occupied by the lungs, although again omitting the centre of the lungs. This results in the penultimate mask depicted in Figure 5.14D. The final step is to compare the original mask with the penultimate mask. Any voxels where they are both equal and non-zero are determined to represent the peripheries. For each acquisition, this results in a mask like that depicted in Figure 5.14E. Although not depicted in this figure, this process is always applied to a base mask with the corners removed (Figures 5.13Aiii). Furthermore, any major pulmonary vessels in this peripheral segment are removed by the same process as detailed previously. This process can also be applied to masks in the sagittal and axial plane, although the first stage is redundant for the sagittal plane as only a single lung is visible. However, in this work, the peripheries of the lung were only investigated in the coronal plane.

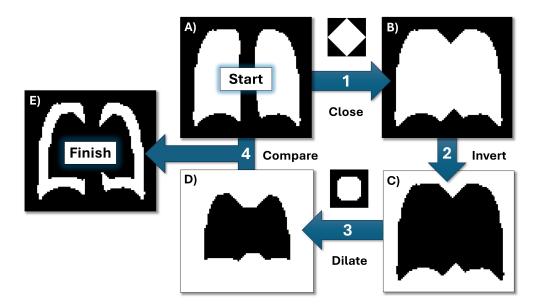


Figure 5.14: Summary of the pipeline employed to segment the peripheral regions of the lung in the coronal plane. Starting with an example coronal binary mask (A), a morphological close operation is applied using a diamond-shaped 2D structuring element to fill in the region between the lungs (B). Following the inversion of this mask (C), a morphological dilate operation is applied using a disk-shaped 2D structuring element, to shrink the size of the of the void region (D). Finally, this mask is compared to the original; all values that are equivalent across both masks and equal to 1 now make up the final peripheral segmentation mask (E).

Finally, in the sagittal plane it was also useful to investigate the apical-to-basal variation within the anterior and posterior lung segments independently. To achieve this, as with the anterior-to-posterior sectioning, the most anterior and posterior voxels were isolated. Then, the central point between these, rounded to the nearest voxel, was used to split the lung mask vertically in two segments.

Note: In this work, the lung mask used for the analysis along each of the axes is displayed alongside the data. This mask includes the removal of the major pulmonary vessels, peripheral segmentation, or the anterior and posterior isolation. However, all masks are depicted omitting the removal of the corners. This was to maintain the shape of the lungs so they would be easily recognisable at a glance. Even though it is not displayed, this

operation was applied to all masks before sectioning the lungs along each axis.

5.2.5.2.3 Normalisation

It is sometimes necessary to normalise the functional maps before participant data can be combined or comparisons are drawn between postures. Both the FV and Jacobian determinant data are sensitive to the magnitude of lung volume change during the acquisition. Not only does this vary between participants, but it is also dependent on the respiratory thresholding chosen during PREFUL processing. Therefore, whilst there is some utility in comparing the magnitude of these functions between postures, to compare the relative variation between them it is useful to normalise the functional data. The perfusion-weighted and exhaled signal maps are based directly on the signal acquired by the scanner, which is in itself dependent on a multitude of factors. Therefore, the relative magnitude of these cannot be compared between acquisitions. These results must be normalised before the participant data can be combined, and then it is only useful for comparison of the relative variation within each posture. Finally, the VTTP and QTTP results do not require normalisation. These functional metrics are already scaled to a point within a single reconstructed ventilation or perfusion cycle.

The normalisation step is achieved by dividing each functional map by the mean value of all voxels within the analysis mask. This converts all voxels of the functional maps into a fractional value, which is subsequently plotted as a percentage, relative to the mean lung value, which is located at 100% of the normalised functional axis.

The perfusion-weighted maps underwent an additional normalisation step.

This was to remove contributions from tissue compression. To achieve this, each voxel is divided by its estimated tissue density. As the data was collected during respiration, the tissue density associated with each voxel is constantly changing essentially independently from the cardiac cycle. Although a high-pass filter is applied to isolate the perfusion variation, it also removes the zero temporal frequency component, which is assumed to be equivalent to the mean tissue density. This is why one cannot simply repeat Equation 4.23 to calculate a fractional perfusion. Instead, the mean value of that voxel throughout the reconstructed registered time-series is assumed to be equivalent to the zero temporal frequency component, which, in turn, represents the mean tissue compression during the perfusion cycle. However, this data must be further normalised by the value of the lung as a whole. This is because the ratio of the average lung signal to the perfusion variation is dependent on the cardiac frequency. This effect is discussed further in Chapter 4; however, it stems from the acquisition time of each frame being significant compared to the cardiac period, resulting in an averaging of signal variation and thus a reduction in apparent amplitude.

5.2.5.2.4 Linear Fitting

Occasionally, to compare functional distributions between postures, it was useful to reduce the dimensionality of the data further to just a single number. In this case, the data were fitted to a linear relationship.

$$y = m \cdot x + c \tag{5.1}$$

where y represents the functional value, x represents the fractional distance, m is the slope and c an offset. All fits were calculated using a weighted non-

linear least squares regression algorithm[112–114], which also returned the standard error $(\delta m, \delta c)$ associated with each parameter. Fits of this type were only calculated for normalised data and therefore the offset value (c) is redundant. However, the resulting slope (m) was used to quantitatively compare the relative change between postures. This value represents an approximation of the total change across the lung along the respective axis, in proportion to the average value of the lung.

Each fit is plotted over its respective data in the form of a dashed line. These fits are based on the assumption that the respective distributions are linear with respect to fractional distance within the lung. However, this is often not the case. This is evident both qualitatively, by comparing the linear fits to the often non-linear data, and quantitatively by the errors associated with the calculated slopes. Therefore, the fit data should not be taken in isolation, but in conjunction with the data it represents.

5.3 Results

5.3.1 General Results

5.3.1.1 Participant Positioning

In this study, the participants were instructed to assume a comfortable position in each required posture. However, this approach often resulted in the long axis of their lungs and associated coronal imaging planes not being precisely aligned parallel or perpendicular to the gravitational vector. This was particularly evident during investigations of the head-down posture as participants would often orientate themselves at a significant angle

to the vertical. Beyond variation in participant positioning, the coronal slices often required angular adjustments to simultaneously capture both the aorta and the diaphragm, whilst avoiding the heart. This resulted in a compounding variation in the angle of coronal slice positioning. The average angle to the vertical of the coronal slices for each posture is presented quantitatively in Table 5.3 and visually in Figure 5.3.

Coronal Slice Positioning		
Posture	Angle to the Vertical (°)	
Sitting	6 ± 4	
Supine	89 ± 3	
Head-Down	22 ± 7	

Table 5.3: Summary of participant positioning in the scanner and resultant orientation of the coronal slice relative to the vertical.

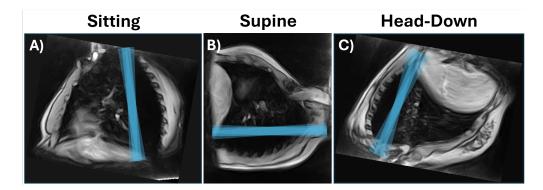


Figure 5.15: Summary of participant positioning in the scanner and resultant orientation of the coronal slice relative to the vertical, including error. All scans are orientated so the direction of gravity is down the page. Slice thickness is not to scale.

5.3.1.2 Sensitivity Gradients

All investigations in this study assume a common proportionality between signal intensity and tissue density throughout each image. However, as evident in Figure 5.16, tissue immediately adjacent to the receive coil can

return a greater magnitude of signal compared to equivalent tissue elsewhere. This effect drops off quickly with distance, however, it may lead to errors during sagittal investigations.

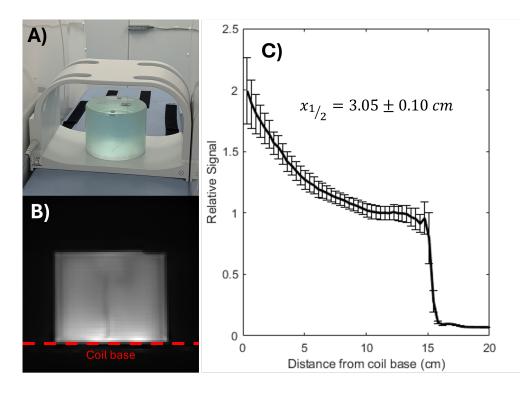


Figure 5.16: This figure demonstrates the non-uniform signal variation due to a sensitivity gradients near the receive coil. A cylindrical phantom containing saline solution (A) was placed within the coil, resting on the back plate (or base). 60 images were captured using the 2D gradient echo sequence (Section 5.2.1.2), with the final 50 combined to create an average image (B). This image was used to investigate the relative signal decay with distance from the coil back plate (C). This distribution was fitted to an exponential function $(y = a \cdot e^{-bx} + c)$ via non-linear least squares regression.

5.3.1.2.1 Investigating B1-Transmit Homogeneity

The B1-transmit field is generated from coils integrated in the scanner. It was preliminarily investigated by using an oblong phantom of uniform spin density throughout, placed within the central region of the body coil. The B1-transmit variation throughout this phantom was investigated using the interleaved double TR method[115]. A multi-slice acquisition was used to

give insight into B1-transmit variation in three dimensions. Figure 5.17 displays B1 variation through the centre of the phantom in the coronal and sagittal plane.

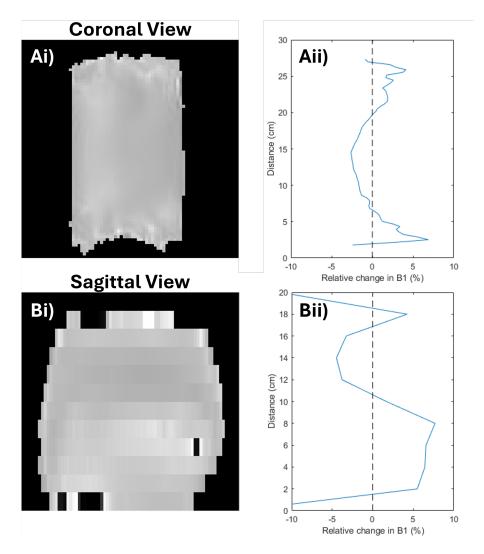


Figure 5.17: Investigation of the B1-transmit homogeneity within an oblong phantom of uniform spin density. The B1-transmit maps are displayed in the central coronal (Ai) and sagittal plane (Bi), with the corresponding mean variation along these axes plotted in Aii and Bii, respectively. The axes are positioned to reflect the spatial variation of the field, relative to the mean field within the plane.

The average variation within the phantom along these planes appears to be limited to approximately $\pm 5\%$. Therefore, this variation is not expected to have a significant impact on the results in this study.

5.3.1.2.2 Investigating B1-Receive Homogeneity

The B1-receive field is due to the body coil. It was preliminarily investigated using the same phantom setup as described for the B1-transmit investigation. 2D Images through the central coronal and sagittal planes were acquired using the same method detailed in Figure 5.16. Figure 5.18 depicts these acquisitions accompanied by the mean variation in signal intensity through these planes.

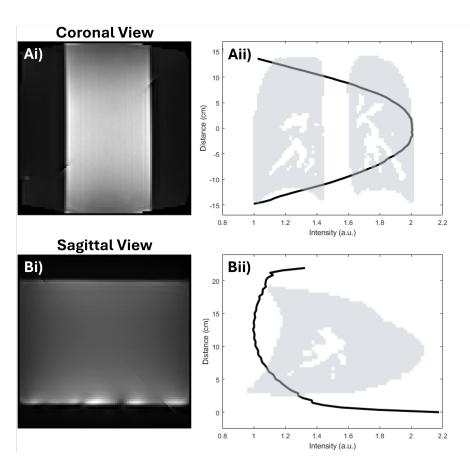


Figure 5.18: Investigation of the B1-receive homogeneity within an oblong phantom of uniform spin density. The B1-receive maps are displayed in the central coronal (Ai) and sagittal plane (Bi), with the corresponding mean variation along these axes plotted in Aii and Bii, respectively. The axes are positioned to reflect the spatial variation of the field, relative to the mean field within the plane. Depictions of the lungs are provided to aid understanding but are not to scale.

Figure 5.18 depicts significant variation in B1-receive field towards the cen-

tre of the coil in the coronal plane and towards the base of the coil in the sagittal plane. These variations may impact the investigation of functional parameters along these axes.

5.3.1.2.3 Estimating the Impact of Sensitivity Gradients on Functional Parameters

Spatial variations in sensitivity will likely have an effect on the observed distribution of the functional parameters discussed in this chapter. The impact on the exhaled tissue density and perfusion variation is intuitive, as these are directly dependent on the signal intensity. Therefore, a greater sensitivity will correspond to apparently greater values of these parameters. When the perfusion is normalised by the tissue density, this effect should be reduced as both these factors will be influenced by the sensitivity gradients.

Fractional ventilation is calculated based on the changing value of lung voxels. However, these voxels are typically located at different spatial positions at different points of the respiration cycle. Therefore, the sensitivity gradients can have an impact on the distribution of this functional parameter. For example, a voxel moving towards the centre of the coil with inspiration will exhibit a reduction in signal due to reduced tissue density, however this will be countered by increased signal due to the increased coil sensitivity in this region. Therefore, the observed fractional ventilation will be suppressed.

To estimate the magnitude of the error due to these effects, one must make assumptions about the size and position of the lungs within the coil. The exhaled lung size in the investigated planes when sitting is estimated from scans in this study to be approximately 18 ± 2 cm in the apical to basal axis, and 13 ± 2 cm in the anterior to posterior axis. When inhaled and

breathing normally, these increase to 20 ± 2 cm and 14 ± 2 cm, respectively. Finally when deep breathing, the apical to basal distance increased to 22 ± 2 cm. Using these values and assuming the lung is positioned in the centre of the coil when fully inhaled, one can simulate the apparent variation of each functional parameter, within the phantom of uniform spin-density throughout. Using the same fitting linear fitting method discussed previously, this results in the gradients presented in Table 5.4.

Effect of B1 on Functional Parameters			
Parameter	Breathing	Apical to Basal	Anterior to Pos-
		slope (%)	terior slope (%)
Fractional Ventilation	Normal	-9 ± 2	1.2 ± 0.4
	Deep	-20 ± 4	
Perfusion & Exhaled	Normal	20 ± 8	16 ±4
Tissue Density	Deep	15 ± 8	

Table 5.4: Estimated impact of the B1 inhomogeneity on the observed gradients of functional parameters along the apical to basal and anterior to posterior directions.

Based on the values in Table 5.4, it appears that the bias field may introduce a significant systematic error to the results in this study. However, these are only preliminary estimates, with the true effect dependent on the size and position of each lung individually. Therefore, these factors are not accounted for in the results presented in this work. In future, bias field correction should be employed on a scan-by-scan basis.

5.3.1.3 Spatial Averaging

When investigated in the sagittal plane, the lung is not perfectly symmetric about either the apical-to-basal or anterior-to-posterior axes. Therefore, measures along each of these axes were not totally independent of each other. This interdependence is plotted in Figure 5.19. This plot follows the same style as the previously presented results; however, the dependent variable has been substituted with the mean position of each segmented zone along the perpendicular axis.

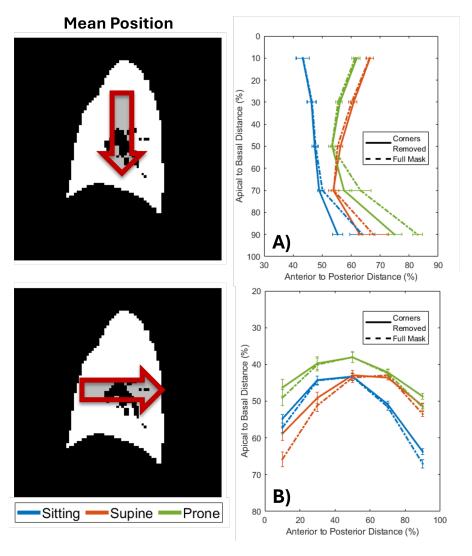


Figure 5.19: Measure of the interdependence of the apical-to-basal and anterior-to-posterior segmentation in the sagittal plane due to the shape of the thoracic cavity. It is composed of data averaged across all participants who took part in protocol 2. The data are plotted along the apical-to-basal (A) and anterior-to-posterior (B) axes, with the dependent variable representing the average position of each zone along the perpendicular axis. The date is plotted for the segmentation of the full mask and the masked used for segmentation with the corners removed (Figure 5.13).

Along the apical-to-basal axis, all postures exhibited a bias towards the posterior in the most basal region. This is due to the non-symmetric shape of the diaphragm as it extends more basally in the posterior region. This may have manifested as an artificial variation in signal intensity in this region due to the above-mentioned sensitivity variations associated with the inhomogeneous B1 receive field. For the sitting and supine postures, this may have increased the signal in this zone due to the proximity of the posterior regions to the receive coil. In the prone posture, this may have reduced the signal in this region due to the greater distance to the receive coil.

Along the anterior-to-posterior axis, the central regions of the lung exhibited an apical bias compared to the peripheral regions. This is likely due to the convex nature of the lung at the apex and the concave nature of the lung at the base. Measurements of the exhaled tissue density, FV and normalised perfusion generally increased along the apical-to-basal axis; therefore, this may slightly reduce measures of these functions in the central regions when measured along the anterior-to-posterior axis.

5.3.1.4 Toleration of the Head-down posture

Most participants appeared to tolerate the head-down posture well, reporting only expected side effects such as light-headedness and dizziness. These subsided within a minute of exiting the posture. However, two participants did not complete the head-down protocol. The first experienced feelings of claustrophobia when practising the posture. The participant and investigators agreed to omit this position. Second, another participant terminated the head-down scanning protocol following the deep breathing scans. They reported feeling uncomfortable and did not want to continue with the pro-

tocol. Notably, these participants were part of the same two-participant session. Therefore, one may conclude that the second participant, upon seeing their peer omit the head-down protocol, felt more comfortable terminating the procedure when they felt too uncomfortable. Worryingly, this may suggest that the other participants experienced greater side effects than they reported, but they didn't want to terminate the protocol for fear of being a burden to the investigators. Importantly, the side effects felt by this second participant subsided within a minute of exiting the posture as per the other participants in the study. Both participants went on to complete the protocol for each of the subsequent postures.

5.3.1.5 Protocol Completion Summary

These completion of each posture are summarised in Figure 5.20.

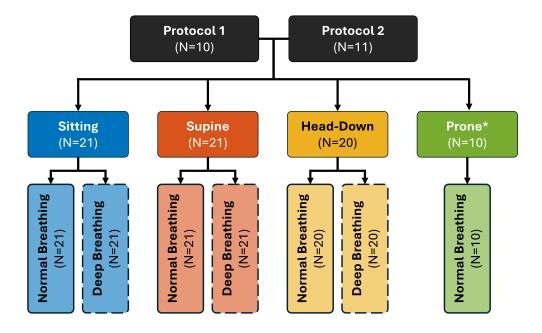


Figure 5.20: Summary of the number of successfully completed postures. * Indicates this posture was only completed as part of protocol 2.

5.3.1.6 Processing Success Rate

5.3.1.6.1 Coronal Acquisitions

The number of successfully processed coronal datasets for ventilation and perfusion is detailed in table 5.6.

Successful PREFUL Processing in the Coronal Plane			
Posture	Breathing	Ventilation	Perfusion
	Regime	N (Rate)	N (Rate)
Sitting	Normal	20 (95.2%)	21 (100.0%)
(N=21)	Deep	21 (100%)	17 (81.0%)
Supine	Normal	20 (95.2%)	20 (95.2%)
(N=21)	Deep	18 (85.7%)	20 (95.2%)
Head-Down	Normal	17 (85.0%)	18 (90.0%)
(N=20)	Deep	19 (95.0%)	18 (90.0%)

Table 5.5: Summary of the number of coronal datasets that have undergone successful processing from the PREFUL pipeline.

5.3.1.6.2 Sagittal Acquisitions

The number of successfully processed sagittal datasets for ventilation and perfusion is detailed in table 5.6.

Successful PREFUL Processing in the Sagittal Plane			
Posture	Ventilation Perfusion		
	N (Rate)	N (Rate)	
Sitting (N=10)	6 (60.0%)	8 (80.0%)	
Supine (N=10)	10 (100.0%)	10 (100.0%)	
Prone (N=10)	9 (90.0%)	10 (100.0%)	

Table 5.6: Summary of the number of sagittal datasets that have undergone successful PREFUL processing.

5.3.2 Variation in Exhaled Signal Intensity

5.3.2.1 Investigations in the Coronal Plane

Prior to analysing functional changes within the lung, it is useful to observe the variation in the proton signal at the point of maximal exhalation. For each participant, this reference state was determined by calculating the mean signal intensity across the three most exhaled frames of the reconstructed ventilation cycle, derived from the PREFUL processing pipeline. Figure 5.21 presents representative images of a single participant in the sitting, supine, and head-down postures at reconstructed maximal exhalation while breathing normally.

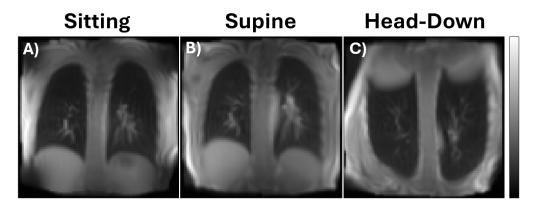


Figure 5.21: Coronal reconstructions of a representative participant at maximum exhalation in the sitting (A), supine (B), and head-down (C) postures, breathing normally.

The absolute change in signal is plotted in Figure 5.22A for each posture, although this result is only anecdotally useful as the signal from each participant's lungs is affected by a multitude of factors. However, we do observe an increase in signal towards the base of the lung for the sitting and supine postures and an increase towards the apex for the head-down posture. Moreover, the signal observed in the supine posture appears to be generally greater than for the other two postures. To provide a more useful comparison between postures, the variation in average signal is normalised using the average signal within the lung. This normalised result along the apical-to-basal direction is plotted in Figure 5.22B. There is a clear trend of increasing relative signal towards the base of the lung for the sitting and supine posture, which is reversed (increasing towards the apex) when in the head-down posture. Additionally, whilst the change in the relative signal is largely continuous throughout the full sitting and head-down lungs, in the supine position this appears to level off beyond approximately 65% of the apical-to-basal distance.

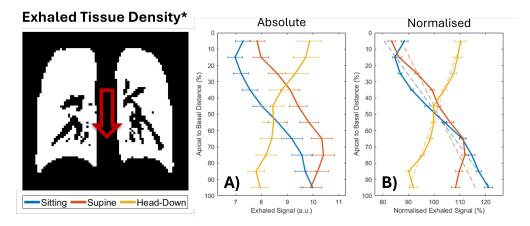


Figure 5.22: Variation in the average Exhaled Signal Intently within ten vertically stacked coronal lung segments, for participants breathing normally. Plotted for both the Absolute Signal Intensity (A) and Normalised Signal Intensity (B). Errors bars represent the standard error on the mean for each lung section. slopes associated with the linear fits plotted on (B) are presented in Table 5.10. The plots are depicted alongside the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Each result in Figure 5.22B is fitted to a linear relationship via linear regression. These fits are plotted on the figure and the resultant slopes along with their associated standard errors are presented in Table 5.7. The greatest magnitude of change occurs within the lung when sitting upright, increasing by approximately 46% of the average from apex to base. This is followed by the supine posture, which sees an increase of approximately 32% from apex to base. Finally, the lowest magnitude change in signal of approximately 23% is observed with the head-down lung, now decreasing towards the base of the lung.

Fitted Linear Slope of Exhaled Signal Intensity		
Posture	Apical to Basal slope (%)	
Sitting	46 ± 4	
Supine	32 ± 5	
Head-Down	-23 ± 2	

Table 5.7: Variation in the exhaled signal intensity along the apical-to-basal axis. slopes associated with the linear fits plotted in Figure 5.22. Shaded cells represent variation within isogravitational planes.

5.3.2.2 Investigations in the Sagittal Plane

As with the coronal experiments, it is useful to first investigate the effect of tissue compression in the exhaled lung. Figure 5.23 presents representative images of a single participant in the sitting, supine, and prone postures at reconstructed maximal exhalation.

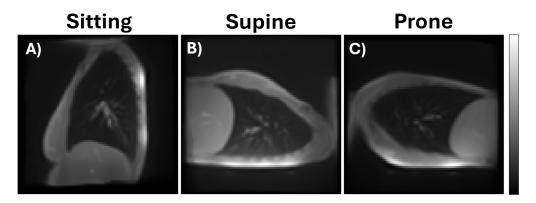


Figure 5.23: Sagittal reconstructions of a representative participant at maximum exhalation in the sitting (A), supine (B), and head-down (C) postures, breathing normally.

Successfully ventilation-proceeded sagittal data sets were used to investigate the variation in exhaled signal along both the apical-to-basal (Figure 5.24A) and anterior-to-posterior (Figure 5.24B) directions. Normalising this variation by the total value of the lung and plotting along the apical-

to-basal direction, yields Figure 5.24A. Note, the absolute variation in tissue density has been omitted for this imaging plane due to the assumed high degree of error from sensitivity gradients. The largest change in the relative exhaled signal intensity was observed in the sitting lung; increasing from apex to base, with the greatest rate of change in the central regions. Similarly, the supine posture exhibits an increase in signal intensity towards the base. Increasing approximately linearly along this axis, although appearing to drop slightly in the most basal zone. Finally, there appears to be little variation in the exhaled signal intensity within the prone posture along this axis, approximately uniform throughout. Each normalised distribution in this plot is fitted to a linear relationship, with the resulting slopes presented in Table 5.8.

Figure 5.24B shows the normalised variation in exhaled signal along the anterior-to-posterior axis. The signal in the supine posture increases towards the back of the lung, with progressively larger changes between zones. A similar pattern is reflected in the prone posture, though directed towards the anterior of the lung and with a slightly reduced magnitude of change. The sitting posture shows a signal magnitude that decreases slightly from the anterior to the centre, before increasing to its largest value in the the posterior. This results in a general trend of increasing signal towards the posterior of the lung, however with a highly reduced magnitude of change compared to the other postures. Each distribution in this plot is fitted to a linear relationship, with the resulting slopes presented in Table 5.8.

