Advanced Spectroscopic Imaging Techniques for X-Nuclei Cardiac Magnetic Resonance Imaging

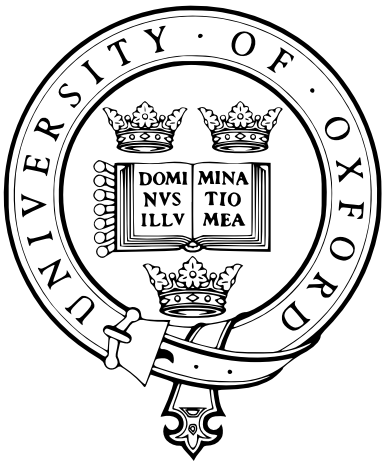

Andrew Tyler

St Edmund Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2021

To Margherita

Acknowledgements

Firstly, I would like to sincerely thank my excellent supervisors Damian Tyler, Justin Lau, Jack Miller and Laci Valkovič. I have learned a huge amount from all of you, and greatly appreciate all the time and energy you have put into encouraging me, imparting your incredible knowledge to me, and teaching me how to be a good researcher.

I am hugely grateful to Peter Jezzard for conceiving and running the ONBI CDT program, as well as the rest of the team and everyone who taught the DTC courses. I would also like to thank my Transfer and Confirmation examiners Liz Tunnicliffe and Karla Miller, whose feedback and advice was invaluable.

Thank you to everyone at CMRG and OCMR, particularly Aaron Hess, Adrienne Siu, Andy Apps, Andy Lewis, Cat Rooney, Charles Hill, Charlie Millard, Claire Trumper, Dragana Savic, Ferenc Mozes, Gabriela Belsley, James Grist, Jane Ellis, Kerstin Timm, Liam Young, Liz Tunnicliffe, Moritz Hundertmark, Paul Bottomley, Steph Anderson, Tony Zhou, Vicky Ball, and Will Watson. I have greatly enjoyed working with you and your suggestions and comments in group meetings have been very useful.

Thank you to my housemates at 50 Windmill Road – I've made friends for life and I'm so grateful for all the cups of red bush tea, TV nights, trips, and the introduction to my (soon to be) wife.

A particular thank you to my family for all of your love, encouragement and support over the last 26 years.

Finally I would like to thank Margherita, whose love and support, through all of the ups and downs of the last three years, has been more than I could have ever wished for.

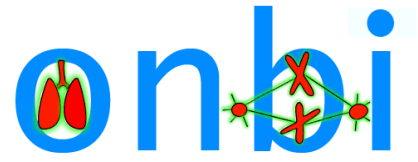

I would like to thank all the organizations which contributed financially to my research, particularly, the Medical Research Council, the Engineering and Physical Sciences Research Council and the British Heart Foundation. Without their support this work would not have been possible.

Abstract

Changes to cardiac metabolism are an important component of heart failure. Understanding these metabolic changes is therefore critical to developing treatments for the failing heart, particularly pharmaceuticals, or other interventions, which aim to restore normal metabolism. A combination of hyperpolarized ^{13}C and ^{31}P magnetic resonance spectroscopic imaging (MRSI) enables quantification of both the substrate usage (^{13}C) and energy output (^{31}P) of cardiac metabolism, allowing a full picture of the metabolic state of the heart to be built up.

Both hyperpolarized ^{13}C and ^{31}P MRSI present significant experimental challenges, which I will address, through the use of advanced acquisition and reconstruction techniques, in this thesis.

The main challenges associated with hyperpolarized ^{13}C MRSI are the short transverse relaxation time of the metabolites, necessitating very rapid readouts, the short half-life of hyperpolarization, which requires the full scan to complete in approximately two minutes, and the bolus injection of hyperpolarized material to the patient, which makes the signal intensity highly dynamic. To address these three challenges, I implemented a hybrid multiple/single-shot spiral sequence, with a rapid readout, which enabled the usual trade-off, inherent to multiple-shot spiral readouts, of temporal and spatial resolution (when selecting the number of shots), to be made at reconstruction, rather than acquisition time.

The principal challenge of ^{31}P MRSI however, is achieving adequate SNR. The low SNR of ^{31}P spectroscopy necessitates significant averaging during the acquisition, which extends scan time and precludes high spatial resolution imaging within a reasonable scan time. In this work I implement two compartmentalized reconstruction techniques, SLAM and SLIM, and apply them at both 3 and 7T, before investigating acquisition parameters, which lead to an improvement in SNR and repeatability over the existing technique used in OCMR.

Together these technological advances have extended our ability to quantify metabolism, which will, through their clinical application, enhance our understanding of heart failure.

Research Outcomes

Conferences

- BC-ISMRM 2018 (talk)- A 3D Hybrid Shot Spiral for Hyperpolarised 13 C Imaging
- ISMRM 2019 (poster)- A 3D Hybrid Shot Spiral for Hyperpolarised 13 C Imaging
- BC-ISMRM 2019 (power pitch)- Multifrequency Interpolation for Hyperpolarised B_0 Inhomogeneity Correction
- PG-BC-ISMRM 2019 (poster)- A 3D Hybrid Shot Spiral for Hyperpolarised 13 C Imaging
- PG-BC-ISMRM 2020 (talk)- Compartmentalised 31 P Spectroscopy at 7T in the Heart - a Feasibility Study
- ISMRM 2021 (poster)- Compartmentalised 31 P Spectroscopy at 7T in the Heart - a Feasibility Study

Academic publications

- A 3D hybrid-shot spiral sequence for hyperpolarised 13 C imaging, **Tyler** et al. Magn Reson Med. 2020; 85:790–801.
- Rapid B_1-insensitive, dual-band quasi-adiabatic saturation transfer with optimal control for complete quantification of myocardial ATP flux, Miller et al. Magn Reson Med. 2021; 85:2978–2991
- Proof-of-Principle Demonstration of Direct Metabolic Imaging Following Myocardial Infarction Using Hyperpolarized 13 C CMR, Apps et al. JACC:CVI 2021; https://doi.org/10.1016/j.jcmg.2020.12.023.
- Compartmentalized Reconstruction of 3D Acquisition Weighted ^{31}P Cardiac CSI at 7T - a Reproducibility Study, **Tyler** et al. In revision for Magn. Reson. Med.

Ethics Statement

All pre-clinical experiments were performed in accordance with relevant UK legislation (with personal, project and institutional licences granted under the Animals (Scientific Procedures) Act 1986), and were subject to local ethical review and an independent cost-benefit analysis.

All human participants (healthy volunteers and patients) gave informed consent for the studies contained in this thesis, which have been approved by The University of Oxford Central University Research Ethics Committee, and are in compliance with the Declaration of Helsinki.

Contents

List of Figures

List of Abbreviations

ATP . . . . . Adenosine triphosphate

AW . . . . . Acquisition weighted

BART . . . . . Berkeley Advanced Reconstruction Toolbox

bpm . . . . . Beats per minute

B_0 . . . . . Amplitude of static magnetic field

B_1 . . . . . Amplitude of RF excitation field

CFT . . . . . Continuous Fourier transform

CoA . . . . . Coenzyme A

CoR . . . . . Coefficient of reproducibility

CoV . . . . . Coefficient of variation

CSI . . . . . Chemical-shift imaging

CT . . . . . Computerized tomography

DC . . . . . Direct current

DFT . . . . . Discrete Fourier transform

DICOM . . . . Digital Imaging and Communications in Medicine

emf . . . . . Electromotive force

EPI . . . . . Echo planar imaging

FDG . . . . . Fluorodeoxyglucose

FFT . . . . . Fast Fourier transform

FOV . . . . . Field of view

FWHM . . . . Full-width half-maximum

GRE . . . . . Gradient echo sequence

HSS . . . . . Hybrid-shot spiral

LV . . . . . Left ventricle

MRI . . . . . Magnetic resonance imaging

NMR . . . . . Nuclear magnetic resonance

NUFFT . . . . Non-uniform fast Fourier transform

OCMR . . . . Oxford Centre for Clinical Magnetic Resonance Research

PCr . . . . . Phosphocreatine

PDE . . . . . Phosphodiester

PDH . . . . . Pyruvate dehydrogenase

PET . . . . . Positron emission tomography

ppm . . . . . Parts per million

PSF . . . . . Pointspread function

RF . . . . . Radio-frequency

RV . . . . . Right ventricle

SD . . . . . Standard deviation

SLAM . . . . . Magnetic resonance spectroscopy with linear algebraic modeling

SLIM . . . . . Spectral localization by imaging

SLOOP . . . . Spectral localization with optimal pointspread function

SNR . . . . . Signal-to-noise ratio

SPECT . . . . Single photon emission computed tomography

SRF . . . . . Spatial response function

SSIM . . . . . Structural similarity index

TE . . . . . Echo time

TR . . . . . Repetition time

T_1 . . . . . Longitudinal relaxation time

T_2* . . . . . Effective transverse relaxation time

UTE . . . . . Ultra-short echo time

2,3-DPG . . . 2,3-diphosphoglycerate

1 Introduction

Contents

1.1 Overview

Since the first clinically useful magnetic resonance imaging (MRI) images were obtained by John Mallard in 1980[1], the technique has found widespread use in medicine, revolutionizing clinical care and medical research. This pioneering work, as well as the majority of modern medical scans to date, utilized proton magnetic resonance, however even since the early days, the potential of x (non proton) nuclei spectroscopy has been apparent, with ^{31}P magnetic resonance spectroscopy (MRS) used to assess the effect of treatment, following myocardial infarction, as early as 1981 in rabbits[2].

Since then a wide range of x-nuclei MRI techniques have been investigated including ^{23}Na, ^{35}Cl, ^{17}O, and ^{129}Xe[3], however, this work will focus on hyperpolarized ^{13}C and natural abundance ^{31}P cardiac MRS. Together, the combination of these two techniques can provide information on metabolic substrate usage in the heart, and the production of ATP, both of which are critical to understanding the progression of heart failure. Unfortunately however, these x-nuclei techniques are technologically demanding, and each has unique issues which still need to be addressed.

In vivo ^{13}C MRS imaging of metabolic substrates requires the injection of a 'hyperpolarized' bolus of ^{13}C labeled metabolites, which produces $> 10,000\times$[4] the MRS signal compared to unlabeled metabolites. This is for two reasons, firstly the low absolute SNR of ^{13}C MRS of metabolic substrates, due to their low

concentrations in the body (relative to water) and the low natural abundance of ^{13}C, and secondly the low concentration of metabolic substrates relative to everything else containing carbon e.g. protein, fat etc. Unfortunately, the short half life of hyperpolarization, and bolus injection, gives rise to two challenges. Firstly the whole experiment must be completed in approximately two minutes, and secondly the scan must have sufficient temporal resolution to capture the dynamics of the bolus, while retaining enough spatial resolution to resolve spatially localized metabolic heterogeneities, such as regions of infarction.

^{31}P MRS also suffers from low SNR, due to the low concentration of phosphorus nuclei, and low gyromagnetic ratio of ^{31}P. As a result, this technique only became feasible for human use with the introduction of high-field 1.5T MRI scanners and localized spectroscopic techniques such as DRESS[5]. Since then, the SNR of ^{31}P MRS has successfully been improved through developments to acquisition technique and scanner design, particularly with the introduction of 3T, and more recently 7T scanners. Spectral SNR, however, still remains a fundamental limitation of ^{31}P MRS, which can lead to variation in measured results.

This body of work details my efforts to address these related challenges using advanced image acquisition and reconstruction techniques. In this chapter I begin by discussing our current understanding of the pathophysiological basis of heart failure, which the techniques developed in this thesis will be used to investigate, then discuss the theoretical underpinnings of MRI and MRS, before providing an over-view of the theory, challenges, and uses of hyperpolarized ^{13}C and ^{31}P cardiac MRS.

Chapter 2 describes my work to implement and evaluate a novel pulse sequence for hyperpolarized [1-^{13}C]pyruvate metabolic imaging, which utilizes a dual density spiral readout, where either a single shot can be reconstructed at low spatial resolution, or multiple shots can be added together to produce a higher spatial resolution image. This has the significant advantage of enabling the trade-off between spatial and temporal resolution to be made at reconstruction time, rather than acquisition time.

Then in Chapter 3 I investigate compartmentalized reconstruction techniques for ^{31}P cardiac spectroscopy in the heart at 7T. These reconstruction techniques improve SNR by segmenting the field of view into homogeneous anatomical compartments such as heart, chest wall, and blood, then reconstructing a single spectra per compartment. These techniques had not been implemented at 7T before this work, and there were significant concerns about variation within compartments, due to B_1 in-homogeneity, which could have affected the reconstructed spectra. These concerns were addressed in Chapter 3 with simulations, phantom experiments, and a healthy volunteer repeatability study, which compared combinations of three acquisition schemes and two compartmentalized reconstruction algorithms, against the current CSI method in use at our center.

Based on the strong potential of the techniques shown at 7T, I then investigated compartmentalized reconstruction at 3T, with the results described in Chapter 4. The chapter begins with a series of pilot studies to optimize the protocol, investigating the use of saturation bands and acquisition scheme. It then details a repeatability study which validates the techniques in healthy volunteers and investigates the effect of cardiac gating. Finally the chapter concludes by comparing compartmentalized reconstructions of clinical research data to the standard CSI technique used in OCMR.

The thesis finishes with a final chapter drawing overall conclusions on the contributions and relevance of this body of work, before discussing the further questions posed by this work, and potential experiments to address them.

1.2 Cardiac metabolism

1.2.1 Metabolism in the healthy heart

The heart pumps blood around our circulatory system twenty four hours a day, for our entire lifetime, with even a small interruption causing severe disease or death. It must contract 60-100 times per minute while at rest, and faster during exercise. This requires a significant energy input in the form of ATP, to power the myosin-II conformational change which produces the contraction, and move

Ca^{2+} (which activates the myosin-II) in and out of the sarcoplasmic reticulum. To function normally, the heart hydrolyses approximately 0.5 μmol of ATP per gram of tissue per second and yet only contains 13 μmol/g of ATP+PCr (ATP is stored as PCr)[6] making the continual generation of ATP by metabolism critical.

The energy to power the function of the heart is provided by the hydrolysis of ATP according to

$$\mathrm{ATP} \rightleftharpoons \mathrm{ADP} + \mathrm{Pi}.$$

The Gibbs free energy for this reaction depends on the concentrations of ATP, ADP and Pi according to the equation[7]

$$\Delta \mathrm{G} = \Delta \mathrm{G}^0 - \mathrm{RT} \ln \left(\frac{[\mathrm{ADP}][\mathrm{Pi}]}{[\mathrm{ATP}]} \right) \tag{1.1}$$

where $\Delta \mathrm{G}^0$ is the standard Gibbs free energy of hydrolysis of ATP ($\Delta \mathrm{G}^0 = -30.5$ kJ/mol), square brackets indicate concentration, R is the molar gas constant and T is temperature.

For a given biochemical reaction to be favorable when coupled to ATP hydrolysis the total Gibbs free energy $\Delta \mathrm{G}_{\mathrm{tot}} = \Delta \mathrm{G}_{\mathrm{ATP}} + \Delta \mathrm{G}_{\mathrm{proc}}$ must be negative. This means that the greater the energy demand of a process (more +ve $\Delta \mathrm{G}_{\mathrm{proc}}$) the further towards ATP the equilibrium of ATP and ADP + Pi must be. This makes it critical that [ATP] is maintained, as if it drops too far, the ATP consuming processes in the heart begin to stall, preventing normal cardiac function.

The heart combats a decrease in $\Delta \mathrm{G}_{\mathrm{ATP}}$ in two ways. Firstly with an ATP buffer, in the form of phosphocreatine (PCr), which 'stores' ATP for use in short term demand spikes

$$\mathrm{ADP} + \mathrm{PCr} \rightleftharpoons \mathrm{ATP} + \mathrm{Cr}$$

and secondly by being highly metabolically flexible, consuming a mixture of fatty acids (70-90%) and glucose + lactate (10-30%) with a small contribution from ketone

bodies[6]. For example, this metabolic flexibility allows the heart to respond to the increased energy demand during exercise by increasing the utilization of glucose[8].

Both glucose and fatty acids are converted to acetyl CoA and produce ATP via the citric acid (TCA) cycle[9] and electron transport chain[10]. However, the oxygen efficiency of these processes are different, with fatty acid oxidation requiring more molecules of O_2 per unit of ATP synthesized, making the correct balance of metabolic substrate usage critical to the long term health of the heart.

1.2.2 Metabolic changes in the failing heart

Changes in substrate metabolism are both a cause and a consequence of heart failure[11, 12]. During heart failure, both fatty acid and glucose metabolite usage decreases, however in the initial stages, glucose metabolism is decoupled from PDH flux leading to an increase in pyruvate and lactate concentration[13]. Reduced substrate metabolism consequently leads to reduced ATP production in the myocardium, and therefore reduced PCr and increased ADP concentrations, leading to reduced ΔG of ATP hydrolysis. The eventual consequence of which is contractile disfunction and heart failure[14].

Given the significant changes to metabolism during heart failure, it is hypothesized that reversing these changes, through treatment, may be able to alleviate the condition[11], however to properly study these metabolic treatments it is essential to to quantify their effect on metabolism in vivo. ^{13}C or ^{31}P magnetic resonance imaging are techniques which may allow us to do so, by quantifying either PDH flux and pyruvate/lactate pool sizes, or high energy phosphate concentrations respectively (see Figure 1.1). Together these techniques allow us to measure both substrate usage and energy production of the myocardium, and build up a full understanding of the effect of a treatment on metabolism.

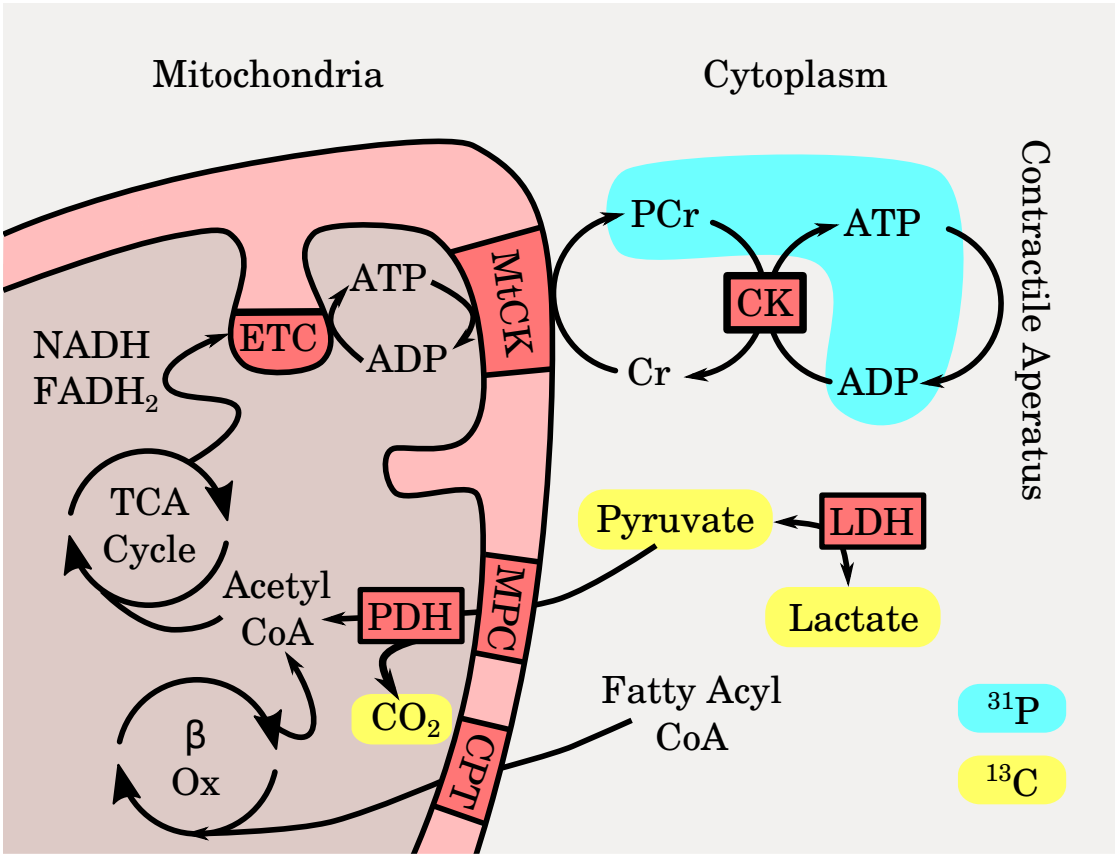

Figure 1.1: Summary of metabolism in the heart, with metabolites we can quantify through ^{13}C and ^{31}P MRI highlighted.

1.3 The theoretical basis of nuclear magnetic resonance

1.3.1 The NMR effect

In 1922, the Stern-Gerlach experiment[15] demonstrated the quantization of angular momentum. In their experiment they passed neutral silver atoms through a non-uniform magnetic field and noted that they were diverted into two distinct paths. They realised that the unpaired electron in the silver atom's d-orbital was taking one of two spin quantum states.

The allowed values of the spin quantum number for fermions are $s = \frac{n}{2}$ (where n takes integer values) and the spin projection quantum number can then take $m_s = -s, -s+1, ..., s$. Hence in the case of the silver atoms, with one unpaired spin-1/2 electron, $s = \frac{1}{2}$ and $m_s = \pm\frac{1}{2}$.

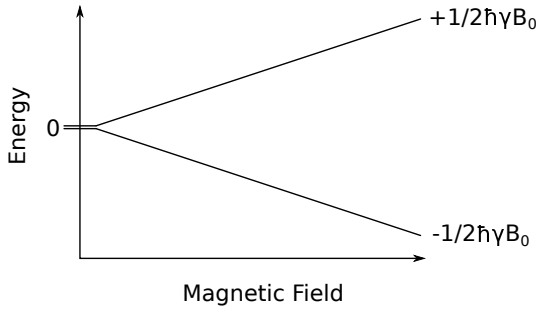

Figure 1.2: The energy of the two quantum spin states of a spin-1/2 nucleus in a magnetic field of varying strength.

This phenomenon is not just limited to electrons; protons and neutrons are both also spin $\frac{1}{2}$, and therefore the nuclei which they make up also have an overall spin. A consequence of this is the NMR effect, which forms the basis of magnetic resonance imaging and will be treated in detail here, first by examining an isolated spin-1/2 nucleus, before expanding this to an ensemble of nuclei.

1.3.2 The isolated spin-1/2 nucleus – a quantum description

1.3.2.1 The Zeeman effect

When a spin-1/2 ($s = \frac{1}{2}$, $m_s = \pm\frac{1}{2}$) nucleus is placed in a magnetic field (B_0) it experiences the nuclear Zeeman effect, which removes the degeneracy of the two m_s states and splits them so that their energy is $\pm\frac{1}{2}\hbar\gamma B_0$ (Figure 1.2). Since there are now two spin states with differing energy, the nucleus can transition between them by absorbing/emitting a photon with energy $\hbar\gamma B_0$. It is these transitions which give us nuclear magnetic resonance.

The Hamiltonian of the Zeeman interaction for a spin in a magnetic field $\mathbf{B}_0$ aligned with the z-axis is defined as

$$\hat{H} = -\gamma B_0 \hbar \hat{I}_z = \omega_0 \hbar \hat{I}_z \tag{1.2}$$

$\hat{I}_z$ is the z angular momentum operator and $\omega_0 = -\gamma B_0$ is an angular frequency. The eigen-states of $\hat{I}_z$ are denoted $|\alpha\rangle$ and $|\beta\rangle$, with eigenvalues of $\hat{I}_z|\alpha\rangle = +\frac{1}{2}\hbar$ and $\hat{I}_z|\beta\rangle = -\frac{1}{2}\hbar$ respectively[16]. For the remainder of this section $\hbar = 1$ by

definition for simplicity. Using these as a basis set, a wave function for the system can be defined as $|\psi\rangle = c_\alpha|\alpha\rangle + c_\beta|\beta\rangle$.

Using these definitions the expectation of the Hamiltonian is:

$$\langle\hat{H}\rangle = \langle\psi|\hat{H}|\psi\rangle = \left[c_\alpha^*\langle\alpha| + c_\beta^*\langle\beta|\right]\hat{H}\left[c_\alpha|\alpha\rangle + c_\beta|\beta\rangle\right] \tag{1.3}$$

which, owing to the orthogonality of the states, simplifies to:

$$\langle\hat{H}\rangle = \frac{1}{2}\omega_0|c_\alpha|^2 - \frac{1}{2}\omega_0|c_\beta|^2. \tag{1.4}$$

This shows that the expectation value for the Zeeman energy of the system is equal to $(1/2)\omega_0$ times the expected difference in populations of the two spin states.

1.3.2.2 Spin angular momentum operators (Pauli Matrices)

By considering the x, y, z angular momentum expectation values ($\langle\hat{I}_x\rangle$, $\langle\hat{I}_y\rangle$ and $\langle\hat{I}_z\rangle$) the state of the spin system can be calculated using the time dependent Schrodinger equation[16]:

$$\frac{\mathrm{d}}{\mathrm{d}t}|\psi\rangle = -i\hat{H}|\psi\rangle \tag{1.5}$$

However we first need expressions for $\langle\hat{I}_x\rangle$, $\langle\hat{I}_y\rangle$ and $\langle\hat{I}_z\rangle$. To simplify these calculations (from NMR: The toolkit[16]) the $\hat{I}_x$, $\hat{I}_y$ and $\hat{I}_z$ operators can be expressed in Pauli matrix form:

$$I_z = \begin{pmatrix} \langle\alpha|\hat{I}_z|\alpha\rangle & \langle\alpha|\hat{I}_z|\beta\rangle \\ \langle\beta|\hat{I}_z|\alpha\rangle & \langle\beta|\hat{I}_z|\beta\rangle \end{pmatrix} = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & -\frac{1}{2} \end{pmatrix} = m_s\sigma_z \tag{1.6}$$

$$I_x = \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} = m_s\sigma_x \tag{1.7}$$

$$I_y = \begin{pmatrix} 0 & -\frac{1}{2}i \\ \frac{1}{2}i & 0 \end{pmatrix} = m_s\sigma_y \tag{1.8}$$

Where $\sigma_{x/y/z}$ are the Pauli spin matrices[17]:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{1.9}$$

Allowing their expectation values to be calculated as:

$$\langle \hat{I}_z \rangle = \begin{pmatrix} c_\alpha^* & c_\beta^* \end{pmatrix} I_z \begin{pmatrix} c_\alpha \\ c_\beta \end{pmatrix} = \frac{1}{2}|c_\alpha|^2 - \frac{1}{2}|c_\beta|^2 \tag{1.10}$$

$$\langle \hat{I}_x \rangle = \begin{pmatrix} c_\alpha^* & c_\beta^* \end{pmatrix} I_x \begin{pmatrix} c_\alpha \\ c_\beta \end{pmatrix} = \mathrm{Re}\left(c_\alpha c_\beta^*\right) \tag{1.11}$$

$$\langle \hat{I}_y \rangle = \begin{pmatrix} c_\alpha^* & c_\beta^* \end{pmatrix} I_y \begin{pmatrix} c_\alpha \\ c_\beta \end{pmatrix} = -\mathrm{Im}\left(c_\alpha c_\beta^*\right) \tag{1.12}$$

1.3.2.3 The evolution of spin angular momentum in a static field

The expectation value of the angular momentum operators in a static field can be calculated using the matrix form of the time dependent Schrodinger equation:

$$\dot{\psi} = -iH\psi \tag{1.13}$$

Substituting in the Hamiltonian for the Zeeman interaction, $H = \omega_0 I_z$, gives:

$$\begin{pmatrix} \dot{c}_\alpha \\ \dot{c}_\beta \end{pmatrix} = \begin{pmatrix} -\frac{1}{2}i\omega_0 c_\alpha \\ \frac{1}{2}i\omega_0 c_\beta \end{pmatrix} \tag{1.14}$$

Which integrate to give equations for c_α and c_β:

$$c_\alpha(t) = c_\alpha(0)e^{-\frac{1}{2}i\omega_0 t}, \qquad c_\beta(t) = c_\beta(0)e^{\frac{1}{2}i\omega_0 t} \tag{1.15}$$

Substituting suitable initial conditions into Equations (1.10) to (1.12) allows the expectation value for the spin state of the system to be calculated at time t. For instance with initial conditions at $t = 0$ of $\langle \hat{I}_x \rangle = 1/2$, $\langle \hat{I}_y \rangle = 0$ and $\langle \hat{I}_z \rangle = 0$, i.e. $c_\alpha = c_\beta = \frac{1}{\sqrt{2}}$ (following a 90° pulse), the angular momentum expectation values are[16]:

$$\begin{aligned} \langle \hat{I}_x \rangle &= \frac{1}{2}\cos(\omega_0 t) \\ \langle \hat{I}_y \rangle &= -\frac{1}{2}\sin(\omega_0 t) \\ \langle \hat{I}_z \rangle &= 0 \end{aligned} \tag{1.16}$$

The implication of this is that the expectation of the spin angular momentum precesses around the z-axis at an angular frequency of ω_0, otherwise known as the Larmor frequency.

1.3.2.4 The evolution of spin angular momentum during a radio frequency pulse

The effect of a radio frequency (RF) pulse on the system can be calculated using the Hamiltonian for RF irradiation along the x-axis[16]:

$$\hat{H} = \omega_0 \hat{I}_z + 2\omega_1 \cos(\omega_{\mathrm{rf}} t + \phi)\hat{I}_x \tag{1.17}$$

In the rotating frame, with $\omega_{\mathrm{rf}} = \omega_0$, (i.e. our frame of reference rotates about the spin at ω_0) the Hamiltonian simplifies to:

$$\hat{H}_{\mathrm{R}} = -\gamma B_1 \hat{I}_x = \omega_1 \hat{I}_x \tag{1.18}$$

Allowing the state of the system at time t to be calculated with the same method given suitable initial conditions. If the spin is in its ground state $|\alpha\rangle$ ($c_\alpha = 1,\ c_\beta = 0$) at $t = 0$:

$$\begin{aligned} c_\alpha(t) &= \cos\left(\frac{1}{2}\omega_1 t\right) \\ c_\beta(t) &= -i\sin\left(\frac{1}{2}\omega_1 t\right) \\ \langle \hat{I}_x \rangle &= 0 \\ \langle \hat{I}_y \rangle &= -\frac{1}{2}\sin\left(\omega_1 t\right) \\ \langle \hat{I}_z \rangle &= \frac{1}{2}\cos\left(\omega_1 t\right) \end{aligned} \tag{1.19}$$

.

These equations show that an RF pulse at frequency $\omega_{\mathrm{rf}} = \omega_0$ rotates the expectation value of the spin angular momentum operators around the x-axis.

The probability of the spin being in the $|\alpha\rangle$ or $|\beta\rangle$ state can be found by calculating $|c_\alpha|^2$ and $|c_\alpha|^2$ to give:

$$\begin{aligned} P_\alpha(t) &= 1 - \sin^2\left(\frac{1}{2}\omega_1 t\right) \\ P_\beta(t) &= \sin^2\left(\frac{1}{2}\omega_1 t\right). \end{aligned} \tag{1.20}$$

This shows that the spin state is oscillating between $|\alpha\rangle$ and $|\beta\rangle$ states with frequency ω_1. The implication of this is that the spin is in resonance with the RF pulse, i.e. is undergoing nuclear magnetic resonance.

This result is a special case of the Rabi cycle, where the RF-pulse has exactly the same frequency as the Larmor frequency, more generally the probability of a transition from $|\alpha\rangle$ to $|\beta\rangle$ is given as:

$$P_\beta = \frac{|\omega_1|^2}{|\Omega^2 + \omega_1^2|}\sin^2\left(\frac{\sqrt{\Omega^2+\omega_1^2}}{2}t\right) \quad \text{where} \quad \Omega = \omega_0 - \omega_{\text{rf}}. \tag{1.21}$$

1.3.3 Through the classical limit - the vector model

1.3.3.1 Calculation of the bulk magnetization

In the previous section we calculated the behavior of a single spin-1/2 nuclei, using quantum theory, now we will consider the bulk properties of an ensemble of non-interacting nuclei, as would be found in a typical MRI experiment, using classical mechanics.

The bulk magnetization of the sample (M) can be calculated as the sum of the individual nuclei's magnetic moments (μ):

$$\mu_\alpha = \gamma\hat{I}_z|\alpha\rangle = +\gamma\frac{1}{2}\hbar, \qquad \mu_\beta = \gamma\hat{I}_z|\beta\rangle = -\gamma\frac{1}{2}\hbar \tag{1.22}$$

$$M = \sum \mu_i = \frac{1}{2}\gamma\hbar\Delta n_{\text{eq}}, \qquad \Delta n_{\text{eq}} = n_\alpha - n_\beta \tag{1.23}$$

Where $n_{\alpha/\beta}$ is the number of nuclei in the α/β states.

To calculate Δn_{eq}, and therefore M, we first use Boltzmann statistics[18] to calculate the ratio n_β/n_α

$$\frac{n_\beta}{n_\alpha} = e^{-\frac{\Delta E}{k_b T}} = e^{-\frac{\hbar\gamma B_0}{k_b T}} \tag{1.24}$$

then, note that the polarization of the spins is defined as:

$$P = \frac{\Delta n_{\text{eq}}}{n_{\text{total}}} = \frac{n_\alpha - n_\beta}{n_\alpha + n_\beta} = \frac{1 - n_\beta/n_\alpha}{1 + n_\beta/n_\alpha} \tag{1.25}$$

and substitute in the equation for n_β/n_α using the approximation $\exp(-a/b) \approx (1 - a/b)$ if $b >> a$

$$\frac{\Delta n_{\text{eq}}}{n_{\text{total}}} = \frac{1 - e^{-\frac{\hbar\gamma B_0}{k_b T}}}{1 + e^{-\frac{\hbar\gamma B_0}{k_b T}}} \approx \frac{1 - (1 - \frac{\hbar\gamma B_0}{k_b T})}{1 + (1 - \frac{\hbar\gamma B_0}{k_b T})} \approx \frac{\hbar\gamma B_0}{2k_b T} \tag{1.26}$$

This equation then allows us to calculate M by substituting it back into Equation (1.23) to give

$$M = n_{\text{total}} \left(\frac{\hbar\gamma}{2}\right)^2 \frac{B_0}{k_b T}. \tag{1.27}$$

1.3.3.2 The motion of bulk magnetization in a static field

The bulk magnetization is an ensemble of many spins and therefore obeys classical physics, allowing the classical equations of motion for a magnetic dipole to be used to calculate its properties. The torque (τ) experienced by a spin in a magnetic field is[19]

$$\tau = \mu \times B. \tag{1.28}$$

Noting that torque is the time derivative of angular momentum and therefore[19]

$$\tau = \frac{1}{\gamma}\frac{\mathrm{d}\mu}{\mathrm{d}t} \tag{1.29}$$

allows us to sum over all spins, and produce an equation for the motion of the bulk magnetization[20]:

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \gamma M \times B \tag{1.30}$$

The static magnetic field, B_0, is defined to act along the z axis, so that $B = (0, 0, B_0)$, allowing the equation to reduce to:

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = \gamma M_y B_0, \qquad \frac{\mathrm{d}M_y}{\mathrm{d}t} = -\gamma M_x B_0, \qquad \frac{\mathrm{d}M_z}{\mathrm{d}t} = 0 \tag{1.31}$$

which can be solved (Appendix A.2) to give[20]:

$$\begin{aligned} M_x &= M_x(0)\cos(\omega_0 t) + M_y(0)\sin(\omega_0 t) \\ M_y &= M_y(0)\cos(\omega_0 t) - M_x(0)\sin(\omega_0 t) \\ M_z &= M_z(0). \end{aligned} \tag{1.32}$$

Yielding the equations of motion for magnetization in the presence of a magnetic field (B_0), noting that for the classical derivations in this thesis $\omega_0 = \gamma B_0$.

If the same initial conditions are substituted in as for Equation (1.16) ($M_x(0) = 1/2$, $M_y(0) = 0$, $M_z(0) = 0$), which describes the motion of the angular momentum expectation values $\langle \hat{I}_{x/y/z} \rangle$ we recover the same result (noting that in Equation (1.16) $\omega_0 = -\gamma B_0$):

$$\begin{aligned} M_x &= \frac{1}{2}\cos(\omega_0 t) \\ M_y &= -\frac{1}{2}\sin(\omega_0 t) \\ M_z &= 0 \end{aligned} \tag{1.33}$$

1.3.3.3 The motion of bulk magnetization during an RF pulse

An RF-pulse, perpendicular to B_0 and oscillating at ω_0 can be represented by the sum of two counter rotating circularly polarized sub-fields. Of these two sub-fields, only the one rotating in the same direction as ω_0 has an effect on the system, with the energy from the other sub-field wasted, and is described as[20]:

$$\mathbf{B}_1 = B_1 \begin{pmatrix} \cos(\omega_0 t) \\ -\sin(\omega_0 t) \\ 0 \end{pmatrix} \tag{1.34}$$

Substituting this into Equation (1.30) gives:

$$\begin{aligned} \frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_z B_1 \sin(\omega_0 t) + M_y B_0] \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[M_z B_1 \cos(\omega_0 t) - M_x B_0] \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= -\gamma[M_x B_1 \sin(\omega_0 t) + M_y B_1 \cos(\omega_0 t)] \end{aligned} \tag{1.35}$$

Which can be solved (Appendix A.3) to give[21]:

$$\begin{aligned} M_x(t) &= M_0 \sin(\omega_1 t)\sin(\omega_0 t) \\ M_y(t) &= M_0 \sin(\omega_1 t)\cos(\omega_0 t) \\ M_z(t) &= M_0 \cos(\omega_1 t) \end{aligned} \tag{1.36}$$

These equations are plotted for a 90° RF-pulse in Figure 1.3 and show that, as the RF-pulse plays out, the net magnetization processes around the z axis at frequency ω_0 and rotates away from the z-axis about the B_1 field with angular frequency ω_1[20]. If we work in the rotating frame and use the same starting conditions as for Equation (1.19), the quantum derivation, the same result is reached, noting that for Equation (1.19) rotational frequencies are defined as $\omega_n = -\gamma B_n$.

1.3.4 NMR decay constants and the Bloch equations

Once the net magnetization has been perturbed by an RF-pulse, it is no-longer in thermal equilibrium and will therefore decay back to its equilibrium state. An RF-pulse perturbs the spins which sum to give the net magnetization in two ways, firstly a proportion (depending on the pulse) are flipped from one quantum state to another, and secondly they become coherent so that the expectation value for the x/y angular momentum is non-zero. The effect of these processes is to rotate the net-magnetization about the B_1 field of the RF-pulse. Similarly, two separate processes govern decay back to the ground state. The first, T_1-relaxation, is the flipping of the spins so that the proportion of the spins in each state equals the equilibrium state. Since the transition between $|\alpha\rangle$ and $|\beta\rangle$ by spontaneous emission is very unlikely due to the closeness in energy of the states (Fermi's golden rule) the process relies on external oscillating magnetic fields generated by tumbling molecules in the vicinity of the spins driving stimulated emission. The second, T_2 relaxation, is the loss of coherence in the spins and is caused by each spin precessing at a slightly different rate due to the effect of random spin-spin interactions during the lifetime of the excited state. In practice the B_0 field is not completely homogeneous, creating an 'artificial' source of T_2 decay and the combined decay from both the intrinsic properties of the sample and the in-homogeneity of the magnet is termed T_2^* decay.

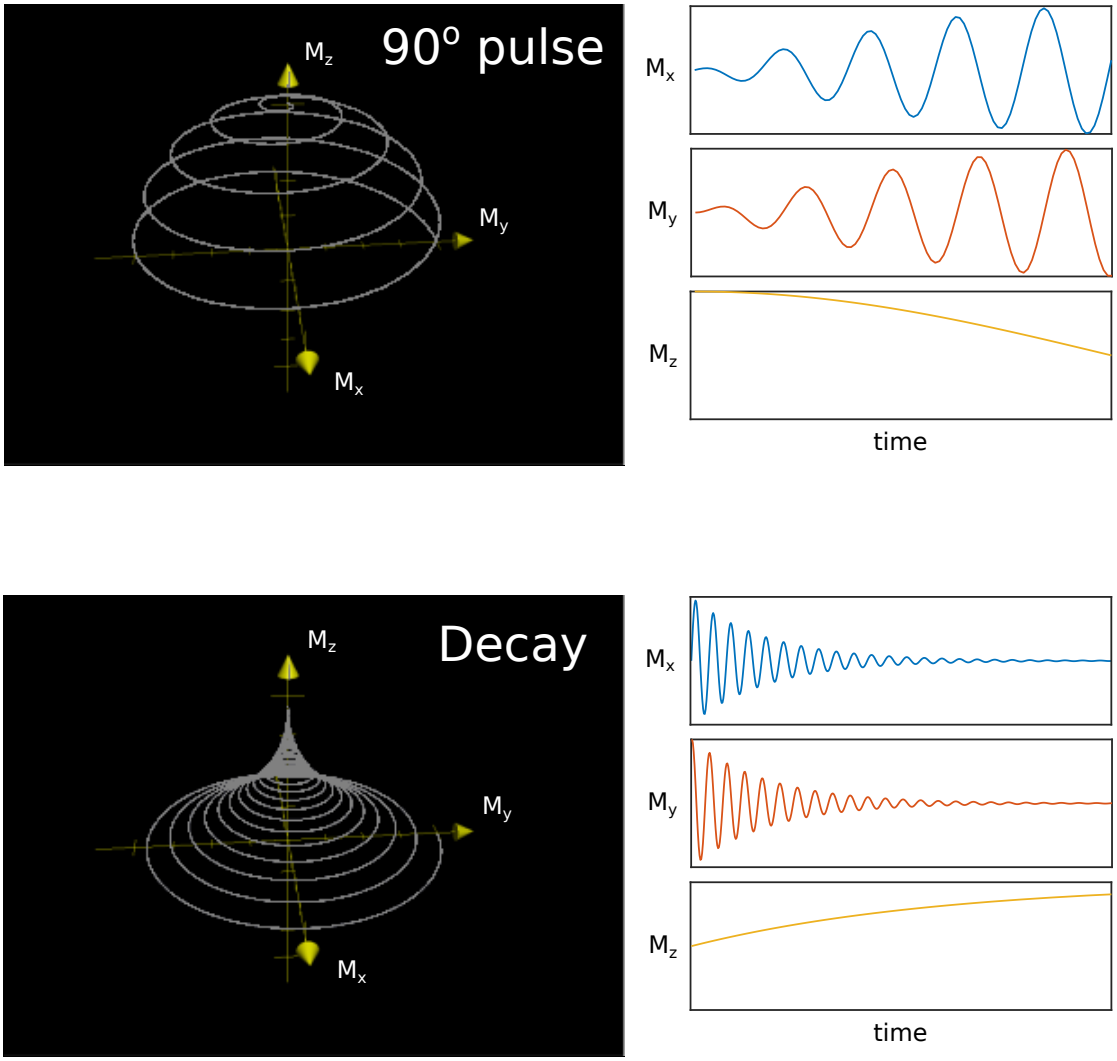

Figure 1.3: The evolution of the net-magnetization during and immediately following a 90° pulse

First order differential equations for the effect of T_1 and T_2 relaxation on the longitudinal and transverse magnetization can be derived theoretically[22], and can be simplified by defining the constant terms as T_1 and T_2 respectively[21]:

$$\frac{\mathrm{d}M_{x/y}}{\mathrm{d}t} = \frac{-M_{x/y}}{T_2}, \qquad \frac{\mathrm{d}M_z}{\mathrm{d}t} = \frac{M_0 - M_z}{T_1} \tag{1.37}$$

These equations can be substituted into Equation (1.30) to give differential equations for the motion of the net-magnetization, which are often termed 'the Bloch equations'[21]:

$$\begin{aligned}\frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_y B_z - M_z B_y] - \frac{M_x}{T_2} \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[-M_x B_z + M_z B_x] - \frac{M_y}{T_2} \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= \gamma[M_x B_y - M_y B_x] + \frac{M_0 - M_z}{T_1}\end{aligned} \tag{1.38}$$

By solving these equations we can effectively simulate MRI pulse sequences allowing for the development of complex imaging schemes, for instance the simple example of T_1/T_2* decay following a 90° pulse is shown in Figure 1.3.

1.3.5 NMR spectra formation

One of the principal applications of the NMR effect is resolving the chemical structure of molecules in a sample and quantifying how much of a given molecule is present. This is made possible because the resonant frequency of a given nucleus depends on the magnetic field experienced by that particular nucleus, and aside from the strength of the B_0 field, this also depends on the molecule's chemical structure. The electrons in a molecule oppose the main magnetic field by Lenz's law, and therefore, nuclei in regions with a greater density of electrons will experience a lower magnetic field and resonate at a lower frequency, allowing the presence of a particular chemical environment to be inferred. The ability of an atom to attract electrons is termed its electronegitivity, so for example, when a proton is adjacent to a strongly electronegitive atom such as chlorine it will resonate at a higher frequency than it would otherwise, allowing that adjacency to be inferred as shown in Figure 1.4. In addition since the resonances for nuclei in different environments are separated, the concentrations of the nuclei can be quantified by the strength of the resonance. For example, in chloroethane the protons adjacent to the chlorine atom experience a greater magnetic field than those on the other carbon allowing the two environments to be distinguished and the strengths of the resonances then depend on the number of nuclei in each environment, and as such have a ratio of 2:3.

The earliest NMR experiments generated spectra by sweeping the RF transmit frequency through a wide range of frequencies and looking for resonances when

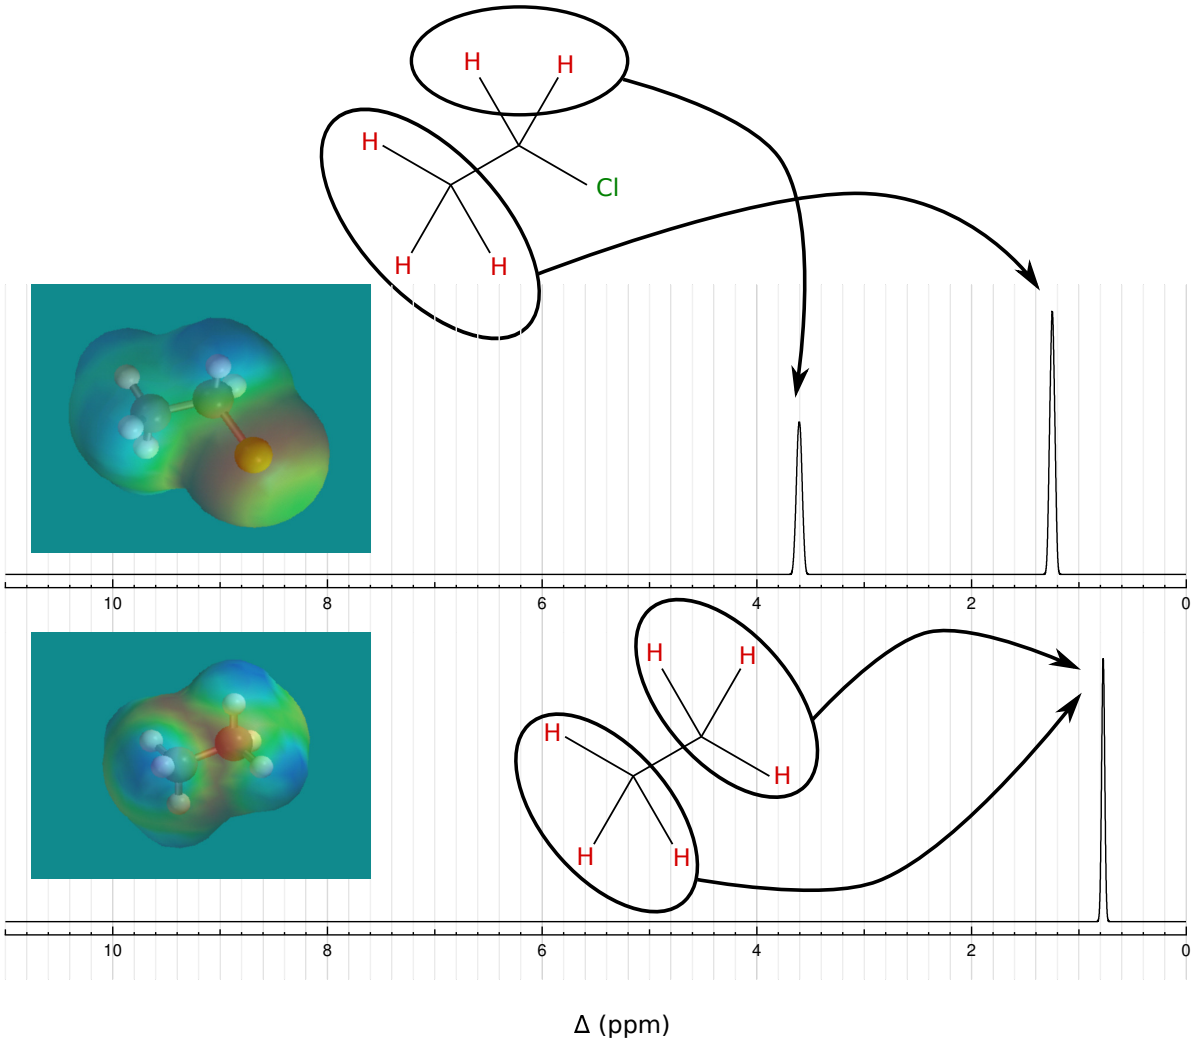

Figure 1.4: Predicted NMR spectra (nmrdb.org) for ^{1}H nuclei in chloroethane and ethane, the electron-withdrawing properties of the chlorine atom causes the adjacent nuclei to resonate at a higher frequency (by convention the x-axis is inverted). The surface on the molecular simulations (Spartan software package) indicates the Van der Waals radii and the color the charge at that point (blue +ve, red -ve). The x-axis is in units of ppm (e.g. for a 300 MHz spectrometer 1 ppm = 300 Hz) and are relative to the commonly used standard tetramethylsilane. Note that in this example the linewidth is too great to observe splitting.

$\omega_{\mathrm{RF}} = \omega_0$ for a given nuclei. This method is intrinsically slow since it must sample each RF frequency in turn, so a more efficient method was devised, namely Fourier analysis. With Fourier analysis, a short RF-pulse excites the sample and once the pulse has finished the scanner receives the RF signal emitted by nuclei decaying back to the ground state. In practice, the received RF is mixed with two reference frequencies (ω_{RF}) 90° out of phase, to allow the phase and magnitude of the signal to be detected. These two signals are then analyzed as follows:

For a given resonance $\Omega = \omega_0 - \omega_{RF}$ represents the frequency of the mixed signal which can be expressed in complex form as[16]

$$s(t) = s(0)\left(\cos(\Omega t) + i\sin(\Omega t)\right) e^{-t/T_{2*}} = s(0)e^{-t/T_{2*}}e^{i\Omega t}, \tag{1.39}$$

assuming that the T_1 decay time is long relative to the experiment. This can be transformed into the frequency domain using the Fourier transform:

$$S(\omega) = \int_{-\infty}^{\infty} s(t)e^{-i\omega t}\,\mathrm{d}t \tag{1.40}$$

$$S(\omega) = \int_{-\infty}^{\infty} s(0)e^{-t/T_{2*}}e^{i\Omega t}e^{-i\omega t}\,\mathrm{d}t = \int_{-\infty}^{\infty} s(0)e^{-[i(\omega-\Omega)+\frac{1}{T_{2*}}]t}\,\mathrm{d}t \tag{1.41}$$

$$S(\omega) = \left[s(0)\frac{e^{-[i(\omega-\Omega)+\frac{1}{T_{2*}}]t}}{-[i(\omega-\Omega)+\frac{1}{T_{2*}}]}\right]_{-\infty}^{\infty} \tag{1.42}$$

Note that there is no signal before $t = 0$, allowing the lower $-\infty$ limit to be substituted for zero:

$$S(\omega) = \left[s(0)\frac{e^{-i(\omega-\Omega)t}e^{-\frac{1}{T_{2*}}]t}}{-[i(\omega-\Omega)+\frac{1}{T_{2*}}]}\right]_{0}^{\infty} \tag{1.43}$$

Giving[23]:

$$S(\omega) = \frac{s(0)}{i(\omega-\Omega)+\frac{1}{T_{2*}}} = s(0)\frac{-i(\omega-\Omega)+\frac{1}{T_{2*}}}{(\omega-\Omega)^2+(\frac{1}{T_{2*}})^2} \tag{1.44}$$

Which can be expressed in terms of $\Delta\omega = \omega - \Omega$ and separated into two parts ($S(\omega) = A(\Delta\omega) - iD(\Delta\omega)$), both of which are Lorentzian functions[16]:

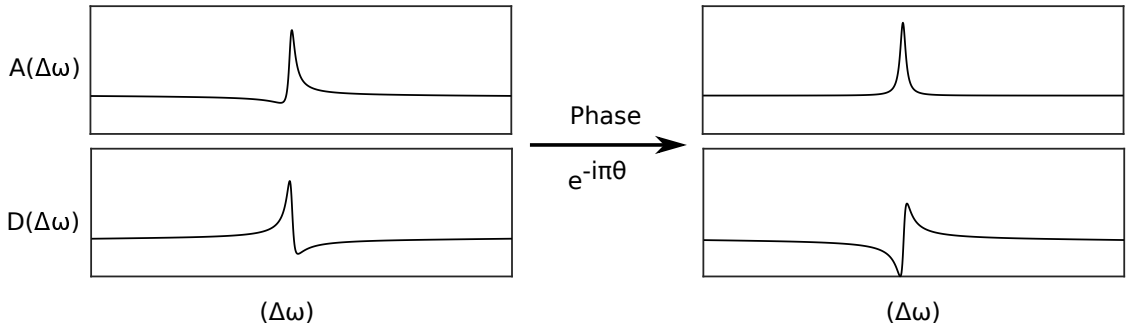

Figure 1.5: Reconstructed $A(\Delta\omega)$ and $D(\Delta\omega)$ spectra before and after phasing.

$$A(\Delta\omega) = s(0)\frac{\frac{1}{T_{2*}}}{(\Delta\omega)^2 + (\frac{1}{T_{2*}})^2}, \qquad D(\Delta\omega) = s(0)\frac{\Delta\omega}{(\Delta\omega)^2 + (\frac{1}{T_{2*}})^2}. \tag{1.45}$$

For display and analysis $A(\Delta\omega)$ is usually taken, however depending on the phase of the signal relative to the reference signal, the initial spectra produced by the Fourier transform will be a mixture of $A(\Delta\omega)$ and $D(\Delta\omega)$ which can be separated by multiplying $S(\omega)$ by a phase term of the form $\exp(-i\pi\theta)$ where θ is determined empirically as shown in Figure 1.5.

1.3.6 The SNR of an NMR/MRI experiment

The SNR of a simple NMR/MRI experiment can be derived by calculating the peak emf induced in the receive coil, using Faraday's law of induction[24] and dividing through by the noise associated with the experiment[25].

Faraday's law states that the emf induced in a coil is

$$\text{emf} = -\frac{\mathrm{d}\Phi_B}{\mathrm{d}t} \tag{1.46}$$

where B is the magnetic field and Φ_B the magnetic field flux passing through the coil.

