Optimising Hyperpolarised Carbon-13 MRI for Metabolic Imaging:

A Combined Approach of Hardware Innovation, Accelerated Acquisitions and Denoised Spectra



Ayaka Shinozaki

Wolfson College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2026

Acknowledgements

This thesis would not exist without the guidance, patience, and generosity of many people, and I am grateful for the opportunity to thank them here.

First and foremost, I wish to thank my Oxford supervisors, Professor Damian Tyler and Dr. James Grist, for believing in me as a researcher and for trusting me with this project without question, even when I wanted to jet off to conferences or compete in university sporting events. Their intellectual rigour, their sharp scientific questioning, and the care with which they think through experiments shaped the way I approach research, and I hope some fraction of their wisdom has made its way into the pages that follow. I am especially grateful for their kindness during the moments when experiments would not cooperate, and when I needed to take unexpected medical leave.

I would also like to thank my industrial advisors, Dr. Chris Randall and Dr. Tony Zhou at PulseTeq, and Dr. Ian Nicholson at MRF Innovations, for sharing their expertise in RF coil development. There are only a handful of people in the world who can build multi-nuclear RF coils, and I have been incredibly fortunate to be supported by three of them. Our discussions meaningfully shaped the direction of this work.

I am grateful to the members of my thesis committee, Dr. Ladislav Valkovic and Professor Martyn Paley, for taking the time to read through this long thesis and for the rare privilege of discussing my DPhil in such depth during the viva. Thank you also to Professor Aaron Hess and Dr. Ferenc Mozes for their impactful guidance during the Transfer of Status viva. I also could not have completed the porcine experiments without the help of Professor Christoffer Laustsen and his group members, especially Lone Hanberg Sørensen, Esben Søvsø Szocska Hansen, and Duy Anh Dang, who all welcomed me into their Aarhus University laboratory.

To my colleagues in the Tyler lab, thank you for the camaraderie, the shared lunches, and the willingness to drop everything to help think through multiple MRI physics topics. Particular thanks to Kylie Yeung and Zhengyang Sun, my DPhil siblings. I could not have asked for better company through this degree.

I gratefully acknowledge the Medical Research Council iCASE scholarship and the Oxford Clarendon Fund, whose generous support made this research possible. I am conscious of how much opportunity that funding represents, and I hope the work presented here justifies the investment.

To my family and friends, both in Oxford (at Wolfson and OUPC) and far beyond, around the world, thank you for the kind words of encouragement when I needed them most. Tiffany Warmuth, Florrie Hutchinson, Dr. Gosia Nguyen, Victoria Albrecht, and Beatrice Marchegiani: thank you for the memorable holidays, for reminding me that a world exists outside academia, and for being the best cheerleaders I could have asked for. And to James R. Hale, who has lived the highs and lows of this thesis alongside me with infinite patience and boundless kindness, thank you.

Finally, to my parents, Dr. Hiroharu Shinozaki and Mrs. Kaori Shinozaki, thank you for a lifetime of unconditional love, for encouraging me to pursue whatever I set my heart on, and for supporting me through so many years of school. We made it. This thesis is as much yours as it is mine.

Abstract

Hyperpolarised (HP) carbon-13 MRI offers a unique, non-invasive window into *in vivo* metabolism by tracking the enzymatic conversion of [1-^{13}C]pyruvate into downstream metabolites such as lactate, alanine, and bicarbonate. However, its clinical translation is constrained by the rapid, irreversible decay of the enhanced signal, the low sensitivity of the ^{13}C nucleus, and the narrow acquisition window within which spatial, spectral, and temporal information must be recovered. This thesis addresses three bottlenecks in the HP ^{13}C imaging chain: receive sensitivity, acquisition speed, and quantification robustness.

First, I developed a four-element multi-layer receive array for cardiac HP ^{13}C MRI at 3 T. Through electromagnetic simulations, bench-top S-parameter measurements, and in-bore Bloch-Siegert B_1^+ mapping, I identified an 11 cm centre-to-centre spacing that prioritised field homogeneity across a heart-sized field-of-view. The array achieved a mean SNR of 14.68 $\pm$ 4.79, outperforming a commercial rigid chest array (12.96 $\pm$ 3.75), and delivered more than twice the mean receive sensitivity at depth.

Second, I implemented SENSitivity Encoding (SENSE) to accelerate 2D chemical shift imaging (CSI) fourfold, reducing acquisition time from 36 s to 9 s per time point. Using a dual-tuned ^{23}Na/^{13}C flexible receive array, I used naturally abundant sodium sensitivity maps as a surrogate for ^{13}C coil sensitivities, validated in phantom and porcine kidney studies.

Third, I developed a dynamic spectral fitting algorithm that leverages high-SNR summed spectra as prior knowledge to constrain the fitting of individual low-SNR time points. Validated across rat, mouse, and human datasets, the dynamic fit reduced the mean Cramér-Rao Lower Bound by up to 78% and kinetic model RMSE by up to 76%, enabling a 79% reduction in the sample size required to detect metabolic differences in a mouse model of obesity.

Contents

List of Abbreviations **viii**

1 Introduction **1**
1.1 Motivation . . . 1
1.2 Theory . . . 3
1.2.1 MRI Physics . . . 4
1.2.2 MR Image Acquisition . . . 7
1.2.3 Radio Frequency Coils . . . 11
1.2.4 Hyperpolarised Carbon-13 MRI . . . 15
1.2.5 Carbon-13 Metabolic Physiology . . . 19

2 Radio Frequency Coil Development for Hyperpolarised Carbon-13 MRI **21**
2.1 Abstract . . . 22
2.2 Introduction . . . 23
2.2.1 Motivation . . . 23
2.2.2 Advantages of Specialised Receive-Only Coils for HP MRI . 25
2.2.3 Hardware Principles of RF Coils . . . 27
2.2.4 The Multi-layer Coil Array Concept . . . 28
2.2.5 Computational Design and Simulation . . . 31
2.2.6 RF Coils Performance Metrics . . . 34
2.3 Characterization of a Single Rx Element . . . 39
2.3.1 LCR Circuit and Hardware Overview . . . 39
2.3.2 Building and Testing the Circuit for the Carbon-13 Frequency 40
2.3.3 Tuning and Matching Results . . . 41
2.3.4 Proton Trap Building and Testing . . . 42
2.3.5 Preamplifier Testing . . . 46
2.4 Geometric Decoupling of Two Rx Elements . . . 51
2.4.1 S21 Scattering Parameters . . . 51
2.4.2 Simulation Methods . . . 51
2.4.3 Simulation Results . . . 53
2.4.4 Bench Top Methods . . . 55

2.4.5 Bench Top Results . . . 57
2.4.6 Bench Top and Simulation Discussion . . . 59
2.4.7 B1 Field Simulations . . . 60
2.5 Simulations of Extended Array Designs . . . 62
2.5.1 Trading Geometric Decoupling with Field Homogeneity . . . 62
2.5.2 Two-Element Spacing Trade-Offs . . . 63
2.5.3 Three- and Four- element Design Development . . . 67
2.5.4 Human Phantom Simulations . . . 73
2.5.5 Discussion . . . 77
2.6 In-Bore Array Validation . . . 79
2.6.1 Scanner Hardware Requirements . . . 79
2.6.2 Multi-layer Array: Proton difference map . . . 85
2.6.3 Multi-layer Array: B1 Field Maps . . . 89
2.6.4 Benchmarking Against Commercial Carbon Coils . . . 94
2.7 Appendix . . . 105

3 Accelerating Translational HP MRI 113
3.1 Abstract . . . 113
3.2 Introduction . . . 114
3.2.1 Motivation . . . 114
3.2.2 Approaches in Fast HP MR Imaging . . . 115
3.2.3 Parallel Imaging . . . 116
3.2.4 Development of 2D SENSE CSI . . . 117
3.3 Methods . . . 122
3.3.1 SENSE Reconstruction . . . 122
3.3.2 Phantom Study . . . 123
3.3.3 In-Vivo Porcine Study . . . 124
3.4 Results . . . 127
3.4.1 Phantom Study . . . 127
3.4.2 In-Vivo Porcine Study . . . 128
3.5 Discussion . . . 130
3.6 Conclusion and Future Directions . . . 133
3.7 Appendix . . . 135

4 Denoising HP MR Spectroscopy using Spectral Fitting 138
4.1 Abstract . . . 139
4.2 Introduction . . . 140
4.2.1 Motivation . . . 140
4.2.2 Spectral Fitting in Spectroscopy . . . 141
4.2.3 Kinetic Modelling and Metabolic Rate Constants . . . 142

4.2.4 Spectral Challenges in HP MRI 142
4.2.5 Current Approaches in Spectral Fitting 146
4.2.6 Proposed Dynamic Fitting Technique 150
4.2.7 Chapter Objectives 151
4.3 Methods 152
4.3.1 Data Acquisition and Study Cohorts 152
4.3.2 Spectral Fitting and Kinetic Modelling Implementation . . . 153
4.3.3 Development of an In-house Graphical User Interface 154
4.3.4 Peak Fit Model Comparisons 157
4.4 Results 165
4.4.1 Rat Heart Study 165
4.4.2 Mouse Heart Study 169
4.4.3 Human Brain 171
4.4.4 Human Liver 174
4.5 Discussion 180
4.6 Conclusion and Future Directions 183

5 Conclusion **186**
5.1 Coil Hardware Design 186
5.2 Accelerating Translational HP MRI 190
5.3 Denoising HP MRI Using Spectral Fitting 193
5.4 The Future of HP MRI 195

References **197**

List of Abbreviations

1D, 2D, 3D . . One-, two-, or three-dimensional

AAB . . . . . Abdominal Aortic Banding

AIF . . . . . Arterial Input Function

ALT . . . . . Alanine Aminotransferase

AMARES . . Advanced Method for Accurate, Robust and Efficient Spectral Fitting

AUC . . . . . Area Under the Curve

BMI . . . . . Body Mass Index

bSSFP . . . . . Balanced Steady State Free Precession

^{13}C Carbon-13

CAD . . . . . Computer-Aided Design

CV . . . . . Coefficient of Variation

CRLB . . . . . Cramér–Rao Lower Bound

CSI . . . . . Chemical Shift Imaging

CST . . . . . Computer Simulation Technology (Microwave Studio)

CT . . . . . Computed Tomography

DC . . . . . Direct Current

DL . . . . . Deep Learning

DNP . . . . . Dynamic Nuclear Polarisation

EM . . . . . Electromagnetic

EMF . . . . . Electromotive Force

EPI . . . . . Echo-Planar Imaging

EPSI . . . . . Echo-Planar Spectroscopic Imaging

FA . . . . . Flip Angle

FDTD . . . . . Finite-Difference Time-Domain

FEM . . . . . Finite Element Method

FID . . . . . Free Induction Decay

FOV . . . . . Field of View

FWHM . . . . Full Width at Half Maximum

GE . . . . . General Electric

GLS . . . . . Generalized Least Squares

GPU . . . . . Graphics Processing Unit

GRAPPA . . . Generalized Autocalibrating Partially Parallel Acquisitions

GUI . . . . . Graphical User Interface

^{1}**H** . . . . . Proton (Hydrogen-1)

HFD . . . . . High-Fat Diet

HMRTC . . . Hyperpolarized Magnetic Resonance Technology Center

HP . . . . . Hyperpolarised

IDEAL . . . . Iterative Decomposition of Water and Fat with Echo Asymmetry and Least-Squares Estimation

IP . . . . . Intellectual Property

ISMRM . . . . International Society for Magnetic Resonance in Medicine

i.v. . . . . . Intravenous

LC . . . . . Inductor–Capacitor (circuit)

LCR . . . . . Inductor–Capacitor–Resistor (circuit)

LDH . . . . . Lactate Dehydrogenase

MASLD . . . . Metabolic Dysfunction Associated Steatotic Liver Disease

MCT . . . . . Monocarboxylate Transporter

ML . . . . . Machine Learning

MRS . . . . . Magnetic Resonance Spectroscopy

MRSI . . . . . Magnetic Resonance Spectroscopic Imaging

MRI . . . . . Magnetic Resonance Imaging

^{23}**Na** Sodium-23

NEX . . . . . Number of Excitations

NMR . . . . . Nuclear Magnetic Resonance

NNFit . . . . . Neural Network Fitting

OXSA . . . . . Oxford Spectroscopy Analysis (toolbox)
PCB . . . . . Printed Circuit Board
PDH . . . . . Pyruvate Dehydrogenase
PET . . . . . Positron Emission Tomography
PH . . . . . Pyruvate-Hydrate
PHIP . . . . . Parahydrogen Induced Polarisation
PIN diode . . Positive–Intrinsic–Negative diode
PK . . . . . Prior Knowledge
PPM . . . . . Parts Per Million
PTFE . . . . . Polytetrafluoroethylene
Q . . . . . Quality factor
RF . . . . . Radio Frequency
RMSE . . . . . Root Mean Square Error
ROI . . . . . Region of Interest
RSS . . . . . Root-Sum-of-Squares
Rx . . . . . Receive
SABRE . . . . Signal Amplification By Reversible Exchange
SAM . . . . . Specific Anthropomorphic Mannequin
SAR . . . . . Specific Absorption Rate
SEM . . . . . Standard Error of the Mean
SENSE . . . . Sensitivity Encoding
SNR . . . . . Signal-to-Noise Ratio
SPSP . . . . . Spectral-Spatial (imaging)
TE . . . . . Echo Time
TG . . . . . Transmit Gain
TOSM . . . . . Through–Open–Short–Match (calibration)
TR . . . . . Repetition Time
Tx . . . . . Transmit
VFA . . . . . Variable Flip Angle
VNA . . . . . Vector Network Analyser
^{129}**Xe** Xenon-129
3Rs . . . . . Replacement, Reduction, and Refinement (of animal research)

1 Introduction

Contents

1.1 Motivation 1
1.2 Theory 3
1.2.1 MRI Physics 4
1.2.2 MR Image Acquisition 7
1.2.3 Radio Frequency Coils 11
1.2.4 Hyperpolarised Carbon-13 MRI 15
1.2.5 Carbon-13 Metabolic Physiology 19

1.1 Motivation

Many human diseases are associated with how the body turns the fuels we eat (e.g. sugars and fats) into the energy needed for growth, movement and repair. These processes collectively are called 'metabolism'. Alterations in metabolism are potential markers of many diseases, from cancer to cardiovascular and neurological disorders. Therefore, imaging metabolism *in vivo* offers a unique window into disease mechanisms and treatment responses. However, existing methods for studying metabolism *in vivo* are limited in their ability to measure metabolic processes non-invasively and in real time. Improved imaging approaches are therefore needed, together with technical advances that enable their translation

into routine clinical practice.

This thesis revolves around a novel *in vivo* approach for measuring metabolism, called Hyperpolarised Magnetic Resonance Imaging (HP MRI). By injecting carbon-13 (^{13}C)-labelled tracers, where the polarisation has been temporarily enhanced, it becomes possible to observe their downstream conversion into other metabolites non-invasively and in real time [1]. This technique offers a >10,000-fold signal increase compared to thermally polarised nuclei, enabling the detection of transient metabolic events with unprecedented sensitivity [1]. Crucially, HP MRI allows the visualisation of altered metabolic pathways, which has been studied as indicators of tumour progression[2, 3] or cardiac diseases[4, 5], within the short-lived window before the hyperpolarised signal decays. This opens up new opportunities for observing disease progression and metabolic imaging beyond the scope of conventional clinical modalities such as proton MRI, CT, or PET.

The clinical adoption of HP carbon-13 MRI remains constrained by several technical challenges. The HP signal decays rapidly, typically within two minutes, and cannot be recovered, placing stringent demands on image acquisition speed, signal sensitivity, and data reconstruction [6]. Achieving high signal-to-noise ratio (SNR) while maintaining adequate spatial, spectral, and temporal resolution is non-trivial. My thesis addresses these challenges through three integrated avenues: first, the design and characterisation of high-sensitivity RF receive coils optimised for HP imaging, particularly for cardiac imaging; second, the development of a robust, easy-to-use, accelerated acquisition strategy using SENSitivity Encoding (SENSE) for 2D Chemical Shift Imaging (CSI); and third, the implementation of an iterative spectral fitting algorithm tailored to inform fast-decaying signals by the characteristics of the high SNR, summed dynamics of the HP data. Together, these hardware and software advances aim to improve the clinical utility of HP carbon-13 metabolic imaging.

In Chapter 1, I review the fundamental science behind my research covering MRI physics, encoding methods, RF coils, HP MRI, and carbon-13 enzymatic conversions. This serves as the general background information for all of the subsequent experimental chapters. In Chapter 2, I present my work undertaken

in partnership with PulseTeq Ltd, an RF coil design company, to develop a novel RF coil array specifically for cardiac HP-MRI to allow rapid clinical translation of this emerging medical imaging technology. In Chapter 3, my work investigates the novel application of Sensitivity Encoding for accelerating two-dimensional chemical shift imaging in HP carbon-13 MRI. As one of the first investigations of SENSE acceleration for hyperpolarised ^{13}C 2D CSI, the findings of my study contributed to the development of a faster, more efficient hyperpolarised ^{13}C MRI technique that offers a promising solution for metabolic imaging in diverse clinical scenarios. My last experimental chapter, Chapter 4, covers how implementing iterative spectral fitting to dynamic ^{13}C MRI can help recover decaying, low level signals in later time points by leveraging information from higher SNR time points. This iterative fit effectively denoises the data and facilitates HP MRI translation into the clinic. All experimental chapters start with an abstract, motivation and in-depth theory section, and are written such that these sections may stand alone. The concluding chapter (Chapter 5) integrates these advances, elaborates on my personal reflections on my DPhil, and discusses my finding's implications for future clinical translation. Together, these technical developments form a multi-faceted framework for advancing HP ^{13}C MRI, from hardware to software, laying the foundation for more reliable and clinically informative metabolic imaging.

1.2 Theory

This section provides a theoretical foundation, beginning with the basic spin physics and signal generation that underpin MRI, before progressing to the principles of MR image acquisition and the critical role of radio frequency (RF) hardware. This chapter then discusses the overview of hyperpolarisation and the enzymatic conversion of the hyperpolarised substrate in the body, to highlight the unique challenges and opportunities presented by HP ^{13}C MRI.

Table 1.1 *NMR properties of commonly encountered nuclei in in vivo NMR*

Isotope	Spin	Gyromagnetic ratio (10^7 rad T^{-1} s^{-1})	NMR frequency at 2.35 T (MHz)	Natural abundance (%)	Relative sensitivity[a]
^{1}H	1/2	26.752	100.000	99.985	1.00
^{2}H	1	4.107	15.351	0.015	1.45×10^{-6}
^{3}He	1/2	−20.380	76.181	1.4×10^{-4}	5.75×10^{-7}
^{7}Li	3/2	10.398	38.866	92.58	0.272
^{13}C	1/2	6.728	25.145	1.108	1.76×10^{-4}
^{14}N	1	1.934	7.228	99.630	1.00×10^{-3}
^{15}N	1/2	−2.712	10.137	0.370	3.86×10^{-6}
^{17}O	5/2	−3.628	13.562	0.037	1.08×10^{-5}
^{19}F	1/2	25.181	94.094	100.000	0.834
^{23}Na	3/2	7.080	26.466	100.000	9.27×10^{-2}
^{31}P	1/2	10.841	40.481	100.000	6.65×10^{-2}
^{39}K	3/2	1.250	4.672	93.100	4.75×10^{-4}
^{129}Xe	1/2	−7.452	27.856	26.44	5.71×10^{-3}

[a] Relative sensitivity is calculated as the product of NMR sensitivity (proportional to $|\gamma^3| \times I(I+1)$) and the natural abundance.

