Development of Upright, Multimodal, Respiratory MRI

by James William Harkin

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy, April 2020

Contents

Abstract

Respiratory diseases can have a severe impact on people's quality of life and life expectancy. To investigate these diseases, CT and spirometry are the current "gold standards" used by clinicians. In tandem, CT and spirometry, offer high resolution images and lung function information. However, CT exposes the patient to ionising radiation, making multiple investigations undesirable. The results of pulmonary function tests, such as spirometry, are also variable and depend on the effort levels of the patient. In addition, the information provided from most pulmonary function tests is a global measure; abnormal values may only occur when a large region of the lung is affected. These issues mean that, using current measures, it can be difficult to diagnose respiratory diseases and track their progression to see if a particular treatment is being effective.

In tissues other than the lungs, ^{1}H MRI can offer highly detailed structural, as well as functional information. However, ^{1}H MRI is difficult to perform in the lungs due to their low ^{1}H density. In order to maximise the information obtainable from respiratory MRI, other techniques need to be employed . Implementing hyperpolarised ^{129}Xe MRI, ^{19}F MRI and ^{1}H diaphragm imaging are the major focuses of this thesis.

The other problem with most conventional MRI for respiratory imaging is that it is performed supine. Lung function is reduced when supine, as opposed to seated or standing. This causes two problems. Firstly, patients with severe lung disease may struggle to lie down for extended periods of time; arguably these are the people who it is most useful to scan. Secondly, most people spend the majority of their lives upright, therefore images acquired upright should be of greater relevance for clinicians. At Nottingham, the 0.5T Paramed MROpen Upright scanner is the first of its kind to offer ^{1}H and multinuclear imaging. This scanner allows for imaging in a variety of orientations. The open design of the scanner also makes it particularly suited to paediatric imaging. Children as young as three have sat comfortably inside the scanner, on a parent or guardian's lap, without need for sedation.

In this thesis, the optimum conditions for polarising ^{129}Xe in a batch mode system are investigated. Higher rates of polarisation build-up were observed at higher temperatures. Vastly higher temperatures were also observed inside the cell (particularly in the first few cm after the laser strikes the cell), than the oven during runaway. Higher polarisations of ^{129}Xe were found in gas mixes with lower concentrations of ^{129}Xe (leaner mixes). However, the largest bulk

magnetisation was found with a balance of ^{129}Xe and N_2. For a given Rb vapour density replacing N_2 with ^{4}He appeared to have little effect on the polarisation build-up rate of ^{129}Xe. The polarisation build-up rate was also seen to be highly variable in a cell with a bead of Rb, using the same external conditions. After "spreading" the Rb bead, by vaporising it and forcing it to condense on the walls of the cell, the build-up rate became more consistent.

Before the Paramed MRI scanner was built, hyperpolarised ^{129}Xe imaging at Nottingham was performed using a 1.5T GE scanner. The original Rapid birdcage ^{129}Xe coil, was not fit for purpose; larger volunteers couldn't be imaged. The testing of a new coil from Clinical MR Solutions L.L.C. is detailed; with the aim of investigating if it was fit for use. In QTAR mode, an inconsistent signal was seen throughout the lungs, making it unsuitable. In QUAD T/R mode, initial tests were encouraging.

On the Paramed, a sequence has been developed to image a single slice in $<$0.5s. This allows for the movement of the diaphragm to be characterised throughout the breathing cycle and to track repeating lung density changes. In addition, a stock Paramed sequence was altered so the duration of the scan was reduced from an average of 35s to 21s; the acquired resolution was maintained at 160x128x7. This meant the sequence could be acquired within a breath hold; allowing the diaphragm morphology at full expiration and inspiration to be characterised. Compressed sensing has also been implemented on the upright. A sequence has been developed, at a resolution of 256x200x22 and undersampling by a factor of 3, to acquire an image in 22.6s. By undersampling a resolution of 256x160x10 by a factor of 4.82 an image can be acquired in 5.1s. With both, the diaphragm can be accurately characterised. The shorter duration sequence is tailored for patients with more severe lung disease, who can't hold their breath for as long. The intention is to apply compressed sensing to hyperpolarised imaging in the future.

Using a small surface ^{129}Xe test coil on the Paramed scanner, a 3D GRE sequence has been developed, which can resolve detail of $<$3mm in $<$20s. Also using ^{129}Xe a flip angle calibration and dissolved phase spectroscopy have been performed. A calibration for an SAR monitoring system has been produced, to allow for in vivo imaging using the test coil once the system has been built. ^{1}H on a hydrofluorocarbon gas has also been imaged, illustrating it should be possible to image ^{19}F using the Paramed.

MRI has excellent potential to be used clinically to diagnose and track the progression of respiratory disease. By imaging in an upright orientation, fewer patients will be excluded from being scanned using the techniques developed and the images should be of greater diagnostic use for clinicians. By using techniques such as hyperpolarisation, the lower field strength of the upright scanner than a conventional scanner can be counteracted.

Acknowledgements

First of all, I must thank the Haydn Green Foundation for funding my exploration into the world of respiratory MRI.

I would like to express gratitude to my supervisor's, Michael Barlow and Ian Hall, for their guidance throughout my studies. Day to day it has been a pleasure to work in the hyperpolarised ^{129}Xe team. Thanks must also go to Andrew Cooper, Sarah Wilson, Brett Haywood, Andrew Prayle, Penny Gowland, Joshua McAteer, Pete Thelwall, Shahideh Safavi, Christoph Arthofer, Benjamin Prestwich and Andrew Peters; for the expertise, support and teaching they kindly provided.

I am hugely grateful to my parents, family and Dervla. Their support while I avoid the "real world" to pursue my interests has been unwavering. Also to Robert Sharp for guiding me in the right direction. If you've made it through the abstract and know me outside of work, I almost certainly owe you thanks also.

Particular gratitude must go towards Robert Irwin, without whom my PhD experience would have been significantly less enjoyable.

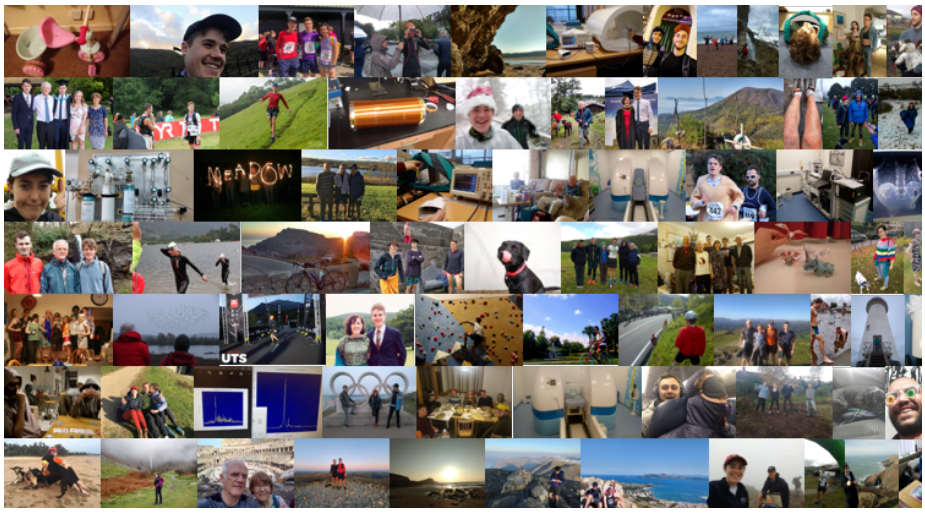

Abbreviations

α	Flip angle
α_E	Ernst angle
γ	Gyromagnetic ratio
μ	Magnetic moment
ω	Larmor frequency
ξ	Oxygen enhancement factor
AAS	Atomic absorption spectroscopy
AM	Alkali metal
B_0	Constant large magnetic field produced for NMR or MRI
B_1	RF excitation pulse
CF	Cystic fibrosis
COPD	Chronic obstructive pulmonary disease
CS	Compressed sensing
CT	Computerised tomography
DNP	Dynamic nuclear polarisation
FA	Flip angle
FGRE	Fast gradient (field) echo
FT	Fourier transform
FWHM	Full width at half maximum
GRE	Gradient (field) echo
Δk_i	Phase encoding step size along ith dimension
HFC	1,1,1,2-Tetrafluoroethane ($C_2F_4H_2$)
L_i	FOV along ith dimension
OP	Optical Pumping
M_0	Net magnetic moment of a nuclear species
M_k	Net magnetic moment of a nuclear species after k RF pulses
MNS	Multinuclear spectroscopy
MRI	Magnetic resonance imaging
NMR	Nuclear magnetic resonance
PFP	Perfluoropropane (C_3F_8)
pO_2	Partial pressure of oxygen
QTAR	Quadrature transmit array receive
Quad T/R	Quadrature transmit and quadrature receive
RF	Radio frequency
ROI	Region of interest
SAR	Specific Absorption Rate
SE	Spin echo
SEOP	Spin exchange optical pumping
SNR	Signal to noise ratio
T1	Relaxation time constant of magnetic moment longitudinally towards B_0
T2	Relaxation time constant of magnetic moment transversely due to reversible contributions
$T2^*$	True relaxation time constant of magnetic moment transversely due to all contributions
TE	Echo time
TG	Transmit gain
TR	Repetition time
VPP	Peak to peak voltage

Chapter 1

Introduction

The majority of respiratory conditions are both chronic and progressive [1]. They can have a severe impact on people's quality of life and life expectancy. Current methods to diagnose and track the progression of respiratory diseases are mainly restricted to providing a global measure of lung function as a whole [2], or detailed structural data at the cost of exposure to ionising radiation (CT) [3].

If only a small region of the lung is affected by a disease, global lung function tests may appear normal. Also clinicians may be hesitant to perform a CT scan because of the risk of ionising radiation exposure. This can make diagnosis of certain lung diseases difficult. Clinicians may also be disinclined to suggest multiple CT scans to track the progression of a disease, or to see if a treatment is effective, because of the elevated risk of multiple exposures to ionising radiation.

Elsewhere in the body, MRI can offer similarly high resolution scans to CT [4]. However, the low proton density of the lungs makes respiratory MRI challenging. Historically, respiratory MRI has not been the first choice for clinicians. However, there are techniques being developed where respiratory MRI could offer as much, or more, information as CT or current lung function tests.

Focusing on protons first, the diaphragm can be clearly defined in an MRI image. There is hope that in the future, diaphragmatic abnormalities may be diagnosed through MRI rather than fluoroscopy. In addition to avoiding the ionising radiation associated with fluoroscopy, MRI also has the potential to offer more information. The image from MRI can be from a well defined slice, unlike fluoroscopy or an X-ray, which offer an infinite slice through the body [5]. At Nottingham the upright Paramed 0.5T MRI scanner has lead to an increased focus on imaging the diaphragm. It is known that the respiratory system is impaired when supine rather than upright [6] and the hope was that the effects of orientation on the diaphragm could be investigated.

Introducing a multinuclear tracer can allow respiratory MRI to be performed more effectively. After breathing in hyperpolarised ^{129}Xe, high SNR ventilation images can be obtained [7]. The higher the levels of polarisation, the higher the signal in the final image. It is important to investigate the optimum conditions

to hyperpolarise ^{129}Xe, because in order for the technique to be rolled out in a clinical setting, high levels of polarised ^{129}Xe will be required within a short period of time (<1hour).

The images obtained from ^{129}Xe imaging can offer more information than a CT scan. Signal only comes from ^{129}Xe breathed in, therefore if a section of the lung has restricted flow, less ^{129}Xe will flow in to it in the breathing cycle and it will appear dark on the final image [8]. ^{129}Xe also diffuses into the blood and around the body [9]. Its large chemical shift enables the operator to know what type of tissue the ^{129}Xe is dissolved in. This allows for a direct measurement of diffusion rates from the lungs into the blood. If a region of lung has impaired diffusion, this could theoretically be detected on a local level. ^{19}F is an alternative multinuclear tracer to ^{129}Xe [10]. It doesn't require a complex and expensive polarisation procedure and unlike ^{129}Xe can be free breathed (^{129}Xe is an anaesthetic). The fact it can be free breathed means wash in/out kinetics can be investigated [11].

The upright MRI scanner, offers the opportunity to image patients at later stages of lung conditions; many of whom can't lie down due to impaired lung function when supine rather than upright [6]. Images acquired in the upright orientation, should also be of more clinical relevance; in general, people spend most of their lives upright and certainly most of the time when their respiratory system is working hard. The upright, open design of the scanner also allows paediatric cohorts to be scanned more comfortably in the magnet; without need for sedation. Children as young as three have sat comfortably in the scanner on their parent's or guardian's lap.

As part of this work, many people have provided help and expertise. Where help was received, or any work performed partially by someone else, they have been acknowledged in the thesis outline below. Throughout my PhD I received guidance from my supervisors; Dr Michael Barlow and Professor Ian Hall. Also, for my MRI work requiring hyperpolarised ^{129}Xe, any hyperpolarised ^{129}Xe was produced by Robert Irwin.

1.1 Thesis Outline

This thesis can be broken down into four broad sections. Background theory, investigations into the optimum conditions for hyperpolarising ^{129}Xe, in vivo ^{129}Xe imaging using a 1.5T flatbed scanner and respiratory imaging using a 0.5T upright scanner.

- **Background -** This section provides background theory to the work I performed. The basics of MRI are first discussed, before background to the MRI sequences are explained in more detail. Next, hyperpolarised MRI is discussed, before the background to ^{19}F imaging is explained. Finally, the background to lung disease states and how hyperpolarisation is performed are discussed.

- **Improving Rate of ^{129}Xe Build-up for Batch Mode SEOP -** While the analysis is my own, the setup and running of the experiment was performed in tandem with a fellow PhD student, Robert Irwin. Running the kit is a time intensive job, with the setup needing constant supervision throughout a run which could last up to 12 hours.

  The experimental setup used previously is first discussed, improvements made to the setup are also explained. Next, how to calculate the polarisation through NMR and the internal temperature through Raman spectroscopy is gone through. The problems with the repeatability of the experiment and the measures taken to improve the lack of repeatability with the old setup and/ or account for it are explained. Three experiments were performed. Firstly, into the effects of oven temperature on the rate of ^{129}Xe polarisation build-up, secondly into the internal temperatures within the cell during SEOP and finally into the effects of changing the gas mix on the internal temperature and ^{129}Xe polarisation.

- **Scanner Hardware -** The hardware of both the 1.5T GE and 0.5T Paramed MRI scanners is discussed in this section. The bespoke hardware required for multinuclear imaging is also explained. Tuning of a surface coil for the Paramed is detailed, as well as how an SAR monitoring system was calibrated. Coil SAR modelling was performed by Benjamin Prestwich and Dr Andrew Peters.

- **1.5T GE ^{129}Xe Imaging -** In vivo, ^{129}Xe imaging on a 1.5T GE flatbed scanner is discussed in this section. Historically, ^{129}Xe imaging at Nottingham was been performed using this scanner. However, the original Rapid birdcage coil was unfit for scanning patients, due to it reducing the effective bore size of the magnet significantly. The attempted implementation of a new flexible chest coil from Clinical MR Solutions L.L.C. is detailed. This new coil would make it more comfortable for the volunteer, allow for a wider variety of volunteers to be scanned and theoretically provide higher quality images. A huge thank you must go to the volunteers, the clinicians who make sure the scanning is safe (Dr Shahideh Safavi and Dr Jonathan Brooke) and also Dr Brett Haywood for training me to use the scanner.

- **Upright 0.5T Proton Imaging -** This section can be broken down into 3 strands of work:

  - The development of a single (thin) slice lung density imaging technique, which allows the lungs and diaphragm to be imaged at a very high time resolution of < 0.5s: Periodic changes in lung density can be extracted from the images, at a frequency of the average breathing rate, through Fourier decomposition. By splicing the images acquired into a video, the diaphragm movement can clearly be seen. The concept of lung density tracking was suggested by Dr Andrew Prayle. Without him, the scanning of children would not have been possible.

He also suggested the image analysis package, ANTsR, and helped with the coding.

  - The addition of a hold into a spin echo sequence: The stock Paramed sequence took a duration of on average 35s seconds. With the addition of a hold, this sequence duration is reduced to ∼21s. The intended use of the sequence was for a diaphragm study lead by Dr Shahideh Safavi; investigating the difference in diaphragm morphology when supine compared to upright. For this study a breath hold was required, so reducing the duration made the scan more comfortable for volunteers. The image analysis was performed by Dr Christoph Arthofer. The sequence has now become the standard SE sequence at Nottingham when the duration of the scan is a consideration.
  - Development of compressed sensing on the upright: The points in k-space to be sampled were chosen by Joshua McAteer. He also reconstructed the images from the raw data. My role was encoding the sampling scheme into the TecMag MRI spectrometer in TNMR, to ensure the correct k-space points were acquired and to run the scanner. The technique was developed using ^{1}H MRI and images have been acquired using the technique in vivo. The intention is to apply the technique to ^{129}Xe in the future. Due to the breath hold requirements of ^{129}Xe MRI, any reduction in the duration of the scan is beneficial; especially for patients with compromised breathing.

- **Upright 0.5T Multinuclear Imaging -** This section can be broken down into 2 strands of work:
  - ^{129}Xe imaging on the upright scanner: The development of a flip angle calibration, dissolved phase spectroscopy and 3D imaging are all detailed.
  - Investigations into the feasibility of ^{19}F imaging on the upright: Dr Pete Thelwall from Newcastle provided expertise, making the investigation much easier.

- **Conclusion -** The broad conclusions from each of the above sections are outlined.

Chapter 2

Background

2.1 Magnetic Resonance Imaging

2.1.1 NMR

In Nuclear magnetic resonance (NMR), an NMR active nuclei species is placed within a constant magnetic field (B_0). Each individual nucleus possesses a magnetic moment (μ) due to their spin ($\vec{S}$). See equation 2.1; where γ is the gyromagnetic ratio of the nuclear species (the ratio of a particle's magnetic moment to its angular momentum).

$$\mu = \gamma \vec{S} \tag{2.1}$$

γ of NMR active nuclei focussed on in this document can be found in table 2.1 [12].

Spin $\frac{1}{2}$ nuclei such as ^{1}H and ^{129}Xe have two independent spin states. When $B_0 = 0$, these two states are degenerate and it is therefore entropically favourable for an equal number of nuclei to possess $\vec{S} = \pm\frac{1}{2}$. In the presence of no external magnetic field, the net magnetic moment of the species is zero. However, when $B_0 > 0$ the degeneracy of the $\vec{S} = \pm\frac{1}{2}$ sublevels is broken in Zeeman splitting [13]; it becomes more energetically favourable for the individual magnetic moments of the nuclear species to align with B_0. The difference

Nuclide	Gyromagnetic ratio(γ) $/MHzT^{-1}$	Natural abundance (%)
^{1}H	42.576	0.9885
^{129}Xe	-11.777	26.4007
^{19}F	40.052	100
^{3}He	32.434	10^{-3}

Table 2.1: Nuclear Spin Properties of NMR active nuclei focussed on in this document.

in energy between the two spin states of a spin $\frac{1}{2}$ system is given by equation 2.2.

$$\Delta E = \gamma \hbar B_0 \tag{2.2}$$

At temperatures above absolute zero, the individual nuclei exhibit thermal motion which opposes the tendency of the nuclei to favour the lower energy state. When the nuclei have reached a stable state, the ratio of nuclei in the higher energy sublevel ($N_\uparrow$), compared to the lower energy sublevel ($N_\downarrow$), is given by the Boltzmann distribution (see equation 2.3) [13].

$$\frac{N_\uparrow}{N_\downarrow} = e^{-\Delta E/kT} \tag{2.3}$$

This stable state isn't reached immediately after the application of B_0. Over time the net magnetic moment of the nuclei species (M_0) will align with B_0 through a process known as spin lattice build-up. Equation 2.4 takes M_{eq} to be the value of M_0 as time (t) tends to ∞ and $T1$ is the characteristic time constant for the nuclear species in question at the field strength of B_0.

$$M_0(T) = M_{eq}(0)(1 - exp(\frac{-t}{T1})) \tag{2.4}$$

In NMR, an oscillating magnetic field Radio Frequency (RF) Pulse, is applied; generally perpendicular to B_0. If the RF pulse resonates at close to the resonant frequency for the nuclear species, there is some finite possibility that each nucleus is excited. This resonant frequency is the Larmor frequency (ω); see equation 2.5 [14].

$$\omega = \gamma B \tag{2.5}$$

In the presence of B_0, the magnetic moments of the individual nuclei precess around B_0 at their Larmor frequency. Therefore, after an RF pulse, M_0 also precesses at the same Larmor frequency.

The precession of M_0 will induce a current in a coil placed close; allowing an NMR signal to be detected. However, over time the magnitude of M_0 in the $\hat{x}$,$\hat{y}$ plane reduces due to two contributions; longitudinal (section 2.1.3) and transverse relaxation (section 2.1.4). This NMR signal is proportional to γ (see equation 2.1) [7].

By using high magnetic fields (B>0.1T), large gaps in Larmor frequency can be created between and within nuclear species. Coils can be tuned to amplify signals with certain frequencies and suppress others; this creates high specificity.

2.1.2 MRI

Conventional Magnetic Resonance Imaging (MRI) exploits protons. Protons are spin $\frac{1}{2}$ nuclei and hence have a non zero magnetic moment. The magnetic field from the protons in water (M_0) in an MRI scanner is $< 1\%$ of B_0, making measurement of M_0 technically challenging. In order to detect M_0, a 90° RF pulse can be applied, tipping M_0 perpendicular to B_0. This means M_0 precesses at the proton Larmor frequency in the $\hat{x}$, $\hat{y}$ plane.

The predecessor of MRI, "NMR phytogeography", was described in 1973 by Paul Lauterbur [15]. In this imaging technique linear magnetic field gradients were applied during signal acquisition so that the Larmor frequency of a specific species was position dependent [16] (see section 2.1.10). A gradient during different parts of an MRI sequence has a different effect, requiring gradient coils which can be switched on and off.

By applying a gradient during signal acquisition, the frequency of resonance will depend on position. Therefore, the current induced by M_0 will be at different frequencies depending on the origin of M_0. However, while frequency encoding allows the information for different spatial points to be acquired simultaneously, the information acquired can only be in one dimension. Even if a gradient was applied along another axis there will be multiple points with the same frequency offset. In order to acquire information along another axis phase encoding is used.

In phase encoding, after the RF pulse and before the signal acquisition, a gradient (perpendicular to the frequency encoding gradient) can be applied. For the period while the gradient is applied the nuclei precess at different rates depending on their position. When the gradient is switched off the nuclei return to precessing at the same speed. However, the nuclei are now out of phase with each other, and nuclei in rows perpendicular to the gradient are in phase with each other. By applying different gradient strengths, the phase difference between the rows will vary. This allows the contribution each row is making to the signal to be extracted [17], [14]. Using phase and frequency encoding a 2D image can be formed.

By applying a gradient during the RF pulse a defined slice can be excited perpendicular to the applied gradient (see section 2.1.8). The excited slice can be altered by applying a gradient during the RF pulse and changing the frequency of the RF pulse slightly. A stack of 2D slices can then be formed, producing an image.

2.1.3 Longitudinal Relaxation

Longitudinal relaxation happens due to spin lattice relaxation (see equation 2.4). Each dipole experiences tiny magnetic fields created by spinning protons. These tiny magnetic fields are unique for each dipole. This acts to reduce the angle of the cone traced out by M_0, meaning over time the component of M_0 in the $\hat{x}$, $\hat{y}$ plane reduces as the component of M_0 in the $\hat{z}$ direction increases. The recovery time constant in the $\hat{z}$ direction is known at T1.

2.1.4 Transverse Relaxation

Different human tissues have different magnetic susceptibilities and their susceptibility is also different from air. Since the individual magnetic moments are in different tissues, they experience slightly different magnetic fields. This means that the magnetic moments precess at different rates and as the magnetic moments go out of phase, the magnitude of M_0 is reduced. These inhomogeneities

don't vary over time and therefore cause a static effect which is reversible (see section 2.2.2). The other reversible contribution to transverse relaxation is introduced through inhomogeneities in the magnetic field of the scanner. Again, the inhomogeneities in the scanner are static and the relaxation they cause is therefore reversible. The net contribution of the reversible contributions to transverse relaxation has a characteristic relaxation time of $T2'$ [18].

The second component to transverse relaxation is spin-spin relaxation. The orientation of the neighbouring spins of protons (which fluctuate) affects the magnetic field experienced by an individual proton and therefore causes its precession rate to fluctuate as the magnetic field changes in a stochastic fashion. These random spin-spin interactions vary over time and are irreversible. The characteristic decay rate from just spin-spin interactions is T2 [19].

The overall characteristic decay rate of transverse relaxation, taking into account both reversible and irreversible processes, is $T2^*$ (see equation 2.6).

$$\frac{1}{T2^*} = \frac{1}{T2} + \frac{1}{T2'} \tag{2.6}$$

2.1.5 Nyquist Criterion

For a signal which consists of discretely sampled points, the Fourier transform (FT) in another domain must be periodic. The Nyquist criterion states that for a signal which repeats over time (T) the sampling rate ($1/\Delta T$) must be at least **twice** as great. In MRI, the signal is discretely sampled. Hence in the Fourier transformed domain it is periodic and therefore replicates. If the signal is undersampled this can lead to aliasing artefacts (see section 2.1.6).

2.1.6 K-space Sampling Strategies

The data acquired during an MRI sequence is acquired in a 2D or 3D space that is the inverse FT of the data in real space. This space is referred to as k-space. Each voxel in k-space contributes information to every voxel in the real space image, with high frequency detail in the periphery of k-space and the signal intensity from the centre [20].

Traditionally a Cartesian k-space filling is used (figure 2.1). However, other methods of filling k space such as a spiral, zig-zag or projection (figure 2.2) are also available.

Undersampling K-space

A sequence can be said to be fully sampled if the Nyquist criterion (equation 2.7) is satisfied; where N is the number of voxels in the $\hat{y}$ dimension, L_y is the FOV in the $\hat{y}$ dimension and Δk_y is the k space phase encoding step size. (see figure 2.3) .

$$\Delta k_y = \frac{1}{L_y} = \frac{1}{N\Delta y} \tag{2.7}$$

In order to reduce scan durations, different undersampling strategies can be

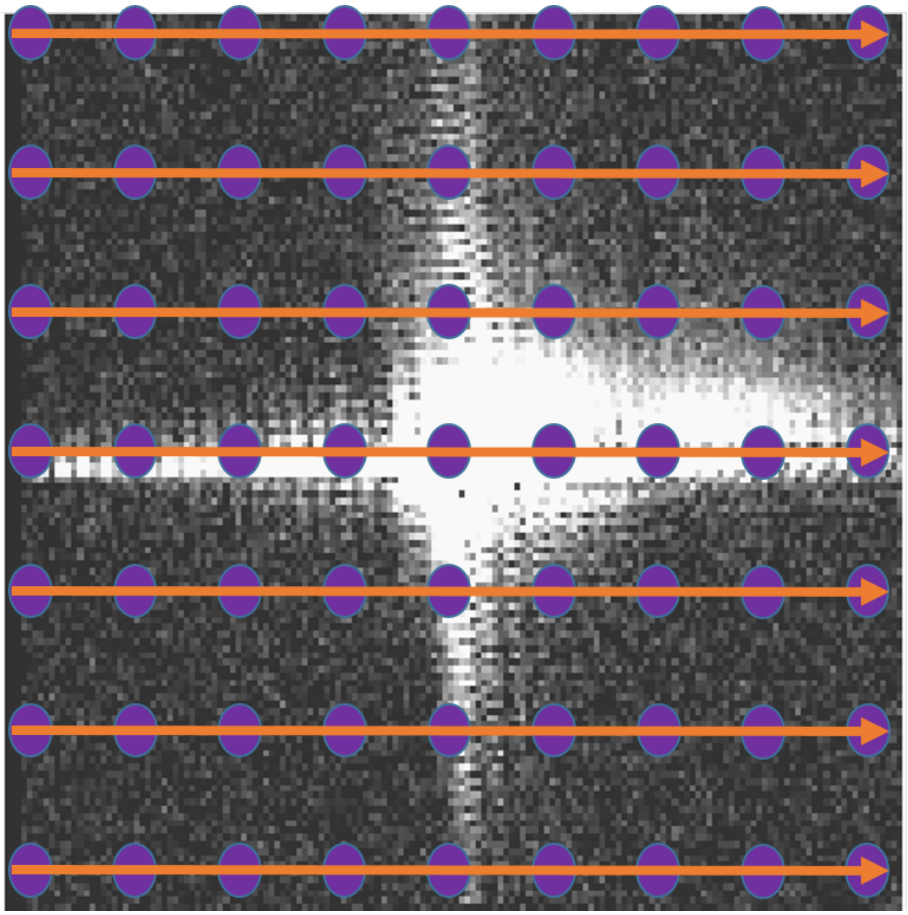

Figure 2.1: Cartesian sampling strategy of k-space.

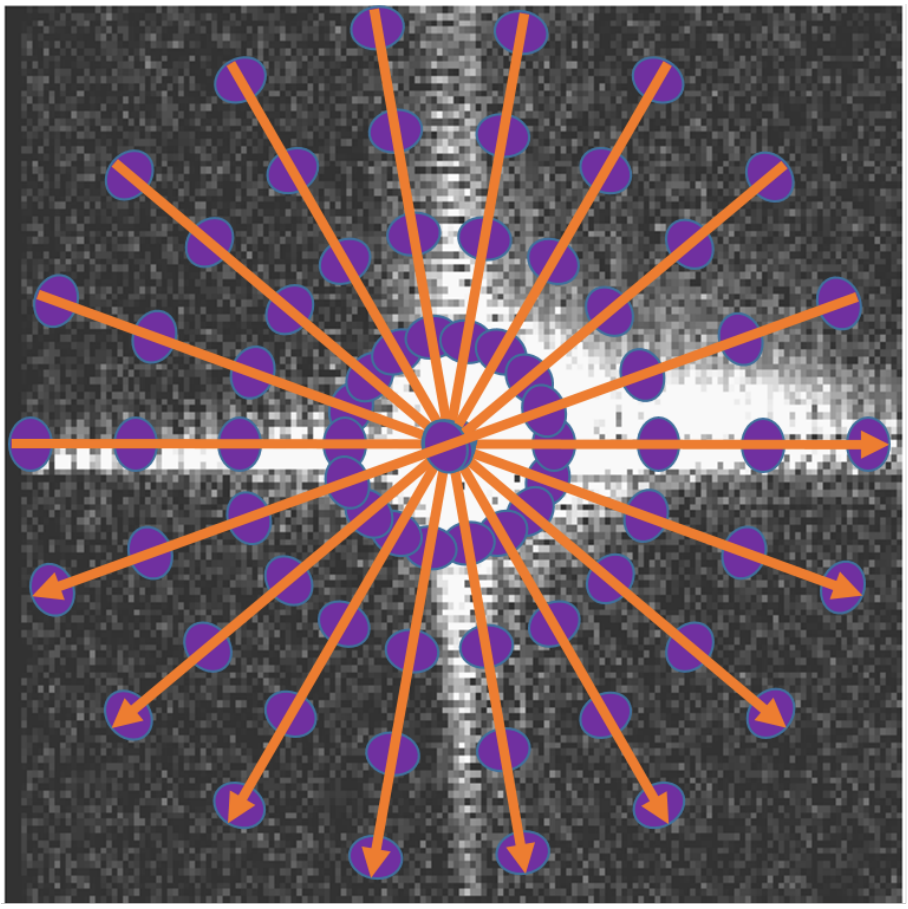

Figure 2.2: Projection sampling strategy of k-space.

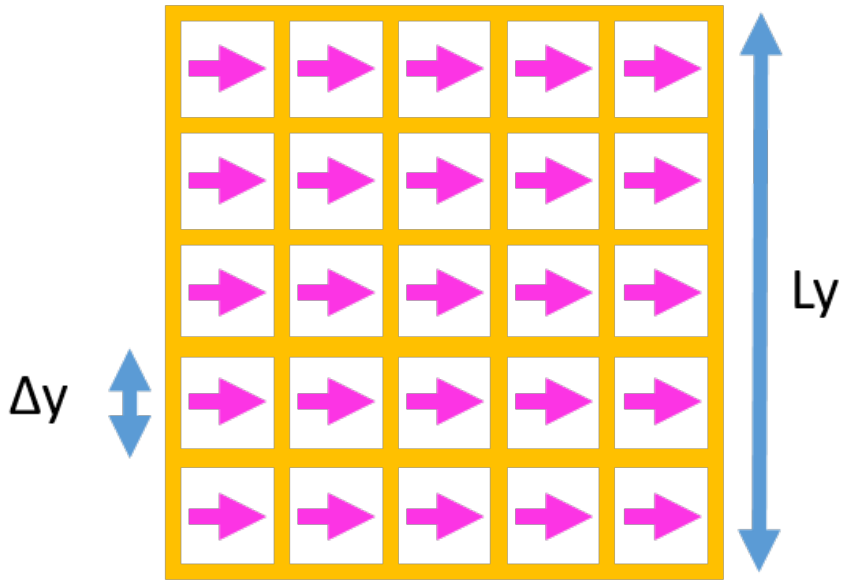

Figure 2.3: Image domain: Ly is the FOV and Δy is the voxel dimension along the y axis.

employed without a loss in image quality. However, undersampling can lead to artefacts such as aliasing. This occurs when the field of view (FOV) acquired is too small and due the Fourier transform being cyclical, information outside the FOV wraps round onto the image.

Aliasing Aliasing (sometimes referred to as wraparound) is due to the periodicity of the Fourier domain of the discretely sampled signal. Aliasing occurs when k-space is too sparsely sampled. If the phase encoding step size Δk_y is too large, the resulting image will exhibit its intrinsic periodicity; whereby signal at the top of the image also has contributions from the bottom and vice versa. In practice, for many methodologies some amount of aliasing can be acceptable. If the region of interest (ROI) is at the centre of the image, then aliasing at the edges shouldn't interfere.

Projection Imaging In conventional projection imaging, the centre points of k-space are oversampled (see figure 2.2). This averaging of the centre of k-space makes the technique particularly robust at dealing with motion artefacts. Especially when compared to Cartesian sampling, where motion causes ghosting in the phase encoding direction throughout the whole image. In projection reconstruction, streaking only majorly affects the edges of the image [21]. Projection sampling can also allow for high rates of undersampling, without a large loss in image quality [22].

Compressed Sensing Compressed Sensing (CS) refers to a technique where images are acquired in significantly fewer measurements than would satisfy the Nyquist criterion. CS is built on the idea that the information acquired in an MRI image is redundant; there are a lot less degrees of freedom than the number of acquisitions. This redundancy can either be modelled before from prior knowledge about the scan, or extracted from the data itself. Redundancy can also be increased by using multiple receive coils; thereby acquiring more data per acquisition. As datasets get larger, the redundancy should increase.

Just like a digital camera raw file being compressed to a JPEG, medical images can be compressed without great loss in image quality [23]. Videos can be compressed to a much greater extent, especially if large portions of the frames change little, or exhibit periodicity e.g. heart temporal images.

For a CS scheme to be successful it must meet 3 criteria [24]:

- Transformed sparsity: The image must be representable by a sparse sampling in a known domain (k-space).
- Any aliasing artefacts in the image must be incoherent.
- The reconstruction of the image from the sampled domain must be non-linear.

A random sampling of k-space is the simplest form of incoherent sampling. This forces any aliasing artefacts to appear as noise in the reconstructed image.

However, random sampling of k-space doesn't account for the fact that the majority of signal in the frequency domain is localised around the centre of k-space. Therefore, in practice, sampling strategies should be designed to have a distribution which matches that of the energy distribution in k-space. This should maximise the achievable imaging acceleration.

2.1.7 RF Pulses

An RF excitation pulse ($B_1(t)$) acts to tip M_0 by a defined flip angle (FA). The FA produced is given by equation 2.8, where t is the time at the end of the RF pulse [14].

$$\alpha(t) = \gamma \int_{t'=0}^{t} B_1(t')dt' \tag{2.8}$$

The simplest form of RF pulse is a rectangular Continuous wave (CW) pulse. A rectangular pulse has a large bandwidth, which means the RF pulse will affect a large range of resonant frequencies. Rectangular pulses can be much shorter than other pulse shapes, but they are not useful where spatial or spectral selection is required.

By applying a gradient during the RF pulse, the Larmor frequency of protons will depend on position. Only protons in a defined frequency band will be excited by the RF pulse, ensuring signal comes from a specific region. Ideally, the FA should be uniform across the slice. Assuming a constant gradient during an R.F pulse, the shape of a slice profile is the same as the frequency profile of the RF pulse. The frequency profile can be found from the FT of $B_1(t)$. The ideal pulse shape is therefore given by a SINC pulse with infinitely many side lobes, since its Fourier transform is a top hat function. In practice, an infinitely long SINC pulse is impractical and consequently some of the side lobes of the pulse should be removed (see figure 2.4). The more side lobes used, the better the approximation. A truncated SINC pulse of length Nt_0 is given by equation 2.9. Where N_L and N_R are the number of zero crossings left and right of the peak respectively, $N = N_L + N_R$, t_0 is the duration of one of the side lobes (see figure 2.4) and A is the RF pulse peak amplitude.

$$B_1(t) = At_0 \frac{sin(\frac{\pi t}{t_0})}{\pi t}) For - N_L t_0 \leq t \leq N_R t_0 \tag{2.9}$$

The bandwidth Δf of an R.F pulse, is the measure of the range of frequencies contained in that pulse. For a symmetric SINC pulse, the bandwidth if given by equation 2.10.

$$\Delta f \approx \frac{1}{t_0} \tag{2.10}$$

The first derivative of a SINC pulse is discontinuous at the edges of the pulse. By apodizing the pulse shape, this problem can be wholly or partially resolved. By replacing A_0 in equation 2.9, with equation 2.11, a Hanning window ($H_\alpha = 0.5$), or Hamming window ($H_\alpha = 0.46$), can be applied. For a symmetric SINC,

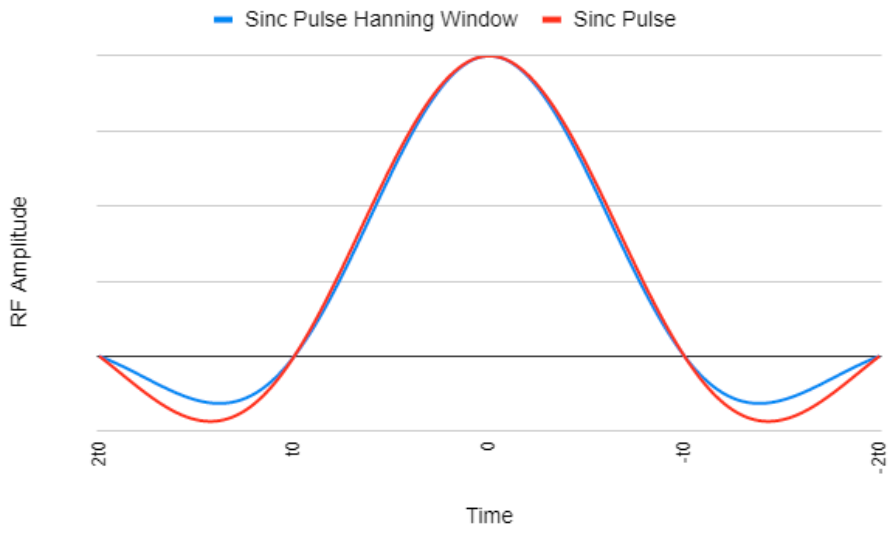

Figure 2.4: Truncated SINC pulse given by equation 2.9 and apodized Sinc with a Hanning window.

the Hanning window eliminates any discontinuity in the first derivative.

$$A = A_0((1 - H_\alpha) + H_\alpha cos(\frac{\pi t}{N t_0})) \tag{2.11}$$

Modern pulse sequences often use Shinnar-Le Roux (SLR) pulses. These tailored RF pulses are computer generated and vary depending on the the pulse duration, bandwidth, FA and ripple desired [25].

2.1.8 Slice Selection

During an RF pulse, only a bandwidth of frequencies Δf are excited. By applying a gradient during this pulse, the Larmor frequency of a nuclear species will be different depending on its location. This allows only a slice of nuclei to be excited. For a gradient of amplitude G_z, the slice excited has a thickness Δz (see equation 2.12) [26][14].

$$\Delta z = \frac{2\pi \Delta f}{\gamma G_z} \tag{2.12}$$

This slice is perpendicular to the direction of the gradient. To adjust the location of the slice excited, the carrier frequency of the RF wave must be adjusted. To offset a slice by distance δz, the carrier wave frequency must by δf (see equation 2.13).

$$\delta f = \frac{\gamma G_z \delta z}{2\pi} \tag{2.13}$$

Cross Talk

Even when using a apodized RF pulse, or an SLR pulse, the excited slice profile still deviates from ideal at the edges of the slice. This means a small proportion of the spins in neighbouring slices are also excited. Unless the time between the neighbouring slices being excited is much longer than the T_1, the magnetisation in the next slice won't have time to return to equilibrium before being excited. This can result in reduced image intensity. To avoid this, a gap can be placed between the slices, or slices can be excited in an interleaved fashion. This can be done by exciting every next but one slice, one by one, before then exciting the remaining slices, one by one. This increases the time between exciting neighbouring slices. Alternatively, a 3D acquisition can be used (see section 2.1.9).

Slice Rephasing

After the slice selection gradient, there is typically some transverse phase dispersion which can result in signal loss. To compensate for this a rephasing lobe of opposite polarity to the slice selection gradient can be used [26]. The isodelay parameter (Δt_I) is used to calculate the optimal strength of the slice rephasing gradient for a particular RF pulse. The isodelay parameter, is the duration from the mid point of the effect of the RF pulse on magnetisation to the end of the RF pulse. For a SINC pulse it is approximately half the duration of the RF pulse. The slice rephasing lobe area (A_r) can be calculated by equation 2.14, where r_z is the duration of the ramp down of the slice selection gradient (assuming the gradient ramps down at a constant rate).

$$A_R = G_z \Delta t_I + \frac{G_z r_z}{2} \tag{2.14}$$

2.1.9 Two-Dimensional vs Three-Dimensional Sequences

The acquisition of a 2D sequence is described in section 2.1; whereby a slice is repeatedly excited to form a 2D image. Multiple 2-D slices from different locations can be stacked up to provide information about a volume.

In a 3D acquisition, a set of slices (slab) are excited all at once. A second phase encoding is then performed, after the RF pulse, by a gradient in the same direction as the slab selection gradient; putting the data in a 3D k-space.

3D sequences produce images which are ideal for calculating volumes; there aren't the same issues with crosstalk in 3D acquisitions as 2D (see section 2.1.8). The best viewing plane for the image can also be chosen after the image has been acquired. 3D acquisitions can allow for thinner slices for a given TE. The SNR of MR acquisitions was shown by Edelstein et al. [27] to obey equation 2.15. Where $\Delta x \Delta y \Delta z$ is the voxel volume and $T_{acq,total}$ is the total amount of time that data is sampled for that voxel.

$$SNR \propto \Delta x \Delta y \Delta z \sqrt{T_{acq,total}} \tag{2.15}$$

Because every voxel in k space in a 3D acquisition provides information about every voxel in real space, $T_{acq,total}$ is simply the total amount of time data is being acquired during the sequence (see equation 2.16).

$$T_{acq,3Dtotal} = N_{phase1} \;\; N_{phase2} \;\; NEX \;\; T_{acq} \tag{2.16}$$

Where NEX is the number of signal averages, and N_{phase1} and N_{phase2} are the number of phase encoding points in 2 different directions. For a 2D acquisition only every voxel in a 2D plane of k space provides information for every voxel in a plane of real space. Therefore, $T_{acq,total}$ is reduced by a factor of N_{phase2}.

$$T_{acq,2Dtotal} = N_{phase1} \;\; N_{phase2} \;\; NEX \;\; T_{acq} \tag{2.17}$$

Therefore, SNR increases by a factor of $\sqrt{N_{phase2}}$ for a 3D acquisition.

However, this ignores that a 2D sequence may take less time. In order to get the desired contrast TE may be much shorter than the desired TR (see section 2.2.1). This means there is a lot of idle time during the sequence. For a 3D sequence, this idle time can't be exploited since the whole volume has been excited. However, for a 2D sequence a different slice can be excited during the same TR while the spins in another slice are recovering; allowing time to be used more efficiently.

3D acquisitions also suffer from more truncation artefacts due to being Fourier encoded in all 3 dimensions. Truncation artefacts can be much more difficult to detect in the slice encoding direction than in plane and the image may need to be reconstructed in another plane to detect them.

2.1.10 Frequency Encoding

By applying a gradient during signal acquisition the frequency of resonance of a particular nuclear species depends on position. By Fourier transforming the signal from the time domain to the frequency domain, the signal from different frequencies now corresponds to different locations along the direction of the gradient.

A de-phasing lobe must first be employed to centre the image. Without this only frequencies exclusively greater than, or exclusively less than the Larmor frequency would acquired; depending on the gradient orientation. The integral of the rephasing gradient up to halfway through the acquisition is chosen so it is of equal magnitude and opposite polarity to the integral of the de-phasing gradient. This ensures a properly centred projection. It should be noted, in some spin echo sequences, the rephasing lobe may require being of the same polarity (see section 2.2.2).

The de-phasing lobe has an effect on the whole sample, analogous to an increasing the rate of $T2^*$ (see section 2.1.4). The re-phasing gradient can then be thought of as reversing this effect until each nuclear species will be back in phase; as if there had been no acceleration in $T2^*$ at the middle of the acquisition time (natural $T2^*$ effects will still occur).

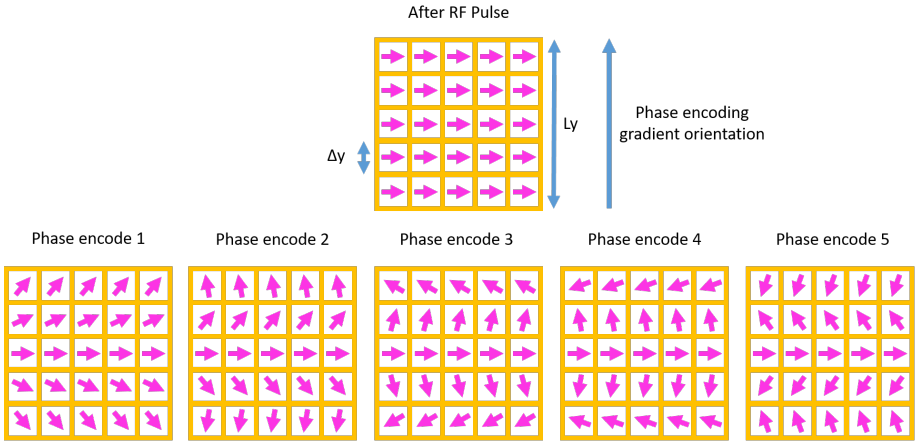

Figure 2.5: Object before (top) and after (right) different phase encoding gradients along the direction of the blue arrow.

2.1.11 Phase Encoding

Phase encoding induces a variation in the phase of the transverse magnetisation parallel to the gradient applied [14] (see figure 2.5).

In order to satisfy the Nyquist criterion and prevent aliasing (see section 2.1.6) the phase encoding step Δk_y should be chosen so that.

$$\Delta k_y = \frac{1}{L_y} = \frac{1}{N\Delta y} \tag{2.18}$$

For a phase encoding gradient in the $\hat{y}$ direction, the precession frequency during a gradient G_y in a B_0 rotating reference frame is given by equation 2.19.

$$\omega = \gamma G_y y \tag{2.19}$$

After integrating over the duration of the phase encoding pulse (T), the phase of the transverse magnetisation is given by equation 2.20.

$$\phi(y) = y\gamma \int_0^T G_y(t')dt' = 2\pi k_y y \tag{2.20}$$

By combining the magnetisation perpendicular to B_0 into a complex number ($M_\perp = M_x + iM_y$), the phase, as well as the magnitude of the signal can be extracted. The projected signal in the $\hat{y}$ dimension of k space can then be given by equation 2.21.

$$S(k_y) = \int M_\perp(y) e^{-i\phi(y)} dy \tag{2.21}$$

Approximating as a discrete sum by setting $y = n\Delta y$, defining N as the number of voxels in the $\hat{y}$ dimension and inserting equation 2.20 into equation 2.21

equation 2.22 is found

$$S(k_y) = \sum_{n=0}^{N-1} M \perp (n\Delta y) e^{-1\phi(y)} dy \tag{2.22}$$

Because each line of voxels in y at point $n\Delta y$ has a different phase, this allows $M \perp (n\Delta y)$ to be extracted. Note, this example has implied Δk_y to be a constant. However, this isn't a requirement, but allows a Fast Fourier transform (FFT) to be applied. For some applications, such as CS Δk_y is not constant (see section 2.1.6).

2.1.12 MRI Imaging Coils

Specific coils can be used to both transmit an RF pulse and receive back the MR signal. Alternatively, sometimes separate coils are used for transmit and receive. It is important that the transmit RF pulse (of many hundred volts peak to peak) is never exposed to the receive electronics, which are used to amplify signals of the order mV. The receive side must be gated to prevent this from happening.

To improve image quality, different coils can be used; depending on the size and shape of what is being imaged. By using a smaller coil, the coil is closer to the origin of the signal and the SNR is therefore greater. Surface coils are also susceptible to suffering from homogeneity issues. Homogeneity issues can result in the coil sending an RF pulse which excites atoms by different amounts depending on their location. The same strength NMR signal from different locations within the region of interest (ROI), may also induce different strength currents within the coil. The MR signal received back may need to be scaled depending on the distance from the coil; by minimising distances it is possible to reduce homogeneity issues.

Surface Coils

A basic surface coil can be produced by creating a coil of wire with a capacitor in parallel (see figure 2.6). The capacitor and the inductance of the coil form a resonant circuit. The resonant frequency of the circuit can be tuned to the frequency of the spins being imaged by changing the value of the capacitor in parallel (C_T). A matching capacitor (C_M) also needs to be added. Without this the coil would have a high input impedance of kilo-ohms. Because of the 50 ohm impedance of the power-amplifier used to amplify the RF pulse, if the impedance of the surface coil circuit is not matched to 50 ohms, then a large proportion of the output power will be reflected directly back.

There is a large drop off with the field produced by surface coils at distances further from the coil; with the homogeneous region of the order of the diameter of the coil. While this limits the size of the region imaged, because the signal is only coming from a small volume; and this volume is the ROI, noise signal from outside the ROI is small. This gives the surface coil a high SNR.

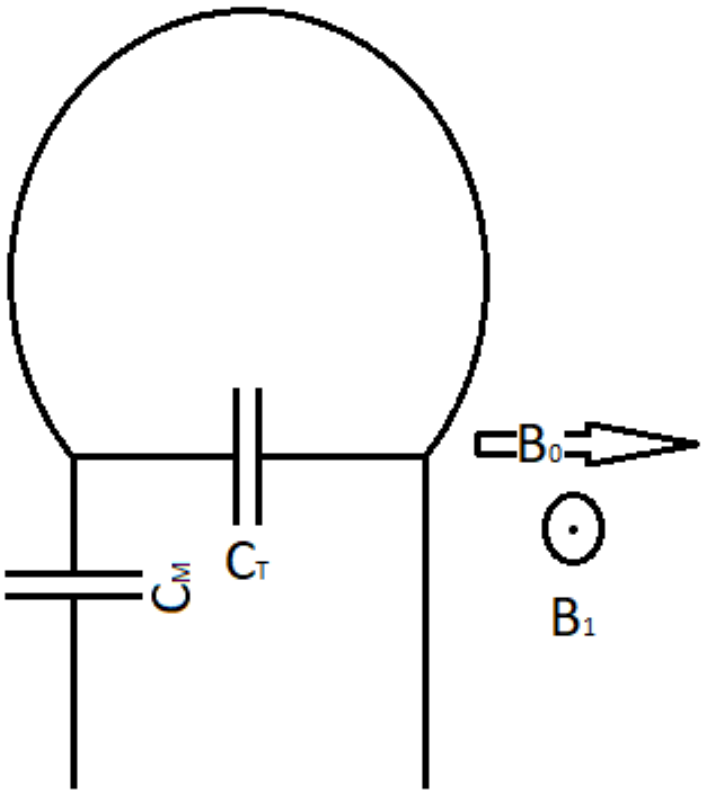

Figure 2.6: Basic surface coil circuit diagram.

Figure 2.7: GRE sequence. RF pulses in Red. Purple are the stimulated echo signals, orange FID signals

Volume Coils

Volume coils have the advantage of surface coils of producing a much more homogeneous field over the region that they encompass. One of the most common volume coil designs is a birdcage coil [28]. In a birdcage coil multiple wires are aligned parallel to B_0 at equal intervals, encompassing the ROI, creating a cylinder. Across the circumference these wires give an sinusoidal current variation, which has been shown to offer an approximately uniform transverse field [29].

2.2 MRI Sequences

The sequence used will depend on multiple factors, such as the resolution required within a specific time period and the weighting of the sequence.

2.2.1 Gradient Field Echo Sequence

In a gradient field echo or gradient echo (GRE) sequence, RF pulses are separated by a defined repetition time (TR). Two of these RF pulses together generates a spin echo and three together will generate a stimulated echo (see figure 2.7). De-phasing gradients and re-phasing gradients can be used to manipulate either the echo or FID, which allows the signal in question to first be suppressed using the de-phasing gradient, before refocussing the signal using a re-phasing gradient. By applying a de-phasing gradient in one dimension, this results in the same nuclear species precessing at slightly different rates across the sample.

The transverse coherences produced during a GRE may either be favourable or unfavourable depending on the purpose of the sequence. Those which exploit these coherences are often referred to as coherent GRE sequences. For the work in this document with GRE sequences (see sections 6.1, 6.3 and 7.1.3), sequences where ideally no transverse magnetisation would remain from one TR to the next were of interest. These are often referred to as spoiled GRE sequences.

Spoiled GRE

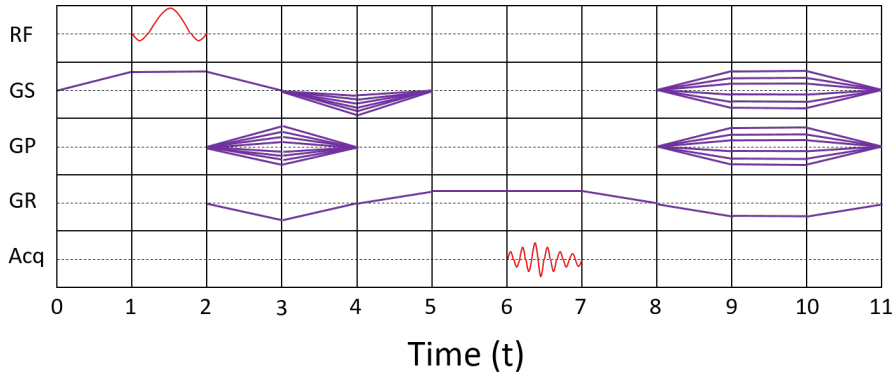

Figure 2.8: GRE sequence, showing the gradient timings during one TR. Note that the event lengths are arbitrary and not necessarily equal.

Figure 2.8 shows an example pulse profile of a 3D spoiled GRE sequence. In this explanation GR, GP and GS are assumed to be orthogonal. GS between time (t)=1 and t=2 determines the thickness of the slice being excited during the RF pulse. Following the RF pulse, there is a phase encoding gradient on GS between t=3 and t=4. After the RF pulse, there is one other phase encoding gradient, this time on GP between t=2 and t=3. The third gradient dimension is the readout gradient (GR). The de-phasing gradient is applied after the RF pulse (t=2 to t=4). Following this, an opposite polarity gradient is applied, the re-phasing gradient (t=4 to t=8) (see section 2.1.10). In the case of a spoiled GRE, the point at which the phases have been refocussed will be half way through the acquisition period (t=6 to t=7).

Spoiling can be employed to prevent build up of transverse magnetisation between RF pulses. GS, GP and GR from t=8 to t=11 are the spoiling gradients in figure 2.8. Spoiling ensures that any transverse components to the magnetisation are removed, even if the magnetisation in the $\hat{G_s}$ axis has not returned to its steady state level.

Spoiling naturally occurs if the TR is very long (TR>>T2*); this allows the transverse magnetisation to decay to zero. However, having a long TR may not be practical. It not only increases the length of the sequence, but also, for a GRE sequence, results in a sequence weighted towards T2* and proton density. For sequences which require a shorter TR, gradient or RF spoiling (or both) can be applied to ensure adequate spoiling.