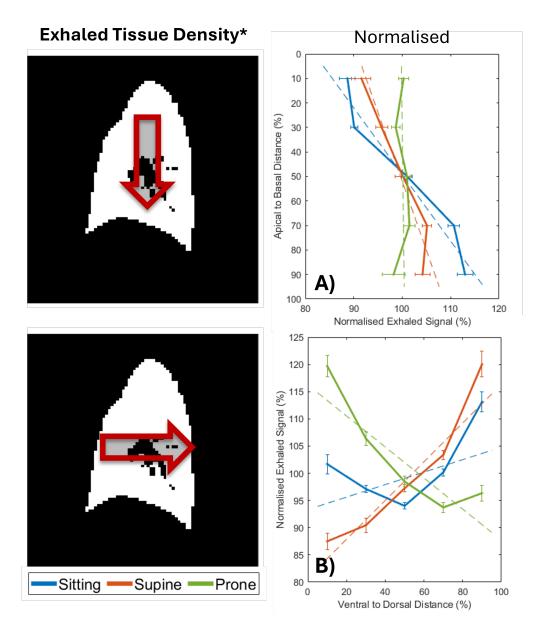


Figure 5.24: Variation in the average Exhaled Signal Intently within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. Plotted for the normalised signal intensity. Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits plotted are presented in Table 5.8. The plots are depicted alongside the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Table 5.8 presents the slopes of the fits plotted in Figures 5.24A and 5.24B. Along the apical-to-basal direction, the lung in the sitting posture demonstrates the greatest slope of increasing in exhaled signal intensity towards

the base. This is followed by the supine posture with a slope approximately half that of the sitting posture. Whereas, the prone lung exhibits an approximately uniform signal intensity throughout and no significant slope. Along the anterior-to-posterior axis, the sitting lung shows a slight increase towards the posterior, although the error on this result is significant. Whereas, the supine and prone lungs exhibit approximately equal but opposite slopes, increasing towards the posterior and anterior, respectively.

Fitted Linear Slope of Exhaled Signal Intensity		
Posture	Apical to Basal slope (%)	Anterior to Posterior slope
		(%)
Supine	18 ± 3	36 ± 7
Sitting	37 ± 6	12 ± 12
Prone	1 ± 3	-29 ± 9

Table 5.8: Variation in the exhaled signal intensity along the apical-to-basal and anterior-to-posterior axis. Slopes associated with the linear fits plotted in Figure 5.24. Shaded cells represent variation within isogravitational planes.

Variability if the apical-to-basal exhaled signal intensity is further visualised in Figure 5.25. First, considering the apical-to-basal variation in the anterior segment in Figure 5.25A. When in the sitting posture, the signal intensity appears to increase towards the base, with the greatest gradient in the central regions. Conversely, this signal appears to be approximately uniform and equivalent throughout this segment for both the supine and prone postures. Although each may decrease slightly towards the base as evident from the drop in signal in these regions. Figure 5.25B shows the apical-to-basal signal variation associated with the posterior of the lung. Again, the sitting posture exhibits an increase towards the base, with the greatest gradient in the central regions. Similarly, this distribution ap-

pears approximately uniform for the prone posture, although now potentially increasing slightly towards the base. However, the supine posture now presents a significant increase in signal intensity towards the base; increasing in all zones save for the most basal, where the value drops slightly.

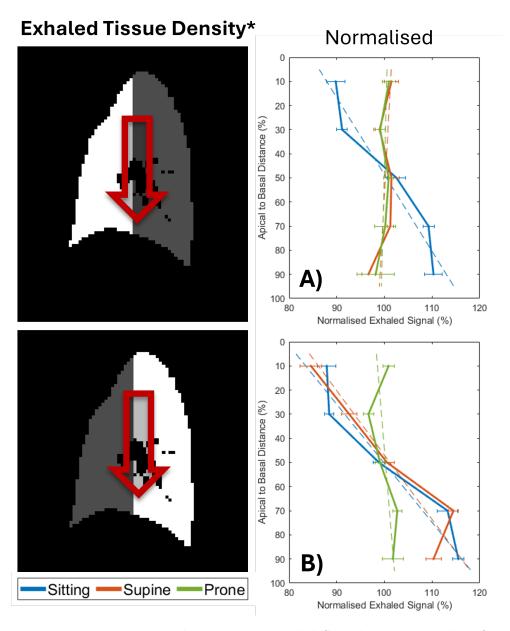


Figure 5.25: Variation in the average Exhaled Signal Intently within five vertically stacked sagittal lung segments at the anterior (A) and posterior (B) of the lung, for participants breathing normally. Plotted for the normalised signal intensity. Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits plotted are presented in Table 5.9. The plots are depicted alongside the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The slopes of the linear fits plotted in Figure 5.25 are presented in Table 5.9. The sitting posture exhibits a significant increase in signal intensity towards the base of the lung in both the anterior and posterior regions, with a slightly greater magnitude in the latter. The prone posture presents a slight increase in signal towards the base in the posterior segment and a slight decrease in the anterior segment. However, these changes are not significant when compared to the error. Finally, in the anterior segment of the supine posture, the lung presents a similar slope to that of the prone posture. However, in the posterior segment, the signal intensity appears to increase significantly towards the base, with a slope comparable to that of the sitting posture in this segment.

Fitted Linear Slope of Exhaled Signal Intensity		
Posture	Lung Segment	Apical to Basal Slope (%)
Supine	Anterior	-3 ± 4
	Posterior	38 ± 9
Sitting	Anterior	32 ± 6
	Posterior	41 ± 7
Prone	Anterior	-4 ± 5
	Posterior	4 ± 4

Table 5.9: Variation in the exhaled signal intensity along the apical-to-basal axis in the anterior and posterior lung segments. Slopes associated with the linear fits plotted in Figure 5.25. Shaded cells represent variation within isogravitational planes.

5.3.3 Variation in Ventilation Dynamics

5.3.3.1 Fractional Ventilation

5.3.3.1.1 Investigations in the Coronal Plane

FV maps were generated for each participant when breathing normally and deeply in each posture. Example maps for a representative participant in each posture while breathing deeply are plotted in Figure 5.26.

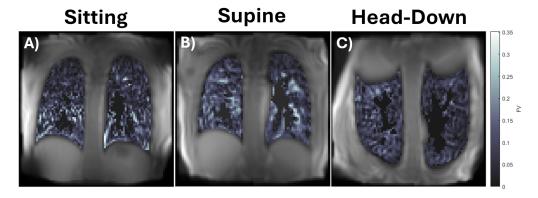


Figure 5.26: Coronal FV maps of a representative participant in the sitting (A), supine (B), and head-down (C) postures, breathing deeply.

The variation in FV for participants breathing normally is presented in Figure 5.27A. For both the sitting and supine postures, the FV appears to increase significantly towards the base of the lung. This change follows a similar gradient for each posture, although the magnitude of the FV appears to be generally greater in the supine position. Additionally, whilst the FV continues to increase for the full length of the sitting lung, this trend appears to plateau and decrease slightly beyond approximately 65% of the apical-to-basal distance within the supine lung. This decrease, combined with the generally larger FV of the supine posture, results in a near intersection of the magnitude of the FV of the sitting lung, in the most basal lung regions. When breathing normally in the head-down posture,

the FV appeared approximately uniform across the apical-to-basal axis.

To compare the relative change in the FV within the lung, this distribution, normalised by the average FV of the lung, is plotted in Figure 5.27B. Intuitively, while the normalised plots are useful for comparing the relative change within postures, they can not be used to compare the relative magnitude between postures. Furthermore, this can be quantified by fitting the dependence of the relative FV on percentage of the apical-to-basal distance, to a linear relationship. The resulting slopes are displayed in Table 5.10. The greatest slope in the apical-to-basal direction is present for the sitting posture. This gradient appears to be greatest in the central regions, levelling off in the most basal zones. The supine posture also exhibits an increase in relative FV, although the magnitude of this change is reduced compared to sitting. Finally, the normalised FV appeared approximately uniform along this axis when in the head-down posture. However, there appeared to be two small maxima in the FV located at approximately 30% and 65% of the apical-to-basal distance.

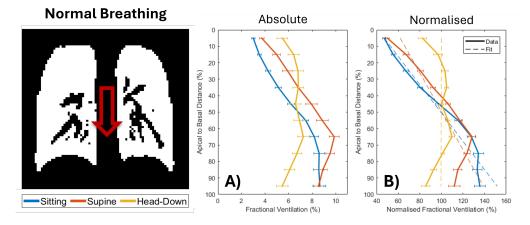


Figure 5.27: Variation in the average FV within ten vertically stacked coronal lung segments, for participants breathing normally. Plotted for both the absolute FV (A) and Normalised FV (B). Errors bars represent the standard error on the mean for each lung section. slopes associated with the linear fits plotted on (B) are presented in Table 5.10. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The distribution of the absolute FV within the lung when breathing deeply is plotted in Figure 5.28A. These plots show a dependence similar to that of normal breathing, although as one might expect, the magnitude of the FV is greater for each posture while breathing deeply. FV consistently increases towards the base of the lung when sitting. When supine, this trend continues, although now with two distinct 'humps' in the distribution of this function. These are centred at approximately 35% and 70% of the apical-to-basal distance, the former of which results in a drop in the FV in the most basal zones. Breathing deeply in the head-down posture, results in an near-consistent increase in the FV towards the apex, only dropping slightly in the most apical zone.

The normalised FV is plotted in Figure 5.28B, with the slopes of the linear fits displayed in Table 5.10. The greatest increase of FV occurred while sitting, again with the greatest gradient in the central zones. This is followed by the supine posture, increasing towards the base but with a reduced linearity in this distribution. The head-down posture exhibits a significant increase in the normalised FV towards the base towards the base.

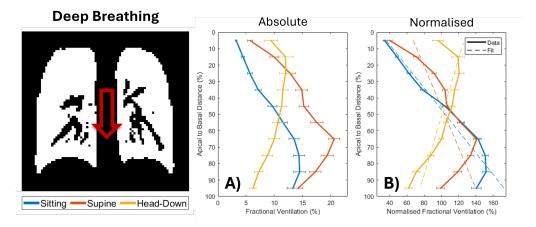


Figure 5.28: Variation in the average FV within ten vertically stacked coronal lung segments, for participants breathing deeply. Plotted for both the absolute FV (A) and normalised FV (B). Errors bars represent the standard error on the mean for each lung section. slopes associated with the linear fits plotted on (B) are presented in Table 5.10. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The slopes of the linear fits plotted in Figures 5.27B and 5.28B are presented in Table 5.10. When breathing normally, the sitting posture appears to experience the greatest relative increase in the FV towards the base of the lung, followed by the supine posture. In the head-down posture, there appears to be no significant variation in the distribution of FV.

When breathing deeply, the apical-to-basal slope of the FV associated with the sitting posture increased compared to breathing normally. In the supine posture, the slope remained approximately equivalent between the breathing regimes. However, in the head-down posture the approximately uniform distribution of FV increased significantly in the apical zones when breathing deeply. This resulted in a significant negative apical-to-basal slope. The sitting posture maintains the greatest magnitude of relative change, followed by supine and then head-down.

Fitted Linear Slope of Normalised Fractional Ventilation		
Posture	Breathing	Apical to Basal slope (%)
Sitting	Normal	116 ± 8
	Deep	152 ± 13
Supine	Normal	86 ± 14
	Deep	80 ± 23
Head-Down	Normal	0 ± 10
	Deep	-57 ± 15

Table 5.10: Variation in the normalised FV along the apical-to-basal axis. slopes associated with the linear fits plotted in Figures 5.27B and 5.28B. Shaded cells represent variation within isogravitational planes.

Figure 5.29 is presented to compare the change in the distribution of FV, for participants breathing normally and breathing deeply. These plots contain the same data as Figures 5.27 and 5.28. However, each posture is now plotted separately within its own row (A,B,C) for their respective absolute and normalised ventilation (i,ii). As one might expect, the absolute FV increases for almost all zones when the participants are breathing deeply. Although, this change is not distributed evenly for each posture. In the absolute FV plots, this increase becomes significantly more pronounced towards the base when sitting. The greatest increase occurs at the base of the lung, reducing to near zero in the most apical zone. When supine, this change is more uniform throughout the lung, although also increases slightly towards the base. Furthermore, there appear to be two distinct 'humps' of a greater increase in the FV, centred a approximately 35\% and 70% of the apical-to-basal distance. Finally for the head-down posture, the change in absolute FV increases towards the apex of the lung, starting at a near zero change at the base. To compare the relative distribution of FV, the normalised FV is presented for each posture. When sitting, there is a redistribution of relative FV away from the apex and to the base. In the head down posture, there is a redistribution approximately opposite to that of the sitting posture, from the base to the apex. Finally, when supine, the relative FV is very similar when normally and deeply. Indeed, the only variation appears to be due to the aforementioned 'humps' in the FV distribution when breathing deeply.

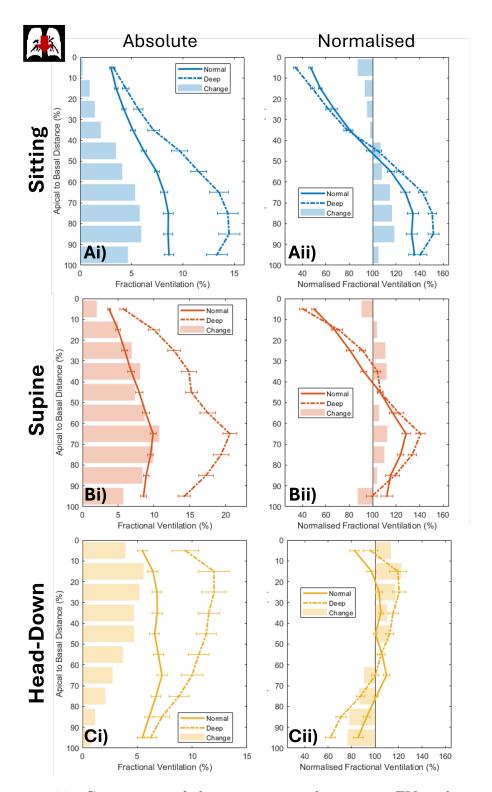


Figure 5.29: Comparison of the variation in the average FV within ten vertically stacked coronal lung segments, between participants breathing normally and deeply. The absolute FV (i) and normalised FV (ii) is plotted for each of the three postures: Sitting (A), Supine (B) and Head-Down (C). The bars represent the change in value from breathing normally to breathing deeply. Error bars represent the standard error on the mean for each lung section.

5.3.3.1.2 Investigations in the Sagittal Plane

Example sagittal FV maps of a representative participant in each posture are plotted in Figure 5.30.

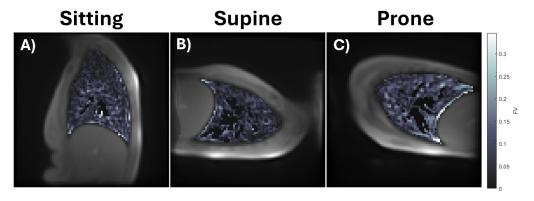


Figure 5.30: Sagittal FV maps of a representative participant in the sitting (A), supine (B), and prone (C) postures, breathing normally.

All sagittal acquisitions that completed successful ventilation processing were used to investigate the variation in FV along both the apical-to-basal (Figure 5.31A) and anterior-to-posterior (Figure 5.31B) directions. Captured while the participants were breathing normally, the variation in absolute FV along the apical-to-basal direction is plotted in Figure 5.31Ai. Each posture sees a near-linear increase in FV towards the base of the lung. This change is greatest while sitting, followed by the supine and then prone postures. The lung in the supine posture appears to have a greater FV throughout the lung compared to the other postures. Figure 5.31Aii is demonstrates this relationship when the FV is normalised by the FV of the lung as a whole. Here, the greatest relative increase in the FV along the apical-to-basal direction occurs when the participants are sitting. The lung in the supine and prone postures exhibits a reduced but near nearequivalent increase in FV towards the base. Each relationship within this plot is fitted to a linear relationship and the respective parameters and their errors are presented in Table 5.11.

Changes in the absolute FV along the anterior-to-posterior direction is plotted in Figure 5.31Bi. This parameter increases towards the posterior of the lung for both the sitting and supine postures, although dropping slightly in the most posterior zone of the latter. Furthermore, the FV is greater within all zones of the supine lung compared to sitting. The prone lung sees the opposite trend, increasing towards the anterior of the lung with similar magnitudes to that of the sitting posture. Although, similar to the supine lung, the prone lung exhibits a drop in the FV within the most anterior region. Figure 5.31Bii depicts these relationships when normalised by the average FV of the lung as a whole. The variations in FV within the lung in the sitting and supine posture exhibit similar relationships to each other, each increasing towards the anterior regions. Although this change appears greatest for the supine posture. This trend is again reversed for the prone lung, increasing towards the anterior of the lung, although decreasing slightly in the most anterior zone.

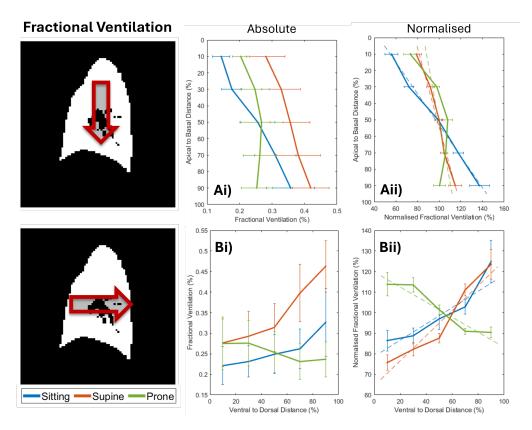


Figure 5.31: The average FV within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. The absolute FV (i) and Normalised FV (ii). Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits are presented in Table 5.11. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The slopes of the fits presented in Figures 5.31Aii and 5.31Bii are presented in Table 5.11. The greatest relative increase in the apical-to-basal FV is observed in the sitting lung, followed by prone and then followed closely by supine. The greatest relative increase in normalised-FV along the anterior-to-posterior direction is observed in the supine lung, followed by sitting. This is reversed in the prone lung, although with a reduced slope magnitude compared to the other postures.

Fitted Linear Slope of Fractional Ventilation			
Posture	Apical to Basal Slope (%)	Anterior to Posterior Slope	
		(%)	
Supine	42 ± 4	61 ± 12	
Sitting	105 ± 6	40 ± 10	
Prone	28 ± 17	-37 ± 8	

Table 5.11: Variation in the normalised FV along the apical-to-basal and anterior-to-posterior axis. Slopes associated with the linear fits plotted in Figure 5.31. Shaded cells represent variation within isogravitational planes.

Similar to the analysis of the exhaled signal intensity, the apical-to-basal variation in FV was investigated in the anterior and posterior segments independently. This is plotted in Figure 5.32 with each distribution normalised by the average FV of the respective segment. In the anterior segment (Figure 5.32A) of the sitting posture, the normalised FV increases towards the base, with the greatest gradient in the central regions. Both the supine and prone postures exhibit a similar relative distribution in the anterior segment, remaining approximately equal throughout, though dropping slightly in the most anterior segment. In the posterior segment (Figure 5.32B), the sitting posture shows a near-linear increase in the relative FV, deviation from the anterior segment (Figure 5.32A) where a reduced gradient was observed the basal regions. The prone posture presents a slight increase in FV towards the base in the apical half of the lung, although this plateaus beyond the central region. Deviating from the trend in the anterior, the posterior segment of the supine lung exhibits a near-linear increase in the FV, only decreasing slightly in the most basal zone.

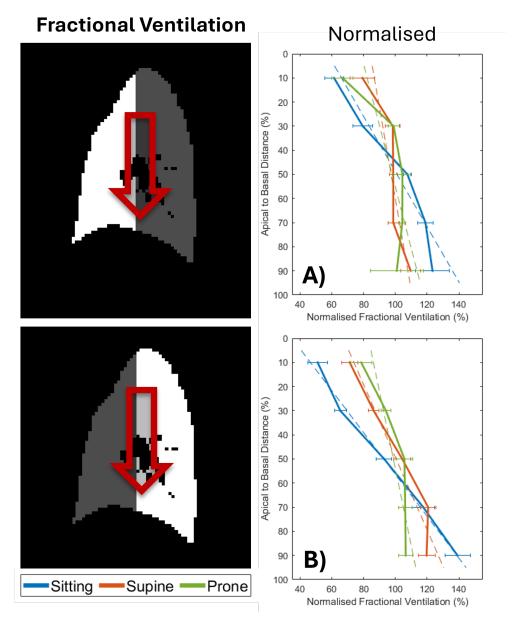


Figure 5.32: Variation in the average FV within five vertically stacked sagittal lung segments at the anterior (A) and posterior (B) of the lung, for participants breathing normally. Plotted for the normalised signal intensity. Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits plotted are presented in Table 5.12. The plots are depicted alongside the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The slopes of the linear fits depicted in Figure 5.32 are presented in Table 5.12. All postures exhibit an increase in the relative FV towards the base of the lung in both the anterior and posterior segments. This increase is greatest in the sitting posture, with the posterior segment presenting the

largest slope. The slope in the anterior segment of the prone lung appears to be slightly greater than that of the posterior segment, although the difference is not significant compared to the large error associated with each value. The FV in the anterior segment of the supine lung increases the least of all the segments investigated. However, the posterior segment presents a significantly greater slope.

Fitted Linear Slope of Fractional Ventilation			
Posture	Lung Segment	Apical to Basal Slope (%)	
Supine	Anterior	26 ± 11	
	Posterior	66 ± 10	
Sitting	Anterior	87 ± 15	
	Posterior	115 ± 7	
Prone	Anterior	39 ± 20	
	Posterior	31 ± 10	

Table 5.12: Variation in the normalised FV along the apical-to-basal axis in the anterior and posterior lung segments. Slopes associated with the linear fits plotted in Figure 5.32. Shaded cells represent variation within isogravitational planes.

5.3.3.2 Ventilation Time-To-Peak (VTTP)

5.3.3.2.1 Investigations in the Coronal Plane

Coronal VTTP maps were generated for each participant while breathing normally and deeply in each posture. Example maps of a representative participant in each posture while breathing deeply are plotted in Figure 5.33.

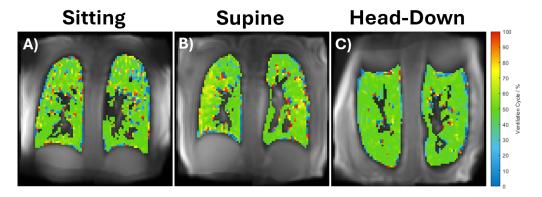


Figure 5.33: Coronal VTTP maps of a representative participant in the sitting (A), supine (B), and head-down (C) postures, breathing deeply.

The variation in the VTTP for each posture while breathing normally is plotted in Figure 5.34A. Both the supine and head-down posture appear similar, near to 50% of the cycle, which is reflective of the point of maximal lung area. Although the VTTP appears slightly earlier in the middle to basal regions for the head-down posture. The VTTP appears to vary significantly for the central regions of the lungs when in the sitting posture, with peak ventilation occurring far later for the central regions. Although at the apex and the base this value is reduced back to values near 50%. Additionally, it appears that all postures see an earlier peak ventilation in the apical regions during normal breathing.

The distribution of the VTTP when deep breathing is plotted in Figure 5.34. All postures appear to reach maximum ventilation near the point of maximum lung volume in the respiratory cycle. When sitting, maximum ventilation is reached constantly later than the point of maximum lung area, although earliest at the base. When supine, maximum ventilation occurs largely at 50% of the cycle, although slightly earlier at the apex. When in the head-down posture, maximum ventilation occurs largely close to the 50% point, although slightly earlier in the middle to basal regions.

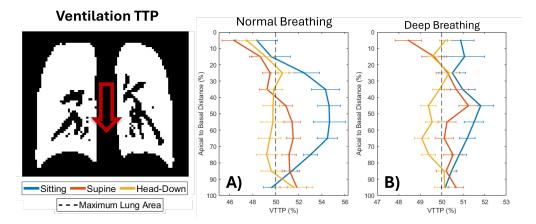


Figure 5.34: Variation in the average VTTP within ten vertically stacked coronal lung segments, for participants breathing normally (A) and deeply (B). Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Figure 5.35 contrasts the change in the distribution of VTTP values between participants breathing normally and deeply. These results are the same as those presented in Figure 5.34; however, with each posture now plotted separately contrasting the normal and deep breathing distributions. For all postures, the peak ventilation is largely reached nearer to the point of maximum lung area. This is most apparent in the sitting posture. When in the sitting and head-down postures the earlier ventilation at the apex observed when sitting is no longer apparent; however, it is still present, albeit with a reduced magnitude when supine.

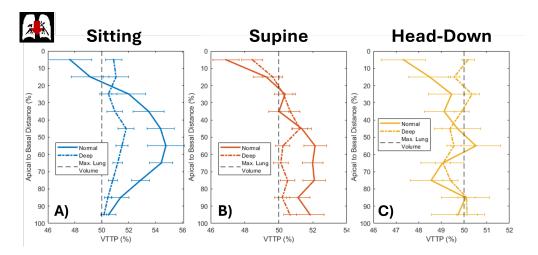


Figure 5.35: Comparison of the variation in the average VTTP within ten vertically stacked coronal lung segments, between participants breathing normally and deeply. Plotted for each of the three postures: Sitting (A), Supine (B) and Head-Down (C). Error bars represent the standard error on the mean for each lung section.

5.3.3.2.2 Investigations in the Sagittal Plane

Sagittal VTTP maps were generated for each participant while breathing normally in each posture. Example maps of a representative participant in each posture are plotted in Figure 5.36.

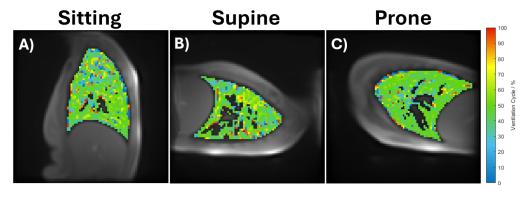


Figure 5.36: Sagittal VTTP maps of a representative participant in the sitting (A), supine (B), and prone (C) postures, breathing normally.