Figure 1.1: Table 1.1 reproduced from De Graaf [9].

1.2.1 MRI Physics

Magnetic Resonance Imaging is an advanced imaging modality extensively employed in clinical diagnostics for its capability to visualise anatomical structures and physiological function without exposure to ionising radiation. Fundamentally, MRI relies on the interaction between magnetic fields and nuclear spins within biological tissues. In clinical imaging, predominantly the magnetic properties of the atomic nuclei of hydrogen protons are acquired for visualisation.

Atomic nuclei comprise protons and neutrons, which possess an intrinsic quantum mechanical property known as spin angular momentum, I. Each of these fermions carries a spin of $I = 1/2$ [7]. The total spin state of a nucleus is determined by the combination of its nucleons. An even number of both protons and neutrons results in anti-parallel pairing and a net spin of zero ($I = 0$). An even total number of nucleons gives integer spins ($I = 1, 2, 3, \dots$). If the numbers of protons and neutrons are even and odd or odd and even, respectively, then the nucleus carries a half-integer spin ($I = 1/2, 3/2, 5/2, \dots$). Carbon-13, which is the main nucleus of interest in this work, has a spin of $I = 1/2$ [8].

For a nucleus to be observable in MRI, it must have a non-zero spin and therefore a nuclear magnetic moment ($\vec{\mu}$), related to the spin angular momentum $\vec{J}$ by $\vec{\mu} = \gamma\vec{J}$. A nucleus with zero spin cannot interact with an external magnetic field B_0. The

proportionality constant γ is the gyromagnetic ratio, a fundamental value unique to each nuclear isotope. The gyromagnetic ratios of the nuclei discussed in this work are $\gamma_{^1\mathrm{H}} = 26.752 \times 10^7\,\mathrm{rad\,T^{-1}s^{-1}}$ with $I = 1/2$, $\gamma_{^{13}\mathrm{C}} = 6.728 \times 10^7\,\mathrm{rad\,T^{-1}s^{-1}}$ with $I = 1/2$, and $\gamma_{^{23}\mathrm{Na}} = 7.080 \times 10^7\mathrm{rad\,T^{-1}s^{-1}}$ with $I = 3/2$ (Figure 1.1).

When these nuclei are placed in a strong external magnetic field B_0, the Zeeman interaction causes the energy levels of $I = 1/2$ nuclei to split into two populations: a low-energy state parallel to the field and a high-energy state anti-parallel to it. The energy difference between these states is $\Delta E = \gamma\hbar B_0$, where $\hbar$ is the reduced Planck constant. The relative populations at thermal equilibrium are governed by the Boltzmann distribution [9]:

$$\frac{N_\uparrow}{N_\downarrow} = \exp\left(-\frac{\Delta E}{k_B T}\right) = \exp\left(-\frac{\gamma\hbar B_0}{k_B T}\right) \tag{1.1}$$

where $N_\uparrow$ and $N_\downarrow$ are the numbers of spins in the higher and lower energy states respectively, k_B is the Boltzmann constant, and T is the absolute temperature. At physiological temperatures the thermal energy far exceeds the magnetic energy, resulting in nearly equal populations with a very small excess in the lower-energy state. This small population difference gives rise to the net magnetisation vector $\vec{M}$, the quantity measured in conventional MRI, and represents the fundamental sensitivity challenge that motivates the carbon-13 hyperpolarisation technique (section 1.2.4) used in this thesis.

The macroscopic behaviour of $\vec{M}$ in a static field B_0 is characterised by Larmor precession. The external magnetic field exerts a torque on $\vec{M}$, defined as $\vec{\tau} = \vec{M} \times \vec{B}$, causing $\vec{M}$ to precess about the B_0 axis at the angular Larmor frequency $\omega_0 = \gamma B_0$ (Figure 1.2). When a transverse oscillating magnetic field B_1 is applied at the Larmor frequency, it tips $\vec{M}$ into the transverse (xy) plane, creating a time-varying magnetic flux that induces a measurable voltage in a conducting loop through Faraday's law of induction.

The classical equations of motion governing this behaviour were derived by Felix Bloch in 1946 [11]. Bloch treated the spin ensemble as a single macroscopic

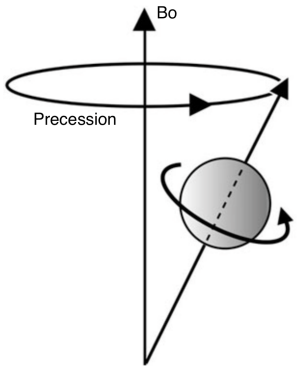


Figure 1.2: In an external magnetic field (B_0), the nucleus precesses about the B_0 axis. The frequency of this precession is the Larmor frequency. Adaptation of Figure 2.19 from reference [10].

magnetisation vector $\vec{M}$, with transverse components $M_x(t)$ and $M_y(t)$ and longitudinal component $M_z(t)$. Bloch recognised that $\vec{M}$ does not precess indefinitely but returns to its equilibrium alignment parallel to B_0 through a process called "relaxation", characterised by two independent time constants T_1 and T_2. Assuming first-order relaxation kinetics, Bloch modified the pure precession equation to give:

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} - \frac{M_x\hat{i} + M_y\hat{j}}{T_2} - \frac{(M_z - M_0)\hat{k}}{T_1} \tag{1.2}$$

Following a 90° excitation, the solutions to Equation 1.2 describe a spiralling precession of $\vec{M}$ about B_0:

$$M_x(t) = M_0\, e^{-t/T_2} \sin\omega_0 t \tag{1.3}$$

$$M_y(t) = M_0\, e^{-t/T_2} \cos\omega_0 t \tag{1.4}$$

$$M_z(t) = M_0\Big(1 - e^{-t/T_1}\Big) \tag{1.5}$$

Longitudinal relaxation, or T_1 relaxation, describes the recovery of M_z towards its equilibrium value M_0 through dissipation of energy to the surrounding lattice; T_1 is defined as the time for M_z to recover to $(1 - 1/e) \approx 63\%$ of M_0. Transverse relaxation, or T_2 relaxation, describes the decay of the transverse components due to the loss of phase coherence from internuclear interactions; T_2 is defined as the time for the transverse magnetisation to decay to $1/e \approx 37\%$ of its initial value. Because

$T_2 \leq T_1$ always holds, the transverse magnetisation typically decays to zero before the longitudinal magnetisation is fully restored. In practice, the signal decays faster still due to additional dephasing from local static field inhomogeneities, characterised by the effective transverse relaxation time T_2^*, where $1/T_2^* = 1/T_2 + 1/T_2'$, where T_2' accounts for the inhomogeneity contribution.

The above description assumes conventional MRI, where the magnetisation begins at its thermal equilibrium value and T_1 relaxation restores longitudinal magnetisation following each RF excitation. Hyperpolarised MRI differs fundamentally because the magnetisation is prepared outside the scanner in a highly non-equilibrium state. This excess longitudinal magnetisation cannot be regenerated *in vivo*; instead, it continuously relaxes back towards thermal equilibrium through T_1 relaxation, which will also be referred to as T_1 decay in this work. Consequently, both RF excitations and T_1 relaxation irreversibly deplete the available hyperpolarised magnetisation, limiting the time available for image acquisition.

1.2.2 MR Image Acquisition

The voltage induced in a conducting loop by the precessing transverse magnetisation constitutes the raw NMR signal. Spatial localisation of this signal requires the MR system to impose additional magnetic field structure beyond B_0, and to sample the signal that is then reconstructed for clinical interpretation. This subsection describes the acquired raw signal, the gradient-based encoding used to map it into an image, and the particular constraints these place on hyperpolarised ^{13}C acquisitions.

Free Induction Decay

Following the application of a B_1 field at the Larmor frequency, which tips $\vec{M}$ into the transverse plane, the precessing transverse magnetisation induces a time-varying voltage in the receive coil [12]. This raw time-domain signal is known as the Free Induction Decay (FID): a sinusoid oscillating at ω_0 whose amplitude decays as the transverse magnetisation loses coherence. The mathematical form of a single-component FID is [12]:

$$S(t) = S_0\, e^{-t/T_2^*}\, e^{i(\omega_0 t+\phi)} \tag{1.6}$$

where S_0 is the initial signal amplitude and ϕ is an initial phase offset. The signal acquired in this way is not localised, representing spin density integrated over the entire sensitive volume of the receive coil. The duration of the FID acquisition window determines the achievable spectral resolution. A longer window yields narrower linewidths in the frequency domain. In hyperpolarised carbon-13 experiments, however, the available magnetisation decreases continuously due to irreversible T_1 relaxation, while transverse coherence is simultaneously lost through T_2^*. Consequently, the FID must be acquired rapidly, creating a fundamental trade-off between spectral resolution and acquisition speed.

Spatial Encoding, Gradients, and *k*-Space

To resolve the spatial distribution of the NMR signal, three orthogonally linear magnetic field gradients (G_x, G_y, G_z) are superimposed on B_0. A gradient applied along a given direction causes the Larmor frequency to vary linearly with position along that direction, so that spins at different locations precess at distinguishable frequencies [13, 14]. This is termed "frequency encoding". To resolve a second spatial dimension, a "phase encoding" gradient pulse is applied between the excitation and the signal readout. This imparts a spatially varying phase shift, encoding position as an accumulated phase rather than an instantaneous frequency[15].

The signals acquired under these gradients are encoded in a grid of raw data within the spatial frequency domain, which is known as k-space[16]. The k-space coordinate $\vec{k}(t)$ is the time-integral of the applied gradient waveform [16]:

$$\vec{k}(t) = \frac{\gamma}{2\pi}\int_0^t \vec{G}(t')\,\mathrm{d}t' \tag{1.7}$$

and the measured signal is related to the transverse magnetisation distribution $M(\vec{r})$ by the Fourier transform[16]:

$$S(t) = \int M(\vec{r})\, e^{-i2\pi\vec{k}(t)\cdot\vec{r}}\,\mathrm{d}\vec{r} \tag{1.8}$$

Reconstruction of the image therefore requires sampling k-space sufficiently to recover $M(\vec{r})$ via the inverse Fourier transform. The spatial resolution is determined by the maximum extent of k-space sampled, and the field of view (FOV) is inversely proportional to the k-space sampling interval. According to the Nyquist–Shannon sampling theorem, the sampling density must be sufficient to avoid aliasing, whereby signal from outside the FOV folds back into the image [17].

In conventional ^{1}H MRI, the high intrinsic SNR permits repeated excitations and allows k-space to be filled incrementally across many acquisitions. In HP ^{13}C MRI, these strategies are not available. Each B_1 excitation irreversibly consumes a fraction of the non-renewable hyperpolarised magnetisation, limiting the total number of excitations, whilst simultaneous T_1 decay and T_2^* dephasing impose a hard time constraint on the acquisition window. Achieving adequate spatial resolution therefore requires a larger k-space coverage, which in turn demands more excitations and longer acquisition times. This acquisition time directly competes with the need to preserve magnetisation and capture rapidly evolving metabolic dynamics.

Spectroscopy and Chemical Shift Imaging

To add another dimensionality to this discussion, the resonant frequency of a nucleus is not determined only by the external magnetic field B_0; it is also influenced by the local electric environment surrounding the nucleus. Electrons orbiting nearby atoms generate small induced magnetic fields that partially shield the nucleus from B_0, reducing the effective field it experiences. The degree of this shielding depends on the molecular structure and the chemical bonds in which the nucleus participates. As a result, chemically distinct nuclei within the same molecule, or identical nuclei in different molecules, resonate at slightly different frequencies. This phenomenon is called the "chemical shift", and the frequency offset is conventionally expressed in dimensionless parts per million (ppm) relative to a reference compound, defined as:

$$\delta = \frac{\omega - \omega_{\text{ref}}}{\omega_{\text{ref}}} \times 10^6 \tag{1.9}$$

where ω is the resonant frequency of the nucleus of interest and ω_{ref} is that of the reference. Expressing the shift in ppm renders it independent of the static field strength, allowing direct comparison between spectra acquired at different field strengths.

Magnetic Resonance Spectroscopy (MRS) leverages the chemical shift effect to identify and quantify different molecular species based on their distinct resonant frequencies. Chemical Shift Imaging (CSI) extends this capability by adding spatial encoding to the spectroscopic acquisition. CSI is a robust multi-voxel imaging technique that obtains a large spectral bandwidth and can capture signal from a broad range of metabolites[18, 19]. Due to its high spectral resolution, CSI does not require prior knowledge of the metabolite spectra and is insensitive to inhomogeneous magnetic fields. Additionally, the ability to acquire full spectra across the imaging field-of-view facilitates retrospective quantification and flexible post-processing pipelines.

In a standard single slice, two dimensional CSI sequence, phase-encoding gradients are applied in two orthogonal directions before the FID is recorded. After a Fourier transform, this yields a dataset with two spatial dimensions and one spectral dimension. For dynamic studies, the CSI acquisition is repeated over time, producing a four-dimensional dataset. The total acquisition time for a CSI grid is proportional to the number of phase-encoding steps, and because each step requires a separate RF excitation, the temporal resolution is governed by the product of the repetition time and the total number of voxels, if all of k-space is fully sampled. It is an RF-intensive and relatively slow acquisition method. This relationship between spatial coverage and acquisition speed is a general constraint of CSI. Optimising the balance between spatial resolution, spectral resolution, acquisition speed, and SNR is consequently a central challenge in HP ^{13}C imaging protocol design, and motivates the hardware and acquisition strategies developed in this thesis.

1.2.3 Radio Frequency Coils

In practice, the conducting loop in which the precessing magnetisation induces a voltage is called a resonant radio frequency (RF) coil. An RF coil is a tuneable electrical circuit engineered to couple efficiently to nuclear spins at a specific Larmor frequency and to interface the NMR signal with the MR system's receiver chain. An RF coil comprises an inductive loop and one or more capacitors, forming a resonant LC circuit. In an MRI system, RF coils serve as the primary transducers between the radio frequency electromagnetic field and the precessing nuclear magnetisation. An RF coil can operate in transmit mode, broadcasting a defined B_1^+ field to tip $\vec{M}$ away from B_0, in receive mode, detecting the weak oscillating flux produced by the precessing transverse magnetisation and converting it into an electrical signal for amplification and digitisation, or in a combined Tx/Rx mode [20, 21]. When considering the overall coil components, the separation of transmit and receive functionalities allows each to be individually optimised. The Tx coil is designed primarily to generate a homogeneous B_1^+ field across the field of view, whilst the Rx coil is optimised to maximise SNR [20, 21]. The design and performance of RF coils directly determine the sensitivity, spatial resolution, and temporal resolution achievable in the MRI experiment.

Signal detection relies on the principle of electromagnetic induction, governed by Faraday's and Lenz's laws. As the net transverse magnetisation vector precesses in the presence of the main magnetic field, it induces an electromotive force (EMF) in the RF receive coil. According to Faraday's law, this induced EMF is proportional to the rate of change of magnetic flux through the coil loop [20, 22]:

$$\text{EMF} = -N\frac{d\Phi}{dt} \tag{1.10}$$

where N is the number of turns in the coil, Φ is the magnetic flux, and $d\Phi/dt$ is its time derivative. The negative sign, as described by Lenz's law, indicates that the induced EMF opposes the change in flux that produced it. This induced voltage is amplified by a low-noise preamplifier and digitised for image reconstruction. The

efficiency of the signal induction process is directly related to the coil geometry, the spatial distribution of the B_1^- receive field, and the electrical properties of the coil circuit [20, 22].

Several factors influence the performance of the coil, including the resonant frequency of the LC circuit as defined as:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \tag{1.11}$$

where L is the inductance and C the total system capacitance [22]. Accurately aligning the coil's resonant frequency to the desired frequency of the experiment significantly affects the performance of the coil. Capacitors are distributed within the circuit to tune the coil to the target Larmor frequency and to match the impedance to the 50 Ω input of the MR system, ensuring efficient power transfer and minimising signal reflection. Additional components such as baluns or cable traps suppress common-mode currents and prevent RF interference [20]. This modular and tuneable architecture enables RF coils to be tailored for specific nuclei, anatomical targets, and imaging applications. A key performance metric is the quality factor Q:

$$Q = \frac{2\pi f_0 L}{R_{\text{total}}} = \frac{1}{R_{\text{total}}}\sqrt{\frac{L}{C}} \tag{1.12}$$

The ratio $Q_{\text{unloaded}}/Q_{\text{loaded}}$ quantifies the relative contribution of sample noise to coil noise; maximising this ratio maximises the achievable SNR [22, 23]. In this equation, the total effective resistance of the coil circuit decomposes as:

$$R_{\text{total}} = R_{\text{coil}} + R_{\text{sample}} + R_{\text{extra}} \tag{1.13}$$

where R_{coil} accounts for conductor losses, R_{sample} for losses induced by RF currents in the surrounding tissue, and R_{extra} for radiative and capacitor losses [22].

The sensitivity $S(\mathbf{r})$ of an RF coil at a spatial location $\mathbf{r}$ is approximately given by:

$$S(\mathbf{r}) \propto \frac{B_1^-(\mathbf{r})}{\sqrt{P_n}} \tag{1.14}$$

where $B_1^-(\mathbf{r})$ is the receive magnetic field at location $\mathbf{r}$ and P_n is the total noise power, which includes both coil and sample contributions [24]. In terms of SNR, this relationship becomes:

$$\mathrm{SNR}(\mathbf{r}) \propto \frac{B_1^-(\mathbf{r})}{\sqrt{R_s + R_c}} \tag{1.15}$$

where R_s and R_c represent sample and coil resistance, respectively [25].

In terms of coil design choices, the fundamental distinction in coil geometry is between volume coils and surface coils. Volume coils, such as birdcage resonators, encompass the anatomy of interest and produce a homogeneous B_1 field across the enclosed volume, but collect noise from that entire volume [20, 21, 26]. Surface coils are planar or conformal loop elements positioned directly over the region of interest, acting as spatial filters that confine sensitivity, and hence noise collection, to the vicinity of the coil [20, 23]. The SNR advantage of surface coils near the conductor is well established, making them well suited to applications where the target tissue lies within a limited depth from the body surface [23]. In most cases, the inherent inhomogeneity of the B_1 field associated with surface coils restricts their use primarily to receive-only operation [20].

Arrays of surface coil elements, introduced as phased-array technology by Roemer and colleagues in 1990 [20], extend the sensitive volume over a larger FOV whilst preserving the high SNR intrinsic to small loop elements. In a phased array, each element connects to an independent receiver channel, and signals are combined with appropriate amplitude and phase weighting to maximise SNR at every spatial location. The spatial heterogeneity of the individual element sensitivity profiles can additionally be exploited to accelerate acquisition through parallel imaging, whereby k-space is under-sampled and the missing data recovered using prior knowledge of the coil sensitivities. Minimising inter-element coupling is essential to preserve both SNR and the accuracy of the sensitivity profiles required for this reconstruction [20, 27]. A detailed quantitative comparison of these coil types and the benefits of array-based configurations can be found in the literature here [28].