Gradient spoiling [30] is performed by applying gradients after the acquisition, but before the next RF pulse. There are two options: Firstly, residual coherent transverse magnetisation can be refocussed. The spoiling gradient's intensity should be equal and of opposite polarity to the phase encoding gradient. Alternatively, the spoiling gradient can be varied (generally randomly or linearly) between pulses. Gradient spoiling can create ideal spoiling in much of the ROI. However, because the gradient isn't the same strength across the

sample, the spoiling isn't uniform across the sample. This leads to banding artefacts where the spoiling fails.

RF spoiling [31] is performed by changing the phase or timing of the RF pulse, to prevent unintended coherences. This can result in a much more uniform spoiling across a volume; assuming uniform transmit.

Optimal Parameters of Sequence, Ernst Angle and Sequence Weighting

For a perfectly spoiled GRE sequence which has achieved a longitudinal steady-state, the signal is given by equation 2.23 [32].

$$S = A[H]\frac{sin\alpha(1 - e^{-TR/T1})}{1 - (cos\alpha)e^{-TR/T1}}e^{TE/T2^*} \tag{2.23}$$

Where A is a constant and [H] represents the proton density. Equation 2.23 is maximised when $\frac{dS}{d\alpha} = 0$ and is obtained with a FA equal to the Ernst angle (α_E) [14] (see equation 2.24).

$$\alpha_E = cos^{-1}(e^{TR/T1}) \tag{2.24}$$

If the FA chosen is too high, then the longitudinal magnetisation won't have enough time to recover and the signal received will drop. Conversely, if a much lower FA is used, the longitudinal magnetisation will have more than sufficient time to relax back and the signal won't be exploited fully.

However, in practice while the Ernst angle may offer the most signal, it may not be the optimum FA to obtain maximum contrast between tissues, even if those tissues have the same proton density (see figure 2.9).

If a sequence is weighted towards a certain parameter (P) it simply means that a small change in that parameter will have a large effect on the signal received back. i.e. maximise $\frac{\Delta S}{\Delta P}$.

[H] and T2* Weighting Looking at equation 2.23, for $\alpha \to 0$, $cos\alpha \to 1$. Therefore, for small alpha:

$$\frac{sin\alpha(1 - e^{-TR/T1})}{1 - cos(\alpha)e^{-TR/T1}} \to sin\alpha \tag{2.25}$$

This removes any T1 effects from equation 2.23. Therefore, for small FAs [H] and T2* weighting dominate and T1 weighting is minimised.

T2* Weighting Looking at equation 2.23 $T2^*$ only appears in the term $e^{\frac{TE}{T2^*}}$. Therefore, T_E controls the $T2^*$ weighting. After absorbing all other elements into constant A^*, equation 2.23 becomes equation 2.26. Differentiating by $T2^*$, equation 2.27 is found.

$$S = A^* e^{-(TE/T2^*)}. \tag{2.26}$$

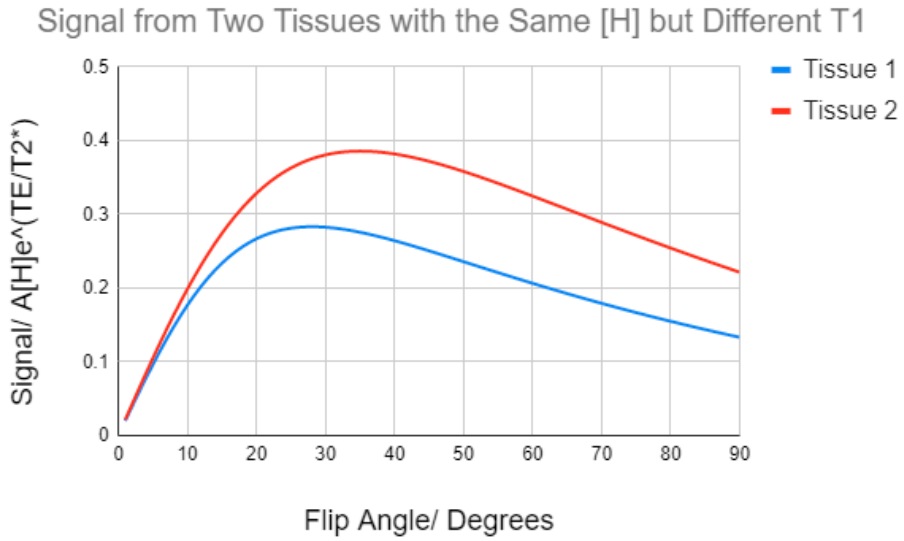

Figure 2.9: Signal comparison between two tissues, with the same proton density, but different T1s, using a perfectly spoiled GRE sequence. $\alpha_E = 28°$ or $\alpha_E = 35°$ for tissue 1 and tissue 2 respectively (Tissue 1 $T1 = (8/5)$Tissue 2 T1). However, the maximum contrast is obtained using a FA of 49° (see equation 2.23).

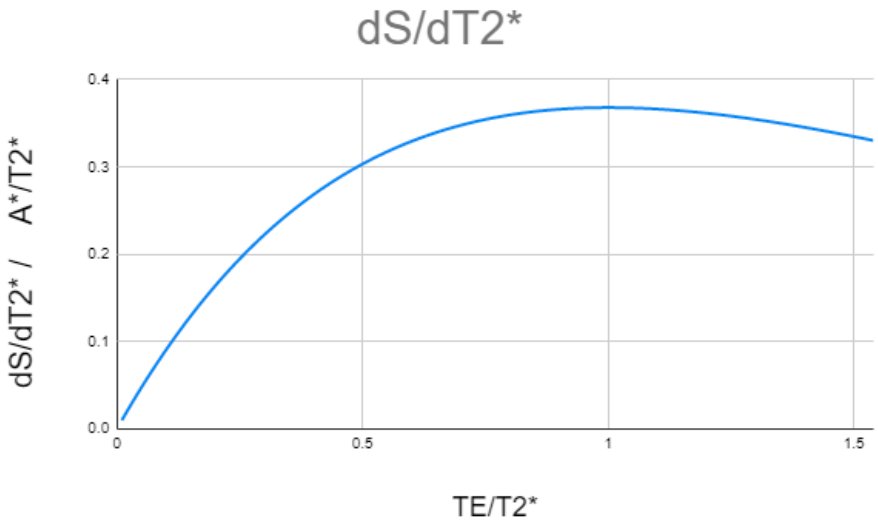

Figure 2.10: Change in signal per small change in the parameter $T2^*$.

$$\frac{dS}{dT2^*} = A^* \frac{TE}{(T2^*)^2} e^{-(TE/T2^*)} \tag{2.27}$$

By plotting equation 2.27 it can be seen that $\frac{dS}{dT2^*}$ is maximised for $TE = T2^*$ (see figure 2.10). However, the preference is for longer values of TE, since the gradient of the line is very large for $TE < T2^*$. For very small values of TE, the $T2^*$ weighting is minimised.

T1 Weighting Looking at equation 2.23, $T1$ Only appears in the term $e^{\frac{TR}{T1}}$. Therefore, TR controls the $T1$ weighting. After absorbing all other elements into constant A_1, equation 2.23 becomes equation 2.28. Differentiating by $T1$, equation 2.29 is found.

$$S = A_1 \frac{1 - e^{-(TR/T1)}}{1 - cos(\alpha) e^{-(TR/T1)}} \tag{2.28}$$

$$\frac{dS}{dT1} = A_1 \frac{(1 - cos(\alpha))}{T1} \frac{\frac{TR}{T1} e^{-(TR/T1)}}{(1 - cos(\alpha) e^{-(TR/T1)})^2} \tag{2.29}$$

By plotting equation 2.29, it can be seen that for larger TR there is less T1 weighting (see figure 2.11 and 2.12). Changing the value of FA (α) also changes the T1 contrast. Ideally $TR << T1$ would be used with a large FA. However, for small TR and large α, there isn't enough time for the polarisation to recover between the TRs. The signal of the overall image is suppressed and so is the T1 weighting.

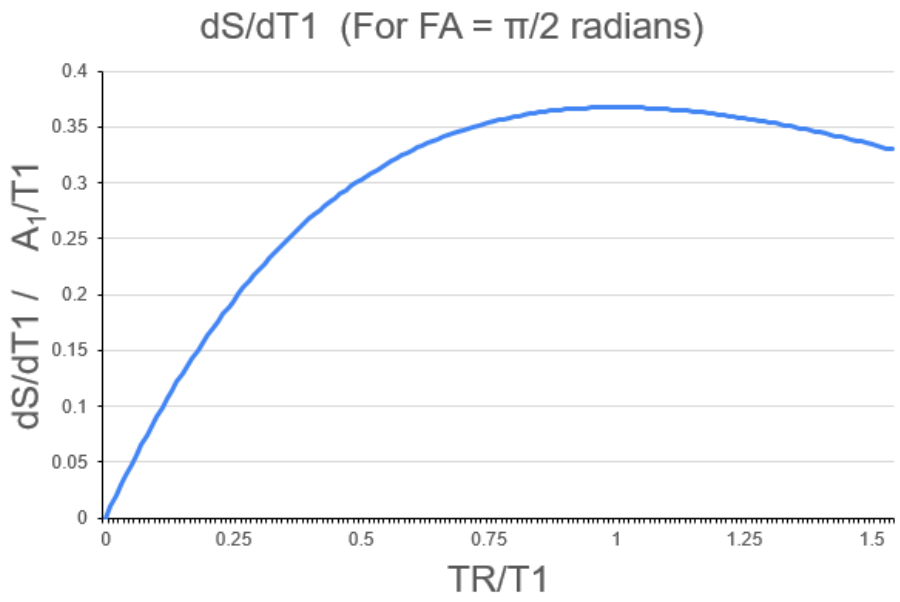

Figure 2.11: Change in signal per small change in the parameter $T1$ for large FA.

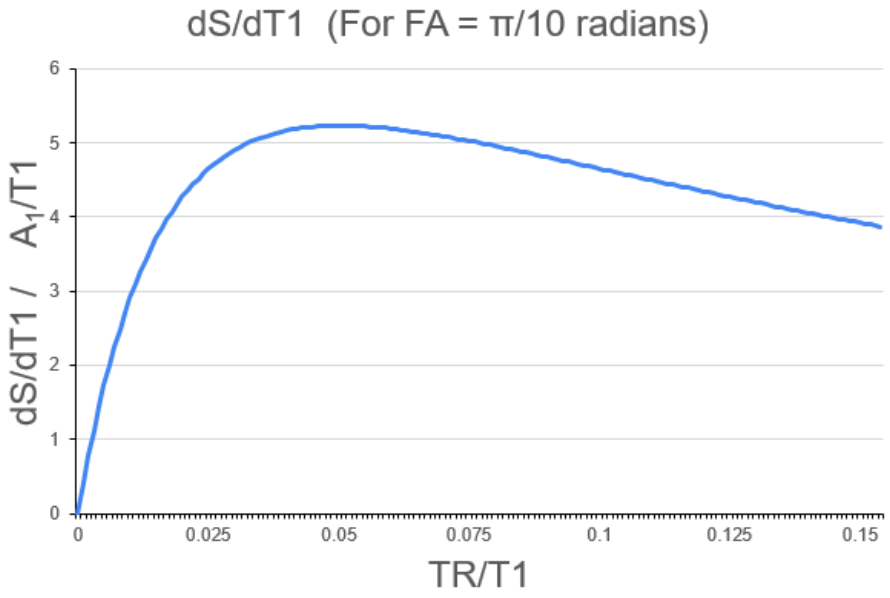

Figure 2.12: Change in signal per small change in the parameter $T1$ for small FA.

2.2.2 Spin Echo Sequence

Spin echo (SE) sequences were first introduced by Hahn in 1950 [33] and have been used in multiple applications, from quantum computing to MRI [34]. In a SE sequence, an RF pulse is first used to excite the magnetic moment of the protons. If a 90° pulse is used, the magnetic moment of the protons is excited into the equatorial ($\hat{x}$, $\hat{y}$) plane. After a period of time (TE/2, measured from the centre of the RF pulse), a 180° pulse is applied which acts to reverse the precession of the spins. TE after the centre of the RF pulse, the spins are refocussed. If spin-spin interactions were discounted and only reversible contributions to transverse relaxation were taken into account (see section 2.1.4), then the spins would perfectly refocus. Instead, after a time TE, the signal has reduced by a characteristic time constant T2 (assuming transverse relaxation dominates). This illustrates why SE sequences were the most common sequences until the early 1980s. At this time the magnetic fields of scanners were relatively inhomogeneous. By using a SE sequence, reversible transverse relaxation was discounted and therefore avoided the effects from having an inhomogeneous magnetic field.

2.3 Hyperpolarised MRI

2.3.1 What is Hyperpolarisation?

In conventional MRI, protons are thermally polarised according to Equation 2.4. After a long period of time (t>>T1) there is a small imbalance in the number of nuclei in the $N_\uparrow$ and $N_\downarrow$ states. This imbalance is referred to as the thermal polarisation. The thermal polarisation is given by equation 2.30

$$P_{thermal} = \frac{N_\uparrow - N_\downarrow}{N_\uparrow + N_\downarrow} \approx \frac{\Delta E}{2kT} = \frac{\gamma \hbar B_0}{2kT} \tag{2.30}$$

With hyperpolarisation, the nuclear species in question is manipulated so a much greater proportion of the magnetic moments are aligned with the magnetic field; in some cases > 90% [35]. This can result in a polarisation more than 10,000 times greater than the thermal polarisation. For example, at the temperature of the human body, the polarisation of protons in water at 1.5T is ∼0.0005%, 10,000 times less than a 5% polarisation of ^{129}Xe which is easily achievable [7]. However, this high level of hyperpolarisation is not stable and will decay back to the thermal polarisation over time.

2.3.2 Why Hyperpolarised MRI?

The reasons for using a hyperpolarised sample fall into two major categories. Either the ^{1}H signal from a tissue is low, or the signal isn't specific; making it difficult to interpret what the signal means.

Not all parts of the body have a high concentration of protons and these regions therefore produce a low MR signal. By hyperpolarising a nuclear species,

even if the hyperpolarised sample is in small concentrations, the signal as a whole can still be large. This is is a consequence of the MR signal being many orders of magnitude greater per atom through hyperpolarisation than thermal polarisation (see section 2.3.1). An example of this imaging is hyperpolarised ^{129}Xe structural lung imaging (section 5).

With ^{1}H MRI theoretical models can be used to track dynamic processes in the body. For example, in BOLD imaging, sequences heavily weighted towards T2* are employed. Because deoxygenated haemoglobin (Hb) is paramagnetic and oxygenated Hb is not, deoxygenated Hb will result in an accelerated T2* [36]. In the brain, after an initial drop in oxygenated Hb, blood flow increases bringing a surplus of oxygenated Hb and a theoretical decrease in T2* [37]. However, other processes in the ROI may also affect T2* and the increase in oxygenated Hb may vary between people. The sequences are only **weighted** towards T2*; some of the signal will also have contributions due to T1 and ^{1}H density. All these factors mean the BOLD signal is not a direct measure of blood flow to the brain and it can therefore be difficult to interpret the signal.

By introducing a hyperpolarised sample of a different nuclear species than ^{1}H, a much higher specificity can be achieved. Because the nuclear species isn't present in its hyperpolarised state elsewhere, the signal should only be obtained from the sample introduced. If the hyperpolarised species stays polarised for long enough it can also be tracked through the body; allowing for dynamic imaging.

Alternatively, molecules which are already present in the body can be hyperpolarised. Using dynamic nuclear polarisation, molecules involved in biochemical processes can be hyperpolarised (see section 2.3.4). Because the signal from the molecules introduced is so much greater than those naturally present, the journey of the hyperpolarised molecules can be tracked through the body and biochemical processes can therefore be investigated.

2.3.3 Hyperpolarised Lung Imaging

Hyperpolarised Lung Imaging conventionally uses noble gases. In order to image the new species, the scanner receive and transmit must be be tuned to the Larmor frequency of the new species (rather than ^{1}H). If a ^{1}H MRI sequence was then converted for use by a different species, the gradients would have to be scaled by the ratio of the gyromagnetic ratios (see table 2.1). For example, for ^{129}Xe the gradients would be increased by a factor of 3.62 to produce an equivalent sequence.

Imaging using a hyperpolarised gas is particularly favourable in the lungs due to the low proton density (about 1/3 of muscle) [38]. This low ^{1}H density is why the lungs appear dark in a conventional MRI sequence. If Xe in its natural form was delivered to the lungs the signal produced would be undetectable. This is because the Xe delivered is in a gaseous form and therefore low density; meaning the signal produced by thermal polarisation alone is too small. Instead the gas is first hyperpolarised.

The polarisation of ^{129}Xe is not induced by the scanner's magnetic field.

Therefore, lower field magnets can be used without the normal loss in SNR [27]. Lower field scanners are generally much cheaper to buy, run, site and maintain. Also, they have less subject induced magnetic field inhomogeneities and produce less acoustic noise [39].

However, the main disadvantage of hyperpolarised MRI is that the polarisation is non recoverable. In conventional MRI, after a group of protons is acted on by the scanner's RF pulse (see section 2.1), they will return back to a thermal equilibrium polarisation in the scanner's magnetic field; at a rate determined by the T1 of the group of protons. However, in hyperpolarised MRI, once the hyperpolarisation has been lost it won't return.

Different Hyperpolarised Gases

There are multiple candidates for the gas which is most favourable for hyperpolarisation; the majority of which are noble gases. However, Radon, due to its radioactive properties is not a candidate. Due to ^{129}Xe's availability, high levels of polarisation achievable and large chemical shift combined with its solubility, it will be the hyperpolarised species focussed on in this thesis.

Helium-3 Helium-3 (^{3}He) is almost completely inert and non toxic, making it an ideal candidate; it can be administered to a volunteer multiple times [40]. ^{3}He also has a larger γ than ^{129}Xe (see table 2.1); making its NMR signal $\sim$ 12 times stronger [7]. High levels of hyperpolarisation of ^{3}He have been available for years, meaning historically most research into medical imaging using hyperpolarised gases has been devoted to ^{3}He [7].

However, there is a global shortage of ^{3}He [41]. The main method of ^{3}He production is from the decay of tritium in nuclear warheads and since 2001 the demand for ^{3}He has skyrocketed as thousands of neutron detectors (used to detect plutonium being imported) were deployed by the US [42]. As the supply has been outstripped by demand, the cost has increased, making the use of ^{3}He expensive. Also the US is regulating supply. This means that while there may be enough to perform research into ^{3}He medical imaging, unless a new source of ^{3}He is found, the supply wouldn't be enough to meet demand if a technique was rolled out into a clinical setting.

Xenon-129 ^{129}Xe may have a lower γ than ^{3}He leading to a lower sensitivity. However, due to its much larger electron cloud, ^{129}Xe has a large chemical shift providing information on ^{129}Xe's local environment [43].

^{129}Xe's historical disadvantage over ^{3}He of a lower polarisation, has now been overcome, with litre quantities of $> 50\%$ polarisation available for ^{129}Xe [44].

^{129}Xe is also available in a relatively high abundance in comparison to ^{3}He, making it not only a cheaper option, but also a more viable option if a technique were to be used clinically on a much wider scale. 26% of naturally occurring Xe is ^{129}Xe, meaning that isotope enriched Xe is not always required if a high polarisation of ^{129}Xe can be achieved.

However, Xe can also act as an anaesthetic in high concentrations. This means there are limits on the dosage that can be given to a patient. S. Sivaram Kaushik et al. [45] found that 91% of their subjects experienced mild side effects after inhaling a 1L dose of Xe; such as dizziness, paraesthesia and euphoria. These all resolved within $2-3$ minutes.

^{129}Xe is also very good at diffusing through capillary walls and has a relatively high solubility in biological tissues [46]. This fact, alongside the large chemical shift, makes ^{129}Xe an ideal contrast agent for investigating gas exchange in the lungs [47]. After ^{129}Xe is breathed in, exchange is continually happening between ^{129}Xe in the lung-space and that dissolved in the blood. This results in multiple peaks in an MR spectrum. The main 3 peaks (in descending magnitude) are attributable to ^{129}Xe in gas in the air spaces, ^{129}Xe dissolved in lung parenchyma/plasma and ^{129}Xe dissolved in red blood cells. Only about 2% of the ^{129}Xe is dissolved in the blood. Wagshul et al. [48] showed that the highest intensity peak shows correlated intensity and frequency oscillations, which are attributable to changes in lung volume during breathing. The other two major peaks are maximal at 5 to 10s and have a T1 of 30s. Because exchange is continually happening between ^{129}Xe in the lung-space and that dissolved in the blood, if the dissolved phase peak is "killed" it will slowly recover as fresh hyperpolarised ^{129}Xe diffuses in from the lungs. This allows the rate at which gas exchange happens between the lungs and blood to be investigated. There is also an additional peak which takes 20 to 30s to reach maximum, attributable to well vascularised tissue such as heart and muscle.

Quadrupolar Isotopes The other three stable isotopes of noble gases that are NMR active are ^{21}Ne (I = 3/2), ^{83}Kr (I = 9/2), ^{131}Xenon (I = 3/2). Additional information on the shape, size and symmetry of void spaces and the chemical composition of surfaces in porous media can be obtained due to the quadrupolar interactions [49][50][51][52][53][54]. However, the quadrupolar gases all exhibit another method of polarisation relaxation; quadrupolar relaxation. In order for a gas to offer a viable solution for imaging the lungs, the relaxation time should be of the same order of magnitude, or greater than, the time scale of gas transport into the lung. For ^{83}Kr, the relaxation times are shorter than that of ^{129}Xe and ^{3}He, but of the same order of magnitude as a breathing cycle of 10s. Pavlovskaya et al. [49] found the T1 of ^{83}Kr at 9.3T to be 10.5s. In order to image the lungs in vivo, O_2 must also be present. While like in ^{129}Xe and ^{3}He the relaxation time of ^{83}Kr is reduced in the presence of O_2, it did so to a much lesser extent, due to the relaxation of ^{83}Kr being dominated by quadrupolar interactions. However, Pavlovskaya et al. [49] did observe a dependence of T1 on the magnetic field, with the T1 reducing to 7s at 1.5T; making the use of ^{83}Kr less viable for lung imaging at lower field strengths.

Propane Through parahydrogen-induced polarisation, hyperpolarised propane gas can be produced. Salnikov et al. [55] produced this gas at 1-2% polarisation. Hyperpolarised propane does carry 2 polarised ^{1}H nuclei compared to one

nucleus for a polarised noble gas. Also the gyromagnetic ratio is greater for ^{1}H than ^{129}Xe; the NMR sensitivity is ∼13 times greater. This puts 1% polarisation for propane equivalent to ∼26% ^{129}Xe polarisation. The hyperpolarised propane can also be produced very quickly; 0.3L in 2s [55]. The major advantage of using propane rather than a noble gas is a conventional MRI scanner can be used at the ^{1}H frequency and no expensive bespoke coils or amplifiers are required. However, propane is flammable and isn't currently approved for use as an inhailable contrast agent.

2.3.4 Dynamic Nuclear Polarisation

Dynamic nuclear polarisation (DNP) allows for the hyperpolarisation of ^{13}C and ^{14}N in liquid substances. By cooling liquids to very low temperatures (∼1K) and irradiating with microwaves, polarisations of 36% of ^{13}C have been achieved in <40 minutes [56]. The sample can be brought up to room temperature quickly by dissolving in a hot liquid. The sample can then be transported to the magnet and injected in vivo. This transfer time limits in vivo use. Bringing the sample up to room temperature reduces the polarisation and the T1 of the sample is also short typically of the order <20s for a urea sample [57].

^{13}C makes up typically ∼ 1% of organic carbon [58] and the organic carbon already in the body is only thermally polarised. Therefore, the ^{13}C signal from the hyperpolarised tracer introduced dwarfs that of the naturally occurrent signal. By polarising substances involved in metabolic processes such as pyruvate, the rate at which metabolic processes occur can be investigated [59].

2.4 Fluorine-19 MRI

^{19}F is a spin 1/2 nuclei with a high gyromagnetic ratio $\gamma = 40.05\text{MHzT}^{-1}$ (see table 2.1). ^{19}F also has a high sensitivity, 0.84 times that of ^{1}H at a constant magnetic field and composes 100% of naturally occurrent fluorine [60].

Unlike ^{1}H, there are negligible concentrations of ^{19}F in the human body; usually less than 103 μmol/g wet tissue weight [61]. This is below the detection limit from MR. Instead a ^{19}F tracer can be introduced. There are multiple fluorinated compounds in the broad group of CFCs which are inert and non toxic. For example, sulphur hexafluoride (SF_6), hexafluoroethane (C_2F_6) and Perfluoropropane (C_3F_8). By mixing with O_2 and breathing in the gas, it is possible to detect a signal in vivo from SF_6 [62], C_2F_6 [63] and C_3F_8 [64]. Just like with hyperpolarised ^{129}Xe imaging, this gives very high specificity, since signal only comes from where the introduced tracer reaches. Because the gas mix can be free breathed (unlike the Xe which is an aesthetic); functional data can be obtained on wash in/ wash out kinetics [11]. The fluorinated compounds are also significantly cheaper than hyperpolarised ^{129}Xe, when polarisation costs are factored in. That is before accounting for the cost of Clinical Trials of Investigational Medicinal Products approval, which is many times more than the cost of the gas itself.

Even though imaging ^{19}F gives high specificity, the signal is very low. Gaseous compounds have a low density of atoms and hence the overall MR signal is low. Fluorinated compounds such as Perfluoropropane (PFP) are chosen, because of the short T_1 of the ^{19}F nuclei ($\sim$ 12ms at 3.0T [64] [65]). This allows for short TR pulse sequences and consequently more time to acquire multiple averages of the ^{19}F sample. By acquiring multiple averages, sufficient signal to form an image can be obtained. However, the short T_1 of ^{19}F in PFP does make PFP impractical to hyperpolarise (the polarisation quickly returns to equilibrium).

Using a 3D spoiled GRE sequence, CS Neal et al. [10] have been able to acquire ^{19}F ventilation images in a single breath hold (18s), at 1cm isotropic resolution. They have also demonstrated that pushing to shorter breath holds (4.5s) is possible.

See section 7.2 for feasibility tests of ^{19}F imaging on the 0.5T upright scanner.

2.5 Lung Disease States

Many lung disease states are characterised as being both chronic, with acute exacerbations and progressive [1]. For example, in a patient with Chronic Obstructive Pulmonary Disease (see section 2.5.2) a sudden worsening of symptoms may be brought on by an infection or airborne pollutant. For many diseases there is no cure currently.

In progressive disease states, early and accurate diagnosis is important. Early diagnosis allows treatment to be sought to either slow down or reverse the progression of disease. Accurate diagnosis makes sure the treatment being provided is the correct one and makes patients less discontent [66]. Accurate imaging, as well as functional lung data, can provide useful information in the diagnosis of a disease. If performed repeatedly over time, they can provide insight into the progression of a disease to investigate if a treatment is working.

2.5.1 Imaging and Quantifying Lung Function

If a new technique is going to be used by clinicians, it must supersede the current methods. Currently the "Gold Standard" for imaging is Computerised Tomography (CT). A variety of techniques are used for quantifying lung function; the most common of which is spirometry. Different methods of imaging and quantifying lung function are compared below:

Pulmonary Function Tests

The results from pulmonary function tests vary depending on effort levels and the extent to which the patient understands the task [2]. Also, while the values obtained can be useful to track the progression of a disease, they are global measurements for the whole lungs and provide no local information about specific parts of the lungs. If only a small region of the lungs is affected, this may not be picked up with pulmonary function tests.

Spirometry is the most common method of obtaining functional data from the lungs [2]. Compared to other pulmonary function tests, it is simple and quick to perform. It can provide information on the forced vital capacity and forced expiratory volume [67]. These measures allow obstructions and restrictive defects to be identified.

Measurement of absolute lung volumes is more difficult and can't be performed through spirometry [68]. Whole body plethysmography performed in an airtight environment (with a tube to breathe through), can be used to determine the forced residual capacity, total lung capacity and residual volume of the lungs. The patient is asked to perform breathing exercises, as the breathing tube is transiently blocked. The change in pressure of the airtight environment and the flow rate through the breathing tube are measured [69]. A high residual volume can be a sign of obstructive lung disease; where the lungs are incompletely emptied. Total lung capacity is sometimes increased in patients with emphysema and decreased in patients with restrictive abnormalities [2].

The rate of diffusion of gas from the lungs into the blood can also be investigated through measurement of the transfer factor. Gas with $\sim$10% He and $\sim$0.3% CO is quickly inspired and the patient asked to hold their breath. The reduction in the concentration of CO is measured in the expired gas. CO diffuses into the blood, but He doesn't, allowing the transfer factor to be calculated. Patients with a reduced surface area of the alveolar membrane, will have a reduced transfer factor [70] [2].

Due to the breathing exercises required, patients may become light headed during pulmonary function tests and particularly ill patients may struggle with the physical requirements of the exercises. If this is the case the patient is unlikely to provide repeatable results [71].

Computerised Tomography

CT scans can provide high contrast, high resolution images of the lungs. CT scans do use ionising radiation, with a full chest CT providing an approximate effective radiation dose of 7mSv; the equivalent of 2 years of background radiation [72]. This additional radiation exposure must be taken into account by medical professionals before performing the scan.

Due to lung diseases often being characterised as chronic with acute exacerbations, multiple investigations are often required. This makes a technique with ionising radiation unattractive; especially one with such a high dose

Fluoroscopy

In fluoroscopy, a continuous X-ray beam is passed through the body allowing motion within the body to be viewed [73].

Using chest fluoroscopy, a sniff test can be performed. The patient is asked to quickly breathe in through their nose, while the diaphragm movement is monitored throughout the inspiration. This allows the absence of muscular contraction of the diaphragm during inspiration, in patients with phrenic nerve

palsy, or those who are suffering with breathing problems post stroke, to be confirmed [74].

^{1}H MRI

In tissues other than the lungs, ^{1}H MRI is able to compete with CT in terms of spatial resolution and beats it in terms of temporal resolution and contrast resolution [4]. However, the signal from ^{1}H MRI is derived from protons in water and fat. Since the lungs are mainly empty space they have low signal on a ^{1}H MRI image (see section 6.1).

2.5.2 Chronic Obstructive Pulmonary Disease

Chronic Obstructive Pulmonary Disease (COPD) is an umbrella term, which encompasses many conditions. Patients may have dominant features of chronic bronchitis, emphysema or asthma, with contributions from one or both of the others [75]. Patients with COPD have both impaired flow as well as alveolar destruction and enlargement. COPD places a heavy burden on health services and is now the third leading cause of death in the USA [76]. Diagnosis is made with spirometry, when the ratio of forced expiratory volume in 1 second, over forced vital capacity (FEV1/FVC), is less than 70% of that predicted for a matched control. It is diagnostic for a significant obstructive defect.

2.5.3 Diaphragmatic Disorders

Diaphragmatic disorders such as hernias [77], tumours [78] and paralysis [79] can have a severe impact on someone's respiratory health. Diaphragm function may also be impaired in other respiratory diseases such as COPD [80]. Fluoroscopy or ultrasound can be performed to investigate the movement of the diaphragm. However, ultrasound is relatively low resolution and very operator dependent [81] and fluoroscopy is both ionising and acquires a continuous slice through the body.

2.5.4 Chronic Bronchitis

Defined clinically as the presence of a chronic productive cough for 3 months during each of 2 consecutive years (other causes of cough being excluded) and airflow obstruction [82]. The obstruction is caused by excessive tracheobronchial mucus production. Damage to the endothelium impairs the mucociliary response. This, in combination with inflammation, obstructs breathing. Smoking is the most important risk factor in developing the disease, with over 90% of patients having a smoking history [83]. Chronic bronchitis is different from acute bronchitis; with respect to the bacterial pathogens found in the lower bronchi. The 3 leading bacterial pathogens, in patients with chronic bronchitis, do not cause acute bronchitis [84]. One theory is that a small colonisation by infectious agents can trigger a small inflammatory response, which itself triggers

subsequent acute exacerbations [85]. Unlike emphysema, the pulmonary bed is largely unaffected.

2.5.5 Emphysema

Defined as permanent enlargement of airspace distal to the terminal bronchioles. This results in a greatly reduced surface area of alveoli, impairing gas exchange. Furthermore, loss of the alveolar walls leads to airway narrowing, due to lack of a support structure, as well as decreased elastic recoil; both of which make breathing harder [86].

2.5.6 Lung Transplant

The main complications that can result from a lung transplant are rejection and infection. The transplant lung is a foreign object in the body and even with attempts to match genetic markers, the body's immune system creates antibodies to fight the foreign object (lung). Even with immune suppressing medication, rejection can happen and it is important to detect this as soon as possible; early diagnosis makes treatment more effective and reduces morbidity [87]. Furthermore, the powerful immune suppressant medication, means the body is more susceptible to infection, which places lung transplant patients at higher risk of other complications.

2.5.7 Cystic Fibrosis

Cystic fibrosis (CF) has an incidence of 1 in 2500 in newborns in white populations; with 7000 people in the UK currently living with the disease [88]. Recent developments in CF treatment has resulted in a predicted median survival rate of 50 years for those born in the 21st century [89]. CF is a genetic, recessive disease caused by mutations in the CFTR gene. 70% of patients have a deletion of phenylalanine at codon 508, but over 1600 mutations of the CFTR gene have been described. Different mutations can result in different phenotypes of the disease. CF affects many organs, but primarily affects the lower airways, pancreas, bowel, and reproductive tracts [90]. For most patients, due to the symptoms and the consequently required treatments, lung disease is the most important. Doing physiotherapy and taking inhaled drugs, such as antibiotics can take more than an hour a day; even in periods of good health. In patients with CF, cilia in the lung fail to clear mucus. It is commonly accepted that this is due to a lower than normal volume of airway surface liquid (the low volume hypothesis) [91]. The failure to remove mucus means patients with CF can't effectively clear bacteria inhaled in. Patients with CF also possess an increased inflammatory response to these bacteria, with up to 10 times more inflammation for a given bacterial load. Imaging a patient's lungs over time could provide information on the effectiveness of a specific treatment for a given patient and could provide additional information of the progression of the disease. CF can be diagnosed using a measurement of sweat electrolyte levels, since virtually

all patients with the disease have raised concentrations of sodium chloride [92]. Newborns can be screened using the Guthrie blood spot test, which has been rolled out in the UK [93].

2.5.8 Asthma

Asthma affects almost 20 million people in the United States and more than 300 million people worldwide [94]. It is characterised by varying levels of bronchoconstriction, airway hyper-responsiveness, mucus secretion, and chronic inflammation. Due to the large number of patients affected (11.6% of 6-7 year old children worldwide [95]), asthma around the world, accounts for billions of dollars of health care expenditure [96]. Asthma is more prevalent in urban areas than rural ones [97] and exposure to pollutants such as carbon particulates is also a risk factor [98]. Ideally, treatment should be targeted and specific, which can make putting all asthma patients into one group problematic. Dividing patients into clusters with different inflammatory or phenotypic profiles has proved difficult, with no precise definitions of the clusters. However, most patients broadly fall into one of five clusters [99]. While treatment of asthma is not currently targeted at patients in different clusters, it is hoped that by better defining the different clusters within asthma, different treatments can be developed.

2.5.9 Idiopathic pulmonary fibrosis

Idiopathic pulmonary fibrosis (IPF) is a chronic lung disease in the interstitial lung diseases category, which are characterised as being diseases with pulmonary inflammation and fibrosis [100]. IPF is more prevalent in males and diagnosis rates in the UK have increased recently, from 9 per 100000 in 2004, to 12 per 100000 in 2012. As of 2012, about 32500 people had IPF in the UK and 5292 people died from IPF in the same year [101]. Diagnosis of IPF is time consuming because it is done by "exclusion", whereby other disease states are discounted before the final diagnosis of IPF is arrived at. Diagnosis is done by a physical examination and IPF typically occurs in adults of over 50 years in age. Symptoms include a dry (non-productive) cough and crackles in the chest, which are most prevalent at the lung bases. Risk factor for IPF include, but are not limited to smoking [102][103], metal dust and wood dust [104].

Patients who present to the doctor virtually all already have an abnormal chest radiograph, with evidence being seen retrospectively in many patient's chest CTs taken years before presentation [105]. Therefore, even if a new technique is developed which allows for earlier diagnosis, the technique would have to be applied to people who have a higher risk factor before they would present on their own in order to be effective. Where a new technique could be useful, is in more effective/ accurate diagnosis and better tracking of the disease. High resolution CT already offers increased accuracy in diagnosis over a chest radiograph [106]. However, MRI would have the advantage of not exposing the patient to radiation.

2.5.10 Lymphangioleiomyomatosis

Lymphangioleiomyomatosis (LAM) affects almost exclusively females. It is rare, but affects almost 30% of females with tuberous sclerosis complex [107] (a genetic disorder characterised by the growth of multiple benign tumours). LAM is characterised by a rapid increase in the number of smooth muscle-like cells. This leads to cystic destruction of the lung and lymphatic abnormalities in the chest and abdomen. Over time, patients suffer from reduced lung function, which can lead to breathlessness and lung collapse [108].

Nottingham is home to the National Centre for LAM. This hopes to offer a comprehensive clinical service for patients, in conjunction with a hub for clinical trials and laboratory research investigating LAM.

2.6 Polarisation Techniques

2.6.1 Brute Force Polarisation

Brute force polarisation is the process that MRI scanners use to polarise Protons (see equation 2.4). By placing a nuclear species in a magnetic field a small amount of polarisation is achieved. By cooling to very low temperatures, much higher levels of polarisation can be achieved.

For a spin 1/2 system (such as ^{129}Xe or ^{1}H) the polarisation is given by equation 2.30. Energy levels are split by the magnetic field in Zeeman splitting (in a spin 1/2 system into 2 levels). In the high temperature limit, these two levels are approximately equally populated. However, as the sample is cooled, electrons favour the lower energy level of this splitting (see section 2.1).

For ^{129}Xe, brute force polarisation isn't a viable option, since the temperatures required are of the order of a few Kelvin at which Xenon is a solid. This not only means Xenon takes a long time to polarise; multiple hours even when using ^{3}He as a substrate to aid relaxation [109], but also makes it unsuitable for use in a medical setting because of the expense of cooling ^{129}Xe to these low temperatures [110].

2.6.2 Metastability Exchange Optical Pumping

Metastability Exchange Optical Pumping (MEOP) is the process by which ^{3}He [111] and ^{21}Ne [112] can be hyperpolarised. In order to polarise ^{3}He, a weak RF discharge ($\sim$ 1W), capable of creating a plasma, is sustained inside the optical pumping (OP) cell. Once the plasma is initiated, the plasma can be sustained by a reduced RF discharge power. The weak RF discharge excites electrons to higher levels and occasionally they decay to the metastable state, 2^3S, in a radiative cascade. Electrons are optically pumped from the 2^3S metastable state, using circularly polarised light tuned to 1083nm, to one of the 2^3P sublevels, which de-excite back to the 2^3S metastable state. After multiple cycles, this creates an imbalance in the hyperfine level populations.

For ^{3}He, the hyperfine interaction results in the nuclear and electronic states being entangled, meaning the orientation created by optical pumping of the electrons induces a nuclear orientation as well. Finally, metastability exchange collisions happen between two ^{3}He nuclei; one in the 2^3S_1 state and the other in the 1^1S_0 state. The states of the two nuclei are exchanged, but the nuclei retain their nuclear polarisation. This transfers a net nuclear orientation to the much more populated ground state [113].

A steady state polarisation is reached within minutes, however MEOP is performed at low gas pressures (10s of mbar) [111]. The gas therefore has to be compressed, up to 1 bar, before is can be used clinically .

2.6.3 Spin Exchange Optical Pumping

Spin Exchange Optical Pumping (SEOP), is a process which can be used to hyperpolarised certain noble gases; including ^{129}Xe. The explanation below takes ^{129}Xe to be the gas in question.

An alkali metal (AM) is first heated so that it forms a vapour. A static magnetic field splits the electron sublevels of the single unpaired electron, in the valence shell of the AM, into two; corresponding to $m_j = \pm 1/2$. Circularly polarised light tuned to the transition between the $5^2S_{1/2}$ ground state to the $5^2P_{1/2}$ excited state (794.77nm for Rb), is shone on the cell parallel (or anti-parallel) to the magnetic field. Because the light is circularly polarised, only electrons in one of the sublevels will be excited. For example, if the light has negative helicity, electrons in the $m_j = 1/2$ sublevel of the ground state will be excited to the $m_j = -1/2$ sublevel of the excited state. From the excited $m_j = -1/2$ sublevel, the electron can decay to the $m_j = -1/2$ ground state, or return to the $m_j = 1/2$ ground state. Since only the $m_j = 1/2$ ground state is being depleted and both sublevels are being filled, the electrons accumulate in the $m_j = -1/2$ ground state. Collisions between the AM vapour and other gas species act to equalise the populations of the sublevels of the excited states. This acts to equalise both ground state repopulation rates and results in a rapid equilibrium of the $m_j = -1/2$ ground state [114]. Some collisions result in spin destruction, including those with ^{129}Xe. This results in a polarisation as a function of position in the cell of $P_{AM}(z,r) = \frac{\gamma_{OP}(z,r)}{\gamma_{OP}(z,r)+\Gamma_{SD}}$ where $\gamma_{OP}(z,r)$ is the Optical Pumping (OP) rate per AM atom at position (z,r) and $\Gamma_{OP}(z,r)$ is the electron spin destruction rate [114]. Some collisions between the AM vapour and ^{129}Xe gas, result in spin transfer between the AM vapour and ^{129}Xe, mediated by Fermi-contact hyperfine interactions [115]. Collisions which result in a mutual spin flip are rare. However, because of the long T1 of ^{129}Xe (minutes-hours) polarisation is allowed to build-up. Relaxation of the ^{129}Xe gas is mainly produced by collisions with the walls of the vessel. With perfect polarisation transfer under ideal conditions, the polarisation of the noble gas approaches that of the alkali metal (see equation 2.31).

$$P_{Xe}(t=\infty) = \frac{\gamma_{SE}}{\gamma_{SE}+\Gamma_{Xe}} < P_{AM}(z,r) > \tag{2.31}$$

γ_{SE} is the XE-AM spin exchange rate and Γ_{XE} is the overall relaxation rate of ^{129}Xe ($1/T1^{XE}$)

2.6.4 Hyperpolarisation History

Kastler's work, in the 1950s, on demonstrating spin order in an AM vapour induced by circularly polarised light, formed the basis for SEOP [116]. With MEOP discovered later in 1963 [117]. Historically, higher polarisations were achieved using MEOP for ^{3}He, than SEOP for ^{129}Xe. This is one of the reasons for a historic focus on imaging using ^{3}He rather than ^{129}Xe [7]. However, recent advances have improved the polarisation achieved through SEOP of ^{129}Xe significantly [44][35].

2.6.5 SEOP Considerations

Alkali Metal used

One consideration is the AM used. Rb is most commonly used due to the availability of lasers tuned to 789.77nm. However, Cs has a larger atomic radius, meaning its electrons are more easily spin polarised. This possibly makes it a better prospect for spin exchange. It also has a higher vapour pressure for a given temperature than Rb or K. Shao et al. [118] measured the spin exchange rates for K (164°C), Rb (136°C) and Cs (122°C) to be 0.031, 0.048 and $0.062s^{-1}$ respectively. Different temperatures were used in order to keep the number density of AM constant.

Pressure

In SEOP, polarisation can be transferred from the AM to the noble gas in either a two body interaction (binary scattering), or three body interaction exploiting van der Waals molecules. The two body process is less efficient, but dominates at higher pressures. Higher pressures have generally been used in order to increase the optical pumping rate of the AM. However, Ruth et al. [119] performed an experiment in which the total pressure in the cell was maintained at 100mbar composed of N_2 and ^{129}Xe. They found that while the light absorption increased with concentration of naturally abundant Xe in the cell, the polarisation achieved was optimal at lower pressures of ^{129}Xe. They achieved 70% polarisation at 10mbar naturally abundant Xe and 90mbar N_2.

Hersman et al. [44] used a regime with low pressure and high velocity. This allowed them to take advantage of the improvement in spin exchange rate, due to van der Waals molecules dominating atomic interactions. They produced 50% polarised ^{129}Xe at rates of $> 1\text{Lh}^{-1}$.

Temperature

As temperature increases inside the cell, the Rb vapour density increases as it vaporises. After the Rb vapour absorbs some of the energy from the laser, a

proportion of this absorbed energy is re-emitted as light at different wavelengths. While some of this re-emitted light ends up leaving the cell, a large proportion of the energy is reabsorbed; raising the internal temperature within the cell. This can have a compounding effect; the higher the Rb vapour density, the more energy absorbed. Consequently, the internal temperature increases, pushing up the vapour density further.

Higher temperatures result in a faster build-up in Rb polarisation and more collisions between Rb and ^{129}Xe. However, many of ^{129}Xe's collisions result in spin destruction. At some tipping point the polarisation of ^{129}Xe will drop, when the rate of spin destruction exceeds polarisation transfer.

At high temperatures, the Rb vapour density will be so high that the majority of the laser light is absorbed in the first part of the cell; resulting in "dark" Rb at the back of the cell. In a continuous flow setup, Antonacci et al. [120] found that, at higher temperatures and low flow rates, dark Rb "can make a sizeable contribution to SEOP inefficiency".

Buffer Gas

As well as the noble gas, other gases may be added to increase the achievable polarisation.

By adding N_2, collisions between N_2 and the AM can de-excite the AM vapour non-radiatively. The energy absorbed by N_2 is stored in the ro-vibrational degrees of freedom of the N_2, in a process known as quenching. Without quenching, the temperature inside the cell increases significantly from the oven temperature and the temperature between different parts of the cell can also vary. Normally when the AM vapour de-excites, radiation is emitted by the AM vapour, which is unpolarised and can be on resonance. If this light is reabsorbed, it can act to reduce the polarisation of the AM vapour [121]. This reabsorption of the unpolarised light is refereed to as radiation trapping.

Historically, ^{4}He would be added to pressure broaden the Rb D1 absorption line. This would allow for the absorption of laser light not on resonance [122]. ^{4}He also has a lower Rb spin-destruction rate than N_2 and for this reason is generally used in higher concentrations than N_2 [123].

For a comparison of different buffer gas concentrations, see section 3.7.

Batch Mode Vs Continuous Flow

In a batch mode system, the ^{129}Xe is contained within the cell, with the Rb, which is sealed. Once the ^{129}Xe has been polarised it can be removed. In a continuous flow system, ^{129}Xe is continually flowed over the Rb; each ^{129}Xe atom has an average residence time in the cell. Once the ^{129}Xe has passed over the Rb, the ^{129}Xe can be frozen; extracting it from the N_2 (^{129}Xe freezes at 161K, in comparison to 77K for N_2), or recirculated through the polariser. The continuous flow system allows for a much larger production of hyperpolarised ^{129}Xe and for this reason is more suited to a clinical setting. However, because in the batch mode the ^{129}Xe atoms can collide with the Rb atoms for a longer

period, the polarisations achieved through this method can be greater. Hersman et al. [44] created a continuous flow polariser which can achieve a ^{129}Xe polarisation of 64% at a flow rate of 0.3 Lh^{-1} (litres per hour) and 50% at 1.3 Lh^{-1}. However, this polariser was on a huge scale, with a meter long cell and an optical arrangement of 5 coherent fibre array package lasers all operating at 25W, in order to uniformly fill the column over its full length. The high cost of this arrangement would make it unsuitable for use clinically. In 2013, Korchak et al. [124] detailed a mobile ^{129}Xe polariser on a more manageable scale; requiring only a supply of pressurised air and two wall sockets. Using a continuous flow setup they achieved a maximum polarisation of 40% at 0.1mbar Xe partial pressure and for a flow of 6.5ml/min achieved 25% polarisation. In 2013 Nikolaou et al. [35] managed to achieve $\sim$ 90% polarisation in the cell with 300 torr of Xe, using a batch system, which produced gas on a scale viable for clinical use.

Chapter 3

Improving Rate of ^{129}Xe Build-up for Batch Mode SEOP

In a clinical setting, the maximum achievable polarisation isn't always the target due to time pressures. It would be more favourable to have sufficient polarisation for imaging, produced in a short amount of time, than a slightly greater percentage of polarisation, produced on an impractical time scale. Investigating how build-up rates and final polarisations vary, with temperature and different gas mixes, should help define optimum conditions for high polarisation rates in a short period of time.

3.1 Experimental Setup

3.1.1 Cell

^{129}Xe is polarised in the cell, constructed from Pyrex glass (see figure 3.1).

The cell has two chambers. An internal chamber (length ∼165mm, diameter ∼25mm), which is separated from the environment by two valves which are designed in a way to allow for the internal chamber to be filled with gas without exposing the cell to the external environment. This internal chamber holds the Rubidium, Xenon and buffer gas. The external chamber is open to the external environment and by controlling the temperature of the air going into the external chamber, this acts as an oven; allowing the Rb to be vaporised.

Cleaning, Loading and Filling Cell

It is very important that impurities are removed from the cell before loading with Rb. Rb is a group 1 metal and highly reactive. In order for the Rb to vaporise and to be polarised, it must be in its elemental form. Therefore, any

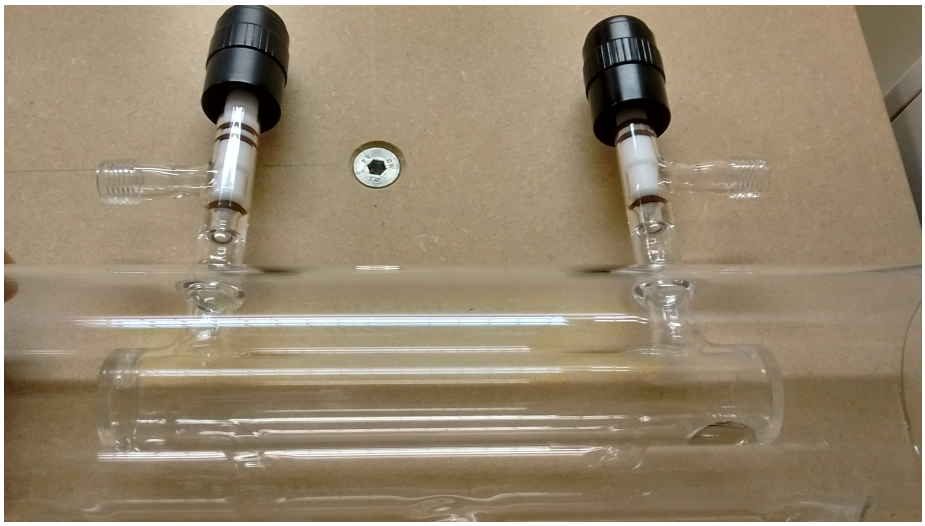

Figure 3.1: Optical cell used in experiment. This is a "spot" cell, with the Rb in a bead rather than deliberately condensed on the walls of the internal chamber. The Rb bead can be seen in the bottom right of the internal chamber of the cell.

impurities such as oxygen, carbon dioxide or water must be removed before loading.

About 0.3 g of Rb is loaded into each cell. Rb is loaded in a controlled environment cabinet with <10ppm O_2 and <2 ppm H_2O (see figure 3.2).

It is important that when loading with Rb, no Rb gets on the front or back windows of the cell, since this is where the laser will impinge upon. Rb on the windows would block the laser light and get very hot.

The cell is evacuated to 10^{-5} Torr, before being filled with a mixture of naturally abundant Xe and a buffer gas of N_2. 26% of the naturally abundant Xe is ^{129}Xe which can be polarised. The N_2 quenches radiative emission and improves the thermal conductivity within the cell; preventing regions within the cell from getting too hot, which can result in Rb "runaway" (see section 3.5). The manifold used to fill the cell can be seen in figure 3.3. Before filling, the manifold and the cell are evacuated. First using a rotary pump, to get the system down to 5 mTorr, before using a turbo pump to get the system down to 10^{-5} Torr. There are 3 cylinders attached to the gas manifold in figure 3.3. The left one is a rich Xe mix with 50% Xe and 50% N_2. The middle white cylinder is a pure N_2 cylinder and the right cylinder is 20% Xe 80% N_2. Using different proportions of gas from these cylinders, different mixes can be created.

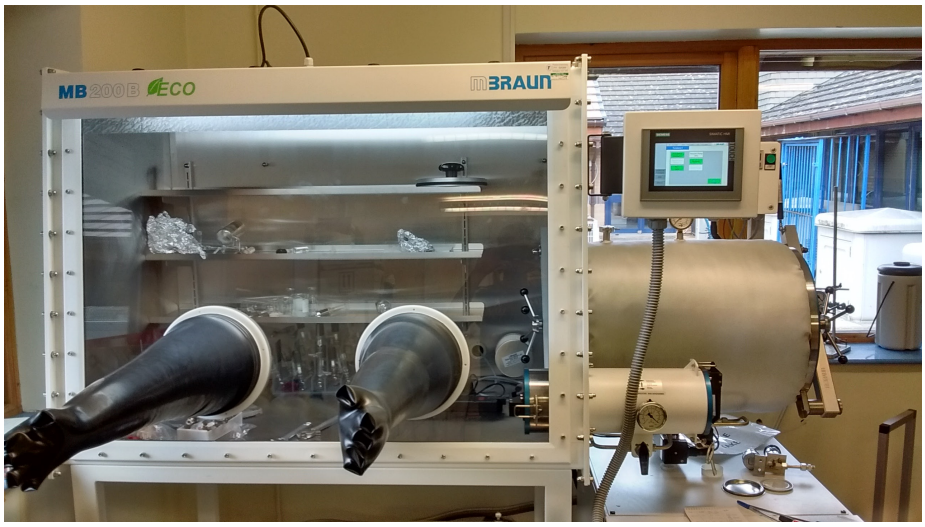

Figure 3.2: MBRAUN 200B controlled environment cabinet used for loading cells in an inert atmosphere.

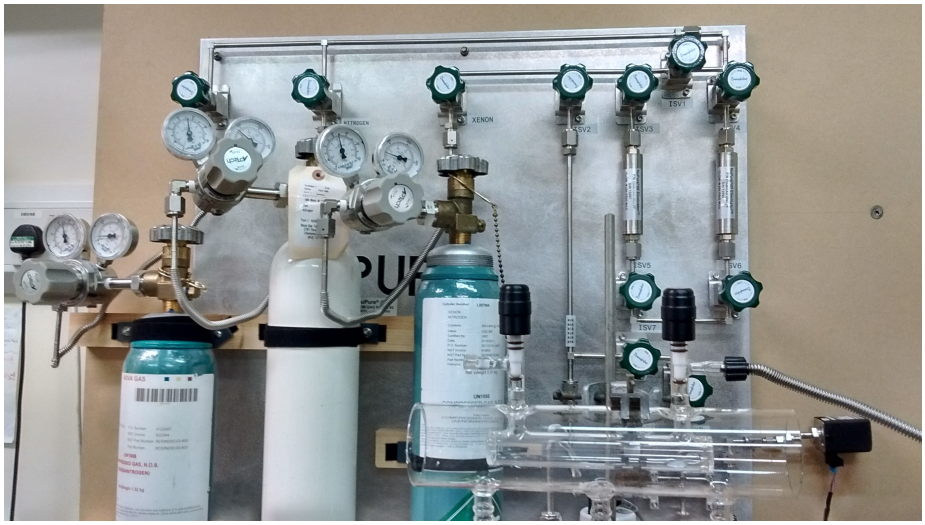

Figure 3.3: Gas manifold used to evacuate cells and fill with Xe and N_2 gas.

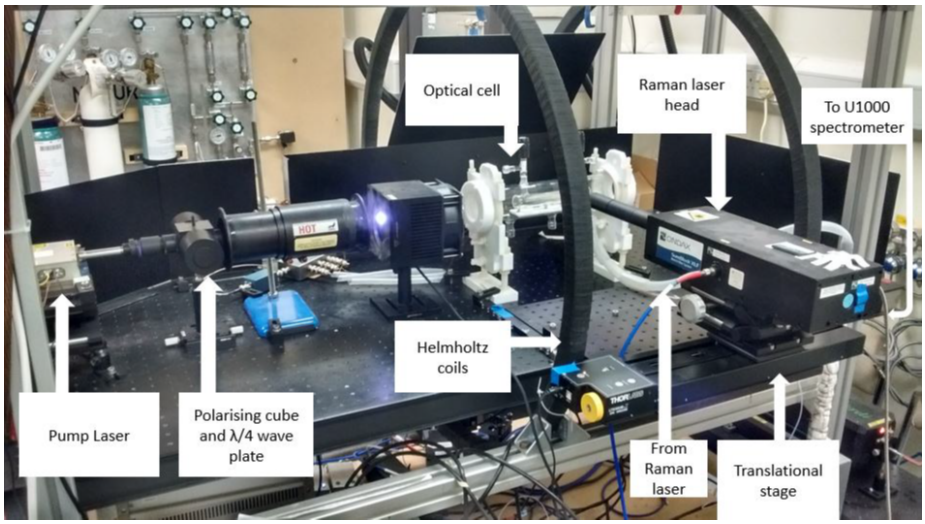

Figure 3.4: Full experimental setup at start of research. The only improvement that had been made from the original setup when this photo was taken was the Translational stage had been moved from 40mm away from the optical cell, to outside the Helmholtz coils (see section 3.1.2).