Along the apical-to-basal axis (Figure 5.37A), the VTTP appears relatively uniform for each posture, although slightly later in the centre when supine. Generally, the lung appears to reach maximum ventilation earliest when

prone, followed by supine and then when sitting. The anterior-to-posterior axis shows more variation within postures. The posterior regions of the sitting lung appear to reach maximum ventilation later, whereas the opposite trend occurs when supine. There appears to be little spatial dependence along this axis for the prone lung. Once again, the lung appears to reach maximum ventilation earliest when prone, followed by supine and finally sitting.

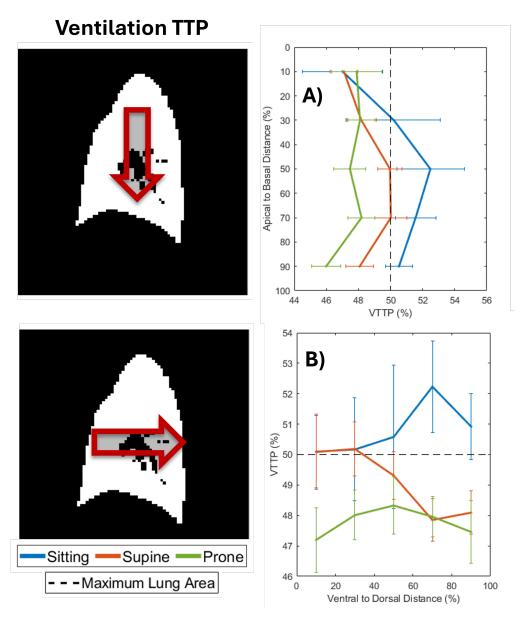


Figure 5.37: Variation in the average VTTP within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

5.3.3.3 Jacobian Determinant

5.3.3.3.1 Investigations in the Coronal Plane

Coronal maps of the Jacobian Determinant were generated for each participant while breathing normally and deeply in each posture. Example maps of a representative participant in each posture while breathing deeply are plotted in Figure 5.38.

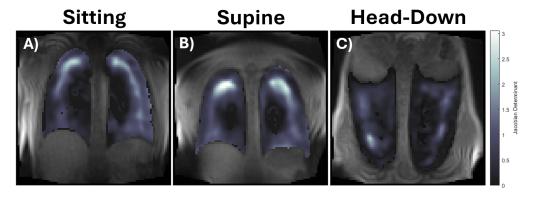


Figure 5.38: Coronal Maps of the Jacobian Determinant of a representative participant in the sitting (A), supine (B), and head-down (C) postures, breathing deeply.

For normal breathing within the full lungs, this distribution is plotted in Figure 5.39A. As one would expect, all zones within all postures (for both breathing regimes) exhibited a positive mean Jacobian determinant, representative of an expansion of the lung parenchyma due to the respiratory cycle. Plotted for normal breathing in Figure 5.39A, all postures present a largely uniform magnitude of expansion within the apical and middle lung sections, although this increases significantly in the most basal regions beyond approximately 65% of the apical to basal distance. This distribution is very similar for the sitting and supine postures, but slightly greater in magnitude when in the head-down posture. Plotting this distribution for deep breathing (Figure 5.39B), each posture exhibits its own distinct trend:

the lung in the sitting posture sees a largely uniform expansion. The lung in the head-down posture sees a steady increase towards the base. Finally the supine lung exhibits a large expansion at the apex, greater than both the other postures. This reduces in the central regions, before increasing again at the base.

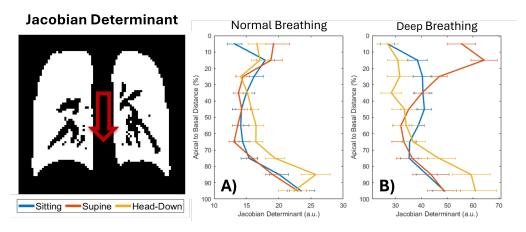


Figure 5.39: Variation in the average Jacobian Determinant within ten vertically stacked coronal lung segments, for participants breathing normally (A) and deeply (B). Error bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Evident in Figure 5.38, when the participant is breathing deeply, the lung sees a reduction in the Jacobian determinant in the central regions. Consequently, it is useful to consider the change in the Jacobian determinant within only the peripheries of the lung. This is plotted in Figure 5.40, for a periphery mask width of seven voxels. When breathing normally (Figure 5.40A), the distribution of the peripheral Jacobian determinant follows a similar trend as for the full lung (Figure 5.39A). Largely constant through the lung, although increasing in the basal region. Once again, this increase is greatest whilst in the head-down posture. Plotting for participant's deep breathing in Figure 5.40B, each posture shows a largely linear change in the distribution of the Jacobian determinant along the apical-

to-basal direction. Increasing towards the base when sitting and, with a greater magnitude of change, when in the head-down posture. Conversely, when supine, an increase is observed towards the apex of the lung, with a similar magnitude of change to sitting.

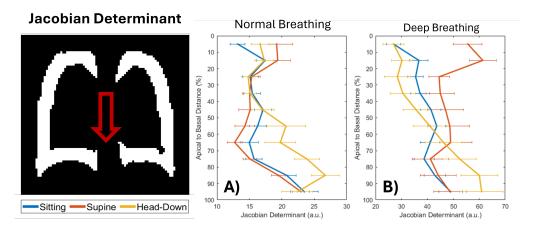


Figure 5.40: Variation in the average Jacobian Determinant within ten vertically stacked coronal lung peripheral segments, for participants breathing normally (A) and deeply (B). Error bars represent the standard error on the mean for each lung section. Depicted alongside the plots is the peripheral mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Figure 5.41 contrasts the difference in the peripheral Jacobian determinant when participants are breathing normally and deeply. Each posture is plotted separately (A, B, C), for the absolute Jacobian determinant (i, containing the same data as Figure 5.40) and the normalised Jacobian determinant (ii). All postures show an increase in the absolute Jacobian determinant when breathing deeply compared to normally. This increase appears largely even throughout the lung when sitting, slightly decreasing towards the base when supine and greatly increasing towards the base when in the head-down posture. Plotting the normalised results, one observes a redistribution in the relative Jacobian determinant from the base of the lung to the middle when sitting and supine, with largely the same distribution for both breathing regimes when in the head-down posture.

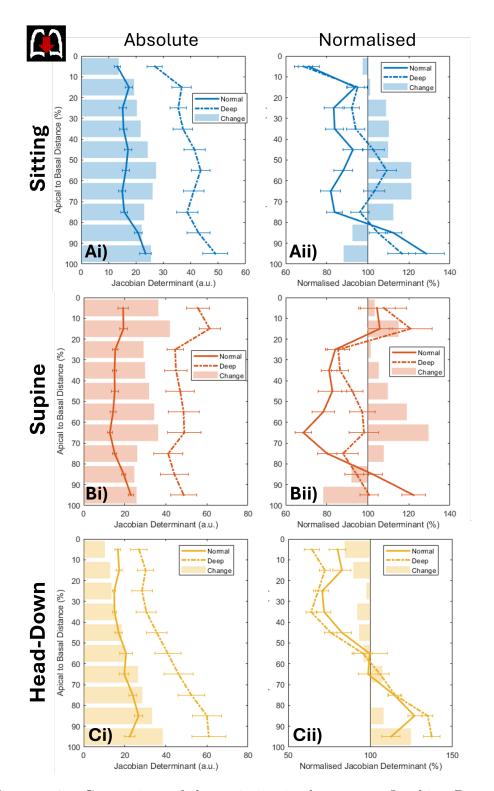


Figure 5.41: Comparison of the variation in the average Jacobian Determinant within ten vertically stacked coronal peripheral lung segments, between participants breathing normally and deeply. The absolute Jacobian Determinant (i) and normalised Jacobian Determinant (ii) is plotted for each of the three postures: Sitting (A), Supine (B) and Head-Down (C). The bars represent the change in value from breathing normally to breathing deeply. Error bars represent the standard error on the mean for each lung periphery section.

5.3.3.3.2 Investigations in the Sagittal Plane

Sagittal maps of the Jacobian Determinant were generated for each participant while breathing normally in each posture. Example maps of a representative participant in each posture are plotted in Figure 5.42.

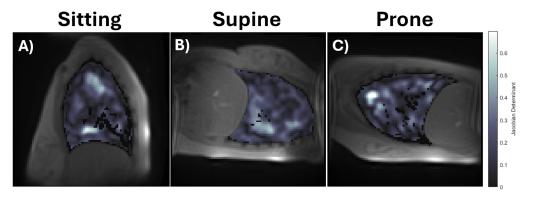


Figure 5.42: Sagittal maps of the Jacobian Determinant of a representative participant in the sitting (A), supine (B), and prone (C) postures, breathing normally.

The absolute value of the Jacobian Determinant along the apical-to-basal axis is plotted in Figure 5.43A. As expected, all postures present a positive Jacobian across all lung regions, representing an expansion due to inspiration. The magnitude of this expansion appears to increase towards the base when in the sitting posture, although with a slight reduction in the central zone. The supine lung also presents an increase in expansion towards the basal regions, however the most apical zone breaks this trend and exhibits the greatest expansion. Appears to expand more uniformly than the other two postures, with a reduced magnitude of change throughout. Although again, the central regions appear to expand less than those at the edges.

Investigating the variation along the anterior-to-posterior axis, the variation in the absolute Jacobian Determinant is plotted in Figure 5.43B. Again, all zones yield a positive Jacobian determinant, thus an expansion due to inhalation. All postures appear to expand to a greater degree at the

posterior, compared to the anterior. This is represented by a positive gradient of the Jacobian determinant along this axis. Furthermore, all postures appear to expand equally in the most anterior zone. The greatest increase in expansion towards the posterior is observed when supine, followed by sitting and then prone.

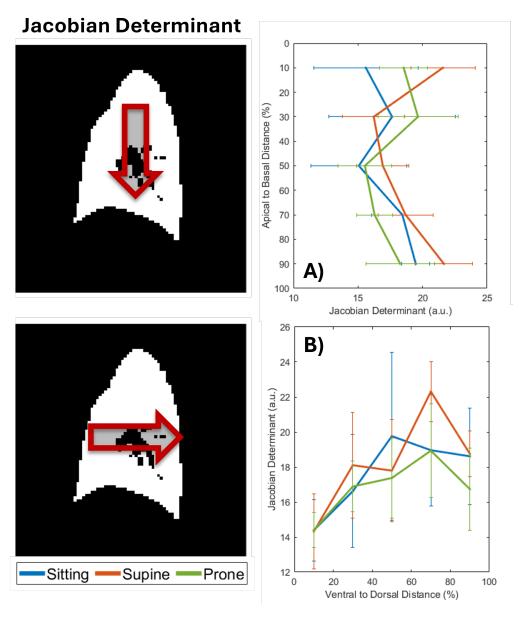


Figure 5.43: Variation of the average Jacobian Determinant within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. Error bars represent the standard error on the mean for each lung section. Depicted alongside the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

5.3.4 Variation in Perfusion Dynamics

5.3.4.1 Normalised Perfusion

5.3.4.1.1 Investigations in the Coronal Plane

Coronal Perfusion Weighted maps were generated for each participant while breathing normally and deeply in each posture. Example maps of a representative participant in each posture while breathing normally are plotted in Figure 5.44.

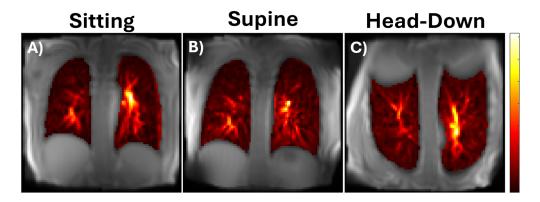


Figure 5.44: Coronal Perfusion Weighted maps of a representative participant in the sitting (A), supine (B), and head-down (C) postures, breathing normally.

For participants breathing normally the perfusion variation is plotted in Figure 5.45. As the perfusion weighted maps are based simply on changing voxel intensity with respect to the cardiac cycle, they must first be normalised before the distributions can be averaged across all participants. Initially, this was done based on the mean value of the lung, denoted here by the full 'volume'. This is plotted in Figure 5.45A. In the sitting posture, the normalised perfusion increases along the apical-to-basal direction, although reduces in the most basal regions. In the supine position, this trend repeats, although with a lower magnitude of change. Finally, in the head-

down posture, the trend is reversed, albeit with a lower magnitude. The perfusion is lowest at the base, increasing towards the centre of the lung from which it remains approximately level, before reducing slightly at the apex.

To remove the contribution of tissue compression to the observed perfusion, the perfusion weighting of each voxel is normalised by its respective tissue density. This tissue density is estimated, based on the average signal intensity of the voxel throughout the reconstructed ventilation cycle. Each voxel within the generated tissue-density-normalised lung is then divided by the average value of the full lung, before analysis with multiple participants. The variation in the tissue-density-normalised lung along the apical-to-basal direction is plotted in Figure 5.45B, for participants breathing normally. Each posture sees a shift in intensity towards the central lung regions. When sitting, the lowest perfusion occurs at the apex, increasing quickly to the centre, before dropping slightly at the base. When supine, the lowest perfusion is observed at the apex, increasing quickly before largely plateauing across the central regions and dropping slightly in the most basal region. Finally, when in the head down posture, the lowest perfusion occurs at the base of the lung, increasing quickly before plateauing across the central regions and dropping again at the apex.

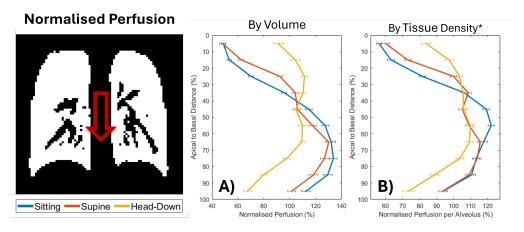


Figure 5.45: Variation in the average Normalised Perfusion within ten vertically stacked coronal lung segments, plotted with normalisation based on volume (A) and tissue density (B). Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results. The * represents that this calculation assumes that the average signal intensity of each voxel within the reconstructed ventilation cycle is representative of the tissue density.

Even with the removal of the central lung vessels, there appears to be in Figure 5.45 a central bias to the perfusion distribution. Therefore, it is useful to consider the distribution of perfusion within only the peripheries of the lung. This is plotted in Figure 5.46, for a periphery mask of width seven voxels and participants breathing normally. The volume-normalised perfusion distribution is plotted in Figure 5.46A. For the lung in both the sitting and supine postures, the perfusion appears to near-linearly increase from the apex to the base of the lung, only dropping slightly in the most basal region. This change is smaller in magnitude for the supine lung compared to sitting. In the head-down lung, this trend is reversed: increasing from apex to base and dropping slightly in the most apical zone. These distributions are fitted to a linear relationship via linear regression, with the resultant slopes presented in Table 5.13.

Normalising the lung peripheries by the respective tissue density of each voxel yields Figure 5.46B. The perfusion in both the sitting and supine

lungs increases towards the base of the lung. Each increasing from the lowest point at the apex; however, levelling out beyond 35% and 65% of the supine and sitting lung respectively. Both postures again see a drop in the most basal region. Although the trend within the head-down lung sees the opposite magnitude of change, increasing from base to apex, the trend also has a different shape. The lowest value is found in the most basal region, then increases slightly before levelling out from the next zone. This continues until approximately 45% of the apical-to-basal distance, before increasing again towards the most apical regions. Finally, the tissue normalised perfusion drops slightly in the most apical region. Each of these trends is fitted to a linear relationship, with the associated slopes presented in Table 5.13.

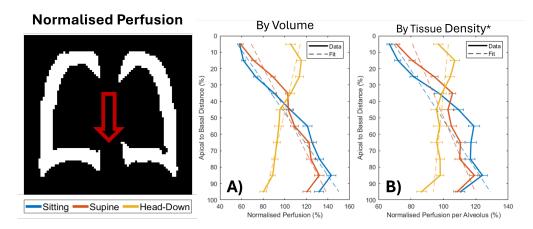


Figure 5.46: Variation in the average Normalised Perfusion within ten vertically stacked coronal lung peripheral segments, plotted with normalisation based on volume (A) and tissue density* (B). Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the peripheral mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results. The * represents that this calculation assumes that the average signal intensity of each voxel within the reconstructed ventilation cycle is representative of the tissue density.

The slopes of the linear fits, plotted in Figure 5.46, are presented in Table 5.13. When normalising by only the mean value of the lung (of the 'volume'), the greatest change along the apical-to-basal axis is observed

when the participants were in the sitting posture, followed by when supine. This change is reversed when in the head-down posture, although to a lower magnitude compared to the other postures. When normalising by the tissue density, the magnitude of the change along the apical-to-basal axis reduced by approximately 50%.

Fitted Linear Slope of Normalised Perfusion			
Posture	Normalisation	Apical to Basal slope (%)	
Sitting	Volume	105 ± 8	
	Tissue Density	66 ± 9	
Supine	Volume	77 ± 11	
	Tissue Density	43 ± 9	
Head-Down	Volume	-35 ± 5	
	Alveolus	-10 ± 5	

Table 5.13: Variation in the peripheral normalised perfusion along the apical-to-basal axis. slopes associated with the linear fits plotted in Figure 5.46. Shaded cells represent variation within isogravitational planes.

All previous perfusion plots have been derived from acquisitions captured while the participants were breathing normally. Figure 5.47 compares the distribution of perfusion within the full lung for participants breathing normally and deeply. Each posture is plotted separately (A, B, C), for the volume-normalised perfusion (i) and the tissue-density-normalised perfusion (ii). The normal-breathing data within these plots are the same as presented in Figure 5.45. All postures show a very similar distribution in perfusion between breathing regimens, for both methods of normalisation. However, all postures show a slight redistribution of the perfusion towards the apex, reducing in magnitude in the more basal regions.

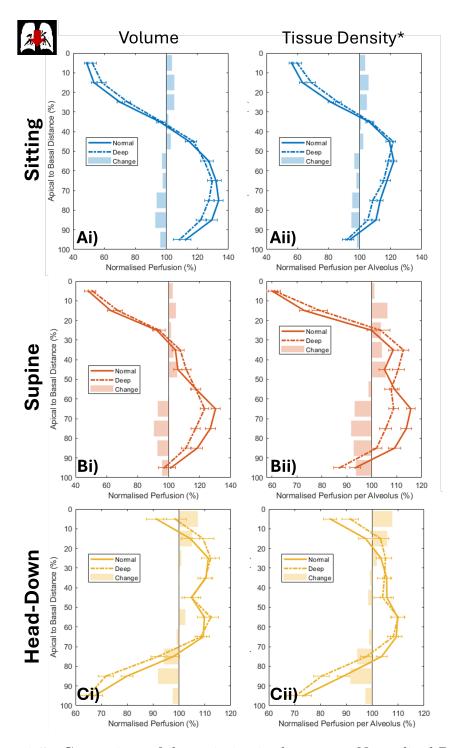


Figure 5.47: Comparison of the variation in the average Normalised Perfusion within ten vertically stacked coronal lung segments, between participants breathing normally and deeply. Normalisation based on the volume (i) and tissue density* (ii) is plotted for each of the three postures: Sitting (A), Supine (B) and Head-Down (C). The bars represent the change in value from breathing normally to breathing deeply. Error bars represent the standard error on the mean for each lung section. The * represents that this calculation assumes that the average signal intensity of each voxel within the reconstructed ventilation cycle is representative of the tissue density.

5.3.4.1.2 Investigations in the Sagittal Plane

Sagittal Perfusion Weighted maps were generated for each participant while breathing normally in each posture. Example maps of a representative participant in each posture are plotted in Figure 5.48.

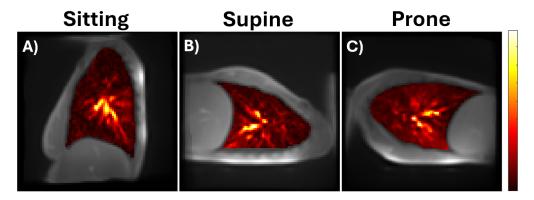


Figure 5.48: Sagittal Perfusion Weighted maps of a representative participant in the sitting (A), supine (B), and prone (C) postures, breathing normally.

The variation in the volume-normalised perfusion along the apical-to-basal axis is plotted in Figure 5.49Ai. For the sitting and supine postures, the volume-normalised perfusion increases towards the base of the lung, although each appears to plateau past the middle lung region. Within the prone lung, the perfusion distribution presents as a central maximum, surrounded by a symmetrically reducing apical and basal distribution. The variation in the tissue-normalised perfusion along the apical-to-basal axis is plotted in Figure 5.49Aii. All postures see a similar perfusion distribution along this axis compared to the previous plot. Although, the sitting lung appears to present a slightly more central bias to the perfusion.

The variation in the volume-normalised perfusion along the anterior-toposterior axis is plotted in Figure 5.49Bi. In the supine lung, the relative perfusion appears to increase towards the posterior. This pattern is reversed with a greater magnitude in the prone lung: increasing towards the anterior. When sitting, there appears to be a lower perfusion in the central zone, increasing towards both the anterior and, to a greater degree, the posterior. Normalising these results by the tissue density and plotting over the anterior-to-posterior axis yields Figure 5.49Bii. The perfusion distribution in the prone lung continues to increase towards the anterior, although to a lower magnitude. In the supine lung, the trend towards the posterior appears to have been largely reduced, no longer presenting any significant change across the lung. The distribution within the sitting lung appears similar to that of the previous plot, although with a reduced magnitude of variation. All distributions in Figure 5.49 are fitted to a linear relationship via linear regression, with the resulting slopes presented in Table 5.14.

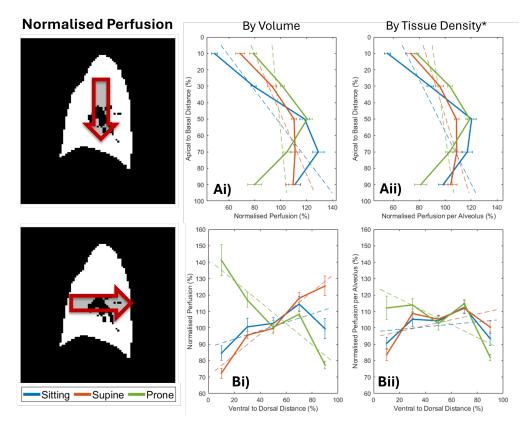


Figure 5.49: Variation in the average Normalised Perfusion within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. Plotted for normalisation based on volume (i) and tissue density* (ii). Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits are presented in Table 5.14. Depicted alongside the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results. The * represents that this calculation assumes that the average signal intensity of each voxel within the reconstructed ventilation cycle is representative of the tissue density.

The slopes of the linear fits along the apical-to-basal and anterior-to-posterior directions in Figure 5.49 are presented in Table 5.14. In all postures, there is an increase in the normalised perfusion towards the base of the lung, with this change greatest when sitting, followed by supine and then prone. Normalising by the tissue density reduces the change in both the sitting and supine lungs, although it increases slightly in the prone lung. Along the anterior-to-posterior axis, the normalised perfusion sees no significant change within the sitting lung, an increase towards the posterior in the

supine lung and a greater but opposite increase towards the anterior in the prone lung. The magnitude of each of these changes is reduced when the perfusion is normalised by the tissue density. This yields no significant change along this axis in the sitting and supine lungs, although the slope of the perfusion towards the anterior of the lung remains when prone.

Fitted Linear Slope of Normalised Perfusion				
Posture	Normalisation	Apical to Basal	Anterior to Poste-	
		Slope (%)	rior Slope (%)	
Supine	Volume	69 ± 18	37 ± 19	
	Tissue Density	51 ± 15	3 ± 26	
Sitting	Volume	98 ± 16	5 ± 39	
	Tissue Density	63 ± 19	1 ± 10	
Prone	Volume	29 ± 20	-68 ± 19	
	Tissue Density	35 ± 18	-43 ± 11	

Table 5.14: Fitted linear slope of the normalised perfusion along the apical-to-basal and anterior-to-posterior axis. Slopes associated with the linear fits plotted in Figure 5.49. Shaded cells represent variation within isogravitational planes.

Figure 5.50 demonstrates the apical-to-basal distribution of the normalised perfusion in the anterior and posterior segments independently. Each relationship plotted below is normalised by the average value of the respective segment. The normalised perfusion distribution within the anterior segment is shown in Figure 5.50Ai. When in the supine position, the volume-normalised-perfusion appears approximately uniform along the apical-to-basal axis of the lung, although it exhibits a slight increase towards the base. In the sitting position, there is a large increase in the volume-normalised-perfusion towards the base in the apical regions; however, this plateaus beyond the centre of the lung. The prone lung presents maximal perfusion

in the central zone, with diminishing values towards both the apex and, to a greater degree, the base. When these results are normalised by tissue density (Figure 5.50Ai), the anterior distribution of perfusion in the supine posture remains largely uniform throughout the segment. The sitting posture continues to exhibit increasing perfusion toward the base in the apical zones; however, rather than plateauing beyond the centre, the perfusion in the basal zones decreases slightly. The prone posture follows a similar distribution to the previous plot, with maximal values in the centre and decreasing toward both the apex and base, though the magnitude of this variation is reduced.

The relative distribution of perfusion in the posterior segment is plotted in Figure 5.50Bi. In this segment of the prone lung, the distribution of perfusion appears largely uniform along the apical-to-basal axis, though slightly elevated in the central zone. Both the sitting and supine lungs exhibit a similar distribution of perfusion in this segment, increasing towards the base before levelling off and dropping slightly in the most basal zone. When these results are normalised by tissue density (Figure 5.50Bii), the distribution within the prone posture remains similar to the previous plot: largely uniform but slightly greater in the centre. The sitting posture demonstrates a similar pattern, though with reduced magnitude of variation. In this segment of the supine lung, the normalised perfusion now increases in all zones towards the base, though the magnitude of change across the lung is reduced compared to the previous plot. All distributions in Figure 5.50 are fitted to a linear relationship, with slopes presented in Table 5.15.