Several parameters affect Rx sensitivity, particularly for surface array coils. Coil geometry and placement relative to the sample significantly influence sensitivity [28]. Electrical tuning and matching ensure that each element resonates at the desired frequency and transfers power efficiently [29]. Variations in patient body types and tissue dielectric properties can pose challenges for clinical matching. Decoupling efficiency in arrays is vital and can be achieved through geometric overlap, preamplifier decoupling, or active detuning circuits [30]. In-depth discussions on decoupling methods are available in [31–36].

In the context of hyperpolarised ^{13}C MRI, these design considerations become particularly critical. The ^{13}C Larmor frequency at 3 T is 32.1 MHz, approximately four times lower than that of ^{1}H, which has consequences for both coil noise characteristics and the required hardware [37]. At this lower frequency, the contribution of coil noise relative to sample noise increases compared to the proton case, although sample noise remains the dominant source for human-scale receive coils [37]. The lower gyromagnetic ratio also has direct consequences for transmit efficiency because the flip angle $\theta = \gamma \int_0^\tau B_1(t)\,\mathrm{d}t$, which is the angle that the magnetisation vector is rotated away from the z-axis by an external radio frequency pulse in relation to B_0, scales linearly with γ. Achieving a given flip angle at the ^{13}C frequency requires approximately four times the B_1 amplitude, which means up to sixteen times the transmit power ($P \propto B_1^2$), relative to proton imaging under otherwise identical conditions.

In hyperpolarised experiments this constraint is compounded by the non-renewable nature of the magnetisation, where each excitation irreversibly consumes a fraction of the available polarisation, so transmit coil efficiency and flip angle accuracy are critical to preserving signal across the time-resolved acquisition [22, 37]. For ^{13}C transmit, volume resonators, most commonly Helmholtz pair or clamshell designs, are widely used in human studies owing to their ease of integration into the scanner bore, despite their inhomogeneous B_1^+ profiles, which can introduce errors in flip angle and hence in quantitative metabolic metrics [37]. On the receive side, dedicated ^{13}C coils must minimise conductor losses whilst ensuring

adequate anatomical coverage. The transient and non-renewable nature of the hyperpolarised signal, combined with the low molar concentration of downstream metabolic products demands receive arrays with high sensitivity to capture rapidly decaying, low-SNR metabolite signals [22, 23, 37].

1.2.4 Hyperpolarised Carbon-13 MRI

The diagnostic value of hyperpolarised ^{13}C MRI derives from the metabolic information that the signal carries. As established in the MRI Physics section 1.2.1, ^{13}C is a nucleus with spin $I = 1/2$ and a gyromagnetic ratio approximately four times lower than that of ^{1}H. Under thermal equilibrium conditions, this combination of low gyromagnetic ratio, low natural abundance (1.1% of all carbon), and the relatively low concentrations of endogenous carbon-containing metabolites results in a signal that is generally too weak to detect with conventional MRI. [38]. To overcome this fundamental sensitivity limitation, the nuclear spin polarisation of ^{13}C-labelled substrates can be temporarily enhanced by many orders of magnitude through a process known as hyperpolarisation, enabling the non-invasive observation of rapid metabolic conversions *in vivo* that occur within the lifetime of the hyperpolarised signal.[1].

Hyperpolarisation Techniques

Three hyperpolarisation techniques have been developed for biomedical applications of ^{13}C MRI. Dynamic nuclear polarisation (DNP) exploits microwave-driven polarisation transfer from unpaired electrons to nuclear spins at cryogenic temperatures [1]. Parahydrogen induced polarisation (PHIP) relies on the chemical incorporation of parahydrogen into a precursor molecule [39]. Signal amplification by reversible exchange (SABRE) facilitates polarisation transfer through transient, non-covalent interactions on a metal catalyst [40]. Whilst each method offers substantial signal enhancement, DNP is currently the only technique to have successfully transitioned to clinical human trials for ^{13}C metabolic imaging, and is therefore the method employed in this work [7].

In a DNP system, the ^{13}C-labelled pyruvate is mixed with a source of unpaired electrons and placed in a superconducting magnet operating above 3 T, cooled to below 1.4 K using liquid helium [1]. At this temperature and field strength, microwave irradiation drives the transfer of the high electron spin polarisation to the coupled nuclear spins, achieving polarisation levels of up to 70% in clinically approved systems [41]. The sample is then rapidly dissolved using superheated water, the free-radical electron source is removed by filtration, and the resulting solution is neutralised with a buffer before intravenous injection [42]. Once injected, the hyperpolarised pyruvate participates in enzyme-catalysed metabolic reactions *in vivo*, and the injected substrate together with its downstream metabolic products can be detected and spatially localised using the MRI acquisition principles described in the Image Acquisition section 1.2.2. However, because hyperpolarisation creates a thermally unstable, non-equilibrium spin state, the enhanced signal decays irreversibly toward thermal equilibrium via T_1 relaxation, allowing an observation window of approximately two minutes *in vivo* before the signal is lost [43].

Challenges Unique to Hyperpolarised Carbon-13 MRI

The non-renewable nature of the magnetisation is a fundamental distinction between conventional ^{1}H MRI and HP ^{13}C MRI. In standard proton imaging, T_1 relaxation continuously restores the longitudinal magnetisation toward thermal equilibrium, allowing repeated excitations and signal averaging. In HP ^{13}C MRI, the hyperpolarised nuclei decay back to their much lower thermal equilibrium state. There is no mechanism by which the enhanced polarisation can be recovered once it has decayed or been consumed by RF excitation. These constraints are compounded by the low molar concentration of downstream metabolic products such as lactate, alanine, and bicarbonate, if the injected substrate is pyruvate, which demands both high receive coil sensitivity and efficient spatial encoding within the narrow acquisition window. In the context of spectroscopic acquisition, the chemical shift separations between the metabolites of interest are large. As an example from Chapter 4 in a representative rat heart spectra where, relative to

the pyruvate resonance, the lactate signal appears at around 12 ppm, pyruvate-hydrate at around 8 ppm, alanine at around 6 ppm, and bicarbonate at around -10 ppm (figure 1.3). These peaks provide sufficient spectral dispersion for their identification, especially when the acquisition is summed across the whole decay time frame. Difficulty occurs in the analysis of the temporally resolved spectra when the large signal from the injected pyruvate substrate obscures the much smaller signals from the downstream metabolic products.

The conversion of pyruvate to its downstream metabolites can be quantified through two approaches. Simple signal ratios such as the lactate-to-pyruvate ratio are straightforward to compute but are sensitive to the timing of the acquisition and the total amount of tracer injected [37]. In contrast, metabolic rate constants such as k_{PL} (pyruvate-to-lactate) and k_{PB} (pyruvate-to-bicarbonate) represent the intrinsic speed of enzymatic flux, independent of these confounding experimental factors, and serve as metabolic biomarkers [37, 44].

In oncology, an elevated k_{PL} is often a hallmark of the Warburg effect, indicating aggressive anaerobic glycolysis even in the presence of oxygen [44, 45]. In cardiology, the rate of conversion to bicarbonate, k_{PB}, provides a direct readout of mitochondrial function and PDH activity, and a reduction in k_{PB} can indicate myocardial ischaemia or the metabolic inflexibility associated with heart failure [46–48]. By isolating these rate constants from the confounding factors of signal decay and injection dynamics, kinetic modelling allows for a more reliable and sensitive assessment of disease progression and response to therapy [37]. Together, these requirements motivate the hardware and acquisition developments described in the experimental chapters of this thesis. Given the rapidly decaying signal, the choice of RF flip angle represents a critical constraint on pulse sequence design. Each excitation pulse converts a portion of the available longitudinal magnetisation into measurable transverse magnetisation. Because the longitudinal pool cannot be replenished, every pulse consumes an irreplaceable fraction of the remaining signal. The magnetisation available at the n-th excitation can be approximated as[49]:

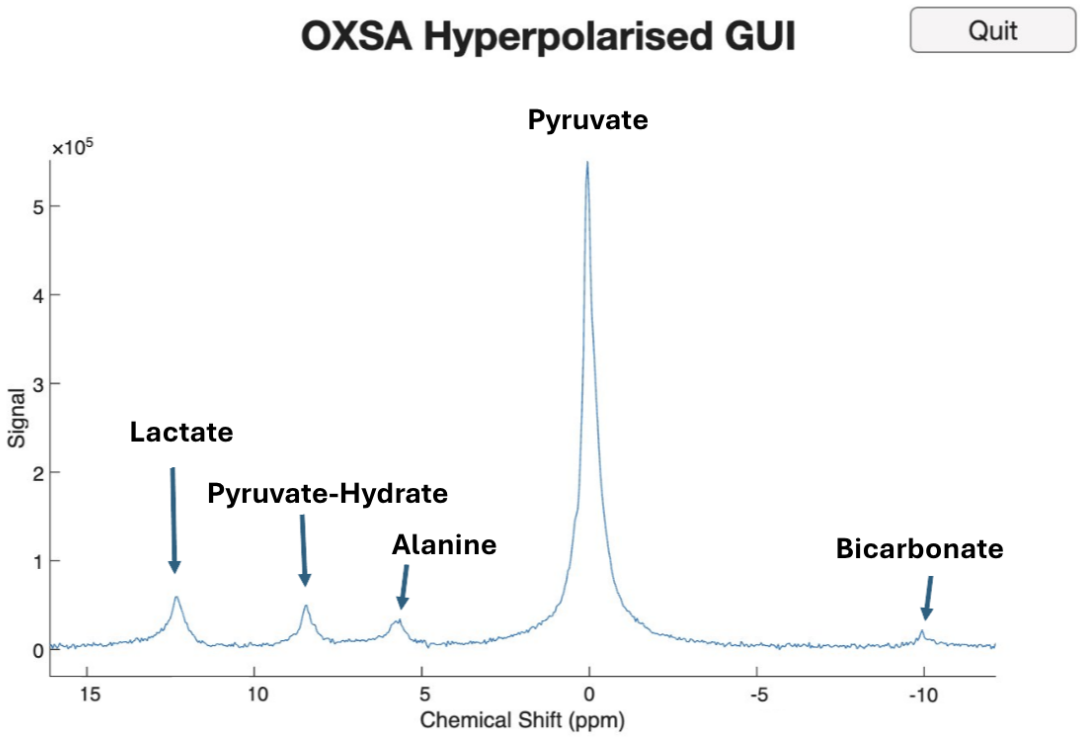


Figure 1.3: An example spectra from a representative rat heart acquired over 60 seconds. Metabolic peaks are labelled to show distinctive and well-separated peaks.

$$M_z(n) = M_z(n-1)\,\cos(\alpha)\,e^{-TR/T_1} \tag{1.16}$$

where α is the flip angle and TR is the repetition time between pulses. A large flip angle yields high instantaneous SNR but rapidly depletes the magnetisation available for subsequent time points. A small flip angle preserves signal across a longer dynamic acquisition but reduces the SNR at each individual time point. This trade-off must be explicitly accounted for in the quantitative analysis of metabolic conversion rates, since failure to correct for the RF-driven depletion of the parent substrate will systematically bias the estimated rate of conversion to downstream metabolites[44, 50]. Further discussion on the metabolic modelling is developed in the introduction of chapter 4.

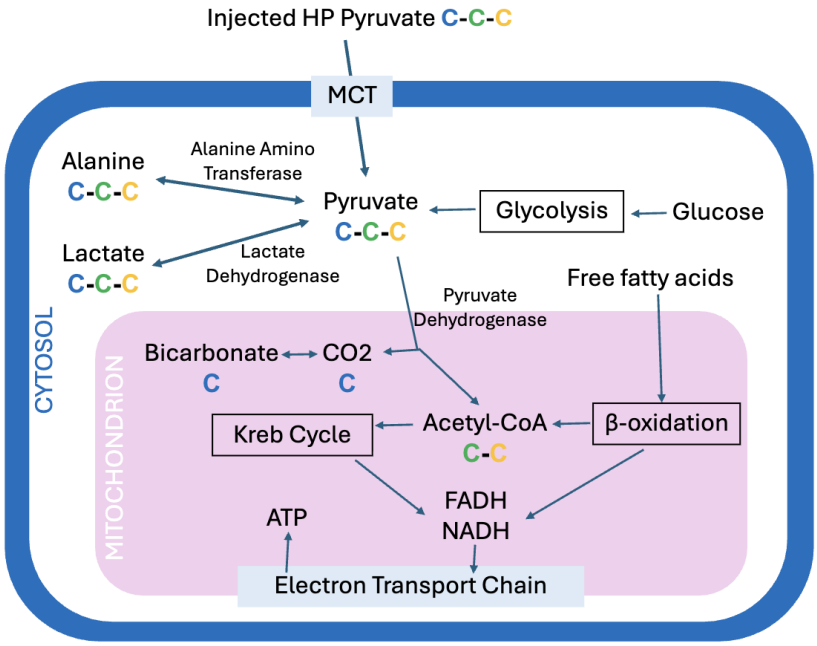


Figure 1.4: A schematic of the principal metabolic pathways of pyruvate.

1.2.5 Carbon-13 Metabolic Physiology

Physiology and Enzymatic Conversion

As introduced in Hyperpolarised Carbon-13 MRI section 1.2.4, pyruvate is an endogenous metabolite. In physiological conditions, this molecule sits at the junction of several cellular metabolic pathways, each governed by rapid enzymatic conversions [37, 51]. By labelling the first carbon position (C1) of the three-carbon pyruvate molecule with hyperpolarised carbon-13, denoted [1-^{13}C] pyruvate, the hyperpolarised carbon label is transferred to downstream metabolic products via these distinct conversion routes [44]. This carbon transfer enables the real-time observation of pathway-specific enzymatic activity *in vivo* [3]. The enzymatic pathways investigated in this work involve the conversion of pyruvate into lactate, alanine, and bicarbonate. In addition to these enzymatic products, pyruvate undergoes a spontaneous, non-enzymatic hydration to form pyruvate-hydrate. Although this reaction carries no direct metabolic significance [52], the pyruvate-hydrate resonance is readily observed in hyperpolarised carbon-13 spectra and are accounted for in the quantitative analysis presented in the following chapters.

Under normal physiological conditions in the cell, when glucose is metabolised via glycolysis, pyruvate is produced [53]. In a HP carbon-13 experiment, the [1-^{13}C] pyruvate is delivered to the tissue via the blood supply and is taken up into cells facilitated by monocarboxylate transporters (MCTs) [44, 54]. The most prominent conversion observed in HP ^{13}C studies is the exchange between [1-^{13}C] pyruvate and [1-^{13}C] lactate, mediated by the enzyme lactate dehydrogenase (LDH) [3, 51]. This reaction is regarded as a key biomarker for anaerobic metabolism and tumour aggressiveness[51, 53, 54]. Alternatively, [1-^{13}C] pyruvate may also be converted to [1-^{13}C] alanine via alanine amino transferase (ALT), a reaction that reflects the amino acid pool and nitrogen metabolism within the cell [51]. The third fate of [1-^{13}C] pyruvate is within the mitochondria, where the pyruvate dehydrogenase (PDH) complex facilitates irreversible decarboxylation of the substrate, yielding acetyl-CoA and carbon dioxide. The carbon-13 labelled $^{13}CO_2$ rapidly equilibrates with [1-^{13}C] bicarbonate. This pathway serves as a direct measure of aerobic carbohydrate oxidation and the entry of carbon into the Krebs cycle[48, 52]. Figure 1.4 shows a schematic of the main pathways of pyruvate as discussed with colour-coded carbons, where the blue "C" indicates the [1-^{13}C]. Each of these downstream metabolites is observable by hyperpolarised carbon-13 MRI and hence provides valuable information about the physiology *in vivo*.

2

Radio Frequency Coil Development for Hyperpolarised Carbon-13 MRI

Contents

2.1 Abstract 22
2.2 Introduction 23
2.2.1 Motivation 23
2.2.2 Advantages of Specialised Receive-Only Coils for HP MRI 25
2.2.3 Hardware Principles of RF Coils 27
2.2.4 The Multi-layer Coil Array Concept 28
2.2.5 Computational Design and Simulation 31
2.2.6 RF Coils Performance Metrics 34
2.3 Characterization of a Single Rx Element 39
2.3.1 LCR Circuit and Hardware Overview 39
2.3.2 Building and Testing the Circuit for the Carbon-13 Frequency 40
2.3.3 Tuning and Matching Results 41
2.3.4 Proton Trap Building and Testing 42
2.3.5 Preamplifier Testing 46
2.4 Geometric Decoupling of Two Rx Elements 51
2.4.1 S21 Scattering Parameters 51
2.4.2 Simulation Methods 51
2.4.3 Simulation Results 53
2.4.4 Bench Top Methods 55
2.4.5 Bench Top Results 57
2.4.6 Bench Top and Simulation Discussion 59
2.4.7 B1 Field Simulations 60
2.5 Simulations of Extended Array Designs 62
2.5.1 Trading Geometric Decoupling with Field Homogeneity 62

2.5.2 Two-Element Spacing Trade-Offs 63
2.5.3 Three- and Four- element Design Development 67
2.5.4 Human Phantom Simulations 73
2.5.5 Discussion 77
2.6 In-Bore Array Validation 79
2.6.1 Scanner Hardware Requirements 79
2.6.2 Multi-layer Array: Proton difference map 85
2.6.3 Multi-layer Array: B1 Field Maps 89
2.6.4 Benchmarking Against Commercial Carbon Coils 94
2.7 Appendix 105

2.1 Abstract

The clinical utility of HP ^{13}C MRI is fundamentally limited by the rapid, irreversible decay of the enhanced signal and the low sensitivity of the carbon-13 nucleus. Conventional surface coils often fail to provide a sufficient balance between peak local sensitivity and the uniform depth penetration required for imaging anatomies such as the heart. This chapter details the development of a novel, industry-driven receive-only (Rx) array composed of four multi-layered coil elements. The primary objective was to leverage the high sensitivity and improved depth penetration of the patented multi-layer architecture to create an array with a larger area of receive field (B_1^-) homogeneity, thereby reducing the sensitivity of metabolic data quality to precise coil positioning during clinical cardiac exams.

The four-element array was developed using an 11 cm centre-to-centre spacing, a configuration identified through electromagnetic simulations and bench-top scattering parameter (S_{21}) measurements to provide the optimal trade-off between inter-element decoupling and field uniformity. In-bore validation on a 3 T clinical scanner, using a phase-based Bloch-Siegert B_1^+ mapping protocol to decouple transmit inhomogeneities, confirmed that the multi-layer design significantly outperformed a commercial rigid chest array. Within a heart-sized field-of-view, the multi-layer array achieved a mean SNR of 14.68 ± 4.79, demonstrating superior performance compared to the commercial reference (12.96 ± 3.75, a 13% improvement). Furthermore, the multi-layer array provided more than twice the mean receive sensitivity at depth,

with its signal intensity at a 12 cm depth exceeding the commercial reference's peak performance at 6 cm. These results demonstrate that the multi-layer array provides a robust design for deep-tissue metabolic mapping, establishing a hardware innovation for the clinical translation of HP ^{13}C MRI.