3.1.2 SEOP Batch Mode Setup

Pump Laser

The pump laser is a 60W QPC Ultra 50 BrightLock device, tuned to 794.77nm. It can be seen on the left of the figure 3.4 and figure 3.5.

The wavelength of the pump laser can be tuned by changing the chiller temperature and current supplied to the laser. A hotter laser cavity increases the wavelength of the emitted photons (red shifted) and a cooler cavity decreases the wavelength (blue shifted). The laser was tuned so it is best absorbed by the Rb D1 line at 794.77nm. A digital thermometer is used to record the temperature of the Laser Diode Array (LDA) and a LDA temperature of 22.2°C was found to be optimal. When running at a current of 35A, a chiller temperature of 17.7°C generally put the LDA at 22.2°C. In Figure 3.5 and 3.6 the optical arrangement of the laser can be seen.

After leaving the LDA, photons first go through a collimating lens, before striking a polarising cube. For a larger beam size, putting the collimating lens after the polarising cube would allow a smaller polarising cube to be used. However, the polarising cube owned was large enough to accommodate a 1 inch beam diameter. The polarising cube rejects horizontally polarised light, directing it towards beam dumps, while vertically polarised light is allowed to pass through. Photons then go through a $\lambda/4$ wave plate at 45°, which changes linear polarisation to circular polarisation. The circularly polarised, collimated light then strikes the internal chamber of the cell, where it optically pumps

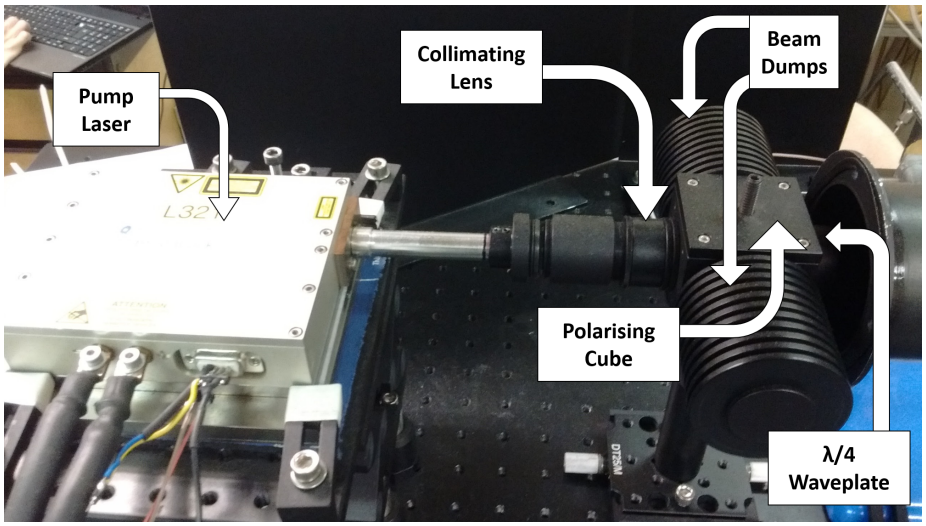

Figure 3.5: Close up of pump laser on the left, with the optical arrangement of a polarising cube, with two beam dumps and a quarter wave plate on the right.

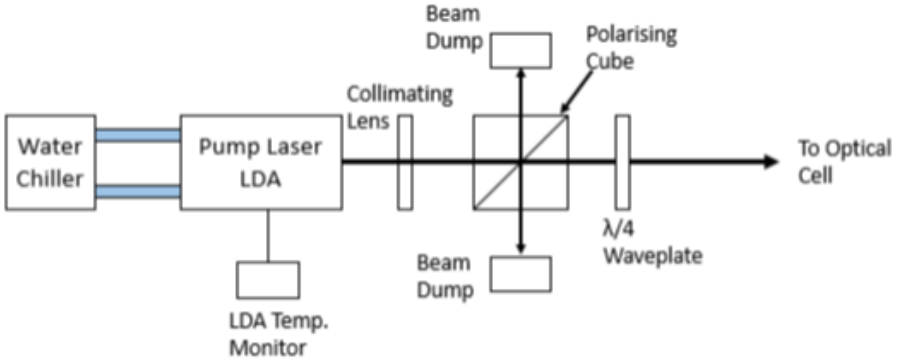

Figure 3.6: Diagram of optical arrangement of experiment.

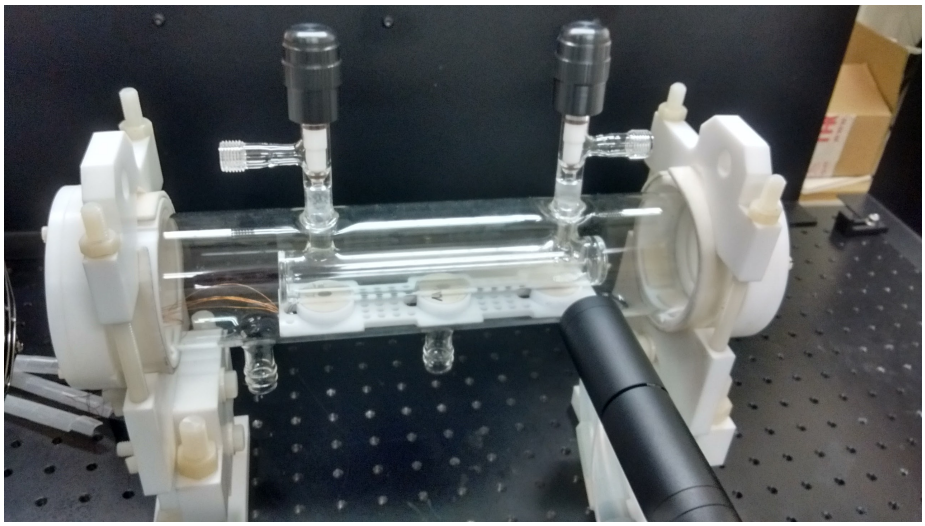

Figure 3.7: Optical cell in situ with NMR coils. Front of the cell on left of image.

the Rb valence electrons (see section 2.6.3). The pump laser and cell must be aligned, ensuring the laser light passes straight through the internal chamber of the cell and doesn't diverge or converge.

NMR coils

A close up of the cell in situ can be seen in figure 3.7. The 3 NMR coils allow the Xe polarisation to be measured at different locations in the cell. From front to back the coils are labelled as; coil 1, coil 2 and coil 3. Each coil is 25.4mm diameter with 350 turns. A Kea^2 NMR Spectrometer, with a simple pulse acquire sequence, is used to obtain a signal. The pulse duration, pulse power and acquisition delay (time between the pulse and collection of the signal) can be set. A pulse power of -6dB ($\sim$ 250mW) was generally used. Every reduction of 6dB equates to halving the pulse power. At this pulse power, it was found that with an acquisition delay of 2500μs, there were no remnants from the initial RF pulse seen. An acquisition delay of 3000μs was generally used, to be certain the signal that was being collected wasn't from the initial RF pulse. This is especially important when working with very weak signals that require multiple acquisitions to be averaged over (see section 3.2.2). The pump laser beam is $\sim$2.54cm in diameter; the same as the internal chamber of the cell.

Helmholtz Coils

The magnetic field is supplied by a pair of Helmholtz coils, with inner diameter of 812.8mm, 179 turns each and powered by a Sorenson XG80-21 power supply unit. The cell is positioned at the centre of these, where the field is most uniform

and strongest. The magnetic field was tuned, so that ^{129}Xe precesses at 34.4 kHz. A current of 8.90A was found to be optimum to ensure ^{129}Xe resonates at the required frequency. The field produced by the Helmholtz coils is 28.23G at 8.90A.

Heating Cell

Hot air is supplied to the external chamber of the cell, by a 400W oven, via a heat pipe. The temperature of the air as it enters the cell, is monitored by a non-magnetic PT-100 resistance sensor (see figure 3.8). Using a negative feedback loop, the temperature of the air (supplied to the external chamber), is controlled using a CAL9500 Temperature Controller in tandem with a Solid State Relays.

Raman Laser and Translational Stage

The Raman Laser (Coherent Verdi V5 532nm) is supplied to the laser head by an optical fibre. The laser impinges on the centre of the cell and is aligned so the focal point of the laser is in the centre of the cell. The head of the Raman laser is mounted on a Thorlabs LTS300/M - 300mm Translation Stage, allowing the focal point of the laser to be moved anywhere (in 5.0μm increments), from the front to back of the cell. Light received back is supplied to a U1000 spectrometer (for more details see section 3.3.1).

In order to maintain a uniform magnetic field, as much ferrous metal as possible, from within the Helmholtz coils, needed to be removed. This meant the stage needed to be mounted outside of the Helmholtz coils; long ferromagnetic objects have a particularly bad effect when orientated parallel to the central axis of the coils.

The NMR signal received, with the stage inside and outside the coils, was investigated. The cell was exposed to the pump laser until the polarisation was at steady state. The cell was filled with a mix of 400 Torr Xe and 1600 Torr N_2 Torr. The signal was taken on coil 2 using a Kea^2 NMR spectrometer (pulse length 300μs, Pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs).

From Table 3.1, it can be seen that there was more than an order of magnitude increase in signal, just by moving the stage outside the Helmholtz coils compared to inside. Once outside, there was little improvement by moving the stage further away, or by removing the stage and Raman head entirely. In fact a greater signal was achieved by having the stage at its furthest extent, rather than with the stage completely removed. One reason for this may be that the magnetic field is not completely uniform. There are small amounts of ferrous metal within the Helmholtz coils and the coils won't be perfectly parallel. Having the stage at its furthest extent, may effectively "shim" the magnet; making the field more uniform.

Figure 3.8: Non-magnetic PT-100 resistance sensor inserted through tube supplying hot, forced air flow.

Stage Location	Peak Signal Amplitude (μV/kHz)	Peak Signal frequency (kHz)
Inside coils 0mm from cell	Noise	-
Inside coils 20mm from cell	Noise	-
Inside coils 40mm from cell	191.94	-1.68
Inside coils 60mm from cell	245.11	-1.34
Inside coils touching coils	568.014	-1.10
Outside coils touching coils	17084.1	-0.13
Outside coils max extent	18756.1	-0.14
No stage	18250.8	-0.01

Table 3.1: Effect of stage location on NMR signal.

With the stage outside the coils there were two options: use a long extension tube and a short focal length lens, or use a long focal length lens and no extension tubes (both with the focal point in the middle of the cell). The setup with the extension tubes was found to produce less extraneous reflections, making it safer and also provided a better Raman signal.

3.2 Calculating Polarisation of ^{129}Xe

In order to calculate the Polarisation of a ^{129}Xe sample, it must be compared to a proton sample. For the purpose of this experiment, the precession frequency of ^{129}Xe was kept the same (34.4kHz) and the magnetic field created by the Helmholtz coils was varied. For a ^{129}Xe precession frequency at 34.4kHz, the optimum current through the coils was 8.90A. Therefore, using the gyromagnetic ratios of ^{1}H and ^{129}Xe (see table 2.1) it was possible to calculate the required current for protons to precess at the same frequency (2.46A) (See equation 3.1).

$$\frac{8.90A \cdot \gamma^{^{129}Xe}}{\gamma^{H}} = 2.46A \tag{3.1}$$

3.2.1 ^{129}Xe Flip Angle Calibration

It is important to decide first what pulse duration will be used to measure the ^{129}Xe polarisation; one that is a good balance of getting a high signal, while not destroying the ^{129}Xe polarisation completely. A spot cell (see figure 3.1), at 160°C was allowed to polarise until it reached steady state. A high temperature was used so that steady state was reached quickly. Using the Kea^2 NMR spectrometer's"1 Pulse Duration Sweep" function, pulse durations from 100μs to 450μs were acquired in 25μs steps on coil 2. (Pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs). Between each acquisition the cell was allowed to recover back to steady state. This doesn't take long, since most of the pulses were at small FAs and even those at high FAs only affect the ^{129}Xe atoms immediately above the coil; only a small proportion of the total ^{129}Xe in the cell is affected. The amplitude of the resulting peaks were recorded and plotted in figure 3.9.

After fitting to a sine curve, an adjusted R^2 of 0.991 was obtained. For pulse durations from 100μs to 450μs, the equivalent FA can be calculated from the fit line. A pulse duration of 300μs was decided upon, equating to a $\sim$70° pulse. This is a good balance between obtaining a high signal, while not killing the signal completely.

3.2.2 Proton Signal

In order to obtain a proton signal, the same design of cell was used as a container and the internal chamber was filled with the water sample. The proton signal is very weak when compared to the signal from the hyperpolarised ^{129}Xe. In order

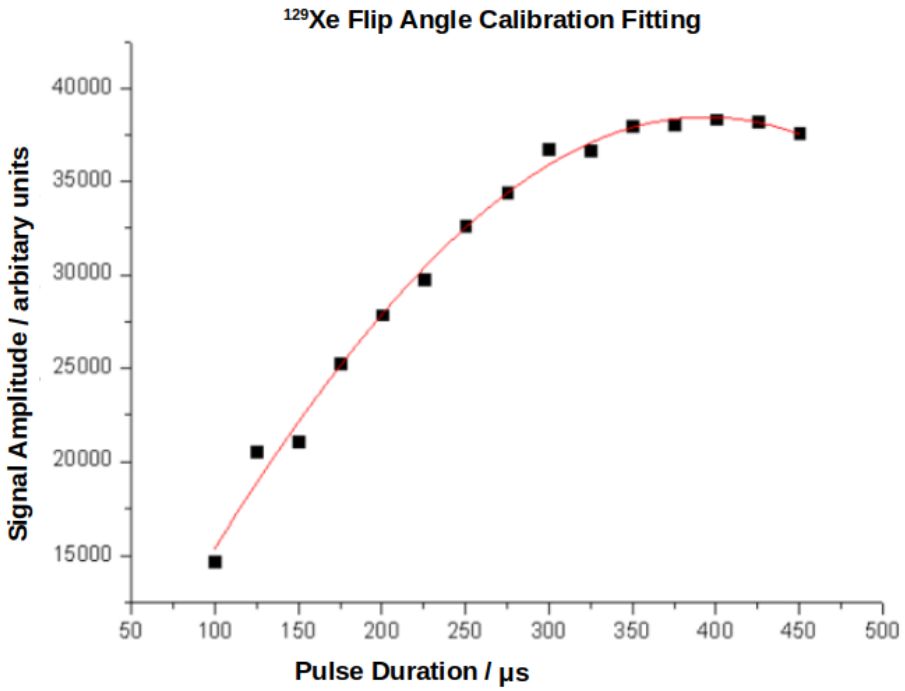

Figure 3.9: Plot of the peak amplitude from ^{129}Xe, at steady state hyperpolarisation, at different pulse durations.

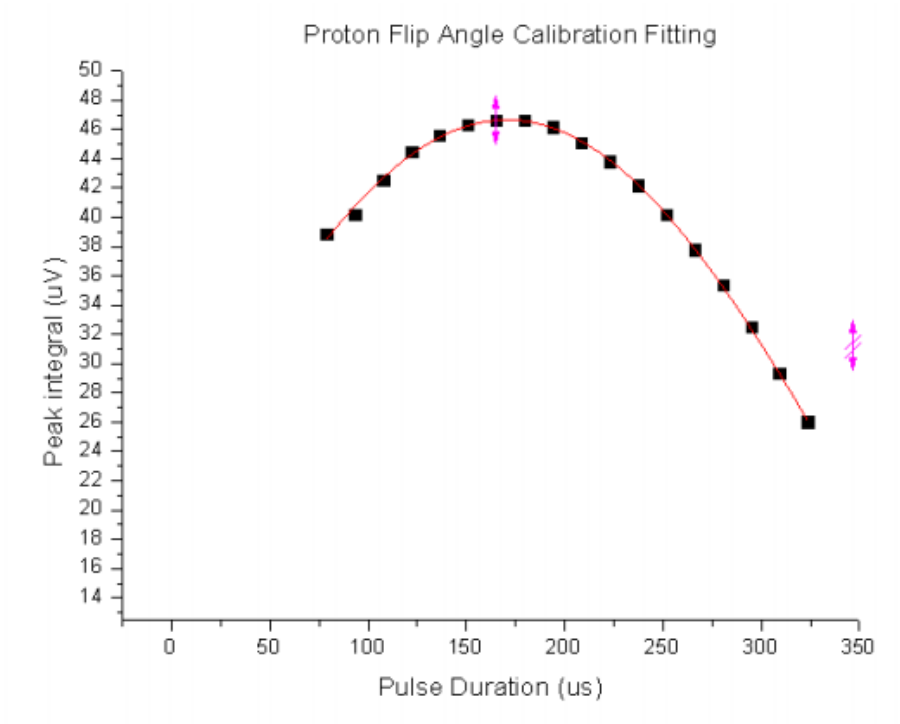

Figure 3.10: Plot of the thermal proton signal at various pulse durations.

to obtain a signal from the noise, multiple repeats must be acquired and averaged. The water sample used was doped with a 10mM concentration of $CuSO_4$. This concentration reduces the T1 of the water from 3-4s to 0.149±0.002s at the Earth's field [125]. This hugely reduces the time needed to obtain a water signal; given the need to wait for approximately 5T1 for the sample to return to thermal equilibrium after an RF pulse.

A FA calibration for the proton signal was performed on coil 2. Using the Kea^2 NMR spectrometer's "1 Pulse Duration Sweep" function, pulse lengths from 78.8μs to 323.6μs were swept through in 14.4 us steps. A repetition time of 1000ms was used to allow sufficient time for T1 recovery. 10000 scans were taken at each point to get sufficient signal. (Pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs). Figure 3.10 shows the resulting points as well as the fitted line (adj. R^2 of 0.99891). Using this line, it was calculated that a 82.98μs pulse length equates to a $\sim$46° pulse.

A proton signal was acquired on coil 2. 50000 scans were averaged in order to create data with sufficient SNR for analysis and a repetition time of 1000ms was used to allow for T1 recovery. (Pulse Length 82.98μs, pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs)(See figure 3.11).

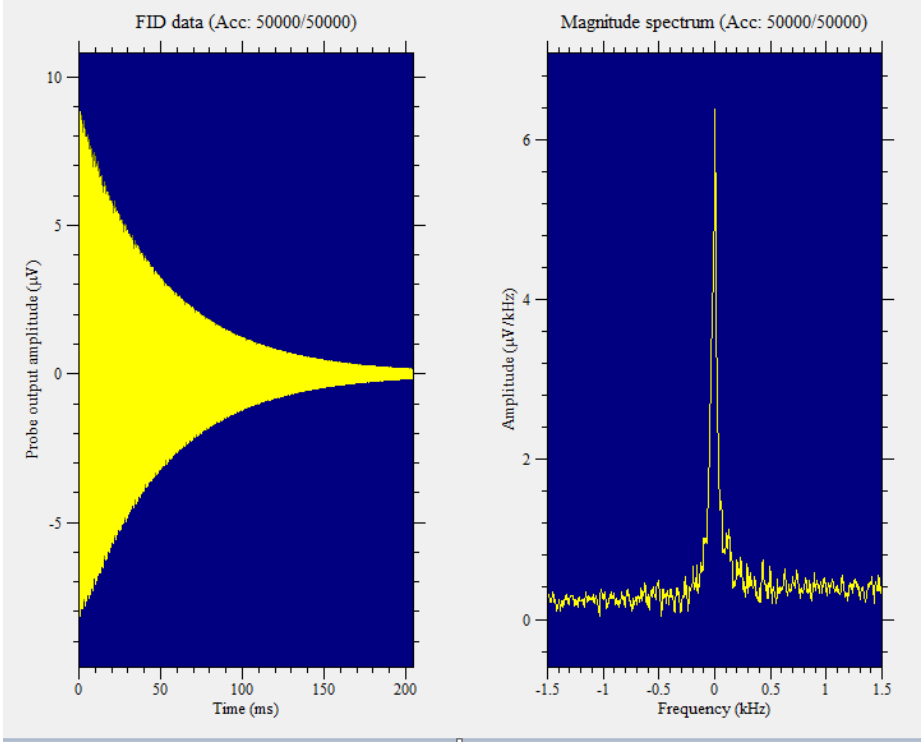

Figure 3.11: 50000 scan average thermal proton signal.

3.2.3 Calculation

After obtaining a proton and ^{129}Xe signal, in order to calculate the polarisation of the ^{129}Xe, the ratio of the number of atoms providing the signal for the ^{129}Xe signal and proton signal must be known.

The molar mass of water is (18g/mol); two hydrogen atoms (1g/mol) and one oxygen (16g/mol). Because there are two hydrogen atoms per water molecule the hydrogen molar concentration (c_H) can be calculated from equation 3.2.

$$c_H = 2\frac{1000g}{18g/mol} = 111.12M \tag{3.2}$$

Since Xe is in its gaseous state, in the temperature range that is used, in order to calculate the concentration of ^{129}Xe (c_{Xe-129}), the molar volume of an ideal gas, at 1 atmosphere, at 273K (V_m) is used; $V_m = 22.414\text{Lmol}^{-1}$. By scaling this to the pressure and temperature of the Xe when loading, c_{Xe-129} can be calculated from equation 3.3.

$$c_{Xe-129} = \frac{\beta_{Xe-129}}{V_m}\frac{P_{Xe}}{P_{760}}\frac{T_{273}}{T_{Xe}} \tag{3.3}$$

β_{Xe-129} is the natural abundance of ^{129}Xe (see table 2.1), P_{Xe} is the pressure of Xe loaded into the optical cell, P_{760} is the pressure at 760 torr (1 atmosphere), T_{273} is 0°C and T_{Xe} is the temperature at which the Xe gas is loaded into the cell.

Combining these elements, a polarisation enhancement factor $e_{enhance}$ can be calculated to determine how much a hyperpolarised ^{129}Xe signal is enhanced when compared to the thermal proton signal (see equation 3.4) [110].

$$\epsilon_{enhance} = \frac{c_H}{c_{Xe-129}}\frac{sin(\alpha_H)}{sin(\alpha_{Xe})}\frac{\gamma_H}{\gamma_{Xe}}\frac{S_{Xe}}{S_H} \tag{3.4}$$

α_H and α_{Xe} are the FAs of protons and ^{129}Xe respectively. S_H and S_{Xe} are the signal intensities produced by the proton and ^{129}Xe samples.

A final compensation factor (C^*_{T2}) is introduced to account for different losses in the NMR signal due to differences in acquisition delay T_{aq} and different $T2^*$ for protons and ^{129}Xe (see equation 3.5) [126].

$$C^*_{T2} = exp(\frac{T_{aqXe}}{T^*_{2Xe}} - \frac{T_{aqH}}{T^*_{2H}}) \tag{3.5}$$

By combining equations 2.30, 3.4 and 3.5, the polarisation (P_{Xe}) can be calculated using equation 3.6.

$$P_{Xe} = 100\epsilon_{enhance}C^*_{T2}P_{thermal} \tag{3.6}$$

3.3 Calculating Temperature Within the Cell

From section 3.5, it can be seen that the oven temperature plays a key role in the speed of polarisation build-up, as well as the final achieved polarisation.

However, using a similar cell design Walter et al. [127] proved that oven temperature isn't representative of the internal temperature of the gas within the cell. Therefore, in order to understand the optimum conditions for SEOP, it is important to know the temperature within the cell. In order to do this, Raman spectroscopy was used to measure the temperature of the N_2 buffer gas. Xe isn't Raman active, but N_2 is due to its change in polarisability (ϕ) after applying an incident electric field, E. μ is the resulting dipole induced by linearly polarised radiation, see equation 3.7.

$$\phi E = \mu \tag{3.7}$$

3.3.1 Raman Spectrum

The interaction between an incident photon and the ro-vibronic states of a Raman active molecule can result in an excitation to a virtual energy state for a short time, before the molecule relaxes to a lower, the original, or higher energy state.

Most of the time Rayleigh scattering occurs. Rayleigh scattering occurs when, after exciting to a virtual state, the molecule returns to its original state (elastic scattering).

Sometimes Raman scattering occurs, producing Stokes or Anti-Stokes lines, through inelastic scattering. When the final state of the molecule is higher than the original state, a photon of lower energy than emitted by the laser is produced, with the energy difference equalling the difference in energy between the ro-vibrational energy levels. The spectrum produced by this process produces Stokes lines. Anti-Stokes lines are produced when the final state of the molecule is lower than the original state.

Only one in 10^6-10^8 incident photons undergo Raman scattering, with the rest undergoing Raleigh scattering or being reflected. In order to detect a Raman spectrum, Raleigh scattering and reflection need to be filtered out. Even after this filtering, the signal from Raleigh scattering and reflection is only attenuated and still dominates.

In the set-up (see figure 3.4 and 3.12) Raman excitation and detection occurs along the same optical path. In order to filter out as much of the original laser light and Rayleigh scattering as possible, two Ondax Sure-Block ultra narrow-band notch filters are used. These allow specific wavelengths to be filtered out (in this case the wavelength of the laser light at 532nm). This means Raman lines only a few wave-numbers from the probe laser can be resolved. An Ondax NoiseBlock ASE filter is also used. On the lasers path to the cell a single frequency is reflected, consequently laser side-bands, spontaneous emission from laser diodes and fibre-induced fluorescence are transmitted to the beam dump. On the return from the cell the ASE filter acts as a spectrally sensitive 90:10 beam splitter; reflecting light of the same wavelength as the laser light (Rayleigh scattering) and transmitting other wavelengths (Raman scattering).

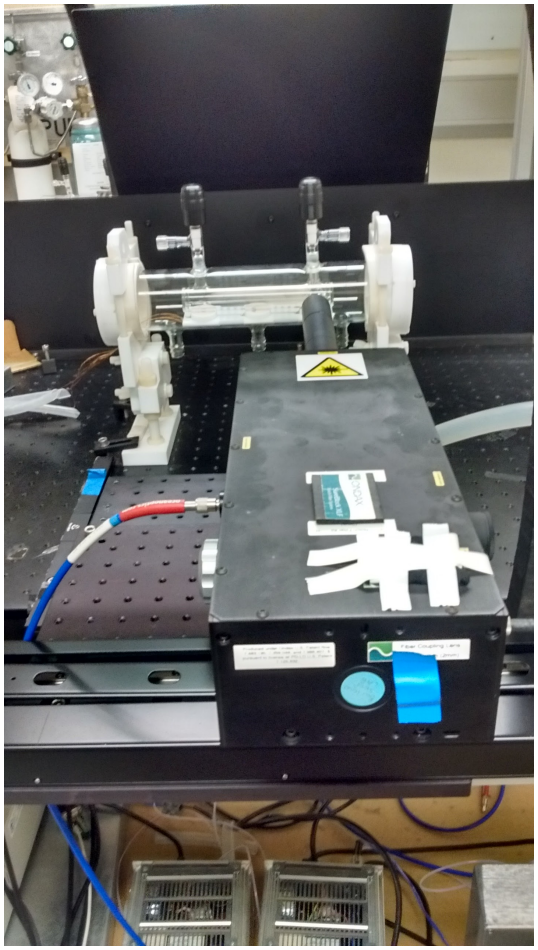

Figure 3.12: Raman head with excitation and detection along same optical path.

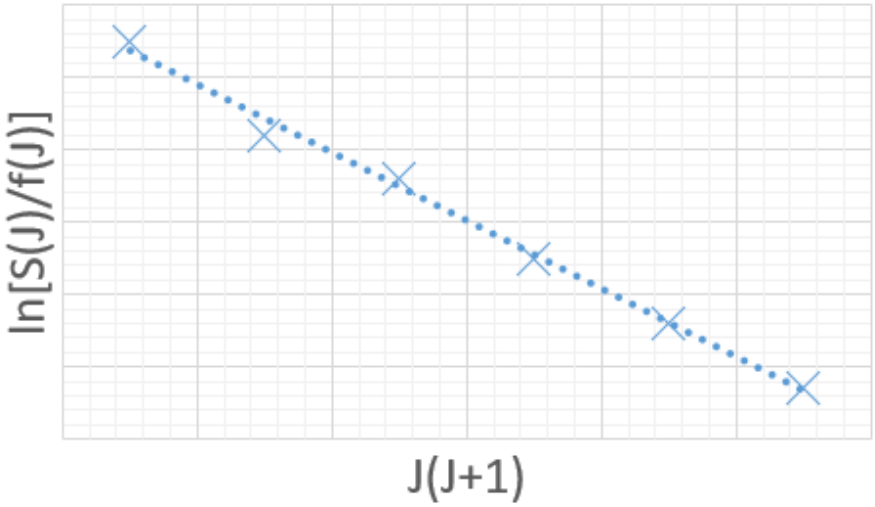

Figure 3.13: As temperature increases, higher energy (higher J) transitions become more likely and the fitted line's gradient becomes less negative.

3.3.2 Calculating Temperature from a Raman Spectrum

The temperature of N_2 gas is calculated by comparing the relative line intensity I(J), of transition J compared to the signal received by the spectrometer S(J). Broadly speaking, at hotter temperatures the N_2 molecules are more likely to remain in a higher ro-vibronic state. This results in transitions with a higher J becoming more likely than at cooler temperatures and transitions with a lower J happen less often than at cooler temperatures. However, transitions with a lower J remain more likely than those with a higher J. A full derivation of how temperature is calculated can be found in Hickman et al. [128]. The temperature can be calculated from figure 3.13; with the gradient of the line equal to $\frac{Bhc}{k_B T}$. f(J) is defined in equation 3.8.

$$f(J) = \frac{3(J+1)(J+2)}{2(2J+3)} \tag{3.8}$$

B is the rotational constant for N_2 (about 2cm^{-1}) and S(J) is the intensity of the line with rotational quantum number J. The steeper this slope the smaller the temperature, since higher J transitions are less likely [128].

3.3.3 Cosmic Rays

Spurious signals from cosmic rays were removed from the Raman spectrum before analysis. If they aren't removed a cosmic ray may be mistaken for an N_2 peak; resulting in an inaccurate temperature being calculated. The U1000 spectrometer uses a charged couple device array detector (CCD), which often

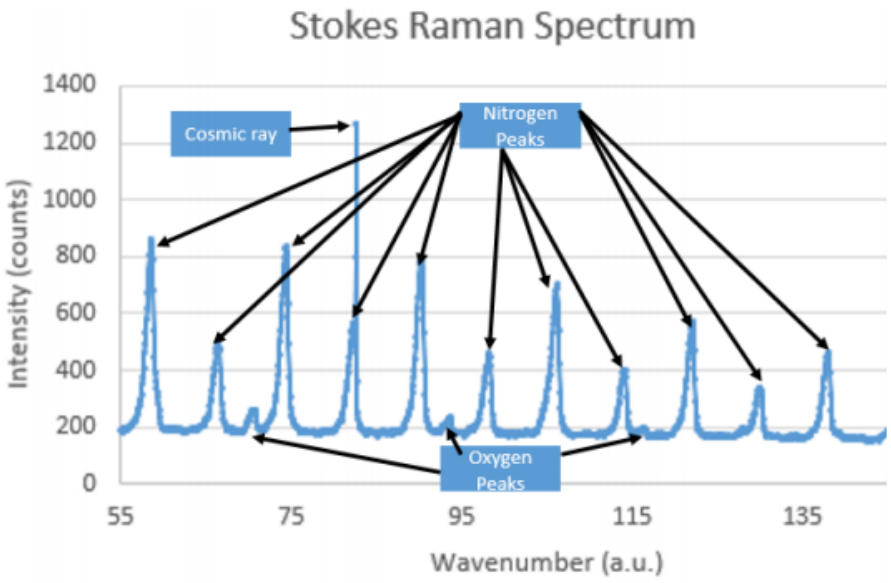

Figure 3.14: Raman spectrum including cosmic ray. Note the N_2 and O_2 peaks at regular intervals.

detects high energy cosmic particles. A cosmic ray can be seen in figure 3.14 as a sharp, high intensity peak; like a delta function when compared to the Stokes lines. The automated cosmic filtering program was also found to be unreliable in removing all the cosmic ray peaks and particular care is required when the cosmic ray is located on one of the Stokes lines. For this reason the cosmic rays were removed manually.

3.3.4 Background Spectra Removal

As can be seen in figure 3.14, although N_2 is the only Raman active species within the cell, there are other contributions to the spectrum. The Raman laser first has to go through air, before it is focussed down to a point within the internal chamber of the cell. Air contains O_2, which is Raman active, as well as N_2. Since only the N_2 in the cell is of interest, it is important to remove the signal produced by the N_2 in the air. Removing signal from the O_2 also helps improve the signal to noise ratio.

In order to do this, a Raman spectrum was acquired with the cell in place, but fully evacuated down to 10^{-5} Torr (referred to as the evacuated spectra). It is assumed that all signal from the O_2 peaks in the evacuated and experiment spectra is from outside the internal chamber of the cell. The evacuated spectra is scaled so the O_2 peaks of the evacuated spectra match that of the experiment spectra. The evacuated spectra can than be subtracted from the experiment spectra before temperature analysis.

3.4 Repeatability of Experiment

Before changing gas mixes, or the external oven temperature, the repeatability of the experiment needed to be checked. In order to do this, the oven temperature was set to 130°C and a cell filled with a mix of 1000 Torr Xe, 1000 Torr N_2. After exposing the cell to the pump laser (at 35A, 22.2°C LDA temp), the signal on coil 2 was measured; every 5 minutes, for 60 minutes (pulse length 300 μs, Pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs). After each run, the pump laser is blocked and the NMR signal crushed by repeatedly running scans with the NMR spectrometer on all three coils using a high FA. This is done until well after the NMR signal has disappeared into noise. After this the experiment was repeated. Two repeats were taken on the first day and six repeats were taken a day later. The cell was then run for two days at a range of temperatures and cooled down to room temperature multiple times. Four days after the first run, a ninth run was taken. Figure 3.15 shows a large variability in the peak signal reached and the rate at which the signal built up over the runs. Figure 3.16 shows the average signal of each run. The average signal increases on all subsequent runs, except between run 2 and run 3. Run 3 was the first run of day two and there is a chance that, although the inlet of the oven was stable at 130°C, the whole cell might not have been at a stable temperature before the experiment was started. This means it might have been cooler than expected in places. Between run 8 and 9 there is a large increase in signal. Over the first 8 repeats, where no runs at higher temperatures were performed between, the NMR signal increased by >4 fold.

Three postulated reasons for lack of repeatability:

- On filling with gas, small concentrations of impurities are introduced into the cell. These impurities react with the surface of the Rb bead poisoning it slightly. After multiple cycles of heating and cooling, the poisoned Rb becomes distributed more evenly throughout the bead, making the surface more pure. Since it is the surface of the bead that is vaporised, the more pure surface Rb will result in a greater vapour pressure at lower temperatures.

- When impurities in the gas mix collide with polarised ^{129}Xe, there is a high probability that the polarised ^{129}Xe will relax [129]. Over time the Rb vapour reacts with these impurities in the gas mix. This removes the impurities from the gaseous phase, resulting in slower relaxation of the polarised ^{129}Xe.

- Upon heating, some of the Rb goes into its vapour state. Upon cooling, Rb condenses. Over time, the surface area of solid Rb increases as Rb plates on the sides of the cell. At a specific temperature, the vapour pressure of Rb should be consistent; provided the pressure and volume of the Rb's surroundings, quantity of Rb and surface area of Rb remain consistent. However, due to plating the surface area isn't consistent.

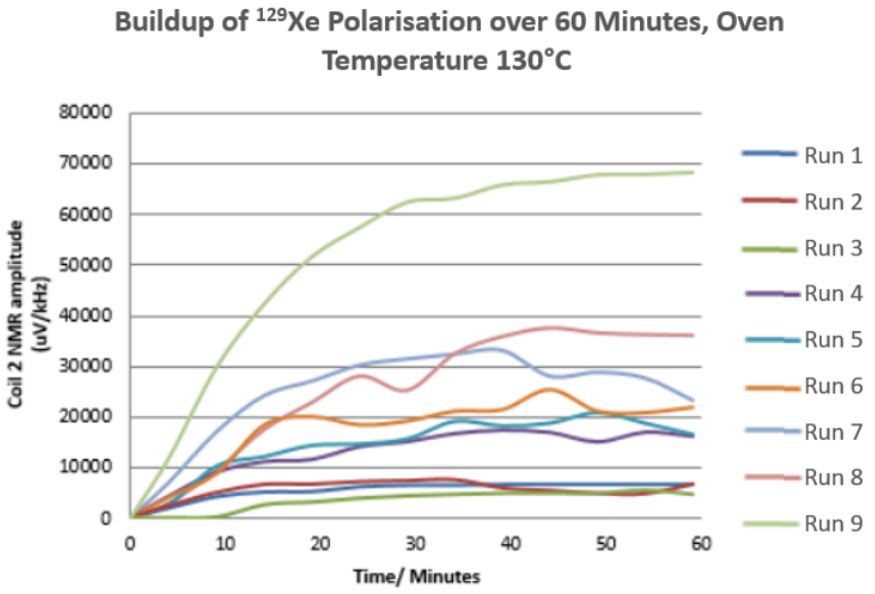

Figure 3.15: Multiple runs taken under the same external conditions: oven temp 130°C.

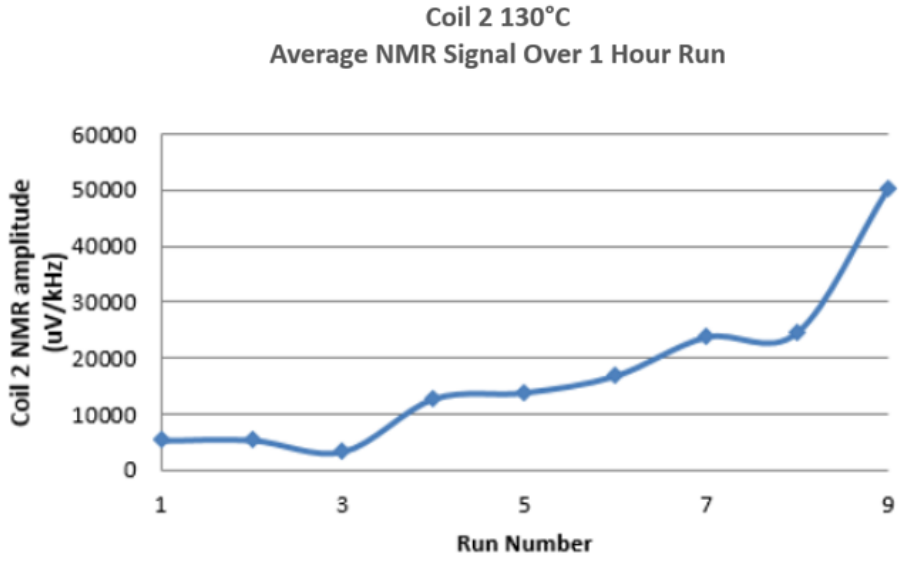

Figure 3.16: Average NMR signal from runs in figure 3.15. These were acquired over three non-consecutive days. A two day break was taken between run 2 and 3, with no other runs done between. Run 9 was taken two days after run 8 but between runs 8 and 9 multiple runs at higher temperatures were performed.

Two suggestions to resolve the repeatability issues postulated: either measure the Rb vapour density directly and scale for this, or pre-plate a cell so further plating has little effect.

3.4.1 Atomic Absorption Spectroscopy

Atomic Absorption Spectroscopy (AAS) allows the Rb vapour density to be measured directly. A stable light source is directed through the centre of the cell, after which it is picked up by an optical spectrometer. By looking at the absorption and the spectral width of the Rb D_1 (795.77nm) or D_2 line (780nm), at a high oven temperature versus room temperature (when it assumed there is no Rb vapour), the Rb vapour density can be measured [120]. Even though it is possible to measure the absorption perpendicular to the pump laser, the pump laser is of such great intensity that any extraneous reflections would cause significant noise if the absorption on the D_1 line was measured. Instead the D_2 line was used for AAS.

Difficulties were found in obtaining a stable source. If the source varies in intensity it will be impossible to discriminate between a reduction in signal due to source fluctuations and one due to light being absorbed by the Rb vapour. Rb D_2 laser diodes were tested, but their power relied on a very stable power supply and the power supplies available were unsuitable. In the end, a Thorlabs Quartz Tungsten-Halogen Lamp was found to offer a suitably stable source using mains power. Any fluctuations in power will affect the full spectrum of wavelengths, so the intensity fluctuations at a specific wavelength are attenuated compared to a laser diode.

An optical fibre with a cosine head (fibre 1 in figure 3.17 and figure 3.18), positioned directly opposite the halogen lamp, directs the light that has been through the optical cell to an Ocean Optics HR2000+ spectrometer.

The optical cell was filled with 400 Torr Xe, 1200 Torr N_2 and 400 Torr He. A reference spectrum was taken at room temperature. The optical cell was then heated to 160°C with no pump laser and allowed to stabilise for 30 minutes, the absorption was measured. Figure 3.19 shows absorption at both the Rb D1 and D2 lines.

The optical cell was then filled with 400 Torr Xe and 1600 Torr N_2. The cell was exposed to the pump laser for one hour and the relative absorption and polarisation measured. This was repeated at three different temperatures on two consecutive days. The same crushing scheme was used as in section 3.4. The cell was allowed to stabilise for 5 minutes after the hot gas inlet reached the desired temperature. Figure 3.20 illustrates how higher temperatures correlate with a higher absorption on the Rb D_2 line. Figure 3.20 also shows the stability of AAS once the polarisation in the cell has reached a steady state. Again, the cell had aged by the second day and on the second run the polarisation at every temperature was higher than the first. This increase in polarisation was matched by the D_2 absorption.

However, ultimately the AAS setup relies on the halogen lamp, fibre 1 and optical cell not moving; see figure 3.17 and figure 3.18. The fibre can be reliably

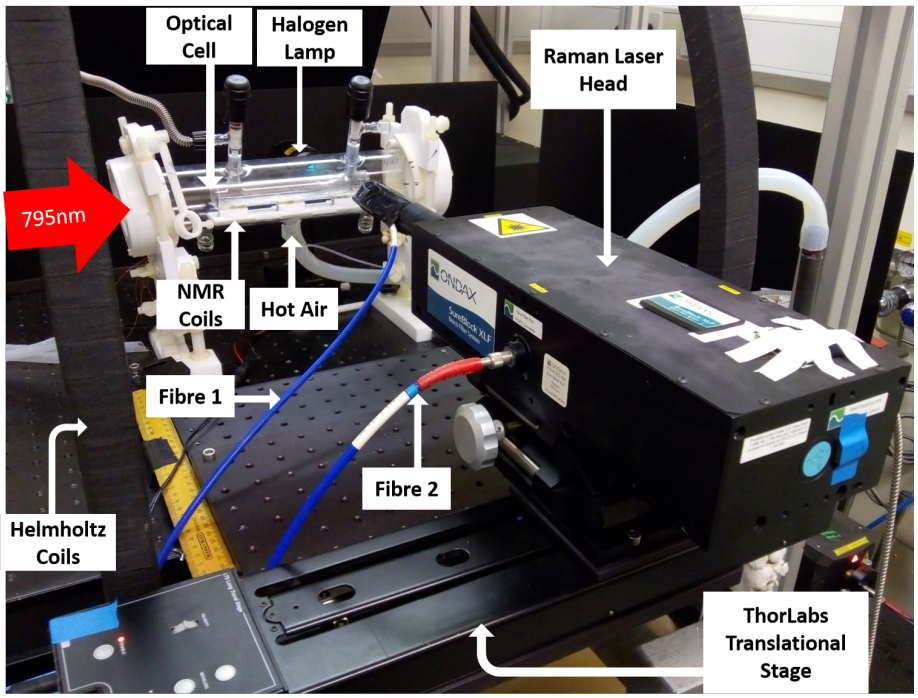

Figure 3.17: SEOP experimental setup. Fibre 1 sees light from the halogen lamp which has passed through the optical cell. Fibre 2 supplies 532nm laser light to the Raman Laser Head. This light continues on to a U1000 spectrometer to perform Raman spectroscopy.

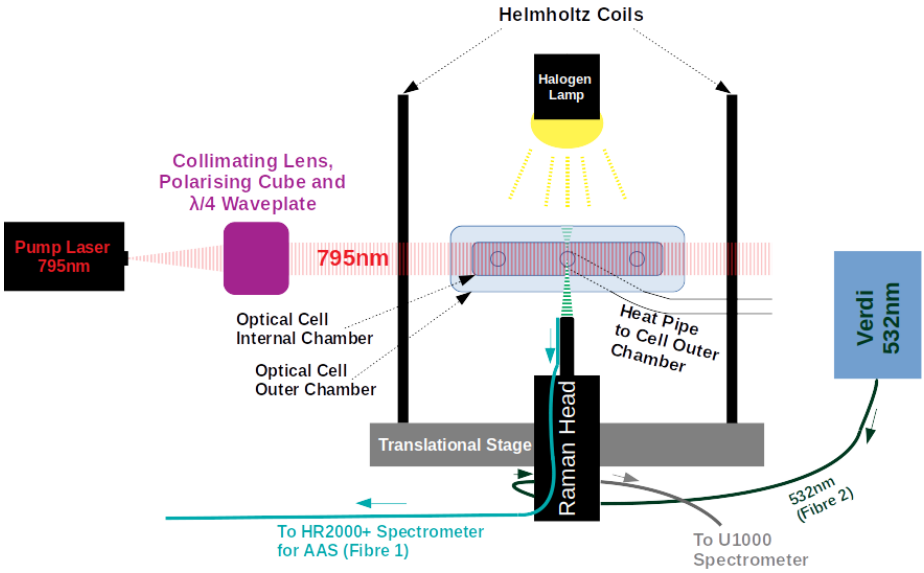

Figure 3.18: Diagram showing SEOP experimental setup from above.

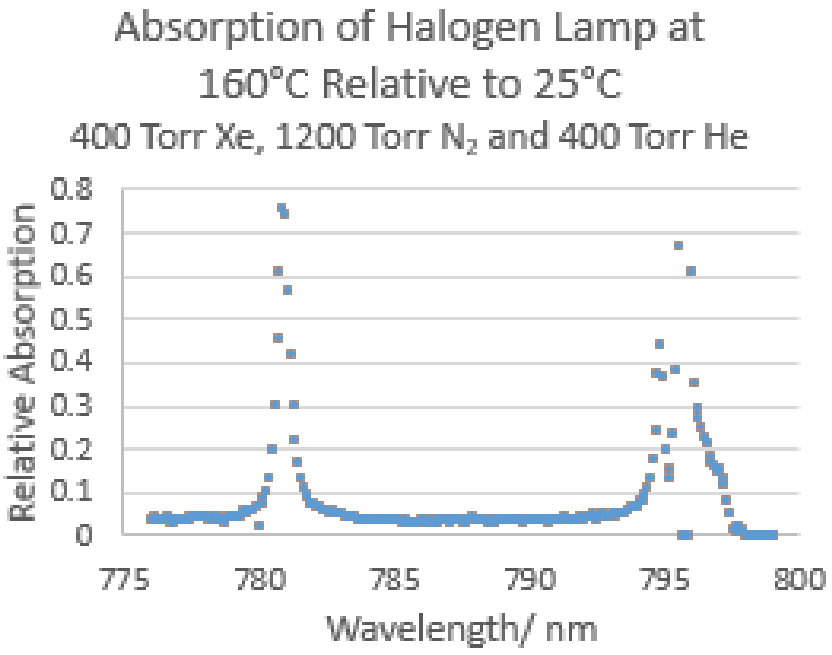

Figure 3.19: Absorption of halogen lamp by Rb vapour. Rb D_2 line left peak, D_1 line right peak.

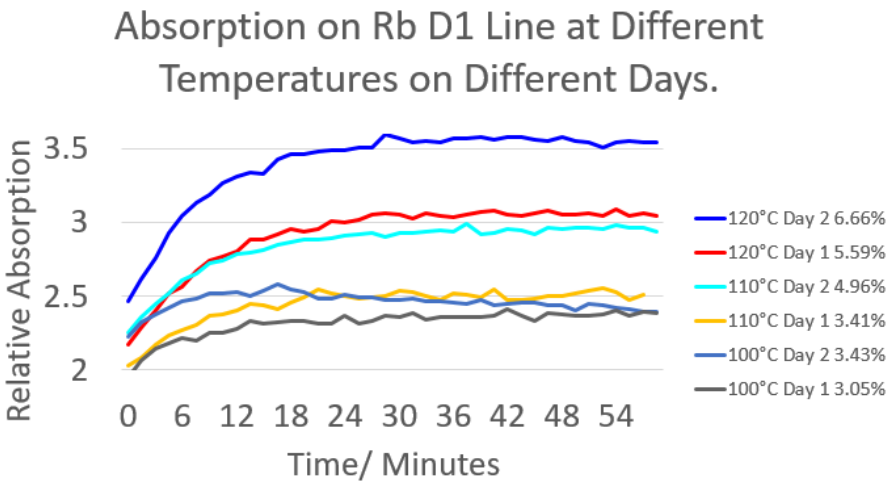

Figure 3.20: The cell was exposed to 795nm to perform SEOP. The hotter the oven temperature, the more absorption on the Rb D2 line. Also the greater the increase in absorption over the run. The percentage ^{129}Xe polarisation achieved for each run can be found in the legend.

returned to the exact location by the translational stage $\pm 5.0 \mu m$. However, the goal is to compare different gas mixes. This requires the valves on the optical cell to be opened and when doing this it was very easy to rotate the cell in the holder. This would make it impossible to compare absorption data from before and after this movement due to the refraction and reflection of the halogen lamp within the optical cell being altered slightly. This is exacerbated by highly reflective plated Rb. There was also the issue of other people with access to the lab knocking the bench slightly. A more stable holder was considered. However, due to so many moving parts and the requirement of a non metallic holder, so as not to interfere with the homogeneous field, another solution to the reliability problems was sought.

3.4.2 Pre-Plated "Spread Cell"

Creating Spread Cell

With the cell fully evacuated, the front and back windows along with one side of the internal chamber of the cell are heated with a heat gun. Direct heat is then applied to the bead of Rb until it vaporises. The cell is left to cool. Any Rb that has condensed where it shouldn't be, is removed by applying heat. As previously explained in section 3.1.1, Rb shouldn't be on the front or back window. Also, because the Raman laser is being used, the laser light must be able to pass into the internal chamber of the cell un-impinged. For this reason, Rb is only plated on one side of the internal chamber. An example plated cell

Figure 3.21: Pre-plated spread cell with one clear side.

(on one side) can be seen in figure 3.21.

Spread Cell Repeatability

With a spread cell, a much higher vapour pressure is achieved at lower temperatures. To test the repeatability of the spread cell, a 400 Torr Xe, 1600 Torr N_2 mix was used and run seven times at 100°C; using the same procedure as in section 3.4 (pulse length 300 us, Pulse amplitude -6dB, pulse acquisition delay time 3000μs, dwell time 20.0μs). An average signal of 34.98 mV/kHz was measured with a standard deviation of 6.554 mV/kHz. Four runs were taken on the first day and three runs on the second day. Figure 3.22 shows there was a downward trend over the runs. The signal changed by less than a factor of two over seven runs. Comparing this to a > factor of four change in signal over seven runs with the spot cell, the spread cell provided a more consistent signal.

3.5 Oven Temperature Effect on ^{129}Xe Polarisation Build-up

Using a spread cell filled with 400 Torr Xe and 1600 Torr N_2, the same procedure and pulse parameters were used as in section 3.4, but this time the temperature was altered between runs. The build-up curves can be seen in figure 3.23. With increasing temperature, the rate at which polarisation builds up increases. However, the highest final polarisation isn't achieved at the highest temperatures. The highest final polarisation was seen at 100°C. At 110°C and 120°C

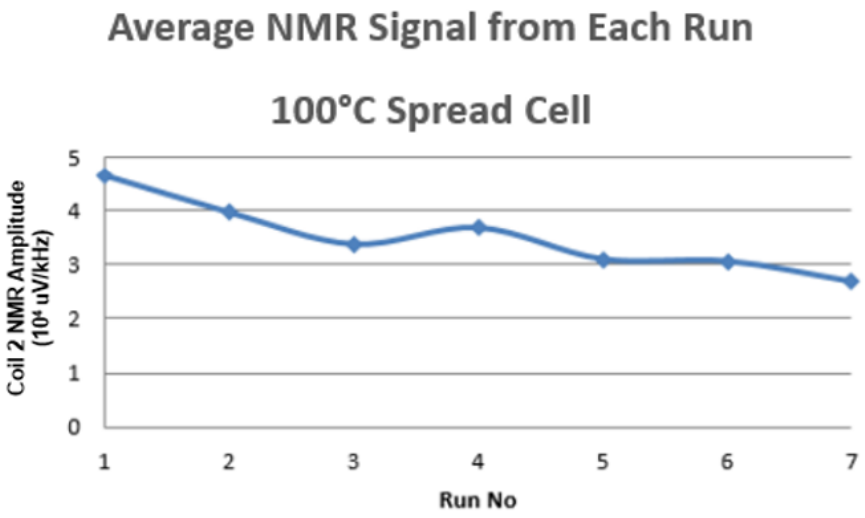

Figure 3.22: Average polarisation of ^{129}Xe over multiple runs. Taken over two consecutive days. First four runs on first day.

runaway was observed where the relaxation processes happen at a faster rate than the build-up of polarisation. Once runaway has happened, the higher the oven temperature, the lower the final polarisation. It is important to understand the optimum conditions for polarisation build-up. In a medical setting the aim would be to achieve the highest possible polarisation in the shortest possible time. Ideally polarisation would start at a high oven temperature, before cooling to avoid runaway; thus achieving a high polarisation in a short time. However, after starting with a high initial vapour pressure it is hard to cool the cell, since the high vapour pressure of Rb means there is a large amount of laser heating. On the setup used, there was no method to actively cool the cell. A much greater thermal mass and a better conductivity between the cell and oven, than a simple air oven, would be required in order to cool the oven sufficiently to avoid runaway [130].

3.6 Raman Experiment

In previous experiments with this setup (before the repeatability issues were discovered), the hot air from the oven came in at either the front or back of the cell, but always below the Rb spot. This creates a small temperature gradient between the front and back of the cell. Even before polarisation the end of the cell with the Rb spot would be hotter. With no pump laser present this difference is around 30°C at oven temperatures between 100°C and 140°C.

Using a spot cell filled with 400 Torr Xe and 1600 Torr N_2, the internal

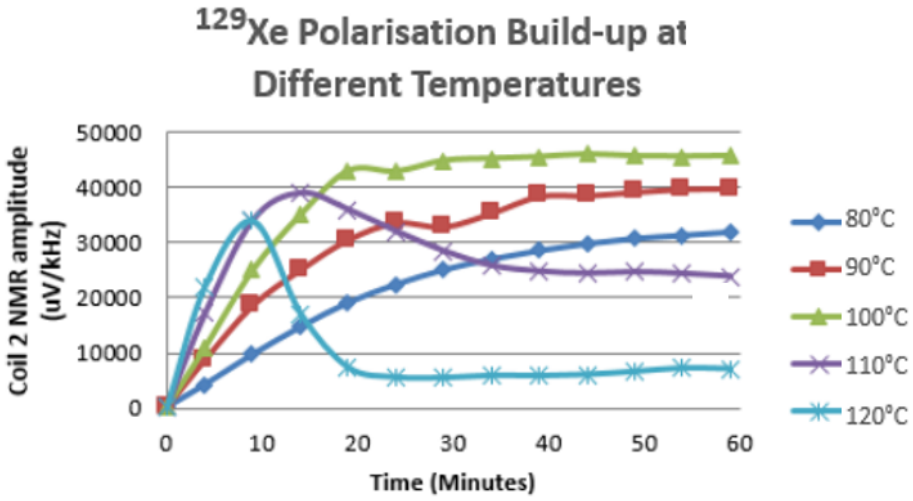

Figure 3.23: Polarisation build-up curves for ^{129}Xe in a spread cell.

temperatures were mapped out at the end of a 60 minute run; after runaway (180°C oven temp) and very steady build-up (80°C oven temperature). The cell was then reversed. The cell was always heated below the spot. Figure 3.24 shows that at 180°C there was a large temperature gradient from the front to the back of the cell (even when heating from the back). However, the front of the cell was cooler when heating from the back than the front. Part of the gradient can be explained by the Rb vapour density being so high, that the majority of the pump laser is absorbed in the first few cm of the cell; causing large amounts of heating. However, the heat pipe location appears to affect the gradient.