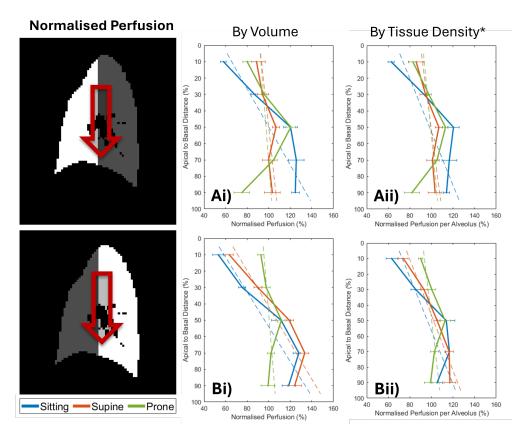


Figure 5.50: Variation in the average normalised perfusion within five vertically stacked sagittal lung segments at the anterior (A) and posterior (B) of the lung, for participants breathing normally. Normalised by the value of the whole lung (i), or by the tissue density of each voxel (ii). Errors bars represent the standard error on the mean for each lung section. Slopes associated with the linear fits plotted are presented in Table 5.15. The plots are depicted alongside the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

The linear fits displayed in Figure 5.50 are presented in Table 5.15. All postures exhibit an increase in perfusion towards the base within both lung segments. This increase is approximately equivalent between the segments when in the sitting posture. This equivalence holds for both normalisation methods, although the slope is reduced when the perfusion is normalised by tissue density. Both segments of the prone lung exhibit a comparatively small but equal increase towards the base of the lung, although this change is not significant compared to the error. Furthermore, there is no significant change in these slopes when normalised by tissue density. The perfusion

distributions in the supine posture share similarities with each of the other postures. In the anterior segment, the distribution resembles that of the prone lung: a small increase in relative perfusion towards the base, which remains unchanged when normalised by tissue density. However, in the posterior segment, the perfusion distribution in the supine lung resembles that of the sitting lung: a significant increase in perfusion towards the base of the lung, which is reduced when normalised by tissue density.

Fitted Linear Slope of Normalised Perfusion				
Posture	Lung Segment	Normalisation	Apical to Basal Slope	
			(%)	
Supine	Anterior	Volume	17 ± 11	
		Tissue Density	20 ± 12	
	Posterior	Volume	90 ± 22	
		Tissue Density	56 ± 10	
Sitting	Anterior	Volume	86 ± 19	
		Tissue Density	62 ± 22	
	Posterior	Volume	94 ± 21	
		Tissue Density	58 ± 22	
Prone	Anterior	Volume	12 ± 32	
		Tissue Density	14 ± 24	
	Posterior	Volume	12 ± 11	
		Tissue Density	17 ± 12	

Table 5.15: Variation in the normalised perfusion along the apical-to-basal axis in the anterior and posterior lung segments. Slopes associated with the linear fits plotted in Figure 5.50. Shaded cells represent variation within isogravitational planes.

5.3.4.2 Perfusion Time-To-Peak (QTTP)

5.3.4.2.1 Investigations in the Coronal Plane

Coronal QTTP maps were generated for each participant while breathing normally and deeply in each posture. Example maps of a representative participant in each posture while breathing deeply are plotted in Figure 5.51.

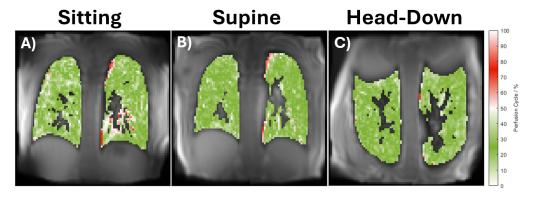


Figure 5.51: Coronal QTTP maps of a representative participant in the sitting (A), supine (B), and head-down (C) postures, breathing normally.

The regional variation in QTTP is plotted for both normal and deep breathing in Figure 5.52. When breathing normally (Figure 5.52A), both the supine and head-down lungs appear to have approximately equal QTTP distributions, which were largely even throughout and slightly later than 25% of the cardiac cycle, which represented the time of maximum aorta signal. The postures deviate slightly in the most apical regions, trending slightly earlier for the head-down lung and later when supine. The sitting lung sees a similar QTTP distribution to that of supine, although consistently later in the cardiac cycle.

The QTTP distribution of participants breathing deeply is plotted in Figure 5.52B. Again, the supine and head-down postures see a similar distribution: generally achieving peak perfusion later, the more basal the lung region.

However, a delayed QTTP is observed in the apex of the supine lung. The QTTP values observed in the sitting lung are consistently later across all lung sections, compared to the other two postures, exhibiting the latest peak perfusion in the apex, which reduces to its lowest point in the central regions before increasing again in the base.

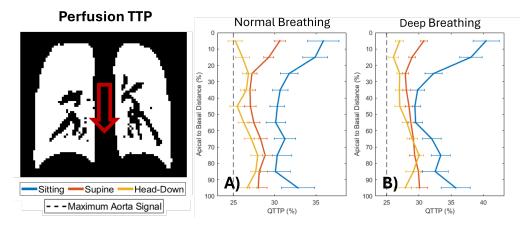


Figure 5.52: Variation in the average QTTP within ten vertically stacked coronal lung segments, for participants breathing normally (A) and deeply (B). Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

Figure 5.53 shows the change in the distribution of the QTTP between participants breathing normally and deeply. These results are similar to those presented in Figure 5.52; however, are regrouped based on posture to compare normal and deep breathing. For all postures, the peak perfusion follows a similar trend between breathing regimes, which appear to be shifted slightly later in the perfusion cycle when breathing deeply.

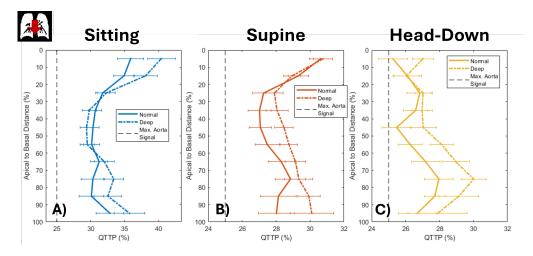


Figure 5.53: Comparison of the variation in the average QTTP within ten vertically stacked coronal lung segments, between participants breathing normally and deeply. Plotted for each of the three postures: Sitting (A), Supine (B) and Head-Down (C). Error bars represent the standard error on the mean for each lung section.

5.3.4.2.2 Investigations in the Sagittal Plane

Sagittal QTTP maps were generated for each participant while breathing normally in each posture. Example maps of a representative participant in each posture are plotted in Figure 5.54.

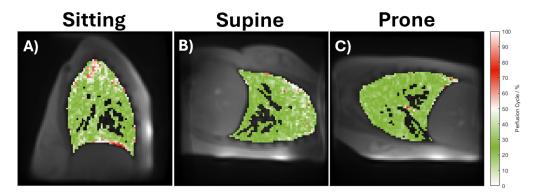


Figure 5.54: Sagittal QTTP maps of a representative participant in the sitting (A), supine (B), and prone (C) postures, breathing normally.

Along the apical-to-basal axis (Figure 5.55A), the QTTP appears largely equivalent between the supine and prone lung, each with no significant variation from 25% of the cardiac cycle, representing the point of max-

imum signal in the central lung vessels. The QTTP within the sitting lung is considerably later in all zones compared to the other postures. The QTTP along this axis is approximately uniform, however is significantly delayed in the most apical zone. Investigating the QTTP along the anterior-to-posterior axis (Figure 5.55B), the supine lung remains approximately equivalent to 25% of the cardiac cycle for all zones. The prone lung exhibits a similar trend, although peak perfusion occurs earlier in the most anterior zones. Once again, the QTTP occurs later in all zones for the sitting lung compared to the other postures. This posture also appears to reach peak perfusion later in the most anterior regions.

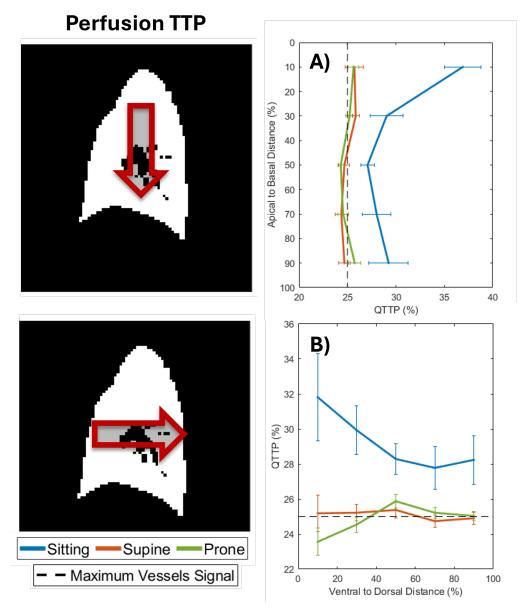


Figure 5.55: Variation in the average QTTP within five vertically stacked (A) and five horizontally adjacent (B) sagittal lung segments, for participants breathing normally. Errors bars represent the standard error on the mean for each lung section. Depicted along side the plots is the mask of an example reconstructed lung with the vessels removed, to aid in the visualisation of the results.

5.4 Discussion

5.4.1 Tissue Compression

5.4.1.0.1 Anterior-to-Posterior

The anterior-to-posterior variation in the tissue density was investigated using experiments in the sagittal plane. This is presented in Figure 5.24 with fit parameters presented in Table 5.8, for the sitting, supine and prone postures. These plots track the variation in signal intensity, which is assumed to be proportional to the tissue density. As expected, in the supine and prone postures, the signal intensity increases along the gravitational vector. Furthermore, the relative change associated with each posture is approximately equal and opposite between the supine and prone postures. Variation in signal intensity along the isogravitational anterior-to-posterior axis of the sitting lung is approximately uniform throughout, albeit with a greater magnitude to the posterior. Each of these trends would be expected due to tissue compression along the gravitational vector, as would be prescribed by Slinky effect[1]. However, in the supine and prone lung, the direction of increasing gravitational force coincides with decreasing distance from the receive coil base. Therefore, while the signal intensity increases along the gravitational vector, this also corresponds to a decreasing distance from the coil. Indeed, the anterior-to-posterior signal intensity distribution for these postures appears to follow a similar pattern to the signal decay with distance due to B1 (receive) inhomogeneity presented in Figure 5.24. Therefore, without per-voxel normalisation, it is difficult to decouple these effects. In future, retrospective bias field correction should be applied to these results before conclusions on tissue density distribution can be drawn.

5.4.1.0.2 Apical-to-Basal

The apical-to-basal variation in exhaled tissue density was investigated in the coronal plane, for the sitting, supine and head-down postures, and the sagittal plane, for the sitting, supine and prone postures. When in the sitting posture, the lung appears to exhibit a significant increase in tissue density with towards the basal regions, in line with the gravitational vector. At first glance, this agrees with the established models of increasing parenchymal density due to compression under its own weight. This distribution is similar across the coronal and sagittal investigations, both in shape and fitted slope. Isolating the anterior and posterior segments, each exhibits a significant increase in signal intensity towards the base, although the magnitude of this relative change is slightly greater in the posterior segment. These results suggest that the coronal slice at the back of the sitting lung provides a good approximation for the lung as a whole.

Evident in both the coronal and sagittal investigations, the trend of increasing tissue density along the gravitational vector in the sitting lung appears to be non-linear; increasing at a greater rate in the central regions. This deviates from the linear relationship expected if this distribution were purely dictated by the Slinky Effect model[1]. One explanation is that the alveoli in the apex have reached near maximum expansion and likewise those in the base maximum compression. Therefore, there would be little variation in signal intensity within these zones, and the central gradient represents the transition between these states. However, if simply the gravitational forces associated with sitting upright were to induce maximal alveolar expansion at the apex in a healthy lung, then it would be highly inefficient to

ventilate this area. Therefore, it is unlikely that this explanation is correct.

An alternative explanation for this phenomenon is a centralisation of mass within the lung, potentially due to the greater density of vessels in this area. This would cause variation in the rate of change of tissue compression along the apical-to-basal axis, amplified in the central regions compared to the peripheries. To visualise this phenomenon, it is useful to imagine an ideal lung of three identical segments, the apex, middle and base. The central segment has the greatest mass, while the apex and base are equal. Orienting the model lung upright as in Figure 5.56, each position within the lung experiences a compression in proportion to the weight of the tissue supported above it. Intuitively, this compressive force is greatest at the base, and thus, so is the tissue density. Moving up the lung, the mass of tissue above decreases, resulting in a decrease in tissue density with respect to vertical position. Upon transitioning to the central segment, the rate of this reduction will increase. This is because this section has a greater mass than that of the base. Therefore, transit up the lung within this segment will return a greater reduction in the weight supported above. This results in a greater reduction in compressive force and thus tissue density with respect to vertical position. Finally, within the apical segment, this rate reduced again as the mass of this segment is lower. To complicate this model slightly, the greater compression of the lower zones will result in a greater change of tissue compression with respect to the vertical height. This manifests as an increased gradient of tissue density in the basal region, and a smoothed transition between the apical and central zone. Substituting the three-segment model for a basally shifted normal distribution, one can create a tissue density to lung position relationship very similar to that of the sitting posture in Figures 5.22 and 5.24A. Therefore, the Weighted Slinky effect appears to be a good approximation for the distribution of parenchyma in the upright lung.

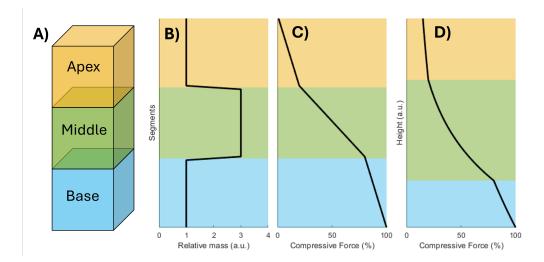


Figure 5.56: The Weighted Slinky Effect describes the redistribution of lung tissue of non-uniform mass deforming under its own weight. When vertically stacked, three ideal segments (A) of differing mass (B) will experience a non-uniform change in compressive force between segments (C). Accounting for this compressive force's effects on the height of each segment (D) will result in a non-linear distribution of compressive force along the vertical axis. The highest regions of the Apex segment (yellow) have been omitted in D.

The variation of exhaled tissue density was investigated along the apicalto-basal axis in the head-down posture only in the coronal plane. This
distribution followed the opposite trend to that of the sitting posture, increasing towards the apex. This variation is in line with the gravitational
vector. Therefore, at first glance, it appears that the distribution of tissue
density in the head-down posture agrees with that of established theory.
However, the relative magnitude variation within the head-down opposite
is reduced compared to that of the sitting posture. This deviates from
the expected redistribution of lung tissue, if this mechanism were purely
dictated by the Slinky Effect model[1]. The simplest explanation for this
may be the non-verticality in participant and slice positioning associated
with the head-down posture. This is expected to have reduced variation in
this posture, although not to the magnitude observed here. If one assumes

the observed non-equivalence to be of physiological origin, two possible explanations present themselves:

• Systematic variation in the mechanical properties. Variations in parenchymal elasticity or compliance would lead to differing degrees of compression in each region while under the same gravitational stress. This may be inherent to the lung tissue itself; indeed, the human body has evolved to spend much of its time upright and therefore may draw a benefit from a particular distribution of mechanical properties. However, this is unlikely, as a non-uniform distribution of mechanical properties is not characteristic of healthy lungs[6]. These results may suggest some form of plasticity in the compliance of the lung. No special consideration was given to the routine of the participants in the hours before this study, and therefore much of this time was likely spent in an upright posture. Consequently, their lungs would have experienced a greater compression at the base than at the apex, for a significant duration; significantly longer than the time spent in the head-down posture. If this had regionally altered compliance within the lung, via a mechanism such as a non-uniform surfactant distribution[81, 82], it could explain the above observations. To test this would require the investigation of a possible temporal dependence on the distribution of tissue density. However, this would require the participants to remain in the head-down posture for an extended period, which might pose risks to their health. An alternative method may be the investigation of the participants in the head-down posture after an extended period while lying down. If the above hypothesis proves true, this could reset any regional variation of mechanical properties due to this mechanism.

• Internal lung structures prevent complete redistribution. The

lung is built from a fractal scaffold of airways and blood vessels [6]. These features may provide an internal mechanical framework and support some of the weight of the parenchyma. This may reduce the compression of the dependent regions due to gravity. This structure would likely stem from the primary blood vessels or early bronchi, particularly where they enter the lung at the hilum. The cartilage present in the early bronchi may also contribute to this scaffold. If this structure is present, then it may not respond proportionally to changes in the gravitational vector, in turn providing different degrees of support to the lung parenchyma[116]. For example, assuming this scaffold has evolved or grown under the conditions of an upright lung, then it may resist movement due to a gravitational vector towards the apex. Imagine the angle of the bifurcation of the trachea into the primary bronchi and how this may change when in the inverted posture. Any resistance would support some of the weight of the basal parenchyma, in turn increasing the expected tissue density at the base and decreasing it at the apex. This would result in a reduction in the variation in parenchymal density compared to the lung in the sitting posture. To test this, one could employ gaseous MRI contrast to map the movement of the primary airways in response to changing posture.

The variation in tissue density was also investigated along the isogravitational coronal plane of supine posture. Within this plane, the model described by the Slinky effect[1] dictates there should be an even distribution of tissue density. However, this is not observed. Instead, the signal increases significantly towards the base of the lung. This pattern manifests in three approximate regions: the signal increases quickly in the apex, occurring at a faster rate than the relative change in the sitting posture. This

change slows in the middle zones before finally decreasing slightly in the basal regions. This relationship is highly unexpected, deviating from the established gravitational models of parenchyma compression[1, 31] which dictate a uniform distribution along isogravitational planes.

The isogravitational apical-to-basal distribution of tissue density was also investigated sagittally for the supine posture. When investigated this way, the lung exhibited an increase in density towards the base, however, this change is significantly reduced compared to that derived from the coronal investigations. This is likely because, while the coronal slice is positioned toward the posterior of the lung, the sagittal plane permits sampling of the full anterior-to-posterior width. When in this posture, the anterior-toposterior axis is in line with the gravitational vector, therefore the posteriorly located coronal slice is positioned within the gravitationally dependent region. When investigating the apical-to-basal variation in just the posterior of the sagittal slice, this segment exhibited a similar distribution to that of the coronal investigation, albeit with slightly greater fitted slope. However, in the anterior segment, the tissue density appears approximately uniform along this axis, in agreement with the gravitationally based models[1, 31]. Therefore, one concludes that there is significant variation in the apical-to-basal distribution of tissue density within the lung, dependent on the anterior-to-posterior position, due to the differing weight supported by the tissue. Furthermore, this means that the observed distribution of tissue density in the coronal slice is not reflective of the lung as a whole and a better approximation may be derived from the sagittal investigation.

The physiological explanations discussed for non-opposite distribution in the sitting and head-down postures could be employed to explain this phenomenon. However, it is likely (at least in part) to be a consequence of non-uniform anterior-to-posterior tissue compression along the apical-tobasal axis. Both the coronal slice and posterior segment are in the dependent lung region, therefore, will experience compression from the anterior parenchyma above. This may contribute to the greater absolute signal intensities observed in the coronal plane for the supine lung. However, as is evident in the sagittal lung images presented in this work, the anterior-to-posterior distance within the lung is not constant. This distance increases gradually towards the base, although it tapers quickly at the apex. Combined with the weight of the heart and major pulmonary vessels, this effect could explain the observed distribution of tissue density along the apical-to-basal distance:

Starting at the apex, the density increases quickly in correlation with the anterior-to-posterior distance. In the central zones, the posterior lung experiences compression from the weight of the major blood vessels and heart. Finally, in the basal zones, the tissue density appears to decrease. This may be due to a reduced weight of vessels[6], narrowing width of the lung behind the diaphragm, or a structural contribution from the diaphragm reducing the weight supported by the lung tissue.

Conversely, when in the supine posture, the anterior segment is located in the gravitationally superior position. Therefore, the lung will experience reduced gravitational loading compared to that of the posterior. This may result in a reduced magnitude of tissue compression and an approximately uniform variation along the apical-to-basal axis, as is observed in this investigation. However, although the change is small compared to the associated error, the tissue density in the anterior segment may drop slightly towards the base. This may reflect a gravitational effect acting in the opposite direction to that observed in the posterior, where the larger mass of the lung towards the base results in a larger expansion of the anterior alveoli towards the base. Thus, reducing the observed tissue density along this

axis.

To summarise, the high variation in apical-to-basal distribution of tissue compression in the supine lung is likely a consequence of non-uniform gravitational loading perpendicular to this plane. This is due to both a variation in the anterior-to-posterior width of the lung and the mass of the blood in the central vessels. The later contribution draws parallels to the centralisation of mass hypothesis employed to describe the distribution associated with the sitting lung.

Finally, the apical-to-basal distribution of tissue density was investigated in the prone posture. One would expect, given the distribution observed in the supine lung, an increasing signal intensity towards the base in the now gravitationally loaded anterior segment. This would be accompanied by a comparatively uniform distribution in the superior posterior segment. However, this is not the case. Both in the full sagittal slice and in each segment independently, the tissue density appears uniform along the apicalto-basal axis. This means that the mechanism dictating the non-uniform parenchymal compression in the dependent region of the supine lung is either not present or countered when in the prone posture. Similar to the explanations of the non-equivalent sitting to head-down apical-to-basal variation, an argument can be made that this is due to structural contributions within the lung or a systematic variation in mechanical properties. However, in this case, there may be an additional explanation. When in the prone posture, the posterior lung often appears to extend further above the diaphragm than is observed when in the other postures. In Figure 5.57, a typical sagittal reconstructed image of the prone lung at an intermediary respiratory state is overlaid with the outline of both the associated prone and the supine masks for this participant. Contrasting the outline of the diaphragm at the base of each of the masks, it appears that in the prone posture the diaphragm has shifted with the gravitational vector towards the anterior of the body. The opposing effect is not apparent in the
supine lung. Combined with the effective immovability of the posterior of
the thoracic cavity, this would result in a drop in pleural pressure in the
posterior-basal region, effectively pulling the lung into the space previously
occupied by the diaphragm and abdominal organs. An alternative way of
thinking about this mechanism is in terms of buoyancy. The gravitational
force will act on the diaphragm and abdominal organs to a greater degree
than on the lower-density lung tissue. Therefore, without constraint on
their movement, the abdominal organs would seek to displace the lungs,
moving them to a more superior position. This may have the effect of
applying a force on the basal lung tissue opposite to the gravitational vector, preventing the non-uniform loading and subsequent compression of the
anterior segment.

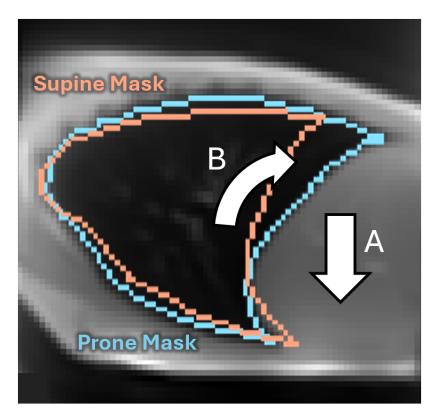


Figure 5.57: Proposed mechanism for the approximately uniform apical-to-basal distribution of tissue density in the anterior of a prone lung. A sample reconstructed prone lung in the exhaled state is overlaid with the outline of the associated prone mask (blue) and the supine mask (orange) of the same participant. The arrows denote the apparent movement of the abdominal organs (A) and lung parenchyma (B) due to this posture.

In culmination, these relationships suggest that whilst the Slinky effect[1] is a useful model for describing the distribution of pulmonary tissue density, it is only a first-order approximation. For a more accurate description, the model may be adapted to reflect the varying weight within the lung. Furthermore, when considering variation in isogravitational planes, one should not omit contributions of gravitational forces perpendicular to the plane. However, in the prone posture, an additional mechanical contribution may be countering this effect. Finally, although correct in sign, this model does not fully describe the change in tissue distribution when maximally altering the gravitational vector of the vertical lung. This may be due to additional physiological mechanisms, such as plasticity of the compliance of

the parenchyma, or internal structural contributions from the early bronchi and major blood vessels. Additional investigation is required.

5.4.2 Ventilation Dynamics

5.4.2.1 Fractional Ventilation

The absolute values of the fractional ventilation were similar to those derived from studies of health participants in the supine posture, acquired at 1.5 T (Section 4.13.4). Current gravitational modes [1, 31], dictate an increase in ventilation along the gravitational vector. Disregarding the above tissue compression results for a moment, this theory dictates there should be an equal and opposite increase of ventilation within gravitationally opposite postures: along the apical-to-basal axis for the sitting and head-down lung and along the anterior-to-posterior axis in the supine and prone lungs. Furthermore, there should be a uniform distribution of ventilation within isogravitational planes: apical-to-basal for the supine lung, and anterior-to-posterior in the sitting lung.

5.4.2.1.1 Anterior-to-Posterior

The anterior-to-posterior variation of the FV was investigated with experiments in the sagittal plane. Measurements of the exhaled tissue density along this axis suffered from the non-homogeneous sensitivity gradients near the receive coil. This is expected to have a reduced impact on the measurement of the FV because this metric is calculated using a per-voxel normalisation step. Therefore, any artificial increase of signal close to the coil should be cancelled on a per-voxel basis. However, if parenchymal tissue moves away from the coil base during respiration, the associated voxel

will exhibit an artificially increased FV. However, as the participants were breathing normally during the sagittal acquisitions, therefore, this effect was assumed to be minimal. As with the exhaled tissue density investigations, this data may benefit from retrospective bias field correction.

If the Slinky effect[1] were to perfectly model the distribution of ventilation, one would expect this function to increase with the gravitational vector. This would result in an equal and opposite magnitude of change in the distribution of FV between the gravitationally opposite supine and prone postures along this axis. Whereas the isogravitational anterior-to-posterior axis of the sitting posture would exhibit a uniform distribution. In actuality, this is close to what is observed. The FV in the supine and prone lung each increased with the respective gravitational vector. However, this change is greater for the supine lung. This matches previous observations in humans[72] and dogs[49], which suggested a more uniform distribution of ventilation in the prone lung compared to the supine lung. However, the sitting lung also exhibits an increase in ventilation towards the posterior. This is approximately equivalent in magnitude to the change along this axis associated with the prone posture.