2.2 Introduction

2.2.1 Motivation

In hospitals, Magnetic Resonance Imaging (MRI) remains a gold-standard diagnostic tool for non-invasively evaluating cardiovascular conditions, offering insights into anatomical structure and functional dynamics [55]. Although conventional proton MRI provides excellent anatomical and functional information, translating HP carbon-13 MRI into the clinic requires dedicated hardware capable of detecting the substantially weaker, unrecoverable, and rapidly decaying metabolic signal. Given the signal constraints established in Section 1.2.4, the use of high-sensitivity carbon-13 RF hardware is necessary to capture metabolic flux before signal depletion. Currently available ^{13}C surface coils, whether in the form of simple loops or small loop arrays, present shortcomings for cardiac imaging. Notably, these coils often struggle with the balance between depth penetration and even field distribution across the chest imaging volume, which results in spatially variable sensitivity and inconsistent image quality [22, 56, 57].

Figure 2.1 shows an existing coil developed by PulseTeq Ltd, for heart imaging. It is a carbon-only dual loop with a 20cm diameter. The intended set up is shown in figure 2.1, where one loop is placed on top of the chest and the other is positioned symmetrically below the torso as a Helmholtz pair. This configuration produces a highly uniform field at the centre of the coils and is highly specialised in capturing the heart volume. However, if the coil is not positioned accurately relative to the heart then performance is poor due to the rapid spatial variation of the RF field profile (Figure 2.2). The trade-off between sensitivity and field uniformity is a fundamental challenge in RF coil design [58, 59], arising from the rapid spatial decay of the B_1 field produced by a current loop, as described by the Biot-Savart



Figure 2.1: The set up for the imaging protocol including the coil, phantom and body loader are shown for the two loop coil from PulseTeq.

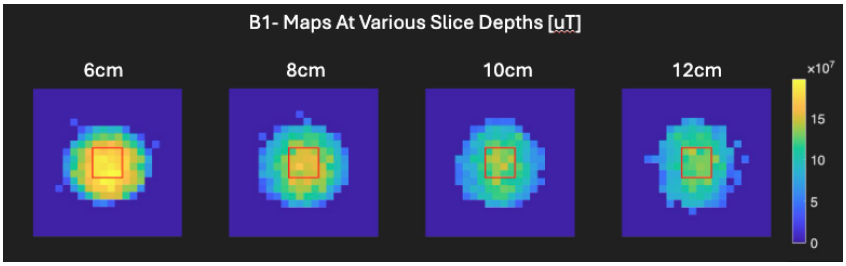


Figure 2.2: Receive B_1 profiles at various depths from the top of the phantom as seen in the set up figure (2.1). At 2x2x2cm voxel resolution, the images show strong localisation of receive field at the centre around the heart ROI (delineated with red squares) with a quick drop off of signal as the depth increases to 12cm.

law. While small surface coils provide strong local sensitivity (high B_1 field near the coil), they lack the uniform field distribution required for whole-organ coverage. Conversely, increasing uniformity typically reduces sensitivity, particularly at depth.

With these shortcomings in mind, I aimed to design a better human torso coil that could support the clinical translation of ^{13}C HP MRI. My goal for this project was to develop a coil that could achieve both 1) the high sensitivity seen in the heart volume offered by the current design and 2) a larger area of homogeneity to make the quality of the acquired data less sensitive to coil positioning.

Recent innovations have shown that multi-layer coil elements (figure 2.3), due to their multiple conductor layers, outperform simple loops in terms of sensitivity and depth penetration, thus presenting a promising avenue for technical improvement[60]. The multi-layer element is looped several times on one channel and has higher

sensitivities with improved penetration depth. However, while the depth sensitivity is excellent, the single-element design does not provide full coverage across the chest, an important consideration for cardiac imaging.

To address this clinical gap, I aimed to develop a cardiac array comprised of multi-layer RF elements. This array design aimed to leverage the high depth sensitivity inherent to the individual multi-layer elements within an array configuration to enhance B_1 homogeneity and provide robust spatial coverage for HP ^{13}C metabolic mapping. In collaboration with PulseTeq Ltd., I investigated whether integrating these high-sensitivity elements within a geometrically decoupled framework could optimise both signal strength and spatial uniformity, thereby enhancing the clinical utility of HP ^{13}C MRI for metabolic imaging.

This chapter details the development and evaluation of a receive-only (Rx) array designed for integration with a commercial PulseTeq transmit (Tx) coil. While my primary focus was the optimisation of the Rx component, I include a characterisation of the Tx field to account for transmit inhomogeneities within the acquired signals. This chapter is structured as follows: I first describe the design rationale of the multi-layer RF coil array, including considerations of hardware requirements, anatomical coverage, and decoupling. Next, I present electromagnetic simulations and bench top characterisation results that informed the final coil configuration. I then detail the in-bore scanner evaluation, covering hardware integration, signal correction methodologies for B_1^- mapping, and a comparative analysis against a commercial ^{13}C coil. Finally, I discuss the design trade-offs and future directions for clinical translation of the multi-layer array.

2.2.2 Advantages of Specialised Receive-Only Coils for HP MRI

Given the non-renewable nature of the HP signal introduced in the Motivation section 2.2.1, maximising receive sensitivity is one of the most critical determinants of image quality [1, 61, 62]. In this project, I focus on developing a receive coil design while adapting a pre-existing transmit coil made by my industrial collaborator.

This is because Rx-only coils offer advantages by decoupling the requirements of transmission and reception, hence reducing the need for balanced power handling, RF safety limits, and transmit performance metrics such as large scale B_1^+ field uniformity. Rx-only is free of these balancing acts and can be optimised for enhanced B_1^- sensitivity in the region of interest [63]. For example, Rx coils can be shaped and positioned to closely conform to the target anatomy. The physical proximity minimises the coil-to-sample distance, maximising the detectable MR signal [22]. Such geometrical adaptability is particularly valuable in translational imaging, where the subject size and anatomy seen in the clinic can significantly vary as opposed to preclinical conditions using standardised test objects.

From a hardware engineering perspective, eliminating the transmit circuitry simplifies the coil's architecture, making bench top testing of the Rx components lightweight. Without the requirement to withstand high-voltage RF pulses, the design can incorporate low-noise components and optimised matching networks. This reduces system noise, improves preamplifier performance, and facilitates tighter integration between coil elements and their respective receivers [61, 63].

In addition to performance, Rx-only designs offer safety advantages. Because they are not used to transmit RF power, they are not directly responsible for generating patient specific absorption rate (SAR), an important consideration for clinical translation of HP ^{13}C studies [22, 64]. However, Rx-only coils must still be adequately decoupled and detuned during RF transmission to prevent induced currents from the transmit coil, which could otherwise lead to localised RF fields, excessive heating, or hardware damage.

From a workflow perspective, the modular nature of separate Tx and Rx systems enhances operational flexibility. A static Tx coil can remain in place while interchangeable Rx arrays are used to accommodate different anatomical targets. This modularity is especially useful in studies investigating multiple organs or applications, where imaging needs can vary significantly between scans, and for imaging protocols requiring quick repositioning[27, 63, 64].

Furthermore, Rx-only systems simplify active detuning and decoupling strategies. Since Rx coils are not exposed to high transmit voltages, the components used to isolate the coil during transmission can be of lower voltage tolerance and reduced complexity, improving overall reliability. Placing the low-noise preamplifiers close to each coil element becomes more feasible and minimises signal losses along the cable, which both improve the overall system noise.

Specialised Rx-only coils are ideally suited for HP MRI applications where the non-renewable and rapidly decaying nature of the HP signal places demands on receive sensitivity. This separation of Tx from Rx functions is already a standard part of clinical proton MRI protocols, where the integrated body coil serves as a uniform Tx source, while specialised surface or array coils act as sensitive receivers. This established workflow ensures optimal performance and extends to HP MRI.

2.2.3 Hardware Principles of RF Coils

All RF coils share basic electronic features, and are based on an LCR circuit, as detailed in the introductory section 1.2.3. In practice, the tuning and matching of an RF coil was one of the most time consuming aspects of this project, and is highly sensitive to the presence of a sample, a phenomenon known as loading. When a subject or a phantom is placed within the coil's field of view, it introduces additional resistance (R_{sample}) and modifies the effective capacitance of the system. This interaction is dictated by the dielectric properties of the load such as its conductivity (σ), which leads to inductive losses through the generation of eddy currents, and its relative permittivity (ϵ_r), which can cause frequency shifts due to capacitive coupling between the coil and the sample.

In the context of phantom studies, these properties are made to mimic biological tissue. A high-conductivity phantom will load the coil more heavily, lowering its quality factor and shifting its impedance away from the $50\,\Omega$ benchmark. Consequently, the matching network must be robust to compensate for varying load sizes and dielectric compositions, ensuring the system remains operational for the varying clinical conditions it may encounter.

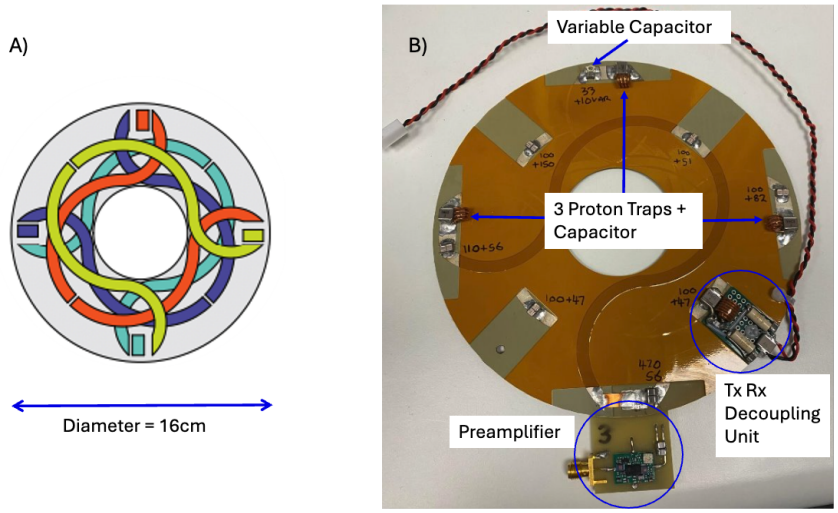


Figure 2.3: (a) As patented in [60], a diagram of the multi-layer coil element illustrating the four stacked conductor loops with an omega-like shape. (b) A photo of the PCB, which houses the multi-layer trace as seen in (a), and further components required for scanner integration.

When Rx coils are used in the presence of a Tx coil, detuning circuits such as PIN diodes are incorporated to disable the Rx pathway during transmission. Active decoupling requires creating an open circuit when engaged. In other words, when the PIN diode is **ON**, the Rx channel is disabled: an inductor becomes parallel to the tuning capacitors, blocking current flow and creating an open circuit in the receive path. This is essential because the transmit power levels are significantly higher than that of the receive and can damage sensitive components in the Rx chain such as the preamplifier [65]. An excellent step-by-step guide for designing and engineering Rx coil components is provided by Fujita et al [30].

2.2.4 The Multi-layer Coil Array Concept

To address the sensitivity-uniformity trade-off described in Section 2.2.1, the multi-layer concept was developed to combine the benefits of strong localised sensitivity with broader spatial coverage [60]. The inspiration for this design stemmed from earlier investigations into petal-shaped resonators and concentric multi-turn loop designs as discussed by Zhou et al. [60, Reference Chapter 4.5.2]. These geometries informed the development of the multi-layer coil element, which consists of four stacked conductor loops with an omega-like shape. These individual

loops are overlapped with a 90-degree rotational offset to reduce coupling and to optimise coverage of the imaging volume [60, Reference Chapter 5.2.1–5.2.2] (shown in figure 2.3A).

Compared to traditional single-loop surface coils and conventional phased array coils, the multi-layer coil demonstrates several improvements. It achieves higher B_1 field efficiency, deeper field penetration, and elevated quality factors (Q-values), all of which contribute to an improved signal-to-noise ratio (SNR) as shown in the literature [60] (reproduced in figure 2.4). In this reference[60], Zhou et al. demonstrated an approximately 1.5-fold increase in SNR compared to conventional surface coils (reproduced in figure 2.5). The figures are reproduced from Zhou et al. with permission from the author.

Building on this development, I designed an array consisting of these muti-layer elements to expand the field of view while preserving the peak performance of the individual multi-layer coil geometry so as to encompass the heart and surrounding anatomy. I drew on the foundational theory introduced by Roemer et al. for optimal signal combination in RF coil arrays [66]. The signal from each coil element S_i is combined using a weighted sum to generate a composite signal:

$$S_{\text{combined}} = \sum_i w_i S_i \tag{2.1}$$

The optimal weights w_i are chosen to maximise the SNR and are calculated by accounting for both the sensitivity profile of each coil and the noise covariance between elements. Most importantly, this additive nature of signal contribution from decoupled coil elements underpins the benefit of array configurations: multiple independently sensitive coils can each contribute signal from different regions, thereby increasing the net SNR and extending the useful FOV [57, 67].

An additional benefit of the multi-layer architecture is its ability to support anatomically conformal arrays. By tailoring the flexible coil layout to follow the shape of the torso, it becomes possible to achieve more uniform signal sensitivity across the imaging region. This is particularly advantageous in cardiac or abdominal imaging, where anatomical variability and motion complicate coil performance.

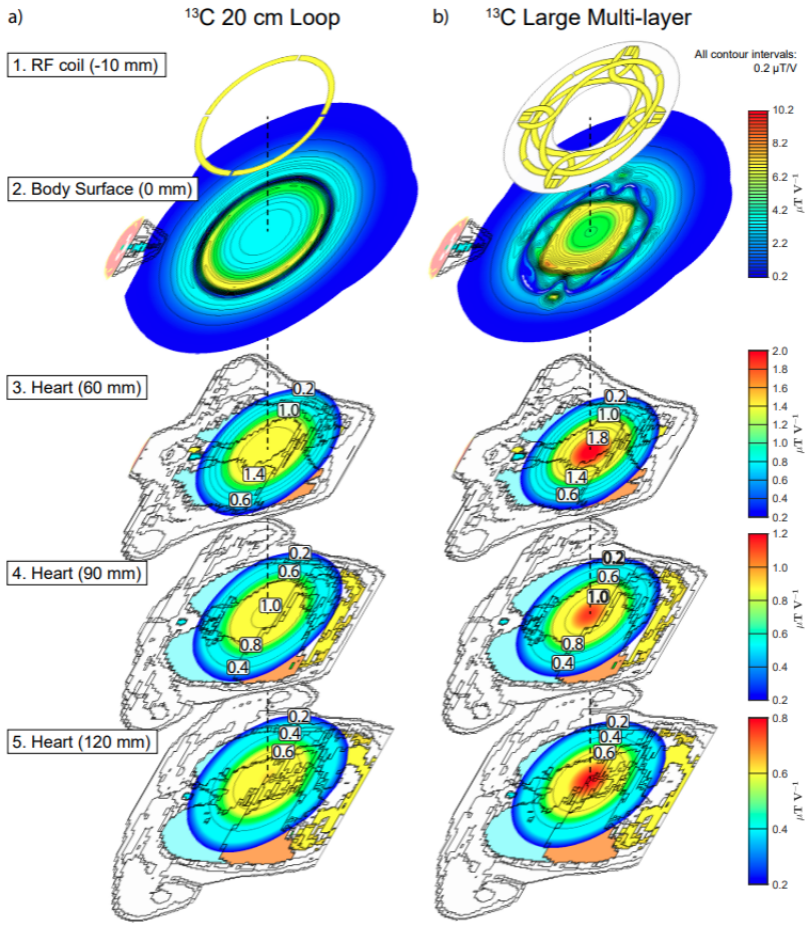


Figure 2.4: 2D simulated B1+ of 20 cm loop (a) and large multi-layer coils (b) at four depths corresponding to the representation in Figure 5.4: 0 mm (body surface), 60 mm (top of heart), 90 mm (middle of heart), and 120 mm (bottom of heart), with 0.2 µT V-1 contours. Reproduced from Zhou et al.[60]

Effective implementation of the multi-layer array requires addressing the issue of inter-element coupling. This is commonly managed through geometric overlap to minimise mutual inductance, which serves to isolate each channel from its neighbours. This geometric decoupling strategy helps ensure that each coil element contributes additively and distinctly to the final image. Together, these features make the multi-layer coil array a promising solution for the next generation of receive coil technology in HP MRI.

In this thesis chapter, I incorporate EM simulations to assess B_1 field distributions and scattering parameters (S-parameters) to quantify coil efficiency. These simulations were complemented by bench top testing using a vector network

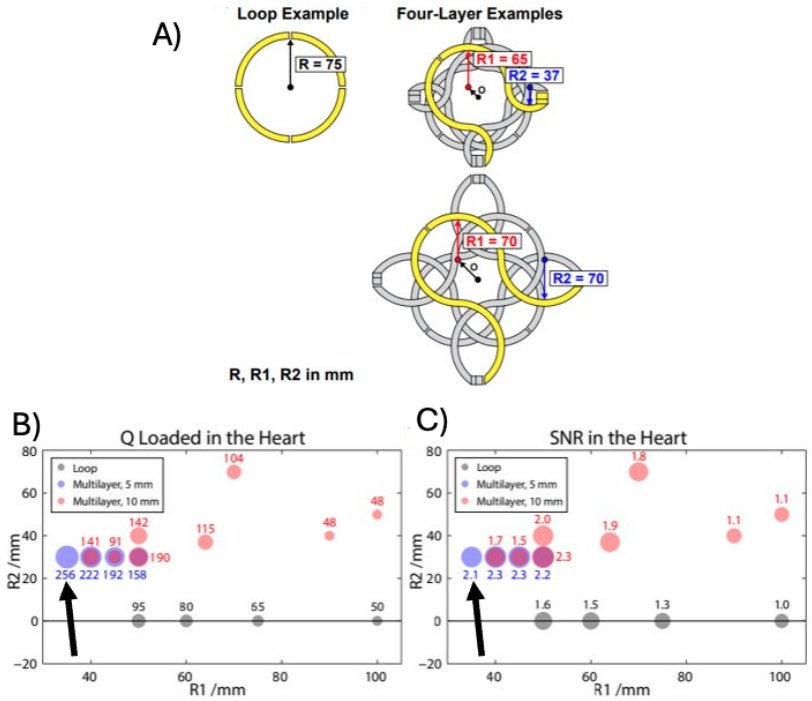


Figure 2.5: (A) Parameterised multi-layer designs. (a) Conventional loop with R (radius) parameter, with example of R = 75mm. (b) Four-layer multi-layer designs with R1 (inner-most radius) and R2 (outter-most radius), with examples of R1 = 65mm, R2 =37 mm and R1 = 70mm, R2 = 70mm. Bubble plots showing (B) QLoaded and (C) SNR values (represented by area) for four layer multi-layer designs of varying R1 (x-axis), R2 (y-axis) and width of 5 and 10 mm, and R of conventional circular loops plotted along the x-axis (shown in black). The multi-layer element dimension used in this thesis is R1 = 35mm and R2 = 30mm and is the blue dot pointed out with the black arrows (QLoaded = 256, SNR = 2.1). Here, the simple loop with comparable radius size is R1 = 50, 60, and 75, where QLoaded = 95, 80, 65 and SNR = 1.6, 1.5, 1.3, respectively. Reproduced from Zhou et al.[60]

analyser (VNA: ZNLE3, Rohde & Schwarz GmbH & Co. KG, Munich, Germany; frequency range 1 MHz–3 GHz) to measure the S-parameters (e.g., S_{11}, S_{21}), which reflect the coil's transmission and reflection characteristics. This two-pronged validation approach, simulation and experiment, supported the development of this new HP RF coil array system.