Using the principal of reciprocity, we can calculate Φ_B as the integral of the B_1 field and net magnetization across the sample[24].

$$\Phi_B = -\int_{\text{sample}} \mathbf{B_1} \cdot \mathbf{M}\, \mathrm{d}V \tag{1.47}$$

If the B_1 field is homogeneous, following a 90° pulse, this gives (using Equation (1.32))

$$\Phi_B = -B_1 V M_0(\cos(\omega_0 t) - i\sin(\omega_0 t)) \tag{1.48}$$

which differentiates to give

$$\text{emf} = B_1 M_0 \omega_0(\sin(\omega_0 t) + i\cos(\omega_0 t)) \tag{1.49}$$

Taking the maximum emf and substituting it into the equations for M_0 and ω_0 gives an equation for the received signal ignoring relaxation effects

$$S = B_1 n_{eq} \frac{\hbar^2 B_0^2 \gamma^3}{4k_b T}. \tag{1.50}$$

The noise can then be calculated by considering the noise generated by each component in the RF chain. The variance of the total noise is given by

$$\sigma^2_{\text{thermal}} = \sigma^2_{\text{body}} + \sigma^2_{\text{coil}} + \sigma^2_{\text{electronics}} \tag{1.51}$$

The RMS thermal (Johnson-Nyquist) noise[26, 27] of the coil and electronics (caused by the thermal agitation of electrons) is given by Equation (1.52)[25].

$$\sqrt{\langle e_n^2 \rangle} = \sqrt{4k_b T R \Delta f} \tag{1.52}$$

where $\sqrt{\langle e_n^2 \rangle}$ is the RMS of the amplitude of the thermal noise (with units of Volts), R the resistance of the coil and electronics and Δf the bandwidth of the signal. Provided any DC offset in the signal is eliminated $\langle e_n^2 \rangle = \sigma^2_{\text{circuit}}$. If Δf is then given in units of ppm the equation becomes

$$\sqrt{\sigma^2_{\text{circuit}}} = \sqrt{4k_b T R \gamma B_0 (\Delta\text{ppm})/10^6} \tag{1.53}$$

The noise associated with the body originates from inductive losses caused by random eddy currents circulating within conductive tissue. The amplitude of these eddy currents depends on the experimental setup, however the frequency received by the coil is restricted to a small range around the Larmor frequency, and therefore the emf induced in the coil is proportional to the Larmor frequency (Faraday's law)[28] giving a B_0 dependence of

$$\sigma^2_{\text{body}} \propto \Delta F \omega_0 = \gamma^2 B_0^2 (\Delta\text{ppm})/10^6. \tag{1.54}$$

If noise from the body dominates, which is the case in almost all clinical MRI imaging, the SNR is governed by

$$\frac{S}{N} \propto B_1 n_{eq} \frac{\hbar^2 B_0 \gamma^2}{4 k_b T} \tag{1.55}$$

Although a number of crude approximations were made in this equation's derivation, it does illustrate the key principles when considering the SNR of an acquisition. Firstly that the amount of signal per unit volume depends linearly on the B_1 field, spin density and B_0, and secondly that it depends on the square of the gyro-magnetic ratio of the nuclei.

1.4 MRI hardware

Figure 1.6 shows an annotated photograph of a typical pre-clinical MRI scanner. In addition to the labeled hardware (magnet, gradient set and RF coil) there are several additional systems critical for the operation of the scanner which drive the RF and gradient systems, and the MRI console which coordinates the various systems to produce the required pulse sequence. Each system will now be discussed in turn.

1.4.1 B_0 field

In order to polarize the nuclear spins found in the sample/volunteer, the scanner must generate a strong B_0 field. This is typically achieved using a superconducting magnet, however, some scanners may use resistive or permanent magnets[29]. The

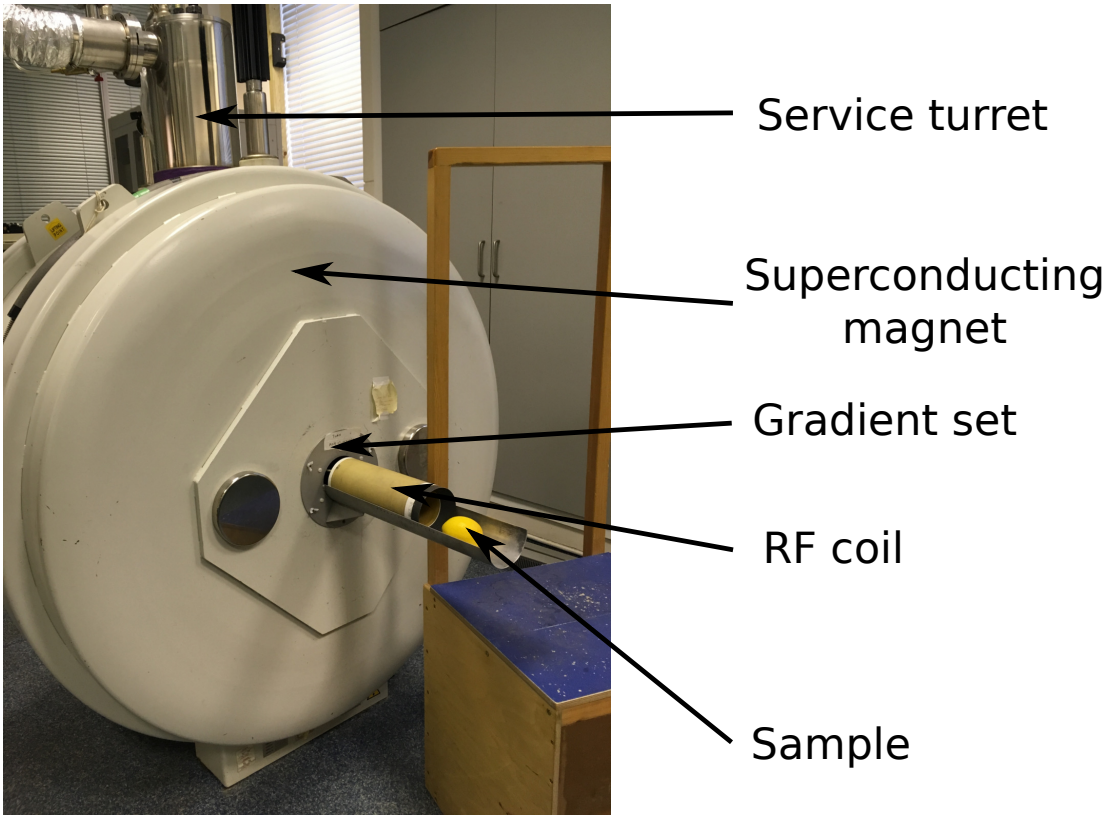

Figure 1.6: Varian 7T DDR pre-clinical MRI scanner

Varian 7T DDR pre-clinical MRI scanner pictured in Figure 1.6 has a 7T actively shielded superconducting magnet. This consists of 6 NbSn superconducting coils carrying a $\approx$ 250A current and arranged to maximize the magnetic field homogeneity in the bore of the magnet. These coils are cooled by a liquid helium cryostat, which is itself contained within a liquid nitrogen cryostat, to minimize helium boil off. All the MRI scanners used in this thesis have active shielding to contain the B_0 field, which is provided by additional superconducting coils, outside of the principal coils, producing a magnetic field opposing the primary field. This serves two purposes, firstly enhancing the magnetic flux within the magnet's bore and secondly reducing the physical extent of the field which would otherwise extend some distance, and cause a safety risk outside of the scan room.

While the magnet is designed to minimize in-homogeneity in the region close to the centre of the field (isocentre), called the designated spherical volume, it will not be perfect. Additionally the sample/volunteer being scanned will induce

their own distortions in the field, as boundaries between regions with dissimilar electrodynamic properties cause magnetic field discontinuities (Maxwell's equations). Therefore, a series of resistive coils are embedded in the magnet, collectively known as the shim set, to allow fine adjustment of the field (some magnets may have factory-set superconducting shims as well). The fields these shim coils generate are defined by spherical harmonics e.g. x, y, z, $x^2 - y^2$ etc. and are adjusted prior to acquisition and then left on continuously throughout.

1.4.2 RF transmit/receive chain

The RF transmit / recieve chain consists of all the hardware necessary to play out the RF pulses used to excite the nuclear spins and receive the RF signal produced by the nuclei as they decay back to their ground spin states. The first component in this chain, the RF waveform generator, synthesizes the analogue RF waveforms to be played out. This is connected to an RF amplifier, which amplifies the signal produced by the waveform generator to the required power and drives the RF coil. The RF coil is the device which transmits and receives the RF field, and is often changed based on the requirements of the experiment, for instance different shaped RF coils are used to scan different parts of the body. An RF switch is used to separate the transmitted and received RF signal and the received signal is mixed with a reference frequency (see section 1.3.5) before digitization with an ADC.

1.4.3 The gradient set

The gradient set is used to apply magnetic field gradients across the bore of the scanner. By using three sets of independent gradient coils, it is able to achieve this in the x, y and z directions independently. The x and y coils are traditionally described as Golay (saddle) coils positioned at 90° to each other; the z (along the bore) gradient coil consists of a pair of Helmholtz coils. These gradient coils require significant power and are driven by a set of gradient amplifiers, which in turn are coordinated by the console. The performance of this system has a significant effect on the utility of the magnet, as it affects the time required to trace out the k-space trajectory

(see Sections 1.5.1 and 1.5.2) of the acquisition, while T_2* signal decay is ongoing. Gradient performance is measured by two key metrics, the maximum available gradient strength (G_{max}) in mT/m and the maximum rate of gradient change, or slew rate, (S_{max}) in mT/m/s. Performance may vary significantly between systems with the Varian 7T DDR pre-clinical scanner shown in Figure 1.6 achieving G_{max} = 1100 mT/m and S_{max} = 6000 mT/m/ms, whereas a clinical Siemens 3T Tim Trio only achieves G_{max} = 40 mT/m and S_{max} = 200 mT/m/ms and a clinical 7T Siemens MAGNETOM Terra achieves G_{max} = 80 mT/m and S_{max} = 200 mT/m/ms, although these clinical systems do have a much larger bore size (60 cm vs 10 cm).

1.4.4 Ancillary systems

In addition to the main MRI hardware, there may be a number of ancillary systems to aid with certain protocols. One example is ECG hardware, which enables cardiac gating, whereby the scan is synchronized to the heart beat so that acquisitions occur in the same cardiac phase. Pre-clinical systems may also include anesthesia and recovery equipment to allow animals to be scanned.

1.5 MRI imaging

Fundamentally an MRI image is a map of the quantity of RF photons emitted by each point in the image due to the NMR effect. If T_1/T_2* and other processes are put aside, this is proportional to the density of spins at each point in the image, which varies by material (e.g. tissue type), thereby providing contrast and producing a useful image. As an example, Figure 1.7 shows an image of an orange, where the different tissue types, e.g. flesh, pith, seed and skin can clearly be differentiated by the ^{1}H spin density, which, in this context, is largely proportional to water content.

The procedure, or pulse sequence, used to perform the MRI acquisition and record the data needed to reconstruct an MRI image can broadly be divided into two parts, RF excitation and readout. Localization of the signal is achieved using linearly varying x, y and z magnetic field gradients. These gradients are used to

Figure 1.7: ^{1}H MRI image of an orange, the difference in ^{1}H spin density (corresponding to H_2O content) between flesh and pith can clearly be seen.

either restrict RF excitation to only a small region, or induce spatial variation in the spins phase during readout, or a combination of both.

Both phases of the pulse sequence can be designed, subject to physical constraints on RF power and gradient strength, by specifying gradient and RF waveforms, calculated using the underlying physical principals of MRI.

An important theoretical concept which underpins our understanding of signal localization using magnetic field gradients during both excitation and readout is the k-space formalism

1.5.1 The k-space formalism (excitation)

To understand the effects of linear gradient fields during RF excitation we will use the excitation k-space formalism[30]. We start with the Bloch equations in the rotating frame (section 1.3.4)

$$\begin{aligned}\frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_y B_z - M_z B_y] - \frac{M_x}{T_2} \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[-M_x B_z + M_z B_x] - \frac{M_y}{T_2} \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= \gamma[M_x B_y - M_y B_x] + \frac{M_0 - M_z}{T_1}\end{aligned} \tag{1.56}$$

Switching to the rotating frame (appendix A.1), adding in terms for the effects of the aforementioned linear magnetic field gradients and ignoring relaxation processes these equations become

$$\begin{aligned}\frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_y \mathbf{G}\cdot\mathbf{r} - M_z B_{1y}] \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[-M_x \mathbf{G}\cdot\mathbf{r} + M_z B_{1x}] \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= \gamma[M_x B_{1y} - M_y B_{1x}]\end{aligned} \tag{1.57}$$

using the small flip angle approximation[30] which approximates, for a small excitation flip angle, that $M_z(t) = M_z(0)$ and $\sin(\phi) = \phi$, the equations reduce to

$$\begin{aligned}\frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_y \mathbf{G}\cdot\mathbf{r} - M_0 B_{1y}] \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[-M_x \mathbf{G}\cdot\mathbf{r} + M_0 B_{1x}] \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= 0\end{aligned} \tag{1.58}$$

If $M_{xy} = M_x + iM_y$ and $B_1 = B_{1x} + iB_{1y}$ these equations can be expressed simply as

$$\dot{M}_{xy} = -i\gamma \mathbf{G}\cdot\mathbf{r} M_{xy} + i\gamma B_1 M_0 \tag{1.59}$$

and a solution found by integrating $\dot{M}_{xy}$ to give

$$M_{xy}(\mathbf{r}, T) = i\gamma M_0 \int_{-\infty}^{T} \gamma B_1(t) e^{-i2\pi \mathbf{k}(t)\cdot\mathbf{r}}\, \mathrm{d}t \tag{1.60}$$

where

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_t^T \mathbf{G}(\tau)\, \mathrm{d}\tau \tag{1.61}$$

is the position in k-space, which represents the number of phase cycles per unit distance in each of the three orthogonal directions.

Further rearrangement of the equation leads to[30]

$$M_{xy}(\mathbf{r}) = i\gamma M_0 \int_{\mathbf{K}} W(\mathbf{k}) S(\mathbf{k}) e^{-i\mathbf{r}\cdot\mathbf{k}} \mathrm{d}\mathbf{k} \tag{1.62}$$

where $W(\mathbf{k})$ is a spatial weighting term

$$W(\mathbf{k}(t)) = \frac{B_1(t)}{|\gamma \mathbf{G}(t)|} \tag{1.63}$$

and $S(\mathbf{k})$ a sampling scheme of $\mathbf{k}$ on the k-space trajectory $\mathbf{k}(t)$ and therefore effectively defines the k-space trajectory. The $|\dot{\mathbf{k}}(t)|$ factor means that the delta function integrates to unity along a unit length path.

$$S(\mathbf{k}) = \int_0^T \delta(\mathbf{k}(t) - \mathbf{k}) |\dot{\mathbf{k}}(t)|\, \mathrm{d}t \tag{1.64}$$

The important consequence of this result can be seen by comparing Equation (1.62) to the definition of the Fourier transform

$$M_{xy}(\mathbf{r}) \propto \mathrm{FT}[W(\mathbf{k}) S(\mathbf{k})] \tag{1.65}$$

which shows that arbitrary excitation localization can be achieved by matching $W(\mathbf{k})S(\mathbf{k})$ to the Fourier transform of the desired excitation localization, either through analytic or iterative means.

1.5.2 The k-space formalism (readout)

Signal localization during the readout requires the introduction of spatial variation in the excited spins in the sample, this is provided by magnetic field gradients inducing spin phase variation across the field of interest. The spatial phase forms the critical link between the time-domain signal, $S(t)$, and the spatial dependence of the signal. In 1D the signal equation, including phase is:

$$S(t) = \int \rho(z) e^{i[\Omega t + \phi(z,t)]} \, \mathrm{d}z \tag{1.66}$$

Where $\rho(z)$ is the density of spins at z, from which we can see that there is a spatial phase dependence. If a linearly varying B field is added (a 'gradient') we get[17]:

$$B(z,t) = B_0 + zG(t), \qquad \omega(z,t) = \omega_0 + \omega_G(z,t) \tag{1.67}$$

with G_z defined as $G_z = \mathrm{d}B_z/\mathrm{d}z$ and $\omega_G(z,t) = \gamma z G(t)$. Assuming that $\omega_0 = \omega_{RF}$ the signal equation reduces to:

$$S(t) = \int \rho(z) e^{i\phi(z,t)} \, \mathrm{d}z \tag{1.68}$$

$\phi(z,t)$ can be calculated as:

$$\phi(z,t) = -\gamma z \int G(t) \, \mathrm{d}t = -2\pi z k(t) \tag{1.69}$$

with $k(t)$ the rate of the phase-change across the sample in units of m^{-1}, $k(t) = \gamma/2\pi \int G(t) \, \mathrm{d}t$, allowing the signal equation to be expressed as:

$$S(k) = \int \rho(z) e^{-i2\pi kz} \, \mathrm{d}z. \tag{1.70}$$

Which matches the definition of the Fourier transform in section 1.3.5 allowing the density of spins ($\rho(z)$) at a given spatial location to be recovered. This is an important equation, which defines the k-space ($k(t)$) formalism underpinning our understanding of all modern MRI pulse sequences. Although this derivation is restricted to 1D, the approach can be simply generalized to 2 or 3 dimensions.

1.5.3 Finite k-space sampling (readout)

The relationship between k-space and real-space defined by Equation (1.70), is equivalent to the continuous Fourier transform, and recovers the spin density (ρ) as a function of position ($\vec{r}$) from the signal (S) as a function of k. In practice we do not have a continuous analytic expression for $S(k)$ as the signal is discretized by the ADC in the RF receive chain, and it would be very time intensive to approximate

a continuous signal from $k = -\infty \longrightarrow \infty$. The signal recorded by the scanner is different to $S(k)$ in two ways, firstly, it is sampled at a regular interval and secondly, only a limited number of samples are recorded. The effect of both of these factors will be discussed below.

We can calculate the effect of sample spacing on the reconstructed image using a comb function made up of Dirac delta functions:

$$\text{comb function} = u(k) = \Delta k \sum_{p=-\infty}^{\infty} \delta(k - p\Delta k) \tag{1.71}$$

The delta function $\delta(x)$ is defined as a function which is zero everywhere $x \neq 0$ and has an area equal to one and Δk is the space between samples.

The result of sampling $s(k)$ with this function is:

$$s_\infty(k) = u(k) \cdot s(k) = \Delta k \sum_{p=-\infty}^{\infty} s(p\Delta k) \cdot \delta(k - p\Delta k) \tag{1.72}$$

By taking the inverse Fourier transform of this we get an expression for $\hat{\rho}_\infty(z)$, the approximation of $\rho(z)$ with discrete k-space sampling:

$$\hat{\rho}_\infty(z) = \Delta k \sum_{p=-\infty}^{\infty} s(p\Delta k) e^{i2\pi p\Delta k z} \tag{1.73}$$

from which it can be seen the reconstructed density is periodic with a period of $1/\Delta k$ giving a field of view of FOV $= 1/\Delta k$. This result is exactly equivalent to the Nyquist criterion.

The effect of only measuring a finite N number of k-space points can also be calculated by considering the multiplication of $s_\infty(k)$ by a rect function, which acts to define the limits of the sampling

$$s_m(k) = s_\infty(k) \cdot \text{rect}\left(\frac{k + \frac{1}{2}\Delta k}{N\Delta k}\right) = \Delta k \sum_{p=-n}^{n-1} s(p\Delta k) \cdot \delta(k - p\Delta k) \tag{1.74}$$

Taking the Fourier transform of this gives

$$\hat{\rho}(z) = \Delta k \sum_{p=-n}^{n-1} s(p\Delta k) e^{i2\pi p\Delta k z} \tag{1.75}$$

By using the well known property that $FT(x(t) \times y(t)) = FT(x(t)) \circledast FT(y(t))$ and that rect Fourier transforms to sinc, the reconstructed signal with finite sampling can be understood as the convolution of the true image with a sinc blurring function i.e.

$$\hat{\rho}(z) = \hat{\rho}_\infty(z) \circledast N\Delta k \,\mathrm{sinc}(\pi N \Delta k z) e^{-i\pi z \Delta k} \tag{1.76}$$

giving a resolution of $1/N\Delta k$.

1.5.4 Non-uniform k-space sampling

Image reconstruction using the DFT assumes that the k-space coordinates are uniformly distributed on a Cartesian grid, however some readout strategies such as spiral readouts (Section 1.7.3) result in k-space sampling without this property. In this case an (inverse) non-uniform FFT (NUFFT) must be used to reconstruct the data. The traditional gridding (inverse) NUFFT consists of four stages, firstly the k-space data is weighted to account for the varying density of k-space points, then gridded on to Cartesian k-space using a convolution kernel, then the inverse FFT is performed and the resultant image is compensated to remove the effect of the convolution kernel. For some trajectories e.g. spiral, the calculation of the density compensation is non-trivial, therefore some more modern reconstructions formulate the problem as a forward NUFFT and minimize the difference between the forward NUFFT of the reconstructed image and the recorded signal, which does not require a pre-calculated density compensation factor .

1.6 RF excitation

1.6.1 Frequency selective excitation

The relationship between the spectral profile of a pulse and the B_1 transmit RF profile is given by the Fourier transform so that $B_1(t) = FT[I(f)]$ where $I(f)$ is the spectral intensity at frequency f. This means that pulses encoding spectral selectivity can be simply designed by taking the Fourier transform of the desired spectral profile. Some examples are shown in Figure 1.8.

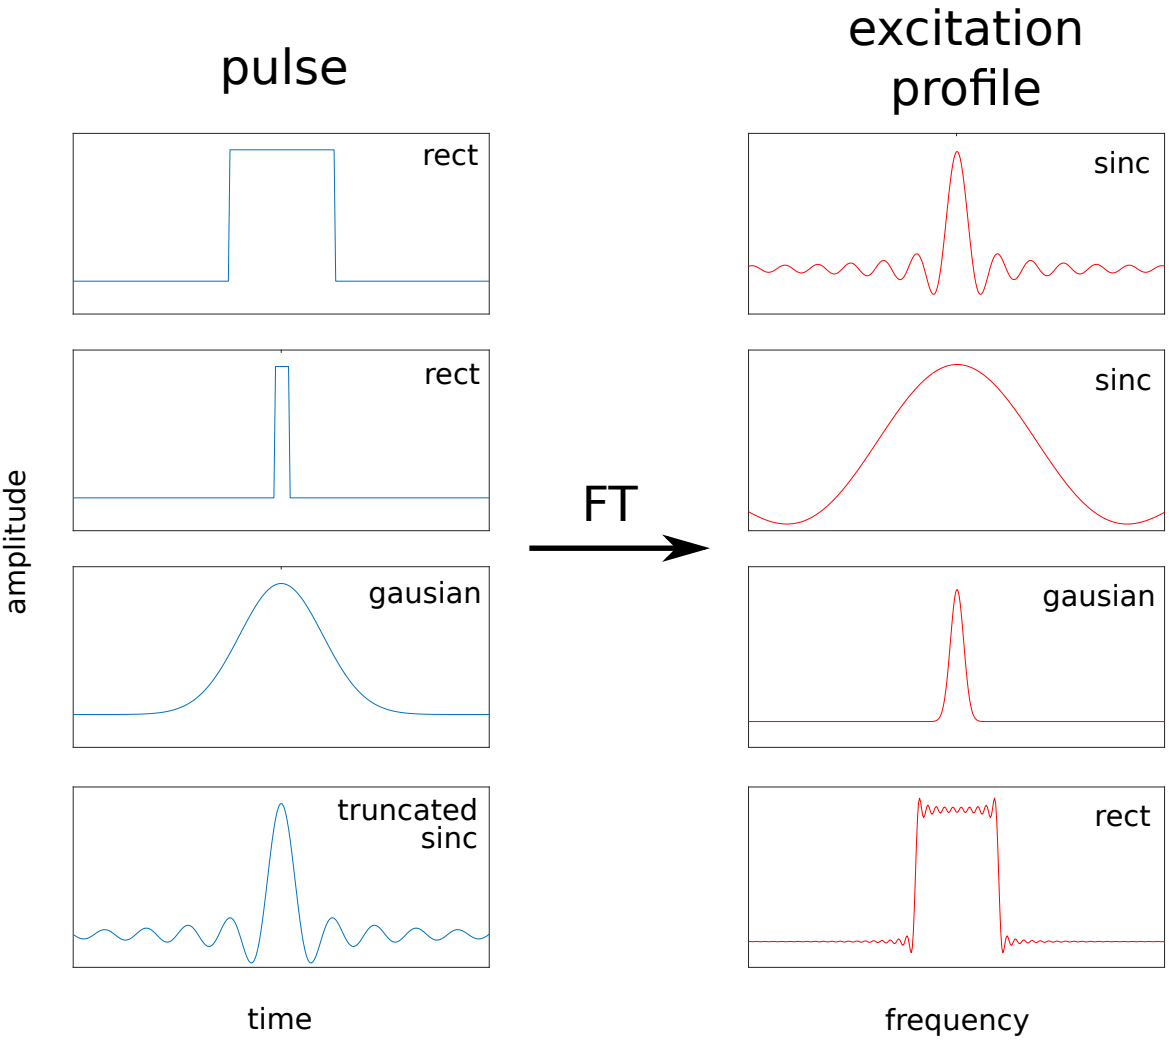

Figure 1.8: Some common pulse shapes and their respective excitation frequency profiles.

1.6.2 Slab select excitation

Slab-selective excitation is achieved using a magnetic gradient field perpendicular to the slab during excitation by a shaped (e.g. sinc) RF pulse.

If the desired slab is perpendicular to the z-axis, i.e. in the x,y-plane, a z-gradient is applied such that[17]:

$$f(z) = f_0 + \frac{\gamma}{2\pi} G_z z \tag{1.77}$$

where $f = \omega/2\pi$ i.e. is frequency in Hz. To excite a slab of thickness Δz centred at z_0 homogeneously, the RF pulse must have even intensity across the range:

$$f_0 + \frac{\gamma}{2\pi}G_z z_0 - \frac{\gamma}{2\pi}G_z\frac{\Delta z}{2} \longrightarrow f_0 + \frac{\gamma}{2\pi}G_z z_0 + \frac{\gamma}{2\pi}G_z\frac{\Delta z}{2} \tag{1.78}$$

and zero elsewhere, giving an RF bandwidth of $\Delta f = (\gamma/2\pi)G_z\Delta z$. Re-arranging this gives a slice thickness of:

$$\Delta z = \frac{\Delta f}{\gamma G_z/2\pi} \tag{1.79}$$

From these equations we know that the spectral profile of the excitation pulse must be proportional to a rect function centred at $f_0 + \frac{\gamma}{2\pi}G_z z_0$, giving:

$$I(f) \propto \text{rect}\left(\frac{f - (f_0 + \gamma G_z z_0/2\pi)}{\Delta f}\right) \tag{1.80}$$

The Fourier transform can then be applied to give the time domain representation of the RF-pulse in the laboratory frame:

$$B_1(t) \propto \Delta f \times \text{sinc}(\pi\Delta f t)e^{-i2\pi(f_0+\gamma G_z z_0/2\pi)t} \tag{1.81}$$

This is effectively an RF-pulse, at the larmor frequency for the centre of the slab, amplitude modulated with a sinc function corresponding to the desired slab thickness. Post excitation, the phase accumulated by the gradient must be undone with a gradient anti-parallel to the first and with an integrated area half that of the first (for symmetric gradients), since the spins only accumulate phase once excited and are excited on average for 1/2 the RF pulse.

The slab selective excitation can also be fully rationalized using the excitation k-space formalism (section 1.5.1) where to produce the desired slice profile the Fourier transform of $[W(\mathbf{k})S(\mathbf{k})]$ must equal a rect function of width Δz. This requires a k-space trajectory ($S(\mathbf{k})$) centered on $k_z = 0$, weighted by a sinc function ($W(\mathbf{k})$) which gives exactly the same gradient and RF forms calculated above. A graphical representation of the excitation can be seen in Figure 1.9.

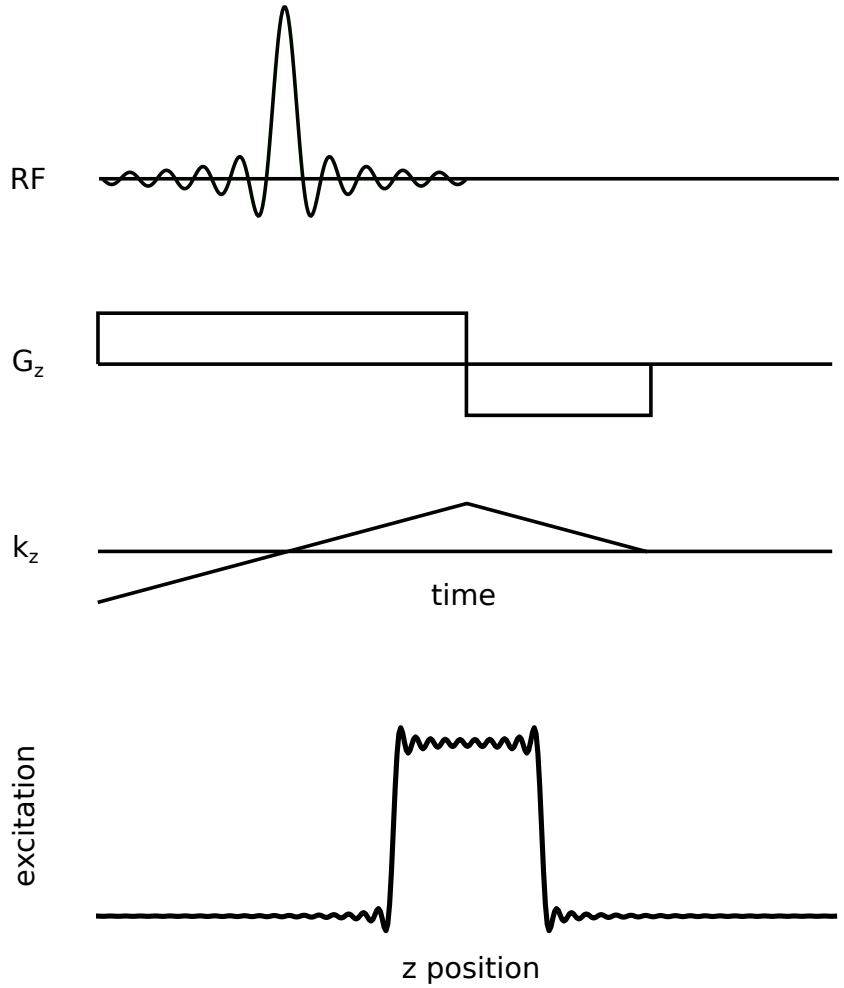

Figure 1.9: Slab selective excitation pulse

1.6.3 Spectral-spatial excitation

By extending the k-space formalism to cover frequency of excitation it is possible to design pulses which are both spectrally and spatially selective. Since frequency and time are related by the Fourier transform, frequency k-space can be defined during excitation as[31]:

$$k_\omega(t) = T - t \tag{1.82}$$

Spatial k-space is then defined as in section 1.5.1

$$\mathbf{k_r}(t) = R \int_t^T \mathbf{G_r}(\tau)\, \mathrm{d}\tau \tag{1.83}$$

Where R is used to account for the differing natural units of spatial and frequency k-space.

Re-deriving the small tip angle solution to the Bloch equations, with the addition of an off-resonance term ω, yields this equation for the derivative of transverse magnetization:

$$\dot{M}_{xy} = -i(\gamma \mathbf{G} \cdot \mathbf{r} + \omega) M_{xy} + i\gamma B_1 M_0 \tag{1.84}$$

solving this equation using the same method as the excitation k-space solution (integrating factor of $\exp(i \int_{\infty}^{t} \gamma \mathbf{G} \cdot \mathbf{R} + \omega \, \mathrm{d}t)$ gives

$$M_{xy}(\mathbf{r}, T) = i\gamma M_0 \int_{-\infty}^{T} B_1(t) e^{-i(\mathbf{k_r}(t) \cdot \mathbf{r} + \mathbf{k}_\omega(t)\omega)} \, \mathrm{d}t \tag{1.85}$$

this can then be re-arranged to

$$M_{xy}(\mathbf{r}, \omega) = i\gamma M_0 \int_{\mathbf{k}_r} \int_{\mathbf{k}_\omega} W(\mathbf{k}_{\mathbf{r},\omega}) S(\mathbf{k}_{\mathbf{r},\omega}) e^{-i\mathbf{k_r} \cdot \mathbf{r}} e^{-i\mathbf{k}_\omega \omega} \, \mathrm{d}k_r \mathrm{d}k_\omega \tag{1.86}$$

where

$$W(\mathbf{k}_{\mathbf{r},\omega}(t)) = \frac{B_1(t)}{|R\mathbf{G}(t)|} \tag{1.87}$$

and

$$S(\mathbf{k}_{\mathbf{r},\omega}) = \int_0^T \{\delta(\mathbf{k_r}(t) - \mathbf{k_r}) |\dot{\mathbf{k}_\mathbf{r}}(t)|\} \{\delta(\mathbf{k}_\omega(t) - \mathbf{k}_\omega) |\dot{\mathbf{k}_\omega}(t)|\} \, \mathrm{d}t \tag{1.88}$$

As for the standard excitation k-space scheme, $S(\mathbf{k}_{\mathbf{r},\omega}(t))$ defines the k-space trajectory (this time in both spatial and frequency dimensions) and $W(\mathbf{k}_{\mathbf{r},\omega}(t))$ defines the RF weighting.

By recognizing that Equation (1.86) is equivalent to a Fourier transform in both k_r and k_ω dimensions we can see that the k-space trajectory and RF weighting must be designed so that the Fourier transform in k_r yeilds the spatial localization required and in k_ω the spectral localization required.

Where a slab and frequency selective excitation is required, this can be achieved by chaining several slab selective excitations together, each of which achieves

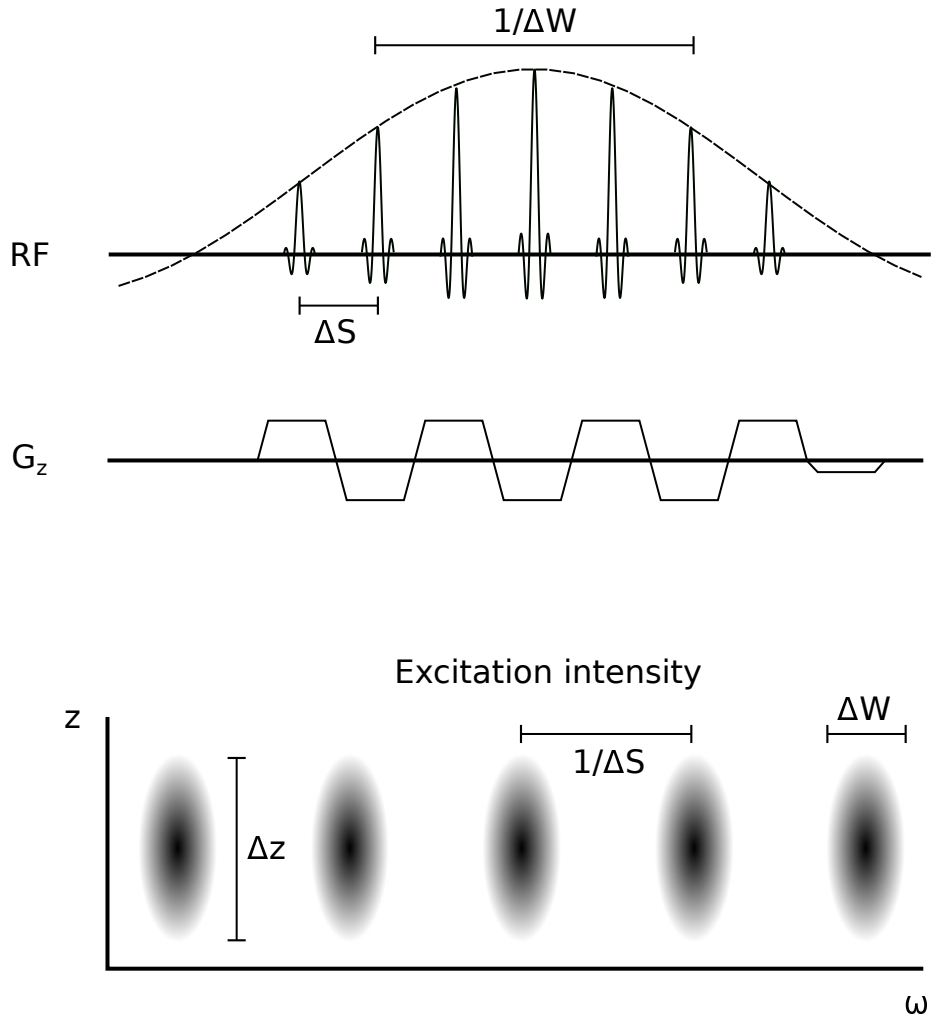

Figure 1.10: Pulse sequence for spectral spatial selective excitation. The sequence excites a slab perpendicular to the z axis of thickness Δz and spectral range ΔW. The individual sinc RF waveforms define the slab thickness ΔW and the overall RF envelope (dotted line) defines the spectral range (ΔW). The sampling of the overall RF envelope by the sinc functions leads to spectral aliasing with a period of $1/\Delta S$.

the spatial localization desired, then weighting the amplitude of the RF pulses by a sinc function to give the spectral localization, albeit aliased by convolution with the periodic RF pulses.

1.6.4 Single voxel excitation

Single voxel excitation, whereby only the volume of interest is excited is used extensively for spectroscopy, as it allows for the FID to be simply analyzed with the basic NMR Fourier method. Examples of single voxel excitation schemes include

ISIS[32], PRESS[33], DRESS[34] or STEAM[35]. Both PRESS and STEAM use similar methods where three slab selective pulses are applied on orthogonal axes, so that only a single central cuboidal region experiences all three pulses. The PRESS sequence consists of 90°–180°–180° pulses, the first 90° pulse tips the magnetization into the x-y plane, which begins to undergo T_2* decay, then the first 180° pulse refocuses the spins and generates a spin echo, which is refocused again by the second 180° pulse to generate a second spin echo. STEAM consists of 90°–90°–90° pulses and generates a stimulated echo at the end of the sequence. Unlike PRESS there is no echo between the second and third pulse, allowing the sequence to be significantly quicker, making it suitable for short T_2 metabolites, however the stimulated echo does mean that there is less signal generated.

1.7 Signal readout

From Equation (1.70) we can see that an image representing spin density can be reconstructed by taking the Fourier transform of received signal as a function of k-space coordinate $S(k)$. Equations 1.73 and 1.76 also show us that to reconstruct an image using the DFT with resolution = R and FOV = F, we must record $N = 1/(R\Delta k)$ k-space points centered on $k = 0$ with spacing $\Delta k = 1/F$.

During the readout phase of the pulse sequence, gradients are applied to manipulate the spin system into states where $S(k)$ can be systematically sampled according to the desired resolution and FOV. A simple sampling scheme could be conceived where each $S(k)$ value is sampled individually by an excitation, 'phase encode', where gradients are played to manipulate the system to a k value, and readout, where the RF signal is recorded. In practice this would be very slow, as the process would need to be repeated many times to reconstruct an image with reasonable resolution, so many schemes exist to increase the speed of acquisition.

1.7.1 Spin warp imaging

Spin warp imaging is a cornerstone of modern MRI pulse sequences, implementing the 'frequency encoding' scheme[36] to significantly accelerate MRI acquisitions

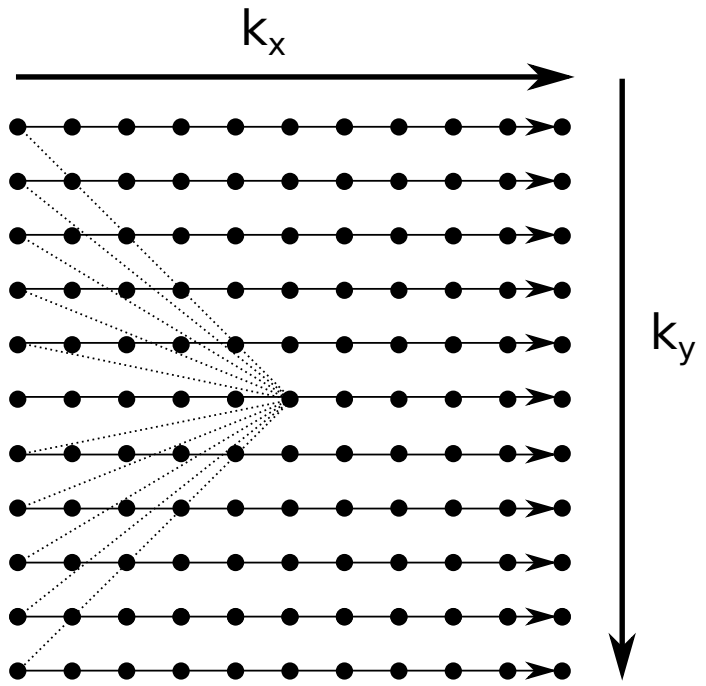

Figure 1.11: k-space scanned out by an multiple spin warp readouts to build up a 2D coverage of k-space

over phase encoding alone. In frequency encoding, a gradient is played during the RF readout, so that $S(k)$ is sampled across a line of k values for each RF excitation (Figure 1.11). This has the effect of drastically reducing acquisition time for a given resolution, however, because only a single value for each $S(k)$ can be recorded, rather than a FID, it does mean that each voxel only gives an overall intensity and not a spectra.

1.7.2 Echo planar imaging

Echo planar imaging (EPI)[37] further accelerates MRI imaging by acquiring all the desired k-space in a single readout. This is achieved by using x and y gradients to 'snake' through k-space, see Figure 1.12. One disadvantage of this technique is that it is susceptible to ghosting artifacts caused by a variety of factors including timing imperfections in the pulse sequence and eddy currents, however this can be ameliorated by a so called 'fly-back' EPI readout.

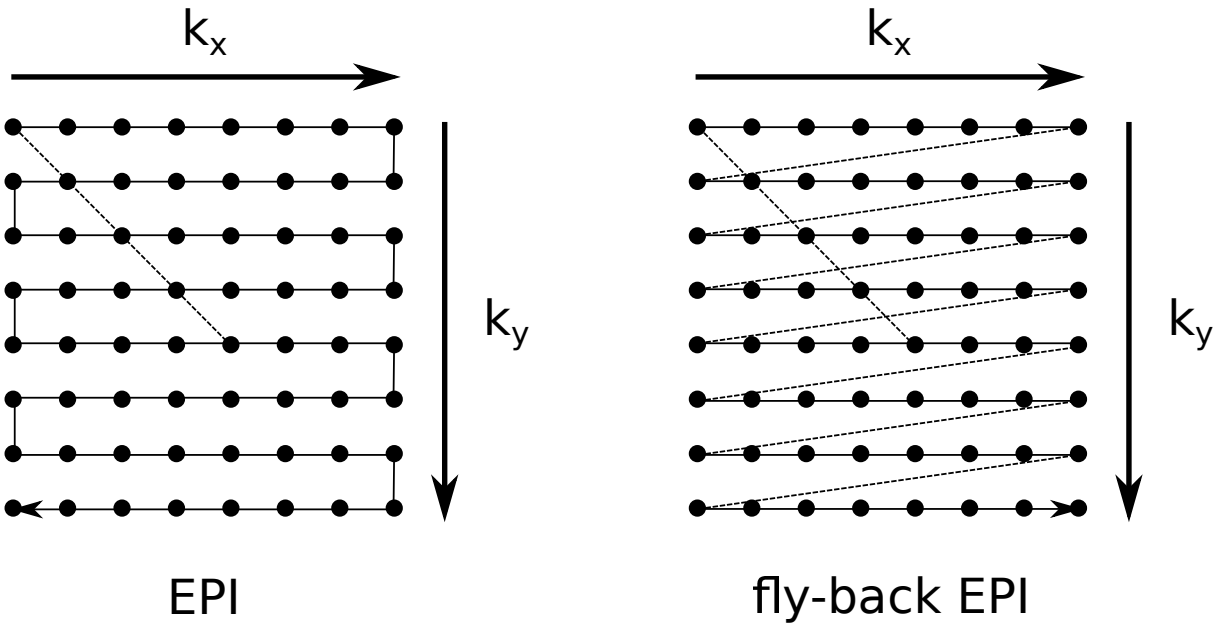

Figure 1.12: k-space scanned out by an EPI readout

1.7.3 Spiral imaging

The previous two readout schemes sample k-space only on a uniform Cartesian grid, corresponding to the points needed to reconstruct the spin density with a fast Fourier transform, however it is possible to reconstruct an image from arbitrary k-space points (subject to FOV and resolution restrictions) using the NUFFT (section 1.5.4). One commonly used non-Cartesian trajectory is the spiral trajectory (Figure 1.13). In spiral readouts, each frequency encode consists of a spiral trajectory, starting from the centre of k-space, according to the polar equation

$$\mathbf{k} = a\theta e^{i\theta}. \tag{1.89}$$

Here $\mathbf{k} = k_x + ik_y$ and

$$a = \frac{\mathrm{d}k_r}{\mathrm{d}\theta} = \frac{\Delta k_{max}}{2\pi} \tag{1.90}$$

where $k_r = |\mathbf{k}|$ and $\Delta k_{max} = 1/\mathrm{FOV}$. If multiple frequency encodes (or 'shots') are to be used, each spiral is rotated by $2\pi/N$, where N is the number of shots, and Δk_{max} is scaled by N. Often, a return trajectory is added to the end of the

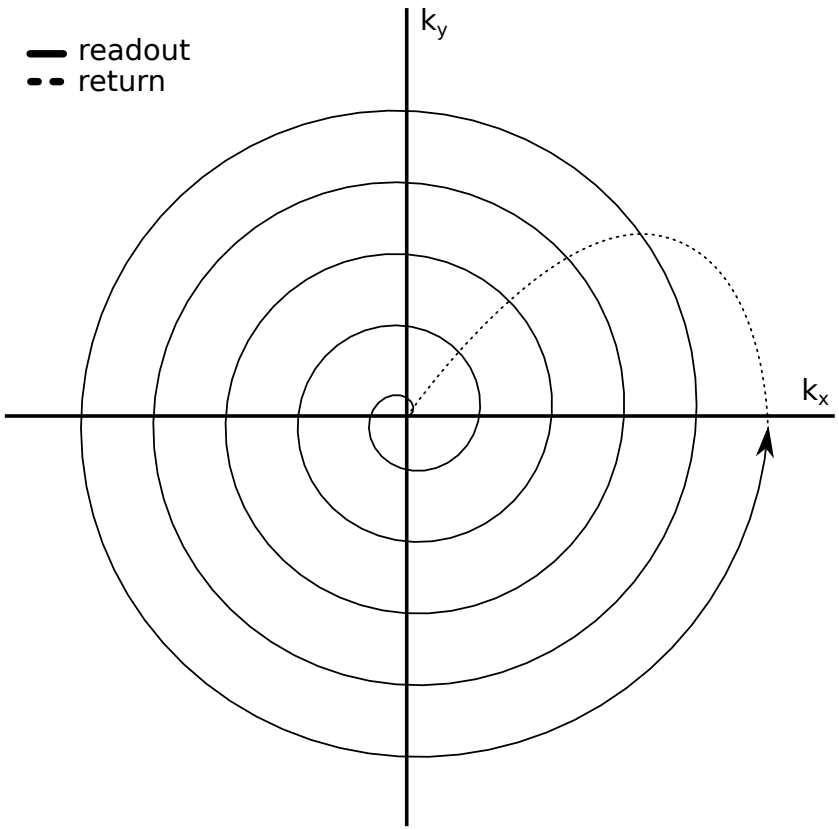

Figure 1.13: k-space scanned out by an spiral readout, dashed section is a return trajectory to prevent interference of accumulated phase with the next acquisition.

acquisition, where a slew rate limited path is plotted back to the centre of k-space, so that there is no accumulated phase left over which could impact the next shot.

Spiral readouts have two advantages over conventional spin warp or EPI readouts, which may be of particular utility for specialist imaging such as hyperpolarized cardiac imaging. First, the centre of k-space, which is the most important region for SNR, is sampled first, whilst signal is highest (if a spin-echo cannot be achieved). Second, the readout is 'symmetric', this results in motion artifacts that manifest as blurring rather than aliasing and minimizes the effects of eddy currents as they approximately cancel themselves out over each rotation.

1.7.4 Chemical shift imaging

Chemical shift imaging adds an extra spectral dimension to the reconstructed spin density to produce localized spectra of the form $\rho(\mathbf{r}, \omega)$. This spectral dimension is orthogonal to the imaging domain, so the k-space formalism remains unchanged,

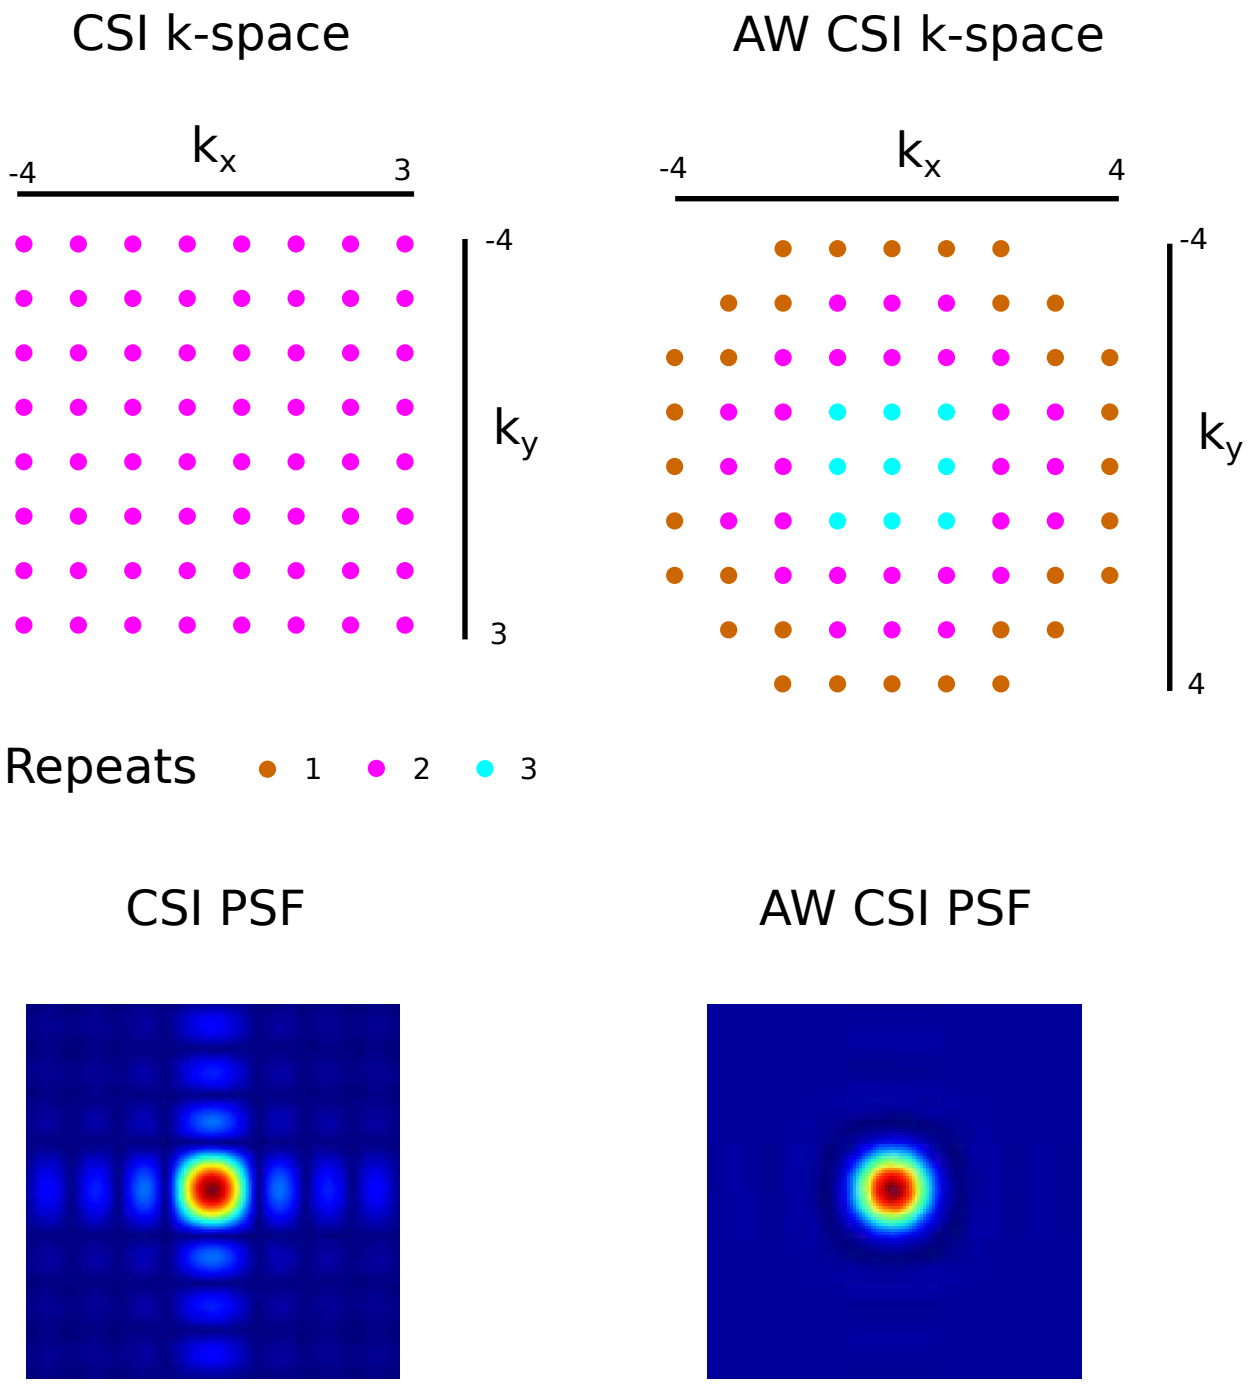

Figure 1.14: k-space phase encodes acquired by a standard CSI acquisition and an acquisition weighted (AW) CSI acquisition with the resultant PSFs.

however a FID (section 1.3.5) is required to reconstruct the spectra which precludes the use of frequency encoding to accelerate imaging[1]. The use of only phase encoding for spatial localization significantly increases acquisition time as each $S(k)$ requires an individual excitation (Figure 1.14). This often results in very low spatial resolution, causing significant contamination of the spectra by neighboring regions, particularly those smaller than one nominal voxel. Acquisition weighting[38] is often used to reduce the extent of the CSI spatial response function and decrease this contamination by weighting the number of repeats of each k-space point by a Hanning window as can be seen in Figure 1.14.

1.8 Hyperpolarized ^{13}C MRI

1.8.1 Dynamic nuclear polarization

The SNR when performing an NMR experiment is directly proportional to the net magnetization of the sample, and therefore polarization of the nuclei. Using Equation (1.25) we can calculate that the nuclear polarization of ^{13}C is only 0.0006% at T=298 K and B_0=7 T. This is due to the very low difference in energy between the excited and ground spin states state ($\approx 5 \times 10^{-26}$J).

This places ^{13}C MRI/MRS firmly in the low SNR regime; the energy gap between spin states for ^{1}H is four times higher and it has a natural abundance of 100% as opposed to 1% for ^{13}C. In addition, carbon based metabolites have a much lower concentration than water in the body (50-100 μM Pyr[39] vs 55 M water in tissues).

DNP overcomes this limitation by transferring polarization from unpaired electrons to ^{13}C nuclei. At $T \approx 1.2K$ (achievable with liquid helium under vacuum) and $B_0 \approx 3.35$ T[2] unpaired electrons have a polarization of ≈95% which can be transferred to the ^{13}C nuclei by irradiation with a microwave source.

Practically, the sample to be polarized consists of a ^{13}C labeled compound and a stable radical, providing the unpaired electrons. This sample is placed into an insert

[1] Unless techniques such as echo-planar spectroscopic imaging, which requires acquisitions at a number of echo times, are used.

[2] B_0 of the α-system polariser used in the experiments described in chapter 2.

in the bore of a superconducting magnet and cooled by the magnet's cryogenic liquid helium in order to achieve a high electron polarization. The sample is then irradiated by a microwave source (near the electron Larmor frequency), to promote the transfer of polarization from the electrons to the ^{13}C. Once the maximum ^{13}C nuclear polarization is reached, a high temperature/pressure pH buffer solution is used to dissolve the sample and force it out of the hyperpolarizer into a collection chamber. The hyperpolarized sample is then rapidly transferred to the MRI experiment[4].

Possible mechanisms for the transfer of polarisation in a DNP experiment include; the cross effect, solid state effect and thermal mixing [40]. For our experimental setups (prototype hyperpolarizer[4], Spin Lab, HyperSense) the dominant effect is thermal mixing[4]. Thermal mixing can be rationalized as equalization of spin temperature of the electron radicals and the ^{13}C nuclei[41].

1.8.2 Imaging metabolism

The power of hyperpolarized ^{13}C as a technique comes from the small molecule metabolites which can be labeled with ^{13}C, injected into a volunteer and tracked through MRSI, giving us immediate information on the metabolic state of the volunteer. Many small molecule ^{13}C labeled metabolites have been proposed[42], and proved their utility in animal models, however at present only [1-^{13}C]pyruvate has been widely injected into human volunteers and patients.