Unfortunately, due to the sensitivity gradients, it was not possible to compare this result to the variation in tissue compression along this axis in this work. However, if one considers the distribution of the Jacobian determinant, it appears that in all postures, the posterior of the lung expands to a greater degree compared to the anterior (Figure 5.43B). This would likely bias the ventilation towards the posterior regions and, in combination with the Slinky effect[1], explain the observed results in Figure 5.31B. In the supine posture, the gradient of expansion matches that of the tissue compression, along the gravitational vector. Therefore, the positive anterior-to-posterior slope of the FV is amplified compared to the slopes

associated with the other metrics. In the prone posture, the lung continues to exhibit a greater expansion towards the posterior, however, this is now opposed by a reversal of gravitational tissue compression compared to the supine posture. Therefore, the resulting variation in FV, although still increasing with the gravitational vector towards the anterior, does so with a reduced magnitude compared to the supine posture. Finally, for the sitting posture, the uniform distribution of tissue density dictated by the Slinky effect[1], combined with the bias in expansion towards the posterior, create an increase in FV towards this region.

These results deviate from those dictated by a purely gravitational explanation of the distribution of ventilation along this axis. Therefore, this may be an indication of additional physiological contributions along this axis. This may be due to a non-uniform distribution of tissue compression. If so, an argument could be made for this being attributed to structural elements within the lung, or a non-uniform distribution of mechanical properties, similar to that previously discussed for the apical-to-basal tissue compression. However, in this case, there may also be a mechanical contribution dictating more expansion in the posterior regions of the lung. Presumably, pulmonary expansion in this plane is predominantly dictated by both the intercostal and accessory muscles and the diaphragm. However, along this axis, the diaphragm may not contract evenly. If it were to contract to a greater degree in the posterior region, then the resulting regionally dependent expansion in this area may explain the observed bias in ventilation. Additionally, when lying on one's front, it may be harder for the intercostal muscles to expand the lungs, due to the weight of the upper torso above. Therefore, respiration in this posture would be more reliant on the diaphragm and this may account for the lower average FV in this posture compared to the supine posture.

5.4.2.1.2 Apical-to-Basal, Normal Breathing

The apical-to-basal variation of FV was investigated in the anterior-toposterior direction in both the coronal and sagittal planes. These investigations encompassed the gravitationally-opposite sitting and head-down
postures, as well as the isogravitational supine and prone planes. The
Slinky effect[1] predicts an increase along the gravitational vector in ventilation, which would result in a near equal and opposite change between the
sitting and head-down postures, and uniform distributions in the supine
and prone posture along this axis. However, this appears to only reflect a
first order approximation of the observed results.

The sitting posture exhibited an increase in this function towards the base of the lung, parallel with the gravitational vector. Furthermore, the magnitude of this change was greater than that of the other postures investigated — approximately equal to the mean FV of the slice. This is reflected in both the coronal and sagittal results, which are within experimental error of each other. This suggests that for the sitting lung, these slices are predominantly representative of the lung function. However, the sitting lung exhibits a slight inequality in the apical-to-basal slope between the anterior and posterior segments, this is likely a result of the difference in tissue density distribution discussed previously. Moreover, the FV within this posture appears to increase at the greatest rate in the central regions, plateauing in the most basal regions. This may also be due to the distribution of tissue compression, potentially due to a non-uniform distribution of mass within the lung. Therefore, these distributions suggests that the Slinky-model[1] of the lung provides an accurate description of the distribution of ventilation. Due to gravity, tissue compression increases towards the base. As a result, these basal alveoli occupy a more compliant part of the pressure-volume curve and expand to a greater degree during inhalation than those at the apex[1, 31]. However, it is notable that the magnitude of the apical-to-basal increase in FV is more than twice that of the observed change in tissue compression. This may suggest further mechanical contributions to the observed FV trend.

When the participants were breathing normally in the head-down posture, the FV appeared approximately uniform along the gravitational vector. This is counter to the expected distribution dictated by established models[1, 31]. The associated slope of tissue density variation is reduced compared to the sitting posture (Table 5.7), however this reduction would likely not solely account for the observed near-uniform distribution of FV. Furthermore, while the sitting variation in FV is amplified compared to the tissue density, in the head-down lung, it is reduced. This would suggest that there is a bias in FV distribution towards the base of the lung. An explanation for this basal bias may be found in the maps of the Jacobian determinant, although the errors with this result were high. In the headdown posture, there appears to be greater parenchymal expansion in the basal regions compared to the apex. Contrary to the Slinky effect[1], this suggests that when in the head-down posture, even though the greatest tissue compression is located at the apex, expansion is biased towards the basal regions. These contributions counter each other, resulting in a more uniform distribution of ventilation compared to the sitting posture. Unfortunately, due to safety limits of time spent in the head-down posture, this posture was not investigated sagittally. Therefore, unlike the other postures, there is no indication whether this coronal slice is representative of the lung function.

The isogravitational apical-to-basal variation of FV associated with the supine posture was investigated both coronally and sagittally for partici-

pants breathing normally. In the coronal plane, two key observations can be made. First, the average magnitude of this function in this slice is greater than for the sitting and head-down postures. Second, we observe a significant gradient towards the base of the lung, which appears to level out in the most basal regions. Both these observations may be explained by the positioning of the coronal slice at the posterior of the lung. As discussed previously, the tissue compression appears greater in this plane, likely due to the weight of the anterior lung above it, which in turn increases towards the base. This compressed tissue occupies the more compliant region of the pressure-volume curve and therefore expands to a greater degree during inhalation. However, the drop in tissue compression observed in the most basal regions is not apparent in the normal-breathing FV distribution in this plane.

Comparing these results with the apical-to-basal variation in the sagittal plane reveals several differences. Across the whole slice, the FV increases towards the base, although with a reduced magnitude. This reflects the observed overall reduction in tissue compression variation in this plane, which, like the coronal slice, is approximately half that of the FV variation. When considering just the variation in the posterior segment, one observes a similar distribution to that of the coronal plane; both in the shape and relative magnitude. The FV in the anterior segment, which exhibited approximately uniform tissue density along this axis, appears to also increase towards the base. However, this increase is reduced in magnitude compared to that of the whole slice. These results give further weight to the hypothesis that the large variation observed in the coronal slice is only representative of the posterior of the lung.

All apical-to-basal FV distributions associated with the supine lung exhibit an increase towards the base. This may at first appear at odds with the distribution dictated by the Slinky effect[1] in this isogravitational plane, although this can be partially attributed to gravity via compression from the anterior of the lung. However, taking into account the associated errors, this does not explain the above observations completely. All postures exhibit an increased slope of FV towards the base compared to that of the exhaled tissue density. This includes the approximately uniform distribution of tissue density in the anterior segment. Therefore, once again it appears there may be an additional mechanism creating a bias in ventilation towards the base of the lung.

Finally, the isogravitational apical-to-basal variation of the FV associated with the prone posture was investigated in the sagittal plane. In this complete slice, and in both segments when investigated independently, the FV appeared to increase towards the base of the lung. This is similar to the anterior of the prone lung, where the associated distribution of tissue density was approximately uniform along this axis. Therefore, this effect cannot be attributed to non-uniform tissue compression and an additional mechanism may be creating a bias of FV towards the base of the lung. However, unlike the other postures, there also appeared to be a slight central bias to the FV distribution along this axis. This trend may be attributed the redistribution of basal parenchyma due to movement of the abdominal organs, discussed previously. This shift in lung tissue may slightly expand the basal alveoli in the expired state, moving them into a less compliant region of the pressure volume curve. This could result in a slight reduction in FV in this region. However, this may be an artifact of the interdependence of investigations along the apical-to-basal and anterior-to-posterior axes. As depicted in Figure 5.19, this basal zone exhibits a significant bias towards the posterior, more so than is observed for the other postures. These posterior regions of the prone posture exhibited a reduction in FV

compared to those at the anterior regions (Figure 5.31Bii); therefore, this may reduce the observed FV in the basal zones when investigated along the perpendicular axis. Further investigation into this effect is required before meaningful hypotheses can be made.

5.4.2.1.3 Apical-to-Basal, Deep Breathing

During the coronal acquisitions of the sitting, supine, and head-down postures, the participants were also investigated while breathing deeply. When sitting, the apical-to-basal slope of the FV increased compared to the distribution associated with breathing normally. In absolute terms, there appeared to be no significant change in the FV of the most apical zone. However, for each subsequent zone along this axis, the FV was amplified by an increasingly greater amount, until the most basal zone. This deep breathing trend represents increasing FV along the gravitational vector, both in absolute terms and the difference from to the distribution associated with breathing normally.

When deep breathing in the head-down posture, the distribution of FV exhibits a similar but opposite trend to that of the sitting posture. The apical-to-basal slope associated with breathing normally became more negative when the participants breathed deeply. In the most basal zone, there was no significant change. However, for each subsequent zone along this axis, the FV was amplified by an increasingly greater amount, until the most apical zone. As in the sitting posture, this trend represents an increasing FV along the gravitational vector, both in magnitude and relative to the distribution associated with normal breathing. Notably, where the magnitude of FV variation associated with normal breathing decreased relative to the variation in exhaled signal intensity, breathing deeply saw this

variation increase.

Finally, the supine posture exhibited a more uniform increase in FV along the apical-to-basal axis, compared to the other two postures. Seemingly the only deviation between the relative FV change of normal and deep breathing in this posture is the introduction of two 'humps' in the FV centred at approximately 35% and 70% of the apical-to-basal distance. These are similar to those observed in the supine exhaled tissue density distribution. This results in a small increase in the apical-to-basal slope; however, this is small compared to the other two postures. Therefore, when breathing deeply, the relative distribution of FV increased along the respective gravitational vector. When in the supine posture, the apparently uniform increase in FV along this axis may be reflective of this being an isogravitational plane.

5.4.2.1.4 Tissue density vs Fractional Ventilation

The results from the apical-to-basal investigations of exhaled tissue density and FV are summarised in Table 5.16. At first glance, the exhaled tissue density increases along the gravitational vector. However, considering the above discussion of FV, it appears that irrespective of tissue density variation in the exhaled state, when breathing normally there is a bias in ventilation towards the base of the lung. This serves to increase the fitted FV apical-to-basal slope in all postures, compared to that of the exhaled tissue density. Amplifying the slope for the sitting and supine postures, inducing a slope for the prone posture, and reducing the slope towards the apex when in the head-down posture. However, when breathing deeply, the relative amplification of FV appears to align with the gravitational vector, increasing towards the base when sitting, towards the apex when in the

head-down posture, and an approximately uniform increasing along this axis when supine. Furthermore, the head-down and supine distribution of the deep breathing FV appears to be more reflective of the distribution of exhaled tissue density, compared to that of the normal breathing distributors. This may be because deep breathing induces a more uniform opportunity for expansion across the lung in these postures, resulting in the FV better matching the areas which can expand the most.

Contribution to Apical-to-Basal Slope				
Posture	Lung Segment	Exhaled TD	Normal FV	Deep FV
Sitting	Full	+	+	+
Head-Down	Full	-	+	-
Supine	Full	+	+	None
	Anterior	None	+	
	Posterior	+	+	
Prone	Full	None	+	
	Anterior	None	+	
	Posterior	None	+	

Table 5.16: Quantitative analysis of cumulative factors affecting the apical-to-basal distribution of tissue density and FV. Starting with the slope of the tissue density in the exhaled state, contributions are relative to that of the previous column. This table highlights the apparent gravitational contribution to tissue density and FV when breathing deeply. However, the contribution when breathing normally is always positive, irrespective of the relative direction of the gravitational vector in that posture.

Mechanical explanation When breathing normally, the lungs in all postures exhibited a bias in the distribution of ventilation towards the basal regions, compared to the associated distribution of exhaled tissue density. This was irrespective of posture and consequently the relative gravitational vector; therefore, it cannot be explained by the

gravitationally dependent model[1, 31]. This observation may instead be due to physiological factors. Mechanical work is required to induce ventilation within the lung. This work must be sufficient to overcome a number of factors, including but not limited to: the displacement of lung tissue; the natural elasticity of the parenchyma[6]; and the changing alveolar pressures associated with inspiration[9].

If the mechanical properties of the acini are uniform along an imagined chain, then for a uniform expansion, each unit must experience the same expansive force. If one assumes that the chain is actuated from only a single end, then some of the force applied to each unit must transfer through to the next. Without any resistance to expansion - from factors such as the alveolar elastic recoil or airway resistance - each unit along the chain would expand sequentially. Starting from the end of the chain experiencing the actuation force, each unit will expand to its maximum volume before the force is transferred to the next acinus. Conversely, if the resistance to expansion is high, the resulting distribution of expansion would be more uniform, however, the process of inspiration would be highly inefficient due to the work required for expansion. Therefore, during inspiration, a fine balancing act must occur within each unit, dictating the transfer of work to the next acinus. However, in addition to the force required to expand each acinus, is the force required to displace the lung tissue (Newton's 2nd Law). This force is reduced for each unit of the chain with distance from the actuator, as the number of subsequent units, and by extension the remaining mass to displace, is reduced.

Considering the consequences of this effect, but only in terms of elastic recoil and Newton's 2nd Law, and assuming an otherwise perfectly uniform distribution of mechanical properties. Then the potential re-

quired to expand each alveolus from minimum volume remains equal along the chain. However, starting from the actuator, the potential required to displace the lung tissue drops with each subsequent unit. Therefore, to balance these forces, those units nearer to the actuator will expand to a greater degree until the elastic recoil force matches the force required to displace the remaining parenchyma. This is depicted in Figure 5.58. Conversely, units further down the chain do not need to expand to the same degree for their elastic recoil force to match the reduced force required to displace the remaining lung tissue. Therefore, lung units nearest to the actuator will expand to the greatest degree. This expansion, and thus ventilation, will drop with each subsequent lung unit [117, 118].

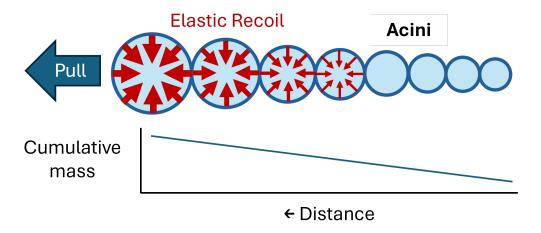


Figure 5.58: Proposed mechanism explaining the apparent bias of FV bias towards the basal regions of the lung. The lung is modelled sequential chain of acini with actuation occurring from a single end. As distance from the actuator increases, each acinus is required to displace progressively less cumulative mass, thereby requiring diminished closure force (in this case elastic recoil) to maintain force equilibrium. Given that elastic recoil force is proportional to the size of the alveoli, the functional units nearest to the actuator expand to a greater degree compared to those further away. The red arrows represent the closure force, with the width representing the relative magnitude. Not to scale.

Assuming that along the apical-to-basal axis expansion is primarily dictated by the diaphragm, then the above model would explain the

observed bias of FV towards the basal regions. However, not discussed above is the effect of changing posture and as a consequence the gravitational vector. Assuming the participant is at rest and under normal gravitational conditions of 1g. When sitting, the lung is oriented upright with the diaphragm at the bottom. The movement of parenchyma during inspiration is assisted by gravity. Therefore, the potential required to move the lung parenchyma during inspiration is now negative. This would lead to a more uniform distribution of ventilation and may contribute to the observed FV distribution better reflecting the shape of the exhaled tissue density compared to the other postures. When horizontal, the diaphragm is located to the side, therefore the movement of lung parenchyma would not be significantly assisted or hampered by the force of gravity. Therefore, this mechanism may create a slight bias in ventilation towards the diaphragm. However, in the head-down posture, the lung is oriented upright with the diaphragm at the top. Opposite to the effect in the sitting lung, now gravity is resisting the movement of lung parenchyma during inspiration, thus increasing the potential required to move it. This would increase the bias in ventilation distribution towards the base of the lung due to this mechanism, as is observed in these results. This hypothesis is further supported by the observation of an apparent non-uniform distribution of lung parenchyma due to compression under its own weight. Therefore, the alveoli closure force appears to be of a similar magnitude to that of the weight of the lung parenchyma.

This model is highly simplified and doesn't consider numerous additional mechanical contributions in healthy lungs[116]. These include the airway resistance, which could temporarily increase the potential

required to inflate the acini as the alveolar pressure becomes more negative [9]. The often-reduced respiration rate associated with deep breathing may reduce the effect of this mechanism as there is more time for the alveolar pressure to equalise [9]. As discussed at length previously, a non-uniform distribution of tissue density will result in a variation in the potential required to expand the alveoli. Those more compressed alveoli occupy a more compliant region of the pressure volume curve and thus reduce the potential required for relative expansion [1, 31]. In addition to this is the effect of surfactant, which may reduce the potential required for expansion at minimal alveolar concentrations [117]. This could increase the relative expansion in areas of greater alveolar compression.

This hypothesis could be tested by comparing the ventilation distribution when participants are breathing normally, or under a positive inspiration pressure. This positive inspiration pressure may assist in the inflation of the alveoli; both reduce the potential required to expand each acinus and dependency on the diaphragm for inspiration. Resulting in a more uniform distribution of ventilation within the lung.

In summary, when breathing normally, the distribution of FV in the lung appears to be both dependent on the tissue compression in the exhaled state, and the position of mechanical input. When sitting, this results in a significant apical-to-basal slope in the FV towards the base. However, when in the head-down posture, the variation in exhaled tissue density is opposite in sign to that of the mechanical bias towards the diaphragm. Therefore, these contributions counter each other, resulting in a largely uniform distribution in the FV along this axis. When in a horizontal posture, the mechanical contribution from the diaphragm likely results in the

observed increase in ventilation towards the base, even with uniform tissue distributions. In the posterior of the supine lung, this contribution combines with the increased exhaled tissue compression along this axis to amplify the slope of the FV towards the base.

When breathing deeply in the supine and head-down postures, the apical-to-basal distribution of FV better reflects the distribution of the exhaled tissue density, compared to breathing normally. As mentioned previously, this may be because the respiration rate associated with breathing deeply is slower, thus the contribution from airway resistance on the bias of ventilation towards the diaphragm is diminished. However, this may also be because when nearing maximal inhalation, the lung tissue has had more of an equal opportunity to expand. The elastic potential of the expanded alveoli exceeds the potential required to displace the lung tissue, thus negating the bias towards the diaphragm. This is reflected in the more uniform distribution of the Jacobian determinant for this breathing regime. As a result, the regions of greatest tissue compression during exhalation return the greatest increases in ventilation.

5.4.2.2 Jacobian Determinant

The variation in the Jacobian determinant was investigated along the apical-to-basal axis in the coronal plane and along the apical-to-basal and anterior-to-posterior axes in the sagittal plane. These results have a number of limitations: Firstly, they are based on two-dimensional registration; therefore, they represent the two-dimensional Jacobian. However, the lung is a three-dimensional system. Any movement perpendicular to the imaging plane will not have been captured by this processing. Secondly, the registration algorithm [119] employed to create these functional maps is based on the

intensity of each respective frame of the reconstructed ventilation cycle. Therefore, signal variation due to B1 (receive) inhomogeneity may have contributed to these results, reducing their accuracy. Finally, the example results in Figures 5.38 and 5.42 depict a somewhat patchy distribution of the Jacobian determinant. This was typical of all participants and represented a highly heterogeneous expansion pattern through the lung. This was physiologically unexpected, as healthy lungs should not possess the necessary variation in mechanical properties in neighbouring lung regions to create such an effect. Instead, this may be due to improperly optimised registration parameters. This algorithm has been designed to work with CT and MRI data captured at higher field strengths. Therefore, a greater number of features (often pulmonary vessels) are usually visible within the slice. These features are then used by the registration algorithm. In contrast, the images captured in this study often exhibited few internal features. Therefore, the intensity-based registration algorithm may primarily have utilised features such as the walls of the thoracic cavity and noise within the lung. This appears to have led to errors in the results, where regions exhibited significant expansion at the expense of their neighbours, creating the patchy maps of the Jacobian determinant. Therefore, we conclude that these maps are not useful for per voxel expansion determination for individual participants. However, when combined across participants, any local variation where regions expand at the expense of their immediate neighbours should average out. Therefore, these distributions remain useful for regional mapping of where expansion takes place. For example, if there is a greater degree of expansion near the base of the lungs compared to that of the apex, this may relate to a greater contribution from the diaphragm compared to the accessory muscles in this plane. This is how results from the Jacobian determinant have been interpreted. This error is not expected to have significantly affected the registration during the PREFUL processing as this was conducted via a different registration method.

5.4.2.2.1 Anterior-to-Posterior

The anterior-to-posterior variation of the Jacobian determinant was investigated in the sagittal plane for the sitting, supine, and prone postures. All postures exhibited a positive Jacobian determinant for all zones along this axis, representing an expansion of the lung tissue with inspiration. Additionally, all postures exhibited a greater degree of expansion towards the posterior compared to the anterior of the lung. This may be due to the greater height of the diaphragm in this region; therefore, there may be more capacity for expansion due to the muscle in this area. This may contribute to the apparent bias towards this region in FV, as previously discussed. Furthermore, this anterior-to-posterior heterogeneity appears to be greatest in the supine posture. The greater degree of tissue compression in this area may explain this phenomenon, increasing the compliance of the lung parenchyma and facilitating a greater degree of expansion in this region compared to the other postures. However, these results suffer from a high degree of variance, as signified by the large error bars in Figure 5.43B. Therefore, further investigation of these results is required.

5.4.2.2.2 Apical-to-Basal

The apical-to-basal variation of the Jacobian determinant was investigated in the coronal plane for the sitting, supine, and head-down postures and in the sagittal plane for the sitting, supine, and prone postures. As for the anterior-to-posterior results, all zones along this axis exhibited a positive Jacobian determinant, representing an expansion of lung parenchyma with

inspiration. When breathing deeply, this function increases further, representing the expected increased degree of expansion. When investigating the whole slice, each posture exhibited a central suppression in the Jacobian determinant, potentially due to a relative reduction in expansion in the areas near the central lung vessels and primary bronchioles. Therefore, the results in the coronal plane were also plotted for parenchyma only in the lung peripheries in Figure 5.40.

When breathing normally, the distribution of the Jacobian determinant appears similar within errors between the sitting and supine postures. These distributions suggest a greater expansion in both the apex and to a greater degree towards the base of the lung (Figures 5.39 and 5.43). This may be representative of the contributions to tidal respiration from the accessory muscles at the apex and the diaphragm at the base of the lung. Notably, the increase at the apex is larger for the supine posture than for the sitting. This may signify a greater contribution to respiration from the accessory muscles when one is lying horizontally compared to sitting upright. When upright, these muscles must work against gravity to expand the lung. However, when in the supine posture, expansion in this plane is perpendicular to gravity, therefore it may require less work to expand the lung by this mechanism. Thus, they may actuate to a greater degree when horizontal.

When breathing deeply, the above effect appears to be amplified. The Jacobian determinant of each zone in the basal regions appears to be approximately equivalent between the sitting and supine postures. However, these deviate significantly in the apical zones. In this breathing regime, the Jacobian determinant appears to increase towards the base when sitting, whereas this function appears largely uniform when supine, although increases significantly in the most apical zones. One may conclude from these results that when supine, the accessory muscles contribute significantly to

the respiratory cycle, particularly when breathing deeply. This contrasts with the sitting posture, where the expansion appears to exhibit a slight bias towards the diaphragm.

The Jacobian determinant associated with the prone posture appears largely uniform along this axis, albeit for the aforementioned reduction in the central zones. As for the supine posture, an increase in expansion is observed in the apex, although this appears reduced in magnitude. The reduction in this region compared to the supine posture may be explained by the positioning of the subjects (Figure 5.8). When prone, the participants were instructed to raise their arms above their head. This may change the contribution to respiration from the accessory muscles. Therefore, future experiments should ensure continuity in participant arm positioning between postures. Furthermore, the variance of the results in the sagittal plane is too large for any meaningful hypothesis to be made.

Finally, the variation in the Jacobian determinant was also measured for the head-down posture. When breathing normally, the distribution of this function in the most apical zones reflected that of the sitting and supine postures. However, in the middle and basal zones, this deviated significantly. Starting from the central region, the Jacobian determinant in these zones appeared to increase significantly, ultimately resulting in a far greater expansion in the basal regions. When breathing deeply, the expansion in the apical region increased compared to breathing normally; however, this was reduced in degree compared to the other postures. This resulted in apparently lower expansion in this region when breathing deeply. In contrast, the average Jacobian determinant towards the base was amplified, resulting in a greater basal expansion compared to the sitting and supine postures. As discussed previously, we do not expect this Jacobian determinant to accurately reflect local expansion of the lung parenchyma, due

to errors in the registration method used to investigate this function. This is reflected in these results, as such variation in this function is physiologically unlikely when breathing deeply and is counter to the results of the FV. Therefore, these distributions of the Jacobian determinant suggest that when in the head-down posture, the expansion of the lung is highly dependent on the diaphragm[120] instead of the intercostal and accessory muscles. This may be explained physiologically by two mechanisms: firstly, similar to the prone posture, participants in the head-down posture have their arms positioned above their head. This may result in reduced actuation of the accessory muscles during respiration. Secondly, the thoracic cavity at the apex of the lung may remain expanded in the expired state due to the force of gravity, resulting in a reduced expansion during inspiration[121]. The process of expiration is passive and dependent on the elastic recoil of the lung alveoli. Therefore, there may be few muscles to resist the expansion of the apical thoracic cavity under its own weight when in the head-down posture. This could also contribute to the apical-to-basal tissue density variation observed previously; if the apical thoracic cavity is expanded, there may be a reduction in tissue compression in this region. Additionally, the weight of the organs on the diaphragm may compress the lung, shifting the diaphragm towards the apex. This may result in a greater potential for actuation of this muscle compared to the other postures, further amplifying the lung-volume dependence on this muscle. This hypothesis requires further investigation of the posture dependence of the shape and size of the thoracic cavity. However, consider the effect of positional asphyxia. When in an inverted posture, the weight of the organs in the torso (positioned above the lung) can restrict the movement of the diaphragm. This is exacerbated in overweight individuals and can eventually lead to asphyxiation. This correlates with the hypothesis above, that when in the head-down posture, respiration is highly dependent on the actuation of the diaphragm. If the diaphragm is restricted in any capacity, the intercostal and accessory muscles may be unable to offset the resulting penalty to respiration. Then ultimately this may reduce ventilation in the lung and in extreme cases lead to asphyxiation.

5.4.2.3 Ventilation Time-to-Peak

The VTTP variation was investigated along the apical-to-basal axis in the coronal plane and along the apical-to-basal and anterior-to-posterior axes in the sagittal plane. This function is based on differences of phase with respect to the ventilation cycle, not signal amplitude. Therefore, assuming minimal anterior-to-posterior movement of lung parenchyma, it is expected that the sensitivity gradients had a minimal effect on the VTTP results.