2.2.5 Computational Design and Simulation

To accelerate the design process and minimise reliance on time-consuming hardware iterations, I adopted an electromagnetic (EM) simulation-led workflow using CST Microwave Studio. This approach enabled the rapid prototyping of complex geome-

tries under biology-mimicking loading conditions, and is commonly employed in industrial RF engineering for coil design validation [60, Chapter 3.2.2]. Simulations incorporated anatomical voxel models and frequency-specific dielectric constants to replicate physiological conditions. Coil geometries were aligned with the B_0 field, and discrete ports were used to simulate tuning and matching networks [60, Chapter 3.2.2].

The CST software package solves Maxwell's equations numerically by discretizing the geometry into a mesh and applying either time-domain or frequency-domain methods. The choice of solver and mesh type directly impacts the simulation's resolution and computational load. In frequency-domain simulations, the finite element method (FEM) is commonly used with unstructured tetrahedral meshes, which are well suited to curved and complex geometries. In contrast, time-domain simulations typically employ finite-difference time-domain (FDTD) methods with structured hexahedral meshes, which are more efficient for voxel-based anatomical models but less accurate at representing smooth surfaces or non-planar conductors [60, Chapter 2.1.2].

The mesh strategy selection in my work was guided by the specific geometric and accuracy requirements of the multi-layer coil designs. Tetrahedral meshing was used to resolve the fine structure of the curved, layered conductors in frequency-domain simulations, while hexahedral meshing was evaluated primarily for anatomical human phantom models.

To evaluate the element performance metrics, field uniformity, and decoupling, I performed systematic parameter sweeps in CST. S-parameters were calculated to evaluate tuning (S_{11}) and inter-element coupling (S_{21}), and these were incorporated into CST's circuit co-simulation environment to generate field maps reflecting realistic electrical loading conditions [60, Chapter 3.2.2.2, 5.4.1.2]. Simulation runtime was dependent on phantom selection; dielectric properties of fat, muscle, and whole-body models (HUGO voxel-based human model details found in [68]) influenced coil loading and computational demand. A full summary of phantom characteristics is provided in Zhou's appendix [60, Chapter 3.2.2.2, 5.4.1.2].

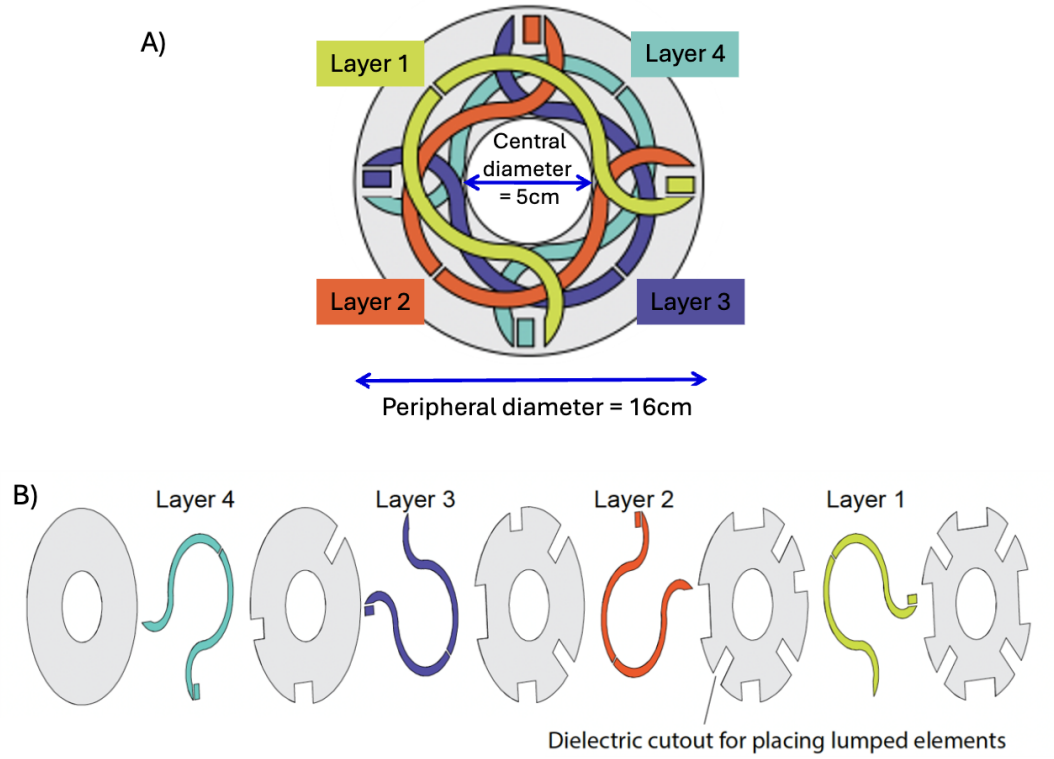


Figure 2.6: A) Diagram of the multi-layer design showing the dimensions and positions of each conductive Omega-shaped layer. The conductive trace is 5mm wide. B) The expanded stack of alternating dielectric and conductive layers. Adapted from figure 5.2 in reference [60]

Simulation results were used to benchmark expected coil performance and guided final design selection. In my work, agreement between CST results and bench top measurements confirmed the viability of simulation as a surrogate for early prototyping [60, Ch. 5.4.2.2]. This enabled iterative refinement of coil geometry and configuration while maintaining engineering control, minimising fabrication cycles, and keeping hardware costs low.

Design parameters are based on prior work with the single-element multi-layer configuration set at central radius = 5cm, peripheral radius = 16cm, conductor trace width = 5mm, and number of layers $N = 4$ [60, Chapter 5.3.2], which is shown in figure 2.6. Field maps and uniformity quantification were used to assess performance across the design space [60, Chapter 5.3.2.1].

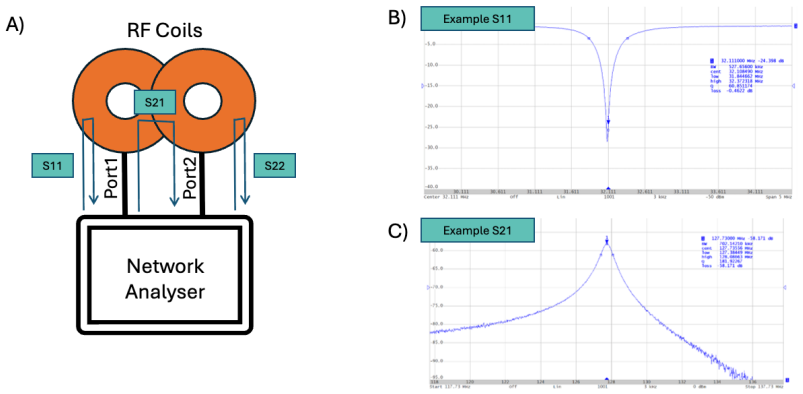


Figure 2.7: A) Diagram of two RF elements connected to a two-port vector network analyser showing the set ups to collect S11, S21, and S22 transmission and reflection coefficients. An example parameter sweep for a (B) S11 and (C) S21 measurement is shown.

2.2.6 RF Coils Performance Metrics

Decoupling Between Array Elements

To ensure minimal electromagnetic interaction that could degrade performance, decoupling between array elements is typically quantified via S-parameters. These parameters monitor performance at desired resonant frequencies while identifying potential interference at undesired ones. While measured as complex impedance on a Smith chart, these values are typically presented as a magnitude in decibels (dB) using the following conversions:

$$S_{ii}(dB) = 20\log_{10}|S_{ii}| \quad \text{and} \quad S_{ij}(dB) = 20\log_{10}|S_{ij}|$$

As illustrated in Figure 2.7A, the reflection coefficients (S_{11},S_{21} and S_{22}) quantify the ratio of reflected to incident power. Ideally, the S_{11} is minimised at the frequency of interest so as to minimise the reflected power, and maximised away from the frequencies of interest. The width of the resonance dip determines the frequency range over which the coil is well matched. A narrow resonance provides excellent matching at the target frequency but may not fully encompass the spectral range of interest, whereas a broader resonance accommodates a wider range of frequencies at the expense of optimal matching. A common benchmark in RF coil construction is to achieve an S_{11} of $-15\,\text{dB}$ or lower at the target frequency, corresponding to

a reflected power of $\leq 3\%$ (example S_{11} measurement in Figure 2.7B). I achieved this benchmark for all RF coils constructed throughout this project. The forward transmission coefficient (S_{21}) measures power transfer between ports. For the array elements, I used S_{21} to evaluate the decoupling between neighbouring elements, adhering to a benchmark of $-15\,\text{dB}$ or lower within the transmit bandwidth to ensure minimal inter-element interference (example S_{21} measurement in Figure 2.7B).

In array coil designs, effective decoupling between adjacent elements is essential to preserve image quality and maximise signal-to-noise ratio (SNR). Without proper decoupling, mutual inductance between neighbouring elements can distort sensitivity profiles, reduce signal uniformity, and introduce artifacts in reconstructed images. Coupling also redistributes current and power among elements, impairing the efficiency of both transmit and receive functions. In the context of clinical applications such as accelerated parallel imaging, these techniques rely on distinct sensitivity profiles for accurate reconstruction [69, 70]. Therefore, insufficient decoupling degrades accelerated parallel imaging performance.

The effectiveness of decoupling is assessed with S_{21} measurements, where lower values indicate minimal cross-talk. Strategies include geometric overlap of neighbouring loops to cancel mutual inductance or simply increasing spacing between elements. Preamplifier decoupling can suppress induced currents. Active detuning using PIN diodes with passive networks isolate receive loops during transmit. Cable traps and baluns suppress common-mode currents along coaxial cables. And transformer/inductive networks can be incorporated to cancel mutual fields. In practice, hybrid solutions are often required, particularly in space-constrained geometries and *in vivo* settings where subject-specific dielectric properties make loading and matching variable. In this work, geometric decoupling was the primary strategy used in the custom multi-layer arrays. A gridding optimisation systematically varied inter-element overlap to minimise coupling while preserving sensitivity and coverage. This array-level isolation underpins the dual-nuclear (carbon and proton) operation addressed next.

Carbon-Proton Decoupling

When ^{13}C receive arrays operate alongside ^{1}H transmit/receive hardware, inter-nuclear isolation is as important as inter-coil interactions. Practical measures include ^{1}H band-stop (LC) traps on ^{13}C elements, ^{13}C band-stop traps on ^{1}H structures when required, active detuning of ^{13}C loops during ^{1}H transmission, and use of baluns/chokes and filtering to suppress common-mode pathways. These steps minimise detuning or heating of ^{13}C circuits by ^{1}H RF and prevent the ^{13}C hardware from perturbing the ^{1}H B_1 field. With both inter-element and inter-nuclear coupling controlled, the spatial performance of the coil can be characterised.

Field Mapping: Transmit and Receive B_1 Uniformity

The coil's transmit field uniformity (B_1^+) and receive sensitivity (B_1^-) are measured using dedicated field-mapping techniques. These maps verify that excitation is consistent and that B_1 field profiles support accurate signal detection across the region of interest. Broadly, the established methodologies for B_1 field mapping can be categorised into magnitude-based and phase-based approaches, each with trade-offs regarding scan time and sensitivity to experimental conditions. The magnitude-based Variable Flip Angle (VFA) method estimates B_1^+ by acquiring two or more images with different excitation angles and calculating the field from the signal ratio. While straightforward, VFA is highly sensitive to T_1 relaxation times and slice-profile inaccuracies. Alternatively, the Sinusoidal CSI method, also a magnitude-based approach often used in spectroscopic imaging, acquires a series of CSI datasets using progressively varied transmit amplitudes. Because the measured signal follows a sinusoidal dependence on the actual flip angle, fitting this signal enables estimation of the local B_1^+ field. However, this typically requires long acquisition times and is prone to errors from B_0 inhomogeneity.

For this project, I utilized the Bloch-Siegert shift method, as it offers several advantages for ^{13}C coil characterization. Unlike magnitude-based methods, the Bloch-Siegert approach is a phase-based technique that is largely independent of T_1 relaxation, pulse repetition time, and underlying spin density [71, 72]. The method

applies two off-resonance RF pulses symmetrically positioned around the Larmor frequency, at $\omega_0 \pm \omega_{\mathrm{RF}}$. These pulses induce a frequency shift in the transverse magnetisation, creating the Bloch–Siegert shift (ω_{BS}), which is proportional to the square of the RF field amplitude:

$$\omega_{\mathrm{BS}} = \frac{(\gamma B_1)^2}{2\omega_{\mathrm{RF}}} \tag{2.2}$$

where γ is the gyromagnetic ratio, B_1 is the RF amplitude, and ω_{RF} is the off-resonance frequency. This shift accumulates as a phase change over the pulse duration T:

$$\Phi_{\mathrm{BS}} = \int_0^T \omega_{\mathrm{BS}}(t)\, dt \tag{2.3}$$

Assuming a known RF pulse shape, the phase shift can be expressed in terms of the peak RF amplitude:

$$\Phi_{\mathrm{BS}} = \left(B^+_{1,\mathrm{peak}}\right)^2 \int_0^T \frac{\left(\gamma B_{1,\mathrm{norm}}(t)\right)^2}{2\omega_{\mathrm{RF}}(t)}\, dt \tag{2.4}$$

Unlike amplitude-based methods, the Bloch Siegert technique is independent of relaxation effects and is resilient to field inhomogeneities and chemical shift, making it particularly attractive for transmit calibration in both clinical and research settings. This method requires relatively high RF power to achieve measurable phase shifts, which was attainable in this experiment.

In this project, the B_1^+ estimation was important in understanding the receive profile of the coil. During CST simulations, the B_1 transmit and B_1 receive are assumed to be complex conjugates of each other because the dual Tx and Rx is modelled in a perfect B_0 field. In the experimental setting, my Tx and Rx consisted of separate hardware placed in an imperfect B_0 field and with inhomogeneous B_1 transmit. This means that the images from the scanner reflect both the B_1 transmit and receive fields, requiring a correction map to extract the B_1^- map. In addition to this, it is important to perform a transmit gain (TG) calibration before each ^{13}C exam to set the RF power for the prescribed flip angle under the current loading,

thereby ensuring accurate B_1^+ and consistent excitation. Establishing the transmit field behaviour provides the basis for decoupling the sensitivity profile.

The quality factor Q, defined as the ratio of stored to dissipated energy per cycle, summarises electrical efficiency in coil design [23]. The unloaded Q is measured in isolation and the loaded Q is measured in the presence of a conductive sample. A higher loaded-to-unloaded ratio indicates greater sample-noise dominance, which is desirable for maximising SNR. This quantity was characterised by Zhou et al. for the multi-layer element (Figure 2.5A). Given the fixed coil geometry and the priority of in-bore validation within the available scanner time, the B_1^- receive profile and SNR benchmarking against a commercial reference were the more clinically informative metrics at this stage; independent Q measurements of the array are identified as a target for future work.

SNR, a central imaging metric, is typically evaluated with phantom studies or *in vivo* by acquiring signal images and noise scans. Because SNR is shaped by both coil sensitivity (B_1^-) and noise correlation between elements, the preceding decoupling and field-mapping steps directly influence the measured performance. Beyond mean SNR, the receive-channel noise covariance matrix, from a noise-only acquisition (no RF excitation, gradients off), quantifies inter-channel noise correlation with an identity matrix indicating perfect decoupling. In this thesis, I mainly focus on the proportionality of signal to noise within the receive B_1 field.

Together, these metrics provide a coherent framework for assessing and optimising RF coil performance. By first ensuring decoupling (element-to-element and dual-nuclear), then validating B_1 fields, and finally quantifying signal homogeneity, the design process targets maximum sensitivity over the desired field-of-view with the goal of optimal coil design for high-quality MRI [30].

2.3 Characterization of a Single Rx Element

2.3.1 LCR Circuit and Hardware Overview

Five multi-layer coil elements were constructed in preparation for the development and construction of the multi-layer array. As a contingency plan for mechanical hardware, I fabricated an extra element for this work. Each element was carefully constructed with various components required for the final unit. This section is an overview of the individual components on the single multi-layer elements and the testing methodology.

Briefly, the single multi-layer element is made up of a large flexible printed circuit board (PCB) with the patented multi-layer trace design, capacitors, proton traps, preamplifier, and a Tx Rx detuning matrix board. A representative PCB with the capacitors, proton traps, disconnected preamplifier, and detuning unit that I built for this project is shown in figure 2.8 (coils 2-5 are shown in the appendix figures 2.82, 2.83, 2.84, 2.85.) The effective use of proton traps allows for better signal detection at the carbon frequency [73], and I positioned three compact solenoids symmetrically at the twelve, three, and nine-o'clock positions around the coil loop to provide uniform proton frequency suppression around the whole coil element. It was critical for the solenoids to be as small as possible to allow adjacent array elements to overlap with minimal vertical separation (Figure 2.9A). I used tuning and matching capacitors (33pF to 470pF +- 5% Non-Magnetic Surface Mount Chip NP0 Capacitor 500 Volts, Passive Plus Inc. New York), where the matching capacitor was at the bottom, placed above the preamplifier, and all other capacitors filling the gaps in the copper trace were tuning capacitors. A single capacitor was set as a variable capacitor (shown figure 2.9B) for flexibility to fine tune during experiments. The variable capacitor is much more expensive than the other fixed capacitors so was minimally used, and placed strategically. Figure 2.9C is a close-up of the detuning unit, which operates to decouple the Tx and Rx elements during scanning in order to protect the receive chain from

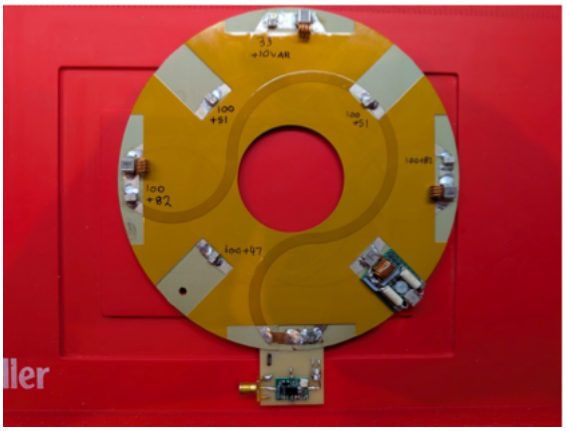


Figure 2.8: Representative receive element (Coil 1) assembly. The PCB incorporates conductive trace gaps bridged by capacitors (values labelled in pF), along with ^{1}H traps, detuning circuitry, and the preamplifier connection point.