Unfortunately, it wasn't possible to compare heating from the front and back to heating from the centre in a non spread cell. The Rb bead needs to be heated directly below the bead in order to generate enough vapour pressure. Due to the cell design, the Rb bead is loaded below the stems and the surface area will increase dramatically if attempts are made to move the bead. Even at the maximum temperature of the oven it wasn't possible to force the cell into runaway when heating from the middle without spreading the bead.

The spread cell also can't be used because when reversing the orientation; a different side of the cell faces the Raman laser and only one side of the optical cell can be kept clear from Rb plating.

Fortunately, the spread cell can be heated from the middle because higher vapour pressures are achieved at lower temperatures. For all experiments on the spread cell, except that in section 3.4.2, the cell was heated from the centre. This should reduce any effects of a thermal gradient induced by the oven.

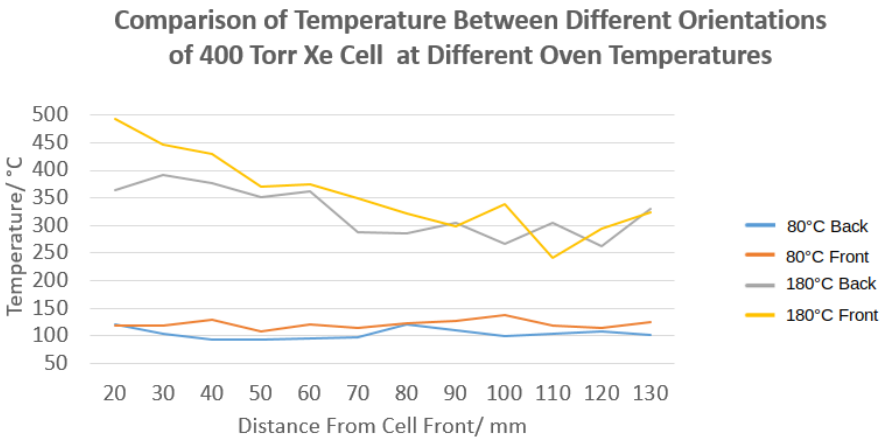

Figure 3.24: Comparison of internal cell temperatures, 60 minutes after exposure to pump laser for different orientations and temperatures.

3.7 Comparing Different Gas Mixes

The choice and quantity of buffer gas is known to have significant effects on the ultimate polarisation achievable (see section 2.6.5). All investigations were performed on a spread cell filled with 2000 Torr of gas and heated from the middle.

3.7.1 Effect of Changing ^{129}Xe to N_2 Ratio

A spread optical cell was heated to three different temperatures and exposed to the pump laser for 60 minutes. Between runs the same crushing policy was adopted as in section 3.4. The internal cell temperature was allowed to stabilise for 5 minutes after the desired oven temperature was reached. This was repeated on concurrent days for two other gas mixes. Runs on the mixes were performed in the order of table 3.2 (starting with higher N_2 concentrations).

Table 3.2 shows, that increasing the N_2 concentration and decreasing the Xe concentration, results in a higher percentage polarisation. As expected, the polarisation is higher at 120°C than 80°C. Both exhibited steady build-up in polarisation over the hour. The 160°C run, for all 3 gas mixes, experienced runaway and consequently the polarisation was lower than the 140°C run.

In a clinical setting, when using a batch mode system, there is no use producing a very high ^{129}Xe polarisation if the concentration of ^{129}Xe in the mix is very low. 10% polarisation from a 200 Torr Xe 1800 Torr N_2 mix is equivalent to 40% polarisation from a 50 Torr Xe 1950 Torr N_2 mix. Looking at the numbers in brackets in table 3.2 there is little advantage in overall signal terms of using one gas mix over another at the optimum temperature. The temperature of the oven had a much greater effect. The 200 Torr Xe mix performed best under non

	80°C	120°C	160°C
50 Torr Xe 1950 Torr N_2	**44.6**(11.2)%	**85.9**(21.5)%	**57.9**(2.94)%
200 Torr Xe 1800 Torr N_2	**15**%	**24.9**%	**16.9**%
1000 Torr Xe 1000 Torr N_2	**1.84**(9.2)%	**4.74**(23.7)%	**2.94**(14.7)%

Table 3.2: Polarisation in different gas mixes after 10 minutes build-up (spread cell). The 160°C cell has already experienced runaway in all three cases. Note bold polarisation values refer to the polarisation of the Xe. Numbers in brackets are the polarisation scaled to the 200 Torr Xe 1800 Torr N_2, to allow for overall signal comparison.

runaway conditions, but at 120°C the bulk magnetisation achieved was similar to the other gas mixes. Figure 3.22 shows there is variability between runs under the same conditions, with a variation of $\sim 19\%$ within one standard deviation. However, at 80°C the 200 Torr Xe mix did performed much better than the other two mixes; both mixes producing a polarisation >1 standard deviation lower.

3.7.2 N_2 vs ^{4}He Buffer Gas

Using 200 Torr Xe and 1800 Torr Buffer gas, three different buffer gas mixes were investigated. The gas mixes were investigated on concurrent days, starting with low ^{4}He concentrations and increasing the ^{4}He concentration each day. A measure of the Rb vapour density was obtained using AAS (see section 3.4.1). The cell was only run in the regime of steady build-up and the polarisation was measured after 60 minutes (after a steady state had been reached). The cell was not forced into runaway conditions. The relative absorption of the Rb D_2 line was plotted against final polarisation (figure 3.25).

For different gas mixes, with comparable Rb vapour densities, the ^{129}Xe polarisation achieved was similar. Figure 3.25 suggests, that the buffer gas mix had little effect on the ultimate ^{129}Xe polarisation achieved after 60 minutes for a comparable Rb vapour density. Figure 3.25 also suggests, the D_2 absorption is approximately linear to the ^{129}Xe polarisation (for the same Xe partial pressure in the gas mix), in the low Rb vapour density low polarisation regime. This would need to be investigated further before firm conclusions can be formed. However, doing this is impracticable on the setup where it is difficult to avoid some small movement to the optical cell (see section 3.4.1 for the more detail of the limitations of the setup).

The experiment providing data for figure 3.25 was performed on three previous occasions, but each time the cell was knocked during one of the fills, causing the AAS data to be useless for future comparison. On the third run at 200 Torr N_2 the heat pipe from the oven became dislodged. It was impossible to return the heat pipe without the cell being moved slightly, so this run was not included in figure 3.25.

Using a 160°C oven temperature, the optical cell was forced to runaway

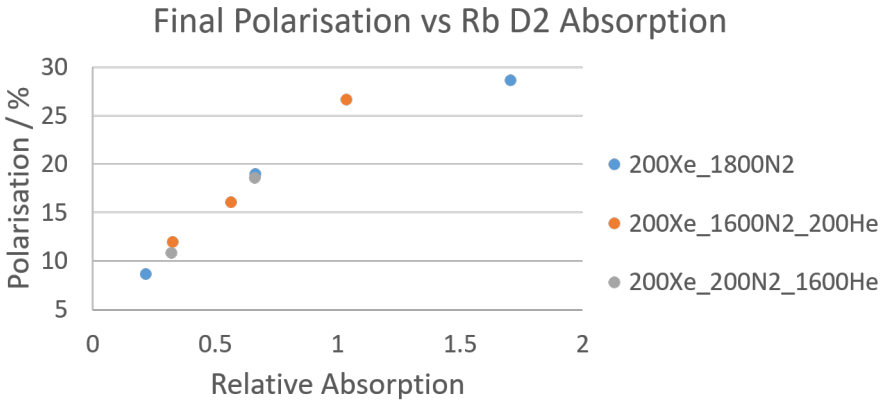

Figure 3.25: Final polarisation achieved after 10 minutes against Rb density at non runaway temperatures with ^{4}He mixes. Relative absorption refers to AAS performed on the D_2 line as a measure of the Rb vapour density. Gas pressures in legend in Torr.

after exposure to the pump laser. After 15 minutes, the internal temperature was mapped using Raman spectroscopy. The Thorlabs Translational stage was stepped from the front of the cell to the back, in 1cm steps. The same three gas mixes were investigated, as well as two different Xe concentrations balanced with N_2. The gas mixes were investigated on consecutive days in the order of the legend in figure 3.26.

The three gas mixes with no ^{4}He had similar profiles from the front to the back of the cell, with the front being slightly hotter than the back. Increasing the ^{4}He fraction appears to cause the cell to run hotter; particularly at the front. This is the reverse of what was expected (see section 2.6.5). However, the spread cell Raman data is much nosier than that of the spot cell. While every effort was made to avoid spreading Rb on one side of the optical cell, there is likely still some spreading present. When the cell is moved, the spread of the Rb on this side will be different to that of the reference spectrum. This makes it difficult to approximate errors.

3.8 Conclusion

In a batch mode setup it is important to pre-plate the Rb before performing SEOP. Not only does this allow the cell to be run at lower temperatures, while still achieving the same vapour density, but it also creates more consistent polarisation build-up curves between runs performed with the same conditions. Lower temperatures mean less time would be required to heat up and cool down the cell, resulting in polarised ^{129}Xe being produced within shorter durations. The consistency is also important. In a clinical setting, polarised ^{129}Xe would

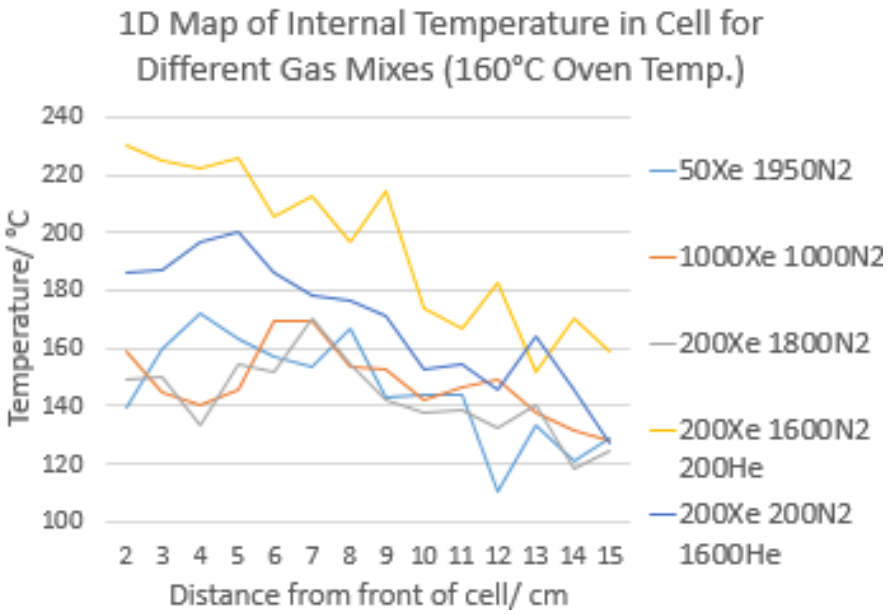

Figure 3.26: Internal cell temperatures for different gas mixes after 15 minute build-up at 160°C oven temperature. Gas pressures in legend in Torr.

need to be produced reliably, at high polarisations, within a known time. A patient can't be left waiting because the polarisation build-up is slower than expected, or because the temperature was too high and the cell was forced into runaway.

As the temperature of the oven increased, faster rates of build-up of polarised ^{129}Xe were observed. At a certain point however, this reverses and the cell goes into runaway; the polarisation destruction rate exceeds the polarisation build-up rate. In a clinical setting it would be advantageous to exploit the very high build-up rates just before runaway occurs. After a short period of fast build-up, the cell could be cooled to prevent the Rb vapour density becoming too high. A period of slower polarisation build-up could then be employed to increase the polarisation further. This would require an oven with a method of active cooling and a high thermal conductivity between itself and the cell [130].

Temperatures within the internal chamber of the cell were observed to be vastly elevated above the oven temperature during runaway. The front of the cell was also observed to be hotter than the back, irrespective of the the location of the heat pipe supplying hot air from the oven. The vastly elevated temperatures can be explained by the large amount of energy absorbed by Rb vapour (at high Rb vapour densities). The fact the front of the cell was seen to be hotter than the back, suggests the majority of the laser light is absorbed by Rb vapour in the first few cm of the cell at high Rb vapour densities.

Higher polarisation percentages were observed at lower ^{129}Xe partial frac-

tions and higher N_2 partial fractions; for each gas mix the total pressure of gas was 2000 Torr. However, the bulk magnetisation was found to be higher with a 200 Torr Xe mix, than a 50 Torr or 1000 Torr Xe mix. In a batch mode setup there is a balancing act, between generating high polarisations of ^{129}Xe with lean mixes of ^{129}Xe and producing low polarisations of ^{129}Xe with richer mixes of ^{129}Xe.

For a given vapour density of Rb, altering the partial fraction of N_2 and ^{4}He in the buffer gas appeared to have little effect on the ^{129}Xe polarisation achieved after 10 minutes. ^{4}He is significantly more expensive than N_2. Therefore, if there is no benefit in using ^{4}He over N_2, it would be logical to opt for N_2 as the buffer gas.

Chapter 4

Scanner Hardware

Within this document, the MRI scanners discussed are a 0.5T MROpen Upright Paramed Medical Systems scanner (where the patient can be positioned both upright and lying down) and a 1.5T GE flatbed scanner.

4.1 0.5T Paramed MRI scanner

Lung function and volume, as well as arterial oxygen levels, are affected by body position [6]. Having the scan volunteer upright should provide images more useful to clinicians because, in general, people spend most of their lives upright (either standing or sitting). Lung function is also reduced with the patient lying down, meaning patients with severe lung diseases may struggle to breathe when lying and even those with minor symptoms may struggle with the breath hold requirement for hyperpolarised ^{129}Xe MRI. Some patients with moderate to severe COPD can't lie down for an extended period of time because of the extra difficulty they have breathing. Having the patient sitting upright while the experiment is performed will make it a much more pleasant experience for the subject. The open design of the Paramed scanner has the additional advantage of being much more acceptable for people with claustrophobia and allows easier access for people with mobility issues. The upright MRI may not have as strong a magnetic field as other scanners, but as previously explained in section 2.3, for hyperpolarised MRI this doesn't result in the same reduction to the SNR of images as with ^{1}H MRI. At Nottingham, children as young as 3 years old have sat comfortably inside the scanner without sedation. Paediatric cohorts can be scanned sitting on a parent's or guardian's lap. This ensures they are more relaxed within the scanner and move little; removing the need for sedation.

4.1.1 Scanner Limitations

The Paramed scanner has a few constraints for acquiring images within a short duration. Some of these constraints are hardware related, but some of them are

due to the software being used to run the sequences.

The Paramed is conventionally run through the NRG clinical interface, using preloaded sequences. When using these sequences, there are limitations set for certain parameters, such as the minimum matrix size, TE and TR. In the clinical interface, the minimum matrix size for a 3D GRE sequence is 160x128x24 and setting the TE and TR to minimum results in an acquisition time of 108 seconds.

The clinical interface is built on top of sequences in TNMR (TecMag spectrometer's software). By running the scanner through TNMR, there is huge flexibility. For example: The gradient and RF amplitudes can be altered for each TR. They can also be given a prescribed shape and the gradient orientation can be altered. The timing of any event within the sequence can be changed and control loops can be inserted to alter the sequence depending on an external event.

Even when running through TNMR, there are still limitations set by the hardware. The maximum gradient is 20mTm^{-1} and the maximum slew rate is $33\text{Tm}^{-1}\text{s}^{-1}$. This means the minimum time to reach the maximum gradient is 0.6ms. To ensure an attempt to exceed the maximum slew rate wasn't inserted in error, all gradient rise and fall durations were set to 0.6ms.

4.1.2 Scanner Hardware Layout

The topology of the Paramed MROpen at Nottingham can be seen in figure 4.1. There are 2 separate RF amplifiers and 8 receive lines on a TecMag Redstone Spectrometer; four for multinuclear and four for proton.

For both multinuclear and proton imaging, the RF pulse generated by the Redstone Spectrometer is attenuated prior to amplification. Through TNMR, the desired attenuation of RF pulse can be selected on a logarithmic (dB) scale, with the power increasing by a factor of 10 every 10dB. See section 4.1.5 for more details on RF pulse power.

The Paramed Scanner is based on a 0.5T, MgB_2 high temperature superconductor, cryogen free, horseshoe magnet. The high temperature superconductor doesn't require liquid He. During maintenance the magnet current can be ramped down before being ramped up again. This can be done within a couple of hours and without the high cost of quenching a conventional MRI scanner. The current in the MgB_2 coils is driven (at all times) by a high stability power supply. This is different from a conventional MRI scanner, where there is a closed loop formed by a superconducting link between the coils.

4.1.3 Coil Design

The Paramed scanner has no built in body receive coil. The built in transmit coil is also unable to act as a receive coil, due to its topology using directional power dividers. There are however, a wide range of different receive coils, optimised for different parts of the body, in a range of sizes. The magnetic field of the Paramed is orientated in a different orientation to that of a conventional MRI scanner (see figure 4.2). This allows different coil designs to be employed. Ideally the

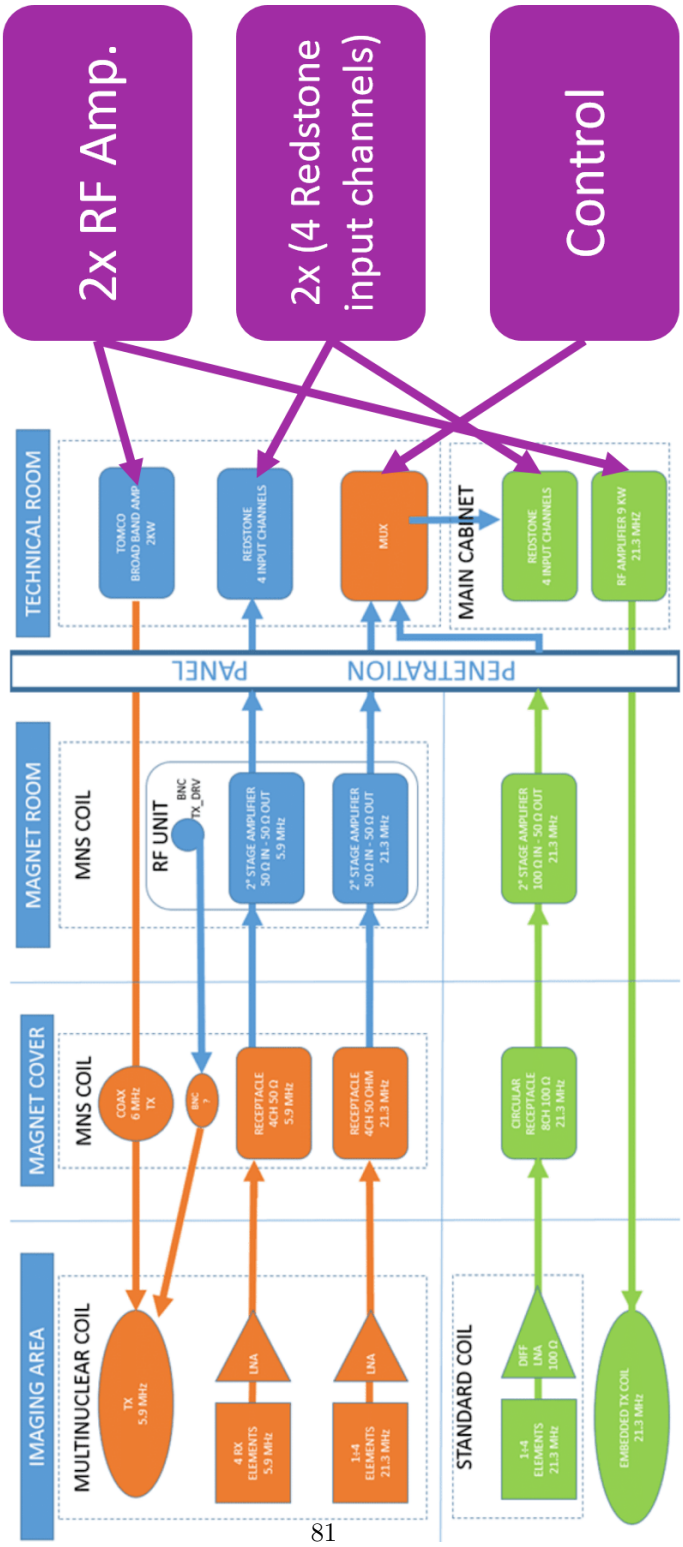

Figure 4.1: The topology of the Paramed MROpen at Nottingham. There are 2 separate RF amplifiers and 8 receive lines; four for proton and four for multinuclear. The unblanking signal allows an RF pulse to be transmitted and is mirrored between the proton and multinuclear side.

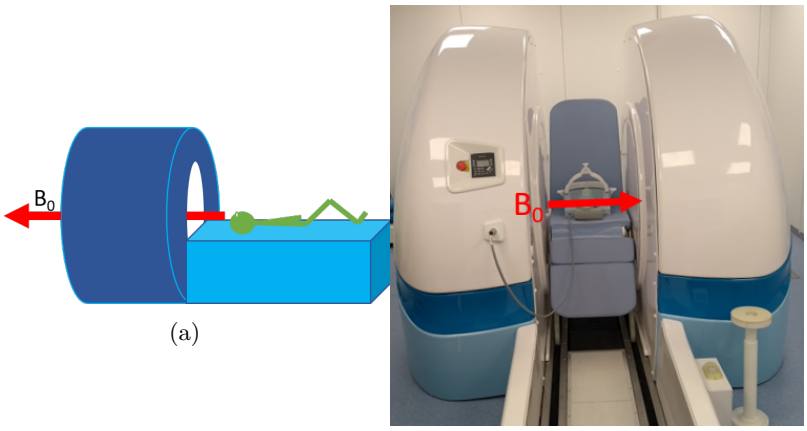

(a) (b)

Figure 4.2: Red arrow direction of B_0: **(a)** Conventional MRI scanner, **(b)** Paramed MRI scanner.

magnetic field produced or received by a coil should be perpendicular to the scanner's magnetic field (B_0).

If a coil is modelled as a simple loop, the magnetic field produced (B_1) by the coil is perpendicular to the loop forming the coil (see figure 4.3). The received magnetic field is also optimally orientated perpendicular to the loop forming the coil. In a conventional scanner, a solenoid loop can't be connected around a volunteer lying in the scanner because B_1 would be parallel to B_0 (see figure 4.4a), instead flat loops are wrapped around the volunteer so the magnetic field runs perpendicular to B_0 (see figure 4.4b). However, by switching the orientation of B_0 to into/ out of the page in figure 4.4; for both designs of coil B_0 is perpendicular to B_1. This is the same situation as with the transverse field of the Paramed.

A bespoke chest coil has been ordered from PulseTeq Ltd. for ^{129}Xe imaging. In the meantime, ^{129}Xe work was performed using a single channel, test, surface coil.

4.1.4 ^{129}Xe Surface Coil

^{129}Xe and ^{1}H resonate at different frequencies in the magnetic field of the scanner; due to their different gyromagnetic ratios (see table 2.1). The coils used to transmit the RF pulse and receive the signal back, have circuitry tuned to the Larmor frequency of the nuclear species in question (see section 2.1.12). A square surface coil, with a side dimension of 10cm, was purchased from Clinical MR Solutions L.L.C. (see figure 4.5). This was intended as a test coil to confirm a ^{129}Xe signal could be received and to determine the optimal frequency to tune a bespoke chest coil to before manufacture.

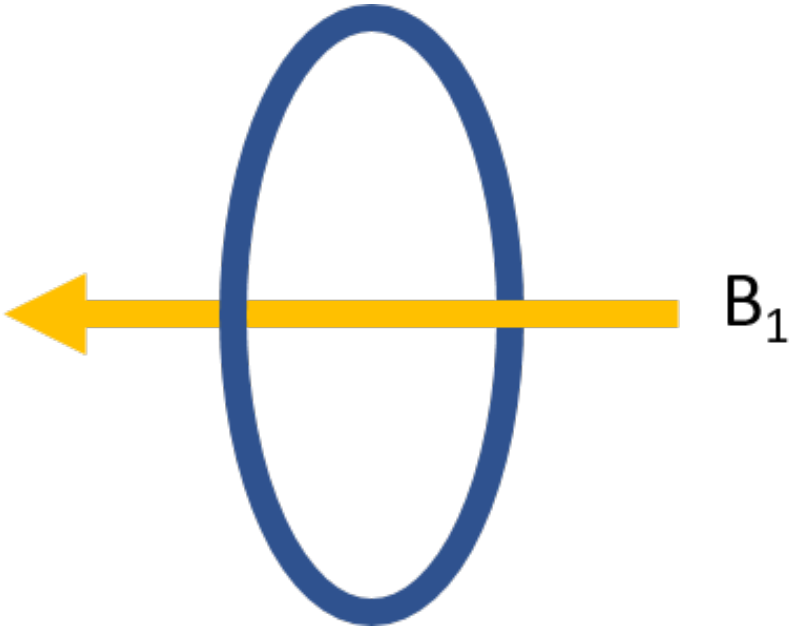

Figure 4.3: Simple loop modelling a coil (blue) producing B_1 (yellow).

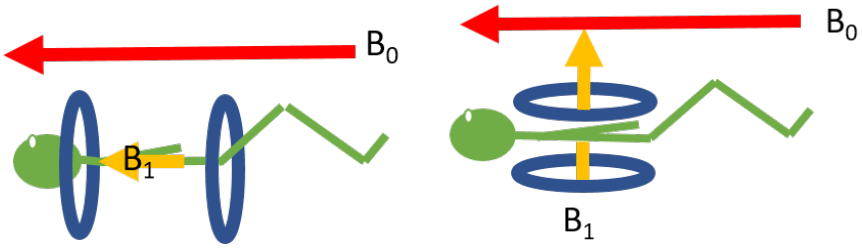

(a) (b)

Figure 4.4: Different coil designs in a flatbed scanner **(a)** B_1 parallel to B_0; coil doesn't perform optimally, **(b)** B_1 perpendicular to B_0; coil performs optimally. By changing the orientation of B_0 to into/out of the page for both situations B_1 would be perpendicular to B_0.

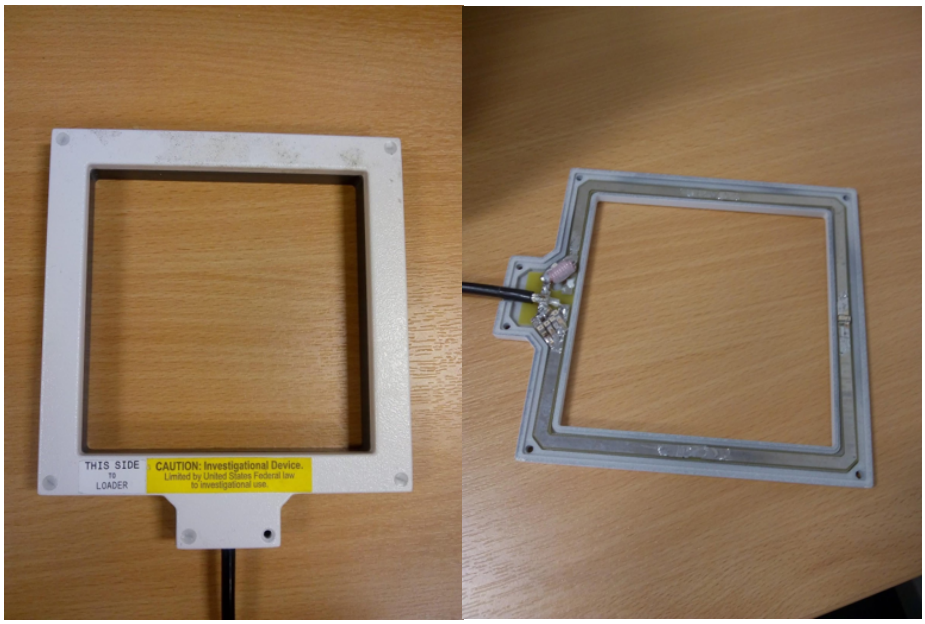

Figure 4.5: Original ^{129}Xe surface coil.

Coil Optimisation

Before Retuning The coil was loaded by placing two, 0.5L saline bags on the coil, see figure 4.6a. The coil was then connected to an Agilent Technologies E5061A 300kHz-1.5GHz ENA Series Network Analyser to determine the peak coil response frequency. The coil had a peak response, at 5.890MHz, of 28.016dB and a full width at half maximum (FWHM) of 43KHZ (see figure 4.2).

A low FA (<1°, see table 7.2), attenuation 70dB, 3ms RF pulse, with the same amplitude modulation as figure 7.10 was used in a basic pulse, acquire sequence, with no gradients and a TE of 30ms. By using a low FA it was possible to interrogate the bag multiple times without reducing the polarisation of the bag significantly, even when at the resonant frequency.

The surface coil exhibits a small DC offset artefact. Even without a Xe phantom close to the coil, after performing an FT on the time domain signal a small peak can be observed at the central frequency. The ^{129}Xe peak was first differentiated from this artefact by changing the pulse length and ensuring the Xe resonant frequency inside the scanner was ∼200Hz from the central frequency of the acquired signal. Changing the RF attenuation level can also change the intensity of artefacts, but changing the pulse length should just affect the ^{129}Xe NMR signal. Investigations were started using a 1L bag of naturally abundant Xe at 9% polarisation.

The ^{129}Xe peak centre frequency was determined at 5.841MHz. Repeated pulses at this frequency saw the intensity of the FID drop by a very small

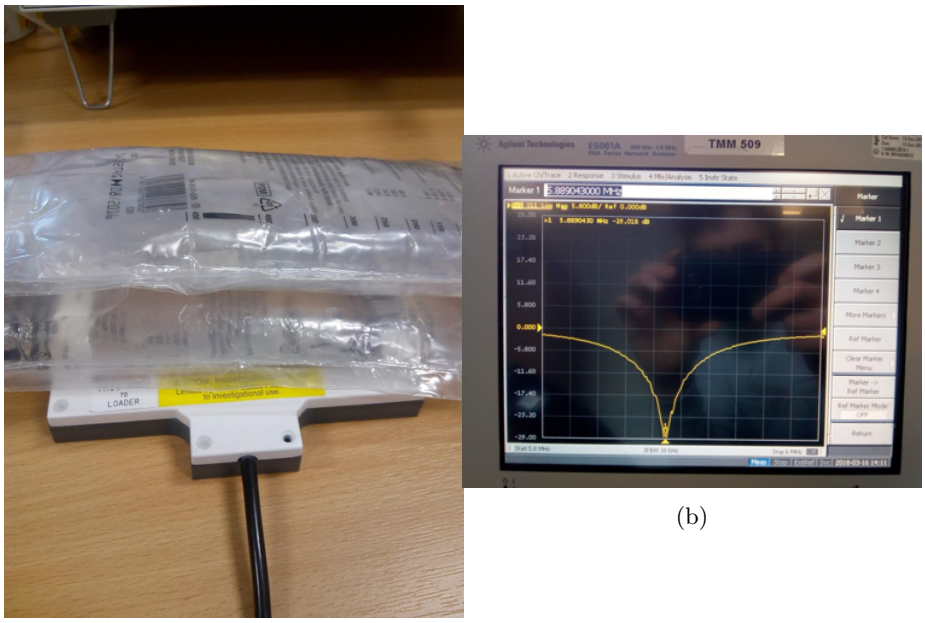

(a) (b)

Figure 4.6: **(a)** Loaded surface coil, **(b)** Peak response of 28.016dB at 5.890MHz.

amount (<1%), confirming the small FA. After confirming 70dB was a high enough attenuation, the bag polarisation was measured and found to be 5.73%. Using the 70dB attenuation RF pulse, a signal with peak intensity of 8222 (using TNMR's arbitrary scale) was received at 5.841MHz. At this frequency, the coil offers less than half its peak response, see figure 4.6b.

After Retuning The Surface coil was retuned to put its peak response at closer to the ^{129}Xe resonant frequency inside the Paramed scanner. By changing the tuning capacitance, it is possible to alter the peak response frequency. A 33pF, high voltage (3KV) capacitor was added in parallel (see figure 4.7a). This changed the peak response to 33.039dB at 5.842MHz and the FWHM to 27kHz. This was with the coil loaded in an identical manner to previously (see figure 4.6a). While the FWHM is reduced from 43kHz to 27kHz, the peak response is now higher and much closer to the ^{129}Xe resonant frequency inside the scanner; with the Larmor frequency comfortably within the FWHM.

Using the same pulse acquire sequence and a new 1L bag with 12.25% polarisation, the peak signal intensity received was 27623. Comparing this to 8222 previously received with a 5.73% polarisation bag and factoring in the polarisation difference; this suggests there is now a **57% improvement in signal** at the centre frequency after retuning the coil.

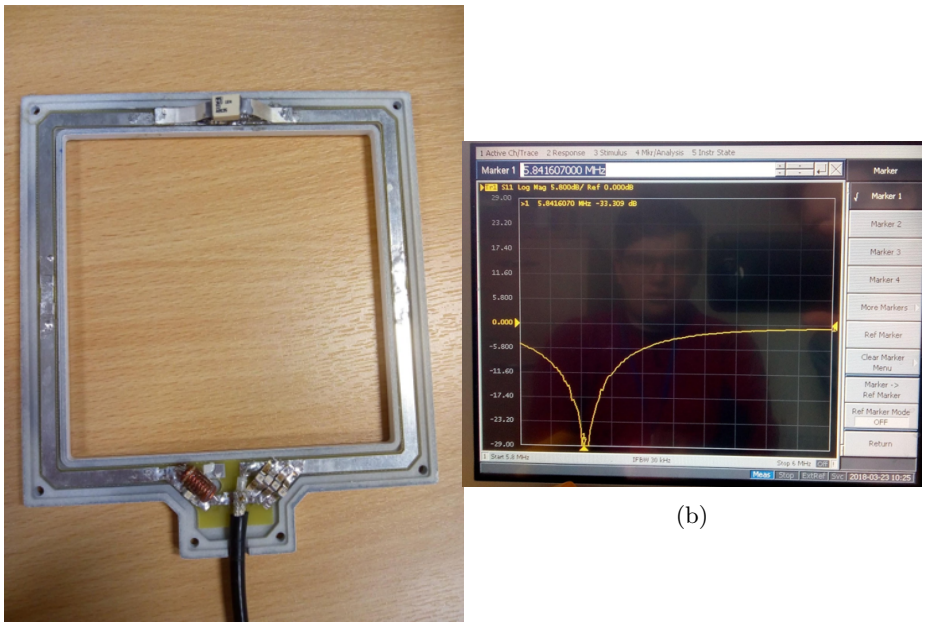

(a)

(b)

Figure 4.7: **(a)** Retuned surface coil with added 33pF capacitor, **(b)** Peak response of 33.039dB at 5.842MHz.

4.1.5 Specific Absorption Rate Monitoring

Specific Absorption Rate (SAR) refers to the RF power delivered to a tissue and is generally expressed in (Wkg^{-1}). The energy absorbed by the tissue, results in tissue heating. Therefore, it is important to set a safe limit for SAR.

The Paramed MRI scanner, in its standard form, doesn't employ any form of active SAR monitoring. Instead, when using the standard sequences and coils, even when using the highest possible tipping angle and smallest FOV, they have all been previously modelled to ensure it is impossible to exceed safe levels.

The power delivered to tissues will not be homogeneous. Instead it will vary depending on the coil architecture. As previously covered in section 2.1.12, surface coils only tend to influence a region with a depth of the order of the diameter of the coil. Unlike volume coils, surface coils have an inherently inhomogeneous B_1 field. This can result in a large proportion of the RF power delivered to the coil, being transferred to small regions of tissue. With the ^{129}Xe surface coil (see section 4.1.4), this needed to be modelled to check if some form of active SAR monitoring is required and to observed how hotspots will vary depending on the location of the coil in relation to the body.

Thanks must go to Benjamin Prestwich and Dr Andrew Peters who kindly performed the coil simulations.

The surface coil was simulated, centred behind the chest. The maximum 10g average SAR was found to be 0.3891 Wkg^{-1}, for 1W of input power. Looking at table 7 and 8 in "Safety Guidelines for Magnetic Resonance Imaging Equipment in Clinical Use" from the MHRA [131]; for "Local Trunk" in "Normal Mode" the SAR is limited to 10 Wkg^{-1}. Normal mode of operation refers to: "Exposure of extended volumes of the body should be such as to avoid a rise of more than 0.5°C in the body temperature of patients and volunteers, including those compromised with respect to their thermoregulatory ability." [131] Keeping within this guideline places a limit of 25.7W on the average forward power over the whole scan.

Power Calibration

The Paramed MRI scanner multinuclear side, employs a Tomco Technologies BT02000-AlphaS, 100kHz-30MHz, 2kW RF amplifier. The maximum pulse width is 10ms, with a rated power of 200W in Continuous Wave (CW) Mode [132].

The Redstone spectrometer was setup so 0dB RF output produces a 1V Peak to Peak voltage (VPP) into 50 Ω. Looking at the data sheet for the scanner's Redstone, this equates to -4dBm of attenuation; where dBm is the calibrated dB scale into a 50 Ω load (from the data sheet) [133]. At 0dBm, the Tomco RF amplifier is able to produce at least 3000W for frequencies up to 20MHz and is broadly consistent across this range (see figure 4.8). This means that a set attenuation of 0dB would force the amplifier into clipping. At 5.8MHz (the ^{129}Xe Larmor frequency), the Tomco RF amplifier produces $\sim 3000W$ [134].

The theoretical power outputs from the Tomco RF amplifier for a 5.8MHz

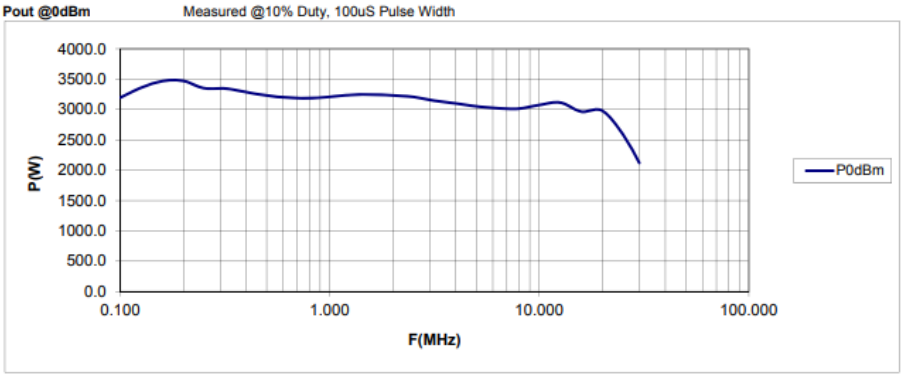

Figure 4.8: Power output for a range of frequencies at 0dBm. Taken from the Tomco Technologies RF Amplifier Factory Performance Certificate for the amplifier used by the Paramed at Nottingham [134].

pulse can be seen in table 4.1. At 4 dB of attenuation, the output of the Redstone is 1mW. At 6 MHz this means the Tomco RF amplifier provides $\sim$ 64.8dB of gain in order to produce 3000W.

Before defining safe levels from data sheets, the actual outputted RF for a given attenuation needed to be investigated.

The Tomco RF amplifier allows the forward power to be sampled using a built in $\sim$50dB output of the transmitted RF pulse. This output is labelled "RF SAMPLE" on the amplifier. By integrating this voltage profile of RF SAMPLE, the average power can then be extracted (see section 4.1.5).

For investigations into the accuracy of the RF SAMPLE profile, the same 3ms RF pulse which is employed in the ^{129}Xe GRE sequence was used (see figure 7.10). The expected pulse profile and the output of RF SAMPLE on the scope can be seen in figure 4.9. There is good agreement between the two profiles, with approximately the same ratio in amplitude between the first and second lobes.

To calibrate the amplitude of RF SAMPLE to a given power, a continuous wave with no amplitude modulation was used for 30s. A Bird RF Power Analyst model 4391 was used to measure the actual forward power and the VPP measured by the scope was recorded. The results can be seen in table 4.2.

Theoretically, the VPP should halve every 6dB decrease in attenuation and the power every 3dB. Therefore, when plotting the log of the power measured against the log of the VPP the gradient should be 2; see figure 4.10. There is good agreement with the data. The fitted line has a gradient of 1.954, but a gradient of 2 is within the 95% confidence interval of the fitted gradient.

When the Bird power meter receives an RF pulse, it hasn't come directly from the Redstone to the Tomco RF amplifier and then on to the Bird Power meter. The pulse has already gone through a duplexer and a 5m BNC cable. These are unlikely to be lossless and may create an additional attenuation offset

Redstone Attenuation/ dB	Attenuation/ dBm	CW Power (W)
0	-4	Clipping with RF ~3200
4	0	3000
14	10	300
15	11	238.2
16	12	189.3
17	13	150.3
18	14	119.4
19	15	94.8
20	16	75.3
21	17	60
23	19	37.8
26	22	18.9
29	25	9.6
32	28	4.8

Table 4.1: Theoretical Power after amplification by the Tomco RF amplifier for a range of attenuations of the Redstone (dB) and the calibrated attenuation scale into a 50 Ω load (dBm). This is using a 5.8MHz RF pulse.

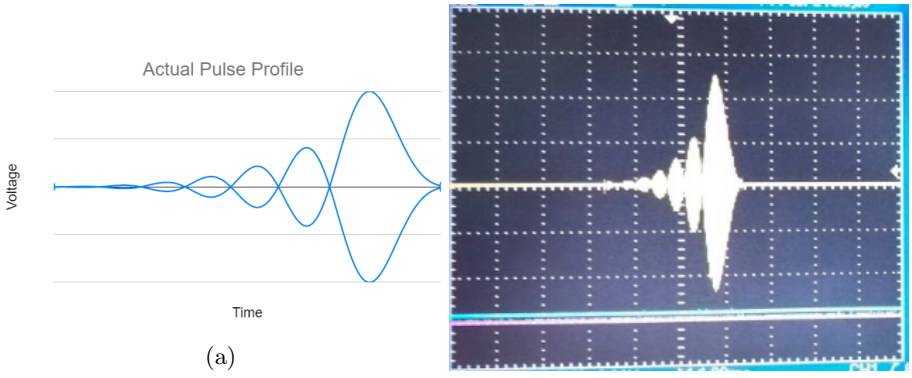

(a) (b)

Figure 4.9: 3ms RF pulse:**(a)** expected pulse profile, **(b)** RF SAMPLE pulse profile on the scope.

Attenuation/ dB	VPP RF SAMPLE (mV)	CW Power (W)
14	1950	200
17	1420	117
20	980	63
23	680	31
26	504	16
29	360	7.6
32	288	5

Table 4.2: Measured Power and VPP of RF sample for a range of attenuations (5.8MHz Pulse).

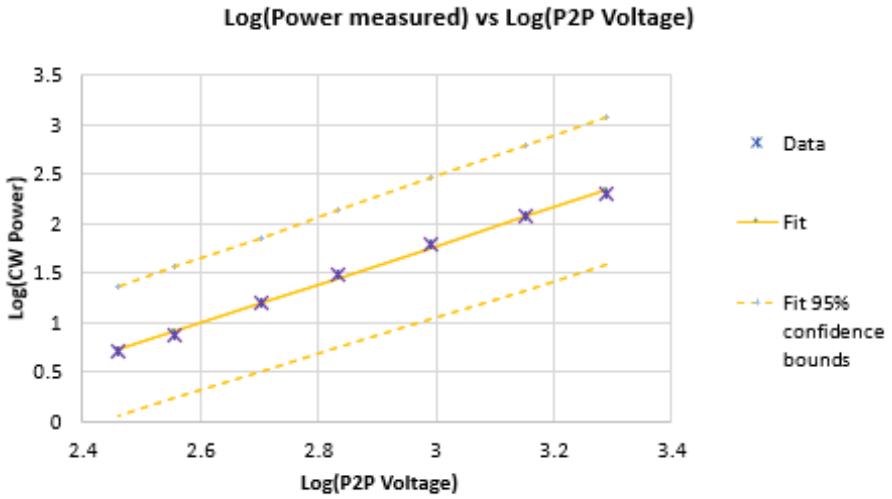

Figure 4.10: Log of the power measured against the log of the VPP fitted to $y = a+bx$, where $a = -4.086$ $(-4.441, -3.732)$ and $b = 1.954$ $(1.83, 2.07)$. Numbers in brackets are each coefficient's max and min values in a 95% confidence interval. The R^2 of the fitted line is 0.996.

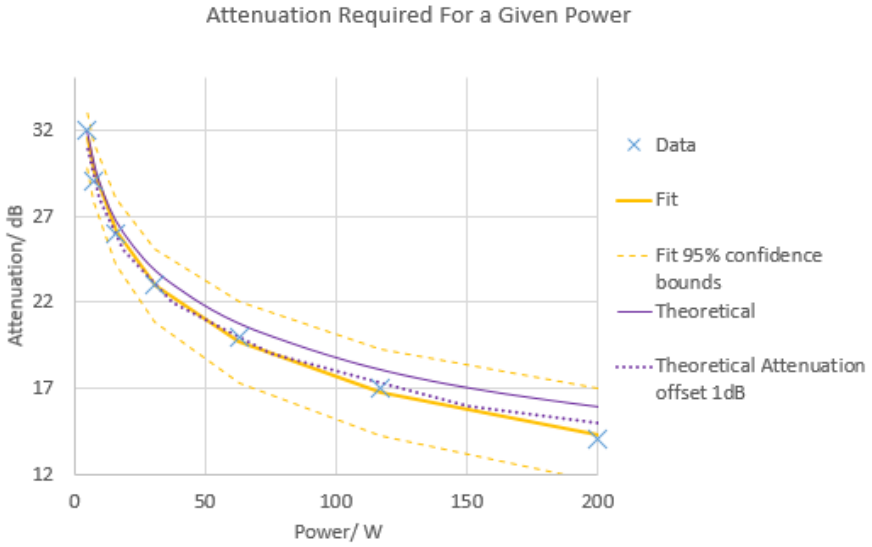

Figure 4.11: Attenuation required for a given power fitted to $y = a + blog(x)$ where $a = 39.07$ $(37.97, 40.17)$ and $b = -10.77$ $(-11.46, -10.08)$. Numbers in brackets are each coefficient's max and min values in a 95% confidence interval. The R^2 of the fitted line is 0.996. The theoretical power for a given attenuation can also be seen; including with a 1dB offset.

$\sim$1db from the dBm scale.

In order to calculate the attenuation offset, the attenuation (y-axis) was plotted against power measured (x-axis) and fitted to $y = a + blog(x)$; see figure 4.11. The theoretical powers from table 4.1 can also be seen.

Offsetting the theoretical attenuations by 1dB provides a good match with the data. This suggests there are $\sim$ 1dB of losses through the duplexer and the BNC cables into the Bird Power Meter.

The exact attenuation of the RF output before it leaves RF sample also needed to be calculated. The attenuation (y-axis) was also plotted against the VPP from RF sample (x-axis) and fitted to $y = a + blog(x)$; see figure 4.12. A 20dB gain from the Redstone output provided the best match with the data; assuming the Tomco RF amplifier provides 64.8dB of gain. This means the RF output of the Tomco RF amplifier has an additional attenuation of $\sim$44.8dB before it leaves RF SAMPLE (see figure 4.12).

SAR Monitor Design

An SAR monitoring circuit has been ordered from the electronics workshop in the physics department. This circuit will average the VPP from RF sample with a timebase of <100ms. If the average VPP goes above a preset limit (which can be changed), then a flag is raised. In the event of a flag the circuit pulls pin 19

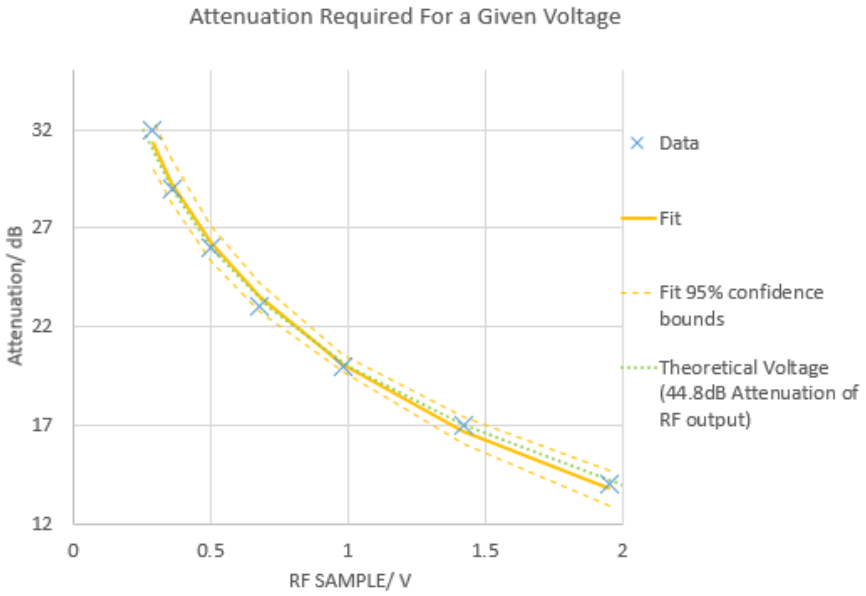

Figure 4.12: Attenuation required for a given VPP fitted to $y = a + blog(x)$, where $a = 19.48$ $(19.48, 20.41)$ and $b = --21.07$ $(-22.53, -19.61)$. Numbers in brackets are each coefficient's max and min values in a 95% confidence interval. The R^2 of the fitted line is 0.996.

on the Tomco RF amplifier 25-pin parallel interface connector to ground (pin 21 is ground) [135].

Pulling pin 19 to ground inhibits transmission of RF, but doesn't switch off the power supply unit. The red SHUTDOWN LED on the Tomco will light red when the Tomco RF amplifier has been shutdown [135].

Summary The first important note is that the multinuclear side should not be run in normal operation at attenuations less than 4dB. This is to avoid exceeding the Tomco RF amplifier's rating of 3000W

For the ^{129}Xe frequency, theoretically the Tomco RF amplifier produces a gain of $\sim$ 64.8dB. Setting the Redstone to 4dB of attenuation, produces a $\sim$3000W forward power at 5.8MHz. The data acquired showed excellent correlation with this, especially if 1dB of additional attenuation due to losses in the system is factored in. The data acquired suggests the theoretical RF power values for a given attenuation can be trusted. Because the theoretical power values are slightly higher than the measured values, by using the theoretical values in the calibration this will provide an extra level of safety.

The RF SAMPLE VPP measured, showed good agreement with a 20dB gain from the theoretical output from the Redstone. After plotting the measured voltages against the theoretical powers, the data can be fitted to equation 4.1.

$$V_{sample} = 10^{alog10(P_{out})+b} \tag{4.1}$$

V_{sample} is the VPP from RF SAMPLE (volts) and P_{out} is the forward power from the Tomco in (KW). There is excellent agreement between the data and equation 4.1, where a = 0.4909 (0.4671, 9.5148) and b = 0.5482 (0.5267, 0.5697); numbers in brackets are each coefficient's max and min values in a 95% confidence interval. Between data and fitted curve there is an R^2 of 0.999. This allows the VPP measured by the SAR monitoring circuit from RF SAMPLE to be converted to a value for forward power.

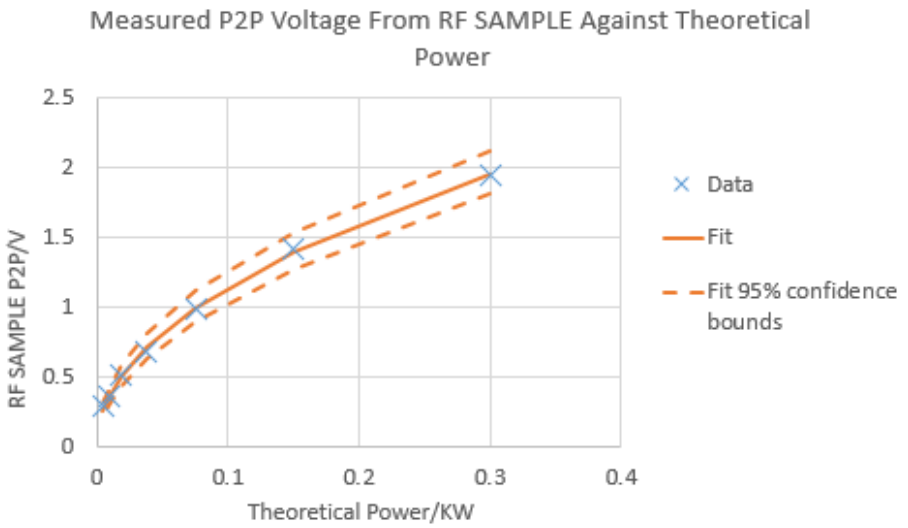

Figure 4.13: The measured RF SAMPLE VPP against the theoretical power outputs from the Tomco. Curve fitted to equation 4.1; R^2=0.999.

4.2 1.5T GE scanner

Previous ^{129}Xe imaging at Nottingham in human cohorts can be seen in Dr James Thorpe's thesis [136] and Dr Shahideh Safavi's thesis [137]. This work was performed on a flatbed GE scanner at 1.5T, using a Rapid Biomedical Coil.

4.2.1 Rapid Biomedical Coil

The Rapid Biomedical coil has a birdcage design and rigid outer shell (see figure 4.14). The birdcage design creates a homogeneous transmitted B_1 field and the rigid outershell ensures the B_1 field alters little between volunteers. There are receive arrays built in to the bottom half of the coil and a flexible blanket is placed on top of the volunteer containing more receive arrays. This ensures the receive arrays are close to the origin of the signal, even if a thin volunteer is scanned. However, the rigid coil effectively decreases the bore size of the magnet from 578mm to 436mm (with the top receive array out) and traps the arms by the volunteer's side. It is impossible for many people to remove themselves from the magnet bore; increasing any claustrophobic effects. Also larger patients simply didn't fit inside the coil.

4.2.2 Clinical MR Solutions Coil

An improved coil design was required, which would allow for the scanning of larger volunteers. A new coil was ordered from Clinical MR Solutions L.L.C.

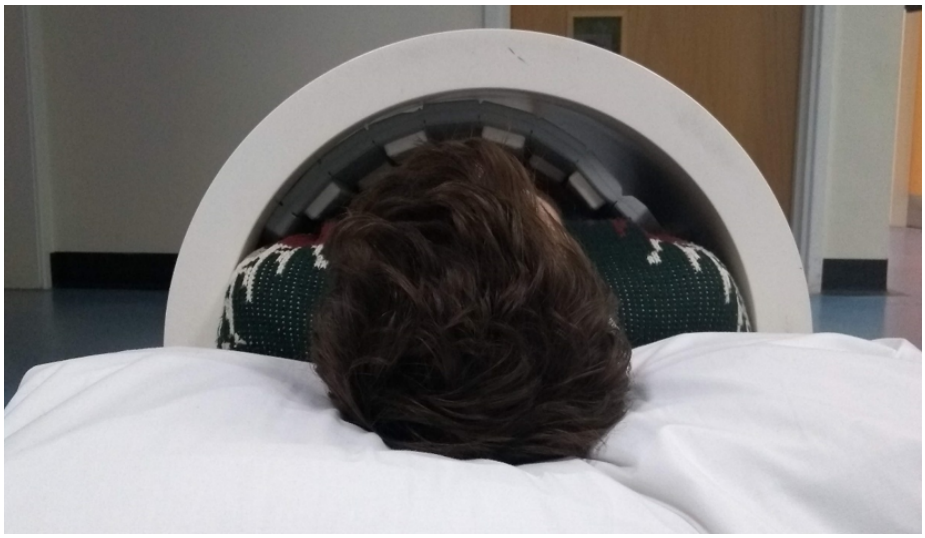

Figure 4.14: Rapid BioMedical coil.

The coil is of the same design as used by Roos et al. [8] (see figure 4.15). This new coil is put on like a jacket and the Velcro straps make it adjustable for the patient. This results in a more pleasant experience for the volunteer while being scanned. The volunteer's arms are no longer trapped within the coil by their side and it is much simpler for someone to remove themselves from the magnet bore.