Pyruvate is a product of glycolysis and a precursor to the citric acid cycle which can be ^{13}C labeled in all three carbon positions. When labelled in the C-1 position to form [1-^{13}C]pyruvate, the ^{13}C label propagates to [1-^{13}C]alanine, [1-^{13}C]lactate and $^{13}CO_2$ which is released upon the conversion of pyruvate to acetyl-CoA. The $^{13}CO_2$ then forms a pH dependent equilibrium with ^{13}C-bicarbonate. This allows quantification of pyruvate dehydrogenase (PDH), lactate dehydrogenase (LDH) and alanine aminotransferase (ALT) flux and the size of the pyruvate, lactate and alanine pools to be estimated (see Figure 1.15). The clinical application of hyperpolarized

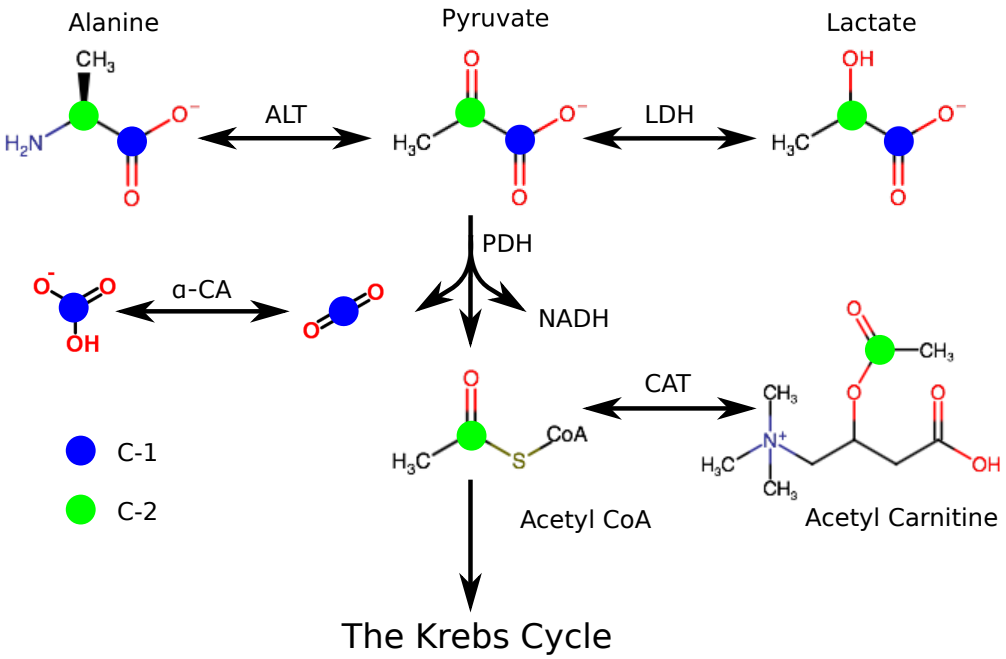

Figure 1.15: The metabolic pathway of pyruvate with the C-1 and C-2 positions highlighted.

[1-^{13}C]pyruvate is a highly active area of research, with the replacement of FDG-PET for cancer screening[43] and assessment of cardiac injury from myocardial infarction[44] forming the main avenues at present.

The mechanism by which hyperpolarized [1-^{13}C]pyruvate is used to screen for cancers has a similar basis to FDG PET. The Warburg effect, which is seen in the majority of cancers, results in a shift from oxidative phosphorylation to glycolysis as the primary means of ATP production in cancer cells. This provides more material for anabolism, however it also increases glucose demand and production of pyruvate. The increased level of glycolysis is detected by a conventional PET scan by increased uptake of ^{18}F labeled (deoxy)glucose to the cancerous cells, however hyperpolarized [1-^{13}C]pyruvate detects instead the increase in pyruvate-to-lactate flux and resultant increase in the size of the lactate pool[43, 45]. This potentially provides greater information on the metabolic state of the tumor and response to intervention[46].

For cardiac hyperpolarized [1-^{13}C]pyruvate scans, it is the ^{13}C-bicarbonate signal, which correlates with aerobic respiration, that is of interest. Early proof of principal studies have shown that in the healthy heart, the wall of the left

ventricle produces a strong ^{13}C-bicarbonate signal[47], that is diminished in regions of ischemia, post-myocardial infarction (Figure 1.16)[44]. This has the exciting potential to improve non-invasive imaging of ischemia in the heart and enable viability testing of heart tissue (as viable tissue must be alive and therefore have PDH flux) before re-vascularization procedures[48].

Both of these research avenues for the applicability of hyperpolarized [1-^{13}C]pyruvate for clinical use have the potential to greatly enhance patient care, with the added bonus that no radiation dose is conferred, unlike other comparable techniques such as CT, FDG-PET or SPECT, however significant challenges still remain to implementation.

1.8.3 Dynamic nuclear polarisation - the challenges

Although dynamic nuclear polarisation presents a wealth of opportunities for studying metabolism *in vivo* and improving patient care, it also has unique challenges as a result of the acquisition being taken in a thermodynamically unstable state.

At 3T, the T_1 relaxation time for a [1-^{13}C]pyruvate solution is approximately 60 seconds[43], meaning that the full imaging experiment must be conducted in approximately two minutes, including any spectroscopy to calibrate the resonant frequency for imaging or B_0/B_1 maps. Furthermore, the observed signal will not reach a maximum immediately, as the bolus of hyperpolarized [1-^{13}C]pyruvate must reach the imaging site, so timing of the acquisition to the signal maximum is critical. The enhanced polarisation is also finite and every RF excitation uses up a portion, making it important to only excite in the imaging region, so that the hyperpolarization of the circulating ^{13}C is not depleted. Together, these limitations call for significantly different priorities when designing pulse sequences and acquisition schemes one of which will be investigated in this thesis.

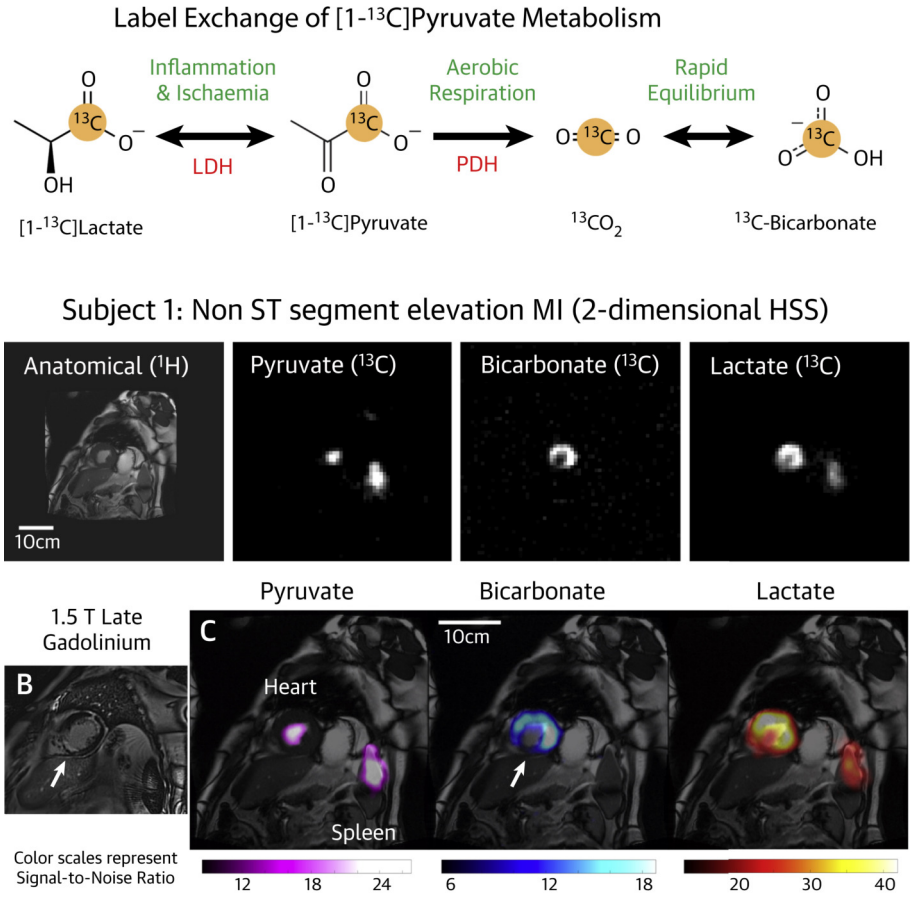

Figure 1.16: Hyperpolarized [1-^{13}C]Pyruvate scan of a patient with myocardial infarction, images depicting the [1-^{13}C]pyruvate, ^{13}C-bicarbonate and [1-^{13}C]lactate signals are shown, with a clear decrease in bicarbonate signal in the vicinity of the infarct. The metabolic pathway for [1-^{13}]Pyruvate in the heart is shown above. Figure reprinted with permission[44].

1.9 ^{31}P MRI

1.9.1 ^{31}P nuclear magnetic resonance

Acting as the 'energy currency' of the body, the production of ATP is critical to the heart's ability to function. ATP is consumed to power both the physical contraction of the heart, and the movement of ions to coordinate its contraction. Heart disease is associated with a fall in the efficiency of ATP production, and therefore a fall in ATP and its storage medium PCr[11]. Both PCr and ATP contain phosphorus atoms in their molecular structure, the only stable, naturally occurring, isotope of which is ^{31}P. Fortunately, this isotope has a nuclear spin of $+1/2$, allowing

it to be detected relatively simply with MRI.

During the initial stages of heart failure the heart attempts to maintain ATP concentration, resulting in a shift in the equilibrium between PCr and ATP, and correspondingly a drop in PCr concentration is observed (30-70%)[11]. It does this because each of the energy consuming processes in the heart require ATP to have a minimum ΔG of hydrolysis, that depends on the ATP concentration (section 1.2.1). Later in heart failure, the ATP concentration decreases as well, but not by the same extent (30-40%)[11]. Conveniently this allows us to calculate a ratio of PCr:ATP, which will correlate with the PCr energy reserve in the heart, where all the unknown scanner based contributions to the measured metabolite signals (e.g. coil sensitivity) cancel out. This metric has been used extensively for clinical cardiac ^{31}P studies[49–51], and has been shown to strongly correlate with cardiac mortality[52].

1.9.2 Spectroscopic imaging of high energy phosphates

The principal challenge associated with ^{31}P MRI is achieving adequate SNR within a time-frame suitable for patient scans. There are three main reasons for this; firstly ^{31}P has a low gyro-magnetic ratio ($\approx 2.5\times$ lower than ^{1}H) which based on the quadratic relationship of SNR to γ (section 1.3.6) results in a $6\times$ decrease in SNR for an otherwise identical experiment. Secondly the human heart is 74% water by mass compared to only 0.11% for phosphorus[53], resulting in a further $\approx 2000\times$ reduction in signal compared to ^{1}H MRI. Finally, short TRs, to allow more rapid repeats, are not possible due to the long T_1 times of phosphorus metabolites (PCr $T_1 = 3.8$ s at 3T[54]), as saturation of the ^{31}P resonances would occur (using the Ernst angle[55] formula, the optimum TR for a 45° flip angle on PCr is 1.3 s). Furthermore, performing localized spectroscopy has the effect of increasing acquisition time, as it precludes the use of frequency encoding to accelerate imaging.

^{31}P pulse sequences therefore represent a compromise, to keep acquisition times reasonable (upwards of 20 minutes per scan is common). Typically localization fidelity is sacrificed, either by using a single voxel technique such as STEAM (PRESS has too long a preperation time for short ^{31}P T_2^*s) or by selecting a

low CSI resolution. Cardiac gating may also be disabled, reducing localization further, and the nuclear Overhauser effect (Section 1.9.4) exploited to increase the polarization of the ^{31}P nuclei.

1.9.3 Calculating the cardiac PCr/ATP ratio

In the ideal experiment the PCr/ATP ratio would be calculated by dividing the area of the PCr resonance by the area of one of the ATP's three resonances. Numerous practicalities exist however which make this more challenging and will be discussed in detail below.

1.9.3.1 Voxel placement

The aim when performing cardiac ^{31}P spectroscopy is to obtain a spectra representative of the myocardium. For this to be the case the voxel must be far enough away from the chest wall to reduce contamination from the strong skeletal muscle PCr resonance, but still representative of the overall myocardium. OCMR studies have shown the midseptal voxel (Figure 1.17), which is centered over the septum to be the optimal placement[56, 57], with deviation from this position having a large effect on measured PCr/ATP ratio. This makes the correct positioning of the CSI grid and selection of the midseptal voxel critical for producing valid results.

1.9.3.2 Saturation correction

The low SNR of cardiac ^{31}P spectroscopy makes TRs of less than the metabolites long T_1 relaxation times (T_1 PCR = 3.8 s, T_1 γ-ATP = 2.4 s[54]) necessary to complete the scans in reasonable time frames. This results in partial saturation of the metabolite resonances, which is different for each metabolite due to the varying T_1 relaxation times, making corrections for this effect necessary.

The equilibrium transverse magnetization, taking into account partial saturation, for each resonance can be calculated with the formula

$$\frac{M_{xy}}{M_0} = \frac{(e^{TR/T_1} - 1)\sin\alpha}{e^{TR/T_1} - \cos\alpha} \tag{1.91}$$

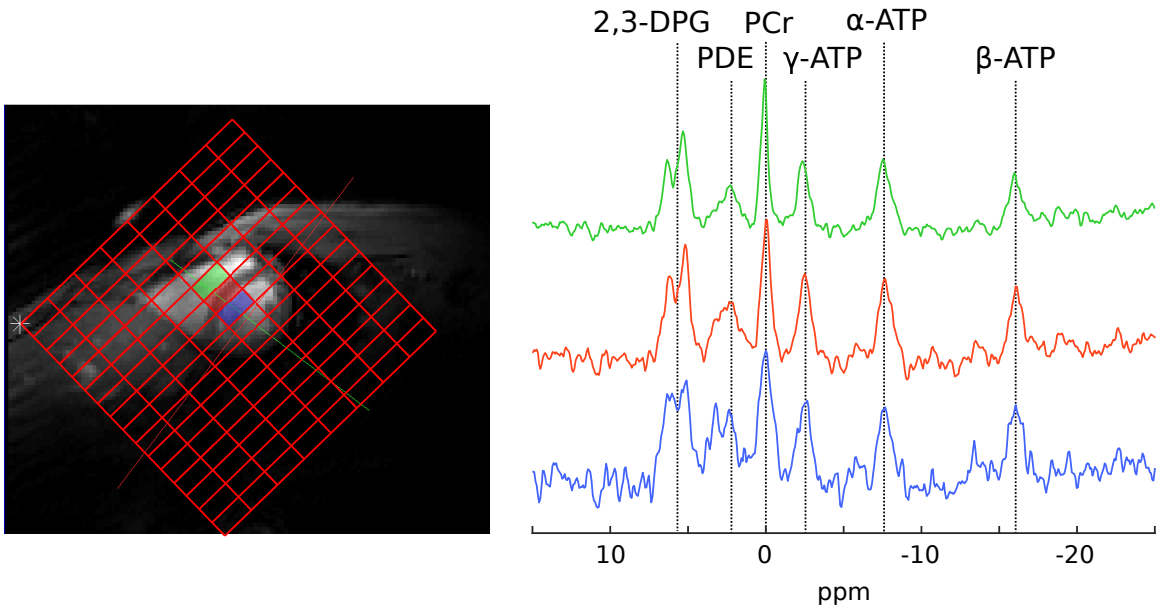

Figure 1.17: The CSI grid overlaid on a localiser scan of a volunteer's heart is shown on the left. The midseptal voxel, placed on the septum is highlighted in red. The spectra generated by the green, red and blue voxels are shown on the right. Each of the resolved resonances are labeled (2,3-diphosphoglycerate, phosphodiester, phosphocreatine, γ,α,β-adenosine triphosphate). Notice the marked difference in PCr/ATP ratio of each spectra.

where α is the flip angle and TR the repetition time. This formula is derived by simplifying the Ernst equations[55] with the assumptions that $TR \gg T_2^*$ and the excitation is exactly on resonance. Thereby allowing this effect to be corrected.

Unfortunately the requirement to know the actual flip angle of the excitation, which varies significantly on a per-subject basis, requires additional scans as part of the protocol.

First the flip angle in a phenylphosphonic acid fiducial is calculated from a series of inversion recovery experiments, then the Biot-Savart law is used to calculate B_1 field maps which are co-registered with the CSI acquisition using a number of fiducials implanted in the coil to allow the flip angle in the midseptal voxel to be calculated.

1.9.4 NOE correction

Depending on the experimental setup, enhancement of the ^{31}P magnetization by the nuclear Overhauser effect (NOE) effect may need to be corrected. The NOE effect refers to a process where protons may transfer their magnetization to ^{31}P

nuclei, to which they are J-coupled, by cross-relaxation when saturated with RF radiation[58]. The ^{31}P nuclei responsible for each resonance is in a differing chemical environment and therefore the degree of NOE enhancement will vary, so it is essential to correct for this where used.

1.9.4.1 Blood correction

Blood contamination of ^{31}P cardiac spectra is very challenging to avoid, since the chambers of the heart are filled with blood. The ^{31}P spectra of human blood contains PDE, 2,3-DPG, ATP and ADP but not PCr[59], whereas the myocardium does not contain 2,3-DPG, which functions to help haemoglobin give up oxygen to oxygen consuming tissues[60]. This makes the removal of the blood contamination relatively straight forward, because the ratio of 2,3-DPG to ATP in blood is relatively fixed, the blood correction can be made by subtracting a proportion (11%) of the 2,3-DPG resonance area from the ATP resonance area.

1.9.4.2 ATP resonance

Unlike PCr, ATP has three ^{31}P atoms, each of which is in a chain with oxygen linkers (Figure 1.18). From the adenosine these are labeled α, β and γ, the α and γ ^{31}P atoms, each have one ^{31}P neighbor and therefore produce doublets in the spectra, whereas the β ^{31}P atom has two, so produces a triplet with a 1:2:1 ratio of intensities. Depending on experimental setup, either the average of all three ATP resonances may be used when calculating the PCr/ATP ratio[54], or just the γ resonance, due to NADH contamination of the α resonance and non-uniformity of the spectral profile of the excitation pulse[56].

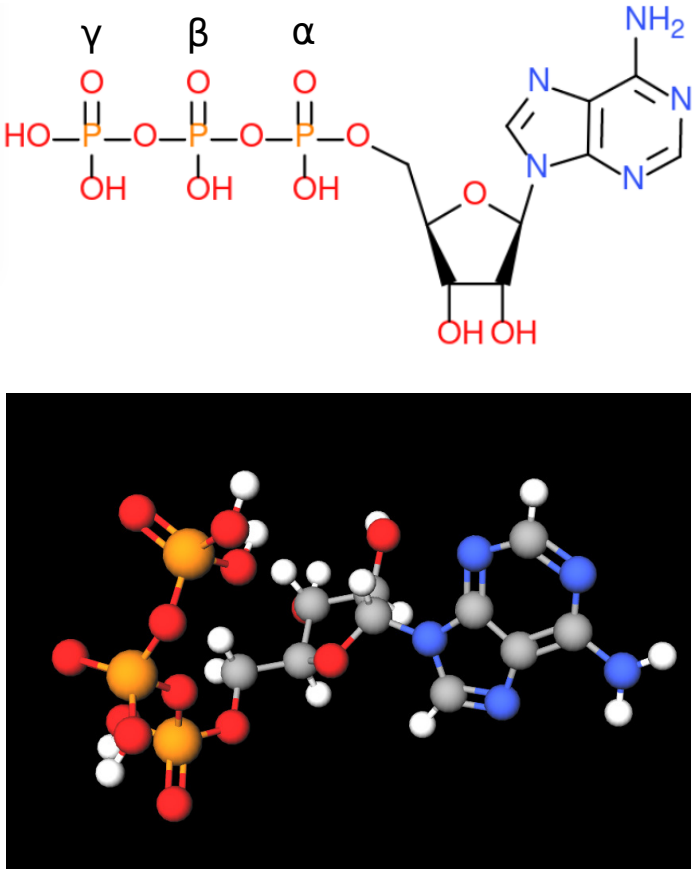

Figure 1.18: The 2D and 3D representations of the molecular structure of ATP, the three phosphate groups, with α, β and γ positions labeled are shown on the left and the adenosine group on the right. 2D representation drawn with ChemDoodle and 3D representation drawn with MolView.

2

Hyperpolarized ^{13}C spiral imaging

Contents

Much of the work contained in this chapter has been published as a full paper, where I was credited as first author, in Magnetic Resonance in Medicine titled 'A 3D hybrid-shot spiral sequence for hyperpolarized ^{13}C imaging'[61] and lead directly to the paper 'Proof-of-Principle Demonstration of Direct Metabolic Imaging

Following Myocardial Infarction Using Hyperpolarized ^{13}C CMR' published in JACC: Cardiovascular Imaging, where I was credited as a co-author.

2.1 Introduction

Hyperpolarized [1-^{13}C]pyruvate is a versatile metabolic probe, which has been used extensively to investigate metabolism and pH in health and disease in vivo[62–64], and is progressing rapidly towards clinical translation [43, 47, 65–68]. Many conditions where metabolic dysregulation is of interest are spatially localized, and therefore it is desirable to image at sufficiently high spatial resolution to appreciate potential metabolic heterogeneities within the tissue, for instance in the myocardium post infarction to assess viability when planning intervention. It is also desirable to perform time resolved imaging to observe temporal metabolic dynamics, which may be significantly altered in pathology, for instance, in ischaemic cardiomyopathies[69]. Furthermore, a growing body of work advocates for the use of time-course data for the determination of metabolic parameters, such as k_{PL} (pyruvate-lactate), either via the fitting of kinetic models[70, 71] or calculation of area-under-the-curve ratios[72], necessitating high quality time-courses with sufficient temporal resolution to properly characterize the shape of the curve. In addition, owing to the inherent physiological variability in the time between the injection of hyperpolarized pyruvate and its arrival in the organ of interest, the use of dynamic imaging strategies is preferred, to ensure that the acquisition does not miss the initial bolus of injected substrate [73]. It is therefore desirable to have both high spatial and temporal resolution in the same acquisition. In practice, there is a trade-off between spatial and temporal resolution, which thus far must be predefined *a priori* at the time of acquisition.

The finite, decaying, and non-renewable magnetization generated by hyperpolarization necessitates rapid, magnetization-efficient imaging techniques. One common strategy is to combine the magnetization-efficiency of spectral-spatial RF excitations with rapid spiral readouts [74–79]. The first hyperpolarized ^{13}C images of the human heart were acquired with a single-shot spiral trajectory [47]. Single-shot spiral-out readouts efficiently encode a plane of k-space with minimal dead time by starting

at the center of k-space with the energy-rich low spatial frequencies. However, for a given field of view (FOV), the maximum spatial resolution that can be achieved with a single-shot spiral is limited by the readout duration allowed by the magnetic field inhomogeneity of the sample[80] and the T_2^* of the hyperpolarized nuclei in question.

One approach to overcome the limiting effect of T_2^* is to use a multi-shot spiral readout [81], which segments the acquisition of k-space into multiple readouts, each temporally shorter and therefore less affected by T_2^* decay than an equivalent single-shot spiral readout. However, as each segment does not typically meet the Nyquist criterion, it is not possible to reconstruct images individually from each interleaf at the prescribed FOV without aliasing. Another strategy to shorten the readout is to design variable density spiral (VDS) trajectories [82], where the sampling density is not constant across k-space. Spiral-out VDS designs with a sampling density that decreases as a function of the distance away from the center of k-space have been used to reduce artifacts due to motion [83] and aliasing [84]. However, there is an SNR penalty associated with VDS designs due to noise amplification from the weighting of the varying sampling density in the image reconstruction process.

In this chapter, the feasibility of adaptive spatiotemporal imaging is demonstrated, using a hybrid-shot spiral (HSS) trajectory which features a constant single-shot density inner region to satisfy the Nyquist criteria up to a minimum desired spatial resolution, followed by a smooth transition to a constant N-shot density sub-Nyquist outer region. By truncating each readout to the inner single-shot region (single-shot HSS), low spatial resolution images can be reconstructed for each TR. By combining non-truncated data which encompass a full set of rotation angles from N acquisitions of the same metabolite, alias-free higher spatial resolution images can be reconstructed (full HSS) at lower temporal resolution. This adaptive spatiotemporal feature of the HSS acquisition allows the trade-off between spatial and temporal resolution to be made *post hoc* at reconstruction time.

2.2 Theory

2.2.1 Spiral trajectory algorithm

Numerous strategies for designing spiral trajectories exist in the literature. Given upper limits for gradient amplitude and slew rate, analytical formulae can be derived to design constant density [85] and variable density [86] spiral trajectories. A hybrid-density trajectory has been proposed for functional MRI with a single-shot inner region up to a fraction of the maximum design k-space radius beyond which the sampling density decreases [87]. In this work, the target HSS trajectory (shown in Figure 2.1) consists of two connected regions of constant sampling density with a transition that is smooth but relatively abrupt as opposed to a gradual change in sampling density common with VDS designs using the aforementioned algorithms.

Iterative algorithms compute each point of the trajectory by solving a coupled set of differential equations with the previous point as the initial condition. A widely used example of an iterative VDS design algorithm represents the sampling density as an effective FOV which is a function of k-space radius [88, 89]. In this algorithm the spiral trajectory is defined as

$$\mathbf{k} = a\theta e^{i\theta} \tag{2.1}$$

where $\mathbf{k} = k_x + ik_y$, θ is the polar angle and a is the rate of increase of k-space radius ($k_r = |\mathbf{k}|$) with respect to θ, which is defined by

$$a = \frac{\mathrm{d}k_r}{\mathrm{d}\theta} = \frac{N\Delta k_{max}}{2\pi} = \frac{N}{2\pi \cdot \mathrm{FOV}} \tag{2.2}$$

If FOV is allowed to vary as a function of k-space radius, a variable density spiral can be constructed. Differentiating Equation (2.1) gives an equation for the gradient vector ($\mathbf{G}$)

$$\mathbf{G} = \frac{2\pi a}{\gamma} e^{i\theta}(\dot{\theta} + i\theta\dot{\theta}) \tag{2.3}$$

Which can be further differentiated to give an equation for the slew rate ($\mathbf{S}$)

$$\mathbf{S} = \frac{2\pi a}{\gamma} e^{i\theta}[(\ddot{\theta} - \theta\dot{\theta}^2) + i(2\dot{\theta}^2 + \theta\ddot{\theta})] \tag{2.4}$$

By taking the magnitude of Equation (2.3) and setting $|\mathbf{G}| = G_{\max}$, the gradient limited value of $\dot{k}_r$ can be found

$$\dot{k}_{\max} = \sqrt{\frac{(\gamma\, G_{\max})^2/4\pi^2}{1 + \left(\frac{k_r}{a}\right)^2}} \tag{2.5}$$

Similarly, solving Equation (2.4) in the same way, gives an expression for the slew rate limited ($|\mathbf{S}| = S_{\max}$) value of $\ddot{k}_r$

$$\ddot{k}_{\max} = \frac{-B}{2A} + \sqrt{\frac{B^2}{4A^2} - \frac{C}{A}} \tag{2.6}$$

Where A, B, and C are the quadratic coefficients used to solve the equation

$$A = 1 + \frac{k_r^2}{a^2} \tag{2.7}$$

$$B = \frac{2k_r\dot{k}_r^2}{a^2} \tag{2.8}$$

$$C = \left(\frac{k_r \cdot \dot{k}_r^2}{a^2}\right)^2 + \frac{4\dot{k}_r^4}{a^2} - \gamma^2 S_{\max}^2/4\pi^2 \tag{2.9}$$

This gives the following equation for $\ddot{k}_r$, subject to the dual constraints of $G_{\max}$ and $S_{\max}$

$$\ddot{k}_r = \begin{cases} \ddot{k}_{\max} & \dot{k}_r < \dot{k}_{\max} \\ \left(\dot{k}_{\max} - \dot{k}_r\right)/\Delta t & \dot{k}_r \geq \dot{k}_{\max} \end{cases} \tag{2.10}$$

Where Δt is the scanner dwell time. $\ddot{\theta}$ can then be found with the expression

$$\ddot{\theta} = \frac{1}{a} \cdot \ddot{k}_r \tag{2.11}$$

To calculate the trajectory, $\dot{\theta}$ and $\dot{k}$ are incremented by $\ddot{\theta}$ and $\ddot{k}$ at each timepoint with θ and k also incremented at each timepoint by the new values of $\dot{\theta}$ and $\dot{k}$. All variables are initialised to 0 at the first timepoint and the algorithm is terminated when k_r reaches the value corresponding to the desired resolution.

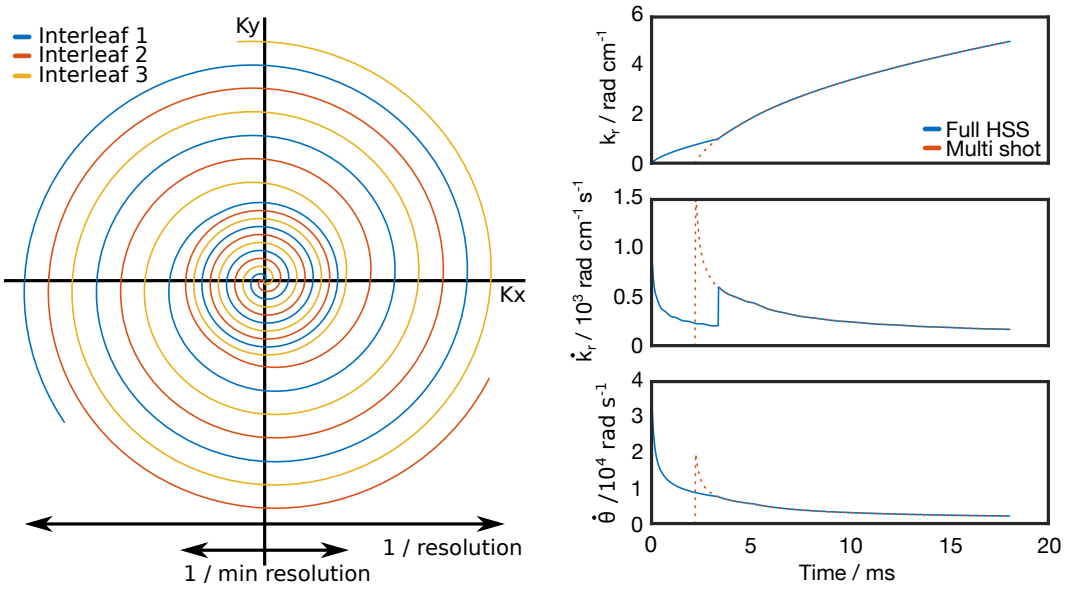

Figure 2.1: Left: The proposed k-space readout trajectory (experimental parameters as in the digital phantom experiment). In this case three shots are required for a full resolution reconstruction, however an image can be reconstructed every shot at a lower spatial resolution.
Right: k_r, $\dot{k}_r$ and $\dot{\theta}$ values as a function of time for the shown trajectory superimposed with an aligned multi-shot spiral with the same imaging parameters.

2.3 Pulse sequence algorithms

2.3.1 Main HSS spiral

The iterative VDS algorithm (described in section 2.2.1) was modified to design HSS trajectories with a transition from single-shot to N-shot sampling density at a k-space radius corresponding to the desired minimum resolution for a high temporal resolution reconstruction.

In its modified form, the algorithm begins as for a single shot spiral, then once the trajectory reaches the k-space radius corresponding to the single-shot HSS resolution, k_{trans}, the FOV parameter is scaled by a factor of $1/N$, reducing the sampling density to that of a N-shot spiral, until the end of the readout. At the transition between the two density regions, $\dot{k}_r$ is set equal to $\dot{k}_{\text{trans}}$, the value of $\dot{k}_r$ for an N-shot spiral at k_{trans}. The algorithm is now in a state identical to an N-shot spiral at $k_r = k_{\text{trans}}$, as shown in Figure 2.1, resulting in a smooth transition between the two densities without violating $G_{\max}$ or $S_{\max}$. The pseudo-code describing this

is shown in algorithm 1. The algorithm was modified from an existing VDS design algorithm[88] and implemented in the C programming language.

$N \leftarrow N$;
$\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, $\theta \leftarrow 0$;
while $k_r < k_{trans}$ **do**
  calculate next $\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, θ;
end
$\dot{k}_{\text{trans}} \leftarrow \dot{k}_r$;
$\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, $\theta \leftarrow 0$;
$N \leftarrow 1$;
while $k_r < k_{trans}$ **do**
  calculate next $\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, θ;
  output $\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, θ;
end
$N \leftarrow N$;
$\dot{k}_r \leftarrow \dot{k}_{\text{trans}}$;
while $k_r < k_{max}$ **do**
  calculate next $\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, θ;
  output $\ddot{k}_r$, $\ddot{\theta}$, $\dot{k}_r$, $\dot{\theta}$, k_r, θ;
end

Algorithm 1: Pseudo-code describing the algorithm used to generate the hybrid shot spiral trajectories in this chapter

The algorithm begins by calculating $\dot{k}_{\text{trans}}$ (the first while loop), before reseting and calculating the single shot trajectory up to $k_r = k_{\text{trans}}$ (the second while loop), at which point N is set to the number of shots in the multi-shot region and $\dot{k}_r$ is set to $\dot{k}_{\text{trans}}$. After which, the algorithm proceeds to calculate the multi-shot portion of the trajectory (the third while loop).

2.3.2 Rewind

The rewind algorithm was designed to return the gradient amplitude **G** and k-space coordinate **k** to zero quickly by computing gradient waveforms for each orthogonal dimension subject to the constraint that the slew-rate $S_{x/y/z} = \pm S_{\max}$ where $S_{\max}$ is a scalar value.

Based on the constraint that $S_{x/y/z} = \pm S_{\max}$ there are two possible generalized gradient waveforms (Figure 2.2) one where the gradient crosses zero and one where

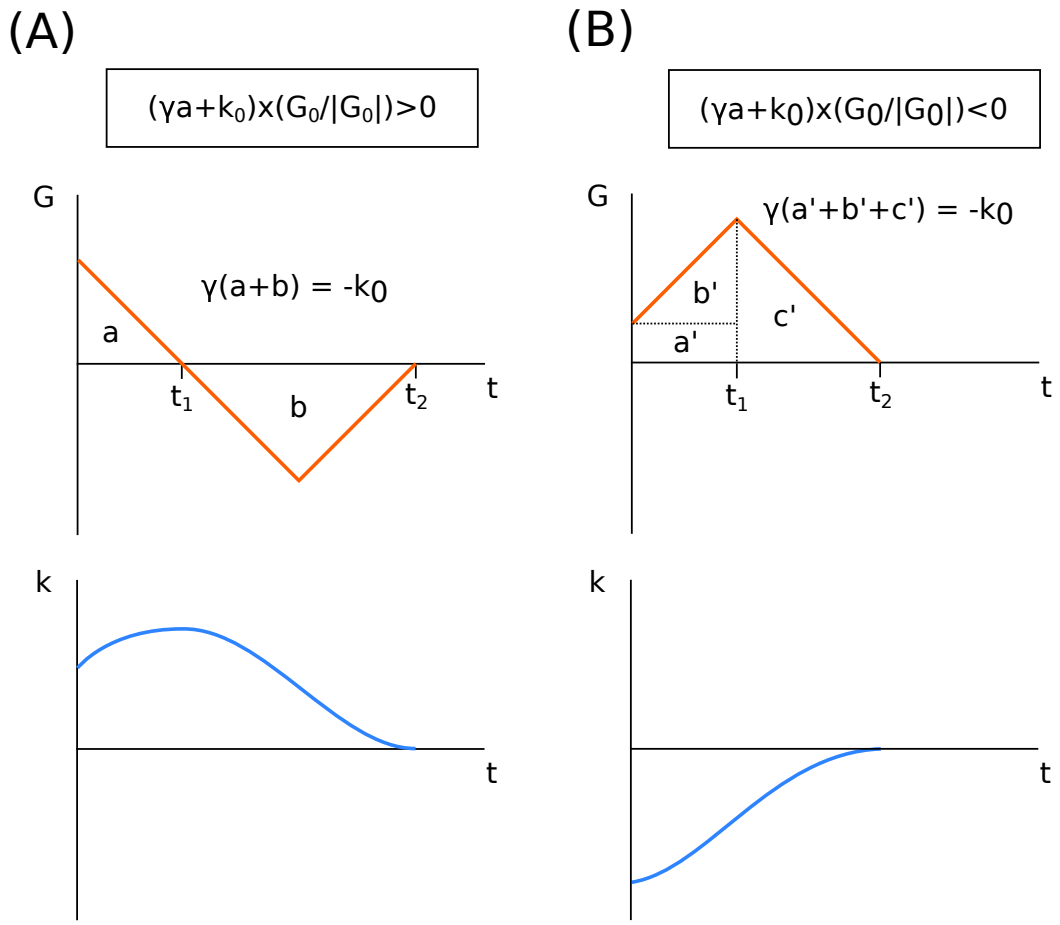

Figure 2.2: The return trajectory for each orthogonal direction is calculated independently. One of two slew-rate limited algorithms are used to calculate the return trajectory, where G crosses 0 (A) and where it does not (B).

it does not. The required waveform can be deduced by calculating the k-space coordinate if the gradients were brought back to zero as quickly as possible ($k_{\min}$), if $k_{\min}$ has the same sign as G_0 the gradient must cross zero (Figure 2.2a) otherwise if the signs are different it must not (Figure 2.2b). Using geometric relations the criterion to select the correct algorithm can be calculated as

$$\text{criterion} = \frac{\gamma}{2}\frac{G_0}{S_{\max}} + \frac{k_0}{|G_0|} \tag{2.12}$$

Where if criterion > 0 the crossing algorithm is used, otherwise the non-crossing algorithm is used.

In both cases the gradient trajectory is defined by the starting values of $k_{x/y/z}$ and $G_{x/y/z}$, $S_{\max}$, and two timepoints t_1 and t_2 which are illustrated in Figure 2.2. For the

crossing algorithm the t_1 value is the slew-rate limited time to bring $G_{x/y/z}$ to zero

$$t_1 = \frac{|G_0|}{S_{\max}} \tag{2.13}$$

and t_2 is calculated such that the trapezium it defines brings $k_{x/y/z}$ to zero giving

$$t_2 = t_1 + \sqrt{2\left(\frac{|G_0|}{S_{\max}}\right)^2 + \frac{4k_0}{\gamma G_0}\frac{|G_0|}{S\max}}. \tag{2.14}$$

For the non crossing case, two simultaneous equations are formed, based on the requirement that at $t = t_2$, $G_{x/y/z} = 0$ and $k_{x/y/z} = 0$. The first equation, based on the requirement that the gradient returns to zero is

$$G_0 + \frac{G_0}{|G_0|}S_{\max}t_1 - \frac{G_0}{|G_0|}S_{\max}t_2 = 0 \tag{2.15}$$

which simplifies to

$$t_2 = t_1 + \frac{|G_0|}{S_{\max}} \tag{2.16}$$

and the second, based on $k_{x/y/z} = 0$ at $t = t_2$ is

$$k_0 + \gamma G_0 t_1 + \gamma\frac{S_{\max}}{2}t_1^2 + \gamma\frac{S_{\max}}{2}(t_2 - t_1)^2 = 0 \tag{2.17}$$

which simplifies by substituting in the first equation to give

$$S_{\max}\frac{G_0}{|G_0|}t_1^2 + 2G_0 t_1 + \frac{G_0}{|G_0|}\frac{G_0^2}{S_{\max}} + \frac{k_0}{\gamma} \tag{2.18}$$

which can be solved using the quadratic formula, yielding the trajectory.

2.3.3 Pulse sequence software structure

An overview of the pulse sequence generating functions are shown in Figure 2.3. The main pulse sequence file calculates the whole pulse sequence including RF excitation and phase encoding, however it uses ancillary functions to calculate non-standard readouts. The `calc2DSpiralGradients` function synthesizes the readout gradients by calling `calc_vds` and `calc_return` which each implement the HSS spiral and return algorithms respectively.

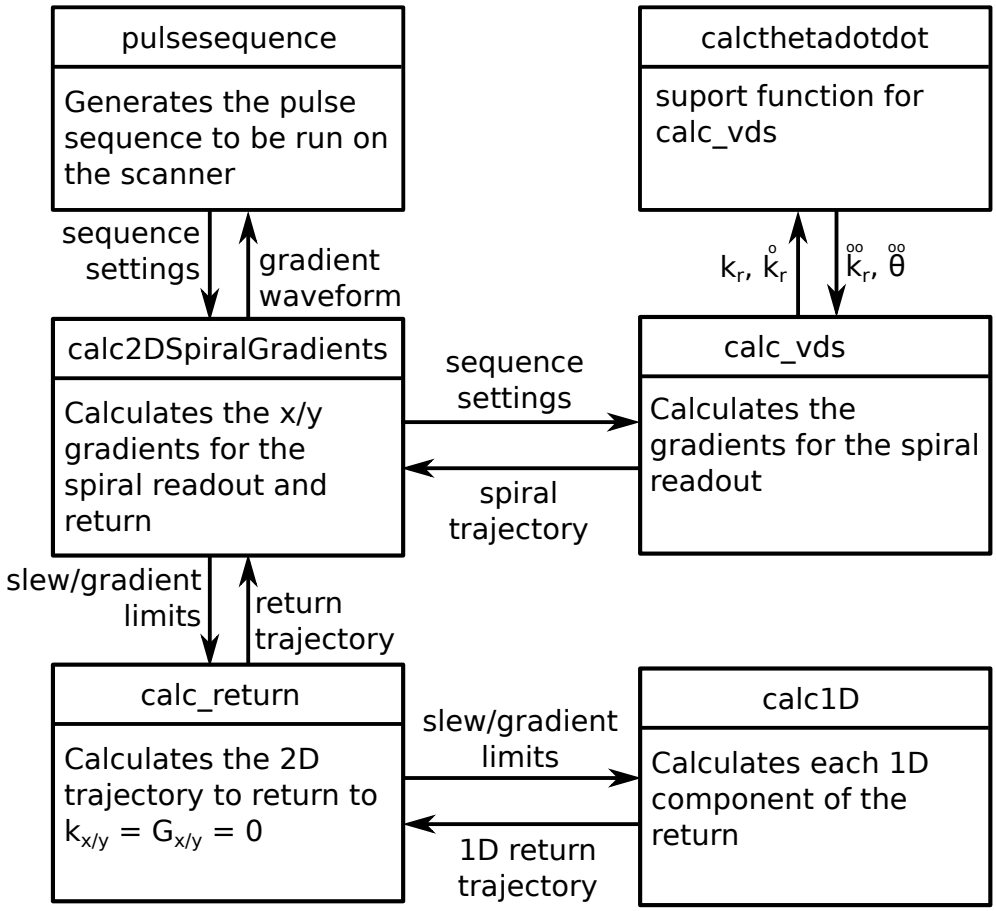

Figure 2.3: An overview of the functions making up the pulse sequence program. See Figure 2.6.

2.4 Methods

2.4.1 Sequence design

The proposed sequence is a 3D hybrid-shot spiral with spectral spatial excitation, centric phase encodes in the z-direction and a hybrid spiral readout with rewind. A simplified representation of the sequence is shown in Figure 2.4.

A spectral spatial excitation scheme was used to provide spectrally selective excitation within the volume of interest, while minimizing excitation outside of the volume of interest, in order to preserve the pool of hyperpolarized magnetization circulating in the animal, as our coil could not provide enough spatial localization by itself. As described previously [90], the spectral spatial excitation pulses were designed under the RF excitation k-space formalism with zipper correction [91].

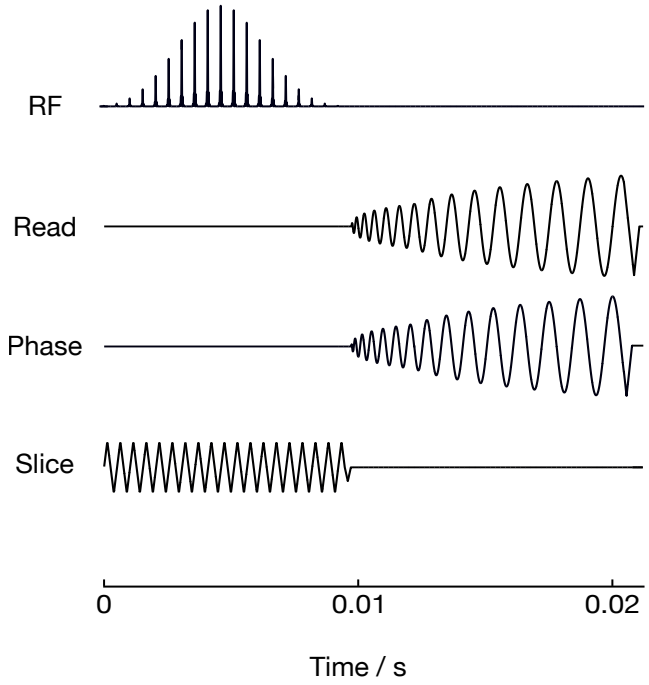

Figure 2.4: A simplified pulse sequence diagram for an HSS in vivo acquisition showing the spectral-spatial excitation (0–0.01 s), z phase encode (0.01 s, slice) and HSS readout (0.01 s – end).

Sub-pulses under a Gaussian envelope were designed with the Shinnar-Le-Roux algorithm [92], with a time-bandwidth product of 3 in the spectral domain and 9 in the spatial domain, and zero amplitude on negative gradient lobes to produce a fly-back pulse. The scheme was originally designed to be slew-rate ($S_{\max}$) limited rather than gradient ($G_{\max}$) limited, as we have empirically found at both 7T and 3T that waveforms played at $S_{\max}$ were reproduced with significantly better fidelity than those at $G_{\max}$[90], when measured by the method proposed by Duyn *et al.* [93].

The pulse was 9 ms long, and possessed 37 sub-lobes each of duration 246 µs, corresponding to an excitation bandwidth of 240 Hz, and a stop-band of ∼2 kHz, sufficient to avoid every unwanted resonance in the spectrum of hyperpolarized [1-^{13}C]pyruvate at 7T, with an excitation bandwidth sufficiently large (3 ppm) to provide relative immunity to transmitting off-resonance and B_0 inhomogeneity. The slab thickness was 45.5 mm.

In the proposed hybrid-shot spiral k-space encoding scheme, a single-shot and N-shot spiral were compounded together and rotated by $360/N^\circ$ so that the central

region of k-space was fully sampled every shot and the outer region was fully sampled every N shots. This allowed for a low spatial resolution image to be reconstructed every shot (single-shot HSS), giving maximum temporal resolution, or for N consecutive shots to be combined giving a higher spatial resolution, albeit at the cost of temporal resolution (full HSS). The spiral readouts were rotated by $360/N°$, as opposed to the golden angle, in order to maximise the SNR of the reconstructed image when N shots were combined[94]. The sequence was implemented using a modified version of the Hargreaves variable density spiral algorithm[88] to generate the two spirals which are smoothly transitioned between.

Figures showing the console output for the sequence and oscilloscope traces of gradients are shown in figures 2.5 and 2.6.

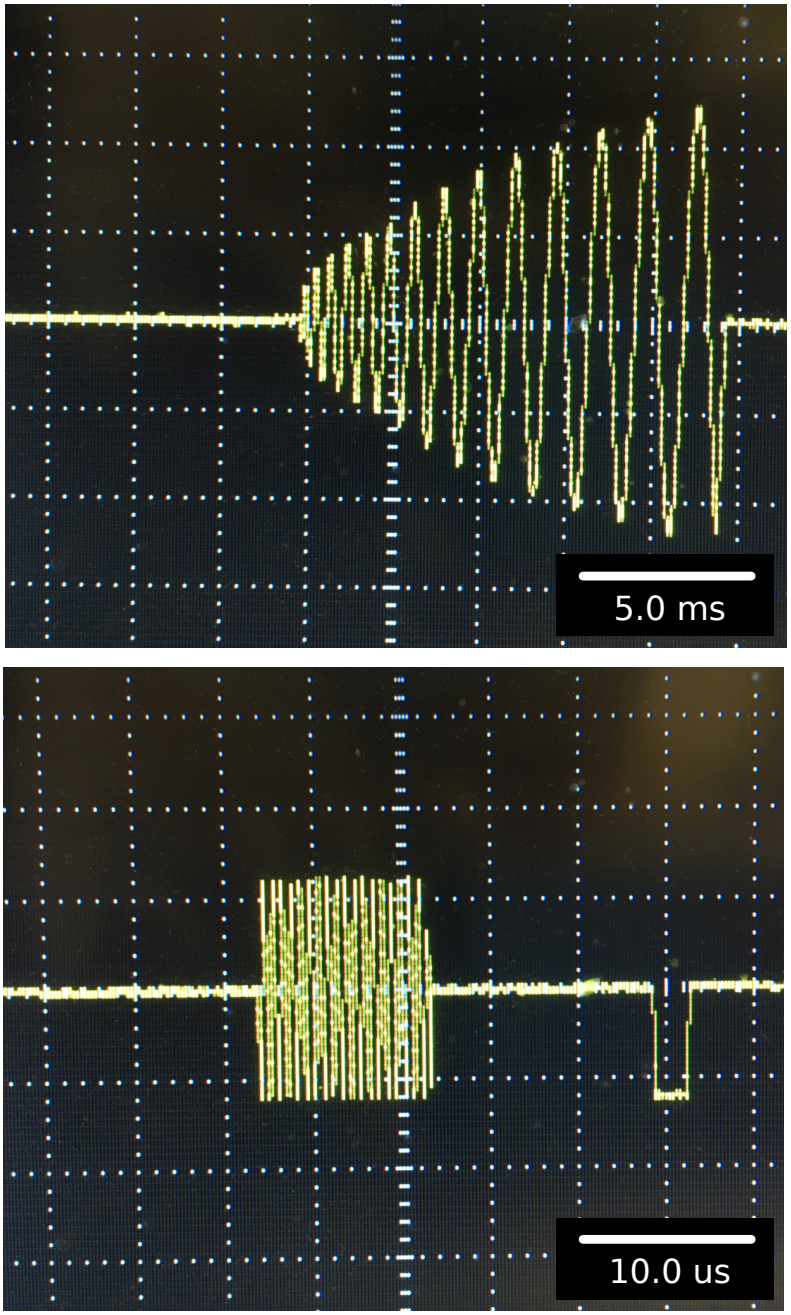

Figure 2.5: Oscilloscope traces showing the x (top) and z (bottom) gradient amplitudes recorded while playing the sequence. The x component of the spiral readout can be seen on the x trace, while the spectral-spatial gradient and phase encode can be seen on the z trace. The depicted z trace immediately precedes the x trace as shown in Figure 2.4.

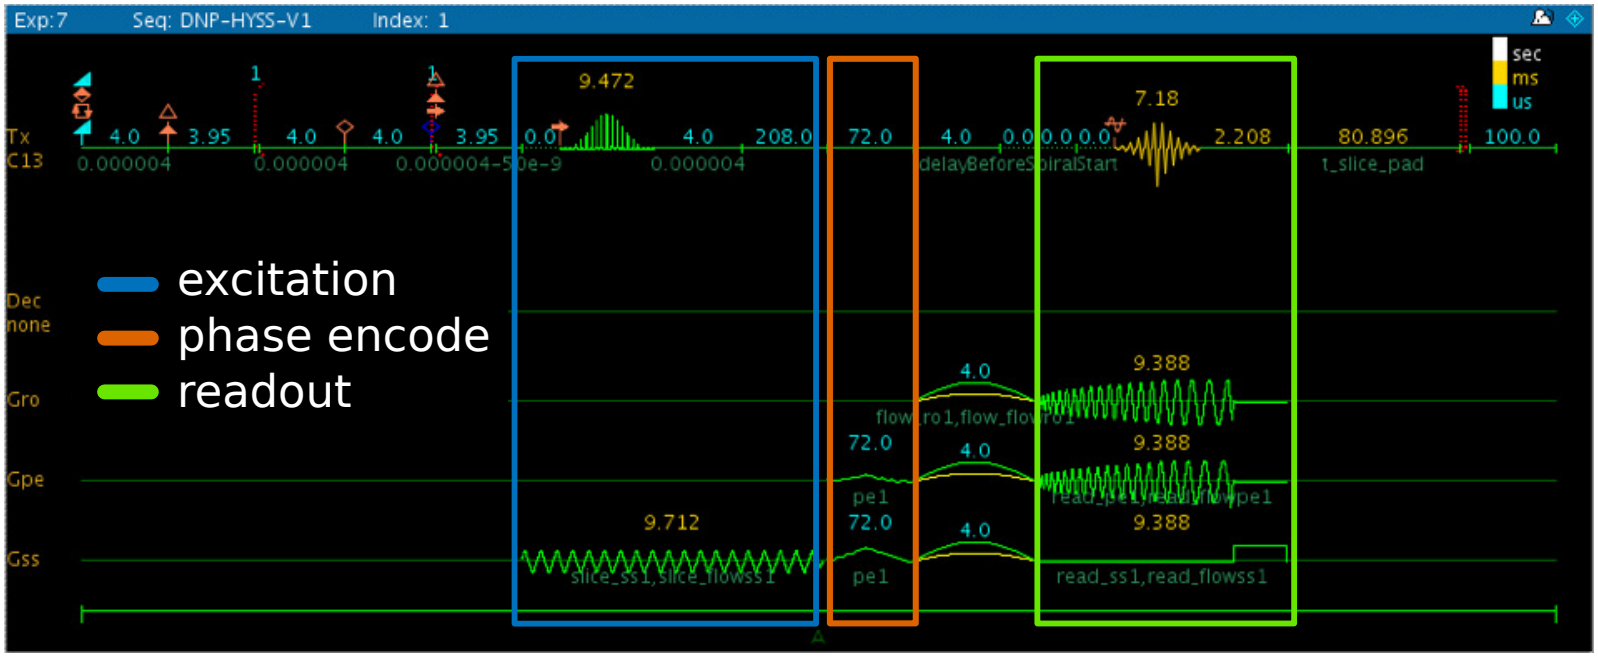

Figure 2.6: Screenshot of the MRI console showing the principal components of the sequence. The Gpe (y) phase encode shown has zero amplitude in this sequence, but was included for flexibility if a different readout scheme were required.

2.4.2 Image reconstruction

Image reconstruction was performed with custom software written in the Matlab programming language (Mathworks, Natick, MA, 2018b). It used the Berkeley Advanced Reconstruction Toolbox [95] implementation of the non-uniform Fourier transform (NUFFT) to reconstruct the data in the x-y plane. The k-space trajectory used in the NUFFT calculation was computed by correcting the prescribed gradients with a previously measured gradient impulse response function[96] and numerically integrating the result.

In this work the bulk B_0 shift (which causes a 'blurring' artifact in image space or 'ringing' in the PSF[97] and is hard to avoid, since the pyruvate center frequency is not known accurately ahead of time) was corrected by a method similar to that used by Lau et al.[98], where the images were reconstructed with a range of different demodulation frequencies and the least blurred image taken, however in this work the selection was done manually, rather than with an automated approach. While this demodulation was necessary from an image quality point of view it also provided an estimate of the actual ^{13}C center frequency, validating that this parameter was set correctly at acquisition.

2.4.3 Digital phantom

To provide initial validation of the sequence, single-shot, multiple-shot and HSS ^{13}C acquisitions were simulated with a synthetic brain phantom[99]. The simulated scanner had $G_{\max} = 500\,\mathrm{mT/m}$, $S_{\max} = 536\,\mathrm{mT/m/ms}$ and a bandwidth of 62.5 kHz. The multiple-shot and HSS trajectories had three interleaves. The trajectories were calculated to achieve an x/y resolution of 1 mm (single-shot HSS 5 mm) and field of view of 60 mm. After calculation, the trajectories were truncated to the length of the multiple-shot spiral (15.9 ms) for comparison with equal readout time, giving resolutions for each method of: single-shot HSS 5 mm, full HSS 1.1 mm, single-shot spiral 1.9 mm and multiple-shot spiral 1 mm. The T_2^* of the simulated phantom was chosen to be 19 ms as a representative value for hyperpolarized

Figure 2.7: Image showing the ^{1}H resolution phantom used in this chapter.

pyruvate in vivo and the simulated excitation flip angle, for depletion of the pool of hyperpolarized magnetization, was 3°.

Two factors affecting the quality of the reconstructed images were simulated: receive off resonance and noise. Receive off resonance was simulated by multiplying the simulated FID by a phase term corresponding to 0-100 Hz off resonance. Noise was simulated by injecting random Gaussian noise, with a FWHM of 0-10% the maximum FID intensity, into the simulated FID prior to reconstruction. A structural similarity index (SSIM) was calculated for each reconstruction, using the ground truth phantom as a reference, to assess the reconstructed image quality.