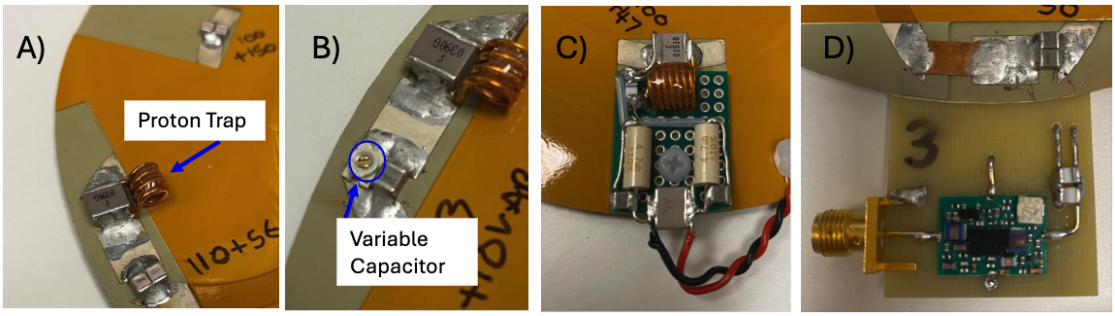


Figure 2.9: Coil elements close ups. A) Proton trap B) Variable capacitor C) Tx Rx decoupling unit D) Preamplifier

high throughput power used by the transmit chain. Each element incorporated a preamplifier (WMA3RA-R1, WanTcom Inc.) (Figure 2.9D).

2.3.2 Building and Testing the Circuit for the Carbon-13 Frequency

To characterise a single receive (Rx) coil element, I soldered capacitors using a temperature-controlled iron set between 280-300 °C. I selected this temperature range to avoid thermal damage to the capacitors. Learning to solder precisely and reliably was a significant component of my technical training during the first year of the DPhil. Early coil construction efforts required iterative practice, attention to thermal tolerances, and understanding of circuit behaviour. This hands-on experience proved essential, not only for successful bench top characterisation, but

also for building confidence in the broader process of hardware engineering.

After building a closed circuit, I tuned and matched my coil elements to 32.11 MHz, the Larmor frequency of ^{13}C at 3 T for the GE Premier system. I took S-parameter measurements using a vector network analyser (VNA), with port 1 connected to the coil to acquire the reflection coefficient, S_{11}. After soldering the coaxial cable across the matching capacitor and connecting the VNA, I measured an initial S_{11} sweep to assess whether the coil was resonating near the target frequency. The coils were placed on a phantom (Pulseteq phantom T-7100-0173). A well-tuned and matched coil exhibited a resonance dip at 32.11 MHz, with $S_{11} < -15$ dB, indicating minimal reflected power and efficient energy transfer to the receive chain.

Capacitor values were iteratively modified. Matching was evaluated using both the magnitude of S_{11} and the Smith chart representation, which provided insight into the resistive and reactive components of the input impedance over frequency. Once optimal tuning and matching was achieved, the final S_{11} curve and Smith chart were recorded for documentation.

2.3.3 Tuning and Matching Results

The final tuned and matched coils exhibited resonance centred at the GE ^{13}C frequency of 32.11 MHz. The S_{11} magnitude at the carbon frequency was measured to be $S_{11,coil1} = -32.51$ dB (Figure 2.10), $S_{11,coil2} = -25.09$ dB(Figure 2.11), $S_{11,coil3} = -24.40$ dB (Figure 2.12), $S_{11,coil4} = -29.48$ dB (Figure 2.13), $S_{11,coil5} = -33.09$ dB (Figure 2.14). These values indicated good impedance matching and minimal power reflection as they satisfied the target threshold of $S_{11} < -15$ dB.

Overall, the resonance dip remained consistently centred at 32.11 MHz, with slight variation depending on the element positioning and contact quality with the phantom. For ^{13}C spectroscopy, where multiple peaks may span several hundred Hz, a moderate bandwidth is essential. To quantify the trade-off between spectral bandwidth and sensitivity, the loaded quality factor (Q_L) provides a good metric :

$$Q_L = \frac{f_0}{\Delta f},$$

where $f_0 = 32.11\,\text{MHz}$ and Δf is the full bandwidth measured between the upper and lower -3dB frequencies. The resulting Q_L values of $Q_{L,coil1} = 63.64$ (figure 2.10), $Q_{L,coil2} = 63.76$(figure 2.11), $Q_{L,coil3} = 60.85$(figure 2.12), $Q_{L,coil4} = 65.27$ (figure 2.13), $Q_{L,coil5} = 65.47$ (figure 2.14). These Q_L values correspond to -3 dB bandwidths of approximately 0.49 MHz to 0.53 MHz, exceeding the frequency span required to encompass the resonances of pyruvate and its downstream metabolites at 3 T while achieving sufficient sensitivity.

These tuning and matching results confirmed the performance of the coil and its suitability for integration into the array, and ready for subsequent multi-channel array characterisation.

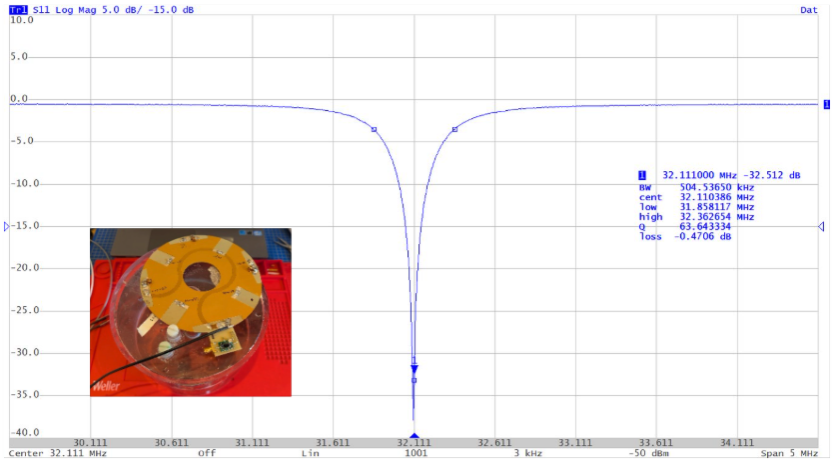


Figure 2.10: S-parameters from Rx Coil 1. Set up with the Pulseteq phantom T-7100-0173 shown on the left-hand side.

2.3.4 Proton Trap Building and Testing

In a ^{13}C receive coil, proton (^{1}H) traps are essential for good signal quality. This is because the ^{1}H hardware, often the body coil, is active during scan sessions for anatomical reference imaging, scanner triggering, and system synchronization. These high-power RF pulses at the proton frequency (127.73 MHz in accordance to 3T GE scanner requirements) can couple with the ^{13}C coil even if it is not resonant at that ^{1}H frequency. The coupling can degrade the transmit efficiency of the proton system;

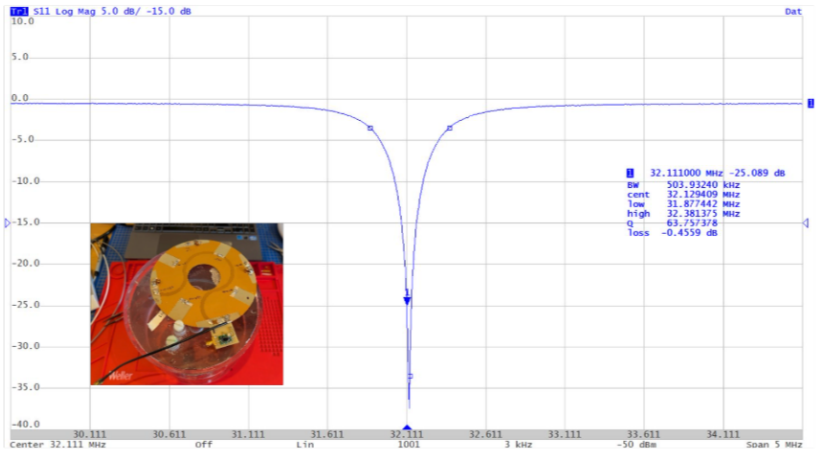


Figure 2.11: S-parameters from Rx Coil 2

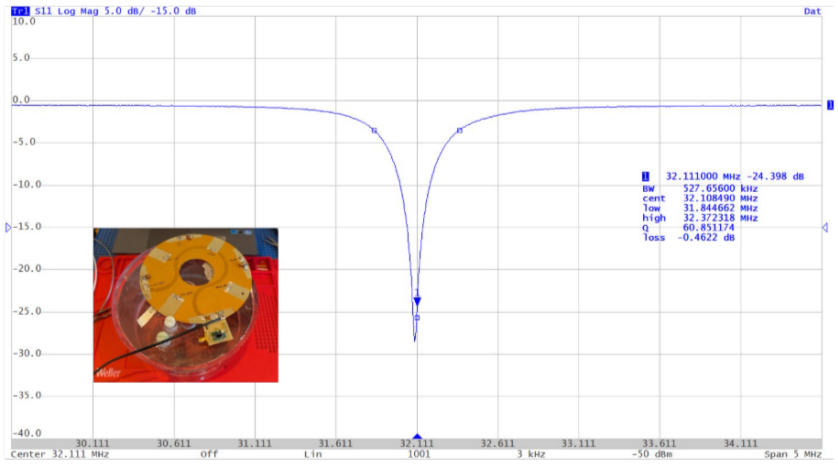


Figure 2.12: S-parameters from Rx Coil 3

introduce noise or artefacts into the ^{13}C receive chain; and more critically, generate unintentional localised RF heating causing burns, which is a serious safety concern.

I incorporated LC traps tuned to the proton frequency on each element. Each trap consists of a solenoid inductor and a series capacitor, with resonance determined by their combined impedance. Because the coil geometry is multi-layered and spatially complex, three proton traps were distributed across each Rx element to suppress current paths as comprehensively as possible.

To assess the aggregate effectiveness of the proton traps, I followed a stepwise bench top protocol. First, I connected a double-loop probe to ports 1 and 2 of the network analyser, ensuring that the probe was well-decoupled, with an expected

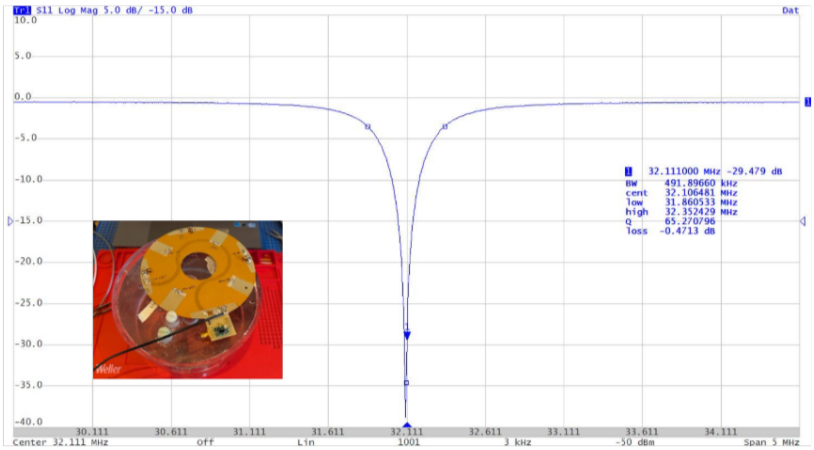


Figure 2.13: S-parameters from Rx Coil 4

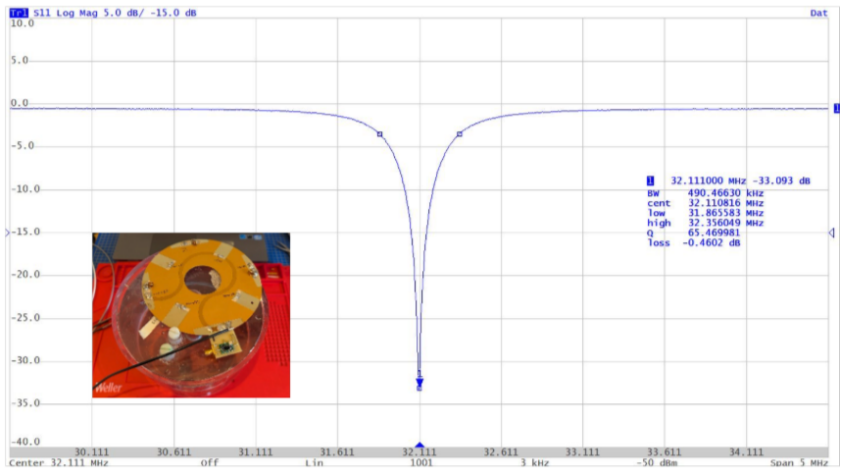


Figure 2.14: S-parameters from Rx Coil 5

baseline S_{21} value near -60 dB, as shown in Figure 2.15. I then positioned the probe above the receive coil at approximately half the diameter of the element (Figure 2.16) and recorded the S_{21}. A peak near 127.73 MHz indicated effective trapping. If no peak was observed, I proceeded to individually test each trap.

For individual trap testing, I soldered a coaxial cable across the trap's matching capacitor and connected it to port 1. I also desoldered the adjacent capacitors to open the circuit to test the proton trap in isolation. A single-loop probe was connected to port 2 and hovered directly over the trap. I then took the S_{21} measurement. A distinct peak at the proton frequency confirmed correct tuning. If necessary, I adjusted the trap by changing the length of the solenoid or replacing

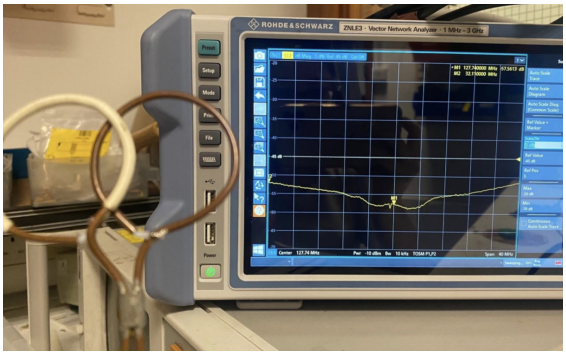


Figure 2.15: The double loop is decoupled and the S_{21} values are as close to possible to -60dB.

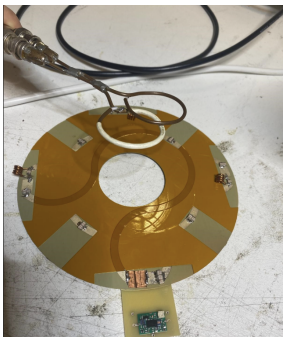


Figure 2.16: Proton trap testing was done by hovering the double loop above the Rx element

the series capacitor to shift the resonant peak. In practical terms, once the solenoid of the proton trap is made, all variables except for the length of the coils and the capacitor's value are fixed as seen in:

$$L = \mu_0 \mu_r \frac{N^2 A}{l} \tag{2.5}$$

where: L is the inductance in henries (H), μ_0 is the permeability of free space ($4\pi \times 10^{-7}$ H/m), μ_r is the relative permeability of the core material (dimensionless), N is the number of turns, A is the cross-sectional area of the coil in square meters (m^2), l is the length of the coil in meters (m). The final proton traps used the following components: a trap capacitor 39 pF (3.6 kV), Knowles Syfer 222523K60390GQTAF9LM (Digikey 1608-1577-1-ND), and a trap inductor 4 turns 1.0 mm enammeled wire wound on 3.9 mm former.

I repeated this on each coil for all three traps to get similar results on each trap.

Figure 2.17 is a representative response from a proton trap isolated from one of the coils showing the max value alignment at the proton's 127.73 MHz frequency.

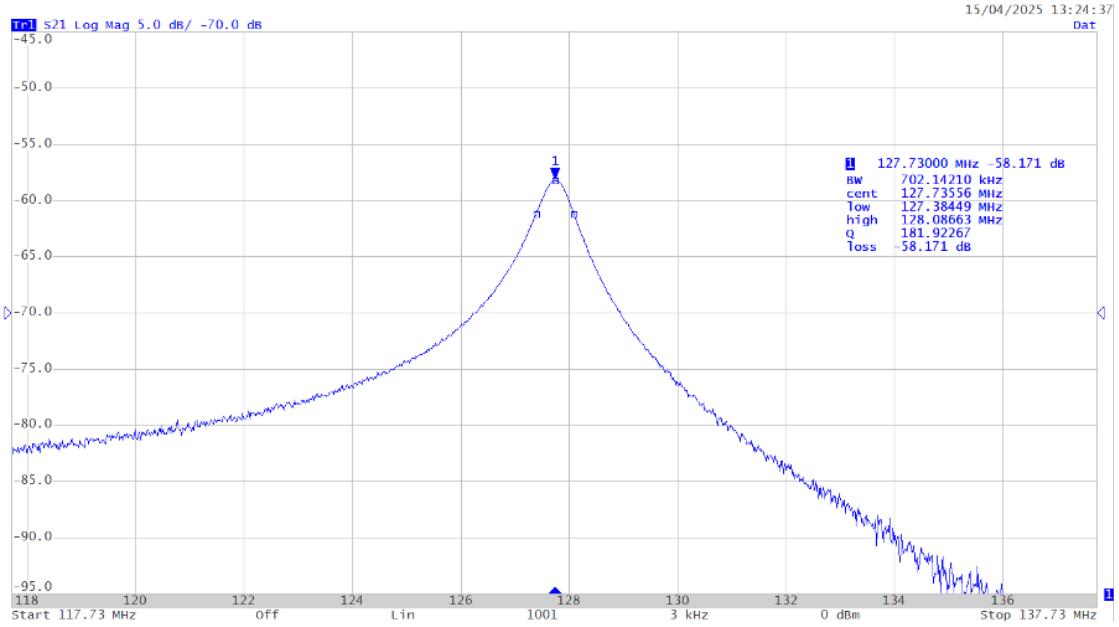


Figure 2.17: Response from updated proton trap build, isolated from a Rx coil

2.3.5 Preamplifier Testing

The preamplifiers (WMA3RA-R1: 3T low input impedance low noise pre-amplifier) used in this project were purchased from WanTcom Inc. and mounted onto the coil PCB by PulseTeq, my industrial collaborator. Each preamplifier was positioned at the immediate output of each individual coil element and was responsible for amplifying the weak MR signal induced in the loop. These signals are on the order of microvolts but the preamplifier boosts the signal. Although these components were supplied pre-built, it was essential to independently verify their specifications, particularly given their critical role in signal amplification.

To verify the manufacturer-provided specifications [74], I tested each preamplifier independently of the coil. The frequency response was measured using a VNA, and the test configuration is shown in Figure 2.18. The setup required a VNA, two coaxial cables (one terminated with dual prongs), a 50 Ω resistor, attenuators, and a constant 10 V DC power supply to drive the preamplifiers. Although MRI

compatibility was not required for bench testing, I selected non-magnetic, MRI-compatible connectors so that the same hardware could be used during subsequent in-bore experiments without introducing safety concerns, magnetic field distortions, or susceptibility-related image artefacts. The wiring and power supply arrangement is detailed in Figure 2.19.

Before conducting S-parameter measurements, I performed a full two-port calibration using the TOSM method (open, short, and match). Each port was individually tested: "open" was achieved by separating the prongs, "short" by soldering them together, and "match" by bridging them with a 50 Ω resistor. This calibration step ensured that the VNA measurements would reflect the actual device response, free from cable- or connector-related artifacts.

Following calibration, I measured the S_{21} for each preamplifier. During preamplifier characterisation, I encountered a persistent "Meas. Rec overload" error on the VNA at 35 MHz. Despite reducing the source power to −50 dB following a full calibration, the S_{21} response exhibited erratic spikes and discontinuities characteristic of receiver saturation (Figure 2.20).

This behaviour results from ADC saturation, where the incoming analogue signal exceeds the maximum voltage range of the Analogue-to-Digital Converter. This clips the signal peaks, distorting the smooth sinusoids that the VNA's digital signal processing expects to see. The resulting loss of phase coherence prevents the calculation of a stable transmission coefficient, resulting in the non-physical jagged traces.