The new coil has two modes that it can be run in "Quadrature Transmit Array Receive" (QTAR) and "Quadrature Transmit and Quadrature Receive" (Quad T/R). The different modes used for scanning can be seen in figure 4.18.

In both modes an outer jacket transmits the RF pulse. The outer jacket is formed of 4 coils and driven in quadrature (see figure 4.16). The RF amplifier output is split, producing 2 outputs of equal magnitude. A phase shifts is introduced to produce 2 identical amplitude signals, with a relative phase of 90°. These signals are sent down 2 separate cables; one supplying coils 1 and 3 and the other supplying coils 2 and 4. A 180° phase shift is then introduced between coils supplied by the same cable.

This results in each coil being 90° out of phase of the adjacent one, with coils diagonally opposite in the closed jacket (coils 1 and 3 or coils 2 and 4 in figure 4.16), being 180° out of phase. The coils with a 180° phase shift produce a magnetic field between them. Theoretically, the field produced by the two pairs of coils produces a more uniform B_1 field, than if just one pair of coils was used. However, because the coil is flexible, depending on its shape, the B_1 uniformity may be impaired significantly.

The receive is handled differently in both modes. In QTAR mode, two pads each with four receive arrays are placed inside the outer jacket. Each of these

(a) (b)

Figure 4.15: **(a)** Rapid BioMedical coil, **(b)** Clinical MR Solutions L.L.C. Coil.

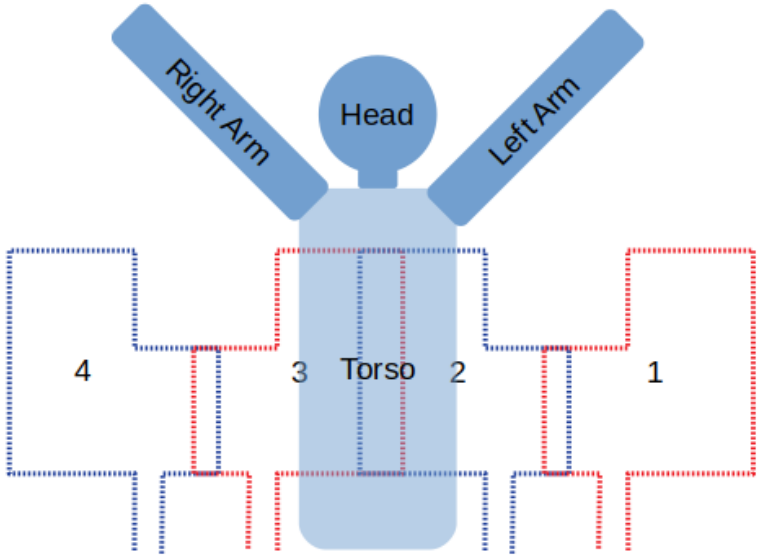

Figure 4.16: Diagram showing volunteer lying inside open outer jacket used for quadrature transmit (and quadrature receive in Quad T/R mode). Coils are numbered 1 to 4. Coils the same colour, have a 180° phase difference, with a 90°phase difference between coils 1 and 2.

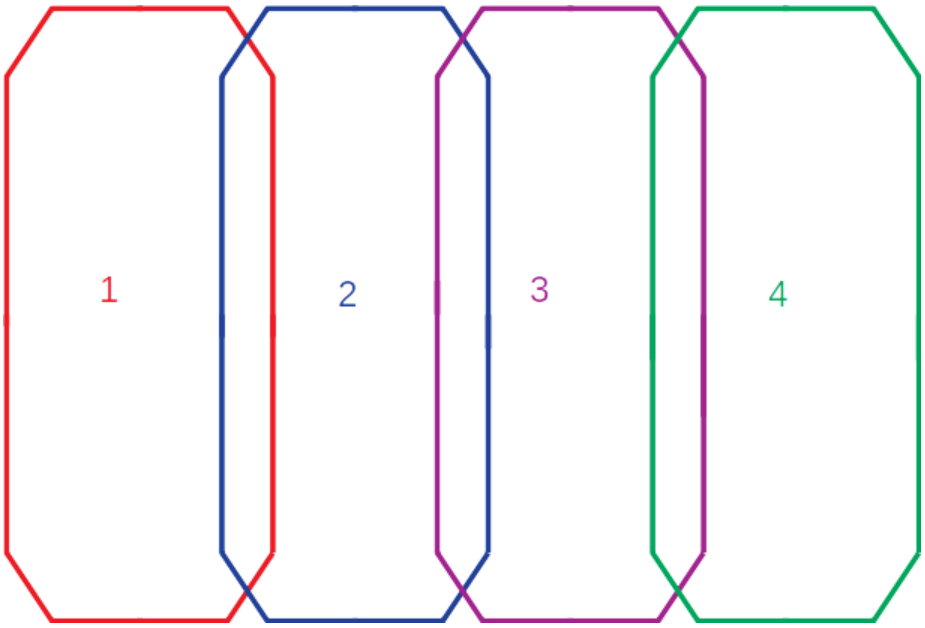

Figure 4.17: Diagram showing the arrangement of the receive arrays inside one of the two pads used in QTAR mode.

receive arrays works independently as a surface coil (see figure 4.17). However, in QUAD T/R mode, the coils used for RF transmit in the outer jacket are also used for receive. Theoretically, using the receive arrays should improve the SNR. However, using them requires more complicated switching to ensure the receive arrays don't interfere with the transmitted RF and ensure any components used to amplify signals in the order of mV are detuned during the RF signal in the order of kV (see section 2.1.12).

The black box in figure 4.18 is a duplexer and also contains the GE coil ID chip. Without this ID chip, the scanner won't recognise the coil and will refuse to run the sequence.

(a) (b)

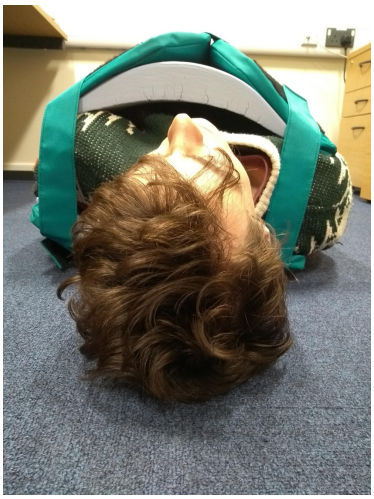

(c)

(d)

Figure 4.18: Clinical MR Solutions L.L.C. Coil depicting different modes used in imaging **(a)** QTAR, **(b)** Quad T/R, **(c)** Quad T/R with foam pads and **(d)** Duplexer with GE coil ID chips enclosed which is used to connect the coil to the scanner.

Chapter 5

1.5T GE ^{129}Xe Imaging

Previous Xenon imaging at Nottingham, in human cohorts, can be seen in Dr James Thorpe's thesis [136] and Dr Shahideh Safavi's thesis [137]. The sequence used for imaging was a GE 2D Fast GRE (FGRE) that was converted to the ^{129}Xe frequency by Dr Steve Hardy and Dr Brett Haywood. This is the same sequence which was used with previous imaging using the Rapid BioMedical coil (see figure 4.15). Details about the new coil from Clinical MR Solutions L.L.C. can be seen in section 4.2. This coil can be run in two modes QTAR and Quad T/R mode (see figure 4.18).

For volunteer imaging using ^{129}Xe, all volunteers gave written informed consent under National Research Ethics Service reference number 13/EM/0401.

5.1 Finding the Centre Frequency

To find the correct frequency for Xe in the scanner, the sequence FID CSI was run with a TR of 3s and FOV of 48cm. Spectro Pre-scan was opened and the transmit gain set to the minimum. A frequency of 17663380Hz was found to be approximately central for Xe. However, the frequency did shift by tens of Hz between days. With a volunteer inside the coil an Xe bag can be placed on top of the coil and the centre frequency found before the volunteer has breathed in any Xe. It was found that the centre frequency shifted by <10Hz between scans, even when the volunteer gets in and out of the magnet. Therefore, the Spectro Pre-scan only needs to be performed before the first Xe image.

5.2 QTAR mode

To begin with, 10mm thick slices were used. The minimum matrix size was also chosen, with 128 frequency encodings and 128 phase encodings, to keep the scan short. With Xe imaging, there is no benefit from retaining polarisation for a subsequent scan, because after breathing out, the Xe is no longer in the lung. Looking at equation 7.2 in the limit T1>>TR, the magnetisation after

pulse number k (M_k) drops according to equation 5.1; assuming the FA (α) is the same for all pulses.

$$M_k = M_0(cos\alpha)^{k-1} \tag{5.1}$$

For a 2D sequence, assuming no diffusion within the bag, each slice of the image is excited separately. Therefore, for 128 phase encodings k=128. It isn't necessarily desirable to deplete all of the polarisation. If the drop in polarisation towards the later phase encodings for each slice isn't accounted for, this will result in the last phase encoding steps having many times less signal than the first. A good ball park to aim for to begin with, would be losing half the polarisation after the scan; a good balance of maximising the polarisation used up and ensuring the polarisation throughout the scan doesn't drop too significantly. For k=128, a half polarisation drop by the end of the scan equates to a FA of 6°. This assumes that depolarised Xe from one slice won't have time to diffuse into the next slice before the next slice is imaged. Even if diffusion happens between excited slices inside the bag, this is less likely to be the case in the human lung. The lung exhibits a microstructure of bronchioles and alveoli which means gas diffusion is greatly inhibited when compared to free diffusion inside the bag. The FA calculation also doesn't account for an imprecise slice selection, whereby some of the Xe is depolarised in the surrounding slices. Effects from diffusion and an imprecise slice selection can both be mitigated by adding a spacing between slices. However, the other assumption made is that the FA throughout the Xe is consistent; this won't be the case if the B_1 field produced by the coil is inhomogeneous. The RF power transmitted could have resulted in $> 6°$ FA in some sections of the bag and $< 6°$ FA in others. This could make selecting the optimum FA difficult, because even if the average FA over the whole sample was optimum, the FA may be too high in certain sections and too low in others. Even if the FA over a sample was calculated for different locations, this will likely change when the coil was loaded with a new sample.

10mm Slice Optimum Flip Angle With 10mm slices, 128 by 128 matrix size and a 48cmx48cm FOV, the optimum strength of the RF pulse was investigated. The coil was loaded with saline bags and a cavity in the centre of the saline bags was created where a Tedlar bag of Xe was placed. After each image, the bag was taken out of the scanner and the polarisation measured before being placed back in the cavity in the same location. Every effort was made to put the bag back in the same location and orientation. The attenuation was altered and the approximate FA calculated (see table 5.1).

A pulse with attenuation 8dB on a 15% polarisation bag offered a poor signal (see figure 5.1a), the drop in polarisation wasn't measured, but the low SNR image (SNR = 11.0 for first slice of figure 5.1a) indicates the RF pulse was too weak. The image produced using an attenuation -2dB RF pulse offered higher SNR (SNR = 34.1 for first slice of figure 5.1b). This was decided upon for imaging.

RF Attenuation/ dB	Bag Polarisation Before/ %	Bag Polarisation After/ %	Approx FA/ degrees
10	7.5	6.5	2.7
6	6.5	4.85	3.9
2	4.85	2.5	5.9
0	16.9	7.5	6.5
-2	18.3	7.4	6.9
-4	7.4	2.2	8.0

Table 5.1: Approximate FA for a range of attenuations; coil loaded with saline. 2D FGRE sequence, 10mm slices.

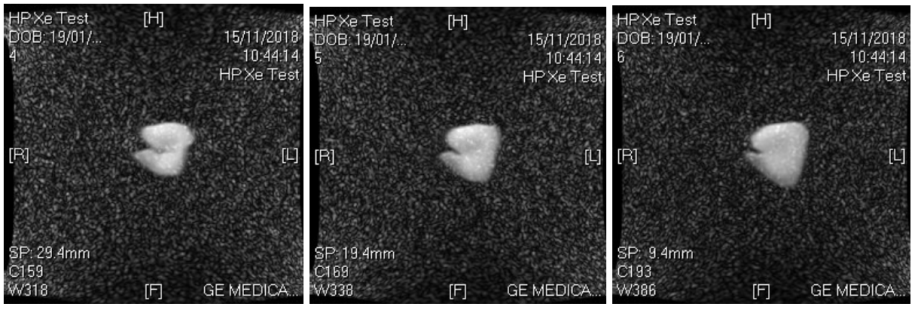

(a)

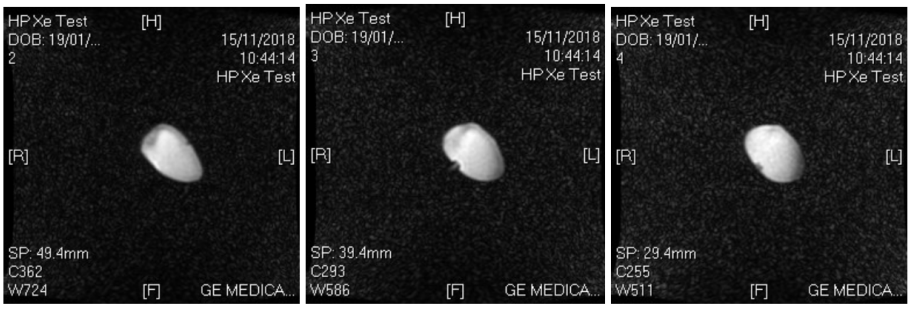

(b)

Figure 5.1: 48cmx48cm 2D-FGRE 10mm slices **(a)** RF attenuation 8dB (slice 1, SNR = 11.0) **(b)** RF attenuation -2dB (slice 1, SNR = 34.1).

Xe Diffusion All FA approximations of this 2D FGRE sequence are based on the assumption that diffusion is slow enough to discount. While the diffusion assumption may be relevant in the lungs, it may not be within the Xe bag.

A 10mm slice was imaged 8 times in the same location, with a 1s gap between the images. The sequence used employs 128 $\sim 6^\circ$ RF pulses per image, which was chosen to deliberately depolarise the excited slice. If diffusion is high compared to the gap between the images, the SNR of the excited slice will reduce by a small amount because fresh hyperpolarised ^{129}Xe will diffuse into the excited slice between the acquisitions However, if diffusion is low over the time gap between the images, the SNR will drop by a large amount. Also under these circumstances only a small proportion of the ^{129}Xe is being depolarised. Therefore, if a period of time is left to allow new hyperpolarised ^{129}Xe to diffuse back into the slice, the SNR of a subsequent image should be of similar SNR to the first image.

With a 1s gap between the images, the SNR was 11.8 and 2.48 for the first and eighth slices respectively (see figure 5.2a and 5.2b). The bag was left in the same location for 5 minutes to allow the polarisation of ^{129}Xe to become homogeneous throughout the bag. No scans were performed during this 5 minutes. The same scan was then repeated but this time with a 3s gap between the images (see figure 5.2c and 5.2d). This time the SNR was 11.5 and 3.80 for the first and eighth slices respectively.

The ratio for the reduction in SNR over the 8 images was 4.76 for the 1s gap set and 3.03 for the 3s gap set. For the 1s gap set there was 7s gap between the first and eighth image and for the 3s gap set there was 21s gap between the first and eighth image. This illustrates that inside the bag there is noticeable diffusion of ^{129}Xe into the excited slice over 18s. However, this was much smaller than the diffusion of ^{129}Xe into the excited slice over 5 minutes; the SNR if the first image in each set reduced little from 11.8 to 11.5.

The time for diffusion will also be greater in the lungs, where the small airways restrict diffusion. With a scan in volunteers, there will be a maximum of 1s between adjacent slices being scanned. This should mean there will be minimal polarised ^{129}Xe diffusing into the next slice when acquiring an image.

Coil Loading With a healthy volunteer inside the scanner, a coronal 2D FGRE sequence was performed, with an FOV of 48cmx48cm and 10mm thick slices. The RF attenuation was chosen to be -2dB. The images were very low in signal (see figure 5.3).

To investigate how a volunteer loads the transmit coil differently to a phantom and subsequently the FA of the R.F pulse, a Xe bag was placed on a volunteer's (V3's) chest inside the coil. It was expected that the FA may be slightly higher with a human loading the coil rather than saline bags because the coil's response at Xe's central frequency was optimised for a human loading the coil.

The polarisation before and after was recorded and the FA estimated (see table 5.2). However, the location of the bag has also changed, from the centre

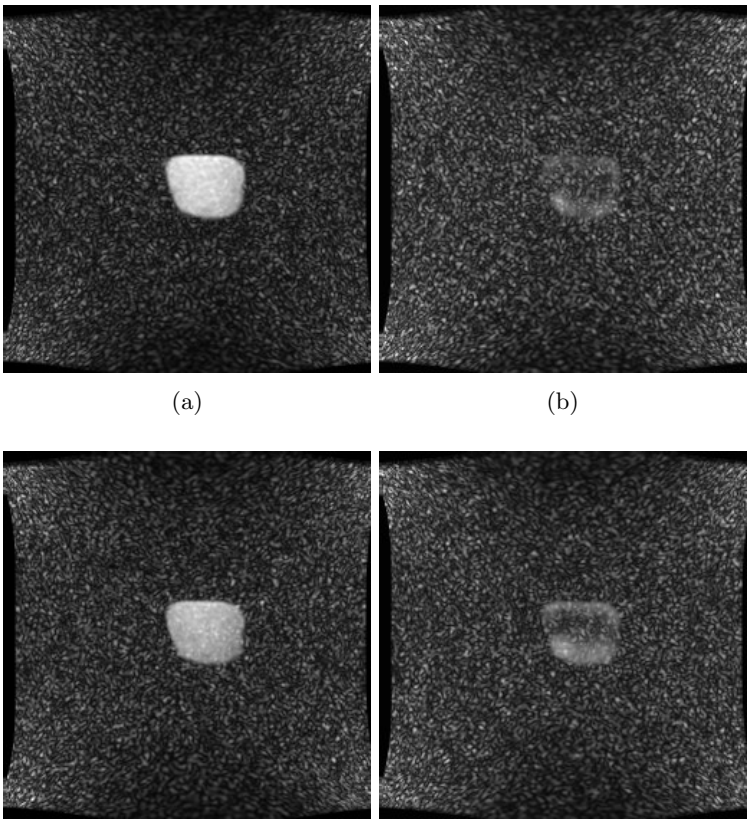

(a) (b)

(c) (d)

Figure 5.2: Two sets of images of same 10mm slice. First set acquired with 1s gap between images. Second set acquired with 5s gap between images. 5 minutes gap left between first set of images and second set. **(a)** first slice of 1s gap set (SNR=11.8) and **(b)** eighth slice of 1s gap set (SNR=2.48). **(c)** first slice of 8s gap set (SNR=11.5) and **(d)** eighth slice of 8s gap set (SNR=3.80).

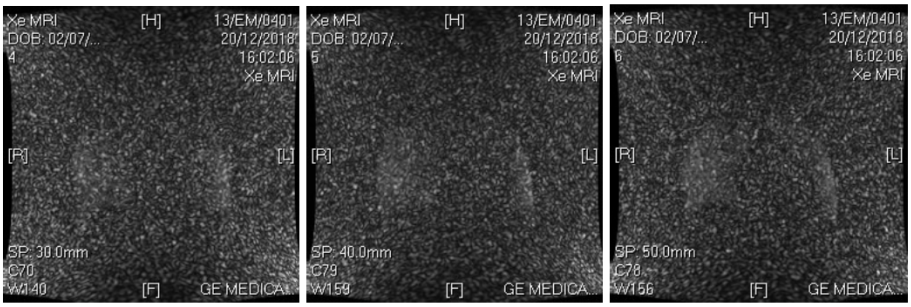

Figure 5.3: Volunteer image of lungs. 1L 15% polarisation Xe.

RF Attenuation/ dB	Bag Polarisation Before/ %	Bag Polarisation After/ %	Approx FA/ degrees
20	13.1	11.8	2.3
18	11.8	10.5	2.5
16	10.5	8.95	2.9
14	8.95	7.3	3.3
14	16	13.1	3.2
12	7.3	5.7	3.6
10	5.7	3.6	4.9
8	3.6	2.1	5.3
8	16.6	9.8	5.3
6	2.1	1	6.2
6	9.8	4.4	6.5
4	4.4	1.7	7.1
4	17	7.1	6.8
2	7.1	2.4	7.5
0	2.4	0.6	8.5

Table 5.2: Approximate FA for a range of attenuations 2D FGRE sequence 10mm slices. Coil loaded with a volunteer (V3) and bag on volunteer's chest

of the coil for table 5.1, to close to the top set of receive arrays for table 5.2. In an ideal coil, the FA should be consistent across the volume, but this may not be the case.

Comparing table 5.2 to table 5.1, the FAs were higher in the setup with the coil was loaded by a human rather than saline bags; for a given RF pulse attenuation. If the FA was too high at an attenuation of 2dB with a human loading, this could explain the poor images from figure 5.3.

RF attenuations of 14dB, 8dB and 4dB, with 16%, 16.6% and 17% polarisation respectively at the start of scanning, can be seen in figure 5.4. Due to the inhomogeneity of the signal in figure 5.4c, average signal values were measured at 2 locations (red and yellow crosses in images), using 25 voxels, for the purpose of SNR calculation. The standard deviation of the noise was calculated using 25 voxels centred around the green cross. SNR values can be found in table 5.3.

The SNR increased by a factor of > 2 between the attenuation 14dB image and 8dB image. The polarisation was higher for the 8dB image, but only by a factor of 1.04. This suggests 8dB is a more optimal than 14dB for this sequence with this loading strategy.

An attenuation of 8dB resulted in higher SNR in both locations than the attenuation 4dB image. The fact the 8dB image started with less polarisation and resulted in less loss in polarisation (see table 5.2), suggests the attenuation 4dB image is acquired with too high a FA; the polarisation is depleted by more than would be optimal before the acquisition is completed. It should be noted that the volunteer was removed and put back in the coil between the attenu-

Attenuation/ dB	Polarisation before imaging/ %	SNR Red Cross	SNR Yellow Cross
14	16	27.9	30.6
8	16.6	56.3	62.1
4	17	29.4	50.5

Table 5.3: SNR values of figure 5.4 images.

ation 8dB and 4dB images. This may have affected the coil response between acquisitions. Laying differently in the coil (with arms tight versus loose) may also affect the loading. However, if the patient lying differently in the scanner affects the coil so significantly, this is a major issue. It will be difficult to design the optimum parameters for a sequence if the optimum parameters shift significantly depending on how the volunteer loads the coil.

The coil was connected to an Agilent Technologies network analyser and the transmit coil response was recorded for a variety of coil loading strategies. The amplitude and frequency of the maximum (peak) response was recorded, as well as the response of the coil at the resonant frequency of ^{129}Xe inside the scanner (17.663MHz) (see table 5.4). Table shows that different loading strategies produce very different coil responses and consequently FAs for the same amplitude RF pulse. Loading not only affected the amplitude, but also the centre of the peak. Having arms tight to the coil makes the coil more cylindrical, this improved the coil response. Having a bag on the chest also reduced the coil response. This is probably because putting a bag on the volunteer's chest introduces some empty space between the coil and the volunteer and therefore the coil's filling factor is less. Heavier volunteers also loaded the coil better, but two similar weight volunteers did not load the coil by the similar amount; factors such as body composition may come into play.

The coil loading was also investigated inside the scanner to see if loading is more consistent at 1.5T and to observe the frequency shift of the coil inside the scanner's magnetic field (see table 5.5). The scanner bore restricts the volunteer so they are forced to have arms tight against the coil when lying inside the scanner.

Inside the scanner, the amplitude of the peak response again varied by a large amount; depending on how the coil was loaded. The amplitude of the peak response was slightly lower inside the scanner than outside, but the amplitude of response at the Xe frequency was higher inside the scanner. Loading the coil with saline bags provided a worse response than any of the volunteers. This goes some way to explaining the data from table 5.2 and 5.1, where for a given attenuation, the FA estimations were higher when V3 loaded the coil than the saline bags. It is worth noting that this effect may be amplified by another volunteer; V2 is the same weight as V3 but loads the coil even better.

Signal Variation Within the Lungs A new healthy volunteer was placed in the scanner and four scans were ran at an attenuation of 12, 8, 4 and 0 dB

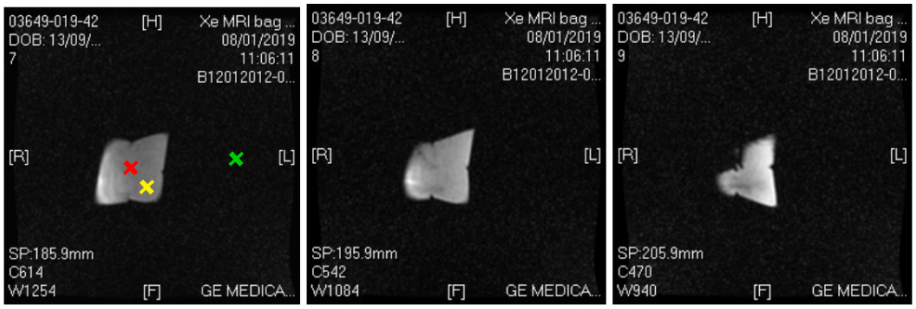

(a)

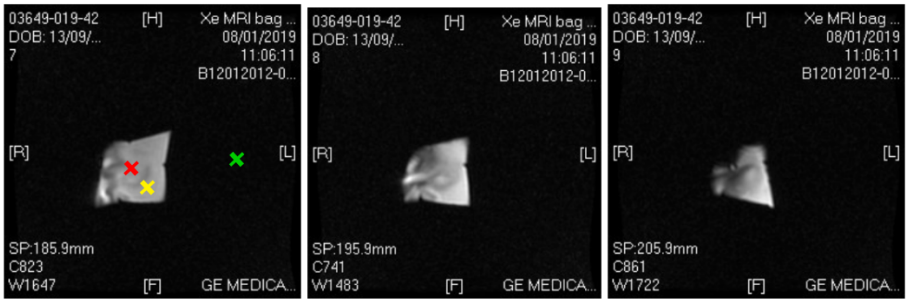

(b)

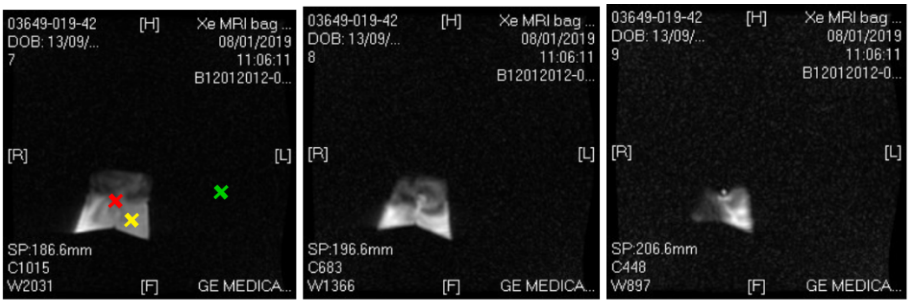

(c)

Figure 5.4: 2D Coronal FGRE 48x48cm FOV 10mm slice. Polarised Xe bag on volunteer chest. **(a)**RF attenuation 14dB Polarisation 16%, **(b)**RF attenuation 8dB Polarisation 16.6% and **(c)**RF attenuation 4dB Polarisation 17%.

Loading Strategy	Peak Amplitude/ dB	Peak Frequency/ MHz	17.Amplitude at 17.663MHz/ dB
Saline Bags	11.5	17.4227	3.4
V1 arms loose	19.7	17.46	5.2
V1 arms tight	21	17.505	6.4
V1 arms loose bag on chest	16.4	17.471	5.0
V1 arms tight bag on chest	17.7	17.501	6.3
V2 arms loose	20.9	17.485	5.9
V2 arms tight	21.9	17.543	8.56
V2 arms loose bag on chest	17.5	17.511	6.3
V2 arms tight bag on chest	18.6	17.541	8.1
V3 arms loose	17.4	17.467	4.9
V3 arms tight	19.0	17.487	5.9
V3 arms loose bag on chest	15.2	17.491	5.5
V3 arms tight bag on chest	17	17.517	6.6
V4 arms loose	20.2	17.513	6.9
V4 arms tight	21	17.586	11.8
V4 arms loose bag on chest	18.5	17.497	6.0
V4 arms tight bag on chest	21.6	17.555	9.5

Table 5.4: Coil response outside the scanner (black quadrature lead), for a variety of loading strategies. V1 is volunteer 1 (∼65kg), V2 approx (∼70kg), V3 (∼70kg), V4 (∼85kg).

Loading Strategy	Peak Amplitude/ dB	Peak Frequency/ MHz	Amplitude at 17.663MHz/ dB
Saline Bags	10.3	17.603	9.1
V2 Arms tight bag on chest	15.5	17.640	15.1
V3 Arms tight bag on chest	14.3	17.755	10.9

Table 5.5: Coil response inside the scanner (black quadrature lead), for a variety of loading strategies. V2 and V3 both ∼70kg.

(see figure 5.5). SNR was calculated for the second slice inside the lungs of each image set. The SNR was 2.92, 6.60, 10.7 and 21.7 for the 12, 8, 4 and 0 dB image set respectively. The higher SNR at the lower attenuations suggesting the FA inside the lungs is much lower than it was in the bag placed on top of a healthy volunteer.

In figure 5.5 it was noted that the brightest slices were at the anterior side of the lungs and by later slices (still inside the lungs), no signal was observed. The posterior slices were imaged later, so this could either be because the posterior slices have a very different FA to the anterior slices, the receive arrays are less able to pick up signal from the posterior slices, or the ^{129}Xe in the posterior slices has depolarised fast enough that by the time they are imaged there is insufficient signal to image.

A new healthy volunteer was imaged at 0dB attenuation. Slices were acquired anterior to posterior (see figure 5.6). The volunteer was imaged with the jacket in different configurations. First, as normal (figure 5.6a), next with the receive pads switched (figure 5.6b) and finally with the whole coil on backwards; both outer jacket front and front receive pad on volunteers back (figure 5.6c). Neither changed that anterior slices were always brightest and the SNR reduced in slices towards the back. In the first few anterior slices, some fine structure can be observed.

Axial images were next acquired, to see if the anterior of the lungs was still brighter than the posterior. The attenuation was set to 0dB and the slices were first acquired starting at the base of the lungs moving towards the apex and then at the apex of the lungs moving towards the base (see figure 5.7). The SNR was less in later acquired slices in both sets of images. This suggests over 20s the reduced T_1 of ^{129}Xe in the lungs compared to a bag is having a noticeable effect on the images. The highest signal was still obtained at the anterior of the lungs with very little signal towards the posterior. (see figure 5.7).

Performance of Other Clinical MR Solutions Coils Other teams in Oxford and Duke also perform hyperpolarised ^{129}Xe imaging using a Clinical MR Solutions 1.5T ^{129}Xe coil. However, both of their coils don't have the additional receive arrays. Instead their coils operate exclusively in QUAD T/R mode.

While some fine structure can be observed in some anterior slices of figure 5.6, the images are of much worse quality than the gas phase images produced at Duke by Cleveland Et. Al. [138] and Driehuys Et. Al. [139]. In Oxford Chen Et. Al. [140] have enough signal to image ^{129}Xe being breathed into the lungs dynamically.

It was postulated that the additional receive arrays are interfering with the outer jacket transmit coil.

Summary The coil global transmit response varies significantly depending on the loading. The signal from polarised Xenon breathed in is very inconsistent across the volume of the lungs and insufficient in QTAR mode to image the

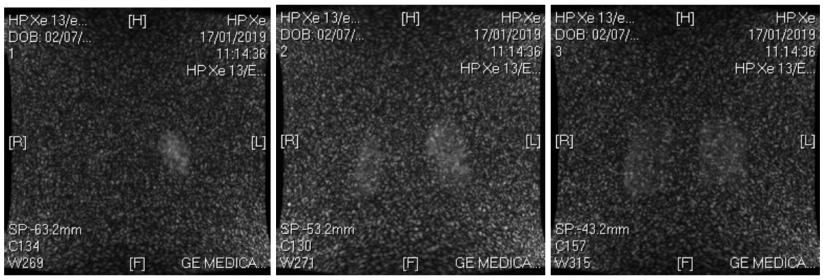

(a)

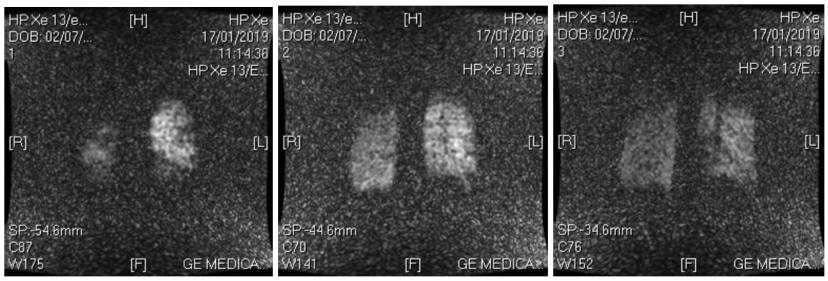

(b)

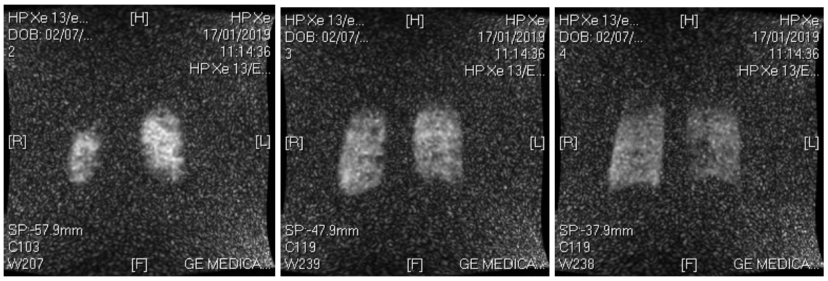

(c)

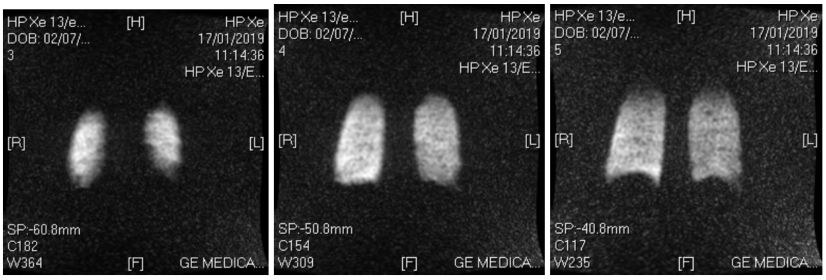

(d)

Figure 5.5: 2D FGRE 48x48cm FOV 10mm slice. Healthy volunteer breath hold at 12dB **(a)**, 8dB **(b)**, 4dB **(c)** and 0dB **(d)** attenuation. The first 3 slices, from anterior to posterior, inside the lungs are shown.

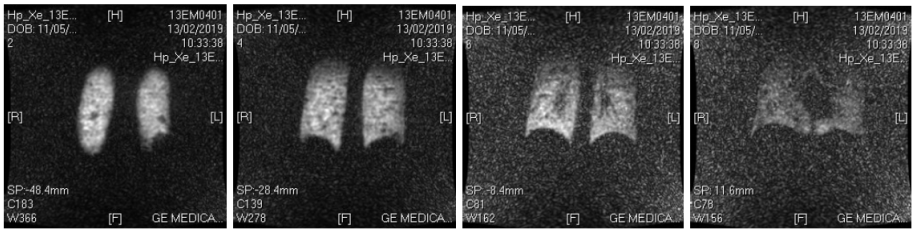

(a)

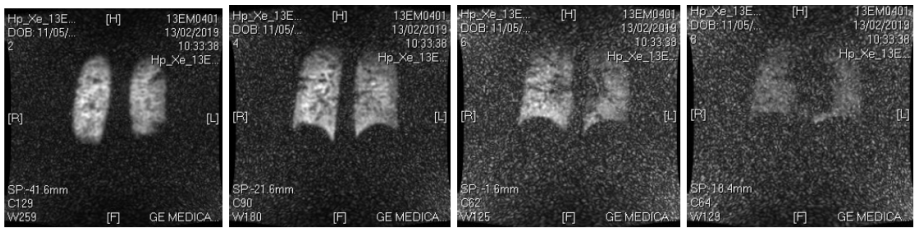

(b)

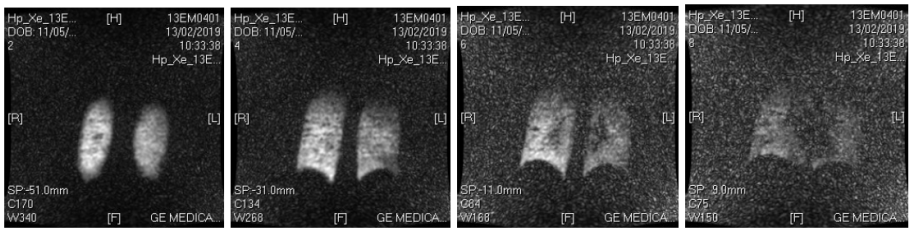

(c)

Figure 5.6: 2D FGRE 48x48cm FOV 10mm slice. Healthy volunteer breath hold at: **(a)** attenuation 0dB, slice 2 SNR 27.4, slice 8 SNR 5.86. **(b)** 0dB receive arrays switched and slice 2 SNR 22.8, slice 8 SNR 4.64. **(c)** 0dB receive and transmit switched slice 2 SNR 25.1, slice 8 SNR 4.80. Slices from anterior to posterior; 2, 4, 6 and 8 are shown.

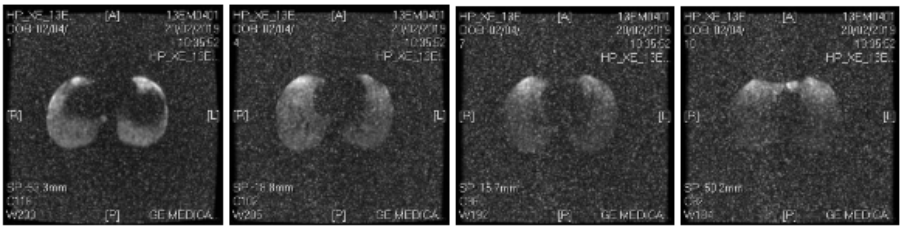

(a)

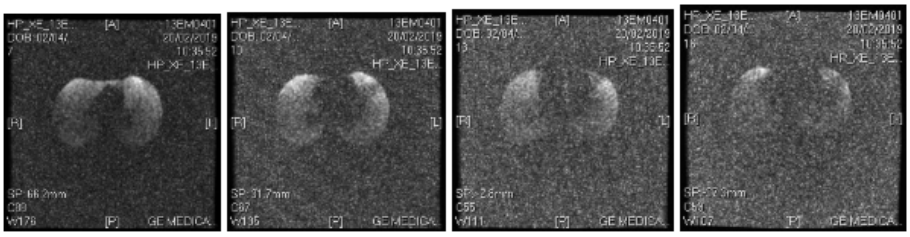

(b)

Figure 5.7: 2D FGRE 48x48cm FOV 10mm slice axial. 1st, 4th, 7th and 10th slices acquired shown (from left to right). Healthy volunteer breath hold at attenuation 0dB. **(a)** Base to apex; slice 1 SNR 5.54, slice 10 SNR 2.81. **(b)** Apex to base; slice 1 SNR 5.10, slice 10 SNR 3.01.

whole lung's outline, let alone resolve fine detail throughout the lung. There is the possibility the two receive arrays are interfering with the outer jacket transmit coil.

5.3 QUAD T/R mode

In this mode the outer jacket is used to both transmit the RF and receive signal from the hyperpolarised ^{129}Xe. However, running in this mode does require some diodes to either be shorted out, or to be biassed by an external power supply.

5.3.1 Optimum Power Supply

With an A.C. to D.C. power supply, specific diodes were biased in order to run in QUAD T/R made. However, at first no images could be produced. After switching to a lower noise power supply, an image could be produced with the coil. The optimum power supply was investigated to see if the images could be improved further.

A low noise A.C. to D.C. power supply and an alkali motorcycle battery were compared to see which produced the best images when biasing the diodes. A very low tipping angle was used to image a polarised Xe bag and the coil was loaded with saline. The same 10mm slice FGRE sequence, as previously used in section 5.2 with an RF attenuation of 40dB was found to reduce the polarisation of the bag from 16.5% to 15% polarisation. The bag was imaged; first using the DC power supply, before using the motorcycle battery. The highest SNR slice was slice 5 for both images. The SNR when using the power supply was 1.95 and with the motorcycle battery was 2.51.

Next, the motorcycle battery was compared to shorting out the diodes. The motorcycle battery image was acquired first and produced an SNR of 2.07. The subsequent shorted out diode image was better, with an SNR of 2.40. Shorting out the diodes produced the highest SNR images. This is also the most stable long term solution and doesn't require a motorcycle battery to be charged before scanning.

5.3.2 Phantom Investigations

A Tupperware box was used as a phantom. The box had a volume of 1.4L and long side of internal length 195mm. The box was evacuated and purged with N_2 gas to remove as much O_2 as possible, so when polarised ^{129}Xe was pushed into the phantom from a bag, the effects of O_2 on the T_1 of polarised ^{129}Xe were minimised. 195mm is longer than the maximum distance from the sternum to the back for most volunteers. The longest side of the phantom was aligned with this orientation and imaged coronally (see figure 5.8). Theoretically, when using 10mm slices there should be signal in 20 or 21 slices. Figure 5.8 shows slices 3 (first slice with signal), 4, 13, 21, 22 and 23 (last slice with signal). If the first and last slice (where part of the slice may have been outside the phantom) are discounted, the signal was seen to be broadly consistent across the full phantom. Good fine detail could also be observed. On slice 13 it is possible to see Xe in the tube coming out of the Tupperware. This tube has an internal diameter of 6mm.

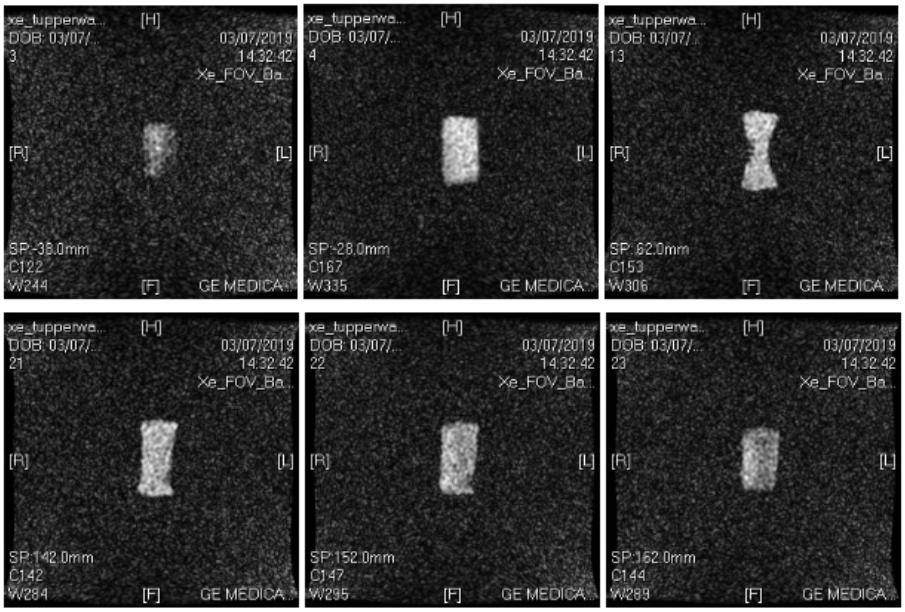

Figure 5.8: 10mm slice coronal 2D FGRE, attenuation 8dB. From left to right, Row 1: Slices 3, 4, 13 and Row 2 slices 21, 22 and 23.

A 0.6L bag was next imaged using a variety of scan parameters (see table 5.6). The estimated FA was similar with scans run at the same attenuation. To allow for comparison between different scans, the SNR was divided by the average Xe polarisation during the scan. Taking scan 1 as the base scan this scan took 17s. The frequency encoding points can be doubled by doubling the acquisition window. This only increases the scan time to 19s. Doubling the phase encoding points, or halving the slice thickness and maintaining the same FOV would double the scan time. Doubling the points in any dimension halves the image voxel size and therefore there is less polarised Xe in that voxel to provide a signal. This can be seen when comparing scan 1 to 2 or 1 to 3. Reducing the slice thickness would be beneficial if there is enough signal, since it would allow for finer detail to be resolved within the lungs. However, it does result in a scan with the same FOV taking more time. To reduce the time of the scan, the FOV in the phase encoding dimension can be reduced and resolution maintained.

Scan No.	Orientation	Matrix Size (Freq. enc. x Phase enc.)	Slice Thickness/ mm	RF Attenuation/ dB	Starting Polarisation/ %	Ending Polarisation/ %	Approx FA/ degrees	Highest SNR Slice	$\frac{\text{SNR}}{<\text{Polarisation}>}$
1	Coronal	128x128	10	20	17.6	14.68	3.1	11.3	0.65
2	Coronal	256x128	10	20	14.68	12.4	2.9	6.9	0.47
3	Coronal	128x128	5	20	12.4	10.67	2.8	4.6	0.36
4	Coronal	78x128	5	20	10.67	9.46	3.2	3.0	0.28
5	Axial	128x128	10	10	9.46	4.09	6.6	15.5	1.6
6	Axial	128x128	10	10	3.6	1.4	7.0	7.0	1.9

Table 5.6: 2D FGRE scans 1 to 6 performed on 0.6L Tedlar bag of polarised Xe.

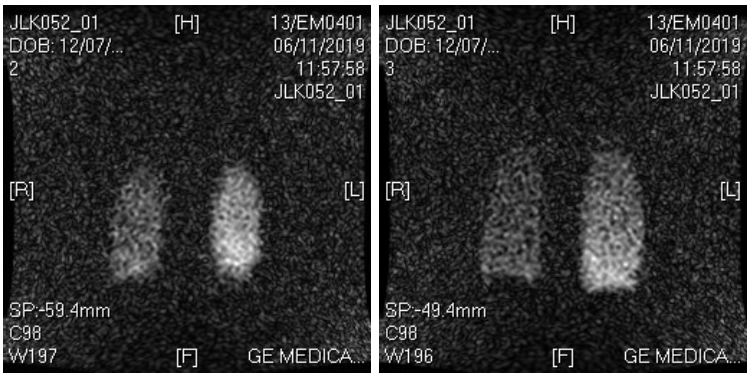

(a) (b)

Figure 5.9: 2D FGRE 48x48cm FOV 10mm slice, coronal, 10dB attenuation, slice order posterior to anterior **(a)** Slice 2 (SNR=10.5) and **(b)** Slice 3 (SNR=7.9).

5.3.3 Healthy Volunteer Investigations

Using a 2D, coronal, 10mm slice, FGRE sequence a healthy volunteer's (V6's) lungs were imaged using hyperpolarised Xe. Given the higher signal in a phantom at 10dB RF attenuation than 20dB (see table 5.6), a 10dB attenuation was decided upon as a starting point. However, it is known (see section 5.2) that saline bags and a volunteer load the transmit coil differently. Therefore, the optimal FA may vary between a volunteer and saline bags loading the coil.

10mm coronal slices were acquired from bottom to top (see figure 5.9). Signal died off after the first few slices; by the 10th slice out of 13, the outline of the lungs couldn't be distinguished. Next, the same sequence was run, but this time the slices were acquired from anterior to posterior and the attenuation was reduced to 0dB. For the first time signal was observed in all slices. The signal was reduced in the central slices (see figure 5.10).

The darker central slices could be explained by proximity effects. The FA may be significantly greater in regions of the lung closer to the jacket and the coil receive may be more sensitive in these regions also. The volunteer scanned also had a slight build. When the volunteer goes inside the scanner, their arms are pushed against the sides of the coil. If the volunteer doesn't fill the volume of the coil, the front of the jacket "tents up" (see figure 4.18); this is likely to significantly effect the loading. When the coil tents up, the volunteer's sternum is further away from the front of the coil than their back is from the back of the coil. This could explain why the signal was high in the posterior half of the lungs in both scans. By increasing the transmit gain, more signal may be obtained from the central portion of the lungs as a higher FA is achieved. However, the inconsistency in the FA, especially close to the coil, may result in too high a FA in outer regions of the lung. The signal was also inconsistent between different lungs; figure 5.10c has greater signal in the left lung than right.

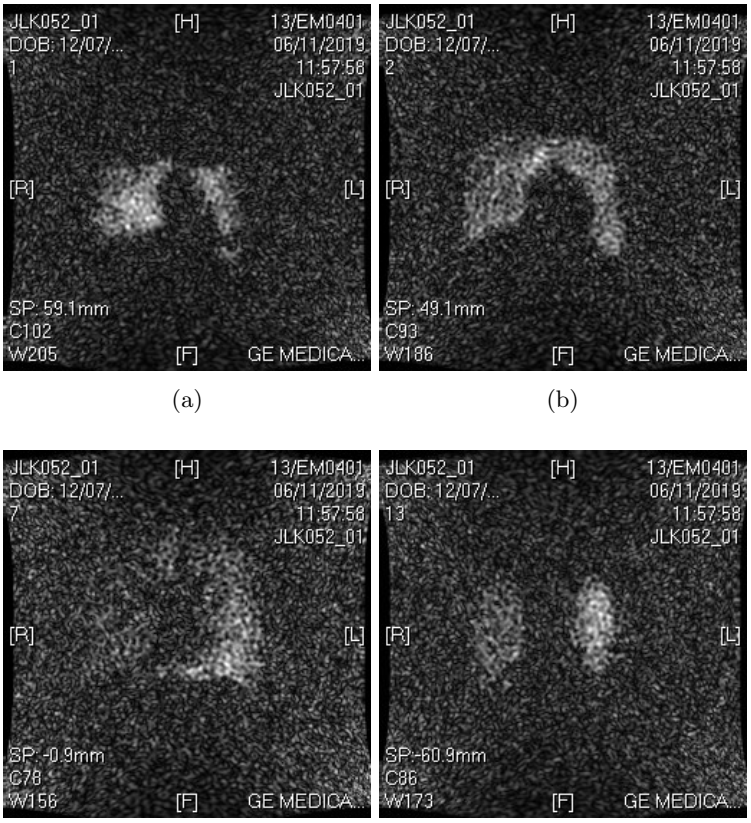

Figure 5.10: 2D FGRE 48x48cm FOV 10mm slice, coronal, 0dB attenuation, slice order anterior to posterior **(a)** Slice 1 (SNR=6.7), **(b)** Slice 2 (SNR=6.8), **(c)** Slice 7 (SNR=4.8) and **(d)** Slice 13 (SNR=7.1).

Volunteer	Coil Configuration	Response Amplitude at 17.67MHz/ dB
V1	Pads in	16.0
V1	Pads out arms loose	17.3
V1	Pads out arms tight	18.6
V1	Pads out leaning right	24.5
V1	Pads out leaning left	28.6
V3	Pads in	16.1
V3	Pads out arms loose	16.6
V3	Pads out arms tight	17.6
V3	Pads out leaning right	17.3
V3	Pads out leaning left	17.8

Table 5.7: Coil response in QUAD T/R mode (black quadrature lead), inside the scanner, for two volunteers in different orientations.

By putting foam pads inside the coil, both behind and on top of the volunteer, the coil is forced to exhibit a more rounded shape (see figure 4.18). The loading should also be forced into being more consistent between scans as there is less freedom of movement for the coil. The pads also add a spacing between the volunteer and the coil. This should attenuate proximity effects and ensures the volunteer's ribcage is at a consistent distance from the jacket.

With this configuration, an identical 0dB attenuation scan was attempted. However, the scanner detected an error and the scan failed to run. The error displayed was "Peak power exceeded on MNS broadband amplifier. Peak power 26W." The scan was reattempted with saline bags and the scan ran as normal. However, when the volunteer then went back inside the coil the scan failed again. It is believed that "Peak power 26W" refers to the maximum average forward power, rather than the maximum instantaneous forward power.

Coil Loading QUAD T/R Mode

A new volunteer was recruited (V1 from QTAR investigations). No polarised Xe was produced because it was believed the scan would fail to run. The volunteer was a similar height and weight to the volunteer on which the scan failed. The same parameters were also used as were when the scan failed (0dB RF attenuation). However, the scan ran as normal with no pads in (the coil tenting up) and pads in (coil exhibiting a more rounded profile). With the volunteer inside the magnet the coil was connected to a network analyser. The investigations were also performed using a volunteer with a larger build (V3 from QTAR investigations). The results can be seen in table 5.7.

The response with V3 inside the coil altered much less than V1 no matter the coil configuration. When V1 was inside the coil with no pads, the coil response exhibited two peaks. This was accentuated by the volunteer leaning to one side (see figure 5.11).

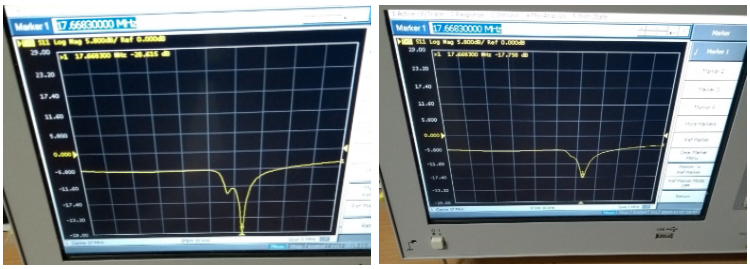

(a) (b)

Figure 5.11: Coil loading in QUAD T/R mode inside scanner. **(a)** V1 leaning left pads out, **(b)** V3 leaning left pads out.

Scan No.	Attenuation/ dB	Breathing pattern	Scan Result
1	0	Breathing Normally	✓
2	0	Holding End Inspiration	✗
3	0	Breathing Normally	✗
4	10	Holding End Inspiration	✓
5	10	Breathing Normally	✓
6	0	Breathing Normally	✗

Table 5.8: Six Scans run at 0dB and 10dB attenuation; ✓scan ran normally, ✗scan failed to run.

Diagnosing the Error

After the successful run with a volunteer of a similar weight to V5, the Xe scan was run on V5; no Xe was produced. The results can be seen in table 5.8. All the failures to run happened at 0dB attenuation and had the same error as previously: "Peak power exceeded on MNS broadband amplifier. Peak power 26W." The error appears to be intermittent; the first scan at 0dB ran as normal.

With a purely resistive load connected, instead of the coil, the sequence ran reliably at 0dB of attenuation. However, at 1dB of transmit gain (TG), the sequence failed. Similarly with a GE multinuclear spectroscopy (MNS) bridge connected and a purely resistive load the same error was recorded, but this time the sequence failed to run at 1dB TG, but not at 0dB TG.

Summary

The coil appears to offer a more consistent signal across the full volume of the lungs in QUAD T/R mode than QTAR mode. However, while the scans ran normally at low RF powers, at higher RF powers the scans intermittently failed. The optimum FA is believed to be at (or above) these RF powers; so in its current state the scanner coil arrangement isn't fit for purpose. It would be

unacceptable for a scan to fail with a patient rather than a healthy volunteer in the scanner. The fact the failure happened at a similar attenuation with a purely resistive load (with both the GE MNS bridge and Clinical MR solutions L.L.C. duplexer connected), points to the fault with the scanner's MNS train.

5.4 Conclusion

In QTAR mode the coil was unable to provide signal throughout the whole lungs. This makes it unsuitable for Xe imaging. Without a consistent response across the whole lungs from the coil, it will be difficult to distinguish between a void in the lungs and a section of the lungs where the signal is less.

Other teams are able to produce high quality ^{129}Xe images using a Clinical MR Solutions coil at 1.5T [139] [138] [140]. However, their coils don't have the additional receive arrays and are instead exclusively ran in Quad T/R mode.

Initial results from QUAD T/R mode were promising; a more consistent signal was seen throughout the lungs than in QTAR mode. However, it is believed the MNS train of the scanner is faulty; resulting in the scanner failing to run at higher RF powers. The optimum RF power for imaging is believed to be at or above these higher RF powers, so in its current arrangement the scanner is unfit for MNS imaging.

GE have conducted tests, which show issues with the MNS RF amplifier, in addition to high levels of noise. This may have been the reason for the scans failing to run. The MNS RF amplifier is currently undergoing repair, after which the SAR safety hardware will need to re-calibrated.

Chapter 6

Upright 0.5T Proton Imaging

In this chapter, proton work on the Paramed 0.5T Upright scanner will be discussed. This is broken down into three sections:

- In section 6.1, a 2D-GRE sequence is detailed. This high temporal resolution sequence was created for the purpose of lung density monitoring.
- In section 6.2, the addition of a hold period at the start of a SE sequence is described. This sequence is designed to image the diaphragm.
- In section 6.3, work performed using compressed sensing is covered.