2.4.4 ^{1}H phantom

Further validation of the sequence was performed with ^{1}H phantom measurements. A resolution phantom (Figure 2.7) was constructed by adding Lego pieces to a tube filled with a Ferumoxytol solution which had been titrated to give a T_2^* of 12.1 ms (elemental iron concentration of 30 µg/mL). The titration was performed by measuring the T_1 and T_2^*, with inversion recovery and gradient echo sequences, of Ferumoxytol solutions that had elemental iron concentrations of 0, 15 and 30 µg/mL.

The scanner was a Varian 7T DDR system with $G_{max} = 1000$ mT/m, $S_{max} = 5000$ mT/m/ms, coupled with a 72 mm dual-tuned proton/carbon birdcage volume coil. A 5° sinc pulse with duration 1000 µs and slab thickness 45.5 mm was used for the excitation and the bandwidth was set to 250 kHz.

Single-shot, multiple-shot and HSS 3D acquisitions were acquired for comparison. The multiple-shot spiral and full HSS trajectories each had 3 interleaves. The trajectories were designed with a field of view of $80 \times 80 \times 45.5$ mm and resolution of $0.5 \times 0.5 \times 1.4$ mm with the single-shot HSS reconstruction giving a resolution of $2.5 \times 2.5 \times 1.4$ mm. The trajectories were then truncated to the same readout time (14.8 ms), corresponding to a full HSS resolution of 0.8 mm (multi-shot spiral 0.7 mm, single-shot spiral 1.3 mm). 32 z phase encode steps were used.

In addition to the spiral acquisitions, a 3D gradient echo sequence was used to acquire a high resolution ground truth image. The FOV was $80 \times 80 \times 45.5$ mm and the resolution $0.4 \times 0.4 \times 1.4$ mm. SSIM indices were calculated for each spiral reconstruction over a region of interest (indicated in Figure 2.10) and an intensity profile was plotted through the phantom for each acquisition. Additionally the HSS spiral was reconstructed without the GIRF correction and using a conventional gridding NUFFT algorithm[100] to provide validation of the reconstruction method using real data.

2.4.5 Preclinical imaging

Preclinical proof-of-concept experiments were performed on a Varian 7T DDR system with $G_{\max} = 1000$ mT/m, $S_{\max} = 5000$ mT/m/ms, with an actively detuned transmit/surface receive setup consisting of a 72 mm dual-tuned proton/carbon birdcage volume coil with a 40 mm two-channel ^{13}C surface receive array with an integrated preamp (Rapid Biomedical GmbH, Rimpar, Germany).

Approximately 40 mg of [1-^{13}C]pyruvic acid doped with 15 mM OX063 trityl radical and 3 µL of a 1-in-50 dilution of Dotarem (Guerbet Laboratories Ltd) gadolinium chelate was polarized and dissoluted in a prototype polariser as described previously [4]. Male Wistar rats were anaesthetised with 2% isoflurane in 90% O_2, 10% N_2O, placed prone onto the imaging coil in a home-built preclinical imaging cradle, and injected with 2 ml of 80 mM hyperpolarized [1-^{13}C]pyruvate over approximately 20 s.

A thermally polarized 5 M ^{13}C urea phantom was included next to the animal to provide a carbon frequency reference. Three dimensional, cardiac-gated volume

shimming was performed using the previously mapped behavior of the magnet's shim set and a spherical harmonic cardiac auto-shimming algorithm as described previously [101]. The ^{13}C center frequency was set by identifying the carbon resonant frequency of the urea phantom, and applying a constant offset of 296 Hz (3.95 ppm). This consisted of the B_0 variation between the urea phantom and the heart (estimated from previously performed ΔB_0 maps), the known chemical shift difference from ^{13}C urea to [1-^{13}C]pyruvate, and empirical optimizations. The ^{13}C center frequency was verified during reconstruction as described in the image reconstruction section.

Once injected with the hyperpolarized [1-^{13}C]pyruvate solution, three rats were imaged using the proposed sequence, consisting of a spectral spatial excitation and a cardiac gated (one shot per RR interval) 3 interleaf HSS readout scheme. The imaging parameters, for the first two rats, consisted of a readout bandwidth of 250 kHz, a field of view of $120 \times 120 \times 45.5$ mm, 12 z phase encode steps, a resolution of $2 \times 2 \times 3.8$ mm and a single-shot HSS resolution of $10 \times 10 \times 3.8$ mm, the nominal minimum TR (time for 1 rotation angle for one metabolite) before gating was 212.4 ms. The excitation flip angles were: pyruvate 3°, bicarbonate 20°, and lactate 20°. One of these rats was injected with hyperpolarized [1-^{13}C]pyruvate a second time in a single session, and imaged with a single-shot spiral with the same FOV and resolution as the full HSS acquisition and a nominal TR of 352.0 ms before gating. The third rat was imaged with a readout bandwidth of 62.5 kHz, a field of view of $60 \times 60 \times 45.5$ mm, 12 z phase encode steps, a resolution of $1 \times 1 \times 3.8$ mm and a single-shot HSS resolution of $5 \times 5 \times 3.8$ mm.

After the hyperpolarized imaging, a set of anatomical images were acquired using a gradient echo cine sequence described previously [102], with the same slice thickness and FOV as the hyperpolarized data.

2.5 Results

2.5.1 Digital phantom

The results of the digital phantom experiment are summarized in Figure 2.8. With no injected noise or simulated off resonance, the full HSS and multiple-shot spiral reconstructions have a similar SSIM metric when using the ground truth as a reference image. The single-shot reconstruction had a lower SSIM than either the multiple-shot or full HSS reconstructions and the single-shot HSS reconstruction had the lowest SSIM. This broadly follows the theoretical resolution achieved by each readout in the acquisition time, although the multiple-shot spiral would have been expected to perform better than full HSS, without the influence of $T_2{}^*$, the depletion of the hyperpolarized magnetization pool and oversampling of the center of k-space for full HSS.

As noise is added, the SSIM for all reconstructions decreases, full HSS, multiple-shot and single-shot, which all have equal readout time, decrease at roughly the same rate. As the injected noise increases, the single shot spiral performs better than either the full HSS or multiple-shot spiral. Single-shot HSS, with a shorter readout time than the other reconstructions, decreases less quickly than the other methods so that it has the highest SSIM once the noise FWHM reaches around 6% of the maximum value of the FID.

The SSIM of all reconstructions decreases with increasing off resonance. The single-shot HSS reconstruction decreases much less quickly than the other readouts, due to its shorter length; however there are differences in the other reconstructions, which each have equal readout time. Of the multiple-shot, single-shot and full HSS reconstructions, the multiple-shot image quality degrades the least quickly with added off resonance. Initially the full HSS reconstruction out performs the single-shot spiral; however at greater than 61.5 Hz, this reverses and full HSS gives the lowest SSIM of all the reconstructions.

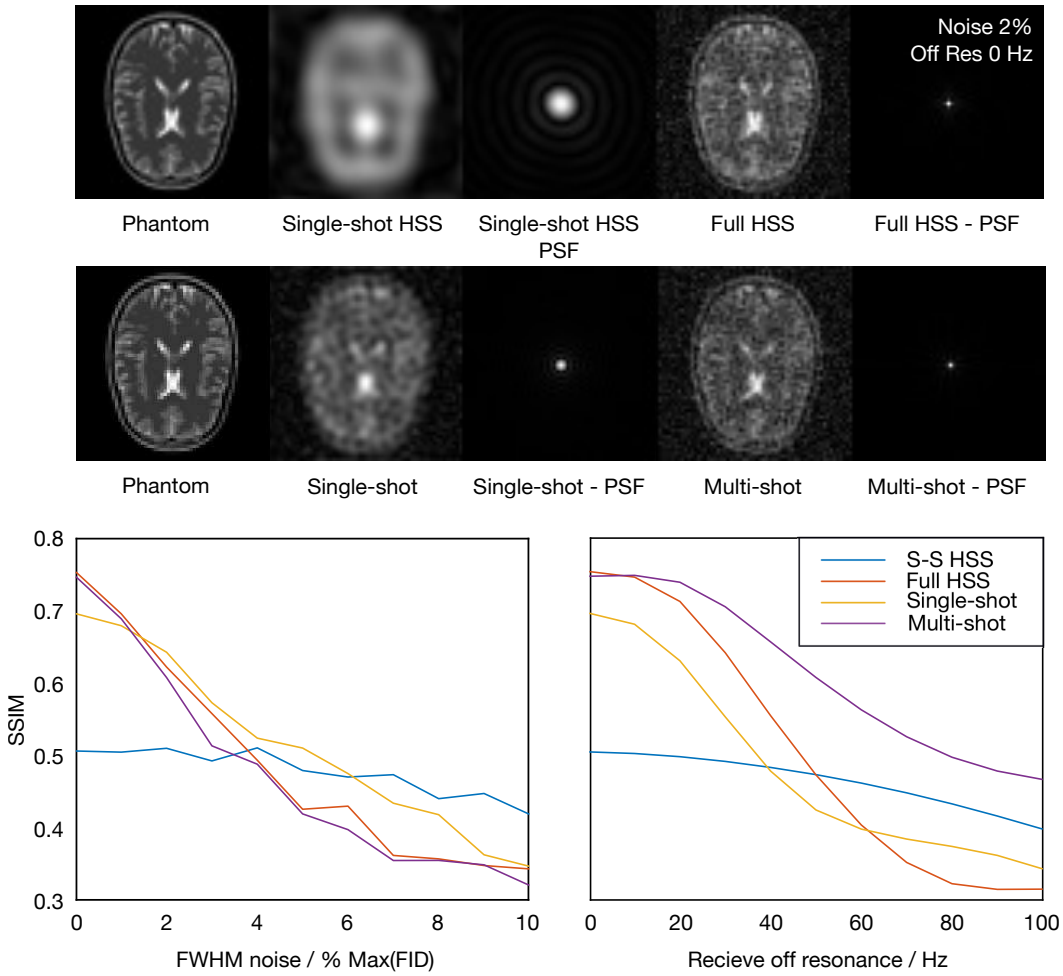

Figure 2.8: Top: Sample images showing reconstructed HSS, single-shot and multiple-shot acquisitions and a ground truth digital phantom with Gaussian noise (FWHM of noise 2% maximum value of FID) injected into the FID.
Bottom: SSIM of the reconstructed images, with the digital phantom as ground truth, as a function of injected noise and simulated receive off resonance.

2.5.2 ^{1}H phantom

2.5.2.1 Phantom calibration

Calibration curves for the T_1 and T_2^* phantom titration are shown in Figure 2.9. An elemental iron concentration of 30 µg/mL gave a T_2^* of 12.1 s, comparable to the in vivo T_2^* of pyruvate and was used for all further experiments.

2.5.2.2 Sequence validation

Reconstructed phantom images and a line plot through the ^{1}H phantom are shown in Figure 2.10. The structural similarity of the reconstructed images to a reference

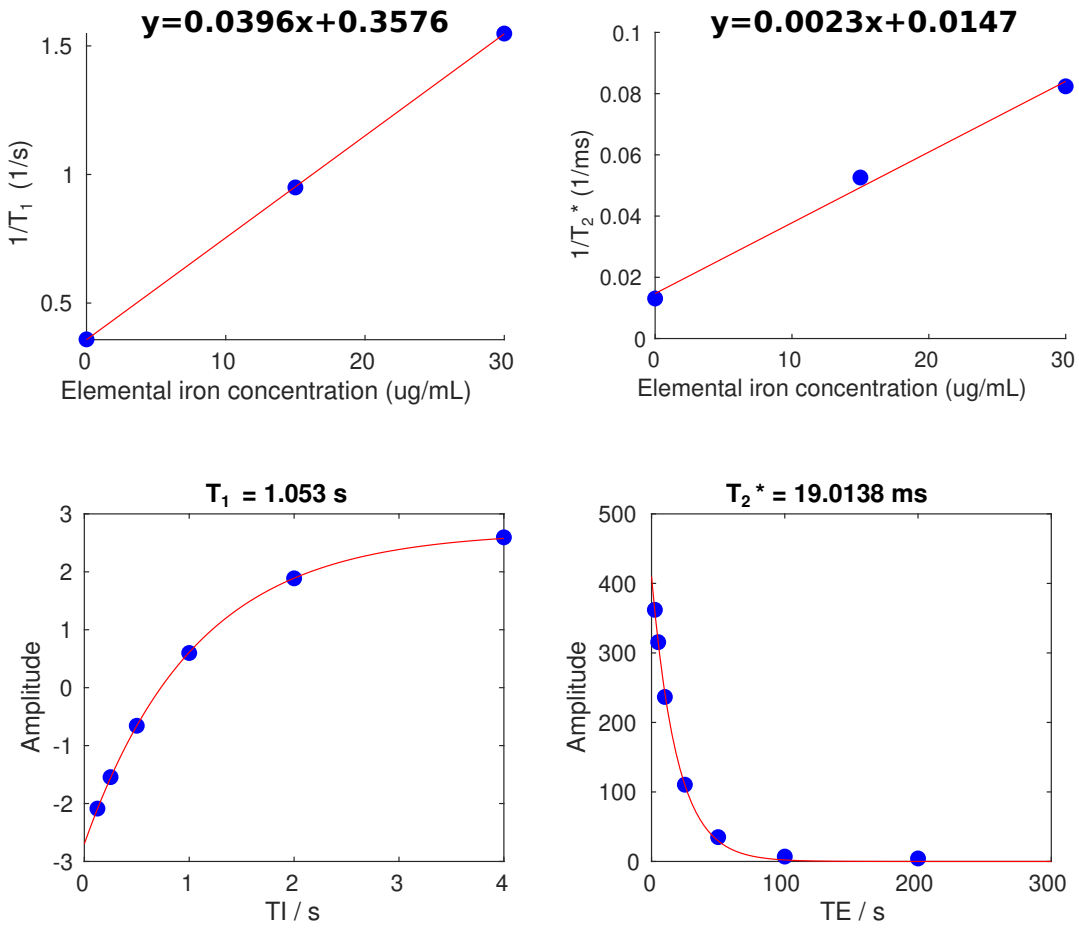

Figure 2.9: Top: Linearized titration curves for T_1 and T_2^* against Ferumoxytol elemental iron concentration. **Bottom:** Example exponential fits to inversion recovery and gradient echo sequences phantom intensities used to calculate the T_1 and T_2^* values.

gradient echo image were calculated and broadly followed the expected trend of decreasing SSIM with decreasing resolution. The single-shot HSS reconstruction had a SSIM equal to the single-shot spiral despite achieving a lower resolution potentially due to the shorter readout time reducing the measured noise.

The line plot through the ^{1}H phantom illustrates the trend quantified by the SSIM indices. The multiple-shot and full HSS reconstructions show more definition than the single-shot or single-shot HSS reconstructions which are comparatively blurred.

Significantly, the HSS acquisition achieved similar image quality to the multiple-shot spiral, when a full reconstruction was performed with the same number of shots, and when only a single-shot was used, it produced images of comparable quality to the single-shot spiral while matching its temporal resolution.

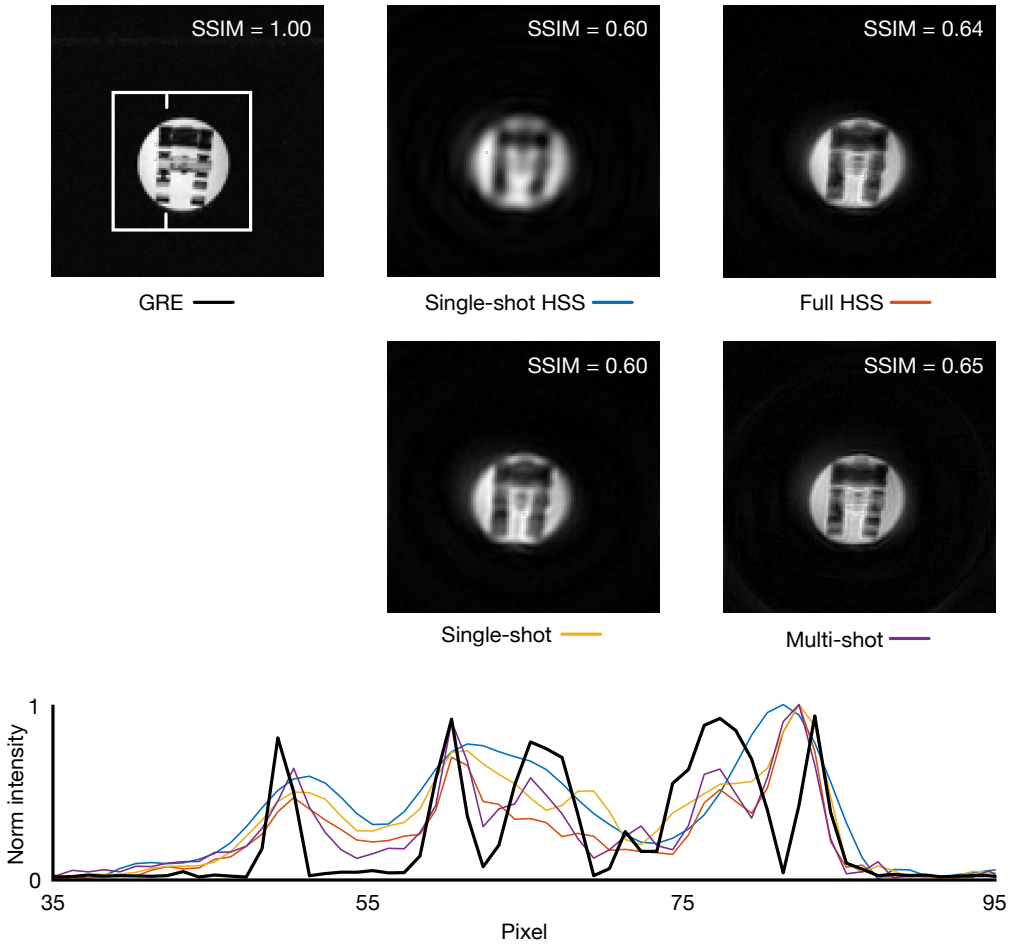

Figure 2.10: Top: Reconstructed images of a resolution phantom (T_2^* = 12 ms) with structural similarity index, over the area in the white box, with the gradient echo (GRE) acquisition shown as reference, in the top right of each image.
Bottom: Normalized line plot through each of the reconstructed images. The position of the line is denoted by the white notches in the gradient echo image.

2.5.2.3 Reconstruction method

The central slice of the full HSS dataset reconstructed with a traditional gridding NUFFT implementation[1][100] and the iterative BART NUFFT implementation, as well as with and without GIRF correction, is shown in Figure 2.11. Using the GIRF correction improved image quality, reducing aliasing of signal around the phantom and improving the SSIM relative to the uncorrected image produced with the same reconstruction. Using the BART NUFFT implementation also improved image quality, resulting in a higher SSIM, so that overall the BART + GIRF

[1]The NUFFT toolbox[100] used for this reconstruction is capable of performing an iterative reconstruction as well.

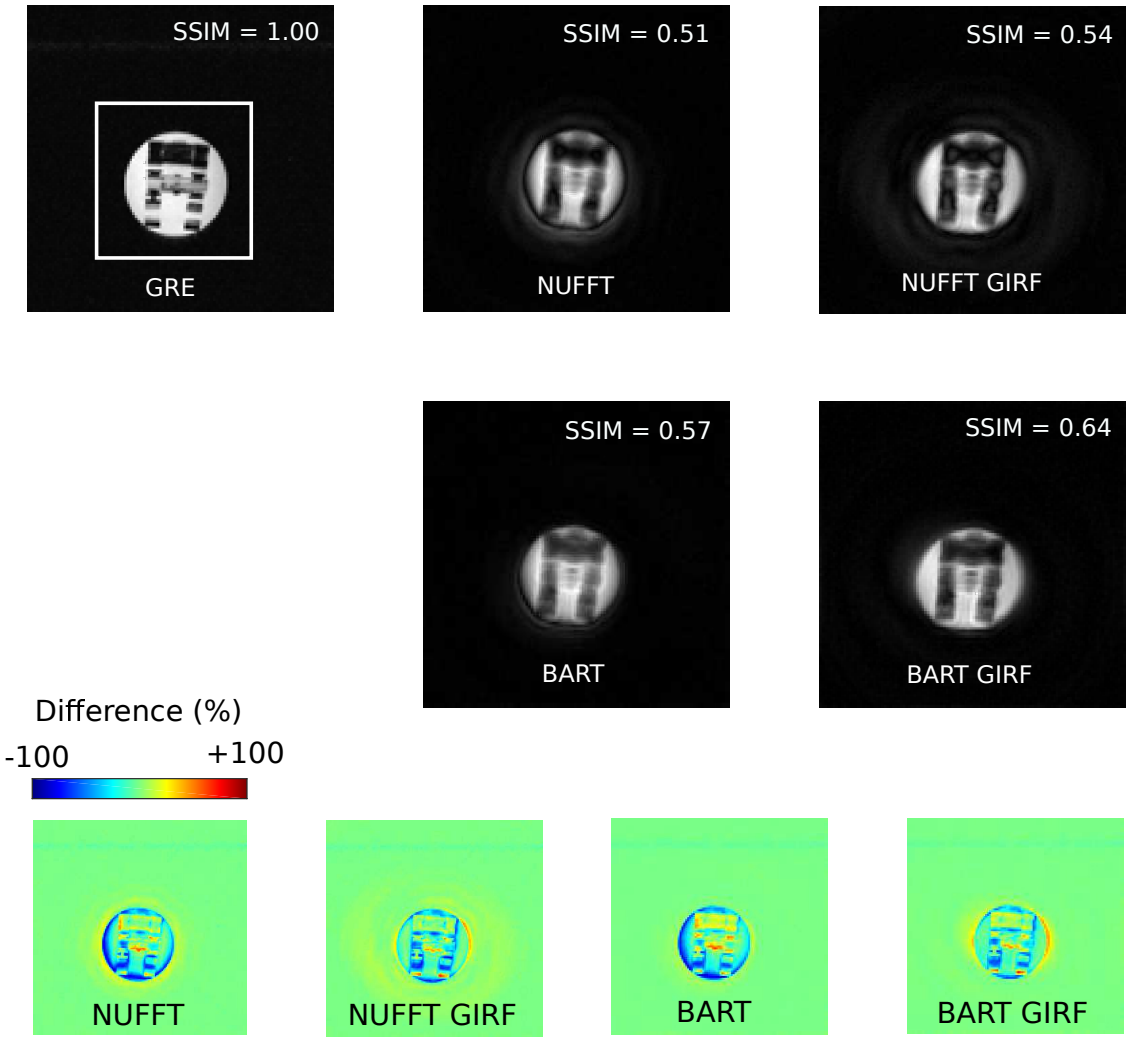

Figure 2.11: Top: Reconstructed images of a resolution phantom ($T_2^* = 12$ ms) for each reconstruction method with structural similarity index, over the area in the white box, compared to the gradient echo (GRE) acquisition shown as reference. **Bottom:** Difference maps for each reconstruction method compared to the reference GRE image.

reconstruction performed the best as measured by SSIM.

2.5.3 Preclinical imaging

The results of the pre-clinical hyperpolarized experiments are illustrated in Figures 2.12 and 2.13. The line plot in Figure 2.12 shows the resolution trade-off made by the HSS sequence. The single-shot HSS reconstruction has a visibly lower spatial resolution than the single-shot spiral, however the full HSS reconstruction, separated the left and right ventricles, demonstrating its higher spatial resolution. Figure 2.12 also shows a time-course of pyruvate SNR for the same hyperpolarized

datasets. In this example, the single-shot HSS reconstruction captured dynamic information about the system, particularly the bolus arrival and second pass of the pyruvate, while the full HSS reconstruction lacked the temporal resolution to do so. Figure 2.13 shows pyruvate, bicarbonate and lactate metabolic images acquired with the HSS sequence for three rats, two imaged with the same imaging parameters and a third with 4× lower bandwidth, 0.5× FOV and 2× resolution. In all three sets of images the signal is localized to anatomically plausible regions, with the full HSS reconstruction revealing anatomical detail, which was not resolved by the single-shot HSS reconstruction.

2.6 Discussion

In this work we demonstrated the feasibility of the HSS pulse sequence for use in hyperpolarized [1-^{13}C]pyruvate cardiac imaging. We also explored the trade-offs which inform the choice of spiral readout when designing hyperpolarized studies. Due to the inherent difficulty in measuring image quality for hyperpolarized reconstructions, as there is no ground truth, we performed validation work in both digital and ^{1}H phantoms in addition to demonstrating feasibility in a pre-clinical rat model.

The main features of the sequence, and the trade-offs it makes, are illustrated clearly by the digital phantom simulation data presented here. In the ideal case, the full HSS reconstruction achieved similar image quality to a multiple-shot readout with the same readout duration, while retaining the ability to reconstruct a low resolution image from a single shot. The main trade-off of the sequence, the single-shot HSS resolution, is clearly demonstrated by its lower SSIM than the single-shot spiral simulation.

Furthermore, the simulated full HSS sequence had broadly similar robustness to noise compared to both single and multiple-shot readouts. However as injected noise increased the single shot spiral outperformed the full HSS and multiple shot readouts, this is potentially because the full HSS and multiple shot readouts move

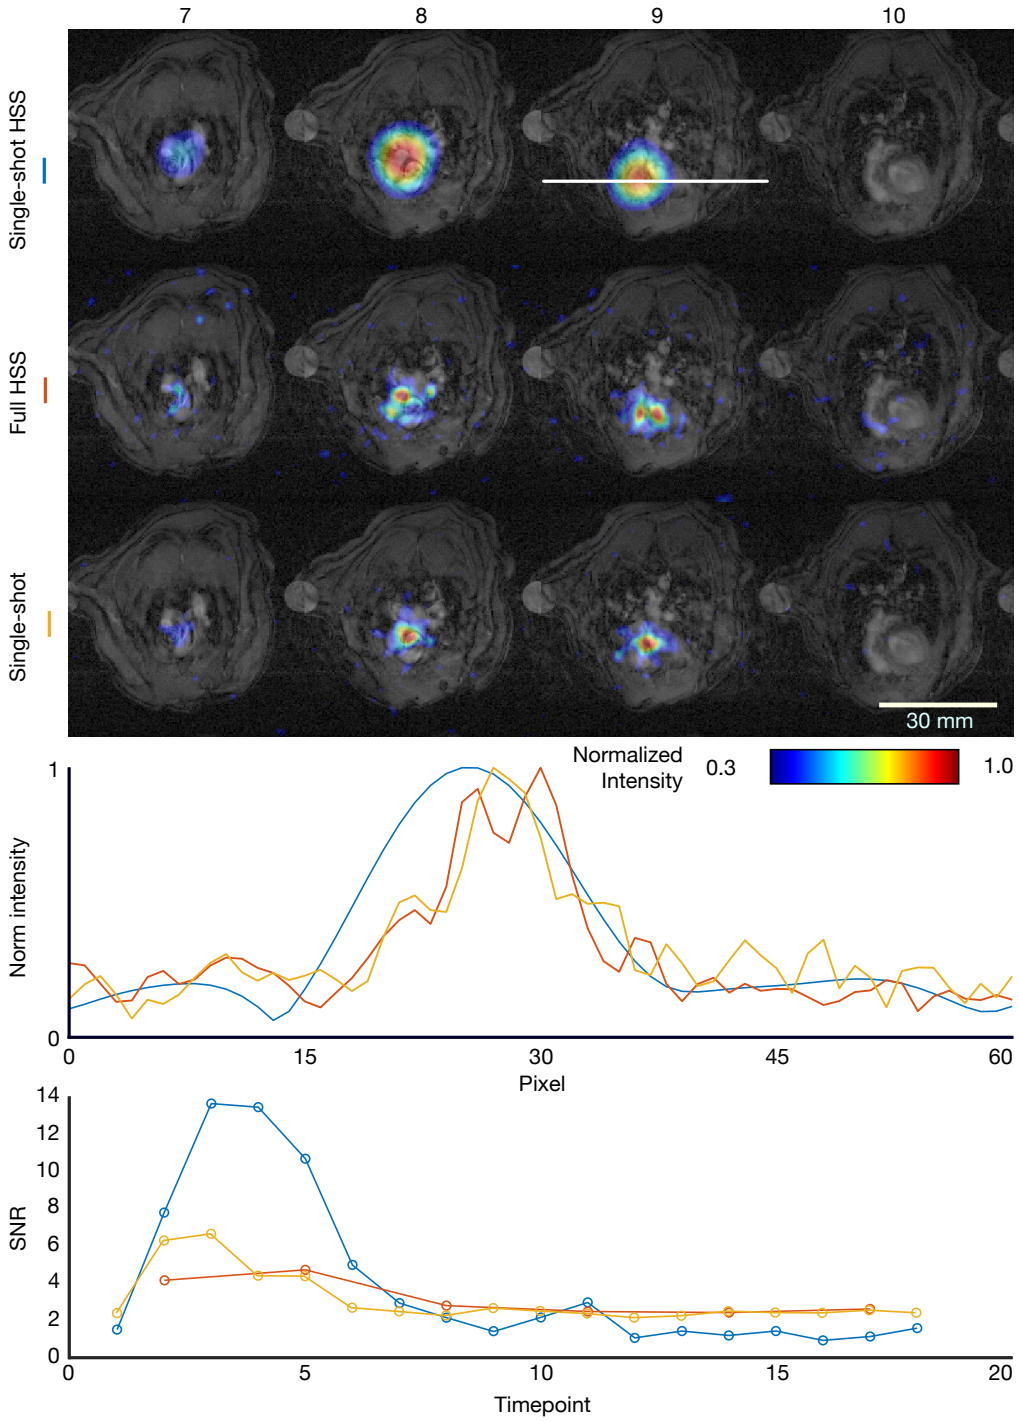

Figure 2.12: Top: Axial slices 7-10 of HSS and single-shot ^{13}C acquisitions of a rat injected with hyperpolarized [1-^{13}C]pyruvate overlaid on to matched ^{1}H anatomical images, excited at the pyruvate resonant frequency.
Middle: Line plots through slice 9 of each reconstruction in the position indicated by the white line in the single-shot HSS image.
Bottom SNR of pyruvate in the rat heart (Figure 2.12, slice 9) as a function of acquisition timepoint for single-shot and full reconstructions of a HSS acquisition in addition to a single-shot acquisition.

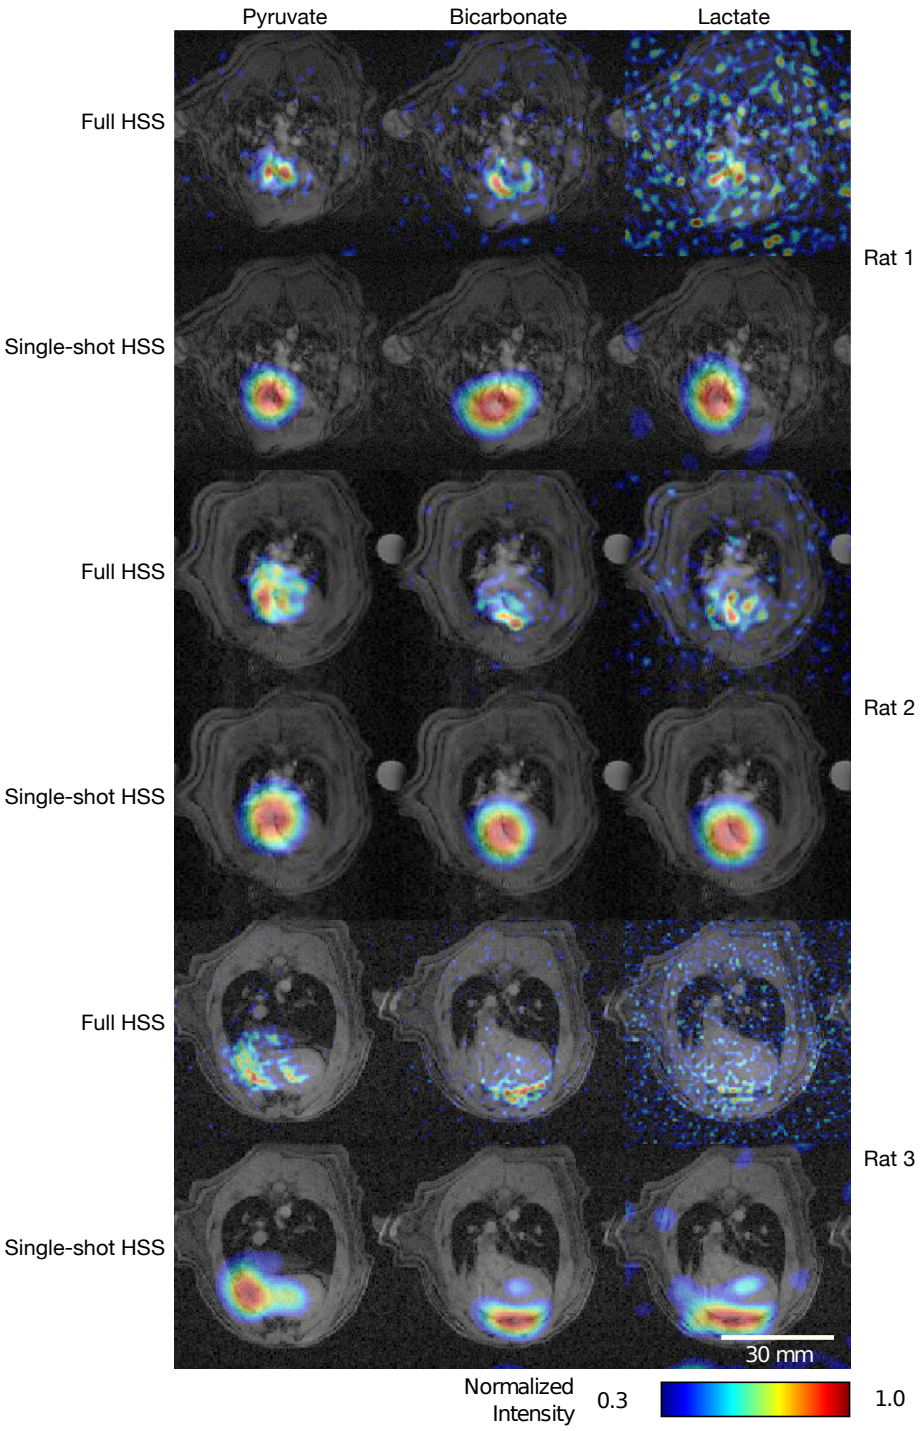

Figure 2.13: Pyruvate, bicarbonate and lactate HSS images for the rat shown in Figure 2.12 (Rat 1) in addition to two additional rats, one (rat 2) imaged with the same imaging parameters and the other (rat 3) with different parameters (bandwidth = 62.5 kHz, FOV = 60 mm, full HSS resolution = 1 mm, single shot HSS resolution = 5 mm).

away from the center of k-space more quickly and sample higher spatial frequency k-space which contributes less signal to the reconstructed image.

Of the equal readout length simulations, (i.e. not single-shot HSS) the multiple-shot simulation had the greatest resilience to off resonance effects, with the full HSS simulation initially performing better than the single-shot simulation before degrading more quickly. Since the readout times were the same, it cannot account for this difference, however it could potentially be in part due to the re-gridding process in the NUFFT and the relative amount of time spent during the readout at high spatial frequencies away from the center of k-space.

The ^{1}H phantom provided quantitative validation of the simulations in a controlled experimental system. While closely mirroring the digital phantom simulations, the similarity in SSIM metric between the single-shot spiral and single-shot HSS acquisition was surprising, however from the line plot through the phantom it is clear that the single-shot spiral did achieve a higher spatial resolution despite a poorer similarity to the reference image, potentially due to greater noise.

The ^{1}H phantom also provided validation of our decision to use the BART toolbox NUFFT implementation and GIRF correction. The improvement in image quality of the BART NUFFT may be due to the algorithm iteratively solving the inverse problem, eliminating the need to compensate for k-space sampling density, which is not well defined for the abrupt jump in sampling density of the HSS sequence at the transition from single to multi-shot.

The pre-clinical results reported here are qualitatively consistent with the digital and ^{1}H phantom validation data, and further illustrate the characteristics of the sequence. The improved resolution of a full HSS acquisition over a single-shot spiral with equal prescribed resolution is demonstrated by the line plot through the heart shown in Figure 2.12, where the pyruvate signal is clearly localized to the two ventricles, unlike for the single-shot spiral. In addition the time course data shown in Figure 2.12 demonstrates that the single-shot HSS reconstruction is able to capture appreciably more of the dynamics of the system, necessary for kinetic modeling,

than the full HSS reconstruction, albeit with a lower spatial resolution than a single-shot spiral with the same temporal resolution. It is also worth noting that for a single-shot HSS reconstruction, or single-shot spiral, each timepoint is reconstructed from consecutive readouts (i.e. for 12 z phase-encodes and a heart rate of 400 bpm, over 1.8 s), unlike for a multi-shot spiral, or full HSS, where each timepoint is a moving average of N shots which may be temporally separated by readouts for other metabolites, making data analysis for kinetic modelling more challenging.

Figure 2.13 shows pyruvate, bicarbonate and lactate images reconstructed from the same data shown in Figure 2.12, as well as images acquired from two further rats. Image quality across the datasets is comparable to prior work imaging the rat heart at 7T[90]. Furthermore, the alternative set of imaging parameters, used for rat 3, resulted in noticeably improved image quality, with better delineation of the myocardium. However the drop-off in signal towards the rear of the heart, particularly noticeable in the bicarbonate images, due to the sensitivity profile of the ^{13}C surface receive coil, demonstrates the challenge of hyperpolarized imaging of the rat heart at 7T. These images demonstrate the reproducibility of the sequence, as well as the feasibility of combining the HSS readout with a spectral-spatial excitation, which allows spatial details, such as the bicarbonate signal from the myocardium, to be resolved while maintaining temporal resolution and spectral selectivity.

Overall, the HSS pulse sequence provides an attractive compromise between spatial and temporal resolution for hyperpolarized ^{13}C MRI, by making a spatial resolution trade-off compared to a traditional multiple-shot acquisition for the ability to reconstruct as a single or multiple-shot spiral from a single acquisition. As such, it forms an attractive candidate for use in clinical imaging, and since the conclusion of this work has been adapted for clinical systems and used to image the human heart post-myocardial infarction[44].

2.7 Conclusion

Hyperpolarized ^{13}C imaging of the human heart has been demonstrated with a single-shot spiral readout, however the maximum resolution, for a given FOV, this

approach can achieve is fundamentally limited by T_2^* decay. This limitation may negatively impact the prospects of developing diagnostic methods for conditions such as ischaemia in the heart, which have shown promise in a pre-clinical setting, that rely on resolving tissue metabolism heterogeneity. The HSS sequence aims to address this limitation of spiral imaging by making a spatial resolution trade-off compared to a traditional multiple-shot acquisition for the ability to reconstruct each shot individually. This allows for multiple-shot images to be reconstructed at a resolution greater than is achievable with a single-shot acquisition, while offering the option to reconstruct images at the temporal resolution of a single-shot spiral, with the trade-off of reduced spatial resolution compared to a standard single-shot sequence.

The experiments presented in this chapter demonstrated the capabilities of the sequence, as well as its relative merits and trade-offs, and validated the benefits it has for clinical hyperpolarized ^{13}C imaging, before it was implemented for clinical studies. In our future work we initially plan to further develop the sequence by implementing a B_0 inhomogeneity correction to reduce the impact, identified by simulations, of off resonance effects in vivo, and implement the GIRF correction, identified by this work as critical to image quality, on the clinical scanner.

3

Compartmentalized ^{31}P spectroscopy at 7T in the heart

Contents

Work contained within this chapter, including the experiments contained within Section 3.5 has been submitted to Magnetic Resonance in Medicine, where I am first author, it is currently under revision.

3.1 Introduction

Magnetic resonance spectroscopy (MRS) of high-energy phosphorus metabolites provides a unique insight into cardiac metabolism (Section 1.9). Of particular interest is the phosphocreatine to adenosine triphosphate (PCr/ATP) metabolite ratio which differs between healthy and diseased hearts[49, 50] and has been shown to be a significant predictor of cardiac mortality[52]. Numerous strategies exist to measure this, all of which involve some form of localization to exclude signal from outside of the heart. Typical localization techniques for the heart include sequences such as ISIS[32], PRESS[33], DRESS[34] or STEAM[35], or CSI spectroscopic imaging, using the Fourier transform to reconstruct phase encoded data[56].

Unfortunately all ^{31}P MRS techniques are limited by the fundamental problems of low ^{31}P concentrations in vivo and a lower gyro-magnetic ratio compared to ^{1}H, making it difficult to achieve adequate SNR. Localization can degrade the achievable SNR further. In the case of single voxel techniques this is due to the potential missmatch between the shape of the anatomical region of interest and the cuboidal voxel, in addition to signal loss during preparation as a result of relaxation processes. For CSI, however, SNR losses are associated with the shape of the spatial response function. Low SNR may reduce the reproducibiliy of ^{31}P MRS, since it negatively impacts the accuracy of fitting a model to the acquired spectra. Although increasing the field strength of the magnet from 1.5 or 3T to 7T can increase SNR and therefore repeatability[56, 103], ^{31}P spectroscopy is still an inherently SNR-limited technique.

Alternative reconstruction strategies, to the Fourier transform based reconstruction of phase encoded data, which separate out the ^{31}P signal into arbitrarily-shaped user-defined compartments that correspond to relatively homogeneous anatomical

structures (e.g. heart, liver etc.), have been under development for several years, but are not widely adopted. The first technique of this 'compartmentalized spectroscopy' family, Spectral Localization by Imaging (SLIM), was developed by Hu et al.[104] in 1988, and soon after extended to Spectral Localization with Optimal Pointspread Function (SLOOP)[105] by Von Kienlin et al. In theory, the underlying reconstructions used for both SLIM and SLOOP have the property that localization is perfect, and there is no inter-compartmental leakage[106], however this assumes perfect compartmental homogeneity, which generally does not hold true in physiological conditions or non-uniform B_1 fields. Unlike SLIM, SLOOP uses numerically optimized phase encoding gradients that relax the requirement for the compartments to be completely homogeneous, but has the disadvantage that the gradient steps must be computed while the volunteer is in the scanner. Assuming sufficient homogeneity (which must be perfect in the case of SLIM) these techniques can provide absolute quantification of metabolite concentrations, because the shape of the spatial response function (SRF) can be defined to exactly match anatomical structures, with theoretically perfect localization, as demonstrated by SLOOP[107, 108].

Recently, Zhang et al. introduced a new compartment-based technique called Spectroscopy with Linear Algebraic Modeling (SLAM)[109], and extended the method to include sensitivity-encoding with multi-receive arrays and CEST acquisitions[110]. Furthermore this work also implemented a method of optimizing the phase encoding gradients using fractional k-space steps, to maximize SNR or minimize the inter-compartmental leakage due to inhomogeneous compartments, called fSLAM[109]. SLAM and fSLAM produce compartment spectra with improved SNR and reduced scan times as compared to co-adding voxels in a regular fully encoded CSI experiment. This is achieved by reducing both the number of phase encodes (allowing for greater averaging within the same scan time) and selecting energy rich phase encodes close to the center of k-space.

For ^{31}P MRS, working at 7T, as opposed to lower fields, has the significant advantages of improved SNR and spectral separation. However it also introduces new

challenges for compartmentalized spectroscopy techniques. Naïve implementations of SLAM and SLIM rely on each segmented compartment being fairly homogeneous, which appears satisfactory for cardiac ^{31}P studies at 3T. At 7T however, the static (B_0) field and RF (B_1) field homogeneity is significantly worse than at field strengths $\leq$ 3T. While these field inhomogeneities have been addressed for SLAM at 3T[111, 112], compartmentalized techniques have not yet been implemented at 7T where conditions are more extreme.

Our center (OCMR) currently performs quantification of mid-septal PCr/ATP using an acquisition weighted CSI acquisition with conventional Fourier based reconstruction for cardiac ^{31}P MRSI at both 3 and 7T. This makes compartmentalized reconstruction techniques, which could further improve SNR over even our current 7T protocol, very attractive avenues of research, especially if they can be used in combination with our existing acquisition, as this would significantly simplify roll-out to our clinical research studies.

The work in this chapter falls into four sections, the first section describes a set of simulations designed to help understand the reconstruction techniques and how best to implement them. The second section details preliminary experiments validating our implementation of the compartmentalized spectroscopy techniques, and their limitations in a real life setting. The third, main, section describes a repeatability study, to assess the applicability of these compartmentalized spectroscopy techniques to cardiac ^{31}P spectroscopy, and the final section describes additional experiments to answer remaining questions posed by the preceding work.

3.2 Theory

Compartmentalized spectroscopy techniques formulate the problem of reconstructing spectra from C compartments as the solution to a system of equations that can be developed from the CSI signal equation.

For one (spatial) dimension, the CSI localization problem can be written as a continuous Fourier transform (FT)[109]:

$$s\left(k,t\right)=\iint\rho\left(x,f\right)e^{-i2\pi(kx+ft)}\,\mathrm{d}f\,\mathrm{d}x \tag{3.1}$$

where $s\left(k,t\right)$ is the received signal at k-space position k and time t and $\rho\left(x,f\right)$ is the spectral density at spatial position x (normalized by 1/FOV) and frequency f.

This section first details how the SLAM and SLIM reconstruction algorithms are derived from this equation, before considering how the optimal k-space sampling scheme can be calculated, and finally deriving the mathematical relationship between the SLAM, SLIM and CSI reconstruction algorithms.

3.2.1 The SLAM reconstruction algorithm

In the SLAM reconstruction, this equation is discretized to give a discrete FT and a time domain FT is applied:

$$s(k)=\sum_{m=1}^{M}\rho(x_m)e^{-i2\pi kx_m} \tag{3.2}$$

For simplicity FT$[\,s\left(k,t\right)]$ is denoted $s(k)$ and $\rho\left(x_m,f\right)$ as $\rho(x_m)$, where each entry is a row vector of length N. The equation can be expanded to:

$$\begin{bmatrix}s(k_1)\\s(k_2)\\\vdots\\s(k_M)\end{bmatrix}_{M\times N}=\begin{bmatrix}e^{-i2\pi k_1x_1}&e^{-i2\pi k_1x_2}&\dots&e^{-i2\pi k_1x_M}\\e^{-i2\pi k_2x_1}&e^{-i2\pi k_2x_2}&\dots&e^{-i2\pi k_2x_M}\\\vdots&\vdots&\ddots&\vdots\\e^{-i2\pi k_Mx_1}&e^{-i2\pi k_Mx_2}&\dots&e^{-i2\pi k_Mx_M}\end{bmatrix}_{M\times M}\times\begin{bmatrix}\rho(x_1)\\\rho(x_2)\\\vdots\\\rho(x_M)\end{bmatrix}_{M\times N}, \tag{3.3}$$

where M is the number of phase encodes/voxels in a theoretical high resolution CSI experiment and N the number of points recorded in the readout. This equation can be written in short hand as $\mathbf{S}_{M\times N}=\mathbf{F}_{M\times M}\times\boldsymbol{\rho}_{M\times N}$ where $\mathbf{S}$ is the acquired data, $\mathbf{F}$ is the discrete Fourier transform operator defined by the second matrix in Equation (3.3), and $\boldsymbol{\rho}$ is the un-encoded spectra for each spatial location x_m.

The SLAM method reconstructs spectra from $C\leq M$ user defined compartments by summing (with a normalization factor) the columns of the $\mathbf{F}$ matrix, which

correspond to x positions within each of the compartments, to give a new matrix $\mathbf{G}$, thereby reducing the dimensions of the reconstruction to[1]

$$\mathbf{S}_{M\times N} = \mathbf{G}_{M\times C} \times \boldsymbol{\rho}_{C\times N}. \tag{3.4}$$

Now that $M \geq C$, and therefore $\mathbf{G}_{M\times C}$ is over-determined, a reduced gradient set of $P \geq C$ phase encodes may be used. To maximize SNR, the phase encodes are chosen from central k-space and equation 3.4 reduces to

$$\mathbf{S}_{P\times N} = \mathbf{G}_{P\times C} \times \boldsymbol{\rho}_{C\times N}. \tag{3.5}$$

This equation is solved by pre-multiplying both sides by the Moore-Penrose pseudo-inverse of $\mathbf{G}_{P\times C}$, to give the least squares solution. Since the number of phase encodes acquired (P) can be set to any value greater than or equal to C, acceleration factors of up to M/C compared to a conventional M gradient step CSI experiment can be achieved.

As a further refinement to the SLAM technique, enabling the use of acquisition weighted data-sets, we introduce R_p, the number of repeats of the p^{th} phase encode. Each row in the $\mathbf{G}_{P\times C}$ matrix is weighted by R_p to account for the effect on the signal magnitude of adding up the varying number of repeats for the different phase encodes. For conventional SLAM implementations $R_p = 1$.

3.2.2 The SLIM reconstruction algorithm

The SLIM (as well as SLOOP) reconstruction algorithms separate the CSI equation (Equation (3.1)) into C integrals, corresponding to C compartments, giving:

$$\mathrm{s}\left(k,t\right) = \int_{\mathrm{comp}\,1} \xi\left(x,t\right) e^{-i2\pi kx}\,\mathrm{d}x + \ldots + \int_{\mathrm{comp}\,C} \xi\left(x,t\right) e^{-i2\pi kx}\,\mathrm{d}x \tag{3.6}$$

where $\xi\left(x,t\right)$ is defined as the time domain FT of $\rho\left(x,f\right)$

$$\xi\left(x,t\right) = \int_{-\infty}^{\infty} \rho\left(x,f\right) e^{-i2\pi ft}\,\mathrm{d}f \tag{3.7}$$

[1]The alternative notation of **PE** for **G** is often used.

If $\xi(x,t)$ depends only on t within each compartment (i.e. the compartment is homogeneous) we can substitute in the compartment signal $\xi_c(t)$ for $\xi(x,t)$, giving

$$\mathrm{s}(k_p,t) = \sum_{c=1}^{C} \mathrm{g}_{pc}\,\xi_c(t), \quad p = 1,\ldots,P \tag{3.8}$$

where P is the number of phase encode steps recorded and g_{pc} is defined in 1D as[104]

$$\mathrm{g}_{pc} = \mathrm{R}_p \int_{\mathrm{comp}\,c} e^{-i2\pi k_p x}\,\mathrm{d}x \tag{3.9}$$

Here R_p is introduced as the number of repeats of the p^{th} phase encode step, in order to account for acquisition weighting, k_p is the p^{th} phase encoding vector; and the integral extends over the c^{th} compartment ($\mathrm{comp}\,c$). Since Equation (3.6) is at no point discretized, the continuous FT (CFT) forms the basis of the SLIM equation, with the step between equations 3.6 and 3.8 introducing a homogeneity condition for each of the compartment spectra.

This gives the matrix equation:

$$\mathbf{S}_P(t) = \mathbf{G}\boldsymbol{\xi}_C(t) \tag{3.10}$$

When solved by pre-multiplying both sides by $\mathbf{H}$, where $\mathbf{HG} = \mathbf{I}$, this gives the separated signal for each compartment which can be Fourier transformed to give each compartment's spectra (ρ_c). Since this equation is over-determined for $C < P$ the number of phase encodes (P) may be reduced to $P \geq C$ as in SLAM.

3.2.3 Calculating optimal k-space sampling for the SLAM and SLIM reconstruction algorithms

The fSLAM and SLOOP methods improves the localization properties of SLAM and SLIM respectively by numerically optimizing the k-space coordinates of the phase encodes[105, 109] based on the compartmentalization mask. Because a localizer must be acquired to segment the compartments, this optimization must be done scanner-side at the start of the exam, increasing time in the scanner.

An essential metric for optimization is the spatial response function (SRF), which gives the contribution of the metabolites at each point in the FOV to each compartment's spectra, and allows the effect on the reconstruction of changing the acquired k-space coordinates to be quantified. SRFs can be calculated for both the SLAM and SLIM reconstructions, by performing the reconstruction for each point in space, with the following formula:

$$\beta_c(x) = \sum_{p=1}^{P} h_{cp} R_p e^{-i2\pi k_p x} \tag{3.11}$$

Here $\beta_c(x)$ is the c^{th} compartment's SRF, h_{cp} is the cp^{th} element of $\mathbf{H}$, and $\mathbf{H}$ is the Moore-Penrose pseudo-inverse of $\mathbf{G}$ for SLAM or SLIM. R_p is the number of repeats of the p^{th} phase encode, to allow for the effect of acquisition weighting to be calculated, for conventional SLAM and SLIM implementations $R_p = 1$.

A number of different cost functions are proposed for fSLAM[109], in this work we implement the cost function that minimizes both the inter- and intra-compartmental errors, which, for the i^{th} compartment, is defined as

$$\Lambda_i = (\phi_i + \psi_i)/I_{\text{cond}} \tag{3.12}$$

The three terms on the right-hand side are calculated as follows. ϕ_i, the inter-compartmental error for compartment i, is:

$$\phi_i = \sum_{j \neq i}^{C} \sum_{m \in \text{comp } j}^{M} |\mathbf{G}^l_{C\times M}(i,m) - \mathbf{M}_C|^2, \qquad \mathbf{G}^l_{C\times M}(i,m) = \mathbf{G}^+_{C\times P} \times \mathbf{G}_{P\times M} \tag{3.13}$$

where $\mathbf{M}_C$ is the row-wise mean of the $\mathbf{G}^l_{C\times M}(i,m)$ matrix and $\mathbf{G}_{P\times M}$ is the subset of the $\mathbf{G}_{M\times M}$ matrix retaining only the rows corresponding to the phase encodes used. This equation is the sum of the squares of the SRF of compartment i that fall outside the area segmented by compartment i. ψ_i is the error arising from inhomogeneity within the compartment and is calculated as the sum of the squares of the SRF for compartment i over the space within the compartment. Finally

I_{cond} is a boolean value set to zero (and thereby making $\Lambda_i = \infty$) if the condition number of $\mathbf{G}$ exceeds a pre-set user-defined threshold[109].

The SLOOP optimization proceeds similarly[105], using a weighted sum of two terms for the cost function, these are:

$$\Omega_i = \int_{\text{outside comp } i} |\mathbf{SRF}_i(x)| \, \mathrm{d}x \,, \tag{3.14}$$

where Ω_i quantifies the contamination of compartment i from outside of compartment i, and

$$N_c = \sqrt{\sum_{p=1}^{P} (R_p |h_{cp}|)^2} \tag{3.15}$$

with N_c the noise intensity in the reconstructed spectrum of compartment c. Optimizing N_c necessarily optimizes the signal efficiency for compartment c which is defined as

$$E_c = \frac{1}{\int_{\text{comp } c} \mathrm{d}x \, \left(\sqrt{\sum_{p=1}^{P} R_p}\right) N_c} \tag{3.16}$$

Unlike fSLAM the optimization is not performed for a single compartment, instead the cost function calculated as

$$f(k) = A \max(\Omega) + B \max(N), \tag{3.17}$$

where the maximum Ω and N values may come from different compartments and A and B are empirically determined coefficients. The SLOOP algorithm is therefore optimizing leakage for the worst compartment by this metric and noise by the worst compartment for that metric, without the user selecting a particular compartment to optimize against, a notable difference to the fSLAM algorithm.

3.2.4 The mathematical relationship between SLAM, SLIM and CSI

3.2.4.1 CSI

The signal equation for a conventional CSI experiment is given by Equation (3.3) and can be summarized as

$$\mathbf{S}_{M\times N} = \mathbf{F}_{M\times M} \times \boldsymbol{\rho}_{M\times N}, \tag{3.18}$$

where $\mathbf{F}$ is the discrete Fourier transform operator. When k-space is fully sampled $\mathbf{F}^H = \mathbf{F}^{-1}$ and $\hat{\boldsymbol{\rho}}$, the discrete approximation of $\boldsymbol{\rho}$, can simply be found by

$$\hat{\boldsymbol{\rho}}_{M\times N} = \mathbf{F}^H_{M\times M} \times \mathbf{S}_{M\times N} \tag{3.19}$$

If less than M k-space samples are acquired a new Fourier operator $\tilde{\mathbf{F}}$ with dimensions $P \times M$ is used. Although this operator is not unitary i.e. $\tilde{\mathbf{F}}^H \neq \tilde{\mathbf{F}}^{-1}$, the equation

$$\tilde{\boldsymbol{\rho}}_{M\times N} = \tilde{\mathbf{F}}^H_{M\times P} \times \mathbf{S}_{P\times N} \tag{3.20}$$

is the zero-filled reconstruction of the data to a full (M) resolution grid.