I observed that while the response appeared smooth at supply voltages below 5 V, the overload error persisted even at 0 V. Using attenuators failed to rectify the condition, but finally I isolated the fault to a defective RF interconnecting cable with a potentially broken outer shield. This break essentially transformed the cable into an unshielded antenna, introducing noise and inducing high-amplitude signal reflections. Replacing the cabling immediately restored a linear and stable response. This experience highlighted the critical vulnerability of passive hardware, especially as ^{13}C coil arrays move toward clinical application. Failures like this

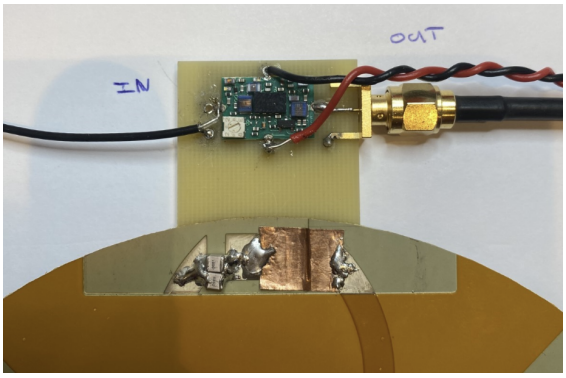


Figure 2.18: The preamplifier was tested in isolation to the rest of the loop as shown. Signal from the VNA comes in from the left and goes out from right in this image. The DC power is supplied through the twisted cables.

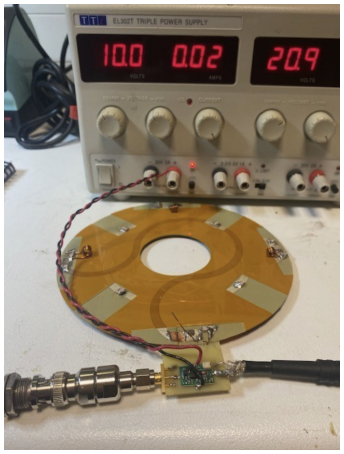


Figure 2.19: Rx element preamp connected to the DC power

underscore the necessity for rigorous daily quality-control protocols to ensure hardware integrity and data reliability.

Aligning closely with the manufacturer's specified nominal gain of 28.0 ± 0.5 dB, my results confirmed consistent gain at the carbon resonance frequency of 32.11 MHz: coil 1 = 27.51 dB, coil 2 = 27.45 dB, coil 3 = 27.50 dB, coil 4 = 27.49 dB, and coil 5 = 27.73 dB. Figure 2.21 illustrates a representative frequency sweep for Coil 1, and the complete characterisation data for Coils 2 through 5 can be found in the Appendix (Figures 2.87–2.90). These measurements validated that all preamplifiers were functioning within specification and provided consistent performance across the elements.

After confirming that the preamplifiers demonstrated a reliable frequency

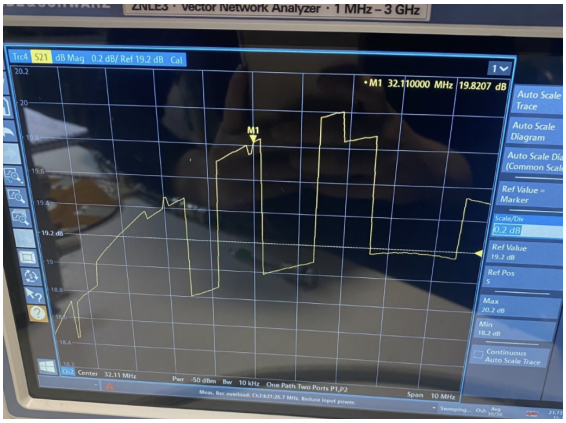


Figure 2.20: Vector Network Analyser S_{21} magnitude response under a receiver overload condition. The discontinuous curve and non-linear spikes indicate Analogue-to-Digital Converter saturation within the receiver, caused here by defective cabling rather than excessive input power.

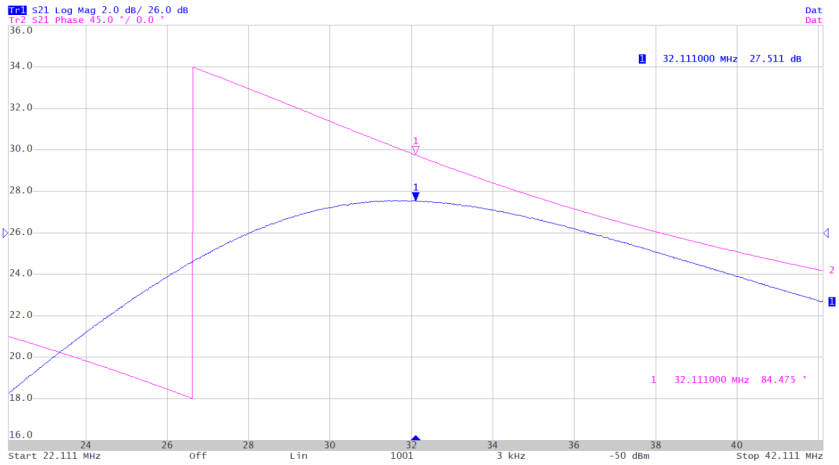


Figure 2.21: Rx coil 1 preamplifier S-parameters.

response within the target range, I proceeded to integrate them with the main coil circuit. In this configuration, the input terminal of the preamplifier (previously used in bench top testing) was connected across the matching capacitors, with wire placed in a PTFE sleeve for electrical isolation. The coil-to-preamplifier connection is shown in Figure 2.22.

To verify the integrity of the integrated system, I conducted a quick sanity check on the interaction between the multi-layer coil and the preamplifier. In this setup, the multi-layer coil was driven using a search coil, while the amplified output was measured using the VNA. The test arrangement is shown in Figure 2.23. The resulting S_{21} measurements exhibited a clear peak at 32.11 MHz, indicating that

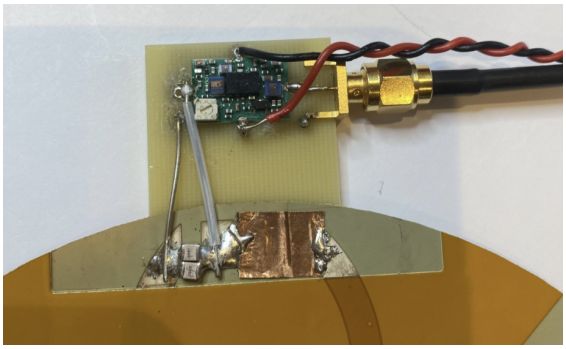


Figure 2.22: Close-up of the coil-to-preamplifier interface. The preamplifier input terminal is connected across the coil matching capacitors with a wire covered in a PTFE sleeve.

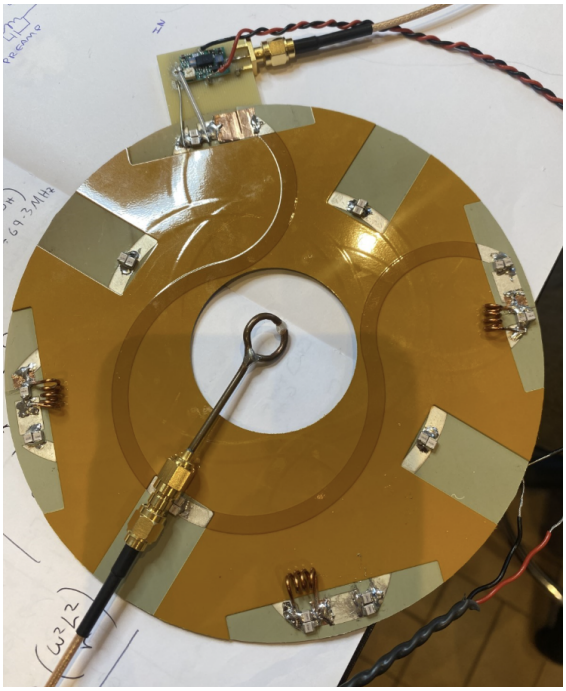


Figure 2.23: To test the preamp to coil connectivity, a search coil drives the multi-layer coil as shown.

the preamplifier and coil were functioning coherently as a single receive unit.

Although this preliminary response confirmed correct integration, a more detailed mapping of the S_{21} response was carried out once the coil elements were assembled into the full array. At that stage, inter-element coupling and interaction with the transmit coil were expected to significantly alter the impedance environment, necessitating re-characterisation under realistic loading and spatial configurations.

2.4 Geometric Decoupling of Two Rx Elements

2.4.1 S21 Scattering Parameters

When two elements resonant at similar frequencies are placed in proximity or with overlap, strong inductive mutual coupling causes the field profile from one element to skew the field of the neighbouring element. It may reduce the depth penetration and receive sensitivity of individual elements due to increased noise picked up by the nearby channels. Geometric decoupling can decrease these coupling effects by partially overlapping the elements to an optimal distance such that the mutual inductance is zero. For two simple loops, literature suggests that an overlap of approximately 20–30% of the coil diameter achieves optimal geometric decoupling [75–77]. Due to the complexity of the multi-layer coil, creating a mathematical framework for geometric decoupling was not practical. However, empirically finding an optimal separation by incrementally changing the overlap distance was very feasible.

My aim for this section was to determine the optimal geometric overlap for two adjacent multi-layer Rx elements placed on a phantom using both simulations and bench top experiments. To achieve this, I systematically varied the centre-to-centre spacing of the coils and measured the resulting S_{21} parameters. The S_{21} parameter quantifies the transmission coefficient, indicating the level of signal leakage or coupling between coils. A strong agreement between simulated and measured S_{21} values would confirm the effectiveness of geometric overlap in reducing mutual coupling for two elements, and supported the final design parameters used in the construction of the extended, multi-element full array.

2.4.2 Simulation Methods

To explore optimal geometric overlap, I performed Electromagnetic (EM) simulations using the CST Microwave Studio package to evaluate the S_{21} scattering parameters of the multi-layer and a simple loop counterpart (Figure 2.24 and 2.25). My goal

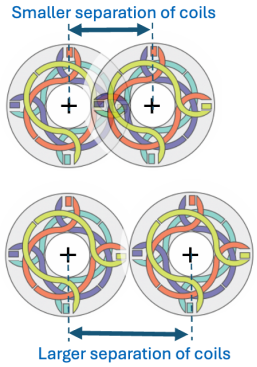


Figure 2.24: Simulated geometric decoupling carried out in CST Microwave Studio. The variable separation between coils is indicated by the blue arrow and represented as the x-axis in subsequent graphs

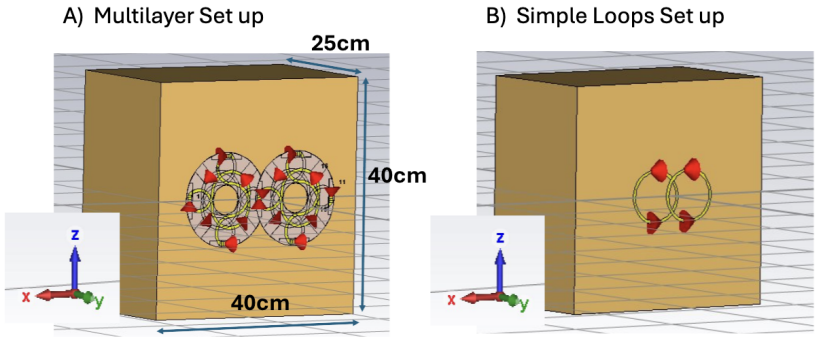


Figure 2.25: Set up in the EM simulations with the coils and fat phantoms for A) two multi-layer coils and B) two simple loops

was to systematically optimise geometric decoupling while preserving the enhanced sensitivity and depth penetration provided by the multi-layer architecture.

For these simulations, I used a homogeneous fat phantom block with dimensions 25 cm × 40 cm × 40 cm. Fat represented the low-loading end of physiological tissue properties that the chest coil will encounter. Given that coupling will degrade the performance of the coil, this low load explores the "best-case" depth penetration, which was ideal for early-stage parametric sweeps. The distance between the top of the phantom and the bottom of the coil surface was 5mm. The distance between the bottom of the lower coil element and the top of the upper coil element was 2mm. The thickness of each element was 0.2mm.

The simple loop design used for these simulation tests was a circular copper trace, 5 mm wide, with an inner radius of 43 mm and an outer radius of 48 mm. This

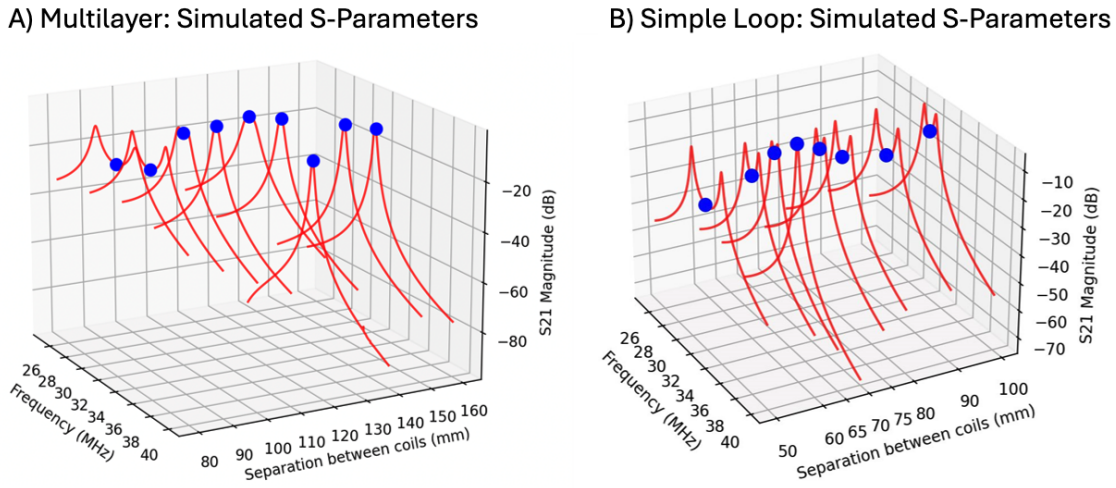


Figure 2.26: A) CST simulations on the multi-layer coil with separations ranging from 80mm to 160mm, with a range of single peaks that minimises at 140mm separation between coils. B) CST simulations on the simple loop coil with separations ranging from 50mm to 100mm with doublet peaks apart from the 70mm separation.

copper trace was the same size and material used in the multi-layer coil elements and the simple loop radius matches the most dense overlap of the four layers in the multi-layer element. Optimal matching and tuning capacitance values were chosen using a single loop on the simulated fat phantom. These values were used for all subsequent simple loop geometric decoupling simulations. There was one matching capacitor port and one tuning capacitor port. The distance between the top of the phantom and the bottom of the coil surface was 5mm. The two simple loops were separated by 2mm. Both coil types, multi-layer and simple loops, preserved capacitor values taken from a matched and tuned single coil on a phantom. I first examined the broadband coupling behaviour using CST simulations in figure 2.26 for both the multi-layer coil design (A) and the simple loop design (B). For the multi-layer coil, I simulated centre-to-centre separations from 80 mm (radius size) to 160 mm (diameter size).

2.4.3 Simulation Results

The S_{21} response of the multi-layer design exhibited a single resonance across centre-to-centre separations of 80 mm to 160 mm, with minimum coupling occurring

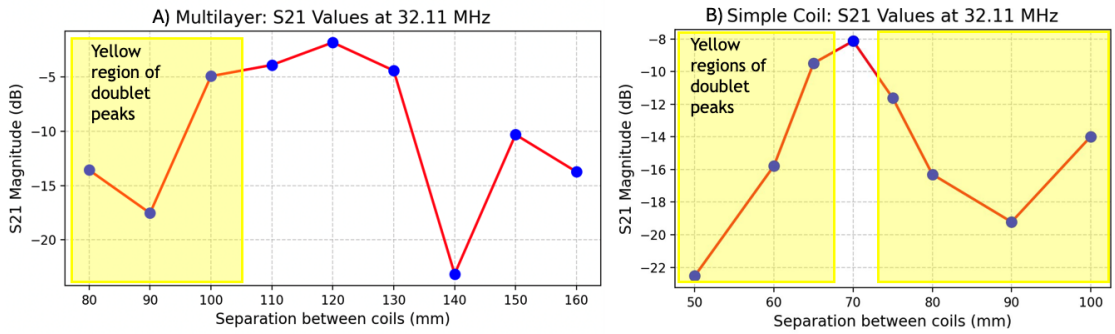


Figure 2.27: Simulated S_{21} values at the ^{13}C resonance frequency (32.11 MHz) as a function of centre-to-centre coil separation for (A) the multi-layer geometry and (B) the simple loop geometry. Yellow shading indicates separations where the simulated S_{21} response exhibited doublet resonances, corresponding to incomplete geometric decoupling.

at a separation of 140 mm. In contrast, the simple loop simulations, performed for separations between 50 mm and 100 mm, produced a characteristic doublet resonance due to residual coupling. Only the 70 mm separation reduced to a single resonance, suggesting improved isolation compared with the other spacings.

To better visualise the geometric decoupling tolerance at the carbon resonance frequency, Figure 2.27 summarises the simulated S_{21} values as a function of coil separation. For the multi-layer coils (Figure 2.27A), the yellow regions indicate separations for which the S_{21} spectrum exhibited doublet resonances and effective isolation could not be achieved. Outside these regions, a broad range of separations maintained a single resonance with low S_{21}, providing flexibility when selecting a practical inter-element spacing. In contrast, the simple loop geometry (Figure 2.27B) exhibited only a single separation that achieved effective isolation at 32.11 MHz, with the doublet resonance persisting at all other spacings. These results demonstrate that, although both geometries can achieve geometric decoupling, the multi-layer design offers a much broader design tolerance, making it less sensitive to manufacturing variation and coil positioning than the simple loop geometry.

From these results, I selected a 14 cm centre-to-centre separation for further evaluation. Figure 2.28 shows the simulated field magnitude profile along the y-axis through the centre of both coils within the phantom. Because the coils have finite thickness, one of the two coils was offset by 0.2 cm away from the phantom in the

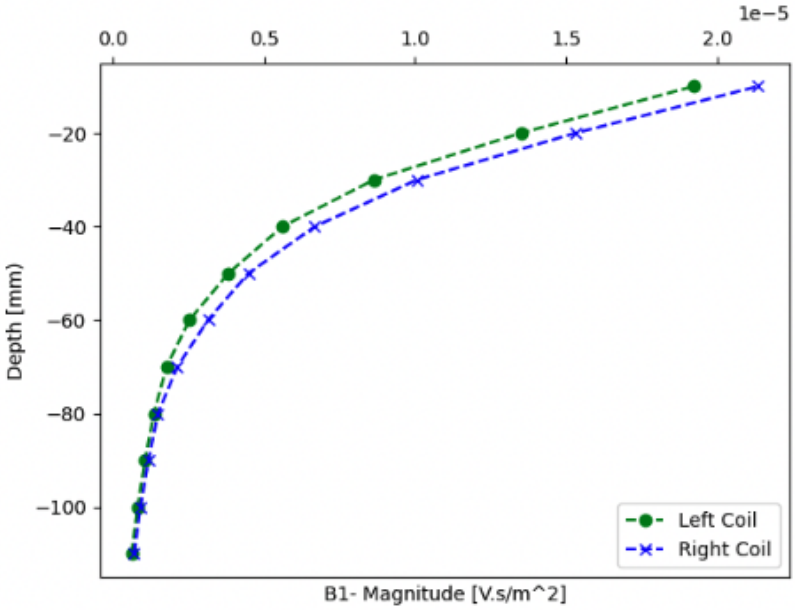


Figure 2.28: Magnitude profile along the centre of the two coils, along the y-axis within the phantom. The 0.2cm separation between coil 1 and coil 2 are negligible at the depths where the heart is located (<-60mm).

model. However, the profile demonstrates that this small separation has a negligible effect at the depths relevant for cardiac imaging (greater than 60 mm into the phantom). In this region, the field strengths from both coils remain well balanced, indicating that the physical layering of the coils does not introduce substantial sensitivity differences at the depth of the heart. These results confirm that the chosen design preserves uniform field penetration at clinically relevant depths while maintaining the decoupling performance identified from the S_{21} simulations.