5.4.2.3.1 Anterior-to-Posterior

The anterior-to-posterior variation of the VTTP was investigated in the sagittal plane for the sitting, supine and prone postures. All zones of all postures appear to reach maximum ventilation near 50% of the ventilation cycle, representing the point of greatest lung volume. However, as indicated by the relatively high error bars compared to the variation within each posture, The variation of these results between participants was large. The prone posture appears to reach peak ventilation approximately uniformly along this axis, slightly earlier than the other two postures. This may be due to the apparently reduced tissue compression in this posture, which may have resulted in a reduced airway resistance. In the anterior zones, the sitting and supine postures each exhibited VTTP values at approximately 50% of the ventilation cycle. However, in the posterior zones, the VTTP these postures deviated to approximately mirror each other about the point

of greatest lung volume. This resulted in the posterior lung ventilating earliest when supine and latest when sitting and may be due to differing action of the diaphragm. However, more investigation is required to confirm these observations.

5.4.2.3.2 Apical-to-Basal

The apical-to-basal variation of the VTTP was investigated in the coronal plane for the sitting, supine, and head-down postures, and in the sagittal plane for the sitting, supine and prone postures. As for the anterior-to-posterior investigation, the prone posture appeared to reach maximum ventilation earlier than the other postures, slightly before the point of maximum lung volume. The supine posture exhibited a slightly delayed VTTP in the central and basal segments compared to those of the apex. This delay may reflect variation in airway resistance due to apparent tissue compression. Therefore, the apical regions, which appear to have experienced the least compression at exhalation, reached their maximum ventilation first.

The head-down posture shared a similar VTTP distribution to that of the supine posture. It appeared largely uniform along this axis, approximately equal to the point of maximum lung area, although slightly earlier in the apical regions. This appears to counter established theory[9] that the VTTP is related to airway resistance due to tissue compression, as in this posture the gravitationally dependent apical regions are compressed to the greatest degree. An explanation may instead be found in the associated distribution of the Jacobian determinant. While most of the expansion appeared to be due to the diaphragm, there is a slight increase in the Jacobian determinant in the most apical zones. If significant, this may be indicative of a slight extension due to the accessory muscles. With the

assistance of gravity, this region may expand early in the respiration cycle, thus resulting in slightly earlier VTTP values in this region.

The greatest variation in the VTTP appears to have occurred in the sitting posture. Here, the central regions appear to reach maximum ventilation significantly later than the position of maximum lung volume. In the basal regions, the distribution returns to 50% of ventilation cycle and in the apical regions, slightly earlier. This is apparent in both the sagittal and to a greater degree in the coronal investigation. The comparison of the most apical and basal points supports the hypothesis that time of maximum ventilation is related to airway resistance due to tissue compression. Furthermore, this observation is in agreement with established theory that the apex of the upright lung ventilates first, due to the reduced airway resistance associated with tissue compression. However, the significantly delayed point of maximum ventilation in the central regions is highly unexpected, especially as this region appears to reach maximum ventilation at a point where the lung should be reducing in size. This may be reflective of a limitation of this analysis method. No distinction is made between air in the airways and those contributing to gas exchange in the alveoli. Therefore, early in expiration, if the alveoli experience a greater closure force compared to that of the airways, then the gas in the alveoli may serve to further inflate the airways. This would result in the central regions, with the highest proportion by volume of airways, appearing to ventilate more at the beginning of the expiration. If true, it is not immediately obvious why this effect is not apparent for the other postures, although may be due to the greatest degree of tissue compression exhibited along this axis for this posture increasing the closure force of the alveoli. However, these plots make no distinction between variation near the hilum and at the lung periphery. Variation due to the above mechanism would be expected primarily near the hilum where the proportion by volume of airways is the greatest. In the example coronal VTTP maps in Figure 5.33, it appears this delay is present even in the periphery of the sitting lung. Therefore, more investigation is required into this observation before meaningful hypotheses can be drawn from these results.

Finally, the VTTP was also investigated whilst the participants were breathing deeply. All postures investigated in this regime exhibited a distribution of this function nearer to the point of maximum lung area compared to when breathing normally. This includes the central regions of the sitting lungs and the apical regions of the head-down and supine lungs. This may be reflective of the slower respiration rate associated with breathing deeply and thus airway resistance has a reduced effect on peak ventilation times. Additionally, this may be because the lungs have a more equal opportunity to expand as one nears one's total lung capacity.

5.4.3 Perfusion Dynamics

5.4.3.1 Normalised Perfusion

The gravitational distribution of perfusion within the lung is described by West's zonal model[2], depicted in Figure 5.3. This model dictates an increase of perfusion along the gravitational vector due to variations in hydrostatic pressures. This occurs in four distinct zones; however, only the latter three zones are expected to be present in the healthy lungs investigated in this study. The variation of normalised perfusion was investigated for all four postures: along the apical-to-basal direction in the coronal and sagittal investigations and along the anterior-to-posterior direction in the sagittal acquisition.

5.4.3.1.1 Anterior-to-Posterior

The anterior-to-posterior variation in the normalised perfusion was investigated for the sitting, supine and prone postures in the sagittal plane. Both the supine and prone postures exhibited a significant increase in perfusion with the gravitational vector. However, this also corresponds to a reducing distance to the coil. Therefore, the B1 (receive) inhomogeneity likely contributed to this distribution, increasing the apparent perfusion towards the coil. The relative magnitude of this effect compared to the true perfusion variation in this plane is unclear. Therefore, without bias field correction, this parameter is only partially useful in this work. Contrasting the observed apparent anterior-to-posterior perfusion variation between the sitting and supine postures, there appears to be a greater variation of perfusion in the supine lung. It is assumed that each of these postures exhibited a similar degree of signal variation due to the B1 (receive) inhomogeneity in this plane, due to the similar positioning of the coil relative to the participants. Therefore, it can be deduced that there may be a greater magnitude of perfusion variation along the gravitational vector in the supine lung compared to this isogravitational plane of the lung in sitting posture. This correlates with established theory [2], that the variation in tissue density and hydrostatic pressure along the gravitational vector contributes to a greater perfusion variation in this plane in the supine lung.

The anterior-to-posterior tissue-density-normalised perfusion distributions, plotted in Figure 5.49Bii, were generated using per-voxel normalisation. Therefore, the effect of the B1 (receive) field inhomogeneity should be reduced. However, as for the FV variation along this axis, these data are reconstructed from multiple respiratory states. Therefore, each voxel in the reconstructed map will take contributions from multiple points in space,

each corresponding to different positions within the B1 (receive) inhomogeneity. Unlike the ventilation reconstruction, and assuming no correlation between respiratory and cardiac phase, the B1 (receive) field inhomogeneity effects should largely cancel out during the interpolation step.

The anterior-to-posterior tissue-density-normalised perfusion variation along the anterior-to-posterior axis appears significantly reduced compared to the associated volume-normalised-perfusion distributions. This may be indicative of two factors: first, the bias in the apparent perfusion due to the B1 (receive) inhomogeneity towards the base of the coil has been reduced. Second, the tissue density appears to account for much of the variation in perfusion variation along this plane. The latter factor is in corroboration with observations by Hopkins et al.[1] that tissue density is a primary determinant of observed perfusion variation along this axis. Although reduced in magnitude, there appears to be an increasing perfusion along the gravitational vector towards the anterior of the prone lung. This may be reflective of the more uniform tissue density along this axis compared to the supine lung. Therefore, normalising by this parameter has a reduced impact on the perfusion distribution. Additionally, there appears to be a slight trend of increasing perfusion along the gravitational vector towards the posterior of the prone lung. However, the fitted slope is small compared to the associated error. Additionally, all postures exhibit a drop in perfusion in the most anterior and posterior zones, which may be indicative of a central bias to perfusion. This could be due to the greater volume of vessels in this area; an observation discussed further in the next section. These observations should be repeated with bias field correction to remove artificial contributions from the inhomogeneous B1 (receive) field.

5.4.3.1.2 Apical-to-Basal

Unlike the anterior-to-posterior investigation, it is assumed that the B1 (receive) inhomogeneity had minimal effect on the normalised perfusion distribution along the apical-to-basal axis. This is especially true for the tissue-density-normalised distributions as, to create these, each voxel was normalised independently. Therefore, any signal variation due to the inhomogeneous B1 (receive) field should cancel out. However, future investigations may benefit from retrospective stray field correction.

The apical-to-basal variation of perfusion was investigated in both the sagittal and coronal planes. For the sitting posture, the coronal apicalto-basal distribution of volume-normalised perfusion is plotted along the vertical axis in Figures 5.45A and 5.49Ai. The analysis yielded a significant increase in perfusion with the gravitational vector. Rapidly increasing within the apical-central zones, before it plateaued at approximately 70% of the apical-to-basal distance and dropped slightly in the most basal zones. This distribution is consistent in both the coronal and sagittal acquisitions and is similar to that presented by Ax et al.[4] (Figure 5.4). Therefore, at first glance, the distribution of perfusion observed in the sitting lung increases along the gravitational vector, in agreement with current gravitational models[2]. This distribution is non-linear and approximately resembles a continuous distribution similar to that dictated by West's zonal model[2]. However, this method of investigation permitted further analysis of the results to isolate contributions to apparent perfusion beyond those due to hydrostatic pressure differences.

The observed apical-to-basal distribution of perfusion in the sitting lung may be the product of three contributions. First, the previously discussed tissue density variation along this axis will result in a greater density of vessels towards the base of the lung. This is reflected in the increased perfusion variation per voxel along the gravitational vector. However, this is not indicative of an increased alveolar perfusion. To remove this contribution, the perfusion was normalised by the estimated tissue density associated with each voxel (Figures 5.45B and 5.49Aii). This resulted in a shift in the position of the plateau to a more central location within the lung. Subsequently, the second contribution to the perfusion distribution may be a central bias. This could be indicative of either the hilum to peripheral gradients described in previous work [43], or insufficient masking of the central lung vessels. However, this effect was not investigated further as part of this work. Assuming this effect to be due to a hilum-to-peripheral gradient, to negate it, only lung areas of a similar hilum-to-peripheral position should be plotted. Therefore, the apical-to-basal variation of perfusion was investigated in only the most peripheral area. Plotting this result for both the volume-normalised and tissue-density-normalised perfusion yields Figure 5.46. In this figure, the perfusion distribution associated with the sitting posture generally appears to increase along the gravitational vector, doing so quickly in the apical-central zones, before levelling off beyond the centre of the lung, and dropping slightly in the most basal zone. This demonstrates a strong correspondence with West's zonal model[2] (Figure 5.3). Consider the tissue-density-normalised perfusion in Figure 5.46B. The steep increase along the gravitational vector in the apical-central area may be reflective of West's 2nd Zone[2]. The observed reduction in gradient between the central and basal regions may be reflective of West's 3rd Zone[2]. Finally, the drop in the most basal zone may be reflective of the 4th Zone in West's model. Additionally, in the most apical zone, the apical-to-basal gradient appears to drop slightly compared to that of the adjacent zones. This may be reflective of a transition to the Zone 1 of West's model. Therefore, the evidence suggests that the third contribution to the apical-to-basal distribution of perfusion in the sitting lung may be hydrostatic pressure gradients, described by West's zonal model[2].

It appears that the apical-to-basal distribution of perfusion observed in the sitting lung can be explained by a combination of contributions, each presented in previous studies. These include the gravitational determinants: tissue compression and hydrostatic effects, and a physiological determinant: an apparent hilum-to-peripheral gradient. These contributions are summarised in Figure 5.59. However, this hilum-to-peripheral gradient may be an artifact of the blood in the pulmonary vessels picked up by the MR acquisition. All blood enters the lung via the pulmonary arteries which themselves enter the lung through the hilum[6]. This is then distributed to all parts of the organ via a dichotomously branching fractal network of pulmonary vessels. On the return, the oxygenated blood exits the lung via the pulmonary vessels, which also pass through the hilum. Therefore, as all blood within the lung must transit this single point of entry and exit, it is intuitive to imagine this will result in a concentration of blood volume towards this region. However, this is not necessarily representative of the alveolar perfusion in these areas. Beyond the segmentation of the major pulmonary vessels, this study made no distinction between the blood transiting the pulmonary vessels and the blood contributing to gas exchange in the alveolar capillaries. Therefore, this may contribute to the observed centralisation of perfusion in this work, and thus the peripheral distribution shown in Figure 5.46B may provide a more accurate representation of alveolar perfusion within the lung.

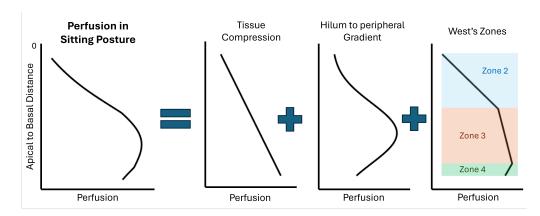


Figure 5.59: Modelled contributions to the apical-to-basal distribution of perfusion in the lung while in the sitting posture, including gravitational contributions from tissue compression and hydrostatic pressure variation, and an apparent hilum to peripheral gradient. Note, the basally shifted Gaussian employed to model the hilum to peripheral gradient is similar to that used to explain the tissue compression in the sitting and supine postures due to a centralisation of mass.

When investigating the anterior and posterior segments of the sitting lung independently, each returns a similar distribution to that of the coronal slice. Indeed, all apical-to-basal slopes of the volume-normalised perfusion are within the margin of error of each other. Each slope is reduced when normalised by the tissue density and remain within the margin of error of each other. These findings suggest that the coronal slice at the posterior of the sitting lung is representative of the lung as a whole.

The distribution of perfusion in the head-down posture was investigated in only the coronal plane. The apical-to-basal volume-normalised variation within the whole slice is plotted in Figure 5.45A. In this plot, there appears to be a reduced variation of perfusion in this posture compared to the lung in the sitting posture, with the head-down posture exhibiting a largely uniform distribution in the central regions, although dropping slightly in the most apical zones and significantly in the gravitationally dependent basal regions. This resembles the variation of perfusion in this posture previously presented in work by Ax et al.[4] (Figure 5.4). Thus, at first glance, the

distribution of perfusion appears to increase along the gravitational vector in agreement with established theories of lung perfusion. However, the magnitude of this change is reduced compared to the sitting posture, and like the sitting posture, the perfusion distribution is non-linear.

As for the sitting posture, the data were analysed in a number of ways to isolate contributions to the perfusion distribution. First, the perfusion variation was normalised by the tissue density variation, plotted in Figure 5.45B. This distribution appears slightly more uniform in the central regions, although otherwise remains largely consistent in shape to that of the volume-normalised distribution. This is consistent with the more uniform distribution of tissue density in the head-down posture compared to sitting, as demonstrated previously, which results in a greater similarity between normalisation methods in this posture. To remove any hilum-to-peripheral gradients, the perfusion distribution was investigated in only the peripheral regions. This is plotted for both the volume-normalised perfusion in Figure 5.46A and tissue-density-normalised perfusion in Figure 5.46B. If one assumes that the head-down apical-to-basal distribution of perfusion is the product of tissue compression, central gradients and hydrostatic forces, and that the prior two are removed by tissue-density-normalisation and peripheral segmentation, then the remaining distribution is due to the hydrostatic gradients. However, this does not match the expected distribution of perfusion due to hydrostatic forces described by West's zonal model. Indeed, this is evident as the shape of the distribution does not reflect that of the gravitationally opposite sitting posture. Instead, the perfusion distribution is largely uniform along this axis, although it increased slightly in the apical regions before dropping again in the most apical zone.

These observations may be explained by the gravitational positioning of the heart relative to the lung. When subjects are sitting upright, the heart is positioned at an approximately gravitationally equivalent location to the base of the lung. Systolic pressure must overcome the hydrostatic pressure required to pump blood to the apical regions. Then, after passing through the alveolar capillaries, the blood in these regions returns to the heart with the assistance of gravity. However, in the head-down posture the lung positioning is inverted. The heart, still located at a gravitationally equivalent location to the base of the lung, is now at a superior position to much of the lung. In this orientation, the systolic pressure required to deliver blood to the lung is reduced, in conjunction with the reduced hydrostatic forces to overcome. This may contribute to the more uniform distribution of perfusion associated with the head-down posture compared to the sitting posture, highlighted in the tissue-density-normalised peripheral apical-to-basal distribution (Figure 5.46B). However, after transiting the alveolar capillaries, the blood in the apical regions must work against gravity, overcoming the hydrostatic pressure required to reach the heart. Analogous to the pooling blood in the head when in the head-down posture, this could result in an increased perfusion within the apical regions of the lung. This may contribute to the increased perfusion in the apical regions in Figure 5.46B. Finally, the decreased perfusion associated with the most apical zone may be due to the weight of the lung tissue above, compressing the vessels and reducing blood flow. This is equivalent to the Zone 4 of West's zonal model.

The isogravitational apical-to-basal variation of perfusion in the supine posture was investigated in the coronal and the sagittal plane. In the coronal plane (Figure 5.45), the supine lung exhibited a significant increase in blood flow towards the base. This perfusion appears to increase fastest in the most apical zones, the rate of this increase is reduced in the central zones, before the perfusion drops slightly in the most basal zones. The

magnitude of this variation is almost equivalent to that associated with the sitting posture. Additionally, the perfusion distribution resembles in shape an amplified version of the associated distribution of tissue density, with two distinct peaks at approximately 35% and 70% of the apical-to-basal distance. This distribution of perfusion is unexpected, as conventional gravitational models of perfusion dictate a uniform distribution within isogravitational planes.

The distribution of perfusion observed in the coronal plane may in part be due to the variation of tissue compression at the posterior of the lung. This would have the effect of an increasing perfusion along the apical-to-basal axis, due to the greater density of vessels. Indeed, the shape of this distribution closely resembles that of the assumed tissue density, giving weight to this theory. Furthermore, considering the apical-to-basal variation of perfusion in the sagittal acquisition, the magnitude of the change along this axis is significantly reduced. When examining only the gravitationally dependent posterior segment, the distribution is similar to that of the coronal slice, while the anterior segment exhibits little variation along this axis. Therefore, as for the discussion of the tissue compression and FV, we conclude that the coronal slice located at the posterior is not representative of the lung as a whole.

When normalised by the estimated tissue density, the resultant perfusion distribution in the coronal plane of the supine lung exhibits a more uniform distribution in the central regions. This is expected with the conventional gravitational models of perfusion. However, in both the basal and, to a greater extent, the apical zones, the reduction in this function remains. This may be due to an apparent hilum-to-peripheral gradient as discussed previously, amplifying the apparent perfusion in the central regions. To abate the effect of this gradient, the perfusion variation in only the most

peripheral regions was plotted in Figure 5.46. In the volume-normalised plot, this yields an increasing trend towards the base, although with a reduced slope compared to the sitting posture. When normalised by the tissue density, as expected, this slope is reduced. However, this plot now resembles a distribution like that dictated by West's zones. Compared to the sitting posture, the magnitude of the variation along this axis is reduced and the zone 2-3 boundary shifted towards the apex. This is highly unexpected. Previously, this work has assumed that the remaining contribution to perfusion in this peripheral tissue-density-normalised plot is purely due to West's zones. However, West's zones are primarily dictated by variation in hydrostatic pressures, which, within isogravitational planes, should be uniform. Therefore, this may suggest an additional contribution to the distribution of perfusion within the lung.

The apical-to-basal perfusion variation in the prone lung was investigated only in the sagittal plane. Unlike for the sitting and supine postures, the peak volume-normalised perfusion in the prone lung appears to be centralised. Following the discussion of the apparently uniform distribution of tissue density within the lung, this result is not unexpected. Without a variation in tissue compression along this axis, the effect of the apparent hilumto-peripheral gradient is readily apparent. This is further demonstrated as the associated tissue-density-normalised plot appears near-identical in shape. When isolating the variation in the gravitationally dependent anterior segment, the central bias remains, although the perfusion appears to drop to a greater degree in the basal region. This may be associated with the previously discussed theorised redistribution of parenchyma from the posterior towards the diaphragm, although this reduction is also observed in the associated tissue-density-normalised distribution. Therefore, there may be an additional contribution to the perfusion distribution. Due to the

shape of the diaphragm in this posture, the most basal zone of the anterior segment will often contain the fewest voxels. This results in a greater error in this datapoint, and thus more investigation is needed before meaningful analysis can be made of this result. In the posterior segment, the distribution of perfusion appears approximately uniform along this axis for both normalisation methods. Notably, the central bias is significantly reduced. This may be because the central lung vessels have shifted with the gravitational vector towards the anterior of the lung, thus reducing the apparent hilum-to-peripheral gradients in this segment.

For all investigations which exhibited a significant positive slope in perfusion along the apical-to-basal axis, the magnitude of this change was reduced when normalising by the assumed tissue density. This includes all analyses of the sitting lung and analyses of the posterior of the supine lung. Intuitively, these also correspond to investigations where the tissue density increased significantly along this axis, which agrees with work from Hopkins et al. [1] that tissue density is a primary contributor to observed perfusion in the lung. All isogravitational investigations which exhibited an approximately uniform distribution of tissue density also exhibited a comparatively small positive slope of perfusion along this axis. However, when each of these variations was normalised by the estimated tissue density, the magnitude of the increase of perfusion along this axis increased. This may suggest an additional non-gravitational contribution to the distribution of perfusion along these axes, increasing blood flow towards the base of the lung. However, each of these variations is small compared to the associated error and therefore more investigation is needed before meaningful conclusions can be drawn from this result.

Multiple apical-to-basal investigations of isogravitational perfusion variation in this study have suggested there may be an additional non-gravitational contribution to the distribution along this axis. This is due to the unexpected increase in blood flow towards the base of the lung. While this has been observed in all investigations in the supine and prone lungs, this contribution may also be present in the sitting and head down lung. However, when the participants were in these postures, it was not possible to isolate these contributions from those due to gravity. Glenny et al. have published numerous papers on non-gravitationally dependent contributions to perfusion distribution[3, 43, 90, 92]. These studies attribute some perfusion variation to the structure of the pulmonary vessels and inequalities at each bifurcation. However, the contribution of this effect remains an area of active discussion. Furthermore, this effect is normally used to explain localised correlation of blood flow, not a systematic increase towards the basal regions as observed in this work. The structure of the pulmonary arteries may favour perfusion towards the basal regions. This could be further amplified in the posterior of the supine lung due to the distribution of lung tissue, and by extension, blood vessels towards this region when in this posture. However, this has not been previously reported in related literature.

Finally, when comparing the perfusion distribution between participants breathing normally and deeply, all postures exhibit a similar perfusion distribution between breathing regimes. However, each distribution appears slightly shifted towards the apex, both when normalised by the mean lung value and by the estimated tissue density. This results in a slightly more uniform distribution of perfusion when in the sitting and supine posture, although increases the variation in the head-down posture. This result was unexpected as the redistribution does not correlate to the gravitational vector. One might expect an amplification in the apical-to-basal variation of perfusion along the gravitational vector, due to greater lung size and thus

increased hydrostatic pressure variation associated with breathing deeply. However, this may be an artifact of the distribution analysis algorithm that sections the lung into zones. Each of these zones is assigned a fractional apical-to-basal distance, scaled between 0 and 100%, irrespective of the size of the lung. Breathing deeply is generally associated with a greater actuation of the diaphragm and larger peak lung volume. Therefore, the intermediary frame used to generate each of the functional maps is likely larger than that associated with breathing normally. If the distribution of perfusion fails to move proportionally with the movement of the diaphragm, then this will result in an artificial distribution towards the apical regions. Imagine a point located at the centre of a reconstructed normal-breathing functional map. This is assigned to a position, 50% of the apical-to-basal distance. If the greater actuation of the diaphragm when breathing deeply results in a reconstructed functional map with an apical-to-basal distance 1.25x that of normal breathing, then this imaginary point in space will now be assigned to 40% of this distance. Any perfusion in this location will have appeared to have shifted towards the apex. Therefore, we conclude that there is little variation in perfusion within the lung between the normal and deep breathing regimes. This result is analogous to that observed in the supine lung, showing that lung volume has little effect on the anterior-to-posterior distribution of perfusion[122].

5.4.3.2 Perfusion Time-to-Peak

The QTTP variation was investigated along the apical-to-basal axis in the coronal plane and along the apical-to-basal and anterior-to-posterior axes in the sagittal plane. Each of these planes employed a different reference for determining the cardiac frequency: the aorta in the coronal plane and the central lung vessels in the sagittal plane. Therefore, because the aorta is

expected to reach peak perfusion before the lung vessels, the QTTP values will be earlier in the sagittal investigations.

As for the VTTP distributions, this function is based on differences of phase with respect to the perfusion cycle, not signal amplitude. Therefore, assuming minimal anterior-to-posterior movement of lung parenchyma, it is expected that the B1 (receive) inhomogeneities had a minimal effect on the QTTP results.

The resting heart rate of a normal adult is between 60 and 100 bpm. However, the sampling frequency employed in this study resulted in a Nyquist frequency of 82 bpm. Therefore, it is likely that for multiple participants in this study, the cardiac cycle was under-sampled. Chapter 4 includes a detailed explanation of how the PREFUL processing pipeline accounts for this.

5.4.3.2.1 Anterior-to-Posterior

The anterior-to-posterior variation in the QTTP was investigated in the sitting, supine, and prone postures. In both the supine and prone postures, this function appears approximately uniform along this axis, reaching peak perfusion at approximately the same point of the perfusion cycle as the central lung vessels. There may be a slight trend that the gravitationally dependent regions of each posture reach peak perfusion first, although this variation is small compared to the associated error. If significant, this observation may be due to gravitational effects, namely variation in hydrostatic pressure. The greater pressure in the gravitationally dependent regions may explain why they might reach peak perfusion first.