3.2.4.2 SLAM

For the SLAM reconstruction algorithm the $\mathbf{G}_{P\times C}$ matrix can be represented as

$$\tilde{\mathbf{G}}_{P\times C} = \tilde{\mathbf{F}}_{P\times M} \times \mathbf{B}_{M\times C}, \tag{3.21}$$

where $B_{M\times C}$ is a matrix operator, which sums columns, and whose rows represent the x coordinates $x_1 ... x_M$ and columns compartments, if x_i is within compartment j then $B_{ij} = 1$ otherwise $B_{ij} = 0$.

Using this relation, and the expanded form of the Moore-Penrose pseudo-inverse, the solution to the SLAM equation can be expanded to:

$$\begin{aligned}\tilde{\boldsymbol{\rho}}_C &= \text{pinv}(\tilde{\mathbf{G}}) \times \mathbf{S} \\ &= (\tilde{\mathbf{G}}^H\tilde{\mathbf{G}})^{-1}\tilde{\mathbf{G}}^H\mathbf{S} \\ &= (\mathbf{B}^H\tilde{\mathbf{F}}^H\tilde{\mathbf{F}}\mathbf{B})^{-1}\mathbf{B}^H\tilde{\mathbf{F}}^H\mathbf{S} \\ &= (\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1}\mathbf{B}^H\tilde{\boldsymbol{\rho}}_M\end{aligned} \tag{3.22}$$

Where $\boldsymbol{\Psi}$ is the PSF operator, $\tilde{\boldsymbol{\rho}}_C$ is the estimated compartment signals and $\tilde{\boldsymbol{\rho}}_M$ is the estimated signal at each position $1...M$.

From the final line of Equation (3.22) we can see the linear relationship between the SLAM and CSI algorithms. To generate the SLAM compartment signals the $\mathbf{B}^H$ operator sums the CSI voxels within each compartment, then the $(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1}$ operator takes linear combinations of these summed signals to compute the final compartment signals. In the case where $P = M$, the matrix $\boldsymbol{\Psi} = I$ (as $\mathbf{F}^H = \mathbf{F}^{-1}$) and the $(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1}$ term simply normalizes each compartment signal by the compartment area. However, where $P < M$, $\boldsymbol{\Psi} \neq I$ and the $(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1}$ term is filtered by the PSF which compensates for signal contamination caused by the size of the PSF, in addition to normalizing by compartment area.

An example scenario consisting of three compartments (see Figure 3.2 for configuration with FOV = 16 cm was simulated to illustrate the effect of P on $(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1}$. Where $P = M = 16$

$$(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{14} \end{pmatrix}, \tag{3.23}$$

and where $P = 8$

$$(\mathbf{B}^H\boldsymbol{\Psi}\mathbf{B})^{-1} = \begin{pmatrix} 3.41 & -2.14 + 0.44i & -0.02 - 0.03i \\ -2.14 - 0.44i & 3.41 & -0.02 + 0.03i \\ -0.02 + 0.03i & -0.02 - 0.03i & 0.07 \end{pmatrix}, \tag{3.24}$$

the difference resulting from the markedly different PSFs of the two situations, which are shown in Figure 3.1. The difference in the off-diagonals of these matrices shows the extent to which the SLAM reconstruction differs from simply summing up CSI voxels in the case $P < M$.

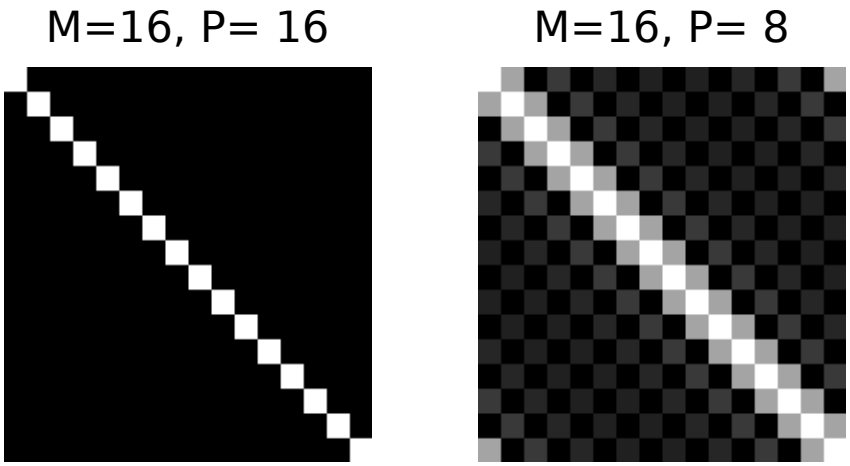

Figure 3.1: Graphical representations of the PSF operator Ψ for two situations where $M = P = 16$ and $M = 16, P = 8$.

3.2.4.3 SLIM

The SLIM algorithm has previously undergone a thorough theoretical analysis by Liang et al.[106], which broke down the SLIM equation to

$$\boldsymbol{\xi}_C(t) = \text{pinv}(\mathbf{G})\mathbf{S}(t) = (\mathbf{G}^H\mathbf{G})^{-1}\mathbf{G}^H\mathbf{S}(t) \tag{3.25}$$

and showed, through simulations, that the off-diagonal elements of $(\mathbf{G}^H\mathbf{G})^{-1}$ compensate for the truncation error present in the Fourier series which make up each element of $\mathbf{G}^H\mathbf{S}(t)$.

By considering the relationship between the SLAM and SLIM equations we show here that firstly, there is a linear relationship between SLIM and SLAM, and secondly that the PSF operator generates the off diagonal elements in $(\mathbf{G}^H\mathbf{G})^{-1}$ in the same way as SLAM.

We start with the definition of $\tilde{\mathbf{G}}$ for the SLAM reconstruction, noting that each element is equal to

$$g_{pc} = \sum_{m \in c} e^{-i2\pi k_p x_m}, \tag{3.26}$$

and if the $\tilde{\mathbf{G}}$ matrix is normalized by $\Delta x = \frac{\text{FOV}}{M}$, that as $M \to \infty$, it approaches

$$g_{pc} = \int_{m \in c} e^{-i2\pi k_p x_m}. \tag{3.27}$$

Which is equivalent to the definition of g_{pc} used in SLIM. This then leads to the conclusion that

$$\tilde{\boldsymbol{\rho}}_C^{\mathrm{SLIM}} = \lim_{M\to\infty} \frac{M}{\mathrm{FOV}} \tilde{\boldsymbol{\rho}}_C^{\mathrm{SLAM}}, \tag{3.28}$$

and therefore

$$\tilde{\boldsymbol{\rho}}_C^{\mathrm{SLIM}} = \lim_{M\to\infty} \left[\frac{M}{\mathrm{FOV}} (\mathbf{B}^H \boldsymbol{\Psi} \mathbf{B})^{-1} \mathbf{B}^H \right] \tilde{\boldsymbol{\rho}}_M, \tag{3.29}$$

the time-domain Fourier transform of which gives

$$\tilde{\boldsymbol{\xi}}_C^{\mathrm{SLIM}}(t) = FT \left\{ \lim_{M\to\infty} \left[\frac{M}{\mathrm{FOV}} (\mathbf{B}^H \boldsymbol{\Psi} \mathbf{B})^{-1} \mathbf{B}^H \right] \tilde{\boldsymbol{\rho}}_M \right\}. \tag{3.30}$$

This equation, together with Equation (3.28), demonstrates the linear relationship between the SLIM, CSI, and SLAM reconstruction algorithms, as well as the origin of the off diagonal terms in the $(\mathbf{G}^H\mathbf{G})^{-1}$ described by Liang et al.

3.3 Simulations

3.3.1 Introduction

In this section, we simulated the SLAM and SLIM methods by generating synthetic CSI data-sets, representing scenarios that challenge the reconstruction, such as touching or inhomogeneous compartments. These simulations allowed us to better understand the theoretical potential of the reconstructions, and to understand how and why the techniques break down as we move away from the idealized scenario of complete compartment homogeneity, which is a key assumption of the SLAM and SLIM reconstructions.

3.3.2 Methods

Synthetic data was simulated for three scenarios, by numerically calculating $s(k,t)$ with Equation (3.1). The data was then reconstructed with the SLAM and SLIM algorithms. The simulated CSI acquisition had BW = 4KHz, FOV = 16 cm, resolution = 1 cm. The CSI signal equation (Equation (3.1)) was evaluated with a

Scenario	Compartment	a	b	c	α	β
1	1	10000	300	30	+1.00	1.00
	2	10000	0	30	+2.00	1.00
2	1	10000	300	30	+1.00	1.00
	2a	15000	0	30	+1.75	0.50
	2b	5000	0	30	+2.25	0.50
3	1	10000	300	30	+1.25	1.00
	2	10000	0	30	+2.25	1.00

Table 3.1: Scenario parameters for Equation (3.31) when calculating $\rho(x, f)$.

trapezoidal numerical integral (MATLAB *trapz* function) with steps of 0.016 cm for the spatial dimension and 4 Hz for the spectral dimension. In all experiments $s(k, t)$ was calculated for $k = (-8 \ldots + 7) \times \text{FOV}^{-1}$ and $t = 0 \ldots 0.1$ s with steps of 250 μs. The signal regions were modeled using a rect function in the spatial domain and a Gaussian function in the spectral domain to give the following equation for $\rho(x, f)$

$$\rho(x, f) = \sum_{\text{compartments}} a_i \exp\left(\frac{(f - b_i)^2}{2c_i^2}\right) * \text{rect}\left(\frac{x - \alpha_i}{\beta_i}\right) \tag{3.31}$$

where a is the maximum amplitude, b is the central frequency of the resonance, c is the linewidth of the resonance, α is the spatial location of the resonant region and β the width of the resonant region.

The three scenarios were: 1) Two touching 1 cm wide signal regions at +1 cm and +2 cm from isocenter with frequency 300 Hz and 0 Hz respectively. 2) two touching 1 cm wide signal regions at +1 cm and +2 cm from isocenter with frequency 300 Hz and 0 Hz, where the 0 Hz signal region is in-homogeneous. 3) Two touching 1 cm wide signal regions at +1.25 cm and +2.25 cm from isocenter with frequency 300 Hz and 0 Hz respectively, with the same compartmentalized reconstruction mask as scenario 1. The parameters used to simulate each of the scenarios are outlined in Table 3.1 and graphical representations of the simulated regions, as well as the compartmentalization masks used for reconstruction, are shown in Figures 3.2 to 3.4.

3.3.3 Results and discussion

The spectra, reconstructed for each of the scenarios with the SLIM and SLAM algorithms, are shown in figures 3.2, 3.3 and 3.4.

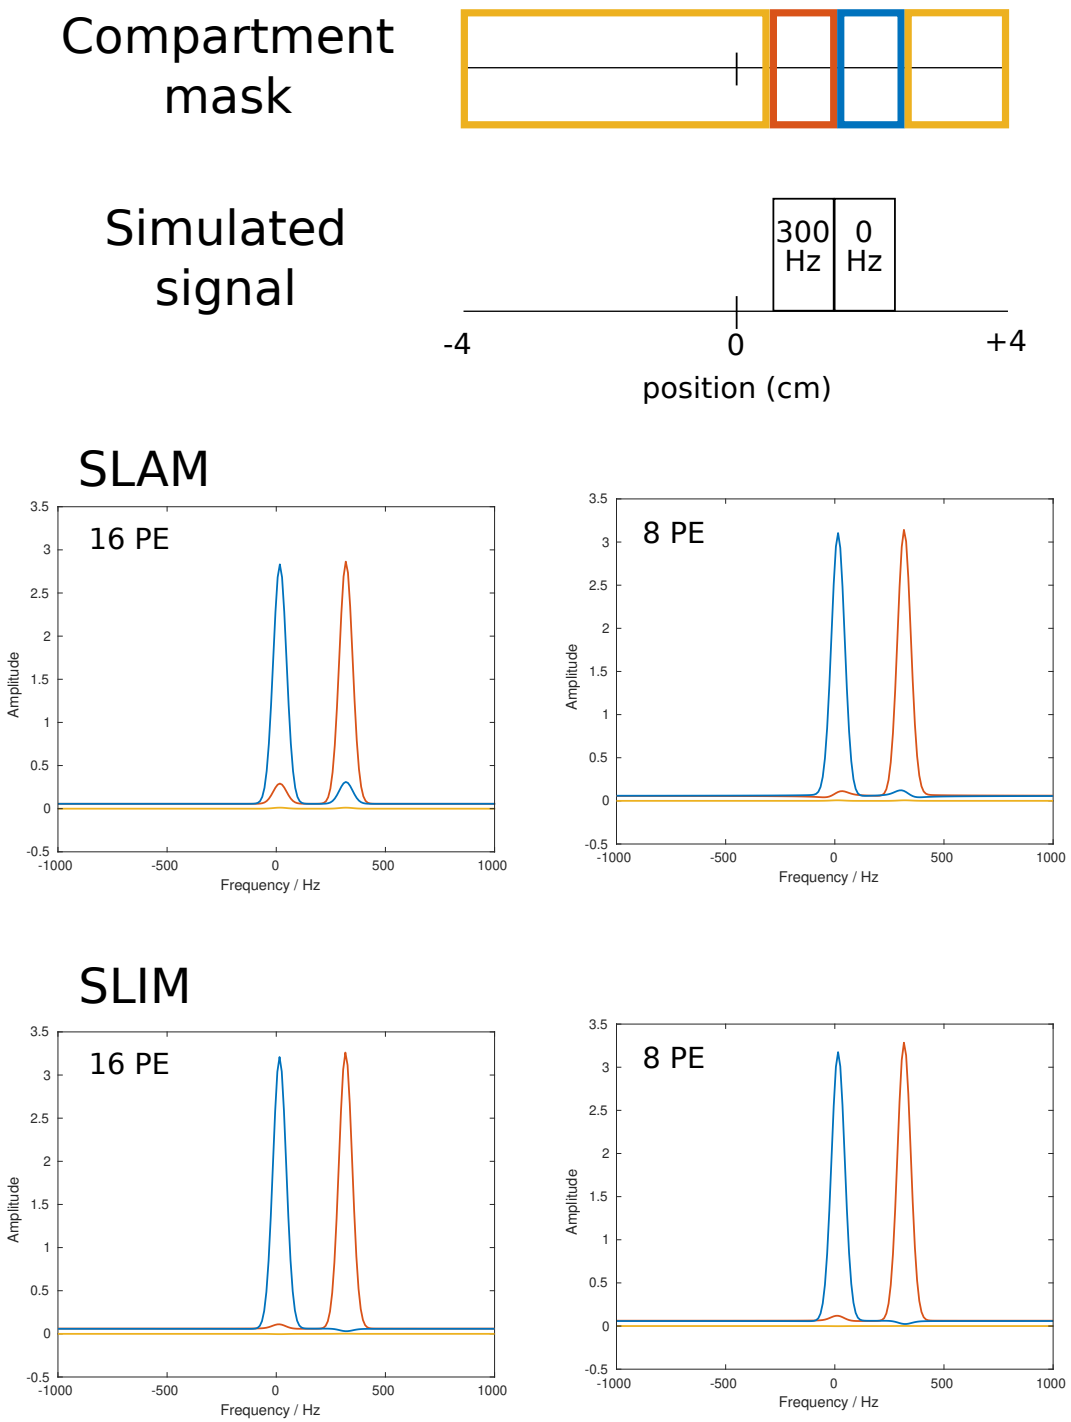

Figure 3.2: SLAM and SLIM reconstructions of 1D 16 and 8 phase encodes synthetic CSI data generated using the CSI signal equation. The synthetic data originated from two touching regions, each 1 cm wide with resonances at 300 Hz and 0 Hz respectively. The compartmentalization mask consisted of 3 compartments, one per signal region and an 'other' compartment. (Scenario 1)

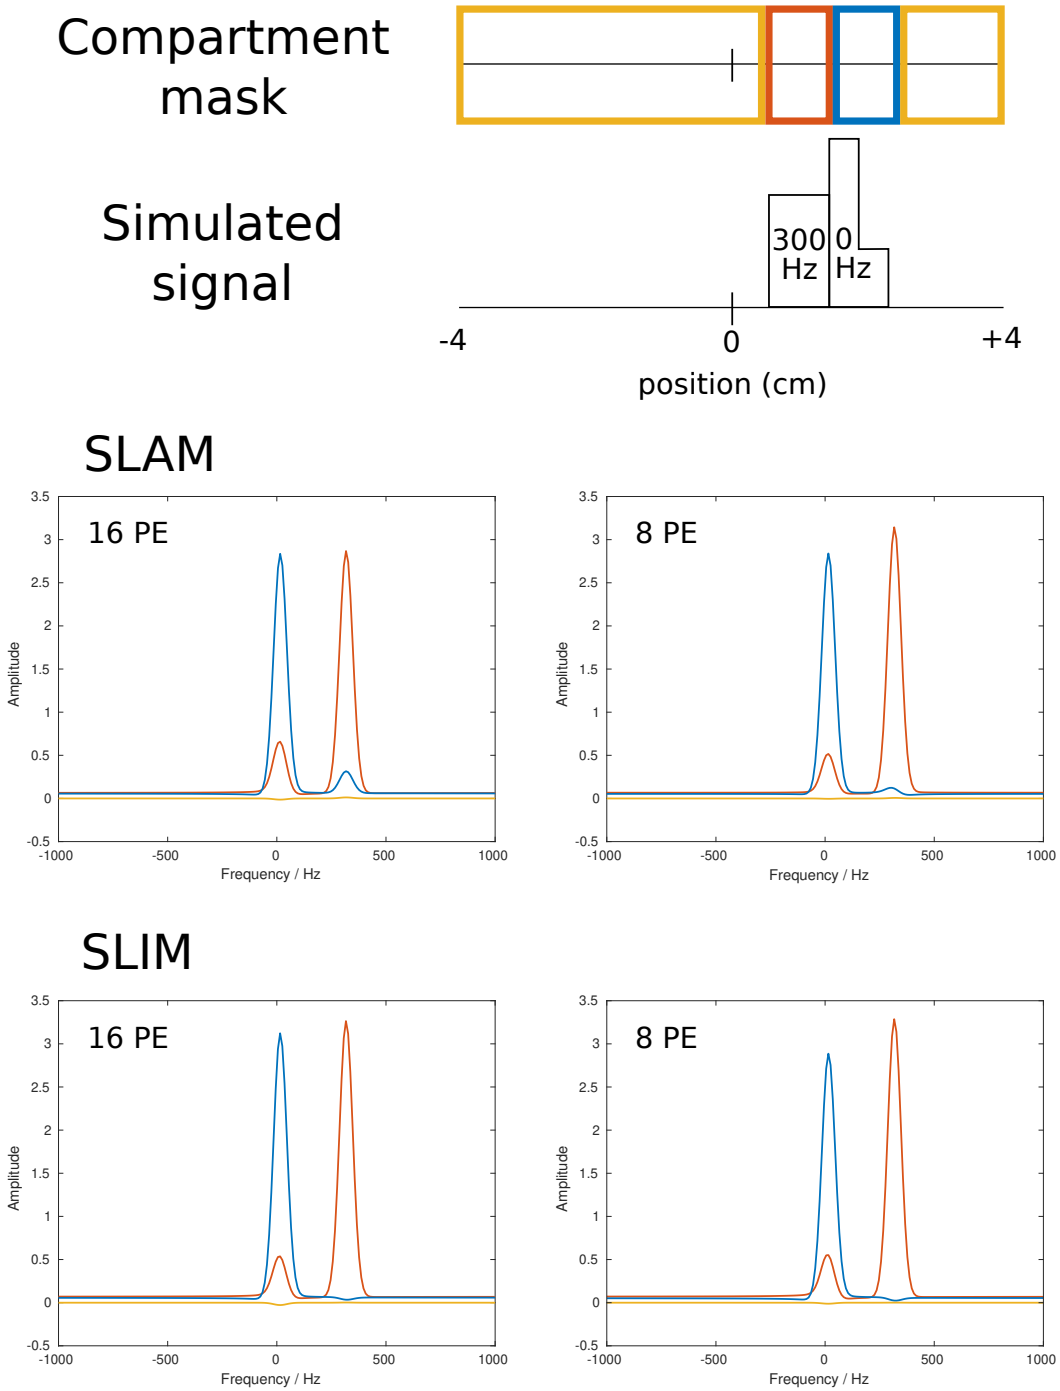

Figure 3.3: SLAM and SLIM reconstructions of 1D 16 and 8 phase encodes synthetic CSI data generated using the CSI signal equation. The synthetic data originated from two touching regions, each 1 cm wide with resonances at 300 Hz and 0 Hz respectively. The 0 Hz compartment was in-homogeneous with the first 0.5 cm at 1.5× intensity and the second 0.5× intensity, so that the integrated signal was equal to the 300 Hz compartment. The compartmentalization mask consisted of 3 compartments, one per signal region and an 'other' compartment. (Scenario 2)

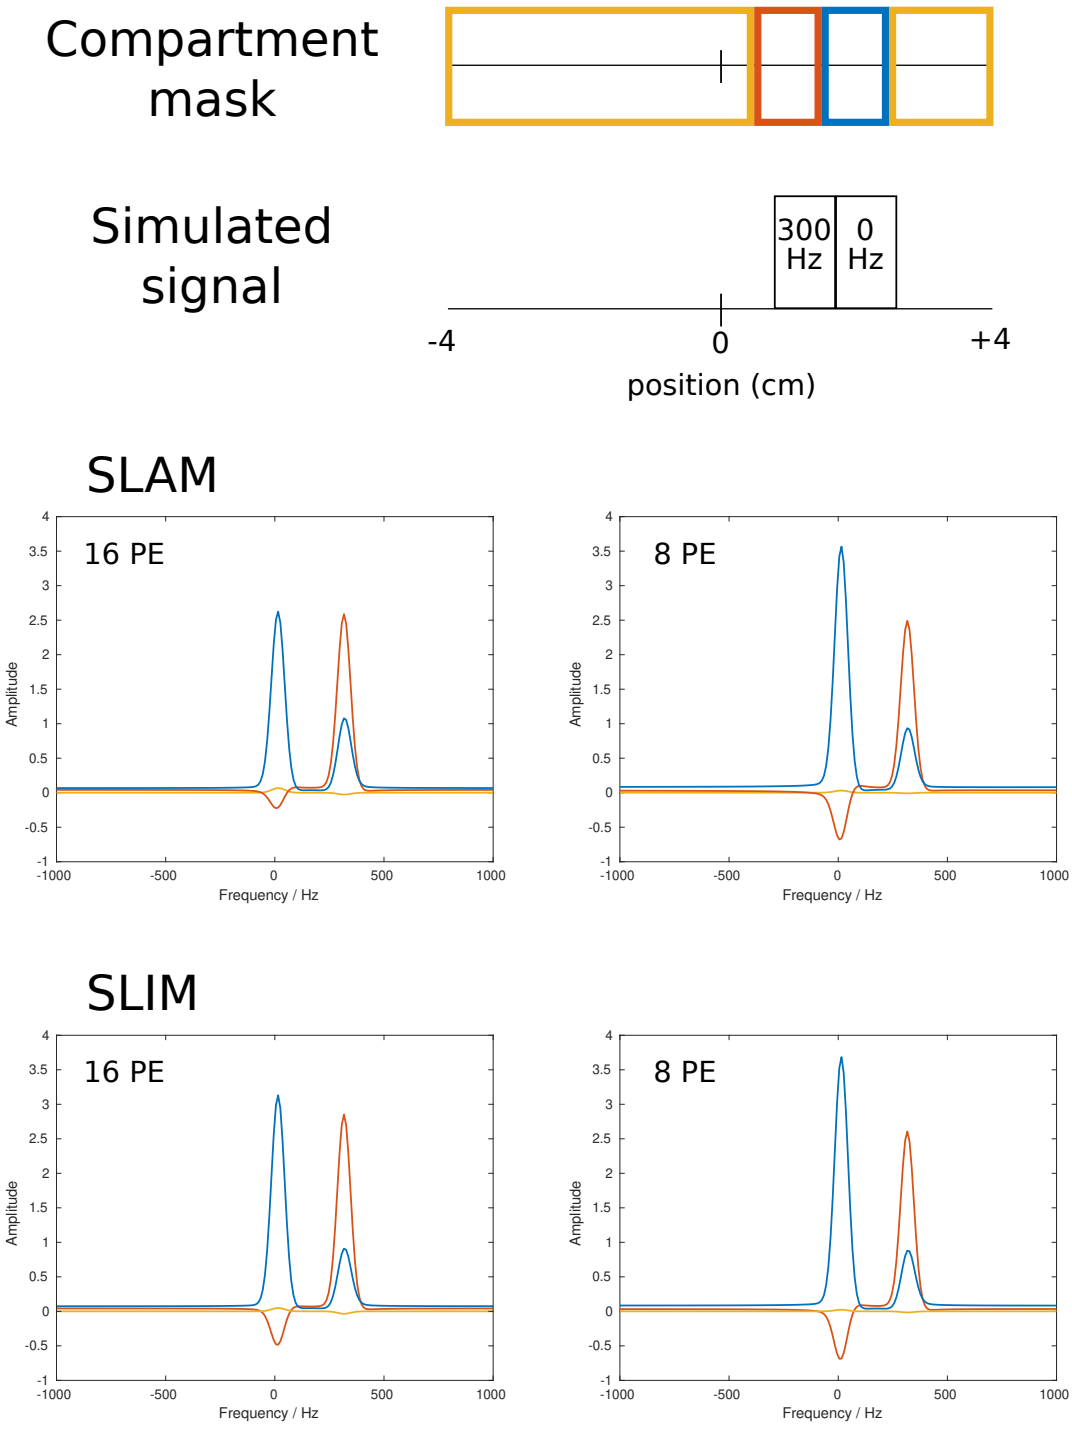

Figure 3.4: SLAM and SLIM reconstructions of 1D 16 and 8 phase encodes synthetic CSI data generated using the CSI signal equation. The compartmentalization mask consisted of 3 compartments, one per signal region and an 'other' compartment. The synthetic data originated from two touching regions, each 1 cm wide, shifted 0.25 cm from the compartmentalization mask positions, with resonances at 300 Hz and 0 Hz respectively. (Scenario 3)

Scenario 1 shows low cross-contamination of the two compartments despite their close proximity, with both 16 and 8 phase encodes. The SLIM reconstruction shows slightly lower cross-contamination, however the two are comparable. The success of the 8 phase encode reconstruction, despite the fact that a Fourier based reconstruction would have insufficient resolution (2 cm voxels) to distinguish the compartments indicates that in idealized situations the SLAM and SLIM reconstructions are able to reduce the number of phase encodes required to produce non-contaminated spectra. Interestingly the 16 phase encode SLAM reconstruction shows greater contamination than the 8 phase encode reconstruction. This is because the 16 phase encode reconstruction has a more compact SRF with smaller negative side lobes, which are used by the 8 phase encode reconstruction to cancel out contaminating signal. The 16 phase encode SLIM reconstruction however is able to cancel out the contamination, as its use of the continuous Fourier transform allows the exact position of the SRF to vary by small increments. While this appears to be a disadvantage for SLAM, in practice it is preferable to minimize reliance on cancellation of side lobes as the compartments are unlikely to be fully homogeneous as is assumed here.

In scenario 2 contamination of the 300 Hz compartment by the 0 Hz compartment is enhanced compared to scenario 1, however contamination of the 0 Hz compartment by the 300 Hz compartment is the same. This illustrates the principal that contamination is caused by in-homogeneity in the contaminating, rather than contaminated compartment. This is because while the SRF of the contaminated compartment sums to zero in the contaminating compartment, it is the product of the SRF and spin density that contributes to the spectra, and this does not sum to zero if the spin density is non-uniform.

Scenario 3 shows the significant contamination that can arise if there is a miss-match between compartmentalization mask and signal regions. The contamination of the 0 Hz spectra is attributable to the compartmentalization mask containing some of the 300 Hz signal region. Interestingly, despite the 300 Hz compartment not containing any of the 0 Hz signal region, there is negative contamination

by the 0 Hz signal, due to the in-homogeneity in the 0 Hz compartment (which now contains a mixture of signal regions), again highlighting the importance of compartmental homogeneity.

3.4 Phantom experiments

3.4.1 Introduction

Once the SLAM and SLIM reconstructions were validated with synthetic data, it was necessary to validate the reconstructions with real acquired 7T data. This was an important step for two reasons, firstly to ensure our analysis protocol was working correctly, and secondly to assess the effect of field and sample in-homogeneity in the real world.

The SLAM and SLIM algorithms require the raw ^{31}P FIDs acquired by the scanner, rather than the reconstructed spectra available in the standard DICOM output, which has undergone significant pre- and post- processing, to reconstruct compartment spectra. The raw FIDs can be acquired by exporting the data in the Siemens TWIX format, however this format does not conform to the same rules as the standard DICOM format, that is used for the localizers that generate the compartmentalization mask. It was therefore critical that the data handling was validated using a phantom with a simple, easy to verify geometry, particularly since differences between the formats, such as the axes 'handedness' or the inconsistent sign of the slice normal, are non-obvious by looking at the form of the data.

The SLIM and SLAM algorithms also assume perfect field and compartment homogeneity. At 7T the B_0 and B_1 fields are highly in-homogeneous[113], and human tissue, particularly when broadly categorized into large compartments (e.g. 'heart, 'liver') is also not homogeneous. It was therefore a significant concern that these in-homogeneities would preclude the use of these algorithms at 7T, so a good understanding of these factors was important.

To confirm the data-handling and reconstruction pipeline was appropriate, data was acquired using a 'tube' phantom, where two tubes of inorganic phosphate, with different ^{31}P resonances, could be moved into different positions. This was a good

test of how well the compartmentalization mask from the DICOM format localizer mapped on to the TWIX format spectroscopic data, because the small size of the tubes relative to the FOV would result in no signal in the tube compartments spectra if there were any misalignment.

Then, to explore the effects of in-homogeneity, we performed a second phantom experiment, where a large inorganic phosphate slab phantom was placed on top of a volunteers leg, in a configuration similar to the heart and chest wall (geometrically, and in terms of relative signal). The leg was expected to be in-homogeneous, and the inorganic phosphate phantom, highly homogeneous, so the effect of the homogeneity could be evaluated.

3.4.2 Methods

3.4.2.1 Tube phantom

The tube phantom consisted of a water tank with movable 50 ml sample tubes suspended in the water. Two of the sample tubes contained K_2HPO_4 disolved in water, one of which was doped with NaOH to alter the ^{31}P resonant frequency, and the others contained agar, with various fat fractions. Two configurations of the phantom were scanned, with the ^{31}P containing tubes in close and distant proximity (see Figure 3.5). The tube phantom was scanned in each configuration with an acquisition weighted 16×16 phase encode 2D ^{31}P UTE CSI sequence with 32 averages at the center of k-space, for an acquisition time of 26 minutes. The CSI scan had an FOV of 240×240 mm, a bandwidth of 8000 Hz and TR of 1 s. The CSI acquisitions were followed by a matched ^{1}H localizer, to enable generation of the compartment mask. After the acquisitions were completed, the localizers were segmented into three compartments, tube 1, tube 2, and 'other' and spectra were reconstructed for each compartment using the SLAM and SLIM algorithms. To investigate the effect of changing the number of phase encodes, spectra were also reconstructed using only the central 8×8 and 4×4 phase encodes.

3.4.2.2 Leg/slab phantom

The leg/slab phantom consisted of a slab containing a K_2HPO_4 solution placed on top of a volunteer's thigh muscle (Figure 3.7), mimicking the arrangement of heart and chest wall seen in the body (Figure 3.11), with the strong signal from the slab representing the strong PCr resonance of the chest wall and weaker PCr signal from the leg representing the heart PCr resonance. The leg/slab phantom was scanned with a 16 × 16 phase encode 2D ^{31}P UTE CSI sequence with 8 averages, for an acquisition time of 24 minutes. The CSI scan had an FOV of 240 × 240 mm, a bandwidth of 8000 Hz and TR of 1 s. The CSI acquisitions were followed by a matched ^{1}H localizer, to enable generation of the compartment mask. After the acquisitions were completed, the localizer was segmented into three compartments, slab, leg, and 'other' and spectra were reconstructed for each compartment using the SLAM and SLIM algorithms. To investigate the effect of changing the number of phase encodes spectra were also reconstructed using only the central 4 × 4 phase encodes.

3.4.3 Results and discussion

3.4.3.1 Tube phantom

The results of the tube phantom experiments are shown in Figure 3.5 (SLAM reconstruction) and Figure 3.6 (SLIM reconstruction). With the tubes in the far position both the SLAM and SLIM reconstructions separated the signal well, with both tube compartments showing a single clear resonance. In the near position there was still good separation when 16 × 16 phase encodes were used (particularly SLAM), however there was significant contamination of the compartment 1 spectra by tube 2 as the number of phase encodes was decreased to 8 × 8.

These results strongly suggest that the data-handling of the TWIX and DICOM format data was correct, since each of the tube compartment spectra contain a single resonance, showing that the compartment mask and signal regions line up well, because any small misalignment would result in very low signal due to the small tube compartment sizes. The contamination seen when the tubes were in

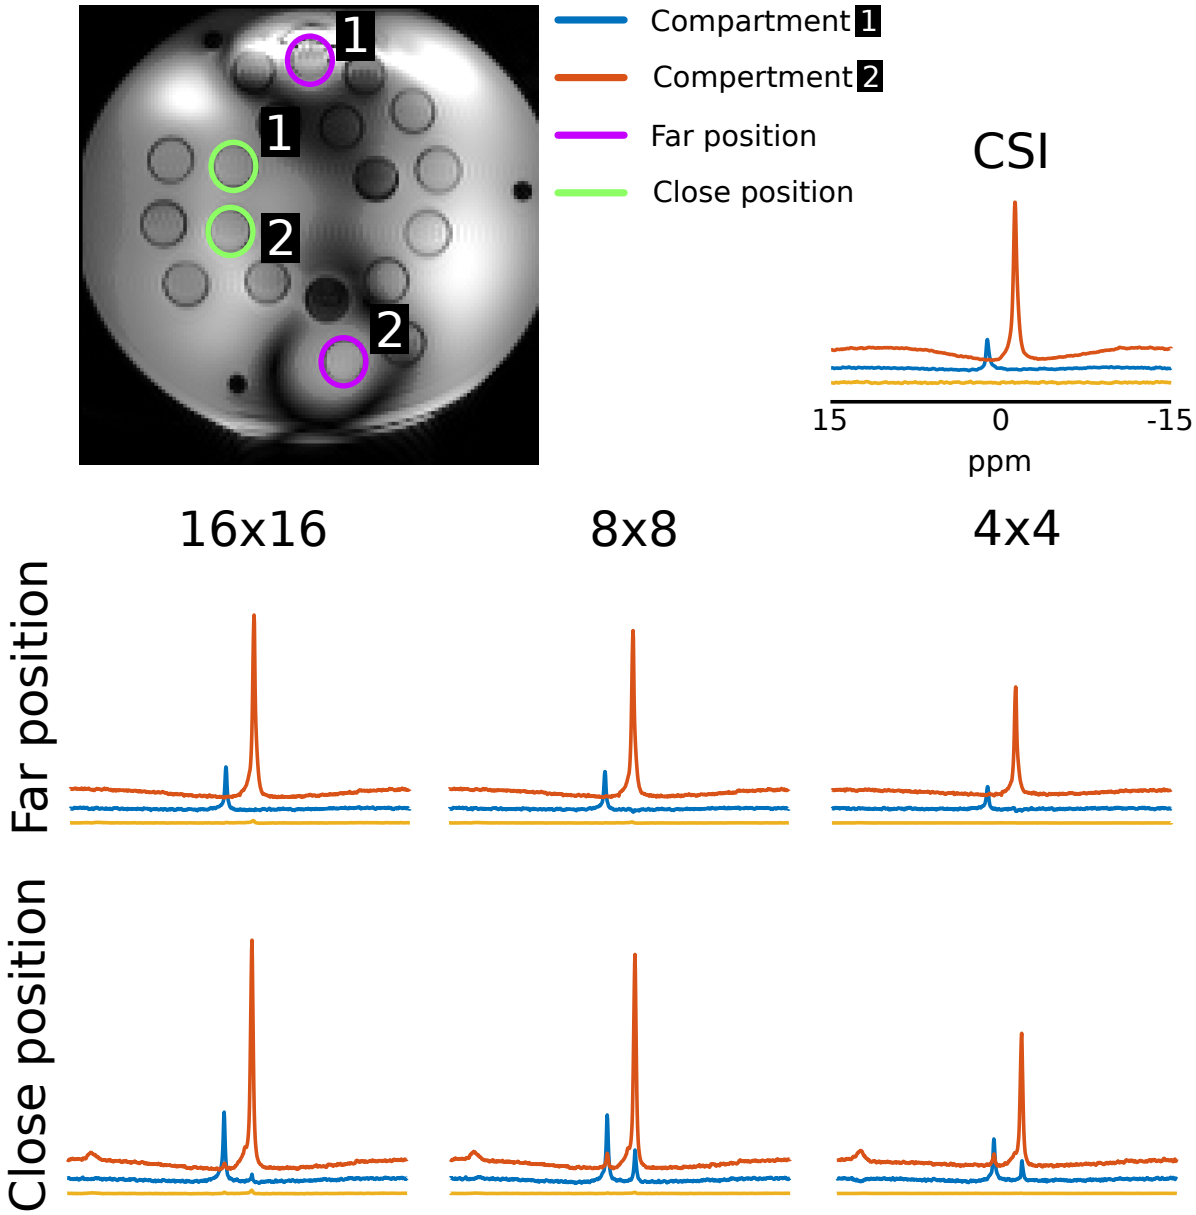

Figure 3.5: 2D SLAM reconstruction generating spectra for two tubes containing inorganic phosphate at differing pH. The experiment was repeated with the tubes in far and close proximity to assess localization. A CSI reconstruction, using 16×16 phase encodes, of the tubes in the far position is shown for comparison. The spectra were normalized by the RMS noise of the spectra from compartment 2 to allow comparison between the reconstructions.

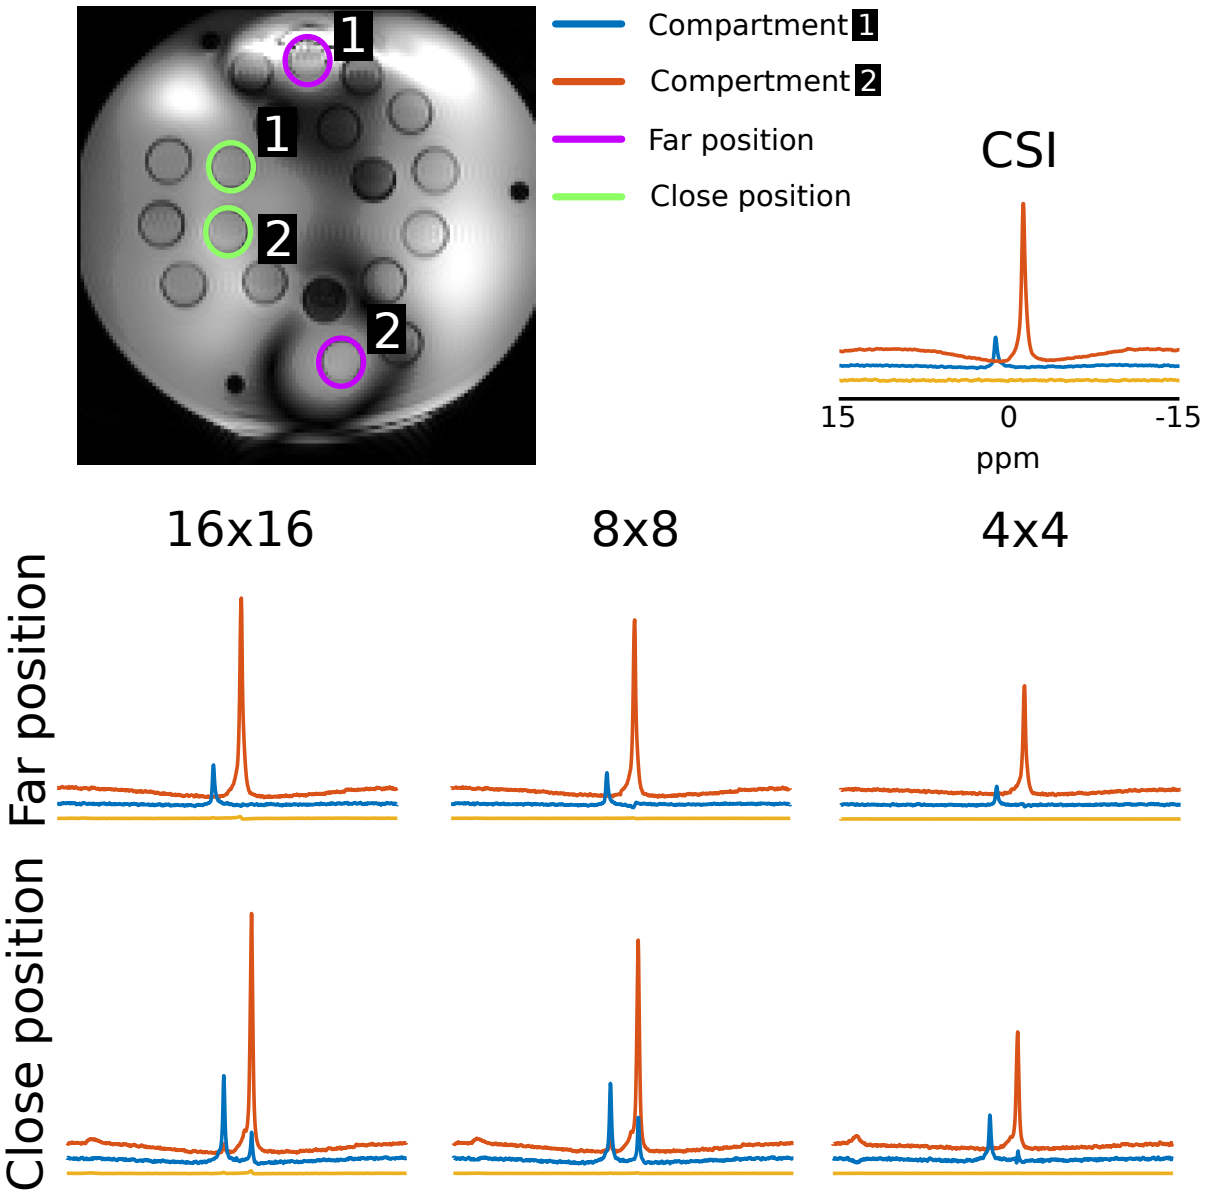

Figure 3.6: 2D SLIM reconstruction generating spectra for two tubes containing inorganic phosphate at differing pH. See Figure 3.5

the close position does however highlight the difficulties of using the technique in the in-homogeneous fields of a 7T scanner.

3.4.3.2 Leg/slab phantom

The leg/slab phantom experiment results are shown in Figure 3.7. When 16×16 phase encodes were used there was very good separation of signal between the leg and phantom compartments, especially given the large magnitude of the phantom signal

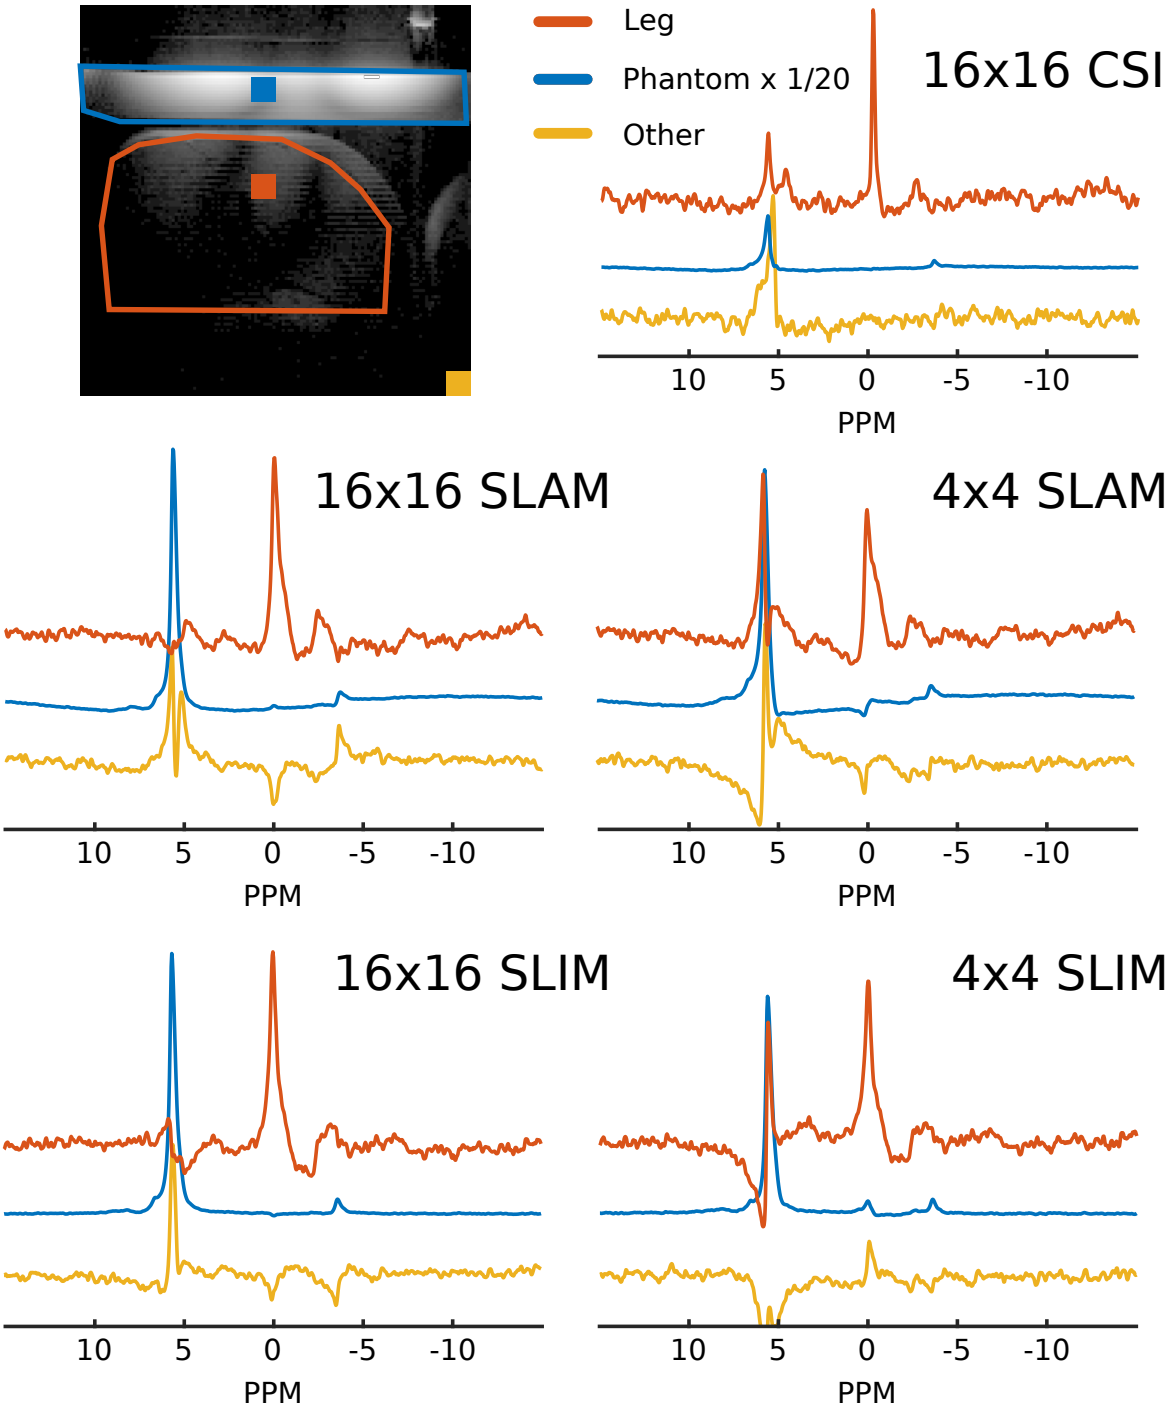

Figure 3.7: To assess the localization of compartmentalized reconstructions a phantom was placed on top of a volunteers leg. The three compartments were phantom (blue), leg (red) and other (yellow). The phantom compartment spectra were scaled by a factor of 1/20 for display purposes. Conventional CSI spectra are also shown for comparison with the location of each voxel depicted by the filled square on the localizer. The CSI spectra were scaled so that the leg voxel had equal noise amplitude to the 16×16 SLAM leg compartment.

compared to the leg signal. This is particularly reassuring because in this setup the slab and leg signals have different chemical shifts so any spectral contamination could be seen easily, unlike in ^{31}P cardiac spectroscopy where the PCr resonance from the heart and chest wall have the same chemical shift. Unfortunately with 4×4 phase encodes there was greater spectral contamination, suggesting that, in the inhomogeneous fields of a 7T scanner, it is important to incorporate larger magnitude phase encodes from further out in k-space to prevent unwanted contamination.

3.5 Repeatability study

3.5.1 Introduction

In this third section, we compare the repeatability of SLAM and SLIM reconstruction algorithms with our existing single voxel CSI reconstruction algorithm, by reconstructing a single acquisition-weighted test-retest data-set with all three algorithms. We also evaluate separate truncated k-space data-sets acquired in the same scan sessions ($4 \times 4 \times 4$ central k-space phase encodes and fSLAM optimized phase encodes) to investigate if using these phase encoding schemes with SLAM and SLIM reconstruction algorithms can yield further SNR improvements, in the same acquisition time, by including more averages of phase encodes close to the center of k-space. SNR, PCr/ATP ratio and coefficients of variation and repeatability were calculated for each acquisition-reconstruction combination, to quantify any differences, and SRFs for each of the acquisition-reconstruction combinations were calculated to rationalize the results.

3.5.2 Methods

Twelve healthy volunteers (6F/6M, age 29 ± 4 years, BMI 23 ± 3 kg/m^2) were recruited and gave informed consent for this study which was approved by The University of Oxford Central University Research Ethics Committee. The repeatability of each method was quantified using a test-retest procedure with the volunteer taking a short break outside of the scan room between tests. Each method consisted of a combination of a phase encoding scheme and reconstruction algorithm.

The phase encoding schemes were; 3D acquisition weighted k-space points (AW), $4 \times 4 \times 4$ central integer k-space points ($4 \times 4 \times 4$) and $4 \times 4 \times 4$ k-space points whose positions were optimized with the fSLAM algorithm (fSLAM), as detailed in Section 3.5.2.2. The reconstruction algorithms were the SLAM, SLIM and CSI algorithms, implemented as described in Section 3.5.2.3. This gave a total of six combinations: AW CSI, AW SLAM, AW SLIM, $4 \times 4 \times 4$ SLAM, $4 \times 4 \times 4$ SLIM and fSLAM SLAM (fSLAM SLAM is henceforth referred to as fSLAM). The scan protocol consisted of a series of 2D localized MRI acquisitions spanning the 3D volume which were used to generate the compartmentalization mask, and three ^{31}P MRS acquisitions, corresponding to the AW, $4 \times 4 \times 4$, and fSLAM phase encoding schemes. An overview of the process is shown in Figure 3.8, with the data pathways for each reconstruction combination detailed in Figure 3.9 and the scans summarized in Table 3.2 and Figure 3.10.

Technique name	Acquisition	Reconstruction algorithm	Phase encodes	Averages	Time (min)
AW CSI	AW k-space	Fourier transform	$8 \times 16 \times 8$	4	6.5
AW SLAM		SLAM			
AW SLIM		SLIM			
$4 \times 4 \times 4$ SLAM	Central integer k-space	SLAM	$4 \times 4 \times 4$	6	6.5
$4 \times 4 \times 4$ SLIM		SLIM			
fSLAM	fSLAM optimized k-space	SLAM	$4 \times 4 \times 4$	6	6.5

Table 3.2: Summary of reconstructions and scans performed.

3.5.2.1 SRF simulations

Spatial response functions (SRFs) for a representative compartmentalization mask (Figure 3.11A) were calculated for each of the acquisition and reconstruction algorithm combinations listed in Table 3.2. The SRFs were calculated using the formula described in Section 3.2, and weighted with receive coil sensitivity (B_1^-)calculated using the Biot-Savart law, which is given by

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_C \frac{I \, \mathrm{d}\mathcal{L} \times (\mathbf{r} - \mathcal{L})}{|\mathbf{r} - \mathcal{L}|^3} \tag{3.32}$$

where r is the position vector, C the path of the coil, $\mathcal{L}$ the position on the coil, and I is the current flowing in the coil. Although this treatment may not be valid for ^{1}H MRI at 7T[113], since the quasi-static approximation breaks down, the Larmor frequency of ^{31}P is sufficiently lower (120 vs 300 MHz) that it may be used here.

In the unweighted (theoretical) case, the sum over the compartment of its corresponding SRF was one and the SRF maps are presented as magnitude and phase images.

3.5.2.2 Data acquisition

All acquisitions were undertaken on a Siemens MAGNETOM 7T whole body MRI scanner (Siemens Healthineers, Erlangen, Germany) with volunteers in the supine position. The RF coil used was a dual-tuned (^{31}P/^{1}H) transmit-receive quadrature, surface coil consisting of a ^{31}P coil (two 15 cm loops, with overlap decoupling) and a single ^{1}H loop (10 cm in diameter)[114]. Two-chamber, four-chamber, short axis and CSI-matched ^{1}H images were acquired with a fast low-angle shot (FLASH) sequence[115]. The coil position was determined from images showing the location of 4 cod-liver oil capsules and one central phenylphosphonic acid fiducial embedded in the coil housing during imaging and the transmit flip-angle was calculated from a series of inversion-recovery free induction decay acquisitions of the central fiducial as previously described[103].

Three different ^{31}P acquisitions, corresponding to the AW, $4 \times 4 \times 4$ and fSLAM phase encoding schemes, were applied for the comparative studies. The

^{31}P acquisitions consisted of a 3D ultra-short echo time (UTE)-CSI sequence[116] The excitation was achieved with a shaped RF pulse, numerically optimized for excitation homogeneity[57]. The excitation had a bandwidth of 8000 Hz and was placed with its center frequency at +266 Hz (towards the γ-ATP resonance) relative to the PCr resonance. The repetition time (TR) was 1 s.

The phase encoding schemes used in the acquisitions are summarized in Table 3.2 and Figure 3.10. All had 3D phase encoding, a FOV of $240 \times 240 \times 200$ mm and took 6.5 minutes to complete. The first phase encoding scheme (AW) was 3D acquisition weighted, with a matrix size of $8 \times 16 \times 8$ (x,y,z) and 4 averages at $k = (0, 0, 0)$. The second scheme, ($4 \times 4 \times 4$) acquired data at phase encodes corresponding to the central $4 \times 4 \times 4$ k-space coordinates of the acquisition weighted scheme, but with 6 averages at all k-space points. The final phase encoding scheme (fSLAM) consisted of $4 \times 4 \times 4$ phase encodes with 6 averages at each k-space point, however the exact fractional k-space values were optimized using the fSLAM algorithm. The k-space coordinates used for each acquisition are detailed Figure 3.10.

To enable calculation of the fSLAM coordinates while the AW and $4 \times 4 \times 4$ phase encoding scheme acquisitions completed, the degrees of freedom in the optimization were reduced to 12, after forming a 1D vector comprised of the k-space coordinates in each orthogonal direction. The I_{cond} parameter was set to 5, ensuring that the condition number of the $G_{P\times C}$ matrix did not exceed 5 and the reconstruction remained a well posed problem. The optimizations were performed with the Matlab programming language and required significant computing power taking approximately 2.5 minutes to run on our workstation (2× Intel Xeon Gold 6252, 48 cores / 96 threads; 768 GiB RAM).

An overview of the complete scan protocol is shown in Figure 3.8A and the data generation and analysis in Figure 3.8B.