6.1 2D-GRE Single Slice Lung Density Monitoring

As someone breathes in, the volume of the lungs increases and the lung density drops. Globally, the lungs have a density fluctuation, which is periodic with the same frequency as the breathing rate. Locally, if the lungs are thought of as lots of sacks of air, as these sacks increase in volume, the density also decreases. Many lung conditions are characterised by heterogeneity of ventilation within the lung. If bronchi are obstructed, regions of the lung will not inflate in the same way. This results in a more fixed density in the impaired region. However, the perfusion of blood through the lung parenchyma will also cause fluctuations in the density of lung tissue.

Current options for non MR real time in vivo imaging, include ultrasound and fluoroscopy (continuous X-ray image). Both allow for the acquisition of high time resolution data. However, ultrasound is very operator dependent and fluoroscopy is both ionising and provides an infinite slice through the body.

Using a 1.5T MRI scanner Bauman Et. Al. [141] have described a method, whereby a 2D GRE sequence with high time resolution is used along with Fourier

decomposition, to separate the signal change due to perfusion and change due to ventilation. Fortunately, the fluctuations due to ventilation and perfusion occur at different frequencies. By Fourier transforming across a time course of images, the perfusion and ventilation components can be extracted. Bauman Et. Al. [141] were able to show clinically relevant information and highlight known ventilation and perfusion defects in humans. Deimling Et. Al. [142] have also used a similar method of Fourier decomposition at 0.35T to produce ventilation images.

The majority of high time resolution MRI scans are currently performed with a horizontal body position. For chest imaging this is far from ideal, because lung function, volume, and arterial oxygen levels are affected by body position [6]. By lying supine lung function is reduced and for patients at the latter stages of chronic conditions, such as COPD, many simply can't lay in a supine position for an extended period of time. Furthermore, it seems logical to extend studies to an upright orientation, as the lung mechanics and function are different, yet represent a majority of time in the woken state.

Developing passive techniques to provide information about lung function in paediatric cohorts would be extremely beneficial. Spirometry is the current gold standard for measuring lung function. However, the results can vary depending on the effort level of the patient. Getting children to perform it properly can be difficult, therefore a passive alternative would be very useful.

It was postulated that, using Fourier decomposition and the Paramed 0.5T, regional lung density changes due to ventilation could be observed. It was also hoped that the 2D GRE sequence required could be used to accurately characterise diaphragm movement throughout the breathing cycle. Respiratory diseases can result in a deformed lung architecture. Understanding how the geometry of the lung changes throughout the respiratory cycle would help in the understanding of these conditions.

6.1.1 Method

A 2D GRE sequence with a single slice was developed in order to acquire images of the lungs and diaphragm. The sequence had a TE = 4.6505ms (from temporal centre of RF, 3.105ms from the midpoint of the effect of the RF pulse on magnetisation) and TR =12.8965ms . The sequence was designed to acquire a scan in the shortest duration possible. An RF pulse taking 3ms was used (see figure 6.1).

A 50ms gap was also inserted between images to avoid the build-up of transverse coherences, which can arise from imperfect spoiling.

Sequence Weighting

Ideally a short TE and long TR sequence with a low FA would be used for proton density imaging. A short TE minimises the T2* weighting, long TR minimises the T1 weighing and low FA maximises the proton density weighting (see section 2.2.1).

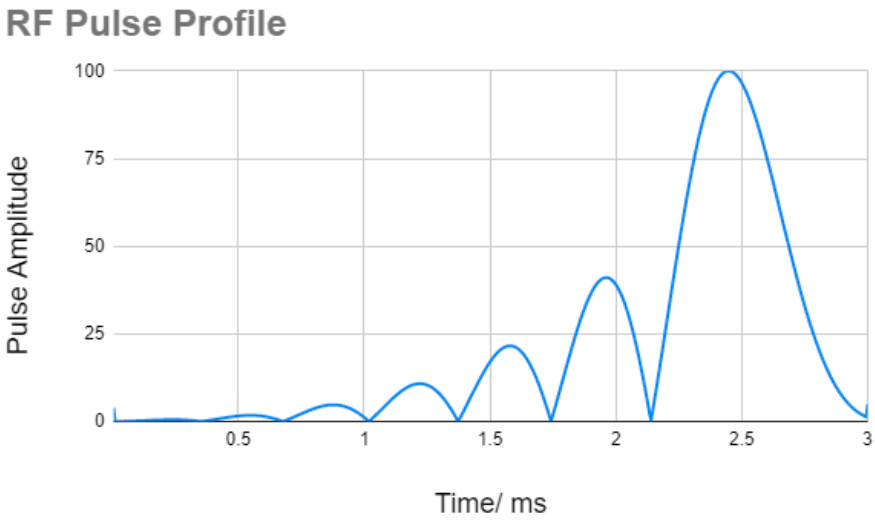

Figure 6.1: RF pulse amplitude profile used in the Paramed GRE sequences.

The TE of the sequence was minimised, by making the phase and frequency encoding gradients ramp up, before immediately ramping down. In order to ensure the maximum slew rate isn't exceeded, the time for gradients to ramp up or ramp down is set at 0.6ms. Within 0.6ms the gradients can ramp from zero to maximum. Setting the minimum ramp duration to 0.6ms ensures that if the FOV of the sequence is changed in future, the maximum slew rate won't be exceeded in error. The theoretical minimum TE, keeping to this 0.6ms ramp duration, is 1.8ms; as the RF duration$\rightarrow 0s$ and acquisition duration$\rightarrow 0s$ (See figure 6.2). However, both of the RF and acquisition must take time. The actual TE chosen was 4.6505ms. The TE could be reduced further if the gradient durations was reduced to less than 0.6ms.

Due to the high time resolution required, the TR chosen was not optimal for producing proton density weighted images. An image needed to be generated within <1s in order to achieve sufficient time resolution to resolve the breathing cycle frequency. Recalling the Nyquist criterion, that for a signal which repeats over time T, the sampling rate ($1/\Delta T$) must be at least **twice** as great. Younger children tend to have a higher respiratory rate than adults, with children of different sexes breathing at a broadly similar rate. Iliff et al. [143] found that of the children studied between the ages of 0-1, the average respiratory rate was 31 ± 1.3 and 30 ± 0.8 breathing cycles per minute, for boys and girls respectively.

Even for a single slice, in order to achieve 64 phase encoding steps, the maximum TR while still acquiring an image in 1s is $\frac{1s}{64} = 15.6ms$. Taking the T1 in the lung to be $\sim 1008ms$ [144], $\alpha_E = 10.05°$ (equation 2.24).

The actual TR = 12.8965ms was chosen to be slightly shorter (because time

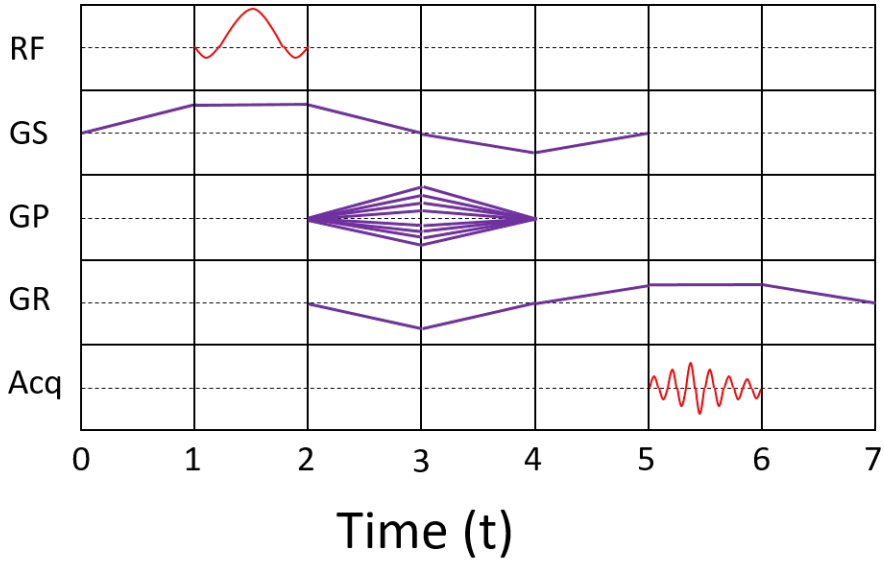

Figure 6.2: 2D GRE Pulse sequence. Event durations not necessarily equal. Minimum TE (while keeping to a minimum ramp duration of 0.6ms), achieved as the duration of event 1-2 $\rightarrow$ 0s and 5-6 $\rightarrow$ 0s and the duration of event 2-3, 3-4 and 4-5 = 0.6ms.

resolution was so imperative). For this $\alpha_E = 9.15°$. A FA of $10° \sim \alpha_E$ was used for imaging.

Slice Thickness

The RF pulse's $t_0 \approx 0.4ms$ for a 3ms pulse where t_0 is the effective duration of one lobe. Therefore, $\Delta f \approx 2.5kHz$ (see equation 2.10). Equation 2.12 can be rearranged to give equation 6.1.

$$G_z = \frac{2\pi\Delta f}{\gamma\Delta z} \tag{6.1}$$

The $\hat{z}$ gradient strength (G_z) for a given slice thickness (Δz) can then be calculated by taking γ as the proton gyromagnetic ratio.

For gradient amplitudes, TNMR uses a scale from 0 to 100; where 0 corresponds to 0mT/m and 100 corresponds to the maximum gradient, 20mT/m. For a 10mm slice thickness, G_z was calculated as 5.87mT/m, therefore the slice thickness gradient strength required is 29.4 in TNMR.

The transverse phase dispersion (see section 2.1.8) must now be accounted for. $\Delta t_I = 0.55ms$ for the 3ms pulse in figure 6.1. Therefore, the required area of the gradient lobe for slice rephasing; $A_R = -5.87(0.55 + \frac{0.6ms}{2}) = -4.99mTsm^{-1}$. Taking the slice rephasing gradient to ramp up and ramp down, both in 0.6ms, the slice rephasing gradient amplitude was calculated to be $-8.31mT^{-1}$, or -41.6 in TNMR. This requires a consistent slew rate of $8.32Tm^{-1}s^{-1}$. The sequence was tested with a variety of slice rephasing amplitudes in order to obtain the maximum signal. The maximum signal was obtained using a slice rephasing gradient of -40.5. This is believed to be due to the imperfect ramp up and ramp down of the gradients, which will not be at a perfectly consistent slew rate. The value of -40.5 was used in the actual sequence. A sequence with three times the slice thickness was also created. In order to do this, both the slice thickness and slice rephasing gradient should be divided by the same factor, 3.

A flat side of a phantom was moved from outside the FOV and slowly moved towards the excited slice until the SNR from the phantom was 2, or a signal from the phantom could be visually observed (whichever was first). The same was repeated from the opposite side. The sample was moved in 1mm increments. The phantom was first observed at two points 10mm apart putting the slice thickness at $10.5mm \pm 0.5mm$. Similarly for the three times thicker slice the phantom was first observed at two points 31mm apart putting the thickness at $31.5mm \pm 0.5mm$.

In Vivo Imaging

For in vivo imaging all volunteers gave written informed consent under University SPMIC ethics. Those under the age of 18 years were scanner under ethics number 8-1804 and those over 18 years were scanned under ethics number N14112016.

Scans were performed at approximately the Ernst angle for lung tissue and imaging was performed for 5 minutes. The scans were ran using either 64 phase encodings (image every 875ms) or 32 phase encodings (image every 463ms).

Scans on paediatric volunteers were performed with them sitting on a parent's lap. This meant they stayed approximately central in the magnet. Volunteers as young as 3 years old have been inside the scanner without need for sedation.

As the volunteer breathes in their lungs expand, causing a reduction in lungs proton density and therefore they produce a lower signal. However, this also results in alveoli shifting between voxels. For lung density variation investigations ATNsR v1.0 software [145] was used to map each image onto the first image, so that the shape of the lungs was the same throughout the time course. This should allow for the assumption that each alveoli stays in the same voxel.

In ANTsR the type of transform used for registration was "SyNAggro" [145]. This is an affine and deformable transformation where mutual information between images is used as an optimisation metric. "SyNAggro" is similar to "SyN", but "SyNAggro" is a more aggressive registration and consequently takes more time.

After registration, the lungs can then be segmented out and an FT performed over the time domain. The highest peak of this should correspond to the average breathing rate throughout the scan.

6.1.2 Results

Figure 6.3a shows one frame from a 32 phase encoding time course (1cm thick slice, 463ms temporal resolution) for volunteer 1 and figure 6.3c shows one frame from the 64 phase encoding time course (3cm thick slice, 875ms temporal resolution) for volunteer 2. For both, the outline of the diaphragm could clearly be seen to rise and fall with breathing.

For the 32 phase encoding time course for volunteer 1, the volunteer was asked to breathe every 4.6s. Nine voxels were chosen centred around the x in figure 6.3a. The mean amplitude of the data in these voxels was obtained for each image and an FT was applied over the time course. The resulting data can be seen in figure 6.3b. The peak at 0.22Hz (every 4.5s) is attributable to the ventilation rate.

For the 64 phase encoding time course for volunteer 2, the volunteer was asked to breathe every 4.4s. Nine voxels were chosen centred around the x in figure 6.3c. The mean amplitude of the data in these voxels was obtained for each image and an FT was applied over the time course. The resulting data can be seen in figure 6.3d. The peak at 0.23Hz (every 4.3s) is attributable to the ventilation rate.

6.1.3 Discussion

The slice thickness of both sequences shows good agreement between the theory and in practice. In real time videos of both 32 and 64 phase encodings it is

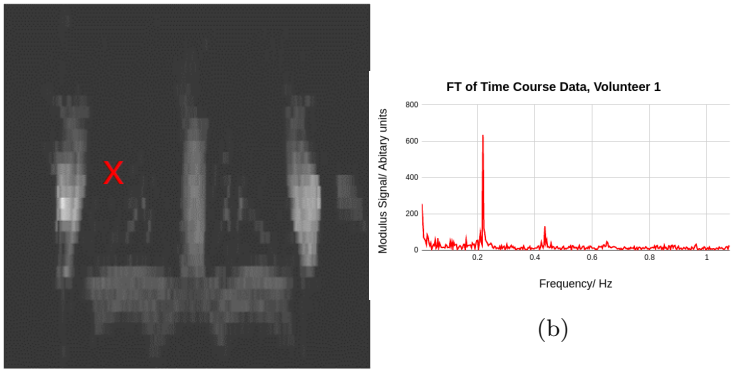

(a)

(b)

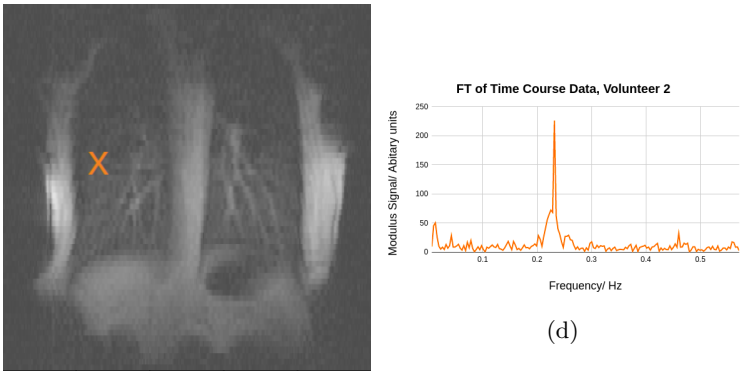

(c)

(d)

Figure 6.3: **(a)** Volunteer 1 32 phase encodings 1cm thick slice centred behind heart, time resolution 463ms. Yellow x marking central location of nine voxels used to produce **(b)**. **(b)** FT of volunteer 1 32 phase encodings time course data. **(c)** Volunteer 2 64 phase encodings 3cm thick slice centred behind heart. Image every 875.6ms. Orange x marking central location of nine voxels used to produce **(d)**. **(d)** FT of volunteer 2 64 phase encodings time course data.

easy to distinguish the diaphragm movement. It would be interesting to image a volunteer with a diaphragmatic abnormality and see if this could be observed in the video.

Over a nine voxel region it is possible to extract the average breathing rate of both the 64 and 32 phase encoding data. Fourier decomposition of the breathing rate is most effective if the breathing is very regularly. Any irregularities in the breathing rate will result in a less pronounced wider peak. Instead a Self-gated Non-Contrast-Enhanced Functional Lung imaging (SENCEFUL) acquisition may be considered [146] [147]. In this a similar sequence is used, but the phase encoding steps are acquired quasirandomly. The direct current is also acquired after each phase encoding step which allows the data to be sorted into individual respiratory and cardiac phases. After registration, this data can be Fourier transformed along the time dimension to produce perfusion and ventilation maps.

While ANTsR should register the images so that cells in the lung parenchyma stay in the same voxel, this can only account for distortions in two dimensions. The slice excited by the RF pulse is fixed and if cells move in an anterior or posterior direction this wont be accounted for. Primarily the volume change within the lungs is due to diaphragm movement, however a central slice posterior of the heart was chosen to minimise the anterior-posterior movement.

As long as the phase encoding gradients are always performed in the same order, the time resolution could theoretically be increased through a sliding window reconstruction [148]. By replacing the first phase encoding line of the first image, with the first of the second image, an image (image a) between image 1 and image 2 can be generated. The next image will be generated by replacing the second line of the first image, with the second line of the second image. This image (image b) is slightly less like image 1 and more like image 2 than image a. Thus, by repeating this until the composite image becomes image 2, the time resolution can be increased. The total number of 2D slices generated would increase from T_{Points} to $(T_{Points} \cdot P_{Points}) - (P_{Points} + 1)$. Where P_{Points} is the number of different phase encoding values and T_{Points} is the number of times a 2D slice is acquired. A golden-angle radial acquisition could be manipulated in the same way [149]. However, in order to apply the sliding window reconstruction to Fourier decomposition, there can be no gap between the images; phase encodings should be equally spaced. This means unless the spoiling used is very effective, transverse coherences will build-up.

6.1.4 Summary

The diaphragm movement throughout the breathing cycle has been imaged at a time resolution of 463ms. The lung density average variation frequency was also extracted at the same as the average breathing frequency. The upright, open design of the Paramed scanner has allowed for volunteers as young as 3 to be comfortable inside the scanner without need for sedation.

6.2 Diaphragm Imaging: Spin Echo Hold Sequence

An investigation was performed into the shape of the diaphragm during a breath hold, on end inspiration and expiration, while upright and supine. Using the minimum TE/TR, 12ms/161.8ms for a 2D SE sequence, 7 slices and the minimum matrix size of 160x128, the duration of the stock 2D SE sequence is on average ∼35s. However, there is an unspecified preparation time before the scan starts; in which no data is being acquired. Because this preparation time varies between scans by up to 5s, the volunteer could be forced to hold their breath by over 40s. This is uncomfortable for most healthy volunteers and definitely too long for patients with respiratory issues. The acquisition period of the scan is only actually ∼21s. This is a much more manageable breath hold.

The TNMR base file was copied and modified. The Scope Trig was forced to trigger when the scan is ready to run; after the initial preparation time (see figure 6.4) and a test was introduced, which waits for an external trigger (4.8V pulse), before continuing on with the scan. This is edge triggered, requiring the voltage to go from 0 to 4.8V in order to proceed with the scan.

Branching tables were introduced at the start of the scan. This ensured the test event was only entered on the first scan of the sequence and on no others (see figure 6.4). These tables are 1000 entries long and assume there will never be more than 1000 1D (frequency encodings), 2D (slices), 3D (phase encodings) or 4D (repetitions) scans. These tables can be thought of as nested loops with 1D incremented first, followed by 2D, then 3D and finally 4D.

6.2.1 Running the Sequence

In order to run the sequence a function generator was connected to Ext Trig and a scope connected to the Scope_Trig. The scan was run like any of the stock Paramed sequences through the NRG clinical interface. By making the sequence run through NRG, rather than just TNMR, the sequence was much simpler to run for someone with only a basic knowledge of the sequence. When the scope trace rose, this illustrates the scan was ready to run. The scanner operator then pushed a button on the function generator to produce a 4.8V square pulse with a frequency <10Hz. The scan now began as soon as the first rising edge of the 4.8V pulse was received by the scanner.

It should be noted that the trigger pulse must be sent within 30s of the trace on the scope rising. If not, the hold will time out and the sequence will suspend. If this happens, the sequence must be deleted from the scanner queue. The sequence can then be run again as normal.

6.2.2 Testing

A Paramed 03-1373-00 148mm phantom was imaged using the original sequence and the modified sequence with the hold. Both sequences were acquired in the clinical interface with the same scan parameters. The signal to noise ratio was

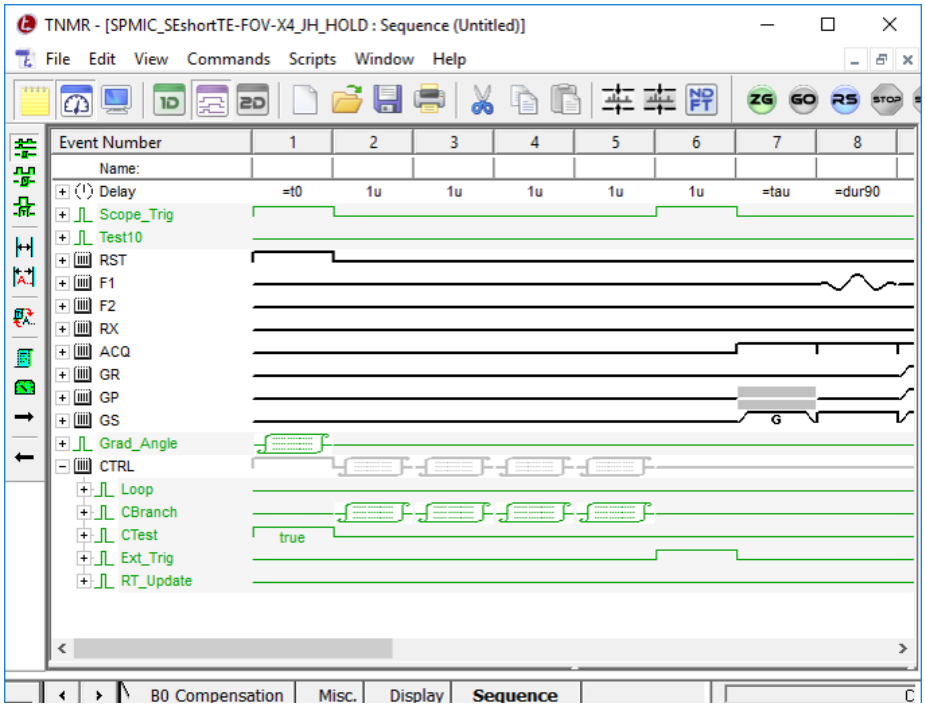

Figure 6.4: Branching tables can be seen in CBranch. The 1D Branching table is in column 2, 2D Branching table is in column 3, 3D Branching table is in column 4, 4D Branching table is in column 5. In the CTest row, the "true" value prepares the sequence to look for an external trigger event. The Ext_Trig, column 6, event waits for the 4.8V pulse to trigger it. The Scope_trig event in column 6 can be probed to query if the scan is ready to run.

Scan Order	1	2	3	4	5	6
Scan Type	New	Original	New	Original	New	Original
Average SNR	39.9	49.3	52.3	59.5	59.3	60.5

Table 6.1: SNR Ratio comparison between the new hold sequence and the original, unmodified sequence.

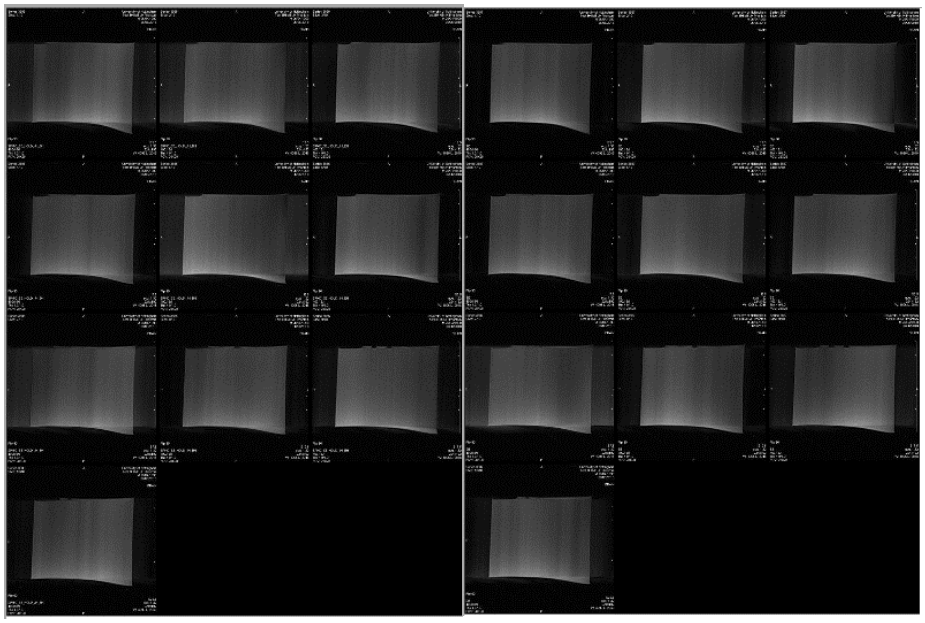

(a) (b)

Figure 6.5: Two consecutive images acquired with the same parameters; **(a)** new hold sequence and **(b)** original unmodified sequence.

measured for all the 10 slices and averaged. Each sequence was repeated 3 times The average SNR can be seen in table 6.1. The SNR varies between runs, even for the same sequence. However, there is similar variability between the two sequences and the SNRs are similar. For comparison, the 1st and 2nd scan from table 6.1 can be seen in figure 6.5.

Using the minimum TE/TR, a matrix size of 160x128 and 7 slices, the sequence is now acquired consistently within 21s ±0.5s after the trigger is pressed.

6.2.3 Work Performed Using Sequence

Using the 2D SE hold sequence, with a matrix size of 160x128x7 and a TE/TR of 12ms/161.8ms, 10 healthy volunteers were imaged. The study was approved by the National Research Ethics Committee and written informed consent was obtained. Images were acquired using the 4 channel BODY MAX coil where

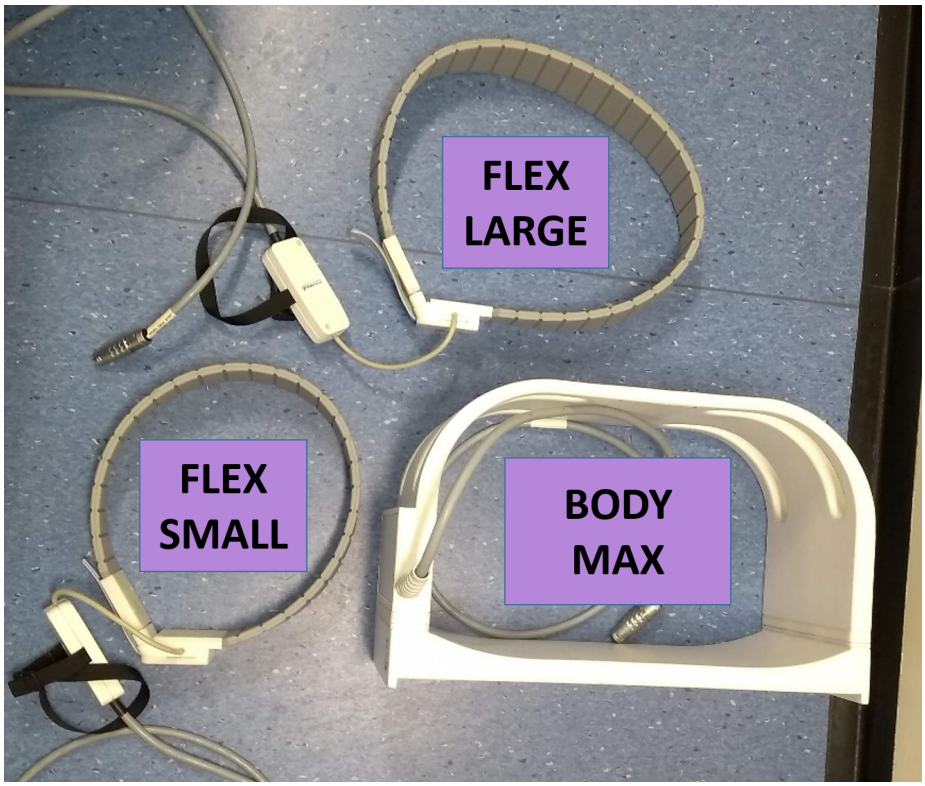

Figure 6.6: BODY MAX, FLEX SMALL and FLEX LARGE coils used for ^{1}H imaging on the Paramed.

possible (see figure 6.6). However, most volunteers were too large to be scanned using this coil, with the rigid shell breaking apart at end inspiration. Instead the FLEX SMALL coil was used for the majority of volunteers, with the occasional volunteer requiring the FLEX LARGE coil (see figure 6.6). Three coronal and three sagittal images were acquired supine and seated, both at full expiration and full inspiration. After manually selecting the diaphragm, the surface of the diaphragm was reconstructed. From the three images for each of the four conditions, an average surface was extracted. From these average surfaces six parameters were measured. These parameters can be seen in figure 6.7. Across 10 healthy volunteers, a larger maximum excursion and volume under the diaphragm was observable at full inspiration than full expiration, in both the seated and supine position. The excursion at full expiration was also larger seated compared to supine [150].

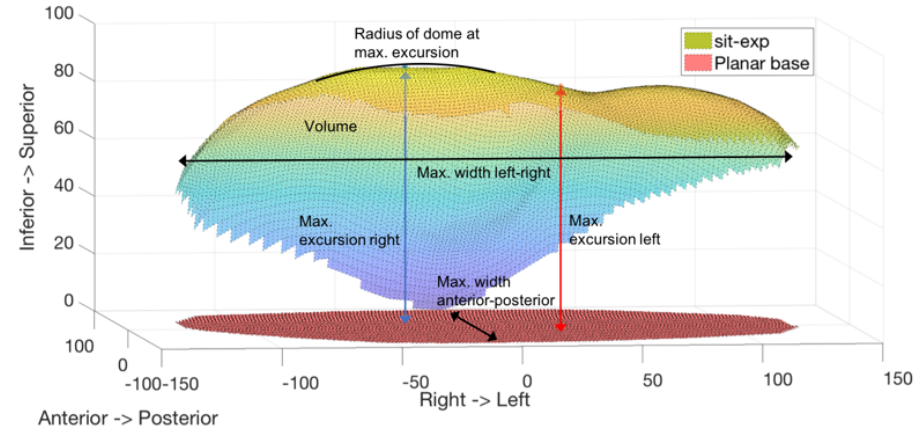

Figure 6.7: Parameters measured from average surface of a typical diaphragm at full inspiration and expiration. Image created by Dr Christoph Arthofer [150].

6.3 Compressed Sensing

Volunteers with diaphragmatic disorders can have impaired respiratory health. Ideally when imaging the diaphragm, the imaging should be performed during a breath-hold (normally at end expiration or end inspiration). This avoids the diaphragm moving during the sequence acquisition. For healthy volunteers, a duration of ∼20s is acceptable. However, for volunteers with respiratory conditions this may be unachievable (even when upright).

The Paramed scanner has limitations that result in the minimum duration of a scan, when running through the clinical interface, being significantly longer than ideal (see section 4.1.1). Running through TNMR however, gives huge flexibility in the matrix size acquired and gradient amplitudes.

Compressed Sensing (CS) refers to a technique where images are reconstructed using an undersampled k-space; only specified phase encodings are acquired (see section 2.1.6).

6.3.1 Method

To investigate how well an undersampling of k-space would perform, compared to a standard Paramed sequence; a 3D GRE with the minimum TR/TE=15.4/8ms, was used as a base. Depending on the sampling scheme, different numbers of the k-space points from this base sequence were acquired. As an example, the Paramed base sequence is acquired in 59 seconds when run with a matrix size of 256x160x24. However, when running through the clinical interface there is an additional unspecified preparation time of up to 15s. This is significantly too long for a breath hold. The FLEX SMALL coil was used for imaging (see figure 6.6). This was the coil used for the majority of previous diaphragm imaging

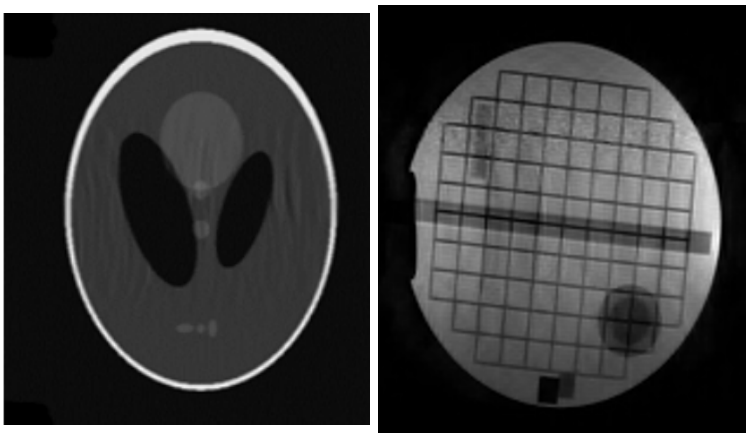

(a) (b)

Figure 6.8: **(a)** Digital phantom undersampled by factor of 3 using the same sampling scheme as in vivo imaging. A ripple artefact is visible. **(b)** ACR phantom used to test the sequence undersampled by a factor of 3.

and has a single channel.

Sampling Pattern

Sampling patterns were developed by Joshua McAteer using a digital phantom (see figure 6.8a). This was cylindrical, like the ACR phantom used for developing the sequence. The ACR phantom was selected because it has complicated fine structure inside and has similar dimensions to an adult's chest cavity (see figure 6.8b).

The undersampling pattern was formed of a fully sampled region in the centre of k-space, a Poisson disk variable density sampling, as well as additional quasi-random points. Using the retrospectively undersampled digital phantom, the size of the fully sampled region, sharpness of the central peak in the Poisson disk variable density sampling and number of random points were altered to maximise the visual information fidelity and structural similarity index [151]. Visual information fidelity was prioritised over the structural similarity index, because while the structural similarity index is a common image quality metric used in MRI, Mason Et. Al. [152] found that visual information fidelity was more closely correlated to radiologist's perceptions of diagnostic image quality.

Writing Sequence

The gradient amplitudes are held in tables in TNMR. These tables can be stacked in loops labelled 0D, 1D, 2D, 3D and 4D. As a convention for the 3D GRE sequences on the Paramed: 1D represents the frequency encoding, 3D the

first phase encoding and 4D the slice encoding. The number of 1D, 3D and 4D points can then be set as required. The 1D points are incremented first. When the 1D points is at maximum, the next table (2D) is incremented by 1 and the 1D points returns to one before incrementing again. However, if this protocol was retained for CS, the value and number of phase encoding points in 3D would need to alter depending on the 4D value. Instead the 4D table is collapsed into 3D. For example, a sequence with 200 3D points and 22 4D points, would result in 200x22=4400 3D points in a fully sampled collapsed table.

In normal mode, TNMR doesn't load tables into the memory chip of the board. Instead each table is read separately, even if tables have the same name and are in consecutive events. Memory limitations meant that in normal mode, the maximum number of phase encoding steps was 1365; for the 3D GRE sequence. This presented a challenge for scans using a larger matrix size.

By running TNMR in shared memory mode, the maximum length of a table would increase to 8192 on this line. However, shared memory mode does increase the "recycling time" at the end of each scan. In normal mode this is 60ns per table, but in shared memory mode this is 1μs per table. The maximum number of tables in a line of the standard 3D GRE sequence is 35. 35μs is insignificant compared to the minimum TR=15.4ms of the clinical sequence.

Phase Cycling

The original Paramed sequence has phase cycling of the RF and acquisition. This is useful in analogue systems where it can remove artefacts in the NMR acquisition from D.C. offsets and random noise. However, the Redstone spectrometer in the Paramed is digital, so the phase cycling serves little purpose. Images were acquired with and without phase cycling and showed insignificant difference (data not shown). The phase cycling was removed for simplicity and to speed up generating new sequences. Doing this removed the complication of deciding whether to increment the phase cycling step as normal, or retain the phase of each k-space point. For example if points 1 and 3 from k space were acquired and the phase cycling step was 180, do points 1 and 3 have phases of 0 and 180 or 0 and 360?

6.3.2 Image Reconstruction

Reconstruction of the images from the k-space data was performed by Joshua McAteer using the BART Toolbox [153]. Three regularisers were used in the reconstruction; a L1 sparsity regulariser in image space, L1 sparsity regulariser in wavelet space and a total variation regulariser.

Comparing Images

Unfortunately, there are differences in the images when the same scan is repeated with a static phantom. The instability issues of the scanner are currently being investigated. The instability can be seen in table 6.1, where the SNR varied by

23% in three identical scans. This is for a spin echo sequence, which should be stable. The instability can make it difficult to calculate computationally which sequence is "better".

Figure 6.9 shows three runs of the same GRE sequence (448x200x22), undersampled by a factor of 3, with a centre out k-space trajectory, 80 dummy lines and a 45° FA. A 2 minute gap was left between scans and the phantom was left in place. The average SNR was calculated for slices 11-14 at 42.1, 34.8 and 27.1 for figure 6.9a, 6.9b and 6.9c respectively. Like with the spin echo sequence there was large variability in the SNR.

There is also large variability in the ghosting artefacts between the images. Figure 6.10 is the same as figure 6.9, but with the square root taken of the voxel intensities. This reduces the contrast in the images and accentuates the ghosting artefacts.

6.3.3 Results

Flip Angle

The average T1 of the ACR phantom was calculated at $\approx 0.147s$. This was measured with a saturation recovery pulse sequence. A 1 minute gap was left between pulses. For a TR of 15.4ms this puts the Ernst angle at $\approx 25.8°$. However, taking the T1 in the lung to be $\approx 1008ms$ [144]. The Ernst angle becomes $\approx 10.0°$. For investigations into the improvements into image quality it may be beneficial to move to higher FAs because it is at higher FAs that image artefacts become more pronounced (see figure 6.11).

Dummy Lines

Within a GRE sequence, before any RF pulse, the longitudinal magnetisation is M_0. However, during faster GRE sequences (TR $\not\gg$ T1), the longitudinal magnetisation won't have time to recover to M_0. Instead it will stabilise at M_{ss}. In general, the higher the FA and shorter the TR, the longer M_{ss} takes to be established and the lower M_{ss}. When operating with low-high k-space trajectories there was less ghosting than when using centre out k-space trajectories. This is believed to be because the majority of the energy density is clustered around the centre of k-space. Because the centre of k-space is sampled first (before the longitudinal magnetisation has reached steady state), the irregularities are more apparent.

"Dummy lines" refer to multiple echos at the start of the sequence where approximately the centre of k-space is sampled. By taking multiple dummy lines before the scan begins, the longitudinal magnetisation is established prior to acquiring data. One way to check if the dummy lines were functioning correctly is to compare the last few dummy lines to the centre of k-space when it is actually acquired in the sequence (this will be more useful for a low-high k-space trajectory when it is acquired later), to assess if the longitudinal magnetisation has reached steady state by the last dummy line. However, because the scanner

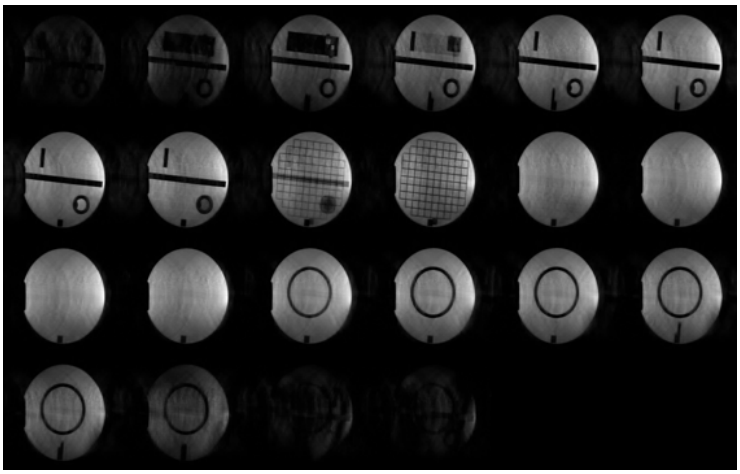

(a)

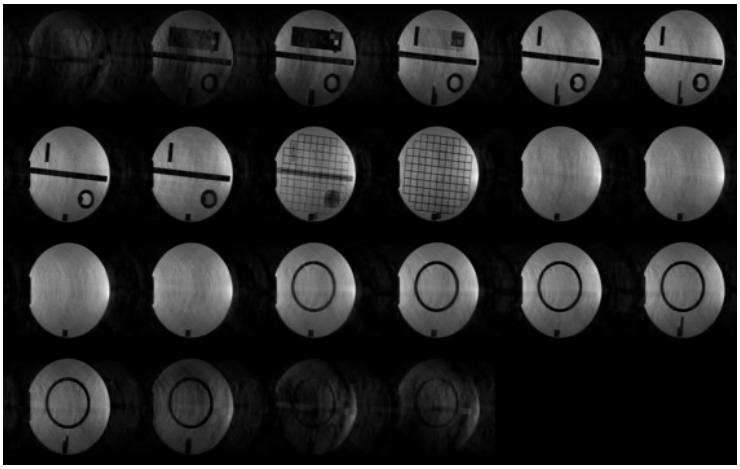

(b)

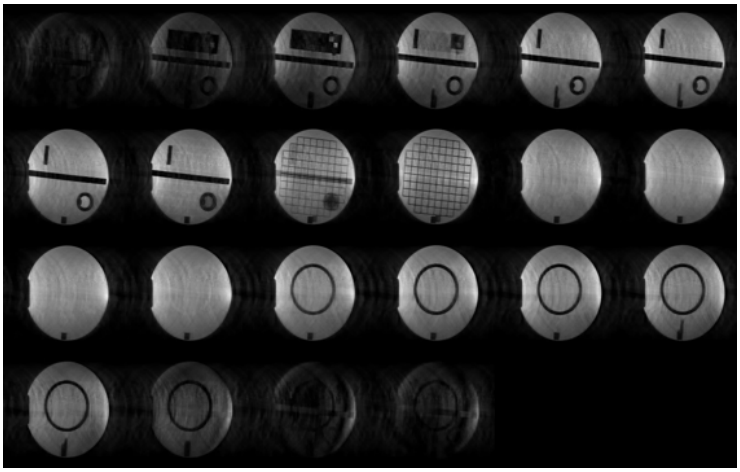

(c)

Figure 6.9: 3 runs of the same GRE sequence undersampled by a factor of 3, centre out k-space trajectory, 80 dummy lines, 45° FA. Displayed in order acquired.

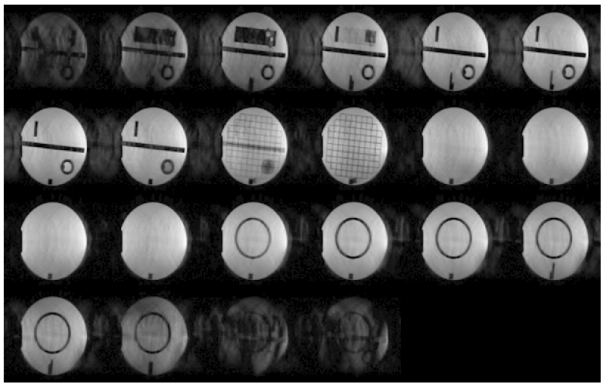

(a)

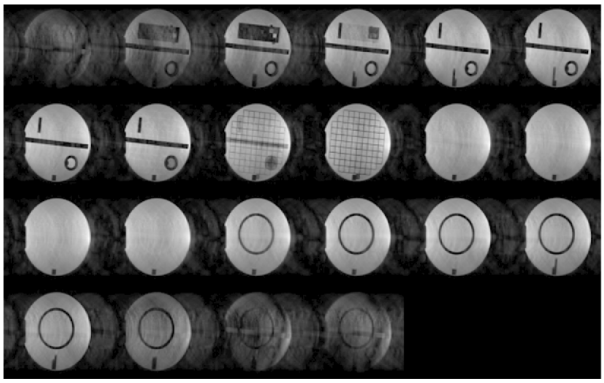

(b)

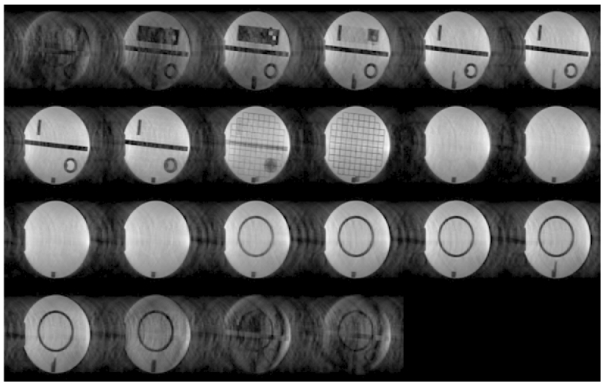

(c)

Figure 6.10: Figure 6.9 with square root taken of voxel intensities. Large variability in ghosting artefacts.

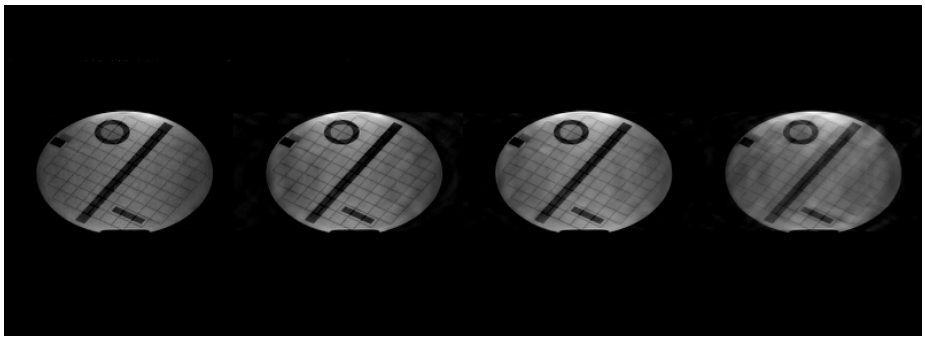

Figure 6.11: Representative slice acquired at four flip angles using a 3D GRE sequence, acquired with a matrix size of 288x200x12 and undersampled by a factor of three. Left to right 10°, 25°, 45° and 90° FA. 90° FA image is less sharp than 10° FA image.

images varied so much between images, and therefore acquisitions, this won't be a relevant comparison to a single acquisition.

Using a 45° FA, centre out trajectory, (448x200x22) undersampling by a factor of 3 and 80 dummy lines, the same sequence was repeated three times with 2 minutes gap between. Contour plots show the magnetisation stabilising to some extent after approximately 30 scans. Prominent repeatable islands can also be seen in the plots (see figure 6.12). A very high intensity spike can also be seen during the first 3 dummy lines. The reduction in intensity shows the magnetisation reducing from M_0 to $\sim M_{ss}$.

Using a 45° FA, centre out trajectory, undersampling by a factor of 3 and 0, 5, 10, 20, 40 and 80 dummy lines, the ACR phantom was imaged (see figure 6.13). SNR was calculated for the bottom two slices of each image in figure 6.13. The SNR was 13.3, 11.7, 18.3, 14.3, 13.9 and 15.4 for the 0, 5, 10, 20, 40 and 80 dummy lines scans respectively. Due to the intra-scan variability it is difficult to say if dummy lines have a significant effect on the SNR of the images. Ghosting artefacts also varied between the images, but they did not become noticeably less prominent as the dummy lines increased.

Dummy Lines Summary The magnetisation appeared to stabilise after about 30 dummy scans. However, dummy lines also increase the duration of the sequence. A low-high, rather than centre out k-space trajectory is recommended, to ensure the centre of k-space is sampled after a large number of TRs. 5 dummy lines are also recommended because a spike was observed in the contour plots (see figure 6.12) during the first 5 dummy lines. When higher rates of undersampling are used in the future the number of dummy lines before the sequence may need to be revisited.

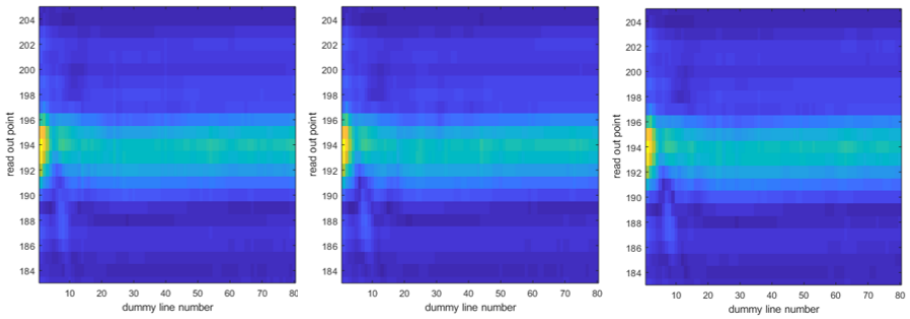

Figure 6.12: Left to right runs 1 to 3 of 80 dummy lines before sequence. x-axis dummy line point, y-axis readout point. A high intensity peak is visible in the first 3 lines and a repeatable shaped island can be seen between the 6th and 11th scan in lower left quadrant of images. Yellow shows a high intensity, blue low intensity.

Undersampling Factor

A stock Paramed 3D GRE sequence, with a 45° FA and a matrix size of 256x200x22 was used as a base to compare to the same sequence undersampled by a factor of 3 and 4. Figure 6.14 also shows fully sampled low-high k-space trajectory images, with the same number of k space points as the undersampled acquisitions (matrix size of 256x67x22 and 256x50x22). The FOV was maintained. The FA of 45° was chosen to accentuate image artefacts. For images acquired within the same duration, the under sampled sequences produced sharper images than the full sampled sequences (comparing figure 6.14b to 6.14c and figure 6.14d to 6.14e).

Approval was obtained to use the sequence with a human volunteer from Andrew Peters (centre manager of SPMIC). The volunteer gave written informed consent and was scanned under the University SPMIC 0.5T development ethics (ethics number H14028-004). All images were acquired at end inspiration using a 10° FA. A 3D GRE sequence with matrix size 256x200x22, undersampled by a factor of 3, was used coronally. This puts the scan duration at 22.6s. The undersampled images couldn't be compared to the fully sampled data because the volunteer couldn't be asked to hold their breath for > 1 minute. However, images can be compared to a fully sampled scan with a matrix size of 256x67x22; scan duration 22.7s (see figure 6.15). A 256x160x10 undersampled by a factor of 4.82 image was also acquired (scan duration 5.1s). For volunteers with more developed respiratory conditions a 22s breath hold may be unachievable.

6.3.4 Discussion

Without a stable image quality between runs of the same sequence, it is difficult to determine if variation between images is to do with the intra-scan variability,

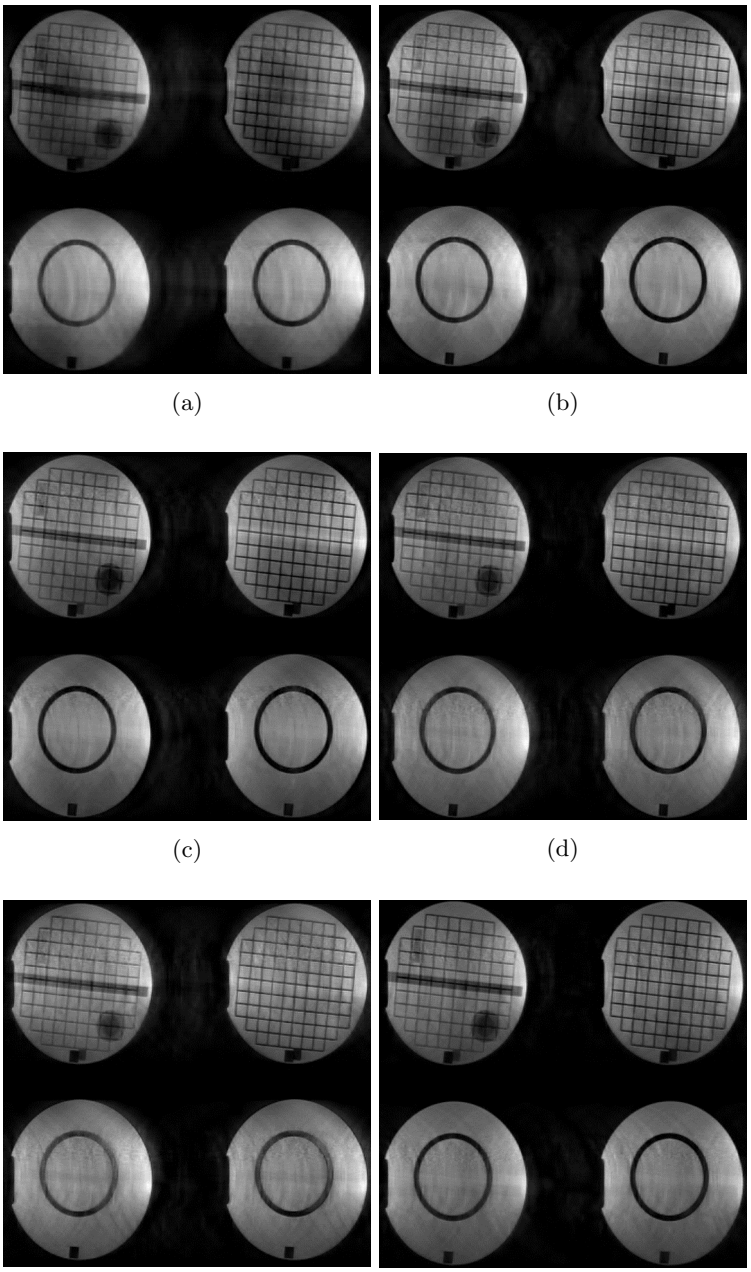

(a) (b)

(c) (d)

(e) (f)

Figure 6.13: 0, 5, 10, 20, 40 and 80 dummy lines for a,b,c,d,e and f respectively. 4 non-consecutive slices of phantom.

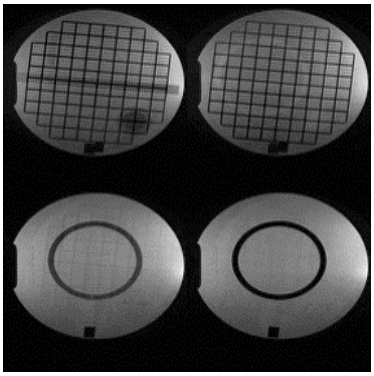

(a)

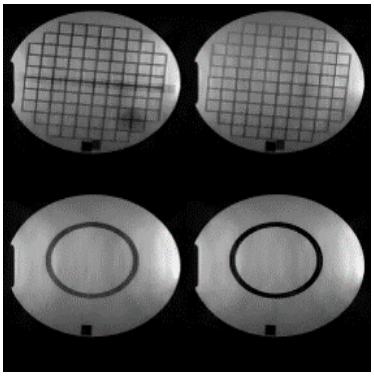

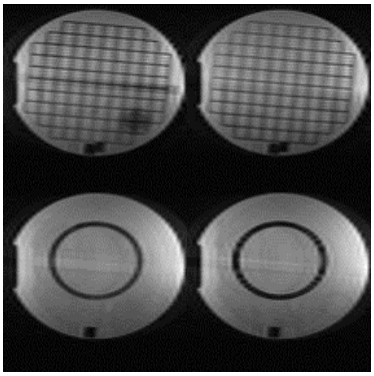

(b) (c)

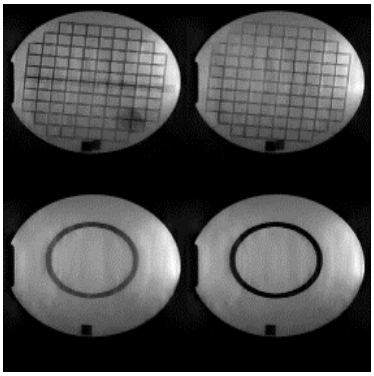

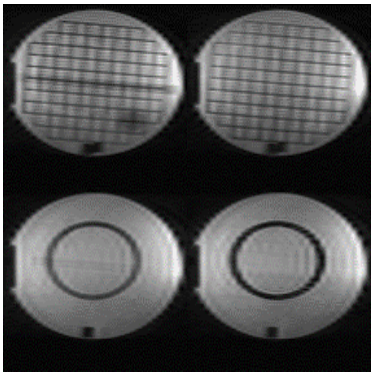

(d) (e)

Figure 6.14: All undersampled images undersampled the 256x200x22. **(a)** 256x200x22 fully sampled, **(b)** undersampled by factor of 3, **(c)** 256x67x22 fully sampled, **(d)** undersampled by factor of 4, **(e)** 256x50x22 fully sampled. **(b)** and **(c)** were acquired in 22.7s, and **(d)** and **(e)** were acquired in 17.0s.