Contrasting these results with the sitting posture, all regions of the lung

appear to reach peak perfusion significantly later than the other two postures. There is a significant variation in this result between participants, signified by the large error bars. This comparatively delayed peak perfusion may be due to the positioning of the heart at a position gravitationally equivalent to that of the base of the lung. Compared to the other postures, the heart must overcome a greater hydrostatic pressure to perfuse the entire lung. This may result in delayed peak perfusion times when in this posture. Additionally, the anterior regions appear to reach peak perfusion later than those located at the posterior of the lung. If significant, this anterior-posterior difference may be due to the positioning of the heart toward the anterior of the thoracic cavity. The momentum of blood travelling from this anterior position through the centrally located hilum may result in a favouring of perfusion towards the posterior regions. To reach the anterior of the lung, this blood must effectively change direction by 180 degrees, which may result in later peak perfusion times. This hypothesis requires further investigation into pulmonary artery blood flow in the sitting posture.

5.4.3.2.2 Apical-to-Basal

The apical-to-basal variation of QTTP was investigated along the apical-to-basal direction in the coronal investigations for the sitting, supine, and head down postures, and in the sagittal investigations for the sitting, supine, and prone postures. When contrasting the QTTP distribution of the repeated postures, the systematically delayed perfusion in the coronal results is evident.

As for the anterior-to-posterior investigation, all zones in the sitting posture reach peak perfusion significantly later than the other postures investigated. Again, this may be due to the basal positioning of the heart and the greater hydrostatic pressures to overcome to perfuse the entire lung. This may delay peak perfusion times in this posture. However, similar to the anterior regions, the apical regions appear to reach peak perfusion significantly later than those at the posterior. Additionally, the central regions appear to reach peak perfusion earlier than the basal regions. One explanation for these observations is that the blood enters the lung at the centre, and favours filling the gravitationally dependent lung regions over the superior apical regions. Contrasting this result to that of the headdown posture, the time-to-peak perfusion of the whole lung appears earlier than that of the other postures. This may be because in this posture, the heart is located at a gravitationally superior position to most of the lung. Therefore, opposite to the sitting posture, the heart does not need to overcome significant hydrostatic pressures to deliver blood to the entire organ. This may result in observed earlier QTTP values. Additionally, the QTTP appears to increase with apical to basal distance, with potentially a reduction in the central regions. This may be explained using the same gravitational mechanism as in the sitting lung. The blood enters the organ via the central lung vessels, which then favours filling the now dependent apical regions before the superior basal regions.

The supine posture exhibits QTTP values generally between those of the sitting and head down postures. Furthermore, in the sagittal investigation, the QTTP values appear approximately uniform and equivalent to that of the central lung vessels. This would be expected if this function is determined by the hydrostatic pressure as compared to the other two postures, there is no significant difference in filling favourability along this isogravitational plane. In the coronal investigation, although more uniform along this axis compared to the other postures, the earliest time to peak values

appear to be in the centre of the lung. This increases steadily towards the basal regions, potentially reflective of the transit time of blood through the organ. The increase in the apical regions is steeper than that observed in the basal regions, resulting in the latest peak perfusion times in this area. As for the basal regions, a delayed peak perfusion may be expected in this area due to the transit time of the blood. However, it is not immediately obvious why this deviates so greatly from the values in the basal regions. The blood may favour transit to the basal regions, which could be related to the greater than expected normalised perfusion values in this area. However, as this is not observed in the sagittal investigation, it may be due to an error in the PREFUL processing and associated cardiac phase assignment. Therefore, more investigation is required before meaningful hypotheses can be made from this observation.

Finally, the sagittally investigated apical-to-basal variation in the prone posture appears very similar to that of the supine posture in this plane. As with the supine posture, there is expected to be very little variation in hydrostatic pressure along this isogravitational plane. Therefore, the observed QTTP distribution matches that which would be expected if the distribution were dictated by this mechanism.

When comparing the QTTP associated with deep breathing, all postures appear to reach peak perfusion slightly later. This may be associated with the increased lung volumes associated with deep breathing and thus it may take longer for the blood to reach the peripheral regions. This may delay the time-to-peak perfusion in all regions of the larger organ.

In culmination, these results suggest that gravity has a dominant effect on time taken to reach peak perfusion in the lung, with dependent regions reaching peak perfusion first. However, additional to this effect, the spatial positioning of the heart relative to the lung combined with the structure of the pulmonary arteries also appears to create spatial inhomogeneities in QTTP results. Furthermore, breathing deeply appears to delay the time taken to reach peak perfusion, potentially due to the associated greater lung volumes. However, due to the potential aliasing errors discussed previously, these results should be corroborated with further studies.

5.4.4 Ventilation and Perfusion Matching

This study did not directly investigate the ratio of ventilation and perfusion within the lung. In its current state, this PREFUL processing pipeline does not return absolute quantities of air and blood variation, instead outputting FV and perfusion weighted maps. Therefore, this work can only contrast the relative change of each of these normalised functions. Nevertheless, observed variation in the FV generally correlates to that of the perfusion for all postures. In the sitting posture, all slopes of relative normal-breathing FV were within the margin of error of the normalised perfusion variation, although the variation of tissue-density-normalised perfusion was reduced. In the head-down posture, the magnitude of the normalised perfusion variation was greater than that of the FV. However, when the perfusion was normalised by the tissue density, these were within the margin of error. A similar trend was observed for all investigations of the supine lung, although not in the anterior segment. Finally, in the prone lung, the slopes of all normalised perfusion distributions were reduced compared to the FV, although each of these was within the margin of error. Furthermore, the shape of each of these distributions is very similar. Therefore, we conclude that in all postures investigated in this study, the relative FV and perfusion distributions remain well matched. This is further confirmed by the toleration of each posture by each of the participants.

After accounting for established gravitational contributions, the apical-tobasal distribution of both the FV and perfusion demonstrated a bias towards the basal regions. This bias appears to create a slope of approximately 20% of the average value of the lung along the length of this axis. The persistence of this bias within isogravitational planes suggests that each of these contributions may be physiological in nature. However, neither the bias in FV nor in perfusion can be explained by established theory. This work proposes the mechanical contribution of the diaphragm as a mechanism for this observed ventilation bias. However, the corresponding mechanism for the perfusion distribution is not immediately obvious. Similar to hypothesis from Glenny et al.[3], the pulmonary arteries may be structured in such a way as to distribute a greater blood flow to the basal regions. Alternatively, the perfusion may share a mechanical contribution with that of the ventilation, thus passively accounting for any functional inhomogeneities. This apparent ventilation-perfusion matching may result from an active mechanism such as HPV. However, this functional variation is relatively small compared to the gravitational contributions, although HPV has been previously observed in healthy lungs[123]. Finally, these gradients may be a shared artefact of the specific imaging and processing methods employed. Therefore, these functional experiments should be repeated in the future with additional imaging modalities in order to corroborate the above observations.

5.4.5 Sources of Error

5.4.5.0.1 Single Slice

This research is limited as it is based on only single slice acquisitions. Therefore, this analysis is based on only a two-dimensional snapshot of a three-dimensional system. Beyond the pitfalls inherent in single slice PREFUL, such as lung parenchyma moving in and out of the imaging plane, slices at a similar position within the thoracic cavity may sample different lung tissue across each of the postures. This is particularly evident in the coronal investigation of the supine posture, where the posteriorly located slice experienced a significantly greater degree of parenchymal compression. For a complete picture of the redistribution of lung function due to varying the relative direction of gravity, one should employ a 3D imaging modality. 3D PREFUL MRI for ventilation investigations has been demonstrated at 1.5T[124]. This technique utilises a continuously acquired stack-of-stars 3D gradient echo acquisition, followed by the retrospective binning of each radial spoke based on the determined respiratory phase. However, this has yet to be demonstrated at 0.5T and thus does not yet facilitate application with low-field Open-MRI. Furthermore, to date an equivalent method for perfusion processing has not been reported, due to the need for registration of each image based on the respiratory phase. A compromise may be found in multi-slice acquisitions, if one can accurately determine the respiratory and cardiac phases for each slice. This could be achieved using external measures of these variations instead of relying purely on observations within each image. However, these results should be repeated with contrast-based techniques for measuring three-dimensional variations in ventilation and perfusion.

5.4.5.0.2 Fractional Ventilation vs Absolute Ventilation

The use of FV as a metric in this study is likely to have altered the observed variation in ventilation distribution within the lung. One should not conflate absolute ventilation with FV. By definition, this metric is fractional; any result is relative to the signal intensity in the exhaled lung. Specifically, the observed changes in signal intensity in the exhaled lung state will impact this metric, even when the absolute change in air volume remains constant across regions. If tissue density is uniform throughout the organ, then the FV may serve a better representation of ventilation within each posture. However, observations in this study suggest that this is not the case. For example, in the sitting posture, a significant increase in signal intensity was observed along the apical-to-basal axis. This will have resulted in reduced values of the FV in the basal regions compared to the apex. Therefore, the true variation in lung ventilation along the apical-to-basal axis in this posture may be greater than the observed variation in FV.

5.4.6 Future Work

Beyond repeating these results with additional imaging modalities, future work can generally be split into two camps: further investigating the mechanisms proposes in this work, or expanding the research to different subject groups.

Hypoxic Pulmonary Vasoconstriction (HPV) is thought to occur in all humans, although the extent of this is unclear [96]. Previous studies have employed MRI to map the pulmonary vascular response to hyperoxia in supine participants and determined that HPV optimises ventilation-perfusion matching in healthy individuals [123]. This study observed significant variation in

the FV along the apical-to-basal axis for subjects sitting upright, which was significantly greater than the anterior-to-posterior variation when supine. The resultant drop in ventilation in the apical regions may be sufficient to observe an HPV response. Therefore, we hypothesise that maps of the normalised perfusion will exhibit variation in the apical blood flow when the lung is exposed to normoxic and hyperoxic conditions.

This work discusses variation in the shape and volume of the thoracic cavity between postures, including the differing action of the diaphragm, intercostal and accessory muscles. These variations could be investigated using breath-hold three-dimensional scans in different postures.

All participants in this study were relatively young. This was to ensure their safety while adopting the head-down posture. Future studies should expand the cohort to include a greater age range, even if the head-down posture is omitted. Older subjects may exhibit differing degrees of tissue compression and functional variation between postures. Additionally, investigations of this kind with paediatric subjects may give insight into the gravitational contributions to lung development.

This methodology could be used to investigate participants who regularly experience different gravitational conditions in order to measure if their lungs have adapted to these environments. Most simply, this could include individuals who frequently adopt a head-down posture during activities such as yoga or gymnastics. It could also include individuals who often experience greater gravitational forces, such as RAF pilots. Finally, this methodology could be applied to individuals who have recently spent a prolonged period in microgravity, such as astronauts returning from the space stations.

This study investigated participants in four postures in order to orientate

the lungs such that they experience gravitationally opposite conditions. Whilst intermediary positions should also be investigated, this study only examined the lungs under conditions of 1g. To date, no modern imaging studies have probed the distribution of human lung function in reduced gravitational conditions[61]. Technology permitting, this methodology could be repeated during periods of temporary weightlessness with an MRI scanner on a parabolic flight. Looking further into the future, when we possess the technology to place an MRI in space, these investigations could be repeated under conditions of sustained weightlessness. These investigations should give a direct measure of the immediate gravitational influences on pulmonary function. However, the gravitational effect on lung development should also be considered.

Finally, with improved insight into the gravitational dependence of the distribution of healthy lung function, studies can begin into the effect of disease. Unexpected redistributions of tissue, ventilation and perfusion between postures may prove to be a new marker for disease.

5.5 Conclusions

This study aimed to investigate the gravitational dependence of the distribution of pulmonary function in healthy individuals. To achieve this objective, 23 participants were examined in four postures – sitting, head-down, supine, and prone – which systematically varied the direction of the gravitational vector relative to their lungs. These experiments were facilitated by a 0.5T Open MRI system in conjunction with PREFUL[5] post-processing to generate functional maps of the exhaled tissue density, and both ventilation and perfusion dynamics.

Established gravitational modes[1, 31] appear to provide a useful first-order approximation for describing the distribution of tissue density in the exhaled lung, with compression increasing along the gravitational vector. However, our findings demonstrate that this alone is insufficient to fully explain the observed variation of parenchyma compression. Non-linear distributions of tissue density, particularly in the sitting and supine lungs, suggest a centralisation of mass within the organ. This may be due to a greater density of pulmonary vessels in this region. Furthermore, non-opposite slopes of tissue density variation were observed for gravitationally opposite postures, deviating from the redistribution of parenchyma predicted by the simple gravitational model. This observation along the apical-to-basal axis of the sitting and head-down postures was hypothesised to be due to structural contributions from the primary pulmonary vessels and early bronchi.

In the supine posture, the non-uniform distribution of tissue density along the isogravitational apical-to-basal plane appeared to contradict the Slinky effect model[1]. However, this contradiction was likely attributable to the non-uniform weight of the lung acting perpendicular to this plane. Consequently, special consideration should be made when comparing apparently similar coronal slices between postures, as lung parenchyma may occupy different regions of the thoracic cavity or experience differing loading conditions. In contrast, when examining the gravitationally opposite prone posture, this effect was not observed. We hypothesise that this difference results from diaphragmatic displacement occurring in line with the gravitational vector, due to the weight of the abdominal organs. Such displacement may have produced a reduction in basal-posterior pleural pressure, thereby redistributing lung parenchyma to this region.

Gravitational tissue compression was expected to be a primary determinant

of the distribution of ventilation in the lung. While this mechanism indeed appeared to be a primary contributor, when tidal breathing, all postures exhibited a bias towards the basal regions irrespective of gravitational orientation. This bias was evident even in isogravitational planes, suggesting additional contributions from non-gravitational physiological mechanisms. We hypothesise that this phenomenon is attributed to mechanical factors, such as the actuation of the diaphragm. The work required to displace lung tissue decreases with distance from the actuator. Therefore, alveoli nearest the diaphragm may expand to a greater degree until their closure force surpasses this potential. This phenomenon is particularly evident in the head-down posture, where this mechanism counters the observed distribution of tissue compression, resulting in a near-uniform distribution of FV. When breathing deeply, the ventilation distribution appeared to better reflect the gravitational distribution of tissue density, likely due to the more uniform expansion opportunity associated with this regime.

Investigation of the Jacobian determinant showed that all postures demonstrated varying contributions from the diaphragm, intercostal, and accessory muscles during respiration. This may be due to the differences in the contribution from gravity, but also to variation in participant arm positioning. This was particularly apparent for participants in the head-down posture where respiration appeared highly dependent on the diaphragm[120]. Additionally, systematic variation in the VTTP can be generally attributed to gravitationally induced variations in airway resistance.

The distribution of perfusion in the lung was demonstrated to be primarily dependent on gravitational effects. Namely, the distribution of lung parenchyma and hydrostatic pressure variation. An apparent hilum-to-peripheral gradient was likely an artifact of blood in the pulmonary vessels and not reflective of alveolar perfusion. In the sitting posture, the alveolar

perfusion variation attributed to hydrostatic effects reflected the pattern described in West's zonal model[2]. However, the corresponding distribution for the head-down posture suggested a venous pooling of blood in the apical regions, potentially due to the gravitationally superior position of the heart in this posture. After accounting for tissue density variation, alveolar perfusion variation within isogravitational planes exhibited a slight bias towards the basal regions that cannot be explained by variations in hydrostatic pressure. This may instead be due to an active mechanism of ventilation-perfusion matching, or a consequence of the shape of the vascular tree[90]. Lung volume appeared to have little effect on the distribution of perfusion.

Distributions of the QTTP appear to be primarily dependent on gravitational orientation, with dependent regions reaching peak perfusion first. Additionally, this function also demonstrated a significant influence from the position of the heart relative to the lung. Although breathing deeply appeared to delay the QTTP across the whole organ, potentially due to the greater hydrostatic pressure variation associated with the larger lung.

In conclusion, established gravitational models of the lung appear to describe well the distribution of ventilation and perfusion under normal conditions. However, additional physiological factors may also influence these distributions, with their contributions more evident in non-standard gravitational conditions. Importantly, ventilation and perfusion remained well matched across all postures. Nevertheless, both exhibited a bias towards the basal regions that cannot be attributed to established models of systemic functional distribution in the lung. Further investigation employing different imaging modalities is required to corroborate these results.

Continued advancement in the understanding of pulmonary functional dis-

tribution and the underlying mechanisms will provide more precise respiratory care interventions, ultimately improving patient outcomes through better-informed clinical practices and treatment modalities. Furthermore, deviation from improved models of functional redistribution between postures in healthy individuals may yield novel biomarkers for pathophysiology.

5.6 References

- Hopkins, S. R., Henderson, A. C., Levin, D. L., Yamada, K., Arai, T., Buxton, R. B. & Prisk, G. K. Vertical gradients in regional lung density and perfusion in the supine human lung: the Slinky effect.
 Journal of applied physiology (Bethesda, Md.: 1985) 103, 240–248.
 doi:10.1152/JAPPLPHYSIOL.01289.2006 (July 2007).
- West, J. B. Regional differences in the lung. Chest 74, 426. doi:10.
 1378/CHEST.74.4.426 (Oct. 1978).
- Glenny, R. W. Determinants of regional ventilation and blood flow in the lung. *Intensive Care Medicine* 35, 1833–1842. doi:10.1007/ S00134-009-1649-3/FIGURES/6 (Nov. 2009).
- Ax, M., Sanchez-Crespo, A., Lindahl, S. G., Mure, M. & Petersson,
 J. The influence of gravity on regional lung blood flow in humans:
 SPECT in the upright and head-down posture. *Journal of applied physiology (Bethesda, Md. : 1985)* 122, 1445–1451. doi:10.1152/
 JAPPLPHYSIOL.00887.2015 (June 2017).
- Voskrebenzev, A., Gutberlet, M., Klimeš, F., Kaireit, T. F., Schönfeld, C., Rotärmel, A., Wacker, F. & Vogel-Claussen, J. Feasibility of quantitative regional ventilation and perfusion mapping with phase-resolved functional lung (PREFUL) MRI in healthy volunteers and COPD, CTEPH, and CF patients. Magnetic Resonance in Medicine 79, 2306–2314. doi:10.1002/mrm.26893 (2018).
- 6. West, J. B. (B. West's respiratory physiology: the essentials eng (ed Luks, A.) Philadelphia, 2021.
- 7. Petersson, J. et al. Regional lung blood flow and ventilation in upright humans studied with quantitative SPECT. Respiratory physiol-

- ogy & neurobiology **166**, 54–60. doi:10.1016/J.RESP.2009.01.008 (Mar. 2009).
- 8. West, J. B. (B. West's pulmonary pathophysiology: the essentials eng (ed Luks, A.) Philadelphia, 2022.
- Cotes, J. E., Maynard, R. L., Pearce, S. J., Nemery, B. B., Wagner, P. D. & Cooper, B. G. Lung Function 7th edition. eng (Wiley, Newark, 2020).
- 10. Verbanck, S., Paiva, M., Schuermans, D., Hanon, S., Vincken, W. & Van Muylem, A. Relationships between the lung clearance index and conductive and acinar ventilation heterogeneity. *Journal of applied physiology (Bethesda, Md. : 1985)* 112, 782–790. doi:10.1152/JAPPLPHYSIOL.01221.2011 (Mar. 2012).
- Venegas, J. G. et al. Self-organized patchiness in asthma as a prelude to catastrophic shifts. Nature 434, 777–782. doi:10.1038/ NATURE03490 (Apr. 2005).
- 12. Douglas, G. Asthma eng (ed Elward, K. S.) (Manson, London, 2011).
- 13. Kon, O. M. *Chronic Obstructive Pulmonary Disease (COPD)*. eng (eds Hansel, T. T. & Barnes, P. J.) (Oxford University Press, 2008).
- Stein, P. D. Pulmonary Embolism 3rd edition. eng (Wiley, Newark, 2016).
- 15. Darquenne, C., Fleming, J. S., Katz, I., Martin, A. R., Schroeter, J., Usmani, O. S., Venegas, J. & Schmid, O. Bridging the Gap Between Science and Clinical Efficacy: Physiology, Imaging, and Modeling of Aerosols in the Lung. *Journal of Aerosol Medicine and Pulmonary Drug Delivery* 29, 107. doi:10.1089/JAMP.2015.1270 (Apr. 2016).

- Darquenne, C. Aerosol Deposition in Health and Disease. Journal of Aerosol Medicine and Pulmonary Drug Delivery 25, 140. doi:10. 1089/JAMP.2011.0916 (June 2012).
- Pelosi, P. et al. Personalized mechanical ventilation in acute respiratory distress syndrome. Critical care (London, England) 25. doi:10.
 1186/S13054-021-03686-3 (Dec. 2021).
- 18. Schultz, M. J., De Abreu, M. G. & Pelosi, P. Mechanical ventilation strategies for the surgical patient. *Current opinion in critical care*21, 351–357. doi:10.1097/MCC.0000000000000015 (July 2015).
- Arfin, T., Pillai, A. M., Mathew, N., Tirpude, A., Bang, R. & Mondal, P. An overview of atmospheric aerosol and their effects on human health. *Environmental Science and Pollution Research 2023 30:60* 125347–125369. doi:10.1007/S11356-023-29652-W (Sept. 2023).
- Ferkol, T. & Schraufnagel, D. The global burden of respiratory disease. Annals of the American Thoracic Society 11, 404–406. doi:10. 1513/ANNALSATS.201311-405PS (2014).
- 21. Barnes, P. J., Bonini, S., Seeger, W., Belvisi, M. G., Ward, B. & Holmes, A. Barriers to new drug development in respiratory disease. *European Respiratory Journal* 45, 1197–1207. doi:10.1183/09031936.00007915 (Apr. 2015).
- 22. MARTIN, C. J., CLINE, F. & MARSHALL, H. Lobar alveolar gas concentrations; effect of body position. *The Journal of clinical investigation* **32**, 617–621. doi:10.1172/JCI102772 (1953).
- 23. RAHN, H., SADOUL, P., FARHI, L. E. & SHAPIRO, J. Distribution of Ventilation and Perfusion in the Lobes of the Dog's Lung in the

- Supine and Erect Position. https://doi.org/10.1152/jappl.1956.8.4.417 8, 417–426. doi:10.1152/JAPPL.1956.8.4.417 (Jan. 1956).
- 24. Anthonisen, N. R. & Milic-Emili, J. Distribution of pulmonary perfusion in erect man. https://doi.org/10.1152/jappl.1966.21.3.760 21, 760-766. doi:10.1152/JAPPL.1966.21.3.760 (1966).
- 25. BRYAN, A. C., BENTIVOGLIO, L. G., BEEREL, F., MACLEISH, H., ZIDULKA, A. & BATES, D. V. Factors affecting regional distribution of ventilation and perfusion in the lung. https://doi.org/10.1152/jappl.1964.19.
 19, 395–402. doi:10.1152/JAPPL.1964.19.3.395 (May 1964).
- BALL, W. C., STEWART, P. B., NEWSHAM, L. G. & BATES,
 D. V. Regional pulmonary function studied with xenon 133. The Journal of clinical investigation 41, 519–531. doi:10.1172/JCI104505 (Mar. 1962).
- 27. Hughes, J. M., Glazier, J. B., Maloney, J. E. & West, J. B. Effect of lung volume on the distribution of pulmonary blood flow in man. Respiration physiology 4, 58–72. doi:10.1016/0034-5687(68)90007-8 (1968).
- 28. Amis, T. C., Jones, H. A., Rhodes, C. G., Heather, J. D. & Hughes, J. M. Regional distribution of pulmonary ventilation and perfusion in the conscious dog. *American Journal of Veterinary Research* 43, 1972–1977 (Nov. 1982).
- Bryan, A. C., Milic-Emili, J. & Pengelly, D. Effect of gravity on the distribution of pulmonary ventilation. *Journal of applied physiology* 778–784. doi:10.1152/JAPPL.1966.21.3.778 (1966).
- 30. Glaister, D. H. Distribution of pulmonary blood flow and ventilation during forward (plus Gx) acceleration. *Journal of applied physiology* **29**, 432–439. doi:10.1152/JAPPL.1970.29.4.432 (1970).

- 31. West, J. B. & Matthews, F. L. Stresses, strains, and surface pressures in the lung caused by its weight. *Journal of applied physiology* **32**, 332–345. doi:10.1152/JAPPL.1972.32.3.332; CTYPE: STRING: JOURNAL (1972).
- 32. WEST, J. B. Regional differences in gas exchange in the lung of erect man. https://doi.org/10.1152/jappl.1962.17.6.893 17, 893–898. doi:10.1152/JAPPL.1962.17.6.893 (Nov. 1962).
- 33. Reed, J. H. & Wood, E. H. Effect of body position on vertical distribution of pulmonary blood flow. *Journal of applied physiology* **28**, 303–311. doi:10.1152/JAPPL.1970.28.3.303 (1970).
- 34. Maeda, H. et al. Pulmonary blood flow distribution measured by radionuclide-computed tomography. Journal of applied physiology: respiratory, environmental and exercise physiology 54, 225–233. doi:10.1152/JAPPL.1983.54.1.225 (1983).
- 35. Beck, K. C. & Rehder, K. Differences in regional vascular conductances in isolated dog lungs. *Journal of applied physiology (Bethesda, Md. : 1985)* **61,** 530–538. doi:10.1152/JAPPL.1986.61.2.530 (1986).
- 36. Olson, L. E. & Rodarte, J. R. Regional differences in expansion in excised dog lung lobes. *Journal of applied physiology: respiratory, environmental and exercise physiology* **57**, 1710–1714. doi:10.1152/JAPPL.1984.57.6.1710 (1984).
- 37. Sprung, J., Deschamps, C., Hubmayr, R. D., Walters, B. J. & Rodarte, J. R. In vivo regional diaphragm function in dogs. *Journal of applied physiology (Bethesda, Md. : 1985)* 67, 655–662. doi:10.1152/JAPPL.1989.67.2.655 (1989).