3.5.2.3 Data analysis

The SLAM and SLIM algorithms, in addition to a standard Fourier based reconstruction (CSI), were implemented in the Matlab programming language (Natick

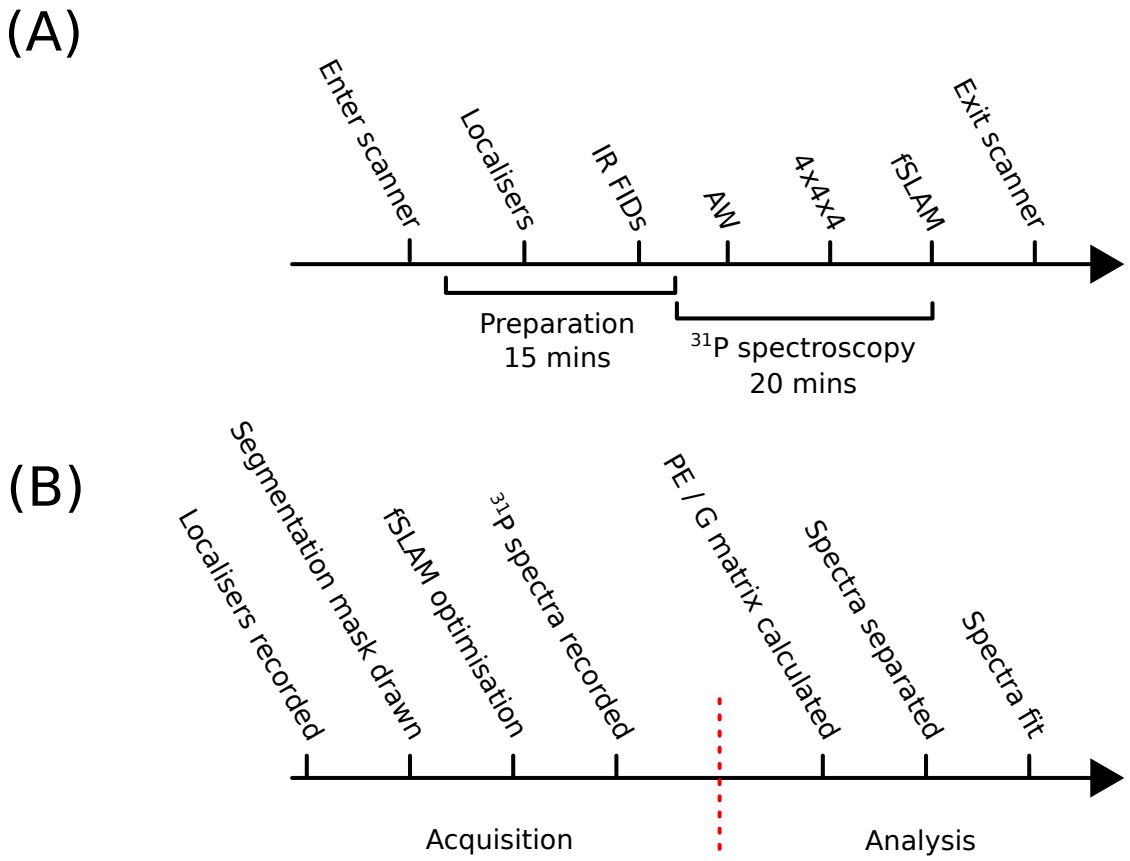

Figure 3.8: (A) Timeline showing the acquisitions performed on a volunteer during each scan. (B) Timeline of data generation and analysis, the segmentation mask is fixed before data analysis.

MA) in accordance with the equations presented in the theory section. An alternative method for solving the SLAM equation (SLAM*)[111], did not produce a significant improvement in the results from our 7T datasets and was not used for our analysis. For cardiac CSI measurements, a single midseptal voxel spectrum was chosen, because summing multiple voxels produced contamination artefacts (see Section 3.6.1). The compartmentalization masks used in the reconstruction were defined during acquisition and consisted of 3 compartments; heart, chest wall, and other (see Figure 3.11A). If present the heart compartment was removed from the final two slices (base of heart) to avoid aliasing of the heart compartment's SRF into the first slice. The raw spectra from each coil element were combined with a noise-whitened singular value decomposition coil combination[117] and fitted with the Oxford Spectroscopy Analysis (OXSA) implementation of AMARES[118, 119]

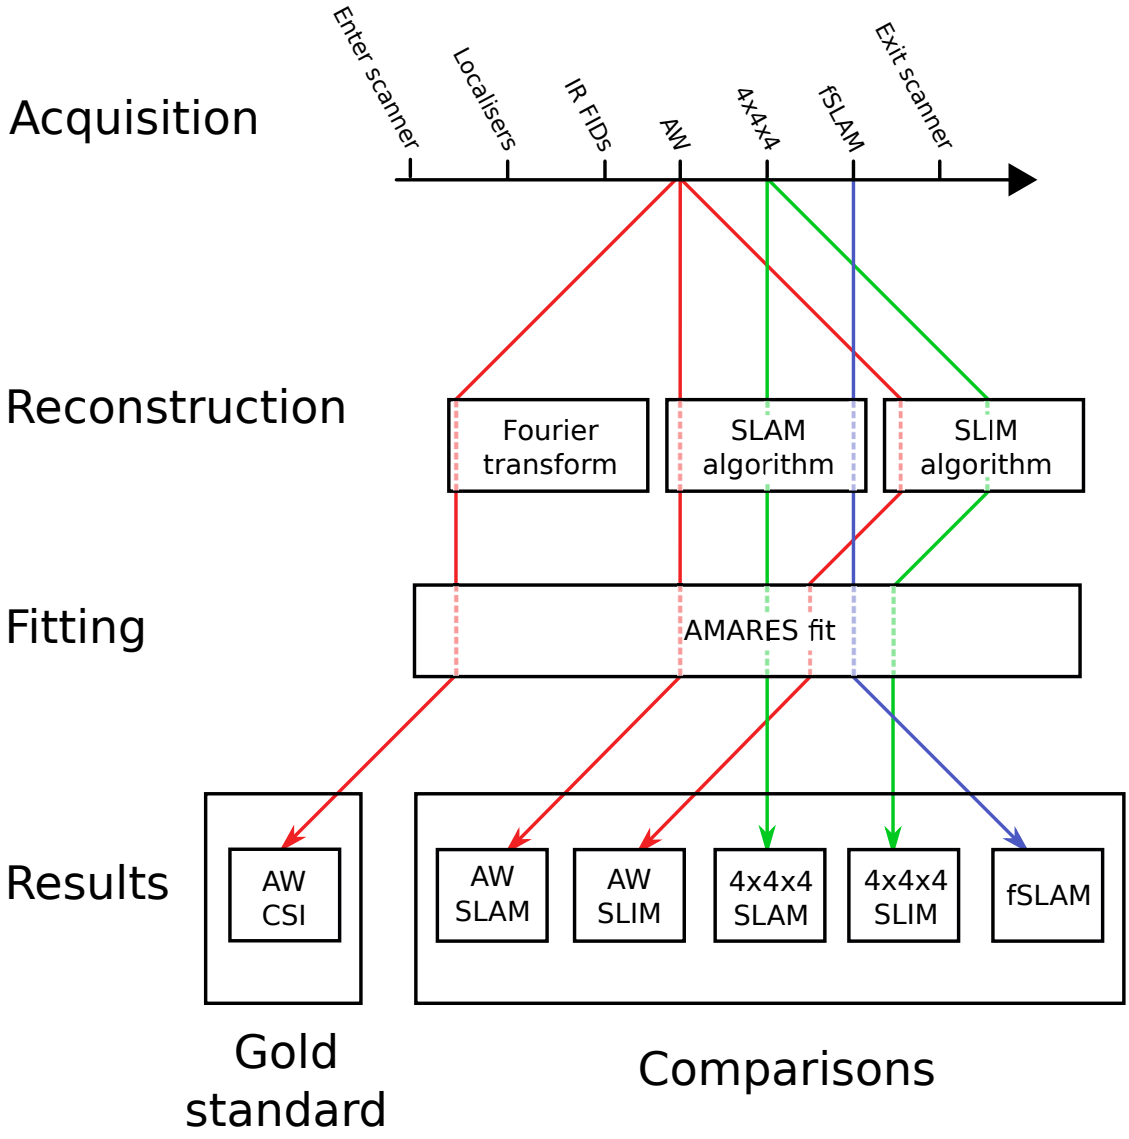

Figure 3.9: An overview of the data pathways for each acquisition/reconstruction combination considered from acquisition through reconstruction and fitting to results.

with an identical protocol used for all spectra. PCr/(γ-)ATP ratios were blood and saturation corrected as previously reported[103]. The SNR of the PCr resonance and the Cramér-Rao lower bound (CRLB) of the PCr fit were calculated. The coefficient of repeatability (CoR) was defined as[56]

$$\mathrm{CoR} = \mathrm{SD}_{\mathrm{intrasubject}} \times 1.96, \tag{3.33}$$

where $\mathrm{SD}_{\mathrm{intrasubject}}$ is the standard deviation (SD) of the signed difference in

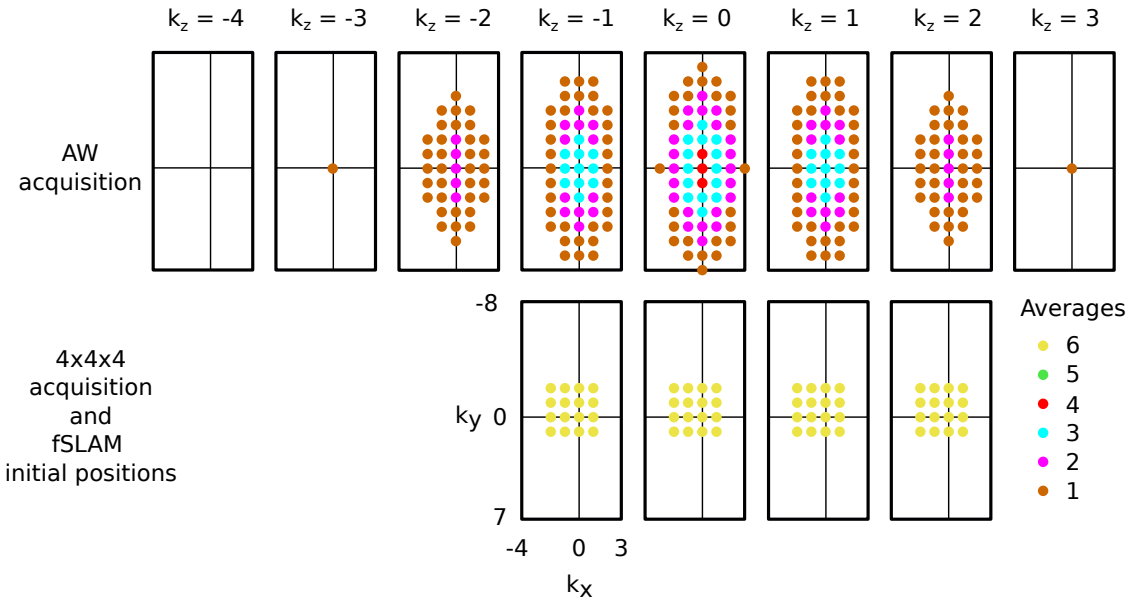

Figure 3.10: k-space coordinates of the phase encodes used in the three ^{31}P acquisitions, all k-space coordinates are normalized by the FOV to give integer values, the fSLAM initial positions refer to the k-space values used as the starting conditions for the fSLAM optimization algorithm. The final fractional fSLAM optimized k-space coordinates are different for each compartmentalization mask and are defined here by a mesh of k_x, k_y and k_z values. For the compartmentalization mask used in Figure 3.11 these are: k_x =-1.94, -0.99, 0.01, 1.08, k_y = -2.19, -1.11, -0.01, 1.13, k_z = -1.83, -0.96, 0.00, 0.93.

PCr/ATP ratio between scans of the same subject with the same technique. The coefficient of variation (CoV) was defined as the SD of a measure across all scans (both sessions) of the same technique divided by the mean value of the measure. Spectra from one participant with excessive movement, and 3 of the remaining data points that showed very high liver contamination were excluded from the analysis. Errors are reported as ± 1 SD. All comparisons were to the AW CSI method, which is considered the standard method at our center. Pairwise two-tailed Wilcoxon signed-rank tests with 5% significance level were used for all comparisons.

3.5.3 Results

An overview of the results at each stage of the reconstruction process is given in Figure 3.11. It shows a representative compartmentalized slice through the heart, the spectra generated by reconstructing the AW data-set corresponding to that volunteer using the SLAM algorithm, a sample fit and comparison of cardiac

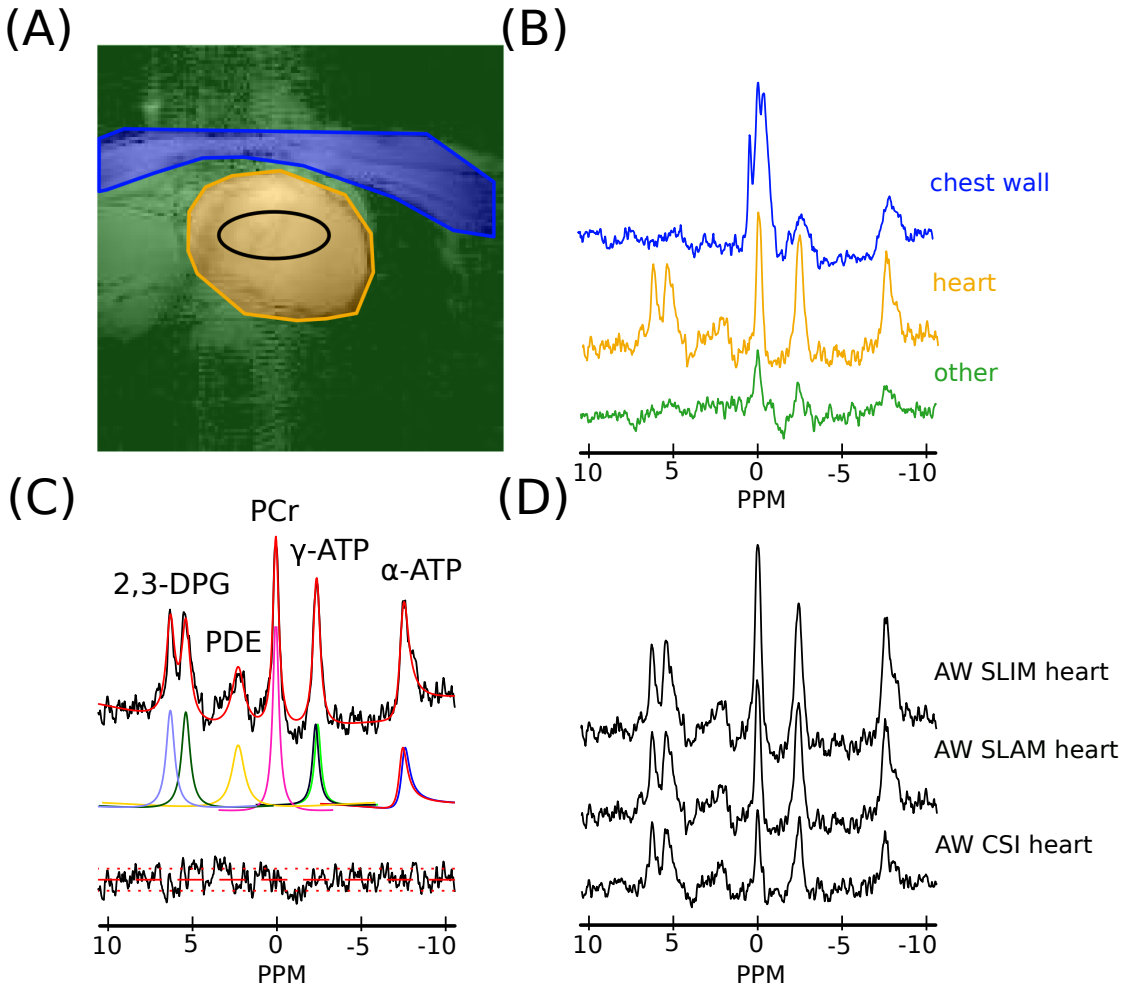

Figure 3.11: (A) A sample segmentation map of one slice of the heart, showing (blue) chest wall, (orange) heart, (green) other. The black ring represents the the midseptal voxel (64% threshold of voxel PSF) of the AW CSI method. (B) Sample spectra for each of the compartments in (A) reconstructed using the SLAM algorithm and AW phase encoding scheme (AW SLAM). (C) Fit, using the OXSA toolbox of the heart compartment spectra in (B). (D) AW SLIM and AW SLAM spectra for the heart compartment in (A) with the AW CSI spectra of the midseptal voxel. Spectra in B and D are normalized by noise.

spectra reconstructed from the same AW data-set using the SLAM, SLIM and CSI reconstruction algorithms. Clear spectral differences between the chest wall and heart compartment can be seen, in addition to the increase in SNR from the SLAM/SLIM algorithm compared to the CSI algorithm.

3.5.3.1 SRF simulations

Representative cardiac compartment SRFs for each of the compartmentalized acquisitions are shown in Figure 3.12 with and without coil sensitivity weighting

(B_1^- Rx). The segmentation of the heart as a spheroid resulted in the maximum of the SRF placed on the midseptal voxel, which has previously been shown to be the optimum voxel for the AW CSI method used here[56], and extending out to the wall of the heart. The AW acquisitions feature an elliptical SRF, with the semi-minor axis in the vertical y-direction, reflecting the higher number of phase encodes, and therefore better spatial resolution in the y-direction. The inclusion of coil sensitivity weighting to the theoretical (un-weighted) SRFs skews the AW SLAM and AW SLIM SRFs towards the top of the heart, in the direction of the coil. The $4 \times 4 \times 4$ SLAM and $4 \times 4 \times 4$ SLIM SRFs both show slightly greater over-spill of the cardiac compartment's SRF into the chest wall than AW SLAM and AW SLIM and a more spherical SRF. The differences between $4 \times 4 \times 4$ SLAM and fSLAM are subtle, however fSLAM does appear to reduce some spill over of the SFR into the chest wall, particularly once coil sensitivity is taken into account.

3.5.3.2 Acquisition weighted k-space acquisitions

Figures 3.13 and 3.14 show box-plots of the PCr/ATP ratio and PCr SNR recorded in the study. Numerical values for the coefficients of repeatability and variation, PCr/ATP ratio, PCr SNR and PCr fit CRLB are given in Table 3.3. The AW CSI method (standard method for our center) measured a PCr/ATP ratio of 1.90 $\pm$ 0.92 with a PCr SNR of 13.10 $\pm$ 8.63. The CoR was 1.79 and CoV 0.48.

In comparison the AW SLAM and AW SLIM both had a significantly higher PCr SNR than the AW CSI method (Wilcoxon, paired, $\alpha = 0.05$). Both had lower CoR and CoV than the AW CSI method. AW SLAM did not have a significantly different PCr/ATP ratio to the AW CSI method, however the AW SLIM PCr/ATP ratio was significantly higher.

3.5.3.3 Truncated k-space acquisitions

The $4 \times 4 \times 4$ SLAM and $4 \times 4 \times 4$ SLIM reconstructions both had a significantly higher PCr SNR than the standard AW CSI method, however the CoV and CoR for the $4 \times 4 \times 4$ SLAM and $4 \times 4 \times 4$ SLIM methods were higher (worse) than for the AW CSI method. The $4 \times 4 \times 4$ SLAM PCr/ATP ratio was not significantly

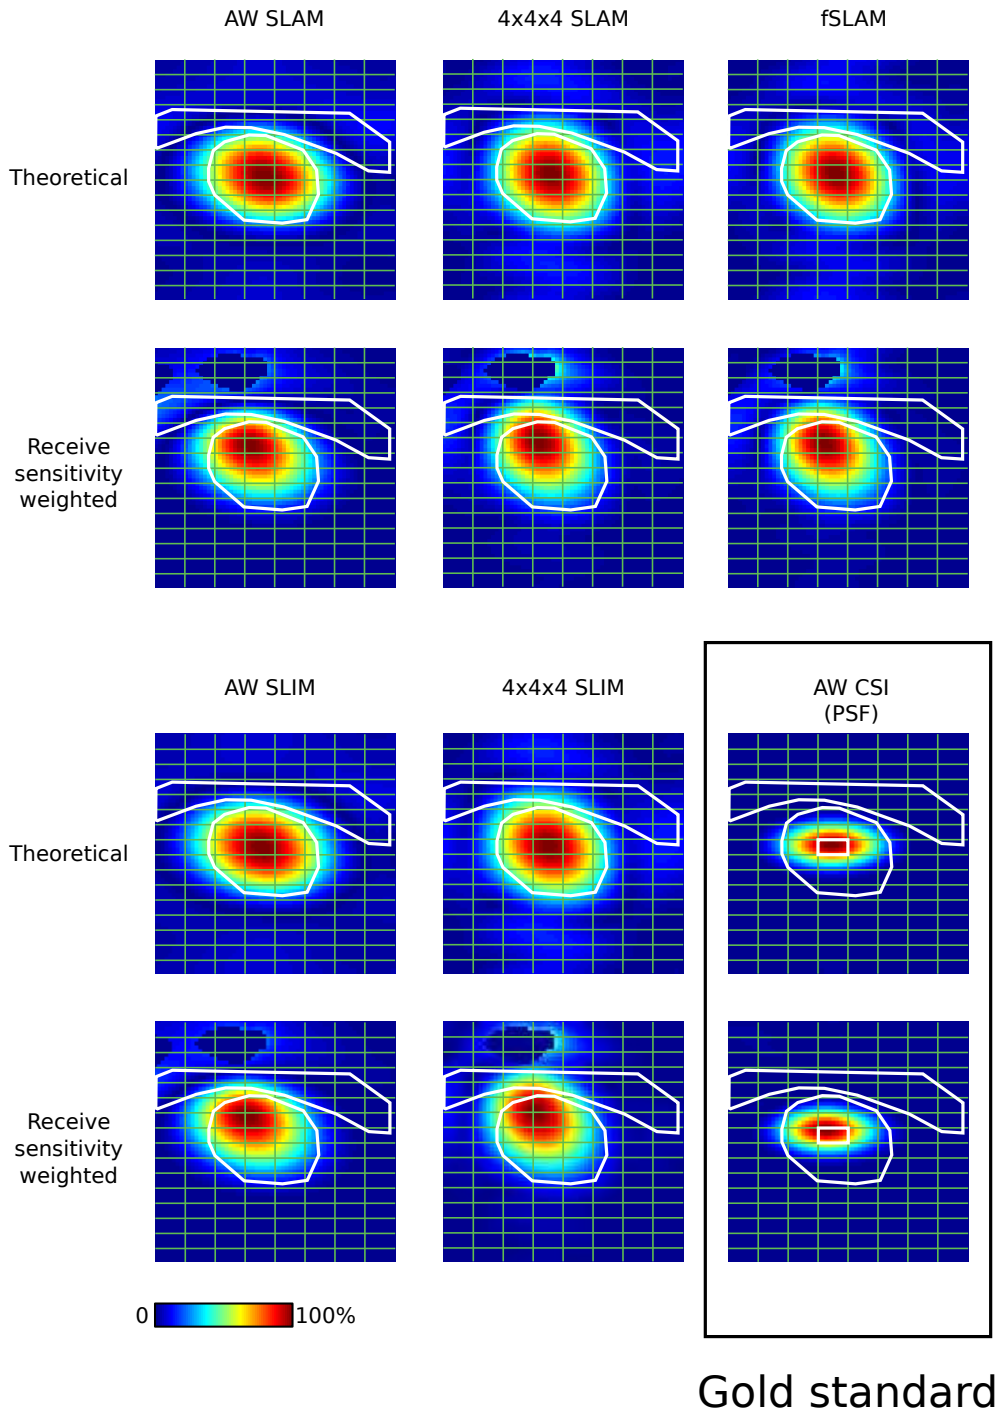

Figure 3.12: Simulated cardiac compartment SRF magnitude images for each SLAM/SLIM algorithm based method for a representitive compartmentalization mask (Figure 3.11A), and AW CSI midseptal voxel PSF (same as Figure 3.11). Both theoretical and B_1 recieve sensitivity weighted SRFs are shown. The CSI grid is shown in green and the midseptal voxel highlighted in white in the AW CSI PSFs.

different from the AW CSI method PCr/ATP ratio, however the $4 \times 4 \times 4$ SLIM PCr/ATP ratio was significantly higher.

The fSLAM method had a significantly higher PCr SNR than the AW CSI method, and a lower (better) CoV, although CoR was higher. The fSLAM method's PCr/ATP ratio was also not significantly different to the AW CSI method's PCr/ATP ratio.

Technique	PCr/ATP Ratio	PCr SNR	PCr CRLB	CoR	CoV	removed
AW CSI	1.90 ±0.92	13.10 ±8.63	13.36±7.37	1.79	0.48	0/22
AW SLAM	1.68 ±0.50	20.75 ±11.18	8.08±4.19	1.22	0.30	0/22
AW SLIM	2.09 ±0.43*	26.93 ±14.95	5.71±2.57	1.05	0.21	0/22
$4 \times 4 \times 4$ SLAM	1.90 ±1.13	25.43 ±11.44	7.60±4.20	3.29	0.59	2/22
$4 \times 4 \times 4$ SLIM	2.46 ±1.46*	35.05 ±18.26	5.57±3.26	3.49	0.59	1/22
fSLAM	2.02 ±0.88	25.05 ±13.29	9.54±10.75	2.43	0.43	0/22

Table 3.3: Mean values of the experimental results derived for each technique, errors are given as ± 1 SD. * indicates PCr/ATP ratio significantly different to the AW CSI method, all PCr SNR values were significantly different to the the AW CSI method. The horizontal lines separate each of the acquisition methods. The removed column is the number of data-points that needed to be removed due to liver contamination.

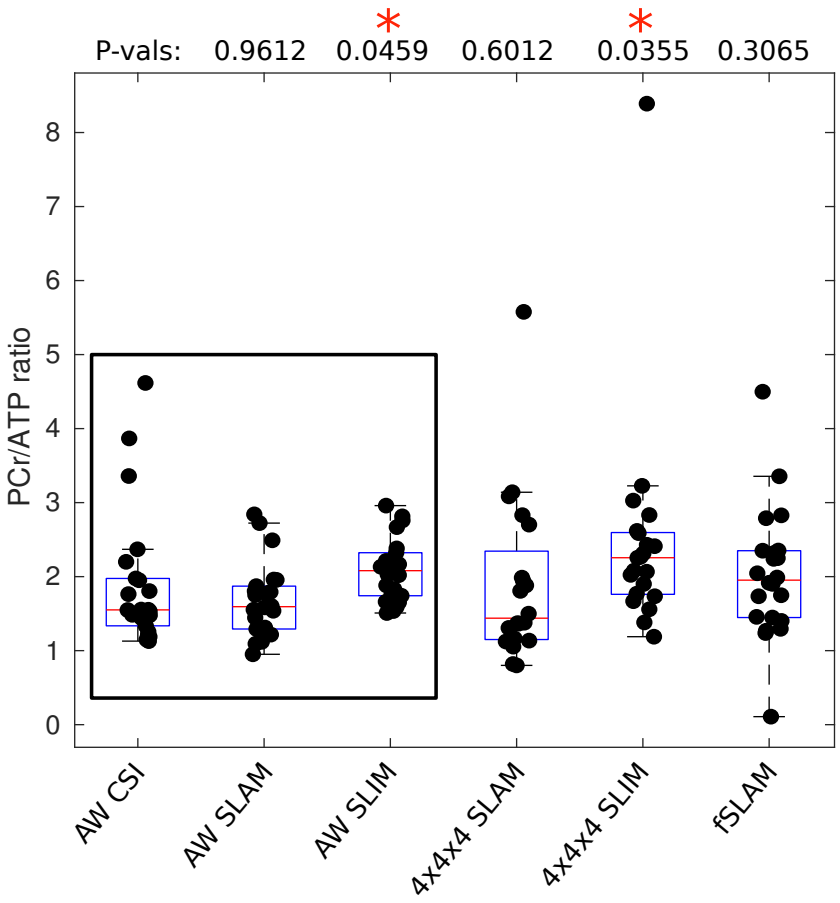

Figure 3.13: Box-plot showing PCr/ATP values for each acquisition, median and IQR indicated by box. * indicates significant difference to the AW CSI method (Wilcoxon signed-rank paired, α=0.05, P-values above box-plot). Reconstructions which use the AW phase encoding scheme are highlighted by the surrounding black box. PCr/ATP values are in the expected range (Table 3.4)

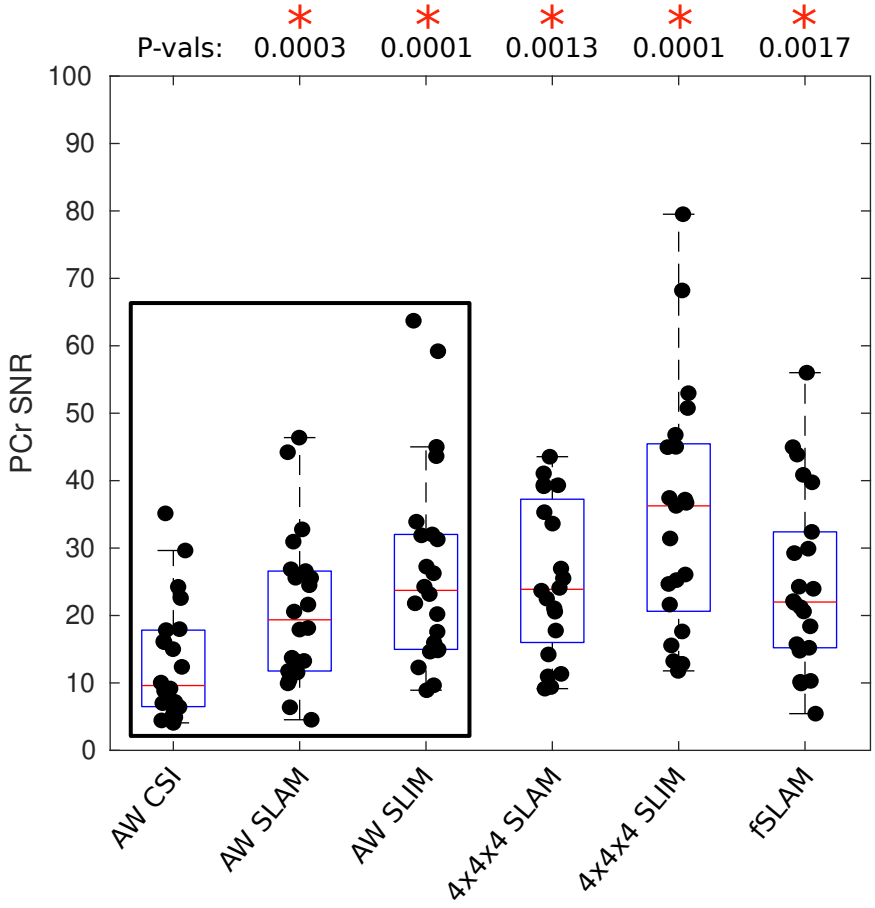

Figure 3.14: Box-plot showing PCr SNR values for each acquisition, median and IQR indicated by box. * indicates significant difference to the AW CSI method (Wilcoxon signed-rank paired, α=0.05, P-values above box-plot). Reconstructions which use the AW phase encoding scheme are highlighted by the surrounding black box.

3.5.4 Discussion

In this work we investigated the use of SLAM and SLIM reconstruction algorithms to reconstruct a standard AW phase encode acquisition (AW SLAM and AW SLIM) at 7T, and compared this to our existing CSI reconstruction of the same AW data (AW CSI). We also evaluated two alternative acquisition phase encode schemes (with equal acquisition time to the AW CSI) for reconstruction with the SLAM and SLIM algorithms; $4 \times 4 \times 4$ k-space phase encodes and fSLAM optimized k-space phase encodes (reconstructed with the SLAM algorithm only). We found that all the compartmentalized acquisition-reconstruction combinations delivered a significant improvement in SNR over AW CSI, and that AW SLAM also improved repeatability while maintaining a not significantly different PCr/ATP ratio. However the increase in SNR, and consequently the lower CRLBs, did not uniformly translate into reduced variance across the three acquisition data-sets, suggesting that the spread in PCr/ATP ratio is dominated by factors other than technique, such as biological variation and motion artifacts, as well as the inherent increase in uncertainty for a ratio compared to a single value, which are areas of future study.

3.5.4.1 Acquisition weighted k-space acquisitions

We recorded a PCr/ATP of 1.90 ± 0.92 using the AW CSI method, which is in good agreement with the previous literature values at 1.5, 3 and 7T (Table 3.4). The AW CSI method had a CoR of 1.79 and CoV of 0.48. The CoR reported by Ellis et al. (0.67)[56], who used the same AW CSI method, is lower than our result, however the CoV (0.38) is comparable. This is probably due to the different receive coil designs used in the two studies, where Ellis et al. used a 16 channel receive array, which would have been expected to deliver better SNR and B_1 homogeneity than the 2 channel quadrature coil used in this study[2]. In both our work and that of Ellis et al.[3] some healthy volunteers had PCr/ATP ratios of between 1 and 1.5, which is lower than expected for a healthy volunteer cohort, however

[2] The 16 channel array coil sustained significant damage prior to the start of this work which was not possible to repair

[3] see Ellis et al.[56] Figure 4

for our data-set these volunteer's ^{31}P spectra commonly had very low SNR, and correspondingly low CRLBs of the fitted peaks, which would have been routinely excluded from analysis in clinical research.

Both the AW SLAM and AW SLIM methods improve SNR over the AW CSI method while improving repeatability. The PCr/ATP ratio is also maintained for AW SLAM, with a ratio of 1.68 ± 0.50, close to the 1.71 ±0.65 midseptal voxel value reported by Ellis et al, suggesting good rejection of out of compartment signals. Although the AW SLIM method returns a significantly higher PCr/ATP ratio than the AW CSI method, it is still well within the range of literature values for healthy volunteers (Table 3.4), including studies at the same site at 3T, and likely due to the subtle differences between the SLAM and SLIM algorithms.

Since the same AW data set is input into both SLAM, SLIM and CSI reconstruction algorithms, it is likely that the SNR increase seen here is due to the increase in coil sensitivity weighted tissue volume (see Figure 3.12) of the cardiac compartment over the single voxel used in the CSI algorithm. This is in contrast to the previous body of work on SLAM[109] which used a reduced set of high SNR central k-space phase encodes to achieve a per unit volume increase in SNR over a higher resolution uniform set of k-space phase encodes. However, the increase in sensitive volume, while maintaining the PCr/ATP ratio is only possible due to the compartmentalized reconstruction producing a cardiac compartment SRF which closely matches the shape of the heart, unlike the summation of AW CSI voxels (see Section 3.6.1).

These results suggest that the AW SLAM and SLIM methods are not compromised by the greater B_0 and B_1 inhomogeneity found at 7T than 3T (though a whole-body transmit coil with greater B_1 homogeneity could provide further improvements[120]) and are therefore able to provide a worthwhile upgrade over the AW CSI method implemented here. This is particularly the case for AW SLAM whose PCr/ATP ratio is not significantly different to AW CSI.

Study	Field Strength	Technique	PCr/ATP Ratio
Ellis et al.[56]	7	CSI	1.71 ±0.65
Rodgers et al.[103]	7	CSI	2.1 ±0.3
Tyler et al.[57]	3	CSI	2.07 ±0.38
			2.14 ±0.46
Lamb et al.[121]	1.5	3D ISIS	1.31 ±0.19
This study	7	AW CSI	1.90 ±0.92
		AW SLAM	1.68 ±0.50
		AW SLIM	2.09 ±0.43

Table 3.4: Select literature PCr/ATP values.

3.5.4.2 Truncated k-space acquisitions

$4 \times 4 \times 4$ SLAM and $4 \times 4 \times 4$ SLIM SNR was significantly increased over the AW CSI method, although neither $4 \times 4 \times 4$ SLAM or $4 \times 4 \times 4$ SLIM showed improved repeatability (CoV and CoR) over AW CSI, and the $4 \times 4 \times 4$ SLIM method had a significantly higher PCr/ATP ratio than the AW CSI method. It is worth noting however, that the $4 \times 4 \times 4$ acquisition was completed after the AW acquisition so the volunteers may have been more restless for this scan, potentially increasing errors from volunteer motion. As expected, optimizing the $4 \times 4 \times 4$ phase encode scheme on a per-volunteer basis with the fSLAM optimization algorithm improved both CoV and CoR, as compared to the $4 \times 4 \times 4$ SLAM and $4 \times 4 \times 4$ SLIM methods, however the CoV and CoR for fSLAM were still higher than for the AW SLAM method, suggesting that the constraints placed on degrees of freedom of the fSLAM phase encode optimization, to achieve a short enough computation time, may have prevented the most optimal solution being found.

Overall we were not able to demonstrate a significant benefit to using $4 \times 4 \times 4$ SLAM, $4 \times 4 \times 4$ SLIM or fSLAM methods over the AW SLAM method, as the improved SNR did not translate to improved repeatability.

3.5.4.3 SLAM and SLIM reconstruction algorithms

Although SLAM and SLIM share mathematical underpinnings, there are marked differences in the results, which must be attributed to the reconstruction, since in the case of the $4 \times 4 \times 4$ and AW phase encoding schemes, the data and

compartmentalization masks are identical. Over all the acquisition data-sets, the SLIM algorithm based reconstructions gave a significantly higher PCr/ATP ratio than the AW CSI method whereas the SLAM algorithm based methods did not, indicating lower PCr contamination of the heart compartment by the chest wall.

The calculated SRFs suggest that cardiac/chest wall movement through breathing may be an important factor to consider for both SLAM and SLIM algorithm based methods. Since the SRF is much larger than for a conventional CSI, it is more likely that the chest wall will stray into the cardiac SRF during breathing, potentially leading to contamination of the cardiac spectra. While this may present an issue, AW SLAM and AW SLIM had a lower CoV and CoR than the AW CSI method, and cardiac/respiratory gating may be able to ameliorate this concern further.

3.6 Further experiments

3.6.1 Summation of CSI voxels

Summation of spectra, reconstructed with the CSI method, corresponding to voxels located within the heart, may increase the SNR of the reconstructed cardiac spectrum, by expanding the size of the sensitive volume, however taking voxels close to the chest wall may lead to increased chest wall PCr contamination.

To test the effect of summing cardiac voxels the following experiment was performed. The AW CSI method was modified so instead of taking just the midseptal voxel, all voxels corresponding to the heart compartment were individually phased and summed together. The calculated statistics were as follows: PCr/ATP = 2.22 ± 0.49*, SNR = 34.32 ± 16.61*, CRLB = 4.42 ± 1.86, CoR = 0.66, CoV = 0.22. (* indicates significantly different to the AW CSI method – paired Wilcoxon signed-rank test p = 5%)

The high SNR and low CoR and CoV, relative to AW CSI can be explained as follows. When the spectra are summed, the spectra originating from voxels close to the chest wall, which have high SNR due to the coil B1 profile, dominate the resultant spectra, giving it high SNR and consequently low CoR and CoV. However, the spectra from voxels close to the chest wall are also contaminated by the strong

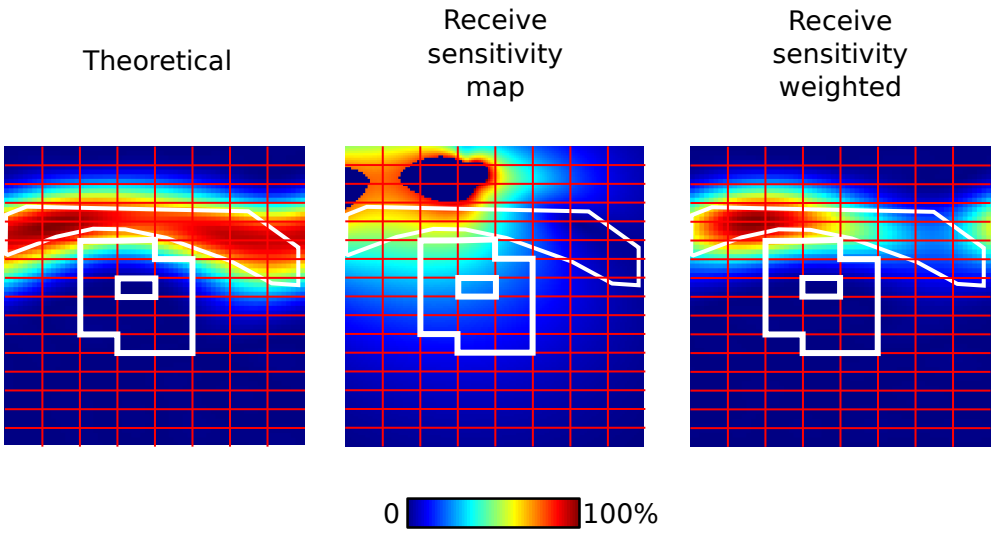

Figure 3.15: The summed point spread function (PSF) of voxels corresponding to the chest wall are shown on the left and right (note difference to SRFs shown in main body text). The receive sensitivity map is shown in the center. The cardiac voxels are highlighted with the white bounding line and the midseptal voxel is additionally highlighted with a white bounding box in the center of the heart. Contamination of the voxels forward of the midseptal voxel by the chest wall signal is evident, as is the increased coil sensitivity in this region.

chest wall PCr signal, meaning that the summed spectra has an artificially inflated PCr resonance, leading to a PCr/ATP which is significantly higher than AW CSI. This makes summation of all (or a subset of high SNR) cardiac voxels inadvisable, as the resultant spectra will not reflect the true underlying cardiac PCr/ATP ratio.

Although the worsening of CoR seen by Ellis et al.[56] when summing up to four voxels, is contrary to this result, we hypothesize that this due to the use of saturation bands in their study, which give a relatively sharp cut-off to chest wall PCr contamination, making the exact voxel position (which would vary between scan 1 and scan 2 of the test re-test procedure) critical.

3.6.2 Acquisition validation

It was important to ensure that there were no systematic differences between the AW and $4 \times 4 \times 4$ acquisitions described in Section 3.5.2.2, other than the different phase encoding schemes. Although the two acquisitions should be directly comparable there were some causes for concern, firstly the AW acquisition was always completed before the $4 \times 4 \times 4$ acquisition, which could lead to systematic

Technique	PCr/ATP Ratio	PCr SNR	Removed	CoR	CoV
AW CSI	1.90±0.92	13.10±8.63	0/22	1.79	0.48
Reduced AW SLAM	1.82±0.55	11.72±3.63	10/22	0.85	0.30
Reduced AW SLIM	2.44±1.07	16.85±8.05	7/22	1.60	0.44
Reduced $4 \times 4 \times 4$ SLAM	1.82±0.87	13.03±3.48	9/22	1.86	0.47
Reduced $4 \times 4 \times 4$ SLIM	1.88±0.61	16.07±5.90	5/22	1.34	0.32

Table 3.5: Computed statistics for the reduced data-sets comprising of one average of $4 \times 4 \times 4$ phase encodes. Errors are given as ± 1 SD. CRLBs are not given as the data is thresholded by PCR CRLB <20%. The horizontal lines separate each of the acquisition methods.

physiological differences between the scans (e.g. heart rate, bladder fullness or movement relative to the localizer) and secondly the $4 \times 4 \times 4$ (and fSLAM) phase encodes were specified by text file, unlike the AW phase encodes which were hard coded, potentially leading to unforeseen differences in the pulse sequence.

To test if these concerns were having a measurable impact on the results we extracted a one average $4 \times 4 \times 4$ phase encode data-set from both the AW and full $4 \times 4 \times 4$ data-sets and analyzed them with the SLAM and SLIM reconstruction pipelines (Section 3.5.2.3). These data-sets were termed 'reduced AW' or 'reduced $4 \times 4 \times 4$ ' respectively. Statistical tests (paired Wilcoxon signed-rank, two tail, p=5%) were used to compare the SLAM and SLIM reconstructions of the two acquisitions. Inspection of the reconstructed spectra indicated that the reduced averaging (the effective scan time was 64 seconds compared to 6.5 minutes for the full scans) caused a large dip in SNR (expected to be 2.5× lower than the full $4 \times 4 \times 4$ SNR) that meant a number of spectra were un-analyzable. We therefore implemented a threshold where spectra with a PCr CRLB of $> 20\%$ were excluded from analysis so that meaningful comparisons could be made between the data-sets.

The results are shown in Table 3.5 and the PCr/ATP ratio and PCr SNR are plotted in Figure 3.16. There was no statistical difference in the PCr/ATP ratio or PCr SNR between 'reduced AW SLAM' and 'reduced $4 \times 4 \times 4$ SLAM' or between 'reduced AW SLIM' and 'reduced $4 \times 4 \times 4$ SLIM', indicating that there was not a significant difference between the two acquisitions. Up to 10/22 (Reduced AW SLAM) of the spectra did have to be removed for excessively high PCr CRLB,

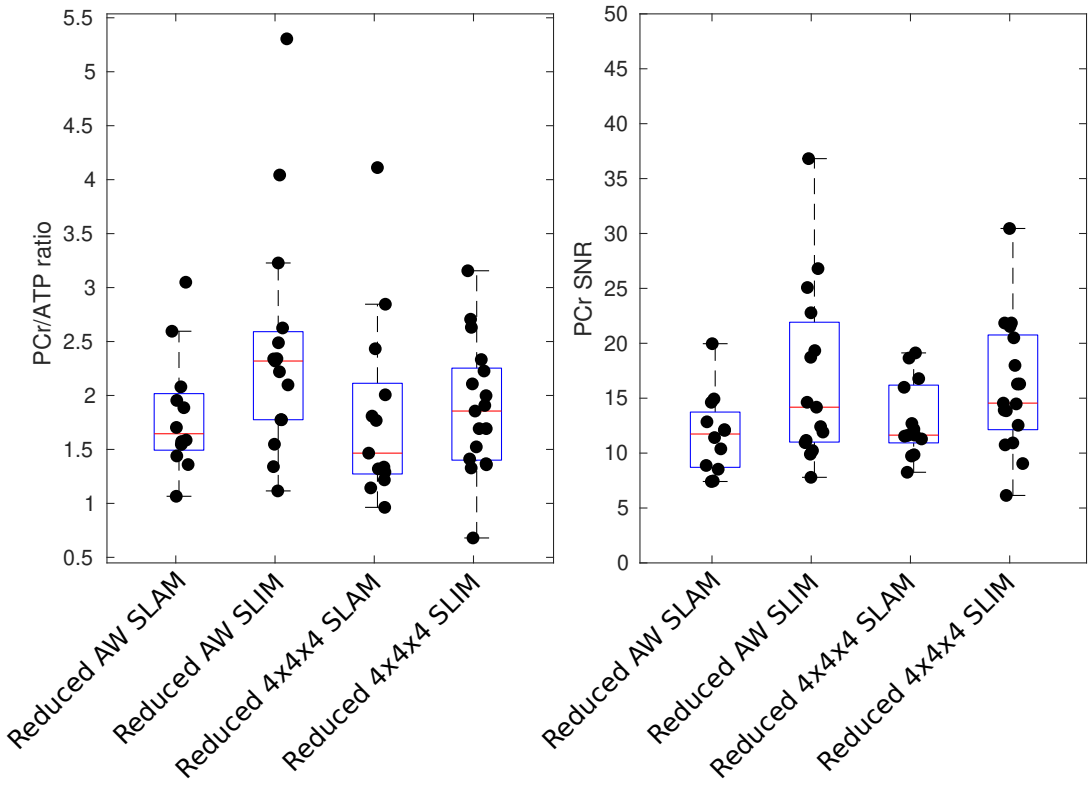

Figure 3.16: Box plot of PCr/ATP ratio and PCR SNR for the reduced data-sets comprising of one average of $4 \times 4 \times 4$ phase encodes. Data is thresholded by PCR CRLB <20%.

indicating that it would not be possible to reduce the number of averages to one and retain a clinically useful technique. However the remaining spectra had comparable CoR and CoV to the full repeats reconstructions, potentially suggesting that once a spectra has above a certain SNR, other factors (e.g. physiological variation) are the driving source for the spread in the data.

The agreement between the data-sets generated by two independently implemented sequences provides good confidence that the acquisitions were implemented and performed correctly, and that we can therefore draw valid comparisons between them.

3.6.3 SLAM* method for solving the SLAM equation

In 2013 Zhang et al. significantly extended the SLAM method by incorporating parallel imaging and an alternative method (SLAM*) for solving the SLAM

equation which can result in greater tolerance to field in-homogeneity[111]. The SLAM* equation is

$$\rho_{C\times N} = B' \times (G_{P\times M})^{+} \times S_{P\times N} \tag{3.34}$$

$B'_{C\times M}$ is a matrix whose columns represent the x coordinates $x_1...x_M$ and rows compartments, if x_i is within compartment j then $B'_{ji} = 1$ otherwise $B'_{ji} = 0$. $G_{P\times M}$ is a subset of $G_{M\times M}$ with only the rows representing the P acquired phase encodes acquired. The success of this equation will depend on the stability of the pseudo-inverse of the $G_{P\times M}$ matrix, which may be of concern for the AW data-set whose $G_{P\times M}$ matrix presents an ill-posed problem.

To investigate the utility of the SLAM* algorithm for our data-sets we substituted the SLAM* reconstruction into our regular reconstruction pipeline (Section 3.5.2.3) and used it to reconstruct each of our acquisition data-sets. The MATLAB *pinv* function, which implements the Moore-Penrose pseudo-inverse using a singular value decomposition, was used to find the inverse of $G_{P\times M}$ and the resultant spectra were analyzed as in Section 3.5.2.3 using the OXSA toolbox.

The results of using the SLAM* algorithm are shown in Table 3.6 alongside the equivalent SLAM reconstruction. The results of using the two reconstruction algorithms are broadly similar, with two notable exceptions. Firstly $4 \times 4 \times 4$ SLAM* had a significantly higher PCr/ATP ratio than AW CSI unlike $4 \times 4 \times 4$ SLAM, and secondly AW SLAM* did not have a significantly higher PCr SNR than AW CSI, unlike AW SLAM.

These differences make the original SLAM algorithm preferable overall for our data-set, particularly for the AW data-set, where the SNR advantage over AW CSI is almost entirely removed by using the SLAM* algorithm. However future work to improve the accuracy of inverting the $G_{P\times M}$ matrix, with regularization and iterative methods, may yield significantly better results with the SLAM* algorithm than were realized here.

Technique	PCr/ATP Ratio	PCr SNR	PCr CRLB	CoR	CoV
AW CSI	1.90 ±0.92	13.10 ±8.63	13.36±7.37	1.79	0.48
AW SLAM*	1.65 ± 0.50	14.04±7.77	12.65 ± 8.64	1.07	0.30
AW SLAM	1.68 ±0.50	20.75 ±11.18*	8.08±4.19	1.22	0.30
$4 \times 4 \times 4$ SLAM*	2.18± 0.86*	28.84 ± 12.72*	6.35 ± 3.67	2.60	0.39
$4 \times 4 \times 4$ SLAM	1.90 ±1.13	25.43 ±11.44*	7.60±4.20	3.29	0.59
fSLAM*	2.26 ± 0.86	26.74±14.95*	8.19± 6.72	2.50	0.39
fSLAM	2.02 ±0.88	25.05 ±13.29*	9.54±10.75	2.43	0.43

Table 3.6: Mean values for SLAM* and SLAM reconstructions with the AW CSI reconstruction included for context, errors are given as ± 1 SD. * indicates PCr/ATP ratio or PCr SNR significantly different to the AW CSI method. The horizontal lines separate each of the acquisition methods.

3.7 Conclusions

In this chapter, we first showed the feasibility of using compartmentalized spectroscopy reconstruction techniques at 7T, with simulations and phantom experiments, confirming that SLAM and SLIM reconstructions were able to effectively separate out signals from different compartments. We then implemented compartmentalized spectroscopic reconstructions for 3D ^{31}P cardiac spectroscopy at 7T, and completed a repeatability study to assess these techniques against our center's current gold standard AW CSI method. Before finally re-validating our acquisitions and examining an extension to the technique using an alternative SLAM reconstruction algorithm.

The main finding of this chapter is that it is possible to achieve an increase in SNR for ^{31}P cardiac MRSI using compartmentalized spectroscopy, without significantly affecting the measured PCr/ATP ratio and improving both the coefficient of variation and coefficient of repeatability. We found that we were able to achieve this using the SLAM algorithm to reconstruct data generated with our existing acquisition weighted ^{31}P acquisition, making the roll out of this technique to clinical studies relatively simple. We also found that there was no benefit to using our implementations of the alternative $4 \times 4 \times 4$ and fSLAM phase encode schemes compared to the acquisition weighted scheme, or using the SLAM* algorithm.

4

Compartmentalised ^{31}P spectroscopy on clinical systems at 3T

Contents

4.1 Introduction

In Chapter 3 we demonstrated the feasibility of using compartmentalized reconstruction methods for cardiac ^{31}P spectroscopy at 7T. These methods, in particular SLOOP and SLAM, have been used in a number of clinical studies for ^{31}P cardiac spectroscopy at 1.5[107, 108, 122, 123] and 3T[112], however the techniques are not currently in widespread use, due to the perceived added complexity of adding the required acquisition, which may need optimizing on a per-patient basis, to clinical research protocols.

The results presented in Chapter 3, showed that AW SLAM improved both SNR and repeatability compared to AW CSI at 7T, which suggests that similar improvements may be possible at 3T. This AW acquisition is a commonly used method for 3T clinical research, making it possible to implement the AW compartmentalized reconstruction methods, while only making minor changes to the clinical research procedure, and fully retaining the ability to use the existing, Fourier based, CSI reconstruction method. At 3T, SNR is lower than at 7T, making the potential SNR increase from using compartmentalized reconstruction techniques at this field strength even more valuable. Furthermore B_1 homogeneity is typically improved at 3T compared to 7T[113], which should translate into improved localization by reducing partial volume effects.

An additional motivation to implement compartmentalized reconstruction at 3T has been highlighted by clinicians who currently use the CSI method for clinical research. The large PCr signal produced by the chest wall means that small shifts in the midseptal voxel position could significantly alter the measured PCr/ATP ratio. While this is an unwanted source of variability for the measurement, it is also a potential source of bias, if a researcher were to sub-consciously shift the voxel position based on the desired result. In contrast, the compartmentalization masks for the SLAM and SLIM reconstructions are clearly defined by anatomical regions, making bias less likely, and provide the potential to remove human interaction all together by using automated segmentation algorithms[124].

In this chapter, the use of compartmentalized reconstruction ^{31}P cardiac spectroscopy at 3T for clinical research is validated by comparing it to our existing AW CSI technique (Chapter 3) with a healthy volunteer repeatability study, and reconstructions of clinical research data. I also evaluate the effect of a number of experimental factors on the repeatability with two preliminary studies, and as part of the main study. The factors I considered with the preliminary studies were the acquisition phase encoding scheme (AW $16\times8\times8$ or uniform $4\times4\times4$), and the use of saturation bands. Cardiac gating, the number and composition of the compartments used in the segmentation model, and reconstruction algorithm (SLAM or SLIM) were then addressed as a part of the main repeatability study, since the choice of cardiac gating, and therefore reconstruction model/algorithm may depend on clinical considerations. Additionally, I also evaluated the effect of errors when setting up the ^{31}P scan, by shifting the main study data-set's FOV, and observing the change in measured PCr/ATP ratio. As in Chapter 3, PCr/ATP ratio, PCr SNR, PCr/ATP ratio Cramer Rao lower bounds, and coefficients of repeatability (CoR) and variation (CoV) were calculated, where appropriate, to allow quantitative comparison.