2.4.4 Bench Top Methods

Given that the geometric decoupling simulations led to promising S_{21} values, I carried out bench top S_{21} measurements using two identical receive-only surface coils, each matched and tuned independently to the ^{13}C resonance frequency (32.11 MHz) in the presence of a 6 L saline phantom (2.34 g/L NaCl in a plastic container, further details on the preparation below). The details of the multi-layer coils and their matching and tuning are as described in the previous section 2.3: Characterization of a Single Rx Element.

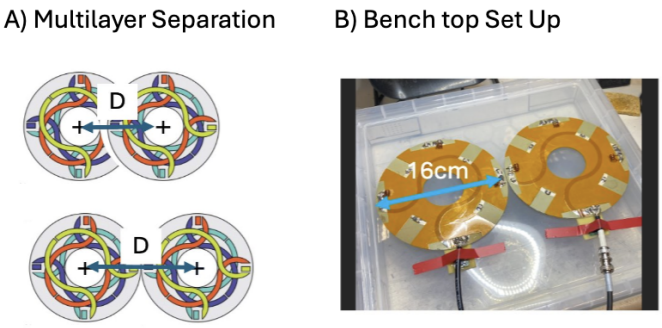


Figure 2.29: A) Two multi-layer coils are set up adjacent to each other, where the distance between the two centres (D) is varied: radius (8cm) < D < diameter (16cm) in increments of 1 cm. There is no rotation of the coils. B) The bench top set up of the two multi-layer coils connected by the coaxial cables placed on the saline tub phantom.

Each coil was connected to a separate port of the VNA via equal-length, phase-matched coaxial cables terminated with connectors. I calibrated the VNA to the end of each cable using the standard open–short–load procedure. The following test arrangement is shown in Figure 2.29. I varied the centre-to-centre distance (D) between the two coil elements from 8 cm (coil radius) to 16 cm (coil diameter) in 1 cm steps. For each spacing, I performed a frequency sweep centred at 32.11 MHz with a span of 10 MHz and a bandwidth of 10 kHz, and recorded both the magnitude and the frequency of S_{21}.

My aim was to find the spacing that gave the minimum S_{21} at 32.11 MHz, adhering to the -15 dB benchmark defined in Section 2.2.6. All measurements were carried out at room temperature (20–22 °C), and the coils were measured without active detuning or preamplifiers to isolate the effect of geometry on coupling.

I used a 6 L non-conductive, flat top plastic tub to house the saline solution, with the volume chosen to ensure complete coverage of the coil's sensitive region and provide a realistic geometric load. Prior to data collection, I prepared the phantom with 2.34 g/L sodium chloride (NaCl) in water, with the concentration serving as a proxy for solution conductivity following the phantom preparation by Giovannetti et al [23]. By monitoring the input reflection coefficient (S_{11}) of the coil on a VNA with the phantom in place, I iteratively adjusted the NaCl concentration until the impedance matched 50 Ω at the carbon frequency. Matching to 50 Ω is essential

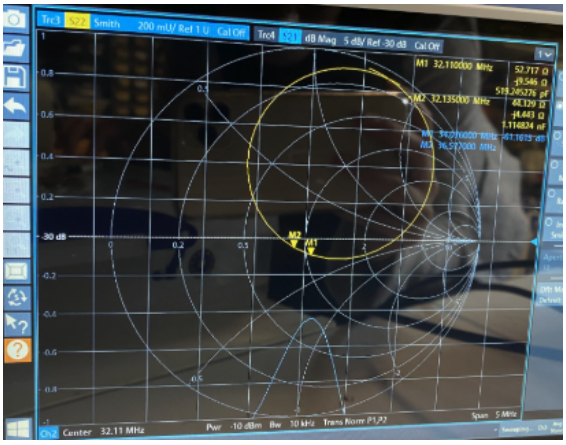


Figure 2.30: Smith chart showing the matched S22 response of the coil with the saline phantom, intersecting near 50 Ω at 32.11 MHz.

to maximise power transfer between the coil and the MRI system's RF chain, thereby minimising reflection losses. The tuning process was guided by Smith chart analysis, and a representative trace is shown in Figure 2.30, where the S_{22} curve intersects the real axis slightly below the 50 Ω point, which I deemed acceptable for experimental purposes. All impedance measurements were performed with the phantom equilibrated at room temperature for at least two hours to ensure thermal stability of the electromagnetic properties. The plastic container was selected for its negligible RF interaction aside from its geometric influence on coil loading.

As a one-time calibration measurement, I measured the S_{21} response with the coils placed at the widest possible distance on the phantom (24cm) while maintaining the same height plane, which exceeded the diameter (16cm). This one-off baseline measurement, where D > diameter served to contextualise the reduction in coupling observed with progressive overlaps in radius (8cm) < D < diameter (16cm).

2.4.5 Bench Top Results

The plot of S_{21} versus centre-to-centre separation in Figure 2.31 shows the isolation values achieved at different spacings between the two elements. I found a broad range of acceptable decoupling values (<-15 dB) across separations of 12 cm, 13 cm, and 14 cm. Not included in this plot, at a maximum achievable separation of

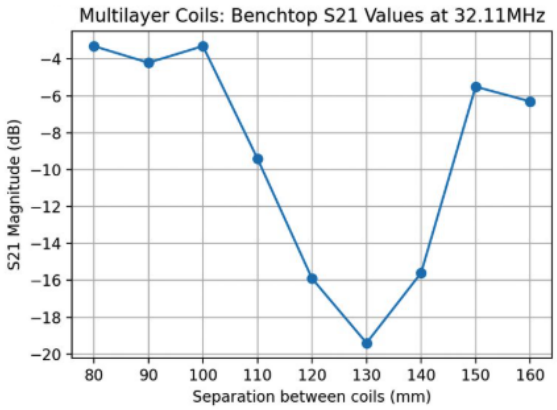


Figure 2.31: Bench top S_{21} values at varying overlap distances at the carbon frequency for the multi-layer two coil configuration.

24 cm, the S_{21} value reached -26.16 dB, which I took as a baseline for what could be achieved with near-perfect decoupling under ideal, non-interacting conditions.

I further explored variations in geometric decoupling beyond the one-dimensional gridding formula I initially used. Given the complex design of the individual coil elements, there was scope to rotate the loops relative to one another to see if this improved isolation. I also tested changes in the relative phase between coils as another way to optimise their interaction. The results from these rotation and phase-shift tests are provided in the Appendix section 2.7: Further Rotational Geometric Decoupling. While these adjustments had some effect and require further exploration, the clear takeaway was that centre-to-centre distance remained the most influential factor in decoupling performance.

One specific configuration I tested was a 180° rotation, shown in Figure2.32, because this arrangement is included in the planned four-element array layout. Placing one element rotated by 180° and directly facing the other did not degrade isolation, suggesting that this orientation could be a useful option when building multi-paddle array designs. The corresponding S_{21} measurements are shown in Appendix Figure 2.78.

I did notice some variability in the S_{21} measurements across repeated trials. This is likely due to the sensitivity of the bench top setup such as small, unintentional

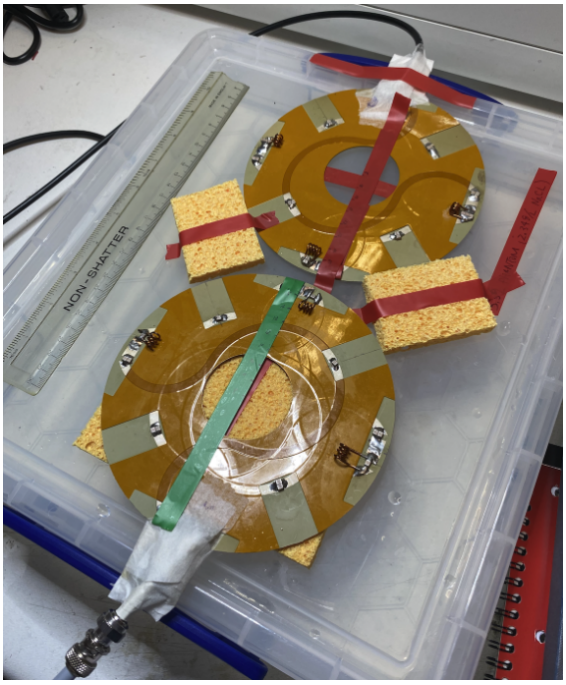


Figure 2.32: Set up of one element rotated by 180° and directly facing the other, which is incorporated in a subsequent four-element design.

bends in the coil conductors or slight shifts in positioning, which could lead to changes of up to $\tilde{5}$ dB.

2.4.6 Bench Top and Simulation Discussion

The CST simulation results for the multi-layer design established a theoretical optimal spacing of 14 cm, while revealing a robust geometric decoupling range from 13 cm to 16 cm, where a single peak was maintained and S_{21} was minimised. In contrast, simulations of the simple loop design identified a highly sensitive geometric decoupling at exactly 7 cm and any deviation from this spacing resulted in doublet peaks. These findings highlighted the inherent robustness of the multi-layer architecture for practical array construction.

Subsequent bench-top measurements confirmed these simulated trends, with the experimental optimal spacing found at 13 cm. This slight 1 cm shift from the simulated optimum (14 cm) was expected, as physical hardware introduced variables such as finite conductor resistance and manufacturing tolerances that are absent in the idealized CST environment. Furthermore, while simulations

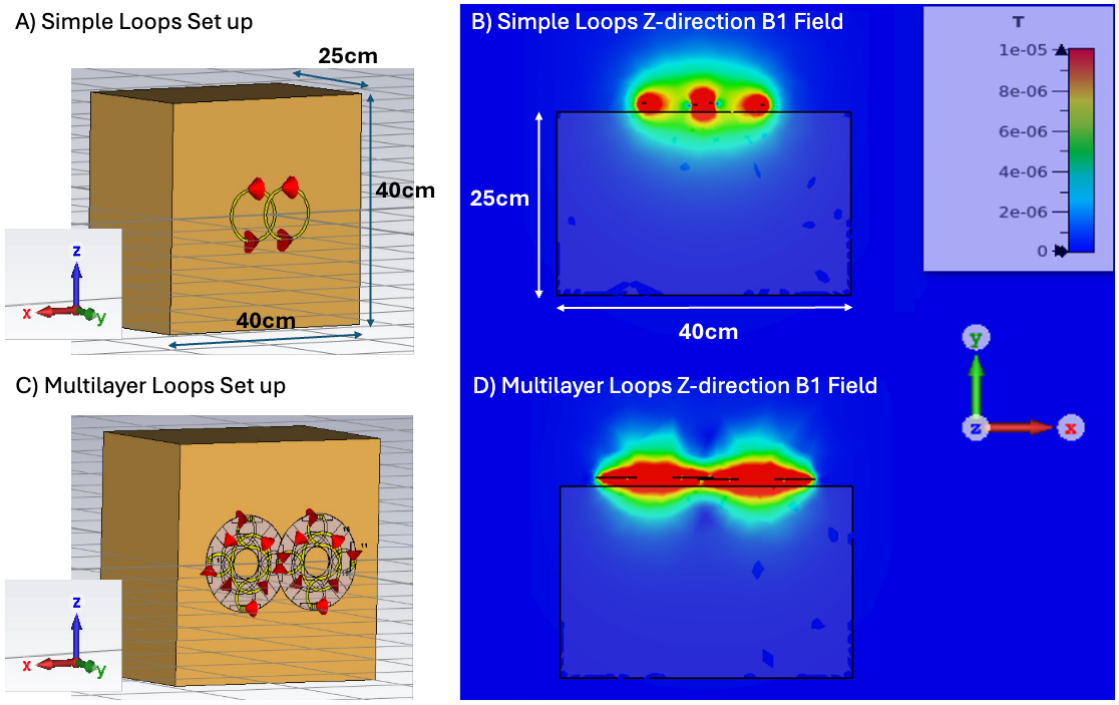


Figure 2.33: Comparison of the optimal S_{21} for the multi-layer with the optimal for the simple loop. All phantoms are fat phantoms (40 x 40 x 25cm) and the B_1 field is on a 1e-5 T scale

allow for infinitesimal capacitor adjustments to achieve perfect resonance, bench-top tuning is constrained by the physical requirement to manually desolder and resolder components. This process limits fine-tuning precision and introduces potential heating damage to the coil. With these minor physical discrepancies, the experimental data showed overall agreement with the simulations. Both approaches demonstrated that the multi-layer array maintains effective geometric decoupling over a wider range of separations than the simple loop, validating the design's suitability for a multi-element cardiac array.

2.4.7 B1 Field Simulations

Having established that the multi-layer and single loop coils achieved effective geometric decoupling, the next step was to characterise how the designs performed in terms of the produced receive field within tissue. Using the fat phantom model, I ran B_1 field simulations for both coil element types (i.e. multi-layer and simple loop). I examined the B_1 field contours as a z-direction slice ($z = 0$) through

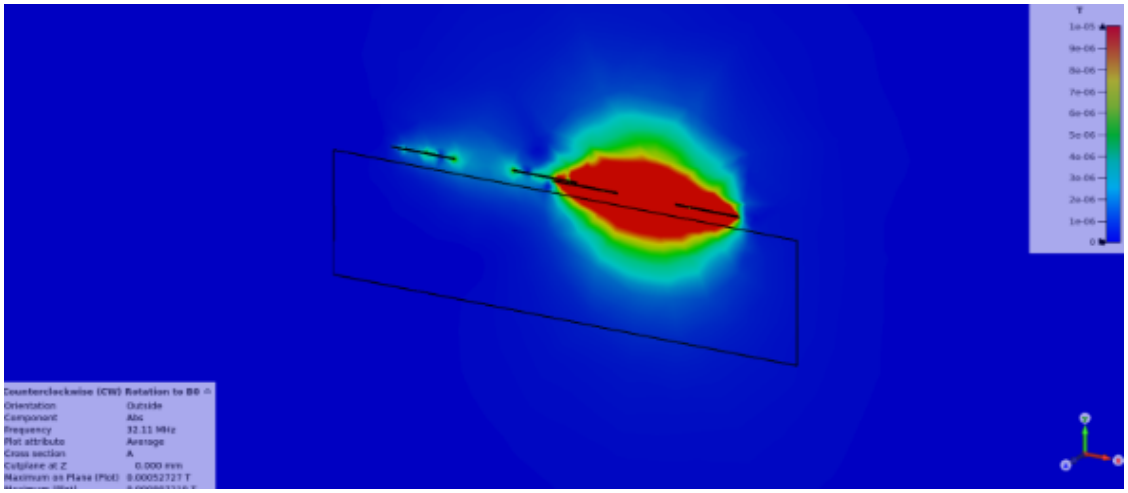


Figure 2.34: Multi-layer 14cm, single coil powered. Sanity check that each coil is powered correctly

the midsection of the coils and phantom. Figure 2.33 compares the B_1 field at optimal geometric decoupling distances for each design: 7 cm for the simple loop and 14 cm for the multi-layer loop. These represent the best-matched and tuned configurations for each array.

The multi-layer array produced a more homogeneous field pattern, appearing as two broad pancake shapes. In contrast, the simple loop array generated three distinct hot spots immediately around the copper traces. However, in the multi-layer configuration I observed a pronounced drop in B_1 field strength between the two elements, where I would expect the heart to be located. To check whether this was an error in power input, I ran single-element excitation simulations (Figure 2.34). After verifying that each element was powered correctly, the drop in the centre persisted, confirming that it is intrinsic to this two-element coil geometry.

I attempted to improve the field profile by changing the relative rotation of the elements and by adjusting the phase difference between them (Appendix subsection 2.7: Further Rotational Geometric Decoupling). Neither approach resolved the central drop, so I returned to the design stage. The resulting solution is described in the next section.

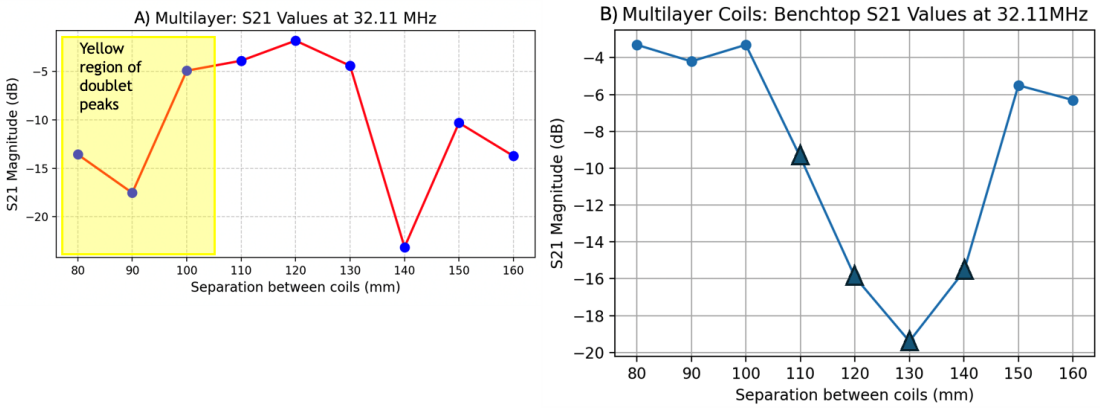


Figure 2.35: S21 values for the multi-layer array as a function of centre-to-centre separation. (A) CST simulation results and (B) bench top measurements both demonstrate a broad range of low S21 values, indicating a range of spacing configurations that achieve effective geometric decoupling. This wide range enables design flexibility, allowing trade-offs between isolation performance and B_1 field homogeneity.

2.5 Simulations of Extended Array Designs

2.5.1 Trading Geometric Decoupling with Field Homogeneity

As outlined in the previous section, the primary limitation of the 14 cm separated multi-layer configuration was a dip in B_1 field strength in the central region between elements, where the heart would be located. To address this, I leveraged the multi-layer design's broad range of low S_{21} values to run EM simulations across narrower centre-to-centre separations than the optimal 14 cm. The objective was to explore the trade-off between reduction in matching and tuning and improvement in field homogeneity.

Discussions with my industrial supervisor highlighted prior, unpublished work on the proton analogue of this project, which also employed a multi-layer design. In that case, moving the elements closer together, thereby accepting a slightly lower geometric decoupling value, resulted in improved uniformity within the high-sensitivity region. This provided the rationale for investigating whether a similar approach could be applied to the ^{13}C coil.

Previous simulation results (figure 2.35A) indicated that 11 cm was the minimum spacing achievable without introducing doublet peaks in the S_{21} response. Bench top