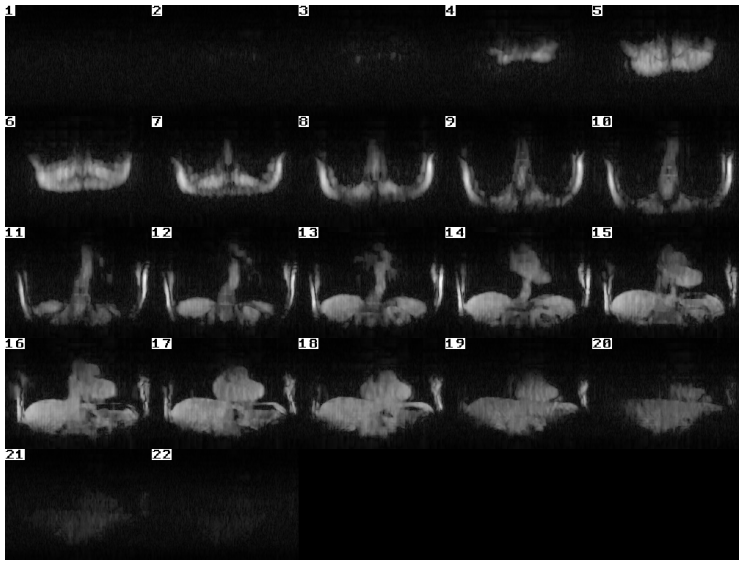

(a)

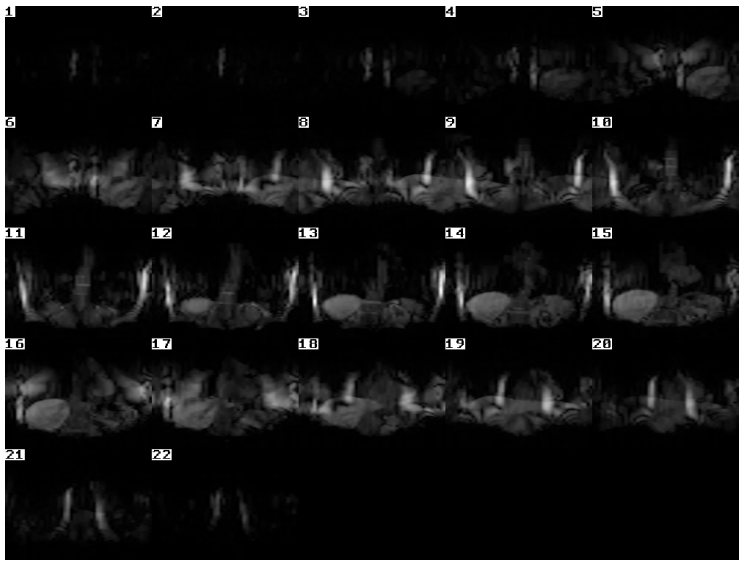

(b)

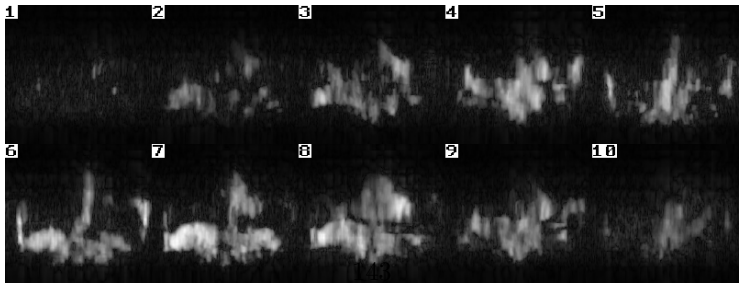

(c)

Figure 6.15: 10° FA coronal images: **(a)** 256x200x22 undersampled by a factor of 3, **(b)** 256x67x22 fully sampled and **(c)** 256x160x10 undersampled by factor of 4.82.

or a difference between the sequences. Although dummy lines don't appear to have a marked effect on the image quality, contour plots do show that the dummy lines become more consistent after 5-30 scans. 5 was chosen because this adds on an insignificant duration to the sequence and after 5, the sharp peak at the start of the dummy scans is avoided. A low-high, rather than centre out k-space trajectory, also ensures the centre of k-space is sampled after the longitudinal magnetisation has become more established.

The origin of the large intra-scan variability was not fully understood at the time of acquiring the images. Investigations have previously been performed, which discounted the lights in the scanner room, or patient intercom, being to blame. Paramed were contacted to perform investigations. They believed the origin may have been the scanner's cold head which has been replaced. At the time of writing, the intra-scan variability has not been reassessed. These investigations will need to be performed before other sources, such as the bed motor or air-conditioning, can be discounted.

By removing phase cycling of the RF and acquisition, sequences were simpler and faster to generate with no noticeable changes in image quality.

From figure 6.14, it can be seen that even fine detail can be distinguished at high rates of undersampling. However, as the amount of undersampling increases wave like patterns become more obvious in the images and the contrast is reduced. The undersampled images are noticeably sharper and the fine structure of the phantom is better resolved than the fully sampled images acquired within the same duration.

Looking at figure 6.15a: By undersampling by a factor of 3, a 256x200x22 matrix size image can be acquired in 22.6s; within a breath hold. This is an increase from the matrix size of 160x128x7 previously acquired in the spin echo hold sequence sequence in a similar duration (21s) (see section 6.2).

The fully sampled low-high k-space trajectory sequence (figure 6.15b), with a matrix size of 256x67x22, can be acquired is 22.7s. The undersampled sequence images are higher resolution, but appear to be slightly nosier when compared to the fully sampled images. However, in the fully sampled images, aliasing (wrap around) can also be observed which covers the ROI. Wrap around does not occur over the ROI in the undersampled images. In order to avoid wrap around, the FOV could be increased, but this would reduce the number of voxels within the dimension of the diaphragm making it more difficult to pick out. The outline of the diaphragm is easier to pick out in more slices in the undersampled image than the fully sampled image.

Even when imaging with a matrix size of 256x160x10 and undersampling by a factor of 4.82, the diaphragm could be distinguished (see figure 6.15c).

It may be useful to apply the CS work to a SE sequence, which should provide more (T2) contrast and hence make comparing some in vivo images an easier task. A 3D GRE sequence was used initially because this sequence can also be used for hyperpolarised ^{129}Xe imaging (see section 7.1.3). Hyperpolarised ^{129}Xe similarly needs to be performed within a breath hold; by undersampling higher resolutions can be achieved. For hyperpolarised imaging there is an additional advantage. The fewer TRs used, the more polarisation "available"

per acquisition; a higher FA can be used. This should increase the SNR of the image.

It would also be useful to use a coil with more than one receiver. The FLEX SMALL coil was used, because historically it has been used for respiratory imaging. However, the BODY MAX coil (see figure 6.6) has four receivers and therefore may perform better for CS sequences. With four receivers, more data is acquired per acquisition; the data should be more compressible.

6.3.5 Conclusion

A CS 3D GRE sequence has been developed, which within a single breath hold can produce an image with matrix size of 256x200x22. This is a large improvement from the previous SE sequence which used a matrix size of 160x128x7 and was acquired in a similar duration; 22.7s (CS 3D GRE) versus 21s (SE). The CS sequence should allow for a more accurate characterisation of the diaphragm. The CS 3D GRE sequence was also performed in vivo using thinner slices (5mm), than previously used by the SE hold sequence (16mm). In vivo, the undersampled sequence allowed for the diaphragm to be more accurately characterised, than a fully sampled 3D GRE sequence acquired in the same duration.

The CS 3D GRE sequence can also be used with a matrix size of 256x160x10 and undersampled to produce an image within $\sim 1/4$ of the duration of previously. This opens the prospect for imaging volunteers with more progressed lung conditions.

These undersampling schemes should next be applied to the ^{129}Xe 3D GRE sequence, where shorter scan durations, with higher resolutions, would also be beneficial.

Chapter 7

Upright 0.5T Multinuclear Imaging

In this chapter multinuclear work on the Paramed 0.5T upright scanner will be discussed. This is broken down into two sections:

- In section 7.1, hyperpolarised ^{129}Xe work on the Paramed scanner is detailed.
- Section 7.2 contains work performed on proton imaging of a hydrofluorocarbon. Although work was performed at the proton Larmor frequency, it is included in this chapter because the work is in preparation for fluorine imaging.

7.1 ^{129}Xe Imaging

Current ^{129}Xe imaging is performed using a small ^{129}Xe surface coil (see section 4.1.4). There is an order in place to PulseTeq Ltd. for a bespoke chest coil.

7.1.1 Flip Angle Calibration

It is important to accurately know the FA for a given attenuation when imaging. Once the FA (α) at a given attenuation is known, the optimal FA(s) can be chosen for imaging (see section 7.1.3).

The magnetisation loss of hyperpolarised ^{129}Xe, during an RF pulse, is non-recoverable. This means even if a second pulse is at a higher FA, the signal received back may be less due to the lower polarisation before the second pulse.

The standard way to extract the FA for a given attenuation, uses multiple RF excitations at the same FA. The transverse hyperpolarised signal (S_k) of the kth pulse decays according to equation 7.1. This assumes there is no excess magnetisation from the previous RF pulse. M_0 is the longitudinal magnetisation

before the first pulse.

$$S_k = M_0 exp[-(k-1)TR/T1] sin\alpha_k \prod_{j=1}^{k-1} cos\alpha_j \tag{7.1}$$

For a constant FA this reduces to equation 7.2, where S_1 is the signal from the first pulse.

$$S_k = S_1(cos(\alpha)exp(-TR/T1))^{(k-1)} \tag{7.2}$$

The transmit gain for a given FA obeys equation 7.3 [154]. TG is in dB and $\lambda \sim 20$ is constant for a given RF pulse shape. Equation 7.3 allows the FA to be extracted for a variety of attenuations.

$$TG_\alpha = TG_\pi - \lambda log_{10}\frac{\pi}{2\alpha} \tag{7.3}$$

By measuring the transverse magnetisation after repeated equal attenuation pulses, the resulting signal can be fitted to equation 7.2. For sufficient RF pulses, the FA can be accurately determined. However, λ must be approximated or calculated separately.

While a FA calibration can be performed in a phantom, a human body loads a coil very differently to a Xe Tedlar bag and Hartmann's solution bags. This means the RF power will be deposited within the human tissue differently to the phantom. In order to obtain a calibration which is accurate in vivo, the FA calibration should be performed in vivo.

In a Tedlar bag at 0.5T, T1 ∼1hour. However, the T1 of ^{129}Xe in the lungs will be significantly reduced compared to in a Tedlar bag. This is due to the presence of oxygen in the lungs (see equation 7.4) [129].

$$\frac{1}{T1} = \frac{pO_2}{\xi} \tag{7.4}$$

Where pO_2 is the partial pressure of O_2 in bars and ξ is the oxygen enhancement factor, which depends on temperature. Most previous studies rely on two measurements [155]; one for FA calibration and one for $T1$ measurement. This isn't ideal because it would rely on two breath holds.

Zhong et al. propose a scheme to determine the FA and $T1$ during a single breath hold [154]. They found that eight RF pulses were sufficient for determining the FA accurately for a given attenuation. They proposed performing eight excitations at different RF pulse attenuations, decreasing the RF attenuation each time.

Position 1 Using the ^{129}Xe surface coil (see section 4.1.4), to get an initial gauge of the RF attenuation, the decrease in attenuation was chosen to be 2dB starting at 50dB. The Tedlar bag was placed inside the coil (see figure 7.1) and loaded with two Hartmann's solution bags. This arrangement will be referred to as position 1. The RF pulse used was the same shape (see figure 7.10) and length (3ms) as that used in the ^{129}Xe imaging sequence (see section 7.1.3).

Transmit Gain (dB)	Fitted Theta (Radians)	R^2
-50	0.1127 (0.002093, 0.2233)	-0.44
-48	0.1143 (8.52E-02, 0.1434)	0.593
-46	0.0781 (0.03264, 0.1235)	0.32
-44	0.1564 (0.1333, 0.1795)	0.82
-42	0.1778 (0.1457, 0.2199)	0.78
-40	0.2422 (0.2319, 0.2525)	0.97
-38	0.2765 (0.2672, 0.2859)	0.98
-36	0.3337 (0.326, 0.3414)	0.99
-34	0.3923 (0.3841, 0.4005)	0.99
-32	0.4789 (0.4729, 0.4848)	0.99
-30	0.4858 (0.4689, 0.5028)	0.99
-28	0.586 (0.5437, 0.6283)	0.95

Table 7.1: Position 1 fitted FA data. Including the 95% confidence bounds of the fitted FA in brackets (see figure 7.3).

The TE was set to 4ms and the TR at 3s to ensure no transverse coherences remained between pulses. At 0.5T, the T1 of ^{129}Xe will be $\sim$1hour and therefore in equation 7.2 $exp(-TR/T1) \rightarrow 1$. The Xe phantom was placed within the coil. Note the surface coil will have an inhomogeneous RF transmit as well as receive (see section 2.1.12). This means the FA of the sample will vary depending on the sample's location.

By part way through the 12th set of RF pulses, the signal had decayed into noise (see figure 7.2). The FA for the later sets was not fitted. Using the Matlab Curve Fitting Toolbox, the data for each set of 8 pulses was fitted to equation 7.2 in the limit $exp(-TR/T1) \rightarrow 1$ and the curve was fixed through the 1st of the 8 pulses (see table 7.1). For the first three sets of RF pulses the FA was not large enough to cause an appreciable loss in polarisation so the fit was poor for those points.

FAs fitted with an $R^2 > 0.7$ are shown in figure 7.3, along with their 95% confidence interval as error bars.

The fitted FAs were fitted using the curve fitting toolbox in Matlab to equation 7.3, using the R^2 values of the fitted FAs as weights. The resulting fit parameters can be founds in table 7.2. Using equation 7.3, the FA for a given RF attenuation can be extracted. This is only valid for the sample pulse length and duration and assumes the phantom is in the same location with the same loading. The validity of the curve at higher attenuations (>44dB) has not been investigated.

Position 2 To get a more accurate measure for the FA in a human, the phantom was then placed in position 2 (see figure 7.4). Here the Tedlar bag is mimicking a human lung and the Hartmann's solution bags the tissues in the back behind the lung.

Figure 7.1: Position 1 Tedlar bag location in relation to coil with loading.

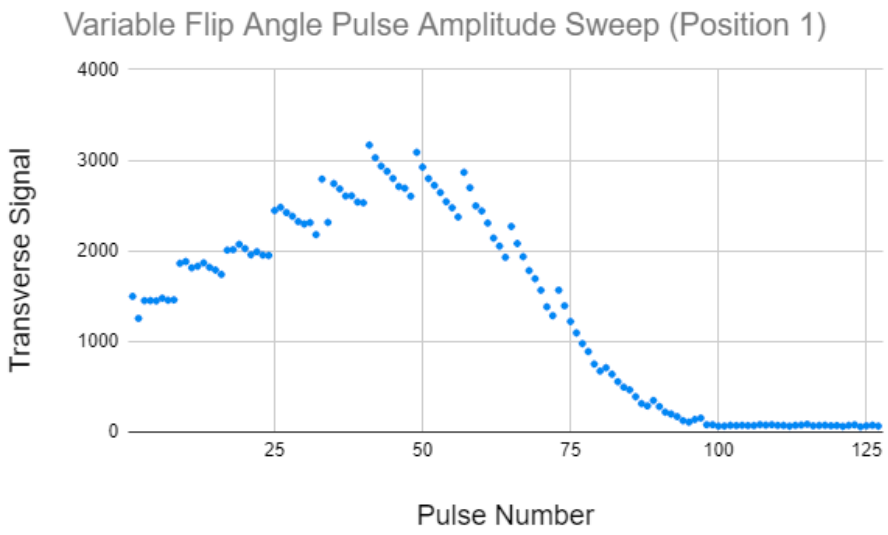

Figure 7.2: Variable FA calibration pulse amplitude sweep. 2dB attenuation decrease every 8th pulse starting at 50dB.

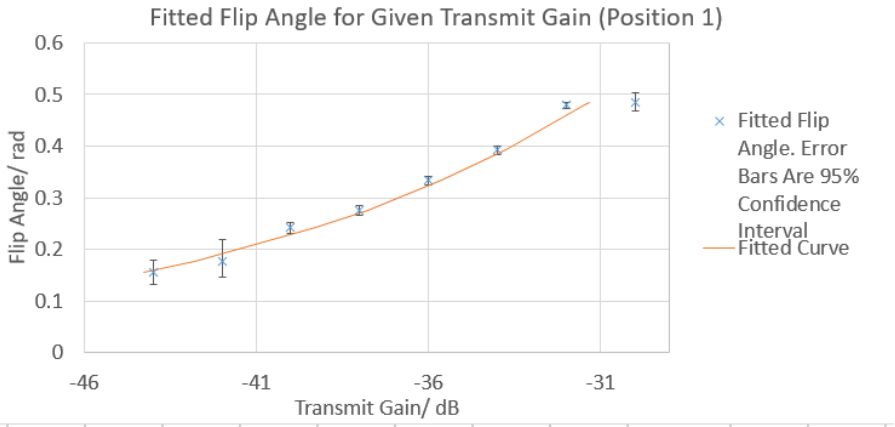

Figure 7.3: FA for a given RF pulse attenuation, position 1. The fitted curve is of the form equation 7.3.

Figure 7.4: Position 2 Tedlar bag location in relation to coil with loading.

Position	λ	TG_π (dB)	R^2
Position 1 (attenuation gap 2dB)	26.29 (22.36, 30.23)	-17.92 (-20.8, -15.03)	0.978
Position 2 (attenuation gap 2dB)	22.71 (16.14, 29.28)	-14.14 (-18.82, -9.459)	0.940
Position 2 (attenuation gap 1dB)	23.88 (18.02, 29.74)	-15.26 (-19.44, -11.07)	0.904

Table 7.2: Equation 7.3 fit parameters. 95% confidence interval of fitted values in brackets.

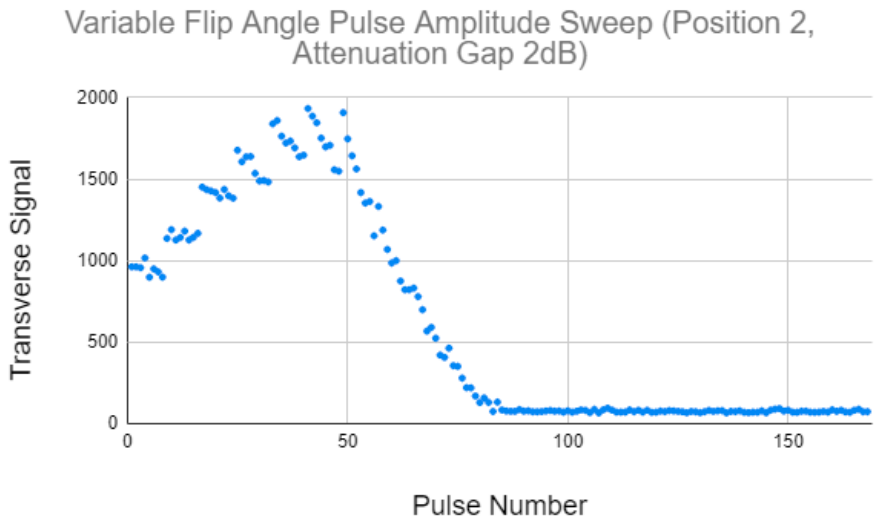

Figure 7.5: Variable FA calibration pulse amplitude sweep. 2dB attenuation decrease every 8th pulse starting at 42dB.

The sequence parameters were kept the same as for position 1. However, in position 2 an attenuation of >44dB wasn't sufficient to cause an appreciable loss in signal. The transmit field will decay as distance from the coil increases, so the first pulse was chosen to be 42dB. The same step size of 2dB was used (see figure 7.5).

By part way through the 11th set of pulses (attenuation = 22dB), the transverse signal was indistinguishable from the noise. Lower attenuations weren't fitted. As with the position 1 data, the first 3 attenuations weren't sufficient enough to depolarise the bag enough to obtain an $R^2 > 0.7$. Table 7.3 shows fitted FAs for position 2 with an attenuation gap of 2dB.

In an attempt to exploit the polarisation better, the experiment was repeated with a new Tedlar bag in approximately the same position and loading. Every

Transmit Gain (dB)	Fitted Theta (Radians)	R^2
-42	0.1138 (0.05246, 0.1752)	0.27
-40	4.77E-07 (-1.07E+04, 1.07E+04)	-0.44
-38	0.1137 (0.09268, 0.1347)	0.57
-36	0.1944 (0.1733, 0.2155)	0.84
-34	0.1846 (0.1684, 0.2009)	0.90
-32	0.2487 (0.2343, 0.2631)	0.96
-30	0.3637 (0.3465, 0.3809)	0.97
-28	0.3959 (0.37, 0.4218)	0.93
-26	0.422 (0.4162, 0.4679)	0.96
-24	0.5663 (0.5354, 0.5973)	0.97
-22	0.5073 (0.408, 0.6067)	0.60

Table 7.3: Position 2 fitted FA data with an attenuation gap of 2dB. Including the 95% confidence bounds of the fitted FA in brackets (see figure 7.7).

effort was made to keep the position consistent, but small changes in location can result in a large change in the resulting FA for a coil with inhomogeneous transmit. This time, the first attenuation used was 37dB and the step size was 1dB. A smaller step size was chosen in order to obtain more points; lower FAs depolarise less than higher ones (see figure 7.6).

By the end of the 11th set of pulses (attenuation = 27dB), the transverse signal was indistinguishable from the noise. Lower attenuations weren't fitted. Table 7.4 shows fitted FAs for position 2 with an attenuation gap of 1dB.

FAs fitted with an $R^2 > 0.7$ for both sets of position 2 data are plotted in figure 7.7, along with their 95% confidence interval as error bars. Each set of data was fitted to equation 7.3 using the curve fitting toolbox in Matlab and the R^2 values of the fitted FAs as weights. The resulting fit parameters can be founds in table 7.2.

Summary In figure 7.8, all 3 curves are plotted from a FA of π to $\frac{\pi}{20}$. The 2 curves showing data from position 2 were very similar. However, the 95% confidence bounds cross with position 1. As expected, the TG_π and λ were lower for the bag in position 1 (closer to the coil), than position 2. The fit for the position 1 data is better with an $R^2 = 0.978$, than that of the position 2 data which for an attenuation gap of 2dB $R^2 = 0.940$ and for an attenuation gap of 1dB $R^2 = 0.904$ (see table 7.2). This is likely because of the inhomogeneous receive of the coil and therefore the SNR was higher for the bag closer to the coil.

When the new PulseTeq coil arrives, this experiment will need to be repeated in vivo. A variable FA sequence, such as that suggested in section 7.1.3, requires an accurate FA calibration in order to exploit the available polarisation fully.

In the FA calibration performed on the surface coil, when fitting to equation 7.2, the fitted line was forced through the first point. If the data perfectly

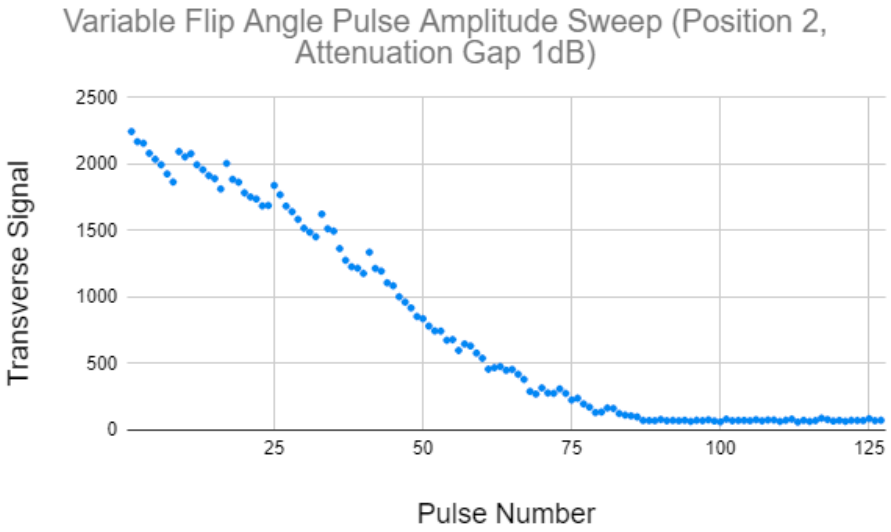

Figure 7.6: Variable FA calibration pulse amplitude sweep. 1dB attenuation decrease every 8th pulse starting at 37dB.

Transmit Gain (dB)	Fitted Theta (Radians)	R^2
-37	0.2235 (0.217, 0.2301)	0.99
-36	0.1884 (0.174, 0.2029)	0.93
-35	0.241 (0.2217, 0.2603)	0.86
-34	0.2675 (0.2594, 0.2756)	0.99
-33	0.3171 (0.3004, 0.3338)	0.96
-32	0.3322 (0.3285, 0.3458)	0.97
-31	0.2932 (0.2737, 0.3127)	0.95
-30	0.3403 (0.3098, 0.3708)	0.90
-29	0.4181 (0.3659, 0.4704)	0.80
-28	0.481 (0.4495, 0.5124)	0.96
-27	0.48 (0.4459, 0.514)	0.95
-26	3.62E-09 (-8.21E+05, 8.21E+05)	-0.09

Table 7.4: Position 2 fitted FA data with an attenuation gap of 1dB. Including the 95% confidence bounds of the fitted FA in brackets (see figure 7.7).

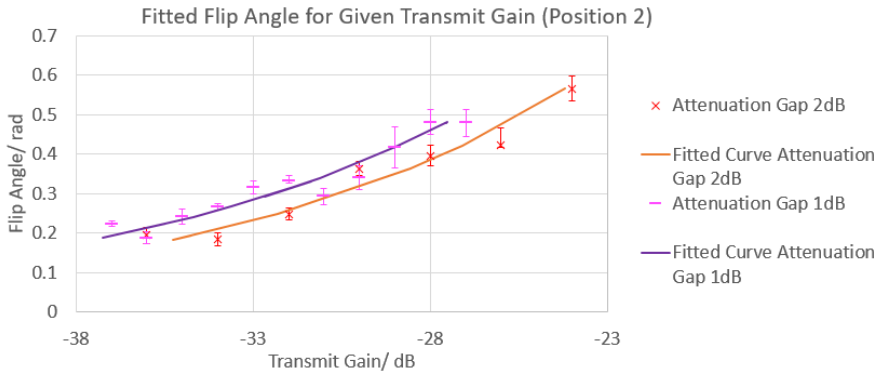

Figure 7.7: FA for a given RF pulse attenuation, position 2. 2 runs were performed with the bag in a similar location with different gaps in attenuation. The fitted curves are of the form equation 7.3. Error bars show the 95% confidence bounds of each fitted value.

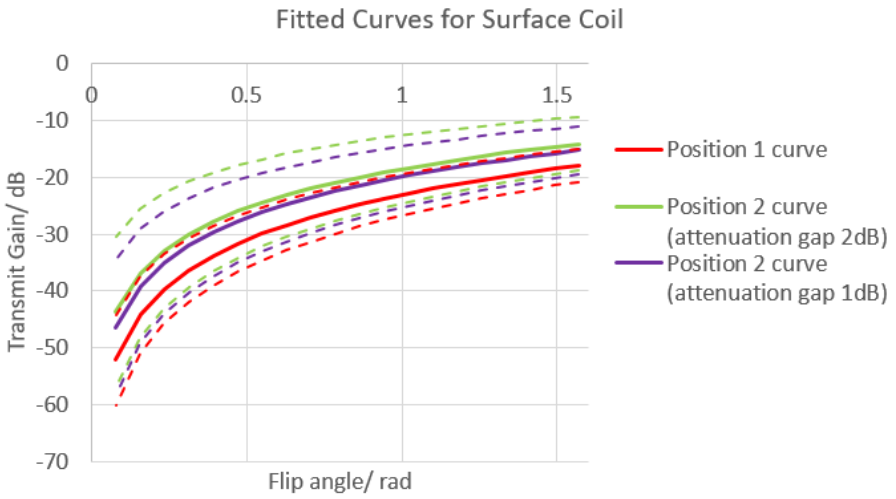

Figure 7.8: Curves fitted to equation 7.3 with parameters and bounds from table 7.2.

follows the theory, then the fitted curve will go through all the points and the amplitude of the first point will equal S_1 in equation 7.2. However,the first point is just as likely to be erroneous as any of the others. Therefore, in future, it is recommended to avoid forcing the curve through a point like this. It should also be noted that the assumption that $exp(-TR/T1) \rightarrow 1$ in the Tedlar bag, is unlikely to be a valid assumption in vivo.

Sample	Chemical Shift/ ppm
Water	190.0
Olive Oil	192.2
tert-Butanol	186.7
Ethanol	160.5

Table 7.5: Dissolved Phase Chemical shift of ^{129}Xe in different liquids. The centre of two peaks was chosen for olive oil.

7.1.2 Dissolved Phase ^{129}Xe Spectroscopy

In order to investigate gas exchange within the lung with hyperpolarised ^{129}Xe, a ^{129}Xe signal needs to be acquired in the dissolved phase. While the surface coil may not provide a homogeneous B_1, it should be very sensitive, making it well tailored for dissolved phase imaging.

Using the same RF pulse shape as in the Xe GRE sequence (see section 7.1.3), a 3ms RF pulse with 30dB attenuation was used. Looking at the FA calibration for position 1 (see table 7.2), this equates to a FA of $\sim 30°$. A $30°$ pulse was decided upon to ensure the FA wasn't $> 90°$ at any location; the FA calibration is very dependent on position. A short TE of 2.2ms was employed to maximise the signal in the acquisition window and the T1 and T2* of the Xe are likely to be very short in the dissolved phase, with high O_2 concentration (see equation 7.4). An acquisition time of 1.31072ms, dwell time 327.68μs and 4000 acquisition points over a spectral width of $\pm$1525.8789063Hz were used.

Dissolved phase spectroscopy was performed on a range of liquids, in which ^{129}Xe is soluble, to compare the chemical shifts. The chemical shift of dissolved phase ^{129}Xe is approximately 200ppm in most liquids. The liquid in question was first sucked into a syringe. ^{129}Xe gas was then sucked from a Tedlar bag into the syringe. The syringe was then shaken to increase the proportion of ^{129}Xe in the dissolved phase compared to gaseous phase. The syringe was then placed inside the surface coil (see section 4.1.4) and the scan was performed. On the day of scanning, the gaseous peak was centred at 5.84035MHz. The resonance of ^{129}Xe dissolved in water was found at 5.84146MHz. This gives it a chemical shift of 190.0ppm. ^{129}Xe dissolved in tert-Butanol, olive oil and ethanol were also scanned. The chemical shifts can be seen in table 7.5. The oil sample offered a double peak (see figure 7.9), separated by 0.765ppm, perhaps due to the oil being made as a mixture of olive oil from multiple sources (the centre of the two peaks was chosen to give a value for the chemical shift). The chemical shifts measured are broadly similar to those recorded by Miller and Williamson et al. [156]. Their ^{129}Xe dissolved phase peak in olive oil had a slightly higher chemical shift than in water at 198ppm and 196ppm respectively. They also used 1-Butanol rather than tert-Butanol and got a chemical shift of 176ppm. Finally, they obtained a 165ppm shift for the ^{129}Xe dissolved phase peak in ethanol.

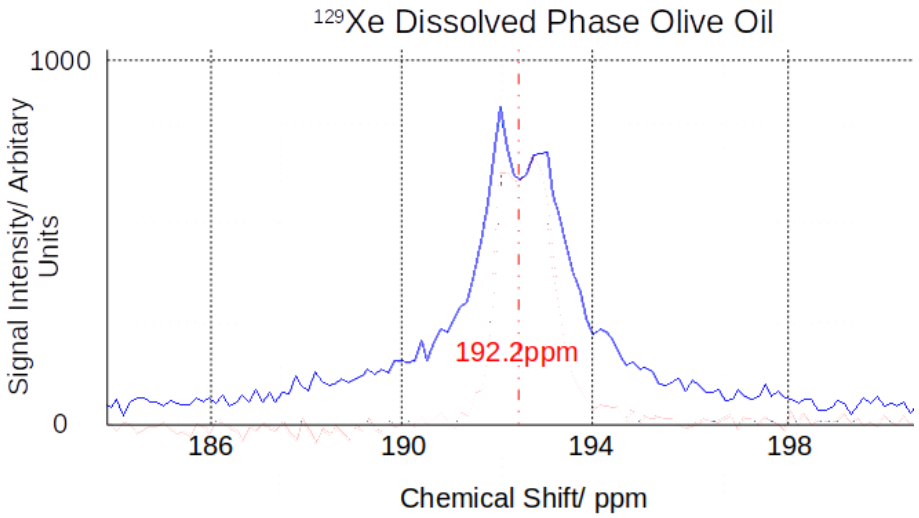

Figure 7.9: Split peak from ^{129}Xe dissolved phase in olive oil.

7.1.3 3D Spoiled GRE ^{129}Xe Sequence

A ^{1}H 3D spoiled GRE imaging sequence was performed using the clinical interface on the Paramed. This acted as the base for the ^{129}Xe sequence. The observed frequency and the frequency of the RF pulse was changed, from 21.114MHz, to the resonant frequency of ^{129}Xe in the scanner's magnetic field (5.842MHz). The amplitude profile of the RF pulse used in both sequences can be seen in figure 7.10.

The FOV of the ^{1}H sequence was ∼30cmx30cmx40cm. Provided the timings of the sequence are maintained, the gradient amplitudes can simply be scaled by the ratio of the ^{129}Xe to ^{1}H gyromagnetic ratios (the Xe gradients should be higher). Provided the gradient amplitude doesn't go over 100 (equating to 20mTm^{-1}), the maximum slew rate won't be exceeded because the rise time of all the gradients is set to 0.6ms. However, if the timings of the gradients is altered, it should be the area, not the amplitude, that is scaled. Lengthening the gradient's duration can allow a smaller FOV to be achieved, while keeping the gradient amplitude to <100 and reducing the gradient's duration can result in a shorter duration sequence.

To alter the FOV in the frequency encoding dimension, the readout gradient is scaled by the ratio of the original dimension size to the new dimension size (smaller FOV has higher gradient). Alternatively, the gradient strength can be calculated by the process outlined in section 2.1.10. To alter the FOV in the phase encoding dimension, the phase encoding gradient step size is scaled by the ratio of the original dimension size to the new dimension size (smaller FOV has larger gradient step size). Alternatively, the gradient strength can be calculated by the process outlined in section 2.1.11.

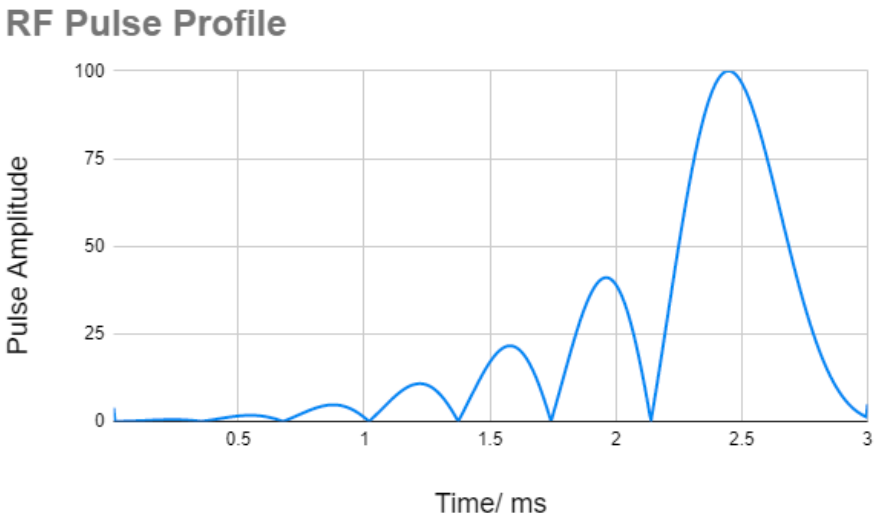

Figure 7.10: RF pulse amplitude profile used in the ^{129}Xe 3D GRE sequence. Also previously used in the Paramed ^{1}H 3D GRE sequence.

For the sequence to be used in humans, it has to be acquired in a single breath hold $\approx 20s$. The minimum number of phase points of the proton sequence in the clinical interface was 128x24. Even using the minimum matrix size and TR, the sequence took >1 minute to run. By reducing this to 128x8, the sequence takes 20s. To reduce the number of phase encoding points, but maintain the FOV, the phase encoding step size was maintained and the same centre of the phase encoding points was used.

Optimal Flip Angle

The 3D sequence has 128x8=1024 RF pulses. The polarised signal will decay according to equation 7.2. In the limit that T1>>TR, this equation becomes equation 7.5, where k is the pulse number and S_k is the signal of the kth pulse.

$$S_k = S_1(cos(\alpha))^{k-1} \tag{7.5}$$

This means that even a 3° FA results in $S_{1024} = 0.25S_1$ and 6° results in $S_{1024} = 0.0036S_1$. A 6° FA would result in insufficient signal for the final RF pulses.

The low-high, fully sampled, Cartesian acquisition pattern results in the first and last parts of the sequence providing the majority of information about the fine structure and the middle part of the sequence the most information about the signal intensity. For FA=6° $S_{512} = 0.060385S_1$; already by the middle part of the sequence the signal is depleted by over an order of magnitude, so the

final image will have little signal intensity. If the k-space was acquired centre out however, the image would be very bright, but there would be little fine structure. A projection sampling scheme would very useful in this scenario. Because each line of k space goes through the centre they all gain an equal amount of information about the outer edges and centre of k-space, thus one isn't favoured over the other.

If an accurate FA calibration has been performed, then the signal of each line of k-space can be scaled by equation 7.5.

Alternatively the FA (α) can be varied so that the signal available from each pulse is of the same magnitude. Looking at equation 7.1, in the limit TR>>T1, equation 7.7 is reached. By setting these equal to each other, equation 7.11 is found.

$$S_k = M_0 sin\alpha_k \prod_{j=1}^{k-1} cos\alpha_j \tag{7.6}$$

$$S_{k+1} = M_0 sin\alpha_{k+1} \prod_{j=1}^{k} cos\alpha_j \tag{7.7}$$

For : k $\geq$ 1 :$S_k = S_{k+1}$:

$$sin\alpha_{k+1} \prod_{j=1}^{k} cos\alpha_j = sin\alpha_k \prod_{j=1}^{k-1} cos\alpha_j, \tag{7.8}$$

$$sin\alpha_{k+1} cos\alpha_k \prod_{j=1}^{k-1} cos\alpha_j = sin\alpha_k \prod_{j=1}^{k-1} cos\alpha_j, \tag{7.9}$$

$$sin\alpha_{k+1} = tan\alpha_k \tag{7.10}$$

$$tan\alpha_k = \frac{sin\alpha_1}{\prod_{j=1}^{k} cos\alpha_j} \tag{7.11}$$

$$\alpha_k = sin^{-1}\left(\frac{sin\alpha_1}{\prod_{j=1}^{k-1} cos\alpha_j}\right) \tag{7.12}$$

However, if $sin\alpha_1$ is chosen to be too large, later FAs may exceed $\pi/2$. In the limit that $\alpha_N \rightarrow \pi/2$, where N is the last RF pulse equation 7.13 is reached [157].

$$\alpha_k = tan^{-1}\left(\frac{1}{\sqrt{N-k}}\right) \tag{7.13}$$

By setting the FA of the kth pulse to equation 7.13, this should normalise the signal across the image, while ensuring the maximum possible signal is extracted from the polarised ^{129}Xe.

Sequence Testing

FA Optimisation A 1L 13.5% polarisation bag of Xe was produced and placed in the scanner. A Tupperware phantom (see figure 7.11), was evacuated. The Tupperware was then placed inside the scanner on top of saline bags which were on-top of the coil, thus simulating a lung (the Tupperware) surrounded by other tissue (saline bags). This is similar to position 2 in the FA calibration with the Tupperware instead of a Tedlar bag (see figure 7.4). The bag of Xe and the Tupperware were connected before Xe was allowed to flow inside the Tupperware. The phantom was then imaged using the ^{129}Xe surface coil and an 8 slice 3D spoiled GRE sequence.

For simplicity, a constant RF attenuation was used throughout the scan. As a starting point, a 6° FA was employed, with the plan to increase the FA and see if the image quality improved. Ideally, a lower FA such as 2.5° would be used, which is the amplitude of the middle RF pulse in the variable FA scheme. However, the FA of the sequence will vary within the phantom and regions of the phantom closer to the coil may have a higher FA and regions further away a lower one; despite the RF attenuation being consistent throughout the scan. Therefore, to ensure a sufficient FA at the extremities of the coil's FOV, a higher FA was chosen. Looking at table 7.2, the "attenuation gap 2dB" fit has a higher R^2 and was therefore used as an approximation of FA. An attenuation of 40dB should produce $\sim 6°$ FA. Immediately after the RF attenuation 40dB sequence, an identical sequences with RF attenuation 35dB ($\sim 11°$) and 30dB ($\sim 18°$) were ran. The resulting images can be seen in figure 7.11.

The attenuation 40dB sequence should have reduced the polarisation to $\sim$0% after 1024 RF pulses. However, the fact it was possible to image the bag twice more at lower attenuations, suggests there are regions of the phantom where the FA was much lower than 6°. The attenuation 35dB sequence provided the highest SNR, despite being performed on a bag which had already been depolarised by 1024 RF pulses. The attenuation 30dB image had lower SNR than the attenuation 35dB image; this could be because the FA was too high, however it could just be because the bag had been significantly depolarised by the previous scans. Even if the ^{129}Xe closest to the coil is depolarised on the first scan, polarised ^{129}Xe will diffuse into this region. On all the images, the indent at the four corners of the phantom could be seen in the 5th slice. This is a feature on the order of 1cm. Finally, the attenuation 30dB image shows a small amount of signal at the top and bottom of the image (outside the Tupperware). This is believed to be from the ^{129}Xe bag just outside the FOV and still with a small amount of ^{129}Xe left in it. Aliasing is observed with the signal wrapping around and appearing at the top and bottom of the image. The attenuation 30dB RF pulse was enough to depolarise the bag >15cm from the coil. For future work with the ^{129}Xe coil and a constant FA scheme, the attenuation 40dB was decided upon. While the SNR may not be the ultimate for the whole Tupperware, it is high enough. This should ensure there is sufficient polarisation for the sequence to be completed, without a large section close to the coil being depolarised fully before all the RF pulses have completed. Consequently, this

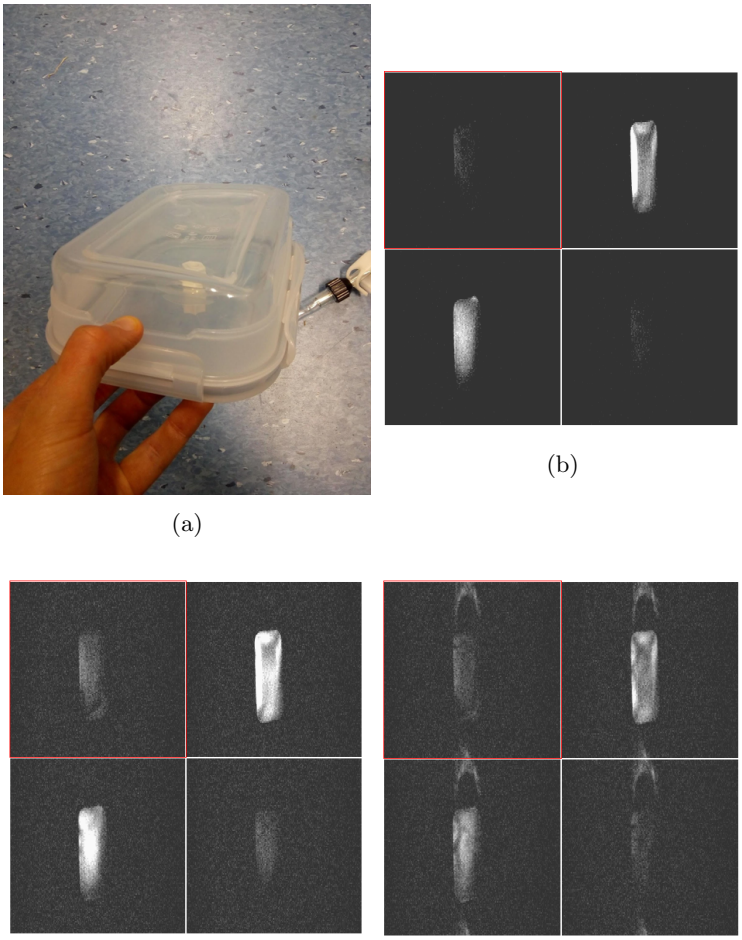

Figure 7.11: **(a)** Filled Tupperware phantom: **(a)** A larger depression can be seen in one corner. **(b)**,**(c)** and **(d)** show slices 3, 4, 5 and 6 from an eight slice, 3D spoiled GRE sequence. **(a)** Run one (attenuation 40dB, SNR 18.43), **(b)** run two (attenuation 35dB, SNR 21.58) and **(c)** run three (attenuation 30dB, SNR 11.22).

should allow fine detail to be resolved.

Fine Structure Capabilities In order to test the ability of the sequence to resolve fine structure, two ^{129}Xe voids were placed in the Tupperware. A glue stick and sealed empty glass vial were placed inside, before evacuating the Tupperware (see figure 7.12).

The images produced can be seen in figure 7.13. The glue stick and glass vial voids can easily be distinguished. In fact, in slice 6, the white plastic base of the glue stick which can be twisted is visible (see figure 7.13b). This is a feature of $\sim$1mm.

Next, a plastic tubing with an internal diameter of 3mm was used as a phantom. The tubing was first open at both ends. A bag of Xe with 14% polarisation was then connected to the tubing and Xe pushed through. Both ends were sealed with plastic clips and the tubing placed in the same location as the Tupperware previously. The same sequence was ran and the resulting image can be seen in figure 7.14.

Summary

A 3D GRE sequence has been shown to produce a high SNR image in a phantom and be able to observe features of <3mm. The location of the phantom was chosen to mimic a human lung and the FOV was similar to the FOV that would be required for adult lung imaging. After the final implementation of an SAR monitoring system, the hope is to produce in vivo images using the surface coil.

A bespoke chest coil has been ordered from PulseTeq Ltd. for ^{129}Xe imaging. The new chest coil should provide a more consistent FA over the full FOV and therefore allow for a more accurate FA calibration. The 3D GRE sequence will then be optimised, using a variable FA; with the FA obeying equation 7.13.

Figure 7.12: Tupperware with a glue stick and glass vial filled with air inside to act as ^{129}Xe phantom.

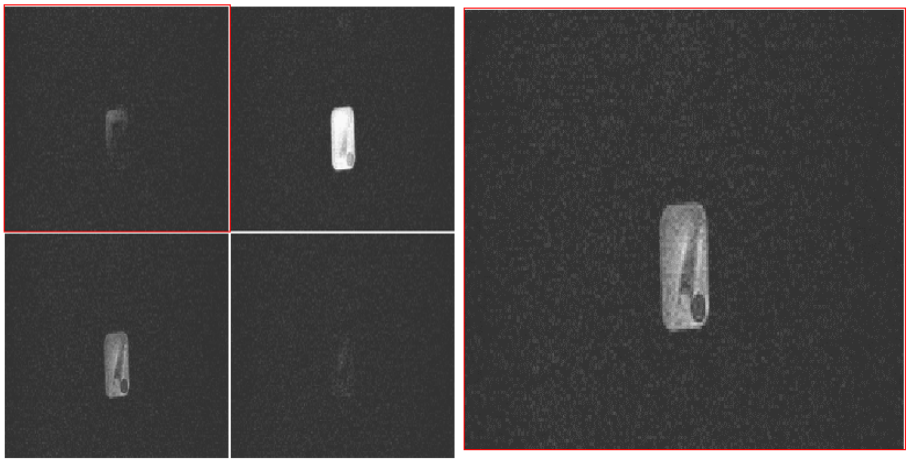

(a) (b)

Figure 7.13: Attenuation 40dB, eight slice 3D spoiled GRE slices of ^{129}Xe phantom **(a)** Slices 4, 5, 6 and 7. **(b)** Slice 6, the glue stick and glass vial voids can clearly be distinguished.

7.2 ^{19}F Feasibility on the Upright Scanner

Through the inhalation of a hydrofluorocarbon, ^{19}F MRI allows for the assessment of ventilation properties; without need for hyperpolarisation [10] [158]. See section 2.4 for background on ^{19}F MRI.

By exploiting the short T1 of ^{19}F in hydrofluorocarbon gases, multiple repeats can be acquired within a short duration. In vivo hydrofluorocarbon gas MRI has previously always been performed with the volunteer horizontal in a conventional scanner. However, lung function and volume, as well as arterial oxygen levels are affected by body position. By lying supine, lung function is reduced and for patients at the latter stages of chronic conditions, such as COPD, many simply can't lie in a supine position for an extended period of time. By imaging volunteers upright, the hope is that the images will be more relevant for clinicians and the volunteers will be more comfortable during scanning.

^{19}F and ^{1}H have very similar gyromagnetic ratios (see table 2.1). In early work done at 0.1T, the difference in Larmor frequency was small enough (and the response of ^{1}H coils used was wide enough), that ^{19}F could be imaged using proton coils with only small adaptations [60]. However, at the magnetic field strength of the Paramed scanner $\sim$ 0.5T the Larmor frequency of ^{19}F is 19.8642MHz compared to proton at 21.1160MHz. This gap of $>$ 1MHz is outside the bandwidth of the Paramed proton transmit body coil. Due to the power dividers used by the Paramed proton transmit body coil, it only has a bandwidth of $\sim$200kHz.

Before buying a dedicated coil tuned to ^{19}F. The feasibility of ^{19}F imag-

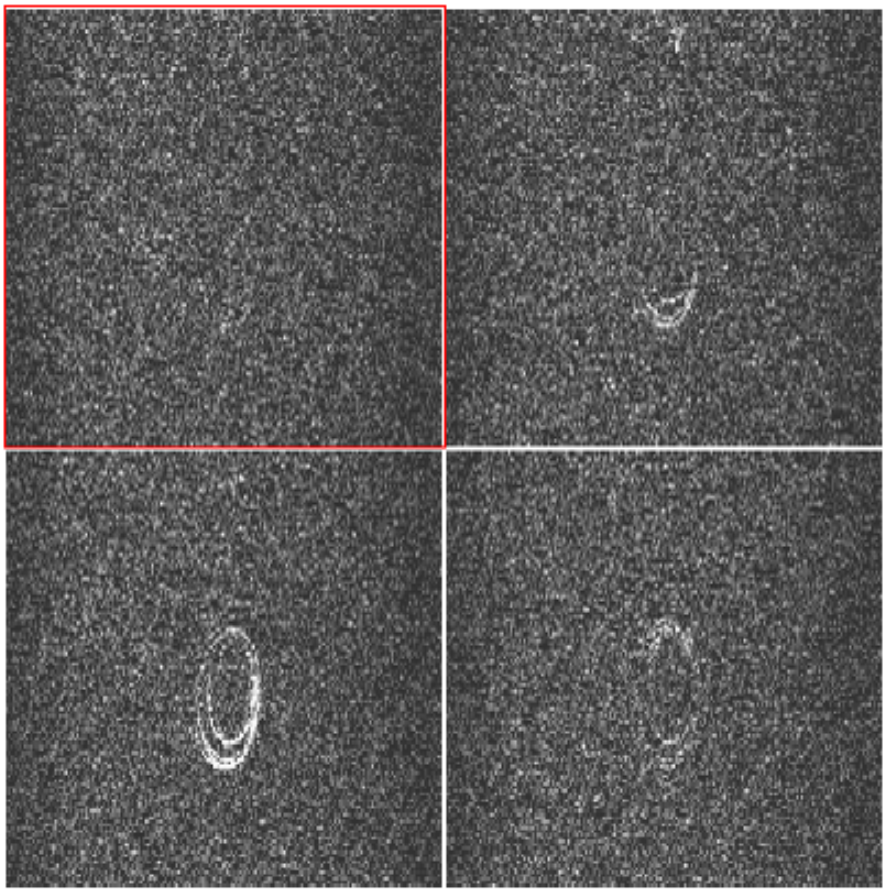

Figure 7.14: Attenuation 40dB, eight slice, 3D spoiled GRE, slices 3, 4, 5 and 6, of ^{129}Xe phantom. 3mm internal diameter plastic tubing used as ^{129}Xe phantom.

Figure 7.15: FLEX SMALL coil used for imaging. The red circle indicates the location of the hard plastics.

ing was assessed, through imaging ^{1}H on a hydrofluorocarbon gas at standard temperature and pressure.

7.2.1 Method

An Electrolube GDP 400 air duster, normally used for cleaning dust from computer fans/ keyboards provided the hydrofluorocarbon gas. This is cheap and contains 60-95% 1,1,1,2-Tetrafluoroethane (HFC) and 5-10% Dimethylether [159].

A 10L Tedlar bag was part filled with $\sim 5L$ of gas from the air duster. From initial pulse acquire tests on the bag, the T2* was seen to be very short; in order to obtain a sufficient signal for imaging, the TE of the sequence would have to be short. The FLEX SMALL coil was used for imaging because it has fairly little hard plastics, when compared with the other Paramed coils. The hard plastics that do form part of the casing are circled in figure 7.15. The T2* of these hard plastics is very short and consequently they provide very little signal at echo times typically used in imaging. However, at short echo times hard plastics will start to offer some signal.

In the 2D GRE imaging sequence used, all gradients were set to ramp up and down in 0.6ms; to reduce the TE as much as possible. The RF duration was

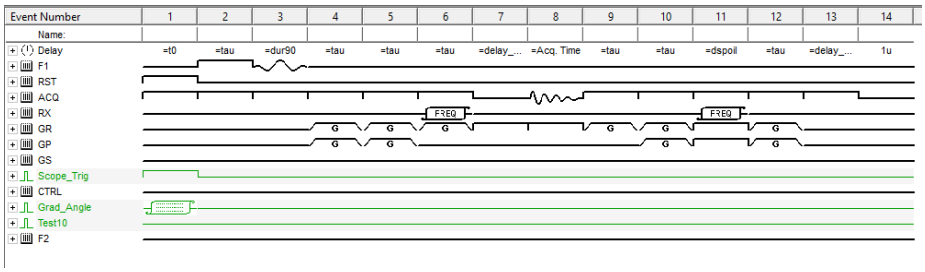

Figure 7.16: GRE sequence used for imaging HFC.

also reduced from 3ms to 0.75ms. The resulting TE was 3.532ms. To increase the signal from each voxel. a large voxel size was used and an infinite slice (no slice selecting gradients) was employed. 128 frequency encodings and 32 phase encodings were used. See figure 7.16 for sequence profile.

Neal et al. [10] work on ^{19}F imaging was taken as a template when working out the feasibility of ^{19}F at 0.5T. Neal et al. imaged PFP at 3T with a TR of 7.5ms. At 3T, the T1 of ^{19}F on PFP in vivo is $\sim$ 12.4ms. After measuring the T1 of ^{1}H on HFC, the TR was calculated using the same ratio of T1 to TR as Neal et al. A flip angle calibration was also be performed to allow the Ernst angle to be used for imaging HFC.

Flip Angle Calibration

In order to keep the TE as short possible, the RF pulse would need to be short. With the requirement that the pulse can still achieve a 90° FA, the 3ms pulse optimised for a 3D-GRE sequence and used in section 7.1.3 wasn't suitable. This pulse has a lot of low intensity side lobes, which mean if the pulse duration is reduced, the RF pulse is unable to still achieve a 90° excitation. The 3ms pulse used in section 7.1.3 was altered (see figure 7.10). The first 5 side lobes were removed to allow for a higher tipping angle from a pulse of the same duration and amplitude. The pulse profile can be seen in figure 7.17. A 1.5ms duration RF pulse was used (rather than a 0.75ms duration as shown in figure 7.17) to ensure a 90° pulse could be achieved.