4.2 Methods

4.2.1 Chapter overview

The experiments reported in this chapter consist of two preliminary studies, a main repeatability study, and reconstruction of clinical research data. In both preliminary studies, three healthy volunteers (3M) were scanned twice with a short break outside of the scanner between scans. In the first preliminary study, each of the scans consisted of two AW ^{31}P acquisitions, with and without saturation bands. In the second preliminary study, each of the scans consisted of two ^{31}P acquisitions, with AW and $4 \times 4 \times 4$ phase encoding schemes. For the main repeatability study, 6 volunteers (2F/4M) were scanned twice, this time with saturation bands and AW phase encoding, based on the previous results. In this case, each scan consisted of two acquisitions, with and without cardiac gating (denoted G and UG respectively). An overview of the three protocols is shown in Figure 4.1 and the phase encodes used

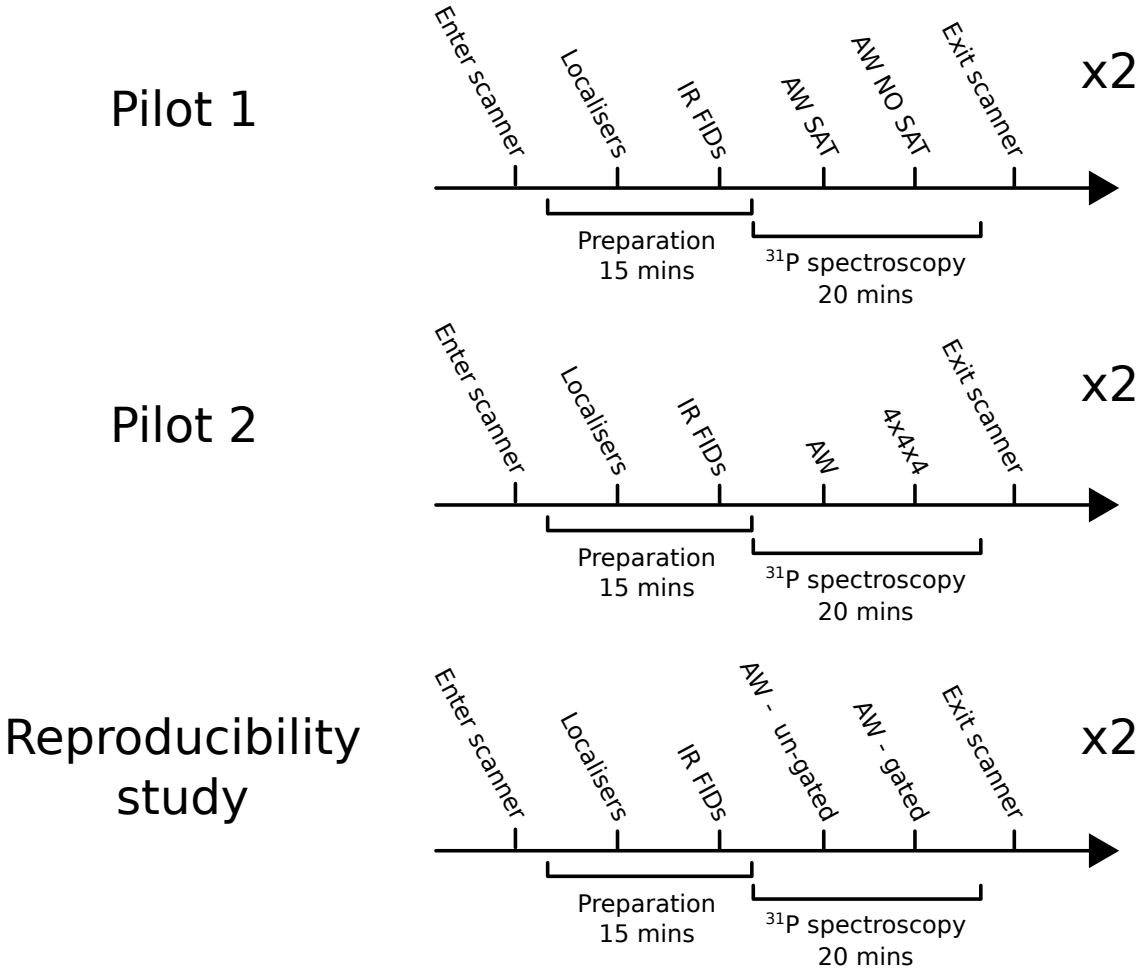

Figure 4.1: Overview of the protocol for each of the preliminary studies and main repeatability study.

for the AW and $4 \times 4 \times 4$ acquisitions are shown in Figure 4.2. The clinical data was acquired with the AW phase encoding scheme without cardiac gating. Five patients (2F/4M, BMI $= 36 \pm 5$ kg/m^2, age $= 69 \pm 9$ years) with heart failure with preserved ejection fraction (HFpEF) were recruited and scanned, by an OCMR clinical fellow.

4.2.2 Data acquisition

All acquisitions were performed using a Siemens TIM Trio whole body 3T MRI scanner (Siemens Healthineers, Erlangen, Germany) with volunteers positioned in the prone position. The coil consisted of a ^{1}H/^{31}P dual tuned 26×28 cm transmit/receive loop and a dual 12×15 cm butterfly ^{31}P receive pair[57]. Two-chamber, four-chamber, short axis (FOV $= 400 \times 340$ mm) and CSI-matched ^{1}H

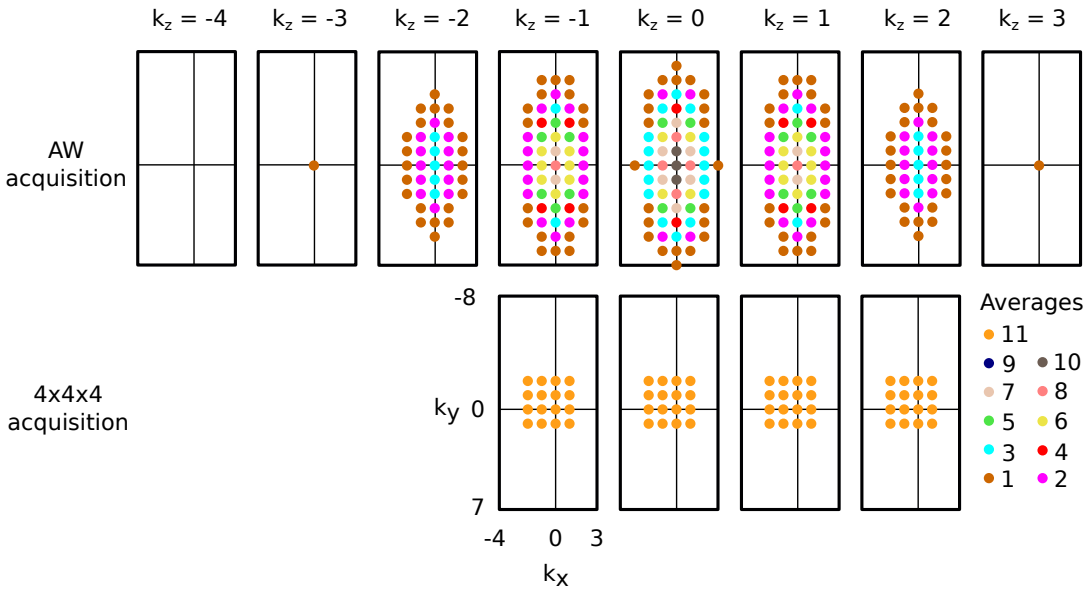

Figure 4.2: k-space coordinates of the phase encodes used in the 3T AW and 4x4x4 ^{31}P acquisitions investigated in this chapter, all k-space coordinates are normalized by the FOV to give integer values.

images were acquired with a fast low-angle shot (FLASH) sequence[115]. The coil position was determined from images showing the location of 4 cod-liver oil capsules and one central phenylphosphonic acid fiducial embedded in the coil housing during imaging and the transmit flip-angle was calculated from a series of inversion-recovery free induction decay acquisitions of the central fiducial as previously described[103].

All ^{31}P acquisitions consisted of a UTE-CSI sequence[116] with a shaped RF-pulse, numerically optimized for excitation homogeneity[57] and NOE enhancement. The readout had a bandwidth of 4000 Hz and was centered at -250 Hz from the PCr resonance, i.e. between the γ-ATP and α-ATP resonances. Cardiac gating (where used) was prospective and triggered the acquisition to coincide with diastole. Unless otherwise stated, 3 saturation bands were used to null the signal from the chest wall and liver (Figure 4.3). All ^{31}P acquisitions had a FOV of 240 × 240 × 200 mm and, in the absence of cardiac gating, took 10.5 minutes to complete with a nominal TR of 0.9 s. Cardiac gating lead to a variable acquisition time, dependent on heart rate; where the R-R interval was less than the nominal TR, i.e. the heart rate was above 67 bpm, the acquisition time was increased by up to 2×.

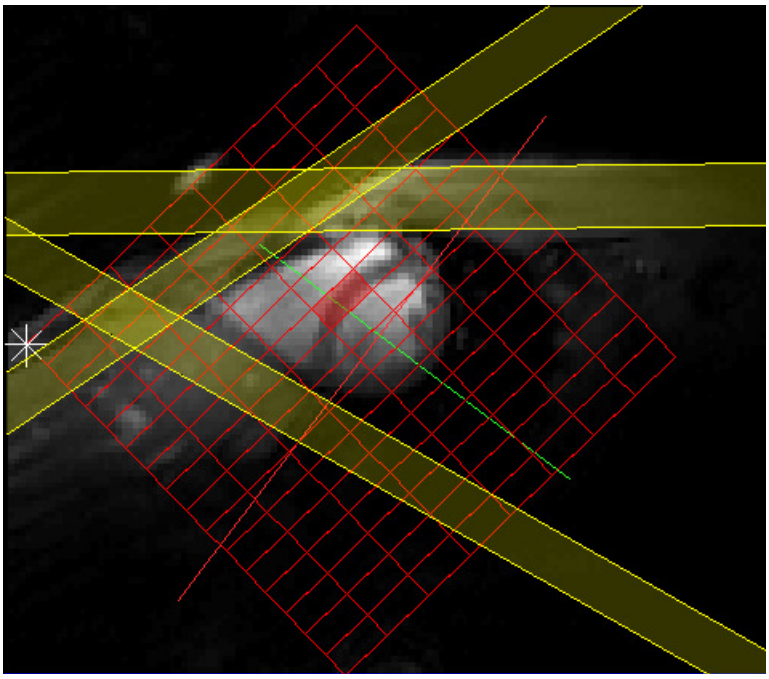

Figure 4.3: Short axis localizer showing the CSI grid in red and 3 PCr saturation bands in yellow, two of the bands saturate the PCr signal from the chest and the third saturates the liver close to the heart.

4.2.3 Data Analysis

The SLAM, SLIM and CSI algorithms and data analysis were implemented in the same way as in Section 3.5.2.3. The two preliminary studies used a three compartment model consisting of heart, chest wall and other. The repeatability study used three different compartmentalization models, the three compartment model and two four compartment models consisting of heart, chest wall, blood, and other, where the blood compartment was defined as either the left ventricle lumen (LV) or the left + right ventricle lumens (LV + RV). An example compartmentalization mask showing the different models is shown in Figure 4.4 with an example fitted cardiac spectra generated using the three compartment model.

Cramer Rao lower bounds (CRLB) were calculated for the PCr/ATP ratio, as is usual at 3T and CoR and CoV were calculated using the definitions in Section 3.5.2.3. All comparisons were made to the reference AW CSI method with pairwise two-tailed Wilcoxon signed rank tests with 5% significance level.

As an addition to the repeatability study, the FOV of the ^{31}P data-sets were

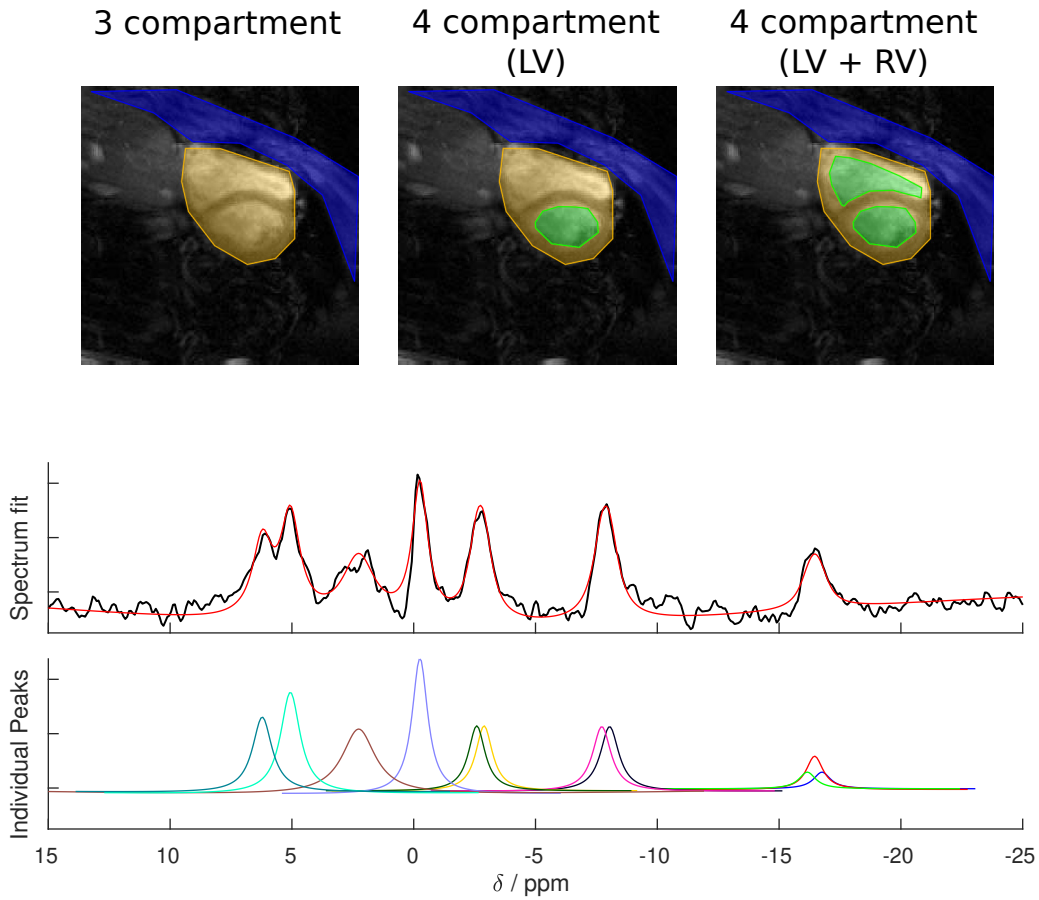

Figure 4.4: Top: Example segmentation masks with chest wall shown in blue, heart in yellow, and segmented ventricle lumen(s) in green. **Bottom:** Example cardiac spectra for a healthy volunteer reconstructed with the 3 compartment (whole heart, chest wall, other) model and the AMARES fitted peaks.

shifted towards and away from the chest wall to simulate the effect of incorrect segmentation/voxel placement and reconstructed using the SLAM, SLIM and CSI methods (see Figure 4.5). The phase shift was effected by applying a phase roll to the ^{31}P data, in the direction of the shift. For the SLAM reconstructions two scenarios were investigated: firstly (SLAM) where the localizer used for the segmentation was left unchanged, which would simulate incorrect segmentation, which may be due to patient movement in the scanner between localizer and ^{31}P acquisition. Secondly (SLAM compensated) where the localizer's spatial position was adjusted to match the new ^{31}P FOV, to investigate if the position of the FOV affects the result, given that the segmentation is performed correctly.

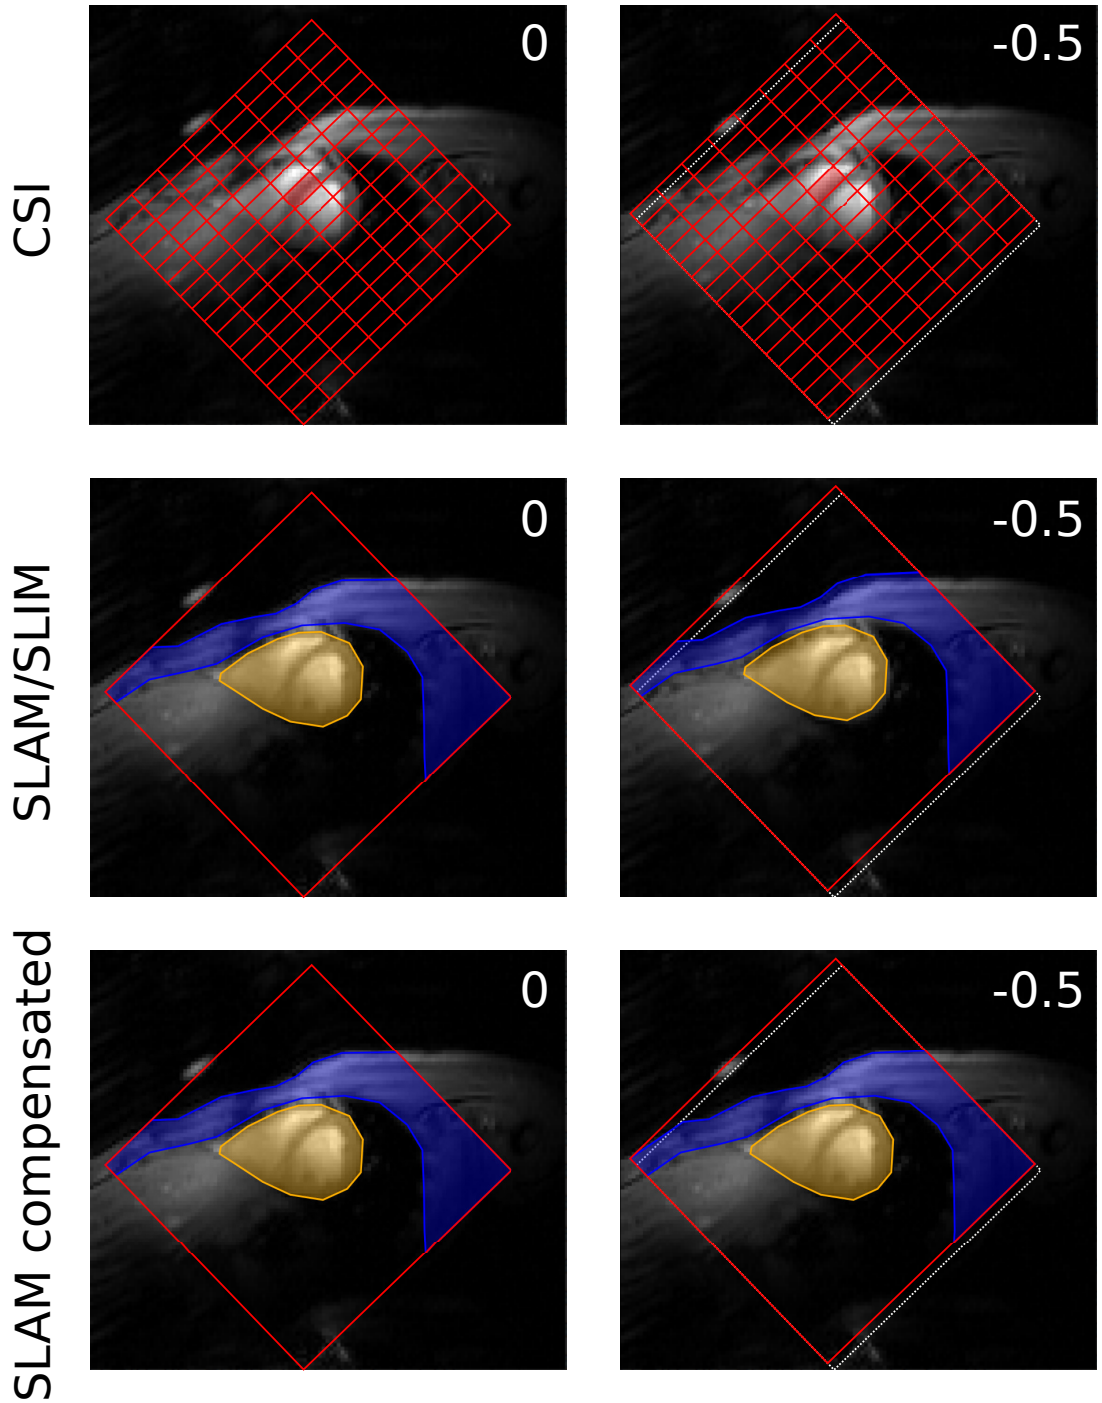

Figure 4.5: Example FOV shifts as applied in the main repeatability study. The left column shows the original FOV position (0 voxel shift) and the right column shows the FOV shifted -0.5 voxels (towards the chest wall). In the case of the CSI and SLAM/SLIM methods this results in the midseptal voxel/heart compartment moving closer to the chest wall. For SLAM compensated, the segmentation mask is updated to reflect the FOV shift applied. The red box/grid represents the FOV/CSI grid, and the white dotted line represents the unshifted FOV. The compartments are highlighted as in Figure 4.4 and the midseptal voxel is highlighted in red.

Technique	Saturation	PCr/ATP	PCr SNR	CRLB (%)	CoR	CoV
AW SLAM	Yes	1.91 ± 0.20	35.7 ± 4.4	11.4 ± 1.8	0.260	0.106
AW SLAM	No	2.13 ± 0.25*	45.4 ± 11.5	11.0 ± 2.3	0.430	0.118
AW SLIM	Yes	2.08 ± 0.21	41.6 ± 4.4	11.6 ± 1.7	0.174	0.100
AW SLIM	No	2.35 ± 0.24*	53.8 ± 13.0	11.3 ± 2.1	0.481	0.104
AW CSI	Yes	2.13 ± 0.21	31.9 ± 11.9	17.3 ± 3.0	0.455	0.100
AW CSI	No	2.41 ± 0.21*	38.3 ± 13.6	19.0 ± 5.0	0.323	0.088

Table 4.1: Mean values of the experimental results derived for each technique with and without saturation bands, errors are given as ± 1 SD. * indicates significant difference in value when saturation bands are disabled. All acquisitions were AW.

4.3 Results

4.3.1 Preliminary study 1 – saturation bands

The computed results of the first preliminary study are shown in Figure 4.6 and Table 4.1. In all cases SNR was adequate for good fits to the spectra. For each reconstruction method turning off the saturation bands significantly increased the PCr/ATP ratio, the Wilcoxon signed-rank p-value was 0.0312 in all cases because the signed difference of all the paired data was positive. For the SLAM reconstruction turning off the saturation bands resulted in a 12% increase in PCr/ATP ratio and for SLIM and CSI the increase was 13%.

4.3.2 Preliminary study 2 – acquisition phase encode scheme

The results of the second preliminary study are shown in Figure 4.7 and Table 4.2. The AW SLAM and SLIM methods had significantly higher PCr SNR than the AW CSI method (p=0.0321), however the $4 \times 4 \times 4$ SLAM and SLIM did not.

Phase encodes	Algorithm	PCr/ATP	PCr SNR	CRLB (%)	CoR	CoV
AW	SLAM	1.75 ± 0.26	18.0 ± 4.9*	19.3 ± 9.7	0.431	0.151
AW	SLIM	1.80 ± 0.24	20.4 ± 7.3*	19.0 ± 9.9	0.727	0.136
$4 \times 4 \times 4$	SLAM	1.73 ± 0.42	18.3 ± 7.8	23.7 ± 13.2	1.268	0.242
$4 \times 4 \times 4$	SLIM	1.64 ± 0.38	18.1 ± 6.9	22.6 ± 7.4	0.306	0.231
AW	CSI	1.73 ± 0.12	13.3 ± 6.6	25.7 ± 10.0	0.488	0.071

Table 4.2: Mean values of the experimental results derived for each acquisition/reconstruction combination, errors are given as ± 1 SD. * indicates significant difference in PCr SNR to the AW CSI method.

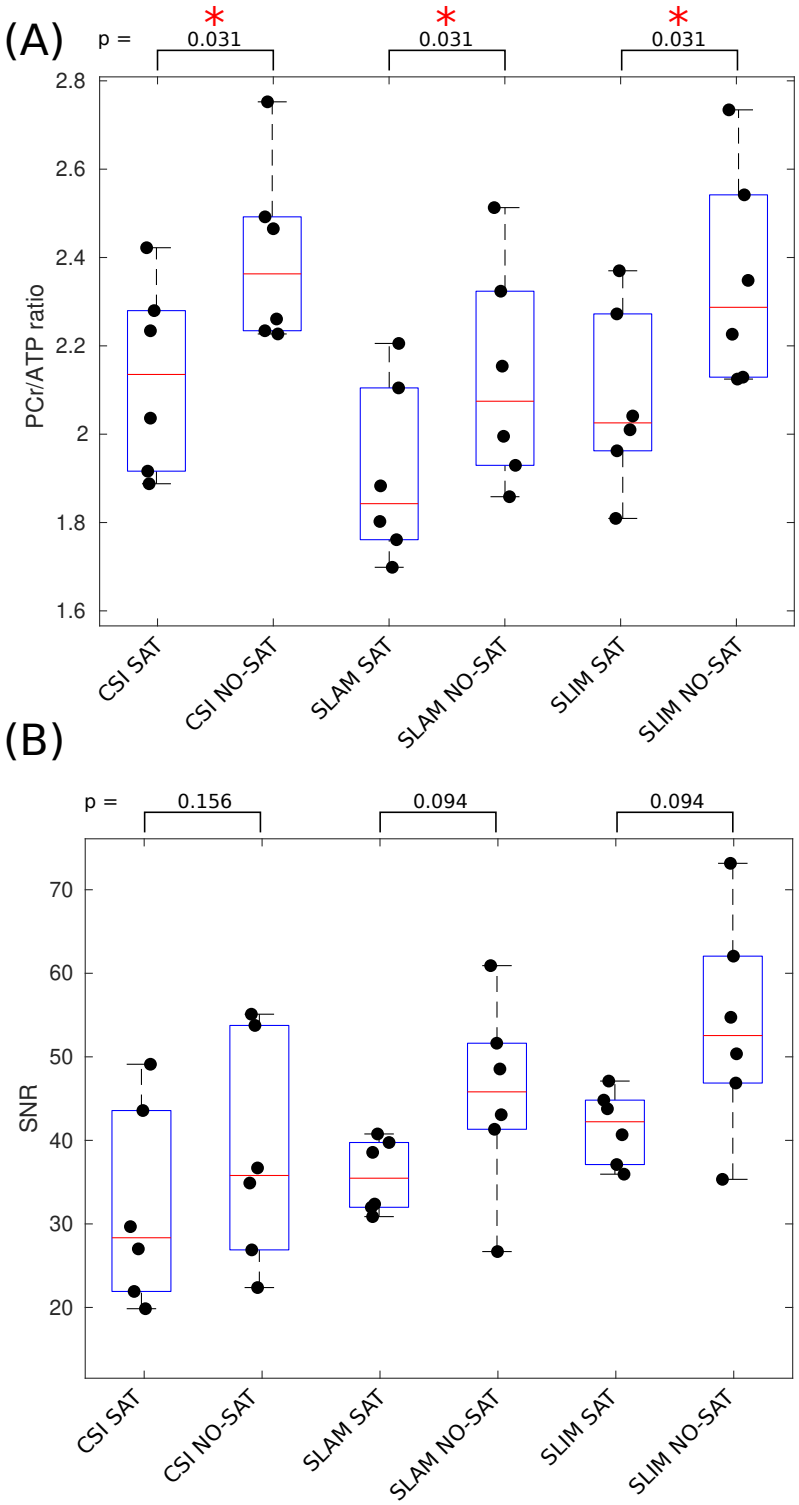

Figure 4.6: Box-plots showing PCr/ATP ratio (A) and PCr SNR (B) for each technique with and without saturation bands, median and IQR indicated by box. * indicates significant difference (Wilcoxon signed-rank paired, α=0.05). p=0.03 for all 3 statistical tests. All acquisitions were AW. NO SAT indicates saturation bands were disabled.

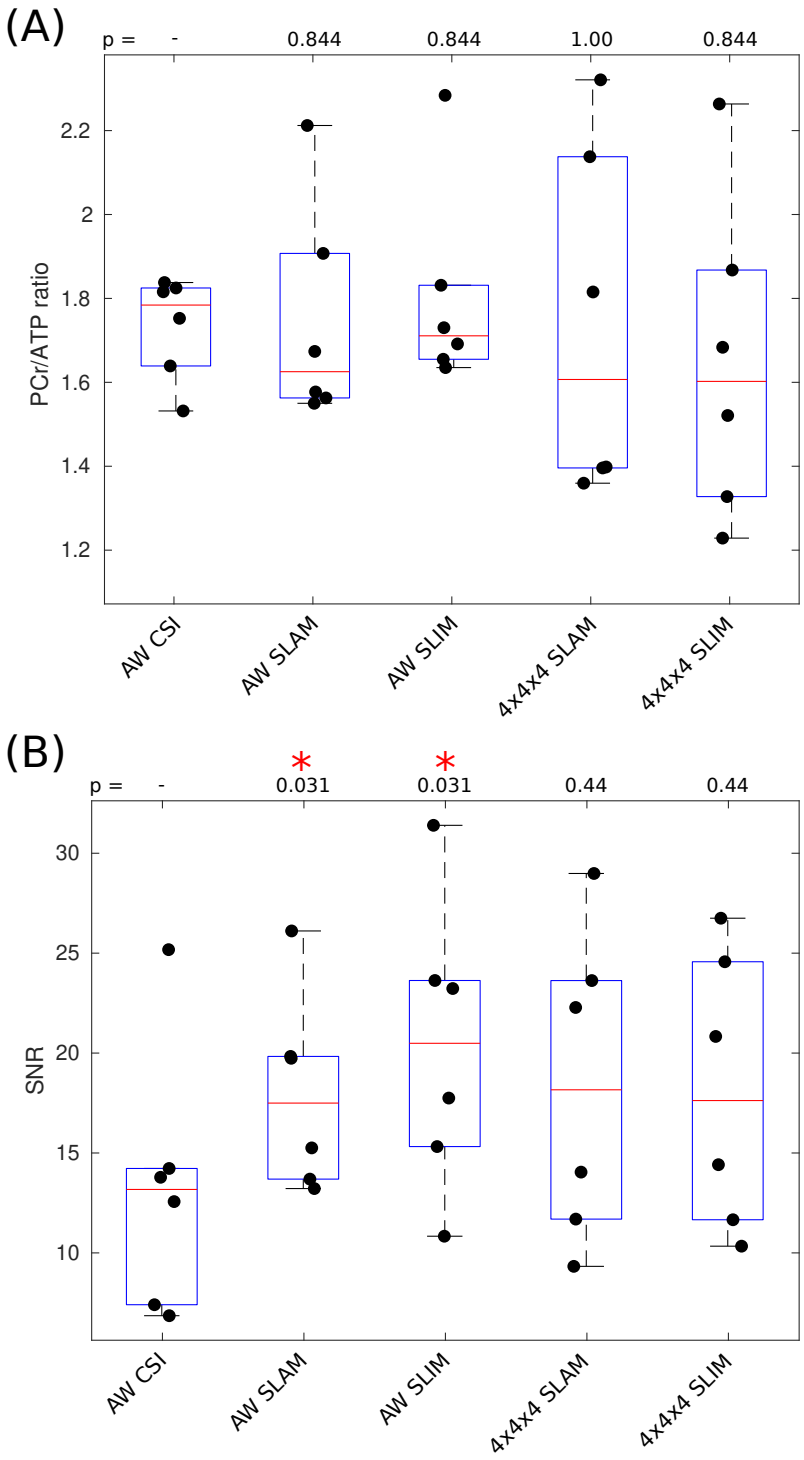

Figure 4.7: Box-plots showing PCr/ATP ratio (A) and PCr SNR (B) for each acquisition/reconstruction combination, median and IQR indicated by box. * indicates significant difference in PCr SNR to the AW CSI method.

4.3.3 Main study – repeatability

The results for this section are presented first for the three compartment model (heart, chest wall and other) with and without cardiac gating, then for the four compartment model where either the left, or left + right ventricle lumens were segmented out as an additional blood compartment. All of the data for the three and four compartment models is shown in Table 4.3.

4.3.3.1 Gated AW CSI

The results of the gated AW CSI method are shown in Figure 4.8, with the computed statistics shown as a part of Table 4.3, Correlation plots comparing the methods to the un-gated AW CSI method are shown in Figure 4.9. The gated AW CSI method did not have significantly different PCr/ATP ratio ($p = 0.850$), PCr SNR ($p = 0.470$) or CRLB ($p = 0.569$) to the un-gated AW CSI method, however both the CoV and CoR were lower, indicating improved repeatability.

4.3.3.2 Three compartment model

Box plots displaying the computed PCr/ATP ratio and PCr SNR for SLAM, SLIM and CSI reconstructions of the the gated (G) and ungated (UG) AW acquisitions are shown in Figure 4.8. The computed statistics for each method are shown as part of Table 4.3 and correlation plots comparing each method to the un-gated AW CSI method are shown in Figure 4.9. Both gated and un-gated AW SLAM and SLIM had significantly lower CRLBs compared to the un-gated AW CSI (G SLAM $p = 0.0010$, UG SLAM $p = 0.0005$, G SLIM $p = 0.0010$, UG SLIM $p = 0.0005$). The gated and un-gated AW SLIM also had significantly higher PCr SNR than the un-gated AW CSI (G SLIM $p = 0.012$, UG SLIM $p = 0.043$). The un-gated AW SLAM had a significantly lower PCr/ATP ratio than the un-gated AW CSI ($p = 0.021$). All of the three compartment SLAM and SLIM reconstructions had lower CoRs and CoVs than the un-gated AW CSI, similar to the gated AW CSI. The reduction in CoV over the un-gated AW CSI is reflected in the trends shown

in the correlation plots in Figure 4.9, where points with high and low PCr/ATP ratios are returned to more normal values.

4.3.3.3 Four compartment model

Segmenting out the left ventricle as a fourth blood compartment resulted in a PCr/ATP ratio not significantly different to the un-gated AW CSI. Additionally both the gated and un-gated acquisitions gave a significantly higher PCr SNR than the un-gated AW CSI (G SLAM p = 0.016, UG SLAM p = 0.012, G SLIM p = 0.009, UG SLIM p = 0.003), although the gated reconstructions did not give a significantly lower PCr/ATP CRLB (G SLAM p = 0.064, UG SLAM p = 0.001, G SLIM p = 0.110, UG SLIM p = 0.002). For both SLAM and SLIM reconstructions using the un-gated acquisition with the LV segmented, resulted in a worse CoR than un-gated AW CSI, although the CoV was better. Using the gated acquisition however resulted in an improvement in both CoR and CoV over the un-gated AW CSI, and the respective 3 compartment SLAM and SLIM reconstructions.

Including the left and right ventricles in the fourth compartment of the SLAM and SLIM reconstructions resulted in PCr SNR not significantly different to the reference un-gated AW CSI. The SLIM reconstruction's PCr/ATP CRLBs were significantly worse than un-gated AW CSI. All the CoVs were higher than un-gated AW CSI, and the gated CoRs were too.

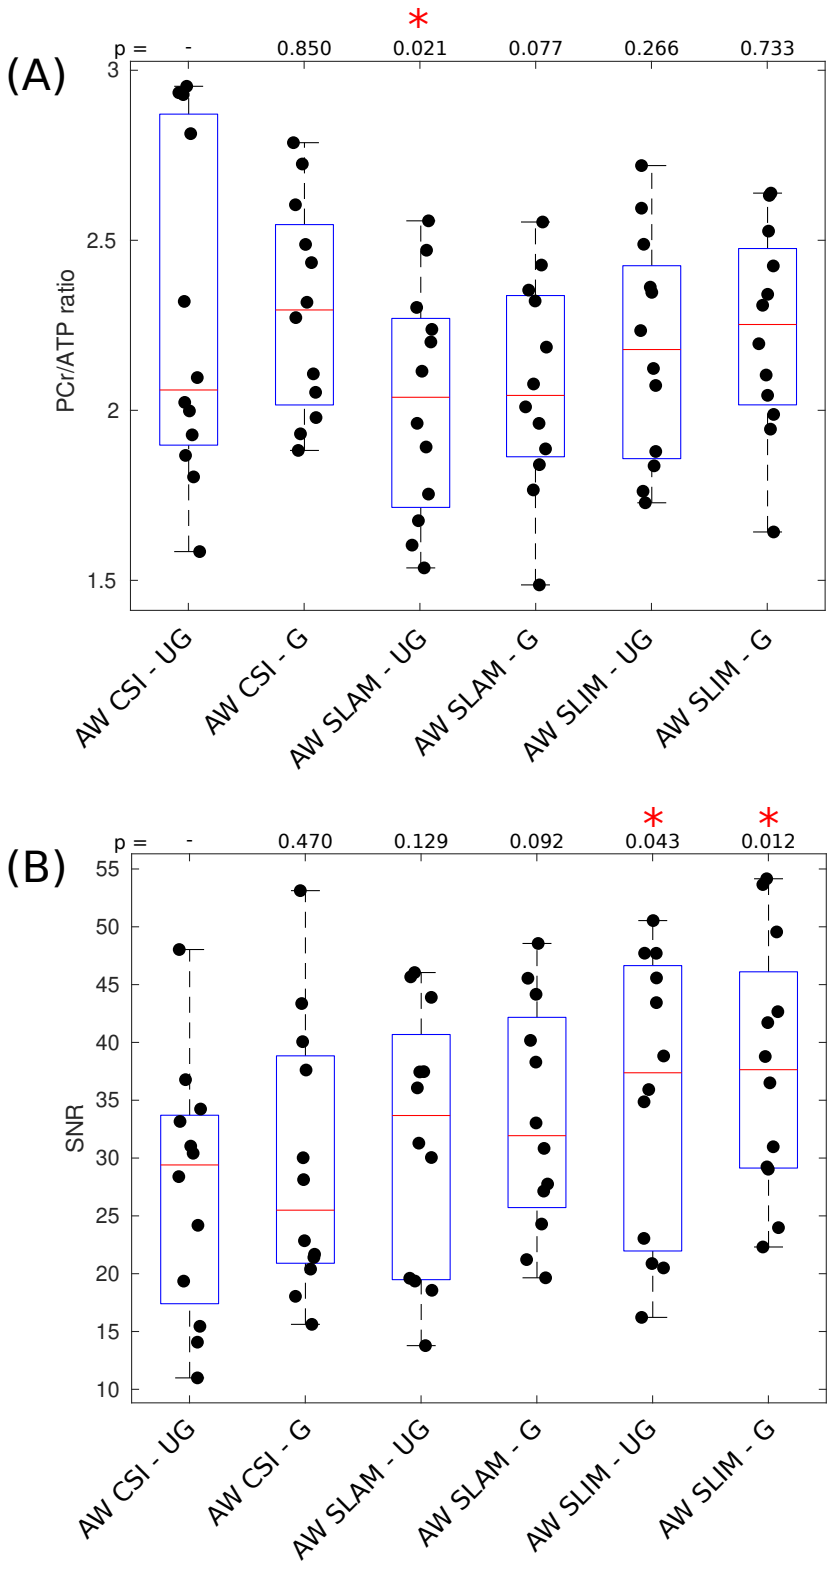

Figure 4.8: Box-plot showing PCr/ATP values (A) and PCR SNR (B) for each acquisition, median and IQR are indicated by box. * indicates significant difference to the un-gated AW CSI method (Wilcoxon signed-rank paired, α=0.05). PCr/ATP values are in the expected range (Table 3.4)

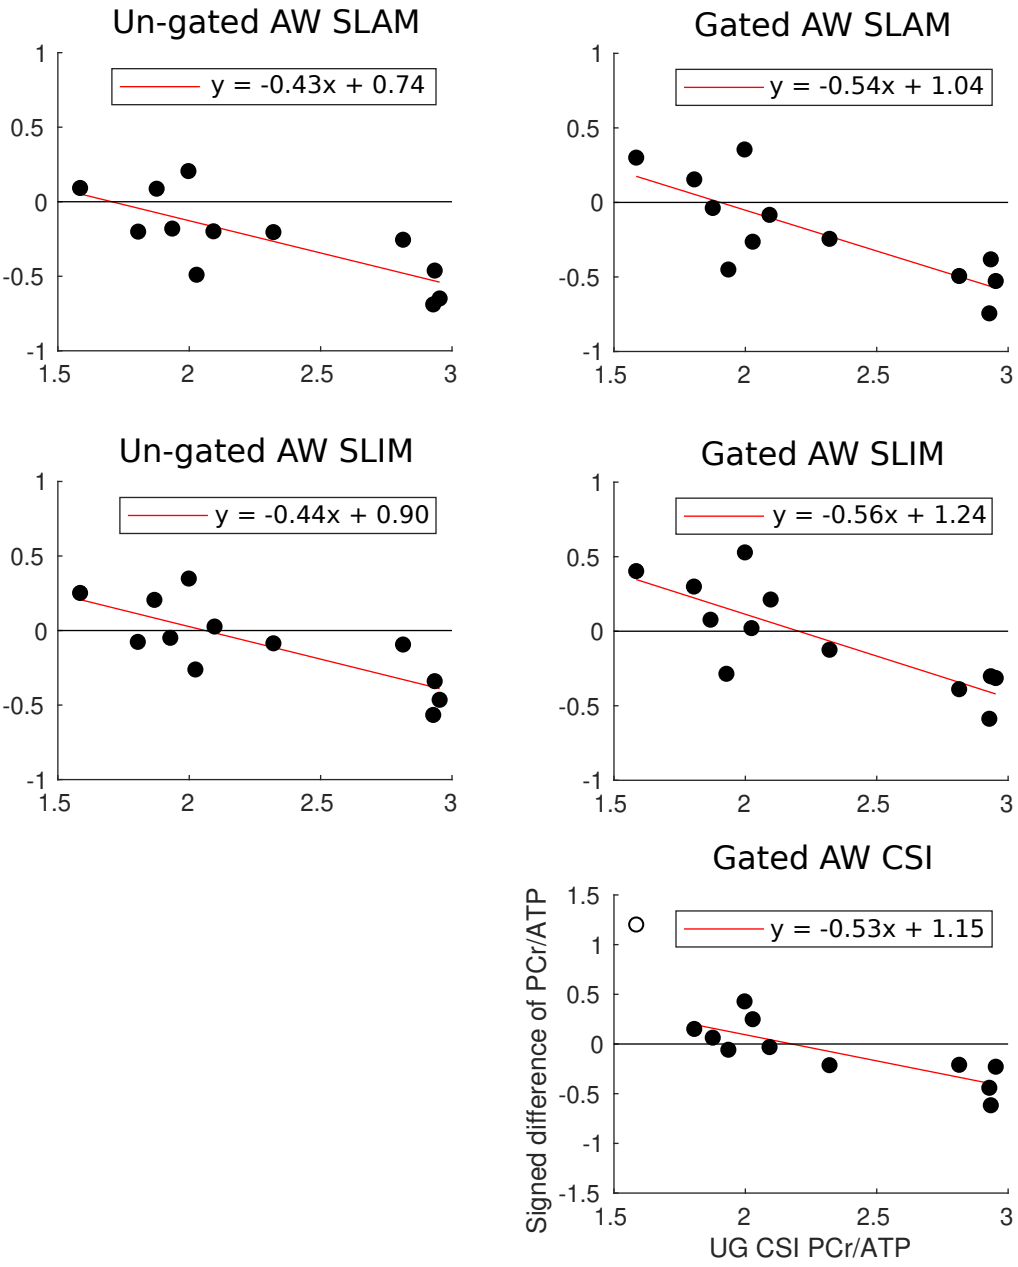

Figure 4.9: Correlation plots comparing each method to the un-gated AW CSI method. y-axis: signed difference in the PCr/ATP ratio of each method and the un-gated AW CSI, x-axis: PCr/ATP ratio of the un-gated AW CSI. An empty circle indicates the data-point was excluded from calculation of line of best fit. A negative gradient corresponds to a reduction in variation compared to the un-gated AW CSI. The gradient of the line is approximately equal for the gated/un-gated methods respectively.

Phase encodes	Algorithm	Lumen	Gating	PCr/ATP	PCr SNR	CRLB (%)	CoR	CoV
AW	CSI	-	UG	2.27 ± 0.50	27.2 ± 10.8	11.6 ± 6.8	1.02	0.22
		-	G	2.30 ± 0.31	29.4 ± 11.7	11.7 ± 8.6	0.77	0.14
AW	SLAM	-	UG	2.03 ± 0.34*	31.6 ± 11.4	7.81 ± 3.29*	0.66	0.17
		-	G	2.07 ± 0.31	33.4 ± 9.8	7.92 ± 3.56*	0.67	0.15
		LV	UG	2.16 ± 0.44	38.9 ± 16.8*	9.37 ± 5.84*	1.20	0.20
		LV	G	2.13 ± 0.29	39.4 ± 18.4*	9.46 ± 5.79	0.61	0.13
		LV + RV	UG	2.20 ± 0.55	25.1 ± 13.7	12.6 ± 7.5	0.89	0.25
		LV + RV	G	2.20 ± 0.52	28.5 ± 17.4	12.1 ± 6.1	1.54	0.24
AW	SLIM	-	UG	2.18 ± 0.33	35.4 ± 12.3*	7.28 ± 3.20*	0.65	0.15
		-	G	2.23 ± 0.30	37.7 ± 11.0*	7.38 ± 3.37*	0.73	0.14
		LV	UG	2.29 ± 0.46	43.1 ± 18.6*	9.32 ± 6.53*	1.21	0.20
		LV	G	2.26 ± 0.29	42.6 ± 19.5*	9.45 ± 6.38	0.49	0.13
		LV + RV	UG	2.23 ± 0.64	23.4 ± 11.8	14.1 ± 9.1*	0.56	0.29
		LV + RV	G	2.22 ± 0.75	24.5 ± 13.9	15.4 ± 11.0*	1.07	0.34

Table 4.3: Mean values of the experimental results derived for each method, errors are given as ± 1 SD. * indicates significant difference to the un-gated AW CSI method (Wilcoxon signed-rank paired, α=0.05). The lumen column indicates the anatomical regions (LV = left ventricle, RV = right ventricle, see Figure 4.4) included in the lumen compartment, - indicates that the lumen compartment was not used. Where the cardiac compartment did not contain blood no blood correction was used. UG indicates the acquisition was un-gated and G that the acquisition was cardiac gated, with the readout timed to coincide with diastole.

4.3.4 FOV shift

The effect of shifting the FOV of the ^{31}P acquisition FOV on the reconstructed PCr/ATP values with each method is shown in Figure 4.10. The change in mean PCr/ATP ratio across the two voxel shift range is 0.82 for CSI, 0.46 for SLAM, and 0.43 for SLIM. Where the SLAM segmentation mask is compensated for the movement in FOV the change in the PCr/ATP values is minimal.

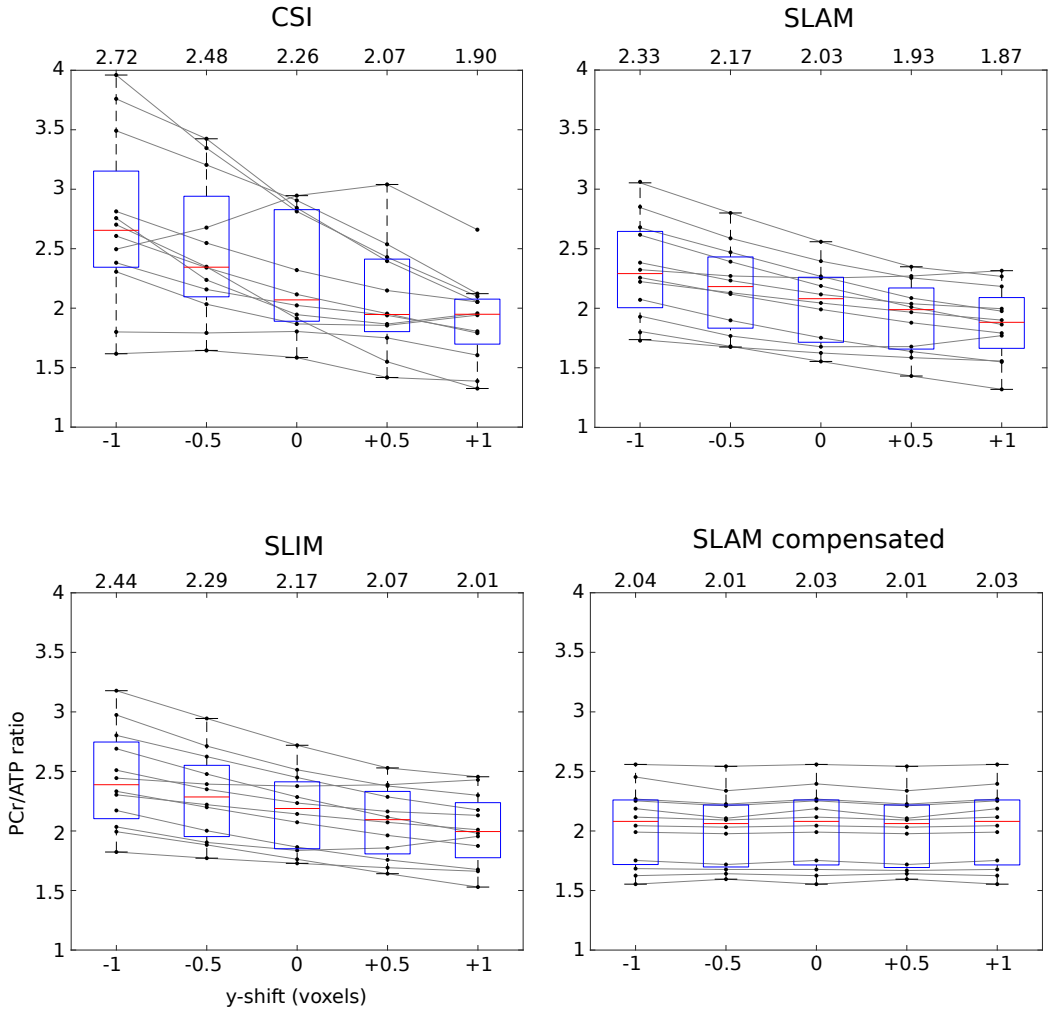

Figure 4.10: Each pane shows how the reconstructed PCr/ATP ratios of the repeatability data-set for each method varies as the data is shifted in the y-direction (up-down on Figure 4.4). The mean PCr/ATP of each shifted data-set is shown above the corresponding box and the points corresponding to a single scan are joined by gray lines. The SLAM and SLIM reconstructions used the three compartment model and show markedly less change in PCr/ATP ratio than the CSI reconstruction. For the SLAM compensated method the location of the compartmentalization mask is compensated by the same amount as the shift applied to the data, demonstrating that the SLAM method is insensitive to the exact position the FOV is placed.

4.3.5 Clinical research scans

Box plots displaying the computed PCr/ATP ratio and PCr SNR for SLAM, SLIM and CSI reconstructions of the clinical HFpEF data-set are shown in Figure 4.11. The computed statistics for each method are shown as part of Table 4.4. All methods achieved similar PCr/ATP ratios, with a lower range of values than the healthy controls. Although not significant there was a trend towards an increase in PCr SNR for both SLAM and SLIM reconstructions, compared to the CSI reconstruction (pairwise two-tailed Wilcoxon, p = 0.3125, 0.1875 respectively). Additionally, both the SLAM and SLIM reconstructions gave significantly lower PCr/ATP ratios than their respective three compartment reconstruction of the healthy volunteer data (Welch's T-test, p = 0.0233, 0.0019 respectively), whereas the CSI reconstruction did not (p = 0.1606).

Method	PCr/ATP	PCr SNR	CRLB (%)	CoV
AW SLAM	1.61 ± 0.27	14.9 ± 4.2	17.3 ± 6.6	0.19
AW SLIM	1.52 ± 0.27	15.0 ± 4.4	16.5 ± 6.0	0.18
AW CSI	1.78 ± 0.61	10.1 ± 4.1	27.5 ± 13.7	0.34

Table 4.4: Mean values of the experimental results derived for each reconstruction method applied to the clinical research HFpEF data-set. Both AW SLAM and SLIM used the three compartment model. Errors are given as ± 1 SD.

4.4 Discussion

In this work we investigated the utility of compartmentalized reconstruction techniques (SLAM and SLIM) for clinical research at 3T. Firstly we optimized acquisition parameters (phase encode scheme, use of saturation bands) with a series of preliminary studies, then we performed a repeatability study, which allowed us to test the techniques against our current reference scan and optimize the compartmentalization model. We additionally investigated the affect of cardiac gating in the repeatability study, which has consequences for both our existing CSI reconstruction and the new proposed reconstructions. Each of the preliminary studies, and the main repeatability study are discussed in detail below.

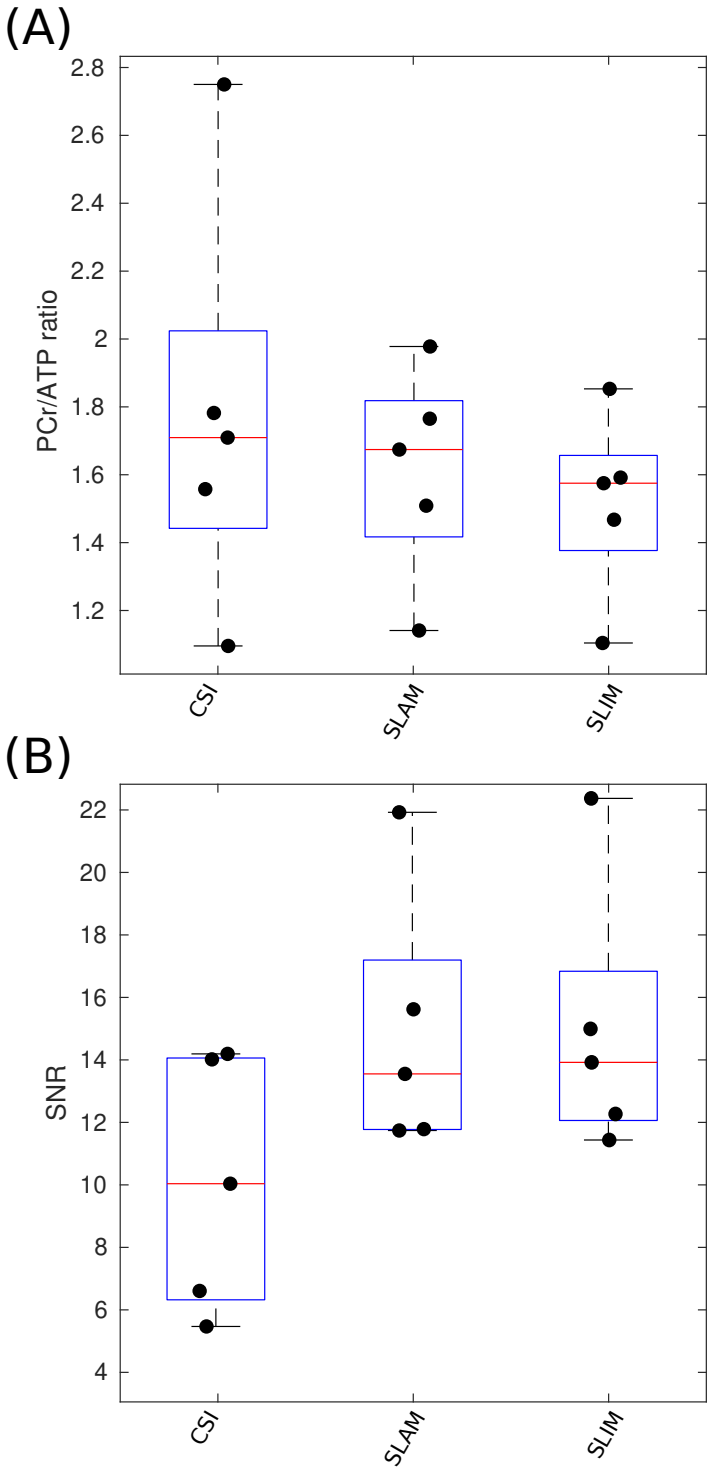

Figure 4.11: Box-plot showing PCr/ATP values (A) and PCR SNR (B) for each reconstruction method applied to the clinical research HFpEF data-set, median and IQR are indicated by box. PCr/ATP values are lower than the healthy control volunteers as is expected.

4.4.1 Preliminary study 1 – saturation bands

The PCr/ATP ratios obtained in this preliminary study, with saturation bands enabled, are in good agreement with the literature (Table 3.4), however with the saturation bands removed the PCr/ATP ratio was significantly increased across all methods. Since the PCr/ATP ratio of the myocardium should be relatively homogeneous (for individuals without comorbidities such as coronary artery disease, infiltrative or inflammatory heart disease, or hypertrophic cardiomyopathy which may lead to regional myocardial differences) this suggests that the saturation bands were suppressing signal from outside of the heart with a higher PCr/ATP ratio, particularly the chest wall, which has a very strong PCr resonance (see figure Figure 3.11). Furthermore while the increase in PCr SNR when the saturation bands were removed was not significant, a trend for an increased PCr SNR was observed, and would be expected if PCr contamination from the chest wall was being suppressed by the saturation bands.

This section therefore strongly supports the use of saturation bands placed on the chest wall and liver, which were used for all further experiments.