- 38. Hubmayr, R. D., Rodarte, J. R., Walters, B. J. & Tonelli, F. M. Regional ventilation during spontaneous breathing and mechanical ventilation in dogs. https://doi.org/10.1152/jappl.1987.63.6.2467 63, 2467–2475. doi:10.1152/JAPPL.1987.63.6.2467 (1987).
- 39. Robertson, H. T. & Hlastala, M. P. Microsphere maps of regional blood flow and regional ventilation. *Journal of applied physiology* (Bethesda, Md.: 1985) 102, 1265–1272. doi:10.1152/JAPPLPHYSIOL. 00756.2006 (Mar. 2007).
- Melsom, M. N., Kramer-Johansen, J., Flatebø, T., Müller, C. & Nicolaysen, G. Distribution of pulmonary ventilation and perfusion measured simultaneously in awake goats. *Acta physiologica Scandinavica* 159, 199–208. doi:10.1046/J.1365-201X.1997.92355000.X (1997).
- 41. MELSOM, M. N., FLATEBØ, T., KRAMER-JOHANSEN, J., AULIE, A., SJAASTAD, V., IVERSEN, P. O. & NICOLAYSEN, G. Both gravity and non-gravity dependent factors determine regional blood flow within the goat lung. *Acta physiologica Scandinavica* 153, 343–353. doi:10.1111/J.1748-1716.1995.TB09872.X (1995).
- 42. Altemeier, W. A., McKinney, S. & Glenny, R. W. Fractal nature of regional ventilation distribution. *Journal of applied physiology* (Bethesda, Md. : 1985) 88, 1551–1557. doi:10.1152/JAPPL.2000.88.5.1551 (2000).
- 43. Glenny, R. W., Lamm, W. J., Albert, R. K. & Robertson, H. T. Gravity is a minor determinant of pulmonary blood flow distribution. Journal of applied physiology (Bethesda, Md.: 1985) 71, 620–628. doi:10.1152/JAPPL.1991.71.2.620 (1991).
- 44. Altemeier, W. A., Robertson, H. T. & Glenny, R. W. Pulmonary gasexchange analysis by using simultaneous deposition of aerosolized and injected microspheres. *Journal of applied physiology (Bethesda*.

- *Md.* : 1985) **85**, 2344–2351. doi:10.1152/JAPPL.1998.85.6.2344 (1998).
- Glenny, R. W., Bernard, S., Robertson, H. T. & Hlastala, M. P. Gravity is an important but secondary determinant of regional pulmonary blood flow in upright primates. *Journal of applied physiology (Bethesda, Md. : 1985)* 86, 623–632. doi:10.1152/JAPPL.1999.86.
 2.623 (1999).
- 46. Glenny, R. W. et al. Selected contribution: redistribution of pulmonary perfusion during weightlessness and increased gravity. Journal of applied physiology (Bethesda, Md. : 1985) 89, 1239–1248. doi:10.1152/JAPPL.2000.89.3.1239 (2000).
- 47. Brudin, L. H., Rhodes, C. G., Valind, S. O., Jones, T. & Hughes, J. M. Interrelationships between regional blood flow, blood volume, and ventilation in supine humans. *Journal of applied physiology (Bethesda, Md. : 1985)* 76, 1205–1210. doi:10.1152/JAPPL.1994.76.3.1205 (1994).
- 48. Kosuda, S., Kobayashi, H. & Kusano, S. Change in regional pulmonary perfusion as a result of posture and lung volume assessed using technetium-99m macroaggregated albumin SPET. European journal of nuclear medicine 27, 529–535. doi:10.1007/S002590050539 (Apr. 2000).
- Treppo, S., Mijailovich, S. M. & Venegas, J. G. Contributions of pulmonary perfusion and ventilation to heterogeneity in V(A)/Q measured by PET. *Journal of applied physiology (Bethesda, Md. : 1985)* 1163–1176. doi:10.1152/JAPPL.1997.82.4.1163 (1997).
- 50. Marcucci, C., Nyhan, D. & Simon, B. A. Distribution of pulmonary ventilation using Xe-enhanced computed tomography in prone and

- supine dogs. Journal of applied physiology (Bethesda, Md. : 1985) 90, 421–430. doi:10.1152/JAPPL.2001.90.2.421 (2001).
- 51. Kreck, T. C. *et al.* Determination of regional ventilation and perfusion in the lung using xenon and computed tomography. *Journal of applied physiology (Bethesda, Md. : 1985)* **91,** 1741–1749. doi:10. 1152/JAPPL.2001.91.4.1741 (2001).
- 52. Jones, A. T., Hansell, D. M. & Evans, T. W. Pulmonary perfusion in supine and prone positions: an electron-beam computed tomography study. *Journal of applied physiology (Bethesda, Md. : 1985)* **90**, 1342–1348. doi:10.1152/JAPPL.2001.90.4.1342 (2001).
- 53. Lisbona, R., Dean, G. W. & Hakim, T. S. Observations with SPECT on the Normal Regional Distribution of Pulmonary Blood Flow in Gravity Independent Planes. *Journal of Nuclear Medicine* **28**, 1758–1762 (Nov. 1987).
- 54. Nicolaysen, G., Shepard, J., Onizuka, M., Tanita, T., Hattner, R. S. & Staub, N. C. No gravity-independent gradient of blood flow distribution in dog lung. *Journal of applied physiology (Bethesda, Md. : 1985)* 63, 540–545. doi:10.1152/JAPPL.1987.63.2.540 (1987).
- 55. Glenny, R. W., Bernard, S. L. & Robertson, H. T. Pulmonary blood flow remains fractal down to the level of gas exchange. *Journal of applied physiology (Bethesda, Md.: 1985)* 89, 742–748. doi:10.1152/JAPPL.2000.89.2.742 (2000).
- 56. Levin, D. L., Schiebler, M. L. & Hopkins, S. R. Physiology for the pulmonary functional imager. *European journal of radiology* **86**, 308–312. doi:10.1016/j.ejrad.2016.09.027 (2017).
- 57. Musch, G., Layfield, J. D. H., Harris, R. S., Vidal Melo, M. F., Winkler, T., Callahan, R. J., Fischman, A. J. & Venegas, J. G.

- Topographical distribution of pulmonary perfusion and ventilation, assessed by PET in supine and prone humans. *Journal of applied physiology (Bethesda, Md. : 1985)* **93,** 1841–1851. doi:10.1152/JAPPLPHYSIOL.00223.2002 (Nov. 2002).
- 58. Petersson, J. et al. Posture primarily affects lung tissue distribution with minor effect on blood flow and ventilation. Respiratory physiology & neurobiology 156, 293–303. doi:10.1016/J.RESP.2006.11.001 (June 2007).
- 59. Montmerle, S., Sundblad, P. & Linnarsson, D. Residual heterogeneity of intra- and interregional pulmonary perfusion in short-term microgravity. *Journal of applied physiology (Bethesda, Md. : 1985)* **98**, 2268–2277. doi:10.1152/JAPPLPHYSIOL.01268.2004 (June 2005).
- 60. Prisk, G. K., Guy, H. J., Elliott, A. R. & West, J. B. Inhomogeneity of pulmonary perfusion during sustained microgravity on SLS-1. *Journal of applied physiology (Bethesda, Md. : 1985)* 76, 1730–1738. doi:10.1152/JAPPL.1994.76.4.1730 (1994).
- 61. Prisk, G. K., Fine, J. M., Cooper, T. K. & West, J. B. Vital capacity, respiratory muscle strength, and pulmonary gas exchange during long-duration exposure to microgravity. *Journal of applied physiology* (Bethesda, Md. : 1985) 101, 439–447. doi:10.1152/JAPPLPHYSIOL. 01419.2005 (2006).
- 62. Cao, J. J. et al. Effects of respiratory cycle and body position on quantitative pulmonary perfusion by MRI. Journal of magnetic resonance imaging: JMRI 34, 225–230. doi:10.1002/JMRI.22527 (July 2011).
- 63. Fan, L., Liu, S. y., Xiao, X. s. & Sun, F. Demonstration of pulmonary perfusion heterogeneity induced by gravity and lung inflation using

- arterial spin labeling. European journal of radiology **73**, 249–254. doi:10.1016/J.EJRAD.2008.11.019 (Feb. 2010).
- 64. Keilholz, S. D., Knight-Scott, J., Christopher, J. M., Mai, V. M. & Berr, S. S. Gravity-dependent perfusion of the lung demonstrated with the FAIRER arterial spin tagging method. *Magnetic Resonance Imaging* 19, 929–935. doi:10.1016/S0730-725X(01)00416-7 (Sept. 2001).
- 65. Edelman, R. R., Hatabu, H., Tadamura, E., Li, W. & Prasad, P. V. Noninvasive assessment of regional ventilation in the human lung using oxygen-enhanced magnetic resonance imaging. *Nature medicine* 2, 1236–1239. doi:10.1038/NM1196-1236 (Nov. 1996).
- 66. Sá, R. C. et al. Vertical distribution of specific ventilation in normal supine humans measured by oxygen-enhanced proton MRI. Journal of Applied Physiology 109, 1950–1959. doi:10.1152/JAPPLPHYSIOL. 00220.2010 (Dec. 2010).
- 67. Geier, E. T., Theilmann, R. J., Darquenne, C., Prisk, G. K. & Sá, R. C. Quantitative Mapping of Specific Ventilation in the Human Lung using Proton Magnetic Resonance Imaging and Oxygen as a Contrast Agent. *Journal of visualized experiments : JoVE* **2019**, 10.3791/59579. doi:10.3791/59579 (June 2019).
- 68. Ouriadov, A. V., Fox, M. S., Couch, M. J., Li, T., Ball, I. K. & Albert, M. S. In vivo regional ventilation mapping using fluorinated gas MRI with an x-centric FGRE method. *Magnetic resonance in medicine* 74, 550–557. doi:10.1002/MRM.25406 (Aug. 2015).
- 69. Gutberlet, M. et al. Free-breathing Dynamic 19F Gas MR Imaging for Mapping of Regional Lung Ventilation in Patients with COPD. https://doi.org/10.1148/radiol.2017170591 286, 1040–1051. doi:10. 1148/RADIOL.2017170591 (Oct. 2017).

- 70. Kauczor, H. U. Hyperpolarized helium-3 gas magnetic resonance imaging of the lung. *Topics in magnetic resonance imaging : TMRI*14, 223–230. doi:10.1097/00002142-200306000-00002 (June 2003).
- 71. Fichele, S. et al. MRI of helium-3 gas in healthy lungs: Posture related variations of alveolar size. Journal of Magnetic Resonance Imaging 20, 331–335. doi:10.1002/JMRI.20104 (Aug. 2004).
- 72. Van Beek, E. J., Wild, J. M., Kauczor, H. U., Schreiber, W., Mugler, J. P. & De Lange, E. E. Functional MRI of the lung using hyperpolarized 3-helium gas. *Journal of magnetic resonance imaging : JMRI* 20, 540–554. doi:10.1002/JMRI.20154 (Oct. 2004).
- 73. Tsai, L. L., Mair, R. W., Li, C. H., Rosen, M. S., Patz, S. & Walsworth, R. L. Posture-dependent human 3He lung imaging in an open-access MRI system: initial results. *Academic radiology* 15, 728–739. doi:10.1016/J.ACRA.2007.10.010 (June 2008).
- 74. Mugler, J. P. et al. MR imaging and spectroscopy using hyperpolarized 129Xe gas: preliminary human results. Magnetic resonance in medicine 37, 809–815. doi:10.1002/MRM.1910370602 (1997).
- 75. Mugler, J. P. et al. Simultaneous magnetic resonance imaging of ventilation distribution and gas uptake in the human lung using hyperpolarized xenon-129. Proceedings of the National Academy of Sciences of the United States of America 107, 21707–21712. doi:10.1073/PNAS.1011912107/SUPPL{_}FILE/PNAS.201011912SI.PDF (Dec. 2010).
- 76. Qing, K. et al. Regional mapping of gas uptake by blood and tissue in the human lung using hyperpolarized xenon-129 MRI. Journal of Magnetic Resonance Imaging 39, 346–359. doi:10.1002/JMRI. 24181 (Feb. 2014).

- 77. Hughes, J. M. B. & West, J. B. Last word on Point:Counterpoint: Gravity is/is not the major factor determining the distribution of blood flow in the human lung. *Journal of applied physiology (Bethesda, Md. : 1985)* 104, 1539. doi:10.1152/JAPPLPHYSIOL.90374.2008 (May 2008).
- Maynard, R. L., Pearce, S., Nemery, B., Wagner, P. D., Cooper, B. & Cotes, J. E. Cotes' lung function eng (eds Maynard, R. L., Pearce, S. (J., Nemery, B., Wagner, P. D. (D., Cooper, B. (G. & Cotes, J. E.) Hoboken, NJ, 2020.
- Milic-Emili, J., Henderson, J. A., Dolovich, M. B., Trop, D. & Kaneko,
 K. Regional distribution of inspired gas in the lung. *Journal of applied physiology* 21, 749–759. doi:10.1152/JAPPL.1966.21.3.749
 (1966).
- 80. OTIS, A. B., MCKERROW, C. B., BARTLETT, R. A., MEAD, J., MCILROY, M. B., SELVER-STONE, N. J. & RADFORD, E. P. Mechanical factors in distribution of pulmonary ventilation. *Journal of applied physiology* 8, 427–443. doi:10.1152/JAPPL.1956.8.4.427 (Jan. 1956).
- 81. Reifenrath, R. The significance of alveolar geometry and surface tension in the respiratory mechanics of the lung. *Respiration Physiology* **24**, 115–137. doi:10.1016/0034-5687(75)90107-3 (July 1975).
- 82. Ardila, R., Horie, T. & Hildebrandt, J. Macroscopic isotropy of lung expansion. Respiration Physiology 20, 105–115. doi:10.1016/0034-5687(74)90100-5 (Mar. 1974).
- 83. Greenblatt, E. E., Butler, J. P., Venegas, J. G. & Winkler, T. Pendelluft in the bronchial tree. *Journal of Applied Physiology* **117**, 979. doi:10.1152/JAPPLPHYSIOL.00466.2014 (Nov. 2014).

- 84. Crawford, A. B., Dodd, D. & Engel, L. A. Changes in RIB cage shape during quiet breathing, hyperventilation and single inspirations. *Respiration Physiology* **54**, 197–209. doi:10 . 1016 / 0034 5687(83) 90057-9 (Nov. 1983).
- 85. Roussos, C. S., Martin, R. R. & Engel, L. A. Diaphragmatic contraction and the gradient of alveolar expansion in the lateral posture. Journal of applied physiology: respiratory, environmental and exercise physiology 43, 32–38. doi:10.1152/JAPPL.1977.43.1.32 (1977).
- 86. West, J. B. (B. *Pulmonary pathophysiology : the essentials* 8th ed. eng (ed West, J. B. (B.) (Wolters Kluwer/Lippincott Williams & Wilkins Health, Philadelphia, 2012).
- 87. Banister, J. & Torrance, R. W. The effects of the tracheal pressure upon flow: pressure relations in the vascular bed of isolated lungs. Quarterly journal of experimental physiology and cognate medical sciences 45, 352–367. doi:10.1113/EXPPHYSIOL.1960.SP001491 (Oct. 1960).
- 88. Hopkins, S. R., Wielpütz, M. O. & Kauczor, H. U. Imaging lung perfusion. *Journal of Applied Physiology* **113**, 328. doi:10.1152/JAPPLPHYSIOL.00320.2012 (July 2012).
- 89. West, G. B., Brown, J. H. & Enquist, B. J. The fourth dimension of life: fractal geometry and allometric scaling of organisms. *Science* (New York, N.Y.) **284**, 1677–1679. doi:10.1126/SCIENCE.284.5420.1677 (June 1999).
- 90. Glenny, R. & Robertson, H. T. Distribution of Perfusion. Comprehensive Physiology 1. doi:10.1002/cphy.c100012 (2011).

- 91. Glenny, R. W. & Robertson, H. T. Fractal modeling of pulmonary blood flow heterogeneity. *Journal of Applied Physiology* **70**, 1024–1030. doi:10.1152/JAPPL.1991.70.3.1024, (1991).
- Glenny, R. W. Spatial correlation of regional pulmonary perfusion.
 Journal of applied physiology (Bethesda, Md.: 1985) 72, 2378–2386.
 doi:10.1152/JAPPL.1992.72.6.2378 (1992).
- 93. Weibel, E. R., Richard Taylor, C. & Hoppeler, H. Variations in function and design: Testing symmorphosis in the respiratory system.

 *Respiration Physiology 87, 325–348. doi:10.1016/0034-5687(92)

 90015-0 (Mar. 1992).
- 94. Dunham-Snary, K. J., Wu, D., Sykes, E. A., Thakrar, A., Parlow, L. R., Mewburn, J. D., Parlow, J. L. & Archer, S. L. Hypoxic Pulmonary Vasoconstriction: From Molecular Mechanisms to Medicine. Chest 151, 181–192. doi:10.1016/J.CHEST.2016.09.001 (Jan. 2017).
- 95. Murray, T. R., Chen, L., Marshall, B. E. & Macarak, E. J. Hypoxic contraction of cultured pulmonary vascular smooth muscle cells. *American journal of respiratory cell and molecular biology* 3, 457–465. doi:10.1165/AJRCMB/3.5.457 (1990).
- Sylvester, J. T., Shimoda, L. A., Aaronson, P. I. & Ward, J. P. Hypoxic pulmonary vasoconstriction. *Physiological reviews* 92, 367–520. doi:10.1152/PHYSREV.00041.2010 (2012).
- Suzuki, H., Sato, Y., Shindo, M., Yoshioka, H., Mizutani, T., Onizuka, M. & Sakakibara, Y. Prone positioning improves distribution of pulmonary perfusion: noninvasive magnetic resonance imaging study in healthy humans. *European radiology* 18, 522–528. doi:10.1007/S00330-007-0808-6 (Mar. 2008).

- 98. Prisk, G. K., Yamada, K., Henderson, A. C., Arai, T. J., Levin, D. L., Buxton, R. B. & Hopkins, S. R. Pulmonary perfusion in the prone and supine postures in the normal human lung. *Journal of applied physiology (Bethesda, Md. : 1985)* 103, 883–894. doi:10.1152/JAPPLPHYSIOL.00292.2007 (Sept. 2007).
- 99. Bolar, D. S., Levin, D. L., Hopkins, S. R., Frank, L. F., Liu, T. T., Wong, E. G. & Buxton, R. B. Quantification of regional pulmonary blood flow using ASL-FAIRER. *Magnetic Resonance in Medicine* 55, 1308–1317. doi:10.1002/MRM.20891 (2006).
- 100. Hopkins, S. R. & Prisk, G. K. Lung perfusion measured using magnetic resonance imaging: New tools for physiological insights into the pulmonary circulation. *Journal of Magnetic Resonance Imaging* 32, 1287–1301. doi:10.1002/JMRI.22378 (Dec. 2010).
- 101. Pai S, A., Zhang, H., Street, J., Wilson, D. R., Brown, S. H. & Oxland, T. R. Preliminary investigation of spinal level and postural effects on thoracic muscle morphology with upright open MRI. *JOR Spine* 4. doi:10.1002/JSP2.1139, (Mar. 2021).
- 102. Charest-Morin, R., Zhang, H., Shewchuk, J. R., Wilson, D. R., Phillips, A. E., Bond, M. & Street, J. Dynamic morphometric changes in degenerative lumbar spondylolisthesis: A pilot study of upright magnetic resonance imaging. *Journal of Clinical Neuroscience* 91, 152–158. doi:10.1016/j.jocn.2021.06.027 (Sept. 2021).
- 103. Steingoetter, A., Fox, M., Treier, R., Weishaupt, D., Marincek, B., Boesiger, P., Fried, M. & Schwizer, W. Effects of posture on the physiology of gastric emptying: a magnetic resonance imaging study. Scandinavian journal of gastroenterology 41, 1155–1164. doi:10.1080/00365520600610451 (Oct. 2006).

- 104. Traser, L., Schwab, C., Burk, F., Özen, A. C., Burdumy, M., Bock, M., Richter, B. & Echternach, M. The influence of gravity on respiratory kinematics during phonation measured by dynamic magnetic resonance imaging. *Scientific Reports* 2021 11:1 11, 1–13. doi:10.1038/s41598-021-02152-y (Nov. 2021).
- 105. LeMarr, J. D., Golding, L. A. & Crehan, K. D. Cardiorespiratory Responses to Inversion. The Physician and sportsmedicine 11, 51– 57. doi:10.1080/00913847.1983.11708682 (1983).
- 106. Sauvageau, A., Desjarlais, A. & Racette, S. Deaths in a head-down position: A case report and review of the literature. Forensic Science, Medicine, and Pathology 4, 51–54. doi:10.1007/S12024-007-0031-4/TABLES/2 (Mar. 2008).
- Andrews, K. Upside-down death the pathophysiology of inversion. Pathology 46, S19. doi:10.1097/PAT.0000000000000056 (Jan. 2014).
- 108. Schäfer, A. T. Death in a Head-Down Position. Forensic Pathology Reviews, 137–154. doi:10.1007/978-1-59259-910-3{_}3 (Nov. 2005).
- 109. Martin, A., Miller, J. B., Walsh, M. & Prahlow, J. A. Case report Positional asphyxia in rollover vehicular incidents. *Injury Extra* 42, 1–3. doi:10.1016/j.injury.2010.09.001 (2010).
- Dixon, A. E. & Peters, U. The effect of obesity on lung function. Expert review of respiratory medicine 12, 755. doi:10.1080/17476348.
 2018.1506331 (Sept. 2018).
- 111. Health Survey for England, 2021: Data tables NHS England Digital
- 112. Seber, G. A. F. & Wild, C. J. Nonlinear Regression. Wiley Series in Probability and Statistics doi:10.1002/0471725315 (Feb. 1989).

- 113. Dumouchel, W. & O'brien, F. Integrating a Robust Option into a Multiple Regression Computing Environment, 41–48. doi:10.1007/978-1-4613-9154-8{_}3 (1991).
- 114. Holland, P. W. & Welsch, R. E. Robust regression using iteratively reweighted least-squares. Communications in Statistics - Theory and Methods 6, 813–827. doi:10.1080/03610927708827533 (Jan. 1977).
- 115. Yarnykh, V. L. Actual flip-angle imaging in the pulsed steady state:
 A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magnetic Resonance in Medicine* **57**, 192–200. doi:10.1002/MRM.21120 (Jan. 2007).
- 116. Mead, J., Takishima, T. & Leith, D. Stress distribution in lungs: a model of pulmonary elasticity. *Journal of applied physiology* 28, 596–608. doi:10.1152/JAPPL.1970.28.5.596; PAGE: STRING: ARTICLE/CHAPTER (1970).
- 117. Bhana, R. H. & Magan, A. B. Lung Mechanics: A Review of Solid Mechanical Elasticity in Lung Parenchyma. *Journal of Elasticity* 2023 153:1 153, 53–117. doi:10.1007/S10659-022-09973-6 (Jan. 2023).
- 118. Al-Mayah, A., Moseley, J., Velec, M. & Brock, K. Toward efficient biomechanical-based deformable image registration of lungs for image-guided radiotherapy. *Physics in Medicine & Biology* **56**, 4701. doi:10. 1088/0031-9155/56/15/005 (July 2011).
- 119. Forsberg, D. Fordanic / Image-registration 2015.
- 120. De Troyer, A. Mechanical role of the abdominal muscles in relation to posture. Respiration Physiology 53, 341–353. doi:10.1016/0034-5687(83)90124-X (Sept. 1983).

- 121. McKeough, Z. J., Alison, J. A. & Bye, P. T. Arm positioning alters lung volumes in subjects with COPD and healthy subjects. *Australian Journal of Physiotherapy* **49**, 133–137. doi:10.1016/S0004-9514(14)60129-X (Jan. 2003).
- 122. Hopkins, S. R., Arai, T. J., Henderson, A. C., Levin, D. L., Buxton, R. B. & Kim Prisk, G. Lung volume does not alter the distribution of pulmonary perfusion in dependent lung in supine humans. *The Journal of Physiology* 588, 4759. doi:10.1113/JPHYSIOL.2010.196063 (Dec. 2010).
- 123. Kizhakke Puliyakote, A. S., Tedjasaputra, V., Petersen, G. M., Sá, R. C. & Hopkins, S. R. Assessing the pulmonary vascular responsiveness to oxygen with proton MRI. *Journal of applied physiology (Bethesda, Md. : 1985)* 136, 853–863. doi:10.1152/JAPPLPHYSIOL. 00747.2023 (2024).
- 124. Klimeš, F. et al. 3D phase-resolved functional lung ventilation MR imaging in healthy volunteers and patients with chronic pulmonary disease. Magnetic resonance in medicine 85, 912–925. doi:10.1002/MRM.28482 (Feb. 2021).

Appendices

Appendix A

Neural Network Training

A.1 Hyperparameters

Hyperparameters for Model Training					
Acquisition Sequence	Coronal 2D GRE	Coronal 2D GRE	Axial 2D GRE	Sagittal 2D GRE	Coronal 2D HASTE
Model Architecture	UNet	SegNet	SegNet	SegNet	SegNet
Pre-Initialised Weights	None	VGG-16	2D GRE Lung VGG-16	VGG-16	2D GRE Lung VGG-16
Solver	Adam	SGDM	Adam	Adam	Adam
Momentum	N.A.	0.9	N.A.	N.A.	N.A.
Gradient Decay Factor	0.9	N.A.	0.9	0.9	0.9
Squared Gradient Decay Factor	0.9	N.A.	0.9	0.9	0.9
Initial Learn Rate	0.002	0.06	0.0075	0.005	0.0007
Learn Rate Drop Factor	0.2	0.5	0.7	0.7	0.6
Learn Rate Drop Period*	12	1	15	15	15
Mini Batch Size***	32	4	32	32	32
Validation Frequency**	50	50	50	50	50
Validation Patience**	5	5	5	5	5
Max Epochs	100	100	300	300	300
L2 Regularisation	0.0005	0.0005	0.0001	0.0001	0.0001

Table A.1: Optimised hyperparameters use to train each of the convolutional neural networks detailed in this work. Omitted parameters used inbuilt MATLAB default values. * Number of epochs, ** number of minibatches, *** number of training images.

Appendix B

Computational Hardware

B.1 PREFUL Pipeline Application

CPU: 13th Gen Intel Core i9-13900K

Memory: 32 GB DDR4

Graphics: NVIDIA T400 4GB

Storage: 1 TB NVMv SSD

B.2 Neural Network Training

CPU: Intel Xeon E5-1620 v3

Memory: 32 GB DDR4

Graphics: NVIDIA TITAN Xp

Storage: 1 TB HHD

B.3 Combustion System Operation

CPU: 11th Gen Intel Core i5-1135G7

Memory: 8 GB DDR3

Graphics: Intel Iris Xe Graphics

Storage: 256 GB SSD