After placing a saline phantom in the FLEX SMALL coil, a range of pulse attenuations were swept through, from 15dB to 2dB, in 0.5dB increments (1.5ms pulse, TE=2.6ms). The gap between pulses was chosen to be 10s; hopefully $> 5T1$. The peak amplitude was reached at a TG=-6dB. An estimation of the FA for each pulse was made by setting $\alpha_{TGi} = \frac{Signal_{TGi}}{Signal_{TG-6}}$. Where α_{TGi}, $Signal_{TGi}$ and $Signal_{TG-6}$ are the FA of pulse i, the signal of pulse i and the signal of a pulse with $TG = -6dB$ respectively. The data with $TG \leq -6dB$ was then fitted to equation 7.3, allowing the TG to be extracted for a given FA (see figure 7.18).

The variable parameters for equation 7.3 can be found in Table 7.6 and give

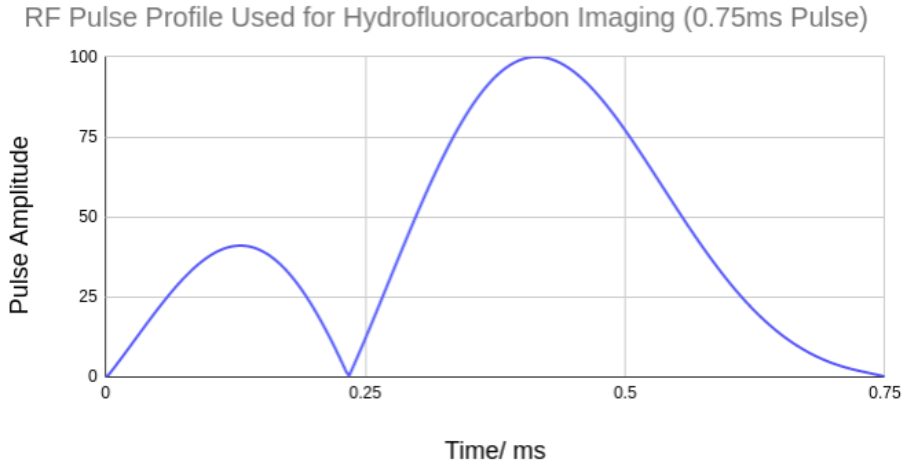

Figure 7.17: RF pulse profile used for hydrofluorocarbon imaging.

TG_π /dB	λ	R^2
-5.533 (-5.679, -5.388)	17.04 (16.6, 17.48)	0.997

Table 7.6: Fit parameters for FA calibration of a 1.5ms pulse on HFC 95% confidence bounds in brackets.

the relationship equation 7.14.

$$TG_\alpha = -5.533 - 17.04log(\frac{\pi}{\alpha}) \tag{7.14}$$

T1 of HFC at 0.5T

The saline phantom was then switched with the HFC phantom. Using two 1.5ms $\approx 90°$ RF pulses (duration 1.5ms attenuation 5.5dB). The drop in signal was measured for a range of TRs. The time between each set of RF pulses was 1 minute to allow for the ^{1}H nuclei to relax back fully. The amplitude of the difference in signal amplitude was fitted to an exponential $\propto e^{(-TR/T1)}$ (see figure 7.19). The fit has an R^2 of 0.998 and gave a T1 = 1600ms $\pm 62ms$, with a 95% confidence interval.

7.2.2 Results

From figure 7.19, the T1 was calculated to be 1600ms ±62ms with a 95% confidence interval. Using the same ratio of TR to T1 as Neal et al. [10]; for a T1=1600ms, the required TR is 968ms.

A TR of 968ms puts the Ernst angle; $\alpha_E \sim 57°$ (see equation 2.24). For a 57° pulse (using equation 7.14), the required TG would be -8.9dB.

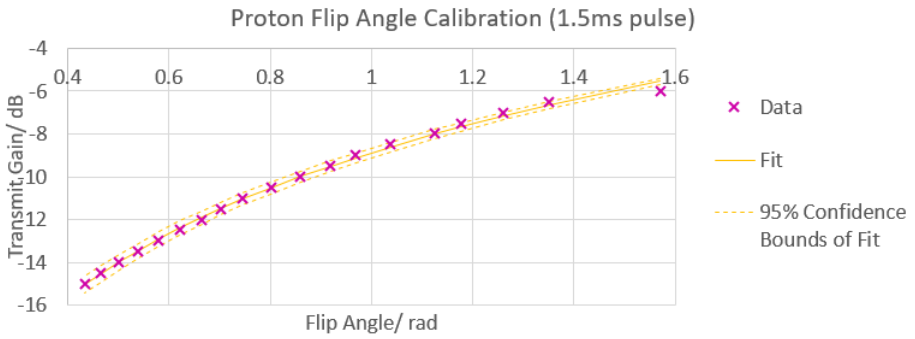

Figure 7.18: Pulse amplitude sweep of 1.5ms pulse. Blue line is the fit, crosses actual data.

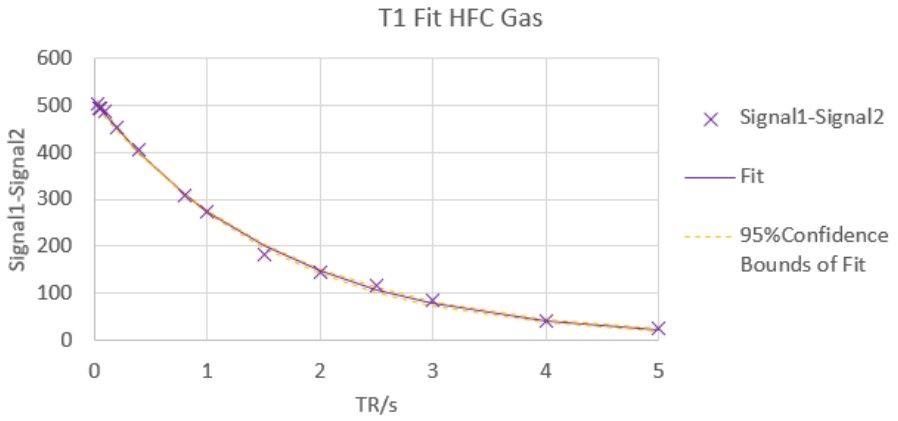

Figure 7.19: T1 fit for HFC gas (Proton)

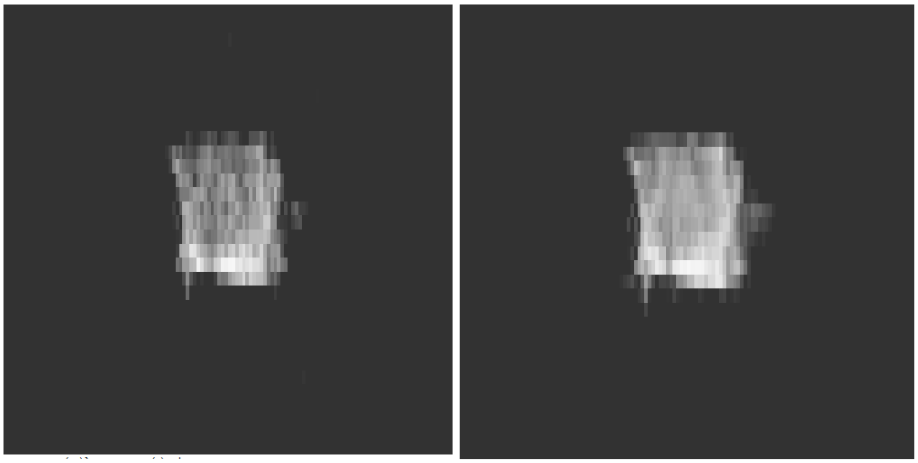

(a) (b)

Figure 7.20: 1mx1m FOV, 10L Tedlar bag: **(a)** 1 repetition SNR=18.67, **(b)** 10 repetitions SNR=56.76.

In order to reduce the TE (and consequently increase the acquired signal), the pulse duration was reduced from 1.5ms to 0.75ms. To double the power of the pulse and equate for this reduction in duration, the TG must be increased by 3dB. A TG= -5.9dB for a 1.5ms pulse with a pulse profile that of figure 7.17 was used for imaging.

The resulting images from imaging a ∼0.5L HFC filled 10L Tedlar bag can be seen in figures 7.20 and 7.21, and an image from a 1L HFC filled Tedlar bag in figure 7.22. To ensure the signal wasn't from the bag itself an empty 10L Tedlar bag was also imaged (see figure 7.23).

7.2.3 Discussion

In the 1L bag image (figure 7.23), a bright spot can be seen outside the bag in the lower left quadrant. This is also visible in the 10L bag images (both when empty/filled). This is believed to be the hard plastics which form the casing where the coil connects to itself (circled in red in figure 7.15); although it could also be the hard plastic clip used to seal the Tedlar bag. Both are made from hard plastics which tend to have a short T1 and hence don't interfere with most imaging. However, due to the short TE employed in the sequence used, the hard plastics register on these images. Alternatively, it could be a flaw in the coil design, that this region is more sensitive. However, this bright spot shouldn't be an issue when moving from the ^{1}H to ^{19}F resonant frequency. A new coil will be used and any hard plastics on the new coil will resonate at a different frequency to ^{19}F.

The fact that a bag of HFC could be imaged using a single scan, with the coil

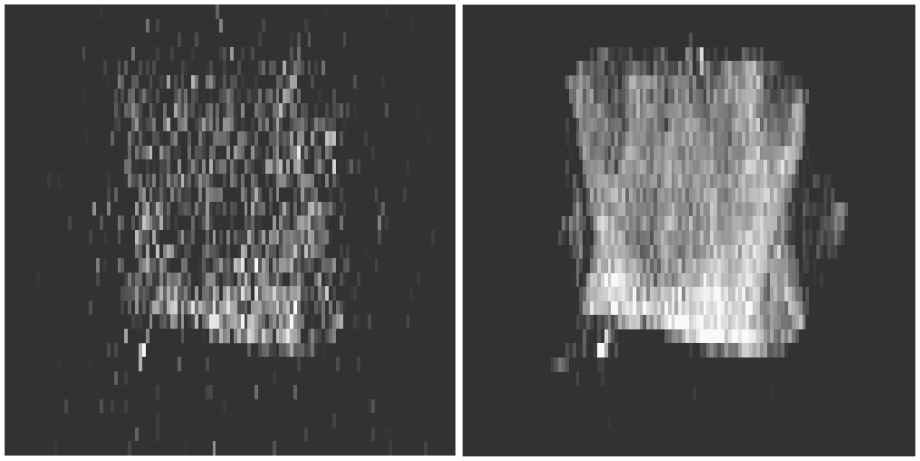

(a) (b)

Figure 7.21: 0.5mx0.5m FOV, 10L Tedlar bag: **(a)** 1 repetition SNR=5.259, **(b)** 10 repetitions SNR=17.89.

suboptimally loaded, was encouraging. The coil couldn't be loaded properly in the tests above, because the signal from loading the coil with a phantom would have dwarfed the signal from HFC; the coil response would have been far from ideal for ^{1}H imaging.

Also the signal is expected to improve when moving to ^{19}F. The higher SNR of the larger FOV images is due to the voxels being larger and consequently there is more HFC in each voxel. For 10 repetitions a ∼3 fold increase in signal was achieved over 1 repetition. This is as expected that the $SNR \propto \sqrt{repetitions}$ [14].

Neal et al. took images at 3T where a ∼2.5 fold improvement in signal would be expected when moving to the ^{19}F frequency with PFP(C_3F_8) from the ^{1}H frequency with HFC($C_2F_4H_2$). ^{19}F nuclei have a relative sensitivity of 0.84 [60], compared to ^{1}H. With HFC, for every molecule there are two ^{1}H and with PFP there are eight ^{19}F atoms. However, PFP has the structure CF_3-CF_2-CF_3 and the CF_3 resonances are at a slightly different frequency than the CF_2 resonances. This chemical shift between these peaks is ∼48ppm [158], ≈5.7kHz at 3T, meaning only the CF_3 resonances are excited and consequently only six ^{19}F excited per molecule; a 2.5 fold improvement in signal. However, at 0.5T the chemical shift will be (∼950Hz); six times smaller. This shift may be small enough to also excite the CF_2 part of each molecule and consequently the improvement in signal may be even greater.

The T_1 of ^{19}F on PFP in vivo, was found by Maunder Et. Al. [65], to be ∼11.9ms at 3.0T and ∼13.0ms at 1.5T. Taking an estimate of the T_1 to be 20ms at 0.5T and keeping the same ratio of T1 to TR; the TR would reduce from 968ms for ^{1}H imaging of HFC, to 12.1ms for ^{19}F imaging of PFP. This would

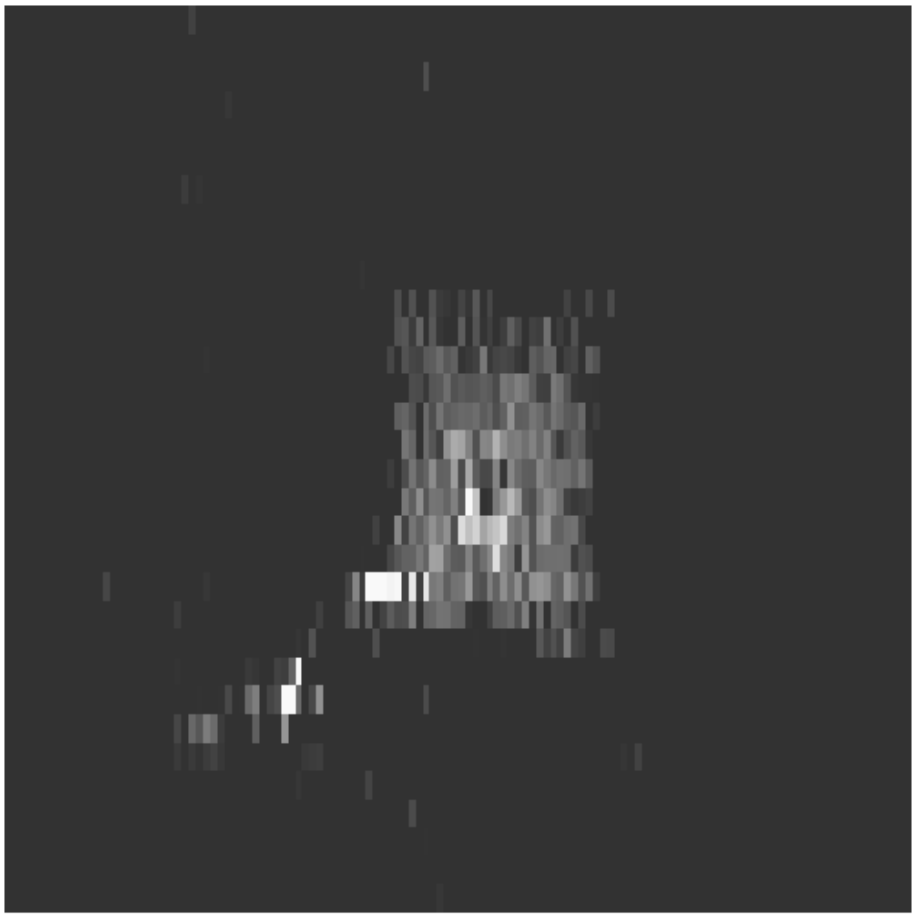

Figure 7.22: 0.5mx0.5m FOV, 1L Tedlar bag, 10 repetition, SNR=8.39.

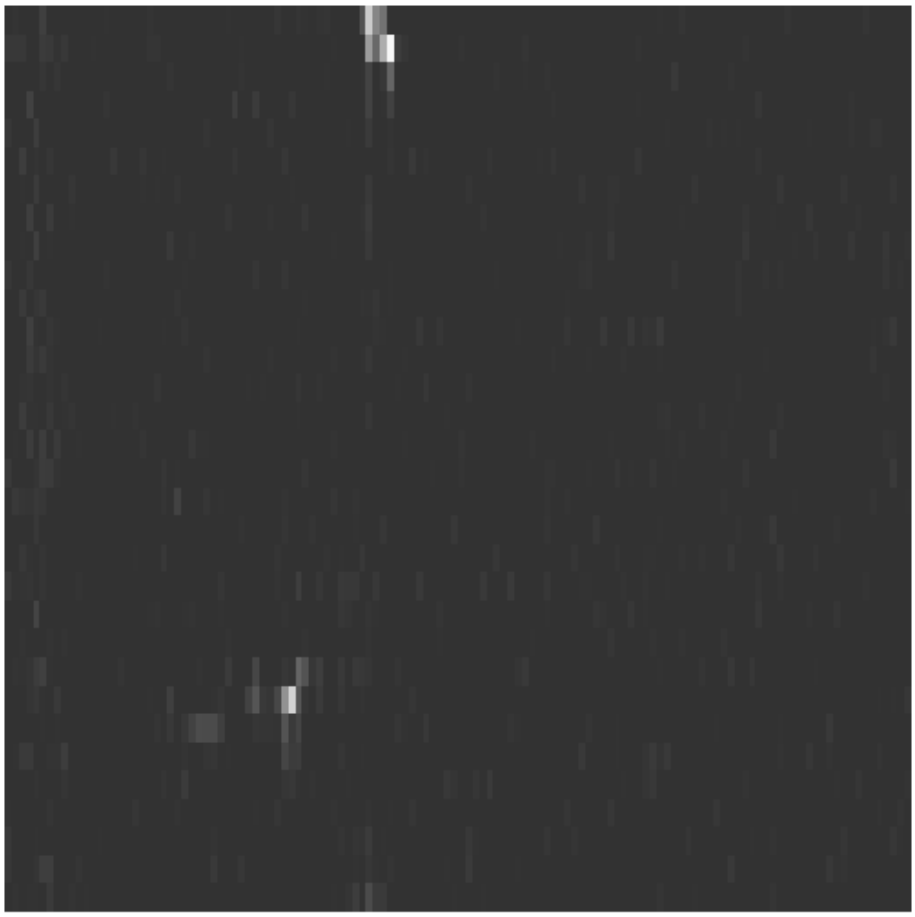

Figure 7.23: 0.5mx0.5m FOV, 10 scans, empty 10L Tedlar bag.

allow for the acquisition of a 3D GRE sequence with 128 phase encodings and 8 slices (the same as in section 7.1.3) in 12.4s; comfortably within a breath hold for healthy volunteers and many patients. Alternatively, a single 2D slice with 128 phase encodings could be acquired within 1.5s

7.2.4 Conclusion

Within a single scan, a ^{1}H image of HFC can be acquired. The signal is also expected to improve by at least 2.5 fold when moving to the ^{19}F frequency. PFP is the proposed gas for initial investigations. Due to the short T1 of PFP, this should allow for multiple averages of an in vivo image to be acquired within a breath hold.

From this preliminary work, a grant submission to MRC CiC has been submitted to support the implementation of a bespoke 0.5T ^{19}F chest coil and gas delivery system. There is the intention to compare ^{129}Xe and ^{19}F imaging in the future.

Chapter 8

Conclusion

Pulmonary MRI has excellent prospects for offering clinicians high resolution structural and functional data. The upright scanner at Nottingham allows for the scanning of volunteers at later stages of respiratory diseases and means the data produced should be of greater diagnostic use for clinicians. The design of the scanner makes it particularly suited to paediatric cohorts. Developing a passive technique to diagnose certain respiratory conditions would be very beneficial, because current pulmonary function tests depend heavily on the effort of the patient. Among paediatric cohorts this is a particular issue. Also, it is far from ideal to expose young patients to ionising radiation from a CT scan. The upright MRI scanner has allowed for children as young as three to be comfortable inside the scanner without need for sedation. The ionising radiation dose from CT also limits longitudinal respiratory imaging studies. By using MRI instead, the ionising radiation dose is avoided.

By doubling the polarisation of ^{129}Xe, the signal from ^{129}Xe MRI images should be twice as great. Although near unity polarisations have been achieved on a clinical scale by Nikolaou Et. Al. [35], these polarisations are achieved over timescales which are unacceptable if the technique was to be used routinely.

Through my work investigating a batch mode system, it was clear that the Rb needed to be spread in order to obtain consistent polarisation build-up curves for the same external conditions. Through Raman spectroscopy, the temperature inside the cell was also seen to be vastly elevated above the oven temperature during runaway; especially at the front of the cell where the laser light impinges upon first. Current polariser technology struggles to utilise the very fast rates of ^{129}Xe polarisation build up that occur at high temperatures and still cool the cell quickly enough to avoid runaway [130]. Clinically this would be useful, because it would allow for the production of high percentage polarised ^{129}Xe on time scales much shorter than currently available.

As expected, higher polarisation percentages of the Xe were observed in gas mixes with higher concentrations of N_2 and lower concentrations of Xe. However, the overall NMR signal was greatest for a mix with more N_2 than Xe, while still not being a very lean mix. This suggests that in a batch mode

system, rich Xe mixes are beneficial, but there is a balance to be struck with the N_2 concentration. In investigations into the buffer gas, no advantage was seen by replacing N_2 with ^{4}He. For a given Rb vapour density, the final polarisation achieved was broadly consistent across gas mixes with the same Xe concentration and different concentrations of N_2 and ^{4}He. Further investigation is needed to see if adding ^{4}He to the buffer gas does offer some benefits. However, because ^{4}He costs more than N_2, if N_2 performs just as well, there is little point using ^{4}He

^{129}Xe imaging is currently on hold on the 1.5T GE scanner, until the Multinuclear RF amplifier and SAR monitoring have been fixed. The new bespoke ^{129}Xe chest coil from Clinical MR Solutions L.L.C. was shown to be unfit for use, in vivo, in QTAR mode. In QUAD T/R mode, initial in vivo results have been promising. Other groups including Duke and Oxford university own a more simple version of the Clinical MR Solutions 1.5T ^{129}Xe coil which only works in QUAD T/R mode. Both groups have produced impressive results with the coil; producing high spatial [139] and temporal resolution [140] ventilation images.

The bespoke chest coil from Clinical MR Solutions L.L.C. makes ^{129}Xe imaging using the 1.5T much more comfortable for volunteers, than the Rapid Biomedical coil previously used. The new coil also allows for the scanning of larger people. However, long term the plan is to move ^{129}Xe imaging onto the 0.5T Paramed scanner. It is hoped the 1.5T will offer a stable platform and a benchmark to compare to in vivo ^{129}Xe images from the 0.5T Paramed scanner in the future.

The Paramed 0.5T scanner's upright, open design, makes it an interesting prospect for respiratory imaging in a wide variety of cohorts. In its stock format there are few sequences which can be utilised for respiratory imaging effectively. A bank of bespoke sequences and techniques tailored for respiratory imaging has started to be built.

A ^{1}H 2D GRE sequence has been developed on the upright. Using this sequence, the diaphragm was imaged at a time resolution of 463ms. In a video of the images, the diaphragm movement throughout the breathing cycle could clearly be characterised. The lung density average variation frequency, was also extracted at the same as the average breathing frequency. By moving to a radial acquisition in the future, the hope is to further improved the time resolution.

A stock SE sequence on the Paramed has also been modified, by inserting a hold at the start of the sequence. This reduced the duration of the scan from on average 35s to $\sim$21s. Inserting the hold also meant the scan is acquired within a much more consistent duration. The variability in duration was reduced from $>$ $\pm$5s to $\pm$0.5s. This made the breath hold requirements of a diaphragm imaging investigation into the morphology of the diaphragm, when upright versus supine, more acceptable.

Through CS, a 3D GRE sequence has been developed, which within a single breath hold (22.6s), can produce an image with a matrix size of 256x200x22; undersampling by a factor of 3. This has been used in vivo and has shown an accurate characterisation of the diaphragm. The GRE sequence can also be used

with a matrix size of 256x160x10, undersampled by a factor of 4.6, to produce an image within a duration of 5.1s; while still resolving the diaphragm. This opens the prospect for imaging volunteers with more progressed lung conditions. In the future, the plan is to apply the undersampling schemes to the ^{129}Xe 3D GRE sequence, where shorter scan durations, with higher resolutions would also be beneficial.

The multinuclear system on the Paramed makes it unique; being the only upright MRI scanner at high field (>0.1T), to offer imaging at the ^{129}Xe Larmor frequency. Before imaging in vivo, a bespoke SAR monitoring system needs to be produced. The monitoring circuit is being built and a calibration has been produced for the circuit.

A surface coil for the upright scanner was tuned to the correct frequency. Using this surface coil a flip angle calibration for ^{129}Xe and dissolved phase spectroscopy of ^{129}Xe in a variety of liquids has been performed. A 3D GRE sequence has also been developed, which within <20s can define structure of <3mm.

The feasibility of ^{19}F imaging on the upright has also been investigated. Images of a hydrofluorocarbon gas were obtained in a single repetition and there is the expectation of at least a 2.5 fold improvement in signal when moving from ^{1}H to ^{19}F. From this preliminary work a grant submission to MRC CiC has been submitted to support the implementation of a ^{19}F chest coil and gas delivery system.

Through this work, different respiratory MRI techniques have been developed, with a particular focus on volunteer acceptability. On the 1.5T flatbed scanner, investigations have been performed into using a ^{129}Xe coil from Clinical MR Solutions L.L.C. This bespoke coil is more comfortable for volunteers and will exclude fewer volunteers due to their size than the coil previously used. The open design of the 0.5T Paramed scanner makes it more comfortable for volunteers than a conventional flatbed scanner and the option to scan upright, makes certain breathing exercises required for respiratory MRI easier. However, the standard sequences on the Paramed aren't designed for respiratory imaging. Sequences have been developed, for both ^{1}H and ^{129}Xe MRI, which are tailored for a variety of respiratory imaging applications. These sequences were developed with the aim to exclude as few people as possible from imaging.

In the future, after an SAR monitoring circuit has been implemented, in vivo ^{129}Xe imaging will begin on the Paramed. To begin with, the surface coil will be used, but imaging will soon transfer to using a bespoke chest coil manufactured by PulseTeq Ltd. Ultimately, the hope is to apply the CS technique, already applied to a ^{1}H 3D GRE sequence, to a ^{129}Xe 3D GRE sequence. The images produced should be compared to in vivo images from the 1.5T scanner. After the implementation of a ^{19}F coil, there is the intention to compare ^{19}F imaging, to ^{129}Xe imaging on the upright.

Bibliography

[1] Michael Kreuter and Vincent Cottin. *The threat in chronic lung diseases: acute exacerbations*. 2017.

[2] Harpreet Ranu, Michael Wilde, and Brendan Madden. "Pulmonary function tests". In: *The Ulster medical journal* 80.2 (2011), p. 84.

[3] Takeshi Kubo et al. "Radiation dose reduction in chest CT: a review". In: *American journal of roentgenology* 190.2 (2008), pp. 335–343.

[4] Eugene Lin and Adam Alessio. "What are the basic concepts of temporal, contrast, and spatial resolution in cardiac CT?" In: *Journal of cardiovascular computed tomography* 3.6 (2009), pp. 403–408.

[5] Shigeru Kiryu et al. "Quantitative analysis of the velocity and synchronicity of diaphragmatic motion: dynamic MRI in different postures". In: *Magnetic resonance imaging* 24.10 (2006), pp. 1325–1332.

[6] Elizabeth Dean. "Effect of body position on pulmonary function". In: *Physical Therapy* 65.5 (1985), pp. 613–618.

[7] John Mugler and Talissa A Altes. "Hyperpolarized 129Xe MRI of the human lung". In: *Journal of magnetic resonance imaging : JMRI* 37 (Feb. 2013), pp. 313–31. DOI: 10.1002/jmri.23844.

[8] Justus E Roos et al. "Hyperpolarized gas MR imaging: technique and applications". In: *Magnetic Resonance Imaging Clinics* 23.2 (2015), pp. 217–229.

[9] Zackary I Cleveland et al. "Continuously infusing hyperpolarized 129Xe into flowing aqueous solutions using hydrophobic gas exchange membranes". In: *The Journal of Physical Chemistry B* 113.37 (2009), pp. 12489–12499.

[10] Mary A Neal et al. "Optimized and accelerated 19F-MRI of inhaled perfluoropropane to assess regional pulmonary ventilation". In: *Magnetic resonance in medicine* (2019).

[11] Wolfgang Günther Schreiber et al. "Dynamic 19F-MRI of pulmonary ventilation using sulfur hexafluoride (SF6) gas". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 45.4 (2001), pp. 605–613.

[12] David R Lide. *CRC handbook of chemistry and physics*. Vol. 85. CRC press, 2004.

[13] James Keeler. *Understanding NMR spectroscopy*. John Wiley & Sons, 2011.

[14] Matt A Bernstein, Kevin F King, and Xiaohong Joe Zhou. *Handbook of MRI pulse sequences*. Elsevier, 2004.

[15] P. C. Lauterbur. "Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance". In: *Nature* 242 (Mar. 1973), pp. 190–191. DOI: 10.1038/242190a0.

[16] Jeffrey L. Duerk and Orlando P. Simonetti. "Review of MRI gradient waveform design methods with application in the study of motion". In: *Concepts in Magnetic Resonance* 5.2 (1993), pp. 105–122. DOI: 10.1002/cmr.1820050202. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cmr.1820050202. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cmr.1820050202.

[17] John P Ridgway. "Cardiovascular magnetic resonance physics for clinicians: part I". In: *Journal of cardiovascular magnetic resonance* 12.1 (2010), p. 71.

[18] I. Solomon. "Relaxation Processes in a System of Two Spins". In: *Phys. Rev.* 99 (2 July 1955), pp. 559–565. DOI: 10.1103/PhysRev.99.559. URL: https://link.aps.org/doi/10.1103/PhysRev.99.559.

[19] N. Bloembergen and L. O. Morgan. "Proton Relaxation Times in Paramagnetic Solutions. Effects of Electron Spin Relaxation". In: *jcp* 34 (Mar. 1961), pp. 842–850. DOI: 10.1063/1.1731684.

[20] Danny Lee et al. "Dynamic keyhole: A novel method to improve MR images in the presence of respiratory motion for real-time MRI". In: *Medical physics* 41.7 (2014), p. 072304.

[21] Gary H Glover and John M Pauly. "Projection reconstruction techniques for reduction of motion effects in MRI". In: *Magnetic resonance in medicine* 28.2 (1992), pp. 275–289.

[22] Klaus Scheffler and Jürgen Hennig. "Reduced circular field-of-view imaging". In: *Magnetic resonance in medicine* 40.3 (1998), pp. 474–480.

[23] Michael G Strintzis. "A review of compression methods for medical images in PACS". In: *International journal of medical informatics* 52.1-3 (1998), pp. 159–165.

[24] Michael Lustig et al. "Compressed sensing MRI". In: *IEEE signal processing magazine* 25.2 (2008), p. 72.

[25] John Pauly et al. "Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging)". In: *IEEE transactions on medical imaging* 10.1 (1991), pp. 53–65.

[26] JMS Hutchinson, RJ Sutherland, and JR Mallard. "NMR imaging: image recovery under magnetic fields with large non-uniformities". In: *Journal of physics E: Scientific instruments* 11.3 (1978), p. 217.

[27] WA Edelstein et al. "The intrinsic signal-to-noise ratio in NMR imaging". In: *Magnetic resonance in medicine* 3.4 (1986), pp. 604–618.

[28] James Tropp. "The theory of the bird-cage resonator". In: *Journal of Magnetic Resonance (1969)* 82.1 (1989), pp. 51–62.

[29] Nicola De Zanche. "Birdcage Volume Coil Design". In: *eMagRes* (2012).

[30] Frederick H Epstein, John P Mugler III, and James R Brookeman. "Spoiling of transverse magnetization in gradient-echo (GRE) imaging during the approach to steady state". In: *Magnetic resonance in medicine* 35.2 (1996), pp. 237–245.

[31] Vincent Denolin, Céline Azizieh, and Thierry Metens. "New insights into the mechanisms of signal formation in RF-spoiled gradient echo sequences". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 54.4 (2005), pp. 937–954.

[32] Richard B Buxton et al. "Contrast in Rapid MR Imaging: T1- and T2-Weighted Imaging". In: *Journal of computer assisted tomography* 11.1 (1987), pp. 7–16.

[33] Erwin L Hahn. "Spin echoes". In: *Physical review* 80.4 (1950), p. 580.

[34] Bernd André Jung and Matthias Weigel. "Spin echo magnetic resonance imaging". In: *Journal of Magnetic Resonance Imaging* 37.4 (2013), pp. 805–817.

[35] Panayiotis Nikolaou et al. "Near-unity nuclear polarization with an open-source 129Xe hyperpolarizer for NMR and MRI". In: *Proceedings of the National Academy of Sciences* 110.35 (2013), pp. 14150–14155.

[36] Linus Pauling and Charles D Coryell. "The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin". In: *Proceedings of the National Academy of Sciences* 22.4 (1936), pp. 210–216.

[37] Kâmil Uludağ, Bernd Müller-Bierl, and Kâmil Uğurbil. "An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging". In: *Neuroimage* 48.1 (2009), pp. 150–165.

[38] Hiroto Hatabu et al. "T2* and proton density measurement of normal human lung parenchyma using submillisecond echo time gradient echo magnetic resonance imaging". In: *European journal of radiology* 29.3 (1999), pp. 245–252.

[39] José P Marques, Frank FJ Simonis, and Andrew G Webb. "Low-field MRI: An MR physics perspective". In: *Journal of magnetic resonance imaging* 49.6 (2019), pp. 1528–1542.

[40] PD Harris and R Barnes. “The uses of helium and xenon in current clinical practice”. In: *Anaesthesia* 63.3 (2008), pp. 284–293.

[41] Dana A Shea and Daniel L Morgan. *The helium-3 shortage: Supply, demand, and options for congress*. 2010.

[42] Adrian Cho. *Helium-3 shortage could put freeze on low-temperature research*. 2009.

[43] K W Miller et al. “Xenon NMR: chemical shifts of a general anesthetic in common solvents, proteins, and membranes”. In: *Proceedings of the National Academy of Sciences* 78.8 (1981), pp. 4946–4949. ISSN: 0027-8424. DOI: 10.1073/pnas.78.8.4946. eprint: https://www.pnas.org/content/78/8/4946.full.pdf. URL: https://www.pnas.org/content/78/8/4946.

[44] F William Hersman et al. “Large production system for hyperpolarized 129Xe for human lung imaging studies”. In: *Academic radiology* 15.6 (2008), pp. 683–692.

[45] S Sivaram Kaushik et al. “Diffusion-weighted hyperpolarized 129Xe MRI in healthy volunteers and subjects with chronic obstructive pulmonary disease”. In: *Magnetic resonance in medicine* 65.4 (2011), pp. 1154–1165.

[46] Michael H Abraham et al. “Solubility properties in polymers and biological media. 2. The correlation and prediction of the solubilities of nonelectrolytes in biological tissues and fluids”. In: *Journal of medicinal chemistry* 28.7 (1985), pp. 865–870.

[47] Kai Ruppert et al. “NMR of hyperpolarized 129Xe in the canine chest: spectral dynamics during a breath-hold”. In: *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo* 13.4 (2000), pp. 220–228.

[48] Mark E Wagshul et al. “In vivo MR imaging and spectroscopy using hyperpolarized 129Xe”. In: *Magnetic resonance in medicine* 36.2 (1996), pp. 183–191.

[49] Galina E Pavlovskaya et al. “Hyperpolarized krypton-83 as a contrast agent for magnetic resonance imaging”. In: *Proceedings of the National Academy of Sciences* 102.51 (2005), pp. 18275–18279.

[50] Thomas Meersmann, Scott A Smith, and Geoffrey Bodenhausen. “Multiple-quantum filtered xenon-131 NMR as a surface probe”. In: *Physical review letters* 80.7 (1998), p. 1398.

[51] Thomas Meersmann, Michaël Deschamps, and Geoffrey Bodenhausen. “Probing Aerogels by Multiple Quantum Filtered 131Xe NMR Spectroscopy”. In: *Journal of the American Chemical Society* 123.5 (2001), pp. 941–945.

[52] Igor L Moudrakovski, Christopher I Ratcliffe, and John A Ripmeester. “131Xe, a new NMR probe of void space in solids”. In: *Journal of the American Chemical Society* 123.9 (2001), pp. 2066–2067.

[53] Yannick Millot et al. “Quantification of electric-field gradient in the supercage of Y zeolites with 131Xe (I= 3/2) NMR”. In: *Comptes Rendus de l'Académie des Sciences-Series IIC-Chemistry* 4.11 (2001), pp. 815–818.

[54] Charlene F Horton-Garcia, Galina E Pavlovskaya, and Thomas Meersmann. “Introducing krypton NMR spectroscopy as a probe of void space in solids”. In: *Journal of the American Chemical Society* 127.6 (2005), pp. 1958–1962.

[55] Oleg G Salnikov et al. “Clinical-Scale Batch-Mode Production of Hyperpolarized Propane Gas for MRI”. In: *Analytical chemistry* 91.7 (2019), pp. 4741–4746.

[56] Aurélien Bornet and Sami Jannin. “Optimizing dissolution dynamic nuclear polarization”. In: *Journal of Magnetic Resonance* 264 (2016), pp. 13–21.

[57] Klaes Golman et al. “Molecular imaging with endogenous substances”. In: *Proceedings of the National Academy of Sciences* 100.18 (2003), pp. 10435–10439.

[58] Seung-Hyun Kim et al. “Carbon isotopes profiles of human whole blood, plasma, red blood cells, urine and feces for biological/biomedical 14C-accelerator mass spectrometry applications”. In: *Analytical chemistry* 83.9 (2011), pp. 3312–3318.

[59] SJ Nelson et al. “DNP-hyperpolarized 13 C magnetic resonance metabolic imaging for cancer applications”. In: *Applied magnetic resonance* 34.3-4 (2008), pp. 533–544.

[60] Ramachandran Murugesan et al. “Fluorine electron double resonance imaging for 19F MRI in low magnetic fields”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 48.3 (2002), pp. 523–529.

[61] Jesús Ruiz-Cabello et al. “Fluorine (19F) MRS and MRI in biomedicine”. In: *NMR in Biomedicine* 24.2 (2011), pp. 114–129.

[62] Ursula Wolf et al. “Subsecond fluorine-19 MRI of the lung”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 55.4 (2006), pp. 948–951.

[63] Dean O Kuethe et al. “Imaging lungs using inert fluorinated gases”. In: *Magnetic resonance in medicine* 39.1 (1998), pp. 85–88.

[64] Marcus J Couch et al. “Pulmonary ultrashort echo time 19F MR imaging with inhaled fluorinated gas mixtures in healthy volunteers: feasibility”. In: *Radiology* 269.3 (2013), pp. 903–909.

[65] Adam Maunder et al. “Optimization of steady-state free precession MRI for lung ventilation imaging with 19F C3F8 at 1.5 T and 3T”. In: *Magnetic resonance in medicine* 81.2 (2019), pp. 1130–1142.

[66] Herbert Y Reynolds. “Diagnostic and management strategies for diffuse interstitial lung disease”. In: *Chest* 113.1 (1998), pp. 192–202.

[67] Allan L Coates et al. “Spirometry in primary care”. In: *Canadian respiratory journal* 20.1 (2013), 1321. ISSN: 1198-2241. DOI: 10.1155/2013/615281. URL: http://europepmc.org/articles/PMC3628641.

[68] J Wanger et al. “Standardisation of the measurement of lung volumes”. In: *European respiratory journal* 26.3 (2005), pp. 511–522.

[69] CP Criée et al. “Body plethysmography–its principles and clinical use”. In: *Respiratory medicine* 105.7 (2011), pp. 959–971.

[70] JE Cotes et al. *Standardization of the measurement of transfer factor (diffusing capacity)*. 1993.

[71] Martin Raymond Miller et al. “General considerations for lung function testing”. In: *European Respiratory Journal* 26.1 (2005), pp. 153–161.

[72] *Radiation Dose in X-Ray and CT exams. 2016. Radiology info for patients*. URL: RadiologyInfo.org.

[73] Yoshikazu Yamao et al. “CT–fluoroscopy in chest interventional radiology: Sliding scale of imaging parameters based on radiation exposure dose and factors increasing radiation exposure dose”. In: *Clinical radiology* 68.2 (2013), pp. 162–166.

[74] Bruno-Pierre Dubé and Martin Dres. “Diaphragm dysfunction: diagnostic approaches and management strategies”. In: *Journal of clinical medicine* 5.12 (2016), p. 113.

[75] Z Mosenifar et al. “Chronic obstructive pulmonary disease (COPD)”. In: *Medscape, USA, Ημερομηνία Πρόσβασης* 18.5 (2016).

[76] Ron Balkissoon et al. “Chronic obstructive pulmonary disease: a concise review”. In: *Medical Clinics* 95.6 (2011), pp. 1125–1141.

[77] Anne M Slavotinek. “Single gene disorders associated with congenital diaphragmatic hernia”. In: *American Journal of Medical Genetics Part C: Seminars in Medical Genetics*. Vol. 145. 2. Wiley Online Library. 2007, pp. 172–183.

[78] Benny Weksler and Robert J Ginsberg. “Tumors of the diaphragm.” In: *Chest surgery clinics of North America* 8.2 (1998), pp. 441–447.

[79] Colin Alexander. “Diaphragm movements and the diagnosis of diaphragmatic paralysis”. In: *Clinical radiology* 17.1 (1966), pp. 79–83.

[80] M Orozco-Levi. “Structure and function of the respiratory muscles in patients with COPD: impairment or adaptation?” In: *European Respiratory Journal* 22.46 suppl (2003), 41s–51s.

[81] Aarti Sarwal, Francis O Walker, and Michael S Cartwright. “Neuromuscular ultrasound for evaluation of the diaphragm”. In: *Muscle & nerve* 47.3 (2013), pp. 319–329.

[82] JM Heath and R Mongia. “Chronic bronchitis: primary care management.” In: *American family physician* 57.10 (1998), pp. 2365–72.

[83] Bartolome R Celli et al. “Standards for the diagnosis and treatment of patients with COPD: a summary of the ATS/ERS position paper”. In: *European Respiratory Journal* 23.6 (2004), pp. 932–946.

[84] Timothy F. Murphy and Sanjay Sethi. “Bacterial Infection in Chronic Obstructive Pulmonary Disease”. In: *American Review of Respiratory Disease* 146.4 (1992). PMID: 1416398, pp. 1067–1083. DOI: 10.1164/ajrccm/146.4.1067. eprint: https://doi.org/10.1164/ajrccm/146.4.1067. URL: https://doi.org/10.1164/ajrccm/146.4.1067.

[85] Ann J Woolcock. “How does inflammation cause symptoms?” In: *American journal of respiratory and critical care medicine* 153.6 Pt 2 (1996), S21.

[86] Amir Sharafkhaneh, Nicola A Hanania, and Victor Kim. “Pathogenesis of emphysema: from the bench to the bedside”. In: *Proceedings of the American Thoracic Society* 5.4 (2008), pp. 475–477.

[87] Rachna Madan, Thanissara Chansakul, and Hilary J Goldberg. “Imaging in lung transplants: checklist for the radiologist”. In: *The Indian journal of radiology & imaging* 24.4 (2014), p. 318.

[88] Jane C. Davies, Eric W. F. F. Alton, and Andrew Bush. *Cystic fibrosis*. Dec. 2007. DOI: 10.1136/bmj.39391.713229.AD.

[89] JA Dodge et al. “Cystic fibrosis mortality and survival in the UK: 1947–2003”. In: *European Respiratory Journal* 29.3 (2007), pp. 522–526.

[90] PB Davis. “Cystic fibrosis since 1938.” In: *American Journal of Respiratory and Critical Care Medicine* 173(5) (Mar. 2006), pp. 475–82.

[91] Hirotoshi Matsui et al. “Evidence for periciliary liquid layer depletion, not abnormal ion composition, in the pathogenesis of cystic fibrosis airways disease”. In: *Cell* 95.7 (1998), pp. 1005–1015.

[92] Beryl J Rosenstein, Garry R Cutting, et al. “The diagnosis of cystic fibrosis: a consensus statement”. In: *The Journal of pediatrics* 132.4 (1998), pp. 589–595.

[93] Kevin W Southern et al. “A survey of newborn screening for cystic fibrosis in Europe”. In: *Journal of Cystic Fibrosis* 6.1 (2007), pp. 57–65.

[94] J Tod Olin and Michael E Wechsler. “Asthma: pathogenesis and novel drugs for treatment”. In: *Bmj* 349 (2014), g5517.

[95] MI Asher et al. “International Study of Asthma and Allergies in Childhood (ISAAC): rationale and methods”. In: *European respiratory journal* 8.3 (1995), pp. 483–491.

[96] Matthew Masoli et al. “The global burden of asthma: executive summary of the GINA Dissemination Committee report”. In: *Allergy* 59.5 (2004), pp. 469–478.

[97] Leonie Kausel et al. “Childhood asthma and allergies in urban, semi-urban, and rural residential sectors in Chile”. In: *The Scientific World Journal* 2013 (2013).

[98] T Nicolai et al. “Urban traffic and pollutant exposure related to respiratory outcomes and atopy in a large sample of children”. In: *European respiratory journal* 21.6 (2003), pp. 956–963.

[99] Sally E Wenzel. “Asthma phenotypes: the evolution from clinical to molecular approaches”. In: *Nature medicine* 18.5 (2012), p. 716.

[100] TE King. “Idiopathic pulmonary fibrosis: diagnosis and treatment. International consensus statement”. In: *Am J Respir Crit Care Med* 161 (2000), pp. 646–664.

[101] British Lung Foundation. *Idiopathic pulmonary fibrosis statistics*.

[102] Kazuro Iwai et al. “Idiopathic pulmonary fibrosis. Epidemiologic approaches to occupational exposure.” In: *American journal of respiratory and critical care medicine* 150.3 (1994), pp. 670–675.

[103] Kathy B Baumgartner et al. “Cigarette smoking: a risk factor for idiopathic pulmonary fibrosis.” In: *American journal of respiratory and critical care medicine* 155.1 (1997), pp. 242–248.

[104] Richard Hubbard et al. “Occupational exposure to metal or wood dust and aetiology of cryptogenic fibrosing alveolitis”. In: *The Lancet* 347.8997 (1996), pp. 284–289.

[105] ID Johnston et al. “British Thoracic Society study of cryptogenic fibrosing alveolitis: current presentation and initial management. Fibrosing Alveolitis Subcommittee of the Research Committee of the British Thoracic Society.” In: *Thorax* 52.1 (1997), pp. 38–44.

[106] PHILIPPE Grenier et al. “Chronic diffuse interstitial lung disease: diagnostic value of chest radiography and high-resolution CT.” In: *Radiology* 179.1 (1991), pp. 123–132.

[107] Peter B Crino, Katherine L Nathanson, and Elizabeth Petri Henske. “The tuberous sclerosis complex”. In: *New England Journal of Medicine* 355.13 (2006), pp. 1345–1356.

[108] S Harari, O Torre, and J Moss. “Lymphangioleiomyomatosis: what do we know and what are we looking for?” In: *European Respiratory Review* 20.119 (2011), pp. 034–044.

[109] EV Krjukov, JD ONeill, and JR Owers-Bradley. “Brute force polarization of 129 Xe”. In: *Journal of low temperature physics* 140.5-6 (2005), pp. 397–408.

[110] Hayley Louise Newton and MSci Hons Dunelm. “Hyperpolarised Xenon Production Via Rb and Cs Optical Pumping Applied to Functional Lung MRI.” PhD thesis. University of Nottingham, 2014.

[111] Anna Nikiel et al. “Metastability exchange optical pumping of 3 He at high pressure and high magnetic field for medical applications”. In: *The European Physical Journal Special Topics* 144.1 (2007), pp. 255–263.

[112] F Sage, JP Lemoigne, and D Lecler. "Optical pumping of the 2p53s3P0 State of 21Ne (I= 32): g factor and metastability exchange cross section". In: *Optics Communications* 30.3 (1979), pp. 332–336.

[113] Guilhem Collier. "Metastability exchange optical pumping (meop) of 3he in situ". PhD thesis. Jagiellonian University, 2011.

[114] Boyd M Goodson et al. "Spin-exchange optical pumping at high xenon densities and laser fluxes: Principles and practice". In: *Hyperpolarized Xenon-129 magnetic resonance*. Royal Society of Chemistry: Abingdon, 2015, pp. 96–121.

[115] W Happer et al. "Polarization of the nuclear spins of noble-gas atoms by spin exchange with optically pumped alkali-metal atoms". In: *Physical Review A* 29.6 (1984), p. 3092.

[116] Alfred Kastler. "Optical methods of atomic orientation and of magnetic resonance". In: *JOSA* 47.6 (1957), pp. 460–465.

[117] FD Colegrove, LD Schearer, and GK Walters. "Polarization of He 3 gas by optical pumping". In: *Physical Review* 132.6 (1963), p. 2561.

[118] Wenjin Shao, Guodong Wang, and Emlyn W Hughes. "Measurement of spin-exchange rate constants between Xe 129 and alkali metals". In: *Physical Review A* 72.2 (2005), p. 022713.

[119] U Ruth et al. "Production of nitrogen-free, hyperpolarized 129Xe gas". In: *Applied Physics B* 68.1 (1999), pp. 93–97.

[120] MA Antonacci et al. "Depolarization of nuclear spin polarized 129Xe gas by dark rubidium during spin-exchange optical pumping". In: *Journal of Magnetic Resonance* 279 (2017), pp. 60–67.

[121] Thomas Meersmann and Eike Brunner. *Hyperpolarized xenon-129 magnetic resonance: concepts, production, techniques and applications*. 4. Royal Society of Chemistry, 2015.

[122] MV Romalis, E Miron, and GD Cates. "Pressure broadening of Rb D 1 and D 2 lines by 3 He, 4 He, N 2, and Xe: Line cores and near wings". In: *Physical Review A* 56.6 (1997), p. 4569.

[123] Graham Norquay et al. "Optimized production of hyperpolarized 129Xe at 2 bars for in vivo lung magnetic resonance imaging". In: *Journal of Applied Physics* 113.4 (2013), p. 044908.

[124] Sergey E Korchak, Wolfgang Kilian, and Lorenz Mitschang. "Configuration and performance of a mobile 129 Xe polarizer". In: *Applied magnetic resonance* 44.1-2 (2013), pp. 65–80.

[125] Meghan E Halse. *Terranova-MRI EFNMR Student Guide*. 2009.

[126] Panayiotis Nikolaou et al. "Multidimensional mapping of spin-exchange optical pumping in clinical-scale batch-mode 129Xe hyperpolarizers". In: *The Journal of Physical Chemistry B* 118.18 (2014), pp. 4809–4816.

[127] DK Walter, WM Griffith, and W Happer. "Energy transport in high-density spin-exchange optical pumping cells". In: *Physical review letters* 86.15 (2001), p. 3264.

[128] RS Hickman and LH Liang. "Rotational temperature measurement in nitrogen using Raman spectroscopy". In: *Review of Scientific Instruments* 43.5 (1972), pp. 796–799.

[129] Samuel Patz et al. "Hyperpolarized 129Xe MRI: a viable functional lung imaging modality?" In: *European journal of radiology* 64.3 (2007), pp. 335–344.

[130] Jonathan Birchall et al. "Batch-Mode clinical-Scale optical hyperpolarization of xenon-129 using an aluminum jacket with rapid temperature ramping". In: *Analytical chemistry* (Submitted for review).

[131] *Safety Guidelines for Magnetic Resonance Imaging Equipment in Clinical Use*. v4.2. Mar. 2015.

[132] *RF Amplifier Data Sheet, BT02000-AlphaS, 100kHz-30MHz, 2kW*. DS006671B, Accessed at `http://www.tomcorf.com/assets/standard_pulsed/DS006671B.pdf`. Dec. 2018.

[133] *dBm - Voltage Chart*. Accessed at `http://www.tecmag.com/mt-content/uploads/2017/11/dbm_v.pdf`. Nov. 2017.

[134] *RF Amplifier Factory Performance Certificate, BT02000-AlphaS*. Serial No. 18545. Mar. 2016.

[135] *Amplifier Parallel Interface Specification*. T002968, Issue K, Accessed at `www.tomcorf.com/pdf/interface.pdf`. Aug. 2018.

[136] James Thorpe. "Lung mechanics and hyperpolarised gas MRI". PhD thesis. University of Nottingham, 2018.

[137] Shahideh Safavi. "Magnetic resonance imaging in respiratory medicine". PhD thesis. University of Nottingham, 2019.

[138] Zackary I Cleveland et al. "Hyperpolarized 129Xe MR imaging of alveolar gas uptake in humans". In: *PloS one* 5.8 (2010).

[139] Bastiaan Driehuys et al. "Chronic obstructive pulmonary disease: safety and tolerability of hyperpolarized 129Xe MR imaging in healthy volunteers and patients". In: *Radiology* 262.1 (2012), pp. 279–289.

[140] Mitchell Chen et al. "Delayed ventilation assessment using fast dynamic hyperpolarised Xenon-129 magnetic resonance imaging". In: *European radiology* 30.2 (2020), pp. 1145–1155.

[141] Grzegorz Bauman et al. "Non-contrast-enhanced perfusion and ventilation assessment of the human lung by means of fourier decomposition in proton MRI". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 62.3 (2009), pp. 656–664.

[142] M Deimling et al. “Time resolved lung ventilation imaging by Fourier decomposition”. In: *Proceedings of the 16th Annual Meeting of ISMRM*. 2008, p. 2639.

[143] Alberta Iliff and Virginia A Lee. “Pulse rate, respiratory rate, and body temperature of children between two months and eighteen years of age”. In: *Child development* (1952), pp. 237–245.

[144] Neville D Gai, Ashkan A Malayeri, and David A Bluemke. “Three-dimensional T1 and T2* mapping of human lung parenchyma using interleaved saturation recovery with dual echo ultrashort echo time imaging (ITSR-DUTE)”. In: *Journal of Magnetic Resonance Imaging* 45.4 (2017), pp. 1097–1104.

[145] Brian Avants. *ANTs in R*. ANTsR v1.0 `https://www.rdocumentation.org/packages/ANTsR/versions/1.0`. Feb. 2012.

[146] L Mendes Pereira et al. “UTE-SENCEFUL: first results for 3D high-resolution lung ventilation imaging”. In: *Magnetic resonance in medicine* 81.4 (2019), pp. 2464–2473.

[147] André Fischer et al. “SElf-gated Non-Contrast-Enhanced FUnctional Lung imaging (SENCEFUL) using a quasi-random fast low-angle shot (FLASH) sequence and proton MRI”. In: *NMR in Biomedicine* 27.8 (2014), pp. 907–917.

[148] JA d’Arcy et al. “Applications of sliding window reconstruction with cartesian sampling for dynamic contrast enhanced MRI”. In: *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo* 15.2 (2002), pp. 174–183.

[149] Shuo Zhang, Kai Tobias Block, and Jens Frahm. “Magnetic resonance imaging in real time: advances using radial FLASH”. In: *Journal of Magnetic Resonance Imaging* 31.1 (2010), pp. 101–109.

[150] C Arthofer et al. “Impact assessment of posture and breath-hold state on diaphragm shape”. In: *Magn Reson Med conference paper* (2019).

[151] Dominique Brunet, Edward R Vrscay, and Zhou Wang. “On the mathematical properties of the structural similarity index”. In: *IEEE Transactions on Image Processing* 21.4 (2011), pp. 1488–1499.

[152] Allister Mason et al. “Comparison of Objective Image Quality Metrics to Expert Radiologists Scoring of Diagnostic Quality of MR Images”. In: *IEEE transactions on medical imaging* (2019).

[153] Martin Uecker et al. “Berkeley advanced reconstruction toolbox”. In: *Proc. Intl. Soc. Mag. Reson. Med.* Vol. 23. 2486. 2015.

[154] Jianping Zhong et al. “Fast Determination of Flip Angle and T 1 in Hyperpolarized Gas MRI During a Single Breath-Hold”. In: *Scientific reports* 6 (2016), p. 25854.

[155] Graham Norquay et al. “Relaxation and exchange dynamics of hyperpolarized 129Xe in human blood”. In: *Magnetic resonance in medicine* 74.2 (2015), pp. 303–311.

[156] Keith W Miller et al. “Xenon NMR: chemical shifts of a general anesthetic in common solvents, proteins, and membranes”. In: *Proceedings of the National Academy of Sciences* 78.8 (1981), pp. 4946–4949.

[157] Lei Zhao et al. *Gradient-echo imaging considerations for hyperpolarized129Xe MR*. 1996.

[158] Marcus J Couch et al. “Inert fluorinated gas MRI: a new pulmonary imaging modality”. In: *NMR in biomedicine* 27.12 (2014), pp. 1525–1534.

[159] *Safety Data Sheet, High Powered Air Duster, GDP, EGDP400, ZE*. Rev. 1. May 2016.