4.4.2 Preliminary study 2 – acquisition phase encode scheme

All the methods had plausible PCr/ATP ratios, with none of the $4 \times 4 \times 4$ or AW SLAM and SLIM reconstructions showing a significantly different PCr/ATP ratio to the reference, confirming good agreement between the compartmentalized reconstructions and the reference AW CSI method. Both the SLAM and SLIM reconstructions of the AW data-set had significantly higher PCr SNR than the AW CSI, whereas the SLAM and SLIM reconstructions of the $4 \times 4 \times 4$ data-set did not. This is surprising as at 7T the $4 \times 4 \times 4$ data-sets gave higher PCr SNR than the AW data-sets, however it is worth noting that the PCr SNR here is similar for both AW and $4 \times 4 \times 4$ data-sets, so it is likely that the study was simply under powered to detect the increase in the case of the $4 \times 4 \times 4$ data-set.

Overall, this data does support the use of AW data-sets for compartmentalized reconstruction at 3T due to the preserved PCr/ATP ratio compared to the reference

AW CSI, and significantly higher PCr SNR. A more compelling case can however be made by considering that while the AW acquisition performs slightly better than the $4\times4\times4$ acquisition, the AW acquisition can be reconstructed with the CSI algorithm to give the reference AW method, making the roll-out of AW SLAM and SLIM for use in clinical research protocols much less risky compared to the $4\times4\times4$ acquisition.

4.4.3 Main study – repeatability

In this repeatability study we compared gated and un-gated versions of AW SLAM and AW SLIM, as well as gated AW CSI to the reference un-gated AW CSI method used for clinical research at our center. We also investigated the effect of adding a fourth compartment, containing either the blood within the left ventricle, or the blood within both left and right ventricles to the reconstruction. Overall we found that both three and four compartment SLAM and SLIM reconstructions were able to deliver a worthwhile improvement in the PCr/ATP CRLB and repeatability over the un-gated AW CSI method currently in use, while removing the possibility of un-intended bias when selecting the midseptal voxel in a conventional CSI.

4.4.3.1 Un-gated AW CSI

The measured PCr/ATP ratio (2.27±0.50) is at the upper end of literature values (Table 3.4), however, given the lack of morbidities and good health of the volunteer cohort, it is not unreasonable. When compared to the repeatability study performed by Tyler et al.[57], using the same hardware, but with cardiac gating and a higher resolution (31 minute) acquisition, the PCr/ATP ratio is slightly higher (Tyler et al. PCr/ATP = 2.07 or 2.14), although this may be influenced by the number of outliers with a PCr/ATP ratio of ≈3. The coefficients of repeatability and variation were closely comparable (CoV = 0.22 (this work), 0.18 (Tyler et al.) and CoR = 1.0 (this work), 1.1 (Tyler et al.)), suggesting the data presented here is broadly representative of the capabilities of the technique.

4.4.3.2 Gated AW CSI

The PCr/ATP ratio, PCr SNR and PCr/ATP CRLB of the gated AW CSI were very close in value to the un-gated AW CSI ($< 10\%$) with no significant differences, however there were larger differences in the CoV and CoR, with the gated AW CSI performing better. This indicates that acquiring data over all cardiac phases does not introduce systematic bias compared to a gated acquisition for AW CSI, but does introduce variance into the result. One explanation for this is that for an un-gated acquisition, over many phase encodes, the mid-septum of the heart will be on average in the position determined by the localizer, however for the critical central phase encodes in an individual scan this is not necessarily true, adding variance compared to the gated scan. This explanation is supported by the relatively small difference gating made to the AW SLAM and SLIM methods, where the larger sensitive volume which encompasses the septum for all cardiac phases makes the reconstruction relatively immune to cardiac motion.

Overall the gated AW CSI data indicates that cardiac gating is worthwhile for ^{31}P CSI, because reducing the CoR and CoV will allow smaller patient cohort sizes for the same statistical power, although compartmentalized methods may be able to deliver a similar improvement without the need to use cardiac gating which can lead to extended (up to 2$\times$ longer) scan times.

4.4.3.3 Gated and un-gated three compartment model

Both AW SLAM and AW SLIM reconstructions had PCr/ATP ratios within the range of literature values (Table 3.4). The un-gated AW SLAM had a significantly lower PCr/ATP ratio (2.03 ± 0.34) than un-gated AW CSI (2.27 ± 0.50), however the un-gated AW CSI had a PCr/ATP ratio at the top end of literature values, which may have been influenced by outliers, and the ratio recorded by AW SLAM is within the range of literature values.

The reduction in both CoR and CoV of the compartmentalized reconstructions compared to the un-gated AW CSI may in part be due to the significantly lower PCr/ATP CRLBs improving the quality of the fit. However, the similarity of the

gradients for the correlation lines (and CoR/CoV) of gated CSI and gated SLAM and SLIM, despite their significantly different CRLBs suggest that this is not the driving factor. Instead, the correlation plots suggest that the main factor driving the improvement of CoR and CoV, is cardiac motion, with the gated acquisition fully ameliorating the effect for gated SLAM, SLIM and CSI, and the increase in the size of the sensitive region, relative to AW CSI, mostly ameliorating the effect for un-gated SLAM and SLIM, as the septum of the heart will always stay within the larger area during the heart beat. Reassuringly this also implies that the reduction in CoV seen with SLAM and SLIM is a real reduction in experimental error, caused by cardiac motion, and not merely a reduction in sensitivity to changes in PCr/ATP ratio.

The implication of this is that the un-gated compartmentalized reconstructions are able to offer most of the benefits of cardiac gating to repeatability, while also giving an improvement in fit quality, as indicated by the CRLBs, without requiring gating.

When comparing the SLAM and SLIM reconstructions, the results here are close enough that it is not possible to draw definitive conclusions on the differences, and indeed any observed differences may depend on coil position/profile, FOV or saturation band placement.

4.4.3.4 Un-gated four compartment model

The PCr/ATP ratios of the SLAM and SLIM reconstructions produced using the four compartment model, were within the range of literature values, however while the LV blood compartment model produced a significantly higher PCr SNR than the un-gated AW CSI, the LV + RV blood model did not. Additionally the LV + RV blood model reconstructions gave higher CoVs and CRLBs than un-gated AW CSI, although the CoRs were lower.

The worsening of the CoV with the un-gated four compartment model, compared to the three compartment model further supports the hypothesis that the improvement in CoV over un-gated AW CSI seen with the three compartment model is due to the increased sensitive volume reducing the impact of cardiac

motion. However, the increase in PCr SNR with the LV blood model suggests that restricting the heart compartment to regions with cardiac signal can reduce noise. However, the reduction in SNR for the LV + RV blood compartment, suggests a limit where as the heart compartment becomes more detailed, especially the narrow right ventricular wall, the resolution of the scan is unable to properly resolve the compartment, so signal is lost.

Overall both four compartment models performed worse when compared to a three compartment model with un-gated data, making them unsuitable for this application.

4.4.3.5 Gated four compartment model

Similarly to the CSI and three compartment SLAM and SLIM reconstructions, the gated LV blood compartment SLAM and SLIM reconstructions have close PCr/ATP ratios, PCr SNRs and PCr/ATP CRLBs to the un-gated scans, but improved repeatability. This is possibly because the LV changes shape radically during the cardiac cycle, so the compartment mask will only be accurate for readouts falling at diastole, and therefore reading out in other cardiac phases, as will be the case if gating is not used, introduces variability.

The LV + RV blood compartment model performs poorly, across all the performance metrics (except CoR) in Table 4.3, however the worsening of performance, particularly the CoR, when cardiac gating is used suggests that the model itself is introducing error as otherwise repeatability would be expected to improve as the actual position of the heart more closely matches the segmentation.

4.4.3.6 Overview

The two main findings of the main repeatability study were: Firstly that the optimized (AW + saturation bands) compartmentalized reconstruction methods were able to offer improved performance over CSI reconstruction with or without cardiac gating. Secondly, that while cardiac gating did improve repeatability for compartmentalized reconstruction methods, the un-gated versions of these methods

were able to provide most of the improvements to repeatability of gating the CSI method, without the time penalty this would entail.

4.4.4 FOV shift

Both the SLAM and SLIM reconstructions show less variation as the FOV was shifted compared to the CSI reconstruction, indicating that they are less sensitive to the exact location of the segmentation mask, than the CSI method is to the position of its midseptal voxel. This is particularly advantageous as is reduces the impact of patient motion between the localizer and ^{31}P acquisition and any errors in placement of the CSI grid.

Significantly, the consistency of the mean PCr/ATP ratio when the FOV shift is compensated for (i.e. the segmentation mask is correctly defined for the FOV acquired), with the SLAM reconstruction, shows that so long as the segmentation mask is correctly defined during analysis, the compartmentalized methods are relatively immune to errors when defining the ^{31}P acquisition FOV. Excitingly this could allow un-intended bias to be eliminated almost entirely, by using a machine learning algorithm to perform the segmentation, after acquisition, knowing that the positioning of the FOV has little impact on the final result.

4.4.5 Clinical research scans

All the scans had plausible mean PCr/ATP ratios, lower than the literature values for healthy volunteers (Table 3.4), as would be expected [11]. Significant differences to the CSI reconstruction in PCr SNR, and PCr/ATP CRLB were not found, however, together with the CoV, were trending towards an improvement, which is supported by the healthy volunteer cohort results.

The benefit to using the compartmentalized methods for these participants can be seen by calculating the sample size required for an experiment, to investigate if HFpEF affects the PCr/ATP ratio compared to healthy controls, to have a statistical power of 0.8 (t-test, α=0.05). Using the means and standard deviations tabulated in Tables 4.3 and 4.4 un-gated AW CSI would require 34 participants

in total, whereas SLAM would require 16 participants and SLIM 8 participants. This is reflected in the significant difference in PCr/ATP ratio, between the healthy volunteer repeatability study cohort and HFpEF cohort, for the SLAM and SLIM reconstructions, but not the CSI reconstruction.

The reduced number of participants required to achieve adequate statistical power, for a given effect size, has the clear clinical research advantage that fewer patient volunteers would need to be recruited for a given study, allowing studies to either be completed more efficiently, or to become feasible, where they were not before, due to small effect size.

4.5 Conclusion

In this chapter I have validated the use of compartmentalized reconstruction techniques for ^{31}P spectroscopy in the heart at the clinical field strength of 3T, building on the work at 7T presented in Chapter 3. I first performed two preliminary studies to optimize our acquisition, before conducting a healthy-volunteer repeatability study, assessing both gated and un-gated acquisitions, investigating the effect of moving the ^{31}P scan FOV, and finally using the reconstruction techniques to analyze clinical research data.

I showed that compartmentalized reconstruction of the AW scan, improves both fitting confidence and repeatability compared to the CSI reconstruction, making it suitable for use in clinical research. I also showed that cardiac gating improves repeatability for compartmentalized reconstructions, although the un-gated compartmentalized reconstructions are able to deliver similar repeatability to gated CSI reconstructions, which allows for a simpler protocol and shorter scan times. Finally, I demonstrated that the compartmentalized reconstructions are independent of the exact position of the CSI FOV acquired, allowing un-intended user bias to be potentially eliminated, by using automated segmentation[124] during reconstruction.

Overall the findings of this chapter make a compelling case for compartmentalized reconstruction of clinical ^{31}P cardiac spectroscopic research scans, as the improved repeatability may allow smaller effect sizes to be observed for a given cohort size,

while the technique ameliorates concerns about un-intended bias, which are even more important as effects become more subtle.

5

Conclusions and future work

Contents

5.1 Conclusions

The main aim of this work was to improve hyperpolarized ^{13}C and ^{31}P cardiac MRSI through the use of advanced spectroscopic imaging techniques.

5.1.1 Specific achievements

5.1.1.1 Development of the HSS sequence

I developed the HSS sequence for pre-clinical hyperpolarized [1-^{13}C]pyruvate imaging, using simulations, phantom experiments and pre-clinical experiments to validate the sequence. I implemented the code on the scanner, including a rewind algorithm and the image reconstruction pipeline, created and calibrated the T_2^* of the phantom and performed all the scans. Additionally I investigated the effect of GIRF correction

and NUFFT algorithm choice. The sequence was implemented and used for clinical scans by the clinical hyperpolarizer team.

5.1.1.2 Development of compartmentalized reconstruction at 7T

I created the full reconstruction pipeline used in Chapters 3 and 4, with the exception of the OXSA toolbox fitting functions, and SLAM algorithm, which was adapted from code supplied by Profs Paul Bottomley and Yi Zhang. I also wrote the code to perform the simulations and calculate the SRFs. I performed a repeatability study to assess the viability of these techniques at 7T, which showed that they could significantly improve SNR while maintaining the measured PCr/ATP ratio, and improving repeatability at this field strength. The AW SLAM method, which proved to be the optimal technique, was a novel addition proposed by this work.

5.1.1.3 Validation of compartmentalized reconstruction at 3T

I performed a two pilot experiments and a repeatability study to optimize and evaluate SLAM and SLIM reconstructions for use at 3T. The experiments investigated the effect of acquisition phase encoding scheme (AW $16 \times 8 \times 8$ or uniform $4 \times 4 \times 4$) and saturation bands, cardiac gating, the number and composition of the compartments used in the segmentation model, and reconstruction algorithm (SLAM or SLIM). These experiments informed optimal acquisition and reconstruction parameters for both gated and un-gated ^{31}P acquisitions, and validated the technique for use at 3T. I also investigated the effect of segmentation/voxel placement error when reconstructing the data, showing that SLAM/SLIM are less susceptible to errors in segmentation/voxel placement. Finally, I reconstructed clinical research data from patients with HFpEF, and showed that the SLAM/SLIM algorithms found a significant difference in PCr/ATP, compared to the healthy volunteers scanned in the repeatability study, whereas the CSI reconstruction did not.

5.1.2 Applications of this work

5.1.2.1 HSS sequence

The experiments detailed in this thesis demonstrate the strengths of the HSS sequence for hyperpolarized [1-^{13}C]pyruvate imaging, particularly the compromise between spatial and temporal resolution, where it incorporates the benefits of both single and multiple shot spirals, with only a small resolution penalty compared to either scheme. As such the sequence forms a very attractive method for clinical research scans and has therefore been implemented on a clinical system and used to image patients[44]. We hope that once the upgrade to the scanner used for clinical hyperpolarized scans is complete, and more hyperpolarized media becomes available, the sequence can become a mainstay of work in OCMR and potentially other sites once more data is gathered.

5.1.2.2 Compartmentalized reconstruction

The main application of this work is to be used as a standard cardiac ^{31}P spectroscopy method for clinical research scans. At both 3T and 7T we were able to show an improvement in repeatability compared to the CSI method currently used for clinical research scans in our center. The reduction in both CoV and CoR we observed will translate into improved sensitivity to changes in PCr/ATP ratio, for instance pre- and post-treatment, which will enable more subtle effects to be observed for a given cohort size. Additionally the in-variance in the measured PCr/ATP ratio as the FOV is shifted for the SLAM reconstruction demonstrates that it has the potential to remove researcher-specific bias.

5.2 Future work

5.2.0.1 HSS sequence

The HSS sequence is under active development by the clinical hyperpolarizer team. The next steps for implementation on the clinical scanner include the addition of a GIRF compensation to the clinical reconstruction pipeline, and transition to the BART reconstruction, both of which were identified by the work contained

within this thesis as methods to improve image quality. Advances in reconstruction may also be possible where, by using an iterative algorithm to solve the inverse reconstruction problem, the full readout could be incorporated into the one shot reconstructions, therefore increasing the resolution of the single shot reconstruction.

5.2.0.2 Compartmentalized reconstruction

There are several avenues of future work I would like to pursue, following on from the compartmentalized reconstruction chapters in this thesis. The first would be to incorporate the reconstruction method into the OXSA toolbox[118], which is currently the main method of analyzing data used by clinicians in OCMR, and would therefore open up the technique to a wider audience who could benefit from it. Secondly, the data presented in Figure 4.10, where the SLAM reconstruction is relatively immune to shifts in FOV, as long as the correct segmentation mask is used, suggests that an automated segmentation approach (e.g. machine learning[124]) could be used to remove bias when analyzing the data. Finally, alternative methods for solving the SLAM/SLIM equation, could be investigated. This would be particularly applicable to the SLAM* equation (Section 3.6.3), where regularization of the matrix to be inverted may be required. It may however, also be possible to use a dictionary based method, similar to magnetic resonance fingerprinting[125] to fit PCr, ATP and 2,3-DPG intensities directly to the data, which could improve the quality of the fit, and ameliorate concerns about the accuracy of inverting the PE/G matrix.

5.2.1 Concluding remarks

The work presented in this thesis makes a significant contribution to the fields of both hyperpolarized ^{13}C and ^{31}P cardiac MRSI, by implementing advanced spectroscopic imaging techniques to maximize the utilization of available signal, in the way that is best suited to the relevant clinical considerations, and as such has the potential to make a large positive impact on clinical research practices.

Appendices

A

A.1 The rotating frame

To simplify calculation the rotating frame is often used, whereby the point of view of the observer rotates around the B_0 field at the frequency $\omega_0 = \gamma B_0$ such that in the absence of RF-pulses or relaxation the net magnetization appears stationary. Here we will derive the relationship between the rotating frame and the stationary laboratory frame. All variables are as defined in figure A.1.

At $t = 0$ the net magnetization vector has position (i, j) which can be expressed as:

$$i = |M|\cos(\theta), \qquad j = |M|\sin(\theta) \tag{A.1}$$

and at time $t = t$ the magnetization processes to (i', j') which can be expressed as:

$$i' = |M|\cos(\theta + \omega_0 t), \qquad j' = |M|\sin(\theta + \omega_0 t) \tag{A.2}$$

These expressions can be expanded using trigonometric identities to:

$$\begin{aligned} i' &= |M|[\cos(\theta)\cos(\omega_0 t) - \sin(\theta)\sin(\omega_0 t)] \\ j' &= |M|[\sin(\theta)\cos(\omega_0 t) + \cos(\theta)\sin(\omega_0 t)] \end{aligned} \tag{A.3}$$

By substituting in the equations for i, j we get expressions to transform between the two frames:

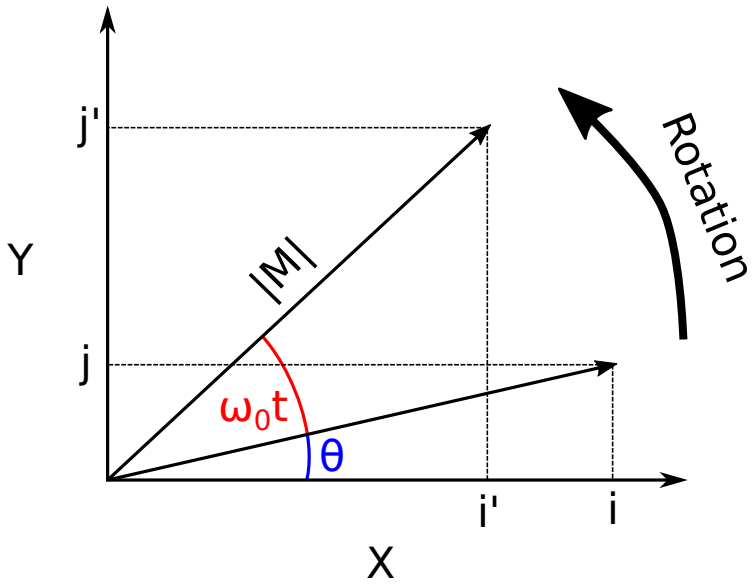

Figure A.1: Diagram of variables used to define the rotating frame, the axes define the laboratory frame. $i = M_x(0)$, $i' = M_x(t)$, $j = M_y(0)$, $j' = M_y(t)$

$$\begin{aligned} i' &= i\cos(\omega_0 t) - j\sin(\omega_0 t) \\ j' &= i\sin(\omega_0 t) + j\cos(\omega_0 t) \end{aligned} \tag{A.4}$$

and by rearrangement

$$\begin{aligned} i &= i'\cos(\omega_0 t) + j'\sin(\omega_0 t) \\ j &= -i'\sin(\omega_0 t) + j'\cos(\omega_0 t). \end{aligned} \tag{A.5}$$

Noting that in this derivation rotation is defined in the anti-clockwise direction. If rotation is defined in the clockwise direction, as is the case for the classical derivations contained within this thesis, the result becomes

$$\begin{aligned} i &= i'\cos(\omega_0 t) - j'\sin(\omega_0 t) \\ j &= i'\sin(\omega_0 t) + j'\cos(\omega_0 t). \end{aligned} \tag{A.6}$$

A.2 Solution to differential equations describing Larmor procession

Solution to the set of differential equations:

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = \gamma B_0 M_y, \qquad \frac{\mathrm{d}M_y}{\mathrm{d}t} = -\gamma B_0 M_x, \qquad \frac{\mathrm{d}M_z}{\mathrm{d}t} = 0 \tag{A.7}$$

Start by finding solution for M_x, differentiate $\mathrm{d}M_x/\mathrm{d}t$ with respect to time to generate equation in terms of M_x only:

$$\frac{\mathrm{d}^2 M_x}{\mathrm{d}t^2} = \gamma B_0 \frac{\mathrm{d}M_y}{\mathrm{d}t} = (-(\gamma B_0)^2 M_x) \tag{A.8}$$

Note that:

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\sin(at) = -a^2\sin(at), \qquad \frac{\mathrm{d}^2}{\mathrm{d}t^2}\cos(at) = -a^2\cos(at) \tag{A.9}$$

By inspection equation A.8 yields the solutions:

$$M_x(t) = \sin(\gamma B_0 t), \qquad M_x(t) = \cos(\gamma B_0 t) \tag{A.10}$$

Which lead to a general solution of:

$$M_x(t) = A\sin(\gamma B_0 t) + B\cos(\gamma B_0 t) \tag{A.11}$$

M_y can then be calculated as $M_y = \frac{1}{\gamma b_0}\frac{\mathrm{d}}{\mathrm{d}t}M_x$ to give:

$$M_y(t) = A\cos(\gamma B_0 t) - B\sin(\gamma B_0 t) \tag{A.12}$$

At $t = 0$ $M_x = B$ and $M_y = A$, hence $A = M_y(0)$ and $B = M_x(0)$ to give the final solutions:

$$\begin{aligned} M_x(t) &= M_y(0)\sin(\gamma B_0 t) + M_x(0)\cos(\gamma B_0 t) \\ M_y(t) &= M_y(0)\cos(\gamma B_0 t) - M_x(0)\sin(\gamma B_0 t) \end{aligned} \tag{A.13}$$

and trivially:

$$M_z = M_z(0) \tag{A.14}$$

A.3 Solution to differential equations describing RF-excitation

Solution to the set of differential equations:

$$\begin{aligned}\frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma[M_y B_0 + M_z B_1 \sin(\gamma B_1 t)] \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma[-M_x B_0 + M_z B_1 \cos(\gamma B_1 t)] \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= -\gamma[M_x B_1 \sin(\gamma B_1 t) + M_y B_1 \cos(\gamma B_1 t)]\end{aligned} \tag{A.15}$$

temporarily switch from the laboratory to the rotating frame

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = 0, \qquad \frac{\mathrm{d}M_y}{\mathrm{d}t} = \gamma B_1 M_z, \qquad \frac{\mathrm{d}M_z}{\mathrm{d}t} = -\gamma B_1 M_y \tag{A.16}$$

As in section A.2 the general solution is:

$$\begin{aligned}M_y(t) &= A\sin(\gamma B_1 t) + B\cos(\gamma B_1 t) \\ M_z(t) &= A\cos(\gamma B_1 t) - B\sin(\gamma B_1 t)\end{aligned} \tag{A.17}$$

if at $t = 0$ the system is at equilibrium, $M_x(0) = 0$, $M_y(0) = 0$ and $M_z(0) = M_0$ we get the following equations:

$$\begin{aligned}M_x(t) &= 0 \\ M_y(t) &= M_0 \sin(\gamma B_1 t) \\ M_z(t) &= M_0 \cos(\gamma B_1 t)\end{aligned} \tag{A.18}$$

Now using the relationship between the stationary and rotating frame derived in section A.1 we reach the equations:

$$\begin{aligned}0 &= M_x \cos(\gamma B_0 t) - M_y \sin(\gamma B_0 t) \\ M_0 \sin(\gamma B_1 t) &= M_x \sin(\gamma B_0 t) + M_y \cos(\gamma B_0 t) \\ M_0 \cos(\gamma B_1 t) &= M_z(t)\end{aligned} \tag{A.19}$$

Which by substitution can be solved to give:

$$
\begin{aligned}
M_x(t) &= M_0 \sin(\gamma B_1 t)\sin(\gamma B_0 t) \\
M_y(t) &= M_0 \sin(\gamma B_1 t)\cos(\gamma B_0 t) \\
M_z(t) &= M_0 \cos(\gamma B_1 t)
\end{aligned}
\tag{A.20}
$$

A.4 C code for HSS rewind algorithm

C code used to implement the rewind algorithm detailed in section 2.3.2.

```
/* Andrew Tyler 16/5/2018 implement algorithm prototyped in rewind.m
   create trapezium gradient to return k and g to zero */

#include <math.h>
#include <stdio.h>
#include <stdlib.h>


void calc_return(double**, double**, int*, double*, double*, double,
      double, double, double);

void calc1D(double** ,int*, double, double, double, double, double,
    double);



void calc_return(xgrad, ygrad, lgrad, K0, G0, gMax, sMax, tStep, gamma)

     double** xgrad;
     double** ygrad;
     int* lgrad;
     double* K0;
     double* G0;
     double gMax;
     double sMax;
     double tStep;
     double gamma;

{

  /* wrapper for calc1D which calculates the seperate x and y gradient
      trajectories */

  double *x1d, *y1d;
  int out_length_x, out_length_y;
  int grad_length;
  double *gxptr, *gyptr;
```

```
  calc1D(&x1d, &out_length_x, K0[0], G0[0], gMax, sMax, tStep, gamma);
  calc1D(&y1d, &out_length_y, K0[1], G0[1], gMax, sMax, tStep, gamma);

  grad_length = (((out_length_x) > (out_length_y)) ? (out_length_x) :
      (out_length_y));


  /* assign memory for gradients */
  if (( *xgrad = (double*) malloc( grad_length * sizeof(double)))==NULL){
    exit (0);
  }
  if (( *ygrad = (double*) malloc( grad_length * sizeof(double)))==NULL){
    exit (0);
  }

  /* fill gradients */
  gxptr = *xgrad;
  gyptr = *ygrad;
  double * write_x1d = x1d;
  double * write_y1d = y1d;

  int idx = 0;

  for (idx = 0; idx < grad_length; idx++){
    if (idx < out_length_x){
      *gxptr = *write_x1d;
      write_x1d++;
    }
    if (idx < out_length_y){
      *gyptr = *write_y1d;
      write_y1d++;
    }
    gxptr++;
    gyptr++;
  }

  *lgrad = grad_length;
  free(x1d);
  free(y1d);
}


void calc1D(out, out_length, K0, G0, gMax, sMax, tStep, gamma)

  double** out;
  int* out_length;
  double K0;
```

```
  double G0;
  double gMax;
  double sMax;
  double tStep;
  double gamma;

{


  /* calculate the trajectory to return both k and g to zero */

  double t_to_zero;
  int p_to_zero;
  double k_to_zero;
  int to_origin;

  /* calculate the k value when gradient is brought to zero as quickly
     as possible */

  t_to_zero = fabs(G0 / sMax);
  p_to_zero = round( t_to_zero / tStep );
  k_to_zero = 0.5 * t_to_zero * G0 * gamma;

  /* decide trapezium type */

  if (G0*K0 > 0){
    to_origin = 1;
  }else if (fabs(k_to_zero) > fabs(K0)){
    to_origin = 1;
  }else{
    to_origin = 0;
  }


  /* do calculation for towards origin trapezium */

  if (to_origin){

    double sign = k_to_zero / fabs(k_to_zero);
    double t_trapz = sqrt((K0+k_to_zero)/(sign*0.25*sMax*gamma));
    int p_trapz = round(t_trapz/tStep);

    if (( *out = (double*)
        malloc((p_trapz+p_to_zero+1)*sizeof(double)))==NULL){
      exit (0);
    }
    double * optr = *out;
```

```
  *optr = G0;
  int idx = 0;
  int p_switch = ceil((double)p_to_zero + (double)p_trapz*0.5 + 1.0);
  *out_length = p_to_zero + p_trapz + 1;


  /* write gradient output to array */

  for (idx = 1; idx < p_switch; idx++){
    *(optr+idx) = *(optr + idx -1) - tStep * sMax * sign;
  }
  for (idx = p_switch; idx < *out_length; idx++){
    *(optr+idx) = *(optr + idx -1) + tStep * sMax * sign;
  }


}else{

  /* Do calculation for away from origin trapz */

  /* use quadratic formual to solve simultanious equations to
     find t1 and t2 for the trapz */

  double t1;
  double t1a;  //two solns to quad eqn we will take biggest
  double t2;
  int p1;
  int p2;

  double sign = G0 / fabs(G0);

  double a = sMax * sign;
  double b = 2 * G0;
  double c = K0 / gamma + 0.5 * sign * G0 * G0 / sMax;

  t1 = (-b - sqrt(b*b - 4*a*c))/(2 * a);
  t1a = (-b + sqrt(b*b - 4*a*c))/(2 * a);

  if (t1a > t1)   // take biggest output
    t1 = t1a;

  t2 = t1 + sign*G0/sMax;

  p1 = round(t1/tStep);
  p2 = round(t2/tStep);

  /* create array to write output to */
```

```
    *out_length = p1 + p2 + 1;
    if (( *out = (double*) malloc((*out_length)*sizeof(double)))==NULL)
      exit (0);
    double * optr = *out;

    /* write output */

    *optr = G0;
    int idx = 1;

    for(idx = 1; idx < p1 + 1; idx++){
      *(optr+idx) = *(optr+idx-1) + sMax * sign * tStep;
    }
    for(idx = p1 + 1; idx < *out_length; idx++){
      *(optr+idx) = *(optr+idx-1) - sMax * sign * tStep;
    }

  }


}
```

References

[1] F.W. Smith et al. “Nuclear Magnetic Resonance Tomographic Imaging in Liver Disease”. In: *Lancet* 317 (1981), pp. 963–966.

[2] R.L. Nunnally and P.A. Bottomley. “Assessment of Pharmacological Treatment of Myocardial Infarction by Phosphorous-31 NMR with Surface Coils”. In: *Science* 211 (1981), pp. 177–180.

[3] Ruomin Hu et al. “X-nuclei Imaging: Current State, Technical Challenges, and Future Directions”. In: *Journal of Magnetic Resonance Imaging* 51.2 (2020), pp. 355–376.

[4] Jan H. Ardenkjær-Larsen et al. “Increase in signal-to-noise ratio of $> 10,000$ times in liquid-state NMR”. In: *Proceedings of the National Academy of Sciences* 100.18 (2003), pp. 10158–10163.

[5] Paul A. Bottomley. “Noninvasive Study of High-Energy Phosphate Metabolism in Human Heart by Depth-Resolved 31P NMR Spectroscopy”. In: *Science* 229.4715 (1985), pp. 769–772.

[6] Torsten Doenst, Tien Dung Nguyen, and E. Dale Abel. “Cardiac Metabolism in Heart Failure: implications beyond ATP production”. In: *Circulation Research* 113.6 (2013), pp. 709–724.

[7] William D. Watson et al. “Use of Cardiac Magnetic Resonance to Detect Changes in Metabolism in Heart Failure”. In: *Cardiovascular Diagnosis and Therapy* 10.3 (2020).

[8] E W Gertz et al. “Myocardial Substrate Utilization During exercise in Humans. Dual Carbon-Labeled Carbohydrate Isotope Experiments.” In: *The Journal of Clinical Investigation* 82.6 (Dec. 1988), pp. 2017–2025.

[9] H.A. Krebs and W.A. Johnson. “Citric Acid in Intermediate Metabolism in Animal Tissues”. In: *Enzymologie* 4 (1937), p. 148.

[10] Peter Mitchell. “Coupling of Phosphorylation to Electron and Hydrogen Transfer by a Chemi-Osmotic type of Mechanism”. In: *Nature* 191 (1961), pp. 144–148.

[11] Stefan Neubauer. “The Failing Heart — An Engine Out of Fuel”. In: *New England Journal of Medicine* 356.11 (2007), pp. 1140–1151.

[12] Qutuba G. Karwi et al. “Loss of Metabolic Flexibility in the Failing Heart”. In: *Frontiers in Cardiovascular Medicine* 5 (2018), p. 68.

[13] M. F. Allard et al. “Contribution of Oxidative Metabolism and Glycolysis to ATP Production in Hypertrophied Hearts”. In: *American Journal of Physiology – Heart and Circulatory Physiology* 267.2 (1994), H742–H750.

[14] Mark A. Peterzan et al. “Metabolic Remodeling in Hypertrophied and Failing Myocardium: a Review”. In: *American Journal of Physiology-Heart and Circulatory Physiology* 313.3 (2017), H597–H616.

[15] W. Gerlach and O. Stern. “Der Experimentelle Nachweis der Richtungsquantelung im Magnetfeld”. In: *Z. Physik* 9 (1922), pp. 349–352.

[16] P.J. Hore, Jones J. A., and Wimperis S. *NMR: The Toolkit*. Oxford University Press, 2000.

[17] E. Mark Haacke et al. *MRI Physical Principles and Sequence Design*. Wiley-Liss, 1999.

[18] Peter W. Atkins. *Physical Chemistry*. Oxford University Press, 1978.

[19] D. Freude. *Spectroscopy*. Chap. NMR. URL: `https://home.uni-leipzig.de/energy/pdf/freuse4.pdf`.

[20] Stuart Clare. “Functional MRI : Methods and Applications”. PhD thesis. University of Nottingham, 1997.

[21] F. Bloch. “Nuclear Induction”. In: *Physical Review* 70.7 (1946), pp. 460–474.

[22] N. Bloembergen, E.M. Purcell, and R.V. Pound. “Relaxation Effects in Nuclear Magnetic Resonance Absorption”. In: *Physical Review* 73.7 (1948), pp. 679–712.

[23] Levitt Malcolm H. *Spin Dynamics : Basics of Nuclear Magnetic Resonance.* Vol. 2nd ed. Wiley, 2008.

[24] D.I. Hoult and R.E. Richards. “The Signal to Noise Ratio of the Nuclear Magnetic Resonance Experiment”. In: *Journal of Magnetic Resonance* 24 (1969), pp. 71–85.

[25] Derek Abbott et al. “Simple Derivation of the Thermal Noise Formula Using Window-Limited Fourier Transforms and Other Conundrums”. In: *IEEE Transactions on Education* 39.1 (1996), pp. 1–13.

[26] J.B. Johnson. “Thermal Agitation of Electricity in Conductors”. In: *Nature* 119 (1927), pp. 50–51.

[27] H. Nyquist. “Thermal Agitation of Electric Charge in Conductors”. In: *Phys. Rev.* 32.1 (1928), pp. 110–113.

[28] T.W. Redpath. “Signal to Noise Ratio in MRI”. In: *The British Journal of Radiology* 71 (1998), pp. 704–707.

[29] Thomas C. Cosmus and Michael Parizh. “Advances in Whole-Body MRI Magnets”. In: *IEEE Transactions on Applied Superconductivity* 21.3 (2011), pp. 2104–2109.

[30] John Pauly, Dwight Nishimura, and Albert Macovski. “A k-Space analysis of Small-Tip-Angle Excitation”. In: *Journal of Magnetic Resonance* 81 (1989), pp. 43–56.

[31] Craig H. Meyer et al. “Simultaneous Spatial and Spectral Selective Excitation”. In: *Magnetic Resonance in Medicine* 15 (1990), pp. 287–304.

[32] Adrianus J. Bakermans et al. “Human Cardiac 31P-MR Spectroscopy at 3 Tesla Cannot Detect Failing Myocardial Energy Homeostasis During Exercise”. In: *Frontiers in Physiology* 8 (2017), p. 939.

[33] Paul A Bottomley. "Selective Volume Method for Performing Localized NMR Spectroscopy". In: *U.S. patent 4 480 228* (1984).

[34] Paul A Bottomley, Thomas B Foster, and Robert D Darrow. "Depth-Resolved Surface-Coil Spectroscopy (DRESS) for in Vivo 1H, 31P, and 13C NMR". In: *Journal of Magnetic Resonance* 59 (1984), pp. 338–342.

[35] Jens Frahm, Klaus-Dietmar Merboldt, and Wolfgang Hänicke. "Localized Proton Spectroscopy Using Stimulated Echoes". In: *Journal of Magnetic Resonance* 72 (1987), pp. 502–508.

[36] W.A. Edelstein et al. "Spin Warp NMR Imaging and Applications to Human Whole-Body Imaging". In: *Physics in Medicine and Biology* 25 (1980), p. 751.

[37] R Rzedzian et al. "Real-Time Nuclear Magnetic Resonance Clinical Imaging In Paediatrics". In: *The Lancet* 322.8362 (1983), pp. 1281–1282.

[38] Rolf Pohmann and Markus von Kienlin. "Accurate Phosphorus Metabolite Images of the Human Heart by 3D Acquisition-Weighted CSI". In: *Magnetic Resonance in Medicine* 45 (2001), pp. 817–86.

[39] J Landon, J K Fawcett, and Victor Wynn. "Blood Pyruvate Concentration Measured by Specific Method in Control Subjects". In: *Journal of clinical pathology* 15 (1962), pp. 579–584.

[40] Kan-nian Hu et al. "Quantum Mechanical Theory of Dynamic Nuclear Polarization in Solid Dielectrics". In: *The Journal of Chemical Physics* 134 (2011), pp. 125105–18.

[41] W De Boer et al. "Dynamic Polarization of Protons , Deuterons , and Carbon-13 Nuclei : Thermal Contact Between Nuclear Spins and an Electron Spin-Spin Interaction Reservoir". In: *Journal of Low Temperature Physics* 15 (1974), pp. 249–267.

[42] Zhen J. Wang et al. "Hyperpolarized 13C MRI: State of the Art and Future Directions". In: *Radiology* 291.2 (2019), pp. 273–284.

[43] Sarah J. Nelson et al. "Metabolic Imaging of Patients with Prostate Cancer Using Hyperpolarized [1-13C]Pyruvate". In: *Science Translational Medicine* 5.198 (2013), 198ra108–198ra108.

[44] Andrew Apps et al. "Proof-of-Principle Demonstration of Direct Metabolic Imaging Following Myocardial Infarction Using Hyperpolarized 13C CMR". In: *JACC: Cardiovascular Imaging* 14.6 (2021), pp. 1285–1288.

[45] Peder E. Z. Larson et al. "Investigation of Analysis Methods for Hyperpolarized 13C-Pyruvate Metabolic MRI in Prostate Cancer Patients". In: *NMR in Biomedicine* 31.11 (2018), e3997.

[46] Rahul Aggarwal, Daniel B. Vigneron, and John Kurhanewicz. "Hyperpolarized [1-13C]-Pyruvate Magnetic Resonance Imaging Detects an Early Metabolic Response to Androgen Ablation Therapy in Prostate Cancer". In: *European Urology* 72.6 (2017), pp. 1028–1029.

[47] Charles H. Cunningham et al. "Hyperpolarized 13C Metabolic MRI of the Human Heart: Initial Experience". In: *Circulation Research* 119.11 (2016), pp. 1177–1182.

[48] Andrew Apps et al. "Hyperpolarised Magnetic Resonance for in vivo Real-Time Metabolic Imaging". In: *Heart* 104.18 (2018), pp. 1484–1491.

[49] A. Hansch et al. “Noninvasive Measurements of Cardiac High-Energy Phosphate Metabolites in Dilated Cardiomyopathy by Using 31P Spectroscopic Chemical Shift Imaging”. In: *European Radiology* 15 (2005), pp. 319–323.

[50] O. J. Rider et al. “Effects of Weight Loss on Myocardial Energetics and Diastolic Function in Obesity”. In: *International Journal of Cardiovascular Imaging* 29.5 (2013), pp. 1043–1050.

[51] Moritz J. Hundertmark et al. “Design and Rationale of the EMPA-VISION Trial: Investigating the Metabolic Effects of Empagliflozin in Patients with Heart Failure”. In: *ESC Heart Failure* 8.4 (2021), pp. 2580–2590.

[52] Stefan Neubauer et al. “Myocardial Phosphocreatine-to-ATP Ratio is a Predictor of Mortality in Patients With Dilated Cardiomyopathy”. In: *Circulation* 96 (1997), pp. 2190–2196.

[53] H.H. Mitchell et al. “The Chemical Composition of the Adult Human Body and its Bearing on Growth”. In: *Journal of Biological Chemistry* 158.3 (1945), pp. 625–637.

[54] Damian J. Tyler et al. “A Comparison of cardiac 31P MRS at 1.5 and 3 T”. In: *NMR in Biomedicine* 21.8 (2008), pp. 793–798.

[55] R.R Ernst and Anderson W.A. “Application of Fourier Transform Spectroscopy to Magnetic Resonance”. In: *The Review of Scientific Instruments* 37.1 (1966), pp. 93–102.

[56] Jane Ellis et al. “Reproducibility of Human Cardiac Phosphorus MRS (31P-MRS) at 7T”. In: *NMR in Biomedicine* 32 (2019), e4095.

[57] D. J. Tyler et al. “Reproducibility of 31P Cardiac Magnetic Resonance Spectroscopy at 3T”. In: *NMR in Biomedicine* 22.4 (2009), pp. 405–413.

[58] Truman R. Brown et al. “NOE Enhancements and T1 Relaxation Times of Phosphorylated Metabolites in Human Calf Muscle at 1.5 Tesla”. In: *Magnetic Resonance in Medicine* 33.3 (1995), pp. 417–421.

[59] Michael Horn et al. “31P-Nuclear Magnetic Resonance Spectroscopy of Blood: a Species Comparison”. In: *Journal of cardiovascular magnetic resonance* 2.2 (2000), pp. 143–149.

[60] Reinhold Benesch and Ruth E. Benesch. “The Effect of Organic Phosphates from the Human Erythrocyte on the Allosteric Properties of Hemoglobin”. In: *Biochemical and Biophysical Research Communications* 26.2 (1967), pp. 162–167.

[61] Andrew Tyler et al. “A 3D Hybrid-Shot Spiral Sequence for Hyperpolarized Imaging”. In: *Magnetic Resonance in Medicine* 85.2 (2021), pp. 790–801.

[62] Klaes Golman and J Stefan Petersson. “Metabolic Imaging and Other Applications of Hyperpolarized $^{13}C^{1}$”. In: *Academic Radiology* 13.8 (2006), pp. 932–942.

[63] Klaes Golman et al. “Metabolic Imaging by Hyperpolarized ^{13}C Magnetic Resonance Imaging for in vivo Tumor Diagnosis”. In: *Cancer Research* 66.22 (2006), pp. 10855–10860.

[64] Ferdia A Gallagher et al. “Magnetic Resonance Imaging of pH in vivo Using Hyperpolarized ^{13}C-Labelled Bicarbonate”. In: *Nature* 453 (2008), pp. 940–943.

[65] Vesselin Z Miloushev et al. "Metabolic Imaging of the Human Brain with Hyperpolarized ^{13}C Pyruvate Demonstrates ^{13}C Lactate Production in Brain Tumor Patients". In: *Cancer Research* 78.14 (2018), pp. 3755–3760.

[66] James T Grist et al. "Quantifying Normal Human Brain Metabolism Using Hyperpolarized [1-^{13}C]pyruvate and Magnetic Resonance Imaging". In: *NeuroImage* 189 (2019), pp. 171–179.

[67] Hans Stødkilde-Jørgensen et al. "Pilot Study Experiences With Hyperpolarized [1-13C]pyruvate MRI in Pancreatic Cancer Patients". In: *Journal of Magnetic Resonance Imaging* 51.3 (2020), pp. 961–963.

[68] Oliver J. Rider et al. "Noninvasive In Vivo Assessment of Cardiac Metabolism in the Healthy and Diabetic Human Heart Using Hyperpolarized 13C MRI". In: *Circulation Research* 126.6 (2020), pp. 725–736.

[69] Angus Z Lau et al. "Reproducibility Study for Free-Breathing Measurements of Pyruvate Metabolism Using Hyperpolarized 13 C in the Heart". In: *Magnetic Resonance in Medicine* 69.4 (2013), pp. 1063–1071.

[70] Benjamin J. Geraghty et al. "Accelerated 3D Echo-Planar Imaging with Compressed Sensing for Time-Resolved Hyperpolarized 13C Studies". In: *Magnetic Resonance in Medicine* 77.2 (2017), pp. 538–546.

[71] Hsin-Yu Chen et al. "Pulse Sequence Considerations for Quantification of Pyruvate-to-Lactate Conversion kPL in Hyperpolarized 13C Imaging". In: *NMR in Biomedicine* 32.3 (2019), e4052.

[72] Deborah K. Hill et al. "Model Free Approach to Kinetic Analysis of Real-Time Hyperpolarized 13C Magnetic Resonance Spectroscopy Data". In: *PLOS ONE* 8.9 (Sept. 2013), pp. 1–9.

[73] Markus Durst et al. "Comparison of Acquisition Schemes for Hyperpolarised 13C Imaging". In: *NMR in Biomedicine* 28.6 (2015), pp. 715–725.

[74] Dirk Mayer et al. "Application of Sub-Second Spiral Chemical Shift Imaging to Real-Time Multislice Metabolic Imaging of the Rat in vivo After Injection of Hyperpolarized $^{13}C_1$-Pyruvate". In: *Magnetic Resonance in Medicine* 62.3 (2009), pp. 557–564.

[75] Dirk Mayer et al. "In vivo Application of Sub-Second Spiral Chemical Shift Imaging (CSI) to Hyperpolarized ^{13}C Metabolic Imaging: Comparison With Phase-Encoded CSI". In: *Journal of Magnetic Resonance* 204.2 (2010), pp. 340–345.

[76] Angus Z Lau et al. "Spectral-Spatial Excitation for Rapid Imaging of DNP Compounds". In: *NMR in Biomedicine* 24.8 (2011), pp. 988–996.

[77] Florian Wiesinger et al. "IDEAL Spiral CSI for Dynamic Metabolic MR Imaging of Hyperpolarized [1-^{13}C] Pyruvate". In: *Magnetic Resonance in Medicine* 68.1 (2012), pp. 8–16.

[78] Jiazheng Wang et al. "Single Shot Three-Dimensional Pulse Sequence for Hyperpolarized ^{13}C MRI". In: *Magnetic Resonance in Medicine* 77.2 (2017), pp. 740–752.

[79] Jack J Miller et al. "^{13}C Pyruvate Transport Across the Blood-Brain Barrier in Preclinical Hyperpolarised MRI". In: *Scientific Reports* 8.1 (2018), p. 15082.

[80] Angus Z Lau, Albert P Chen, and Cunningham Charles H. "Imaging the Circumferential Hyperpolarized 13C-Bicarbonate Distribution in the Normal Heart". In: *Proc Intl Soc Mag Reson Med* 27 (2019), p. 0256.

[81] Craig H Meyer et al. "Fast Spiral Coronary Artery Imaging". In: *Magnetic Resonance in Medicine* 28.2 (1992), pp. 202–213.

[82] Daniel M Spielman, John M Pauly, and Craig H Meyer. "Magnetic Resonance Fluoroscopy Using Spirals With Variable Sampling Densities". In: *Magnetic Resonance in Medicine* 34.3 (1995), pp. 388–394.

[83] Jan-Ray Liao et al. "Reduction of Motion Artifacts in Cine MRI Using Variable-Density Spiral Trajectories". In: *Magnetic Resonance in Medicine* 37.4 (1997), pp. 569–575.

[84] Chi-Ming Tsai and Dwight G Nishimura. "Reduced Aliasing Artifacts Using Variable-Density k-space Sampling Trajectories". In: *Magnetic Resonance in Medicine* 43.3 (2000), pp. 452–458.

[85] Gary H Glover. "Simple Analytic Spiral k-space Algorithm". In: *Magnetic Resonance in Medicine* 42.2 (1999), pp. 412–415.

[86] Dong-hyun Kim, Elfar Adalsteinsson, and Daniel M Spielman. "Simple Analytic Variable Density Spiral Design". In: *Magnetic Resonance in Medicine* 50.1 (2003), pp. 214–219.

[87] Catie Chang and Gary H Glover. "Variable-Density Spiral-In/Out Functional Magnetic Resonance Imaging". In: *Magnetic Resonance in Medicine* 65.5 (2011), pp. 1287–1296.

[88] Jin Hyung Lee et al. "Fast 3D Imaging Using Variable-Density Spiral Trajectories With Applications to Limb Perfusion". In: *Magnetic Resonance in Medicine* 50.6 (2003), pp. 1276–1285.

[89] Brian Hargreaves. "Spin-Manipulation Methods for Efficient Magnetic Resonance Imaging". PhD thesis. Department of Electrical Engineering: Stanford University, 2001.

[90] Jack J Miller et al. "Robust and High Resolution Hyperpolarized Metabolic Imaging of the Rat Heart at 7 T with 3D Spectral-Spatial EPI". In: *Magnetic Resonance in Medicine* 75.4 (2016), pp. 1515–1524.

[91] Angus Z Lau et al. "Rapid Multislice Imaging of Hyperpolarized ^{13}C Pyruvate and Bicarbonate in the Heart". In: *Magnetic resonance in medicine* 64.5 (2010), pp. 1323–1331.

[92] Meir Shinnar et al. "The Synthesis of Pulse Sequences Yielding Arbitrary Magnetization Vectors". In: *Magnetic Resonance in Medicine* 12.1 (1989), pp. 74–80.

[93] Jeff H Duyn et al. "Simple Correction Method for k-space Trajectory Deviations in MRI". In: *Journal of Magnetic Resonance* 132.1 (1998), pp. 150–153.

[94] Stefanie Winkelmann et al. "An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI". In: *IEEE Transactions on Medical Imaging* 26.1 (2007), pp. 68–76.

[95] Martin Uecker, Frank Ong, and Jonathan I Tamir. "Berkeley Advanced Reconstruction Toolbox". In: *Proc Intl Soc Mag Reson Med* 23 (2015), p. 2486.

[96] Signe J Vannesjo et al. “Gradient System Characterization by Impulse Response Measurements with a Dynamic Field Camera”. In: *Magnetic Resonance in Medicine* 69.2 (2013), pp. 583–593.

[97] Kai Tobias Block and Jens Frahm. “Spiral Imaging: A Critical Appraisal”. In: *Journal of Magnetic Resonance Imaging* 21.6 (2005), pp. 657–668.

[98] Angus Z. Lau et al. “Cardiac Perfusion Imaging Using Hyperpolarized 13C Urea Using Flow Sensitizing Gradients”. In: *Magnetic Resonance in Medicine* 75.4 (2016), pp. 1474–1483.

[99] M. Guerquin-Kern et al. “Realistic Analytical Phantoms for Parallel Magnetic Resonance Imaging”. In: *IEEE Transactions on Medical Imaging* 31.3 (2012), pp. 626–636.

[100] Jeffrey A. Fessler and Bradley P. Sutton. “Nonuniform Fast Fourier Transforms Using Min-Max Interpolation”. In: *IEEE T-SP* 51.2 (2003), pp. 560–574.

[101] J.E. Schneider et al. “Automated-Shim Approach to Facilitate ^{1}H-MRS in Mouse Hearts in vivo”. In: *Proc Intl Soc Mag Reson Med* 17 (2009), p. 1782.

[102] Jürgen E Schneider et al. “Fast, High-Resolution in vivo Cine Magnetic Resonance Imaging in Normal and Failing Mouse Hearts on a Vertical 11.7 T System”. In: *Journal of Magnetic Resonance Imaging* 18.6 (2003), pp. 691–701.

[103] Christopher T. Rodgers et al. “Human Cardiac 31P Magnetic Resonance Spectroscopy at 7 Tesla”. In: *Magnetic Resonance in Medicine* 72 (2014), pp. 304–315.

[104] Xiaoping Hu et al. “SLIM : Spectral Localization by Imaging”. In: *Magnetic Resonance in Medicine* 8 (1988), pp. 314–322.

[105] Markus Von Kienlin and Raymond Mejia. “Spectral Localization with Optimal Pointspread Function”. In: *Journal of Magnetic Resonance* 94 (1991), pp. 268–287.

[106] Zhi Pei Liang and Paul C. Lauterbur. “A Theoretical Analysis of the SLIM Technique”. In: *Journal of Magnetic Resonance, Series B* 102 (1993), pp. 54–60.

[107] Martin Meininger et al. “Concentrations of Human Cardiac Phosphorus Metabolites Determined by SLOOP 31P NMR Spectroscopy”. In: *Magnetic Resonance in Medicine* 41 (1999), pp. 657–663.

[108] Meinrad Beer et al. “Absolute Concentrations of High-Energy Phosphate Metabolites in Normal , Hypertrophied , and Failing Human Myocardium Measured Noninvasively With 31P-SLOOP Magnetic Resonance Spectroscopy”. In: *Journal of the American College of Cardiology* 40.7 (2002), pp. 1267–1274.

[109] Yi Zhang et al. “Magnetic Resonance Spectroscopy with Linear Algebraic Modeling (SLAM) for Higher Speed and Sensitivity”. In: *Journal of Magnetic Resonance* 218 (2012), pp. 66–76.

[110] Yi Zhang et al. “Highly Accelerated Chemical Exchange Saturation Transfer (CEST) Measurements with Linear Algebraic Modeling”. In: *Magnetic Resonance in Medicine* 76.1 (2016), pp. 136–144.

[111] Yi Zhang et al. “Highly-Accelerated Quantitative 2D and 3D Localized Spectroscopy with Linear Algebraic Modeling (SLAM) and Sensitivity Encoding”. In: *Journal of Magnetic Resonance* 237 (2013), pp. 125–138.

[112] Guan Wang et al. “High-Resolution and Accelerated Multi-Parametric Mapping with Automated Characterization of Vessel Disease Using Intravascular MRI”. In: *Journal of Cardiovascular Magnetic Resonance* 19 (2017), p. 89.

[113] Mark E. Ladd et al. “Pros and Cons of Ultra-High-Field MRI/MRS for Human Application”. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 109 (2018), pp. 1–50.

[114] Ladislav Valkovič et al. “Adiabatic Excitation for 31P MR Spectroscopy in the Human Heart at 7T: A Feasibility Study”. In: *Magnetic Resonance in Medicine* 78.5 (2017), pp. 1667–1673.

[115] A. Hasse et al. “FLASH imaging. Rapid NMR Imaging Using Low Flip-Angle Pulses”. In: *Journal of Magnetic Resonance* 67 (1986), pp. 258–266.

[116] Matthew D. Robson, Damian J. Tyler, and Stefan Neubauer. “Ultrashort TE Chemical Shift Imaging (UTE-CSI)”. In: *Magnetic Resonance in Medicine* 53 (2005), pp. 267–274.

[117] Christopher T. Rodgers and Matthew D. Robson. “Receive Array Magnetic Resonance Spectroscopy: Whitened Singular Value Decomposition (WSVD) Gives Optimal Bayesian Solution”. In: *Magnetic Resonance in Medicine* 63.4 (2010), pp. 881–891.

[118] Lucian A. B. Purvis et al. “OXSA: An Open-Source Magnetic Resonance Spectroscopy Analysis Toolbox in MATLAB”. In: *PLOS ONE* 12.9 (2017), e0185356.

[119] Leentje Vanhamme, Aad van den Boogaart, and Sabine Van Huffel. “Improved Method for Accurate and Efficient Quantification of MRS Data with Use of Prior Knowledge”. In: *Journal of Magnetic Resonance* 129 (1997), pp. 35–43.

[120] Ladislav Valkovič et al. “Using a Whole-Body 31P Birdcage Transmit Coil and 16-Element Receive Array for Human Cardiac Metabolic Imaging at 7T”. In: *PLoS ONE* 12.10 (2017), e0187153.

[121] Hildo J. Lamb et al. “Reproducibility of Human Cardiac 31P-NMR Spectroscopy”. In: *NMR in Biomedicine* 9 (1996), pp. 217–227.

[122] Markus von Kienlin et al. “Advances in Human Cardiac 31P-MR Spectroscopy: SLOOP and Clinical Applications”. In: *Journal of Magnetic Resonance Imaging* 13.4 (2001), pp. 521–527.

[123] W. Landschütz et al. “Concentration of Human Cardiac 31P-Metabolites Determined by SLOOP 31P-MRS”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 6.2 (1998), pp. 155–156.

[124] Chen Chen et al. “Deep Learning for Cardiac Image Segmentation: A Review”. In: *Frontiers in Cardiovascular Medicine* 7 (2020), p. 25.

[125] D. Ma et al. “Magnetic Resonance Fingerprinting”. In: *Nature* 495 (2013), pp. 187–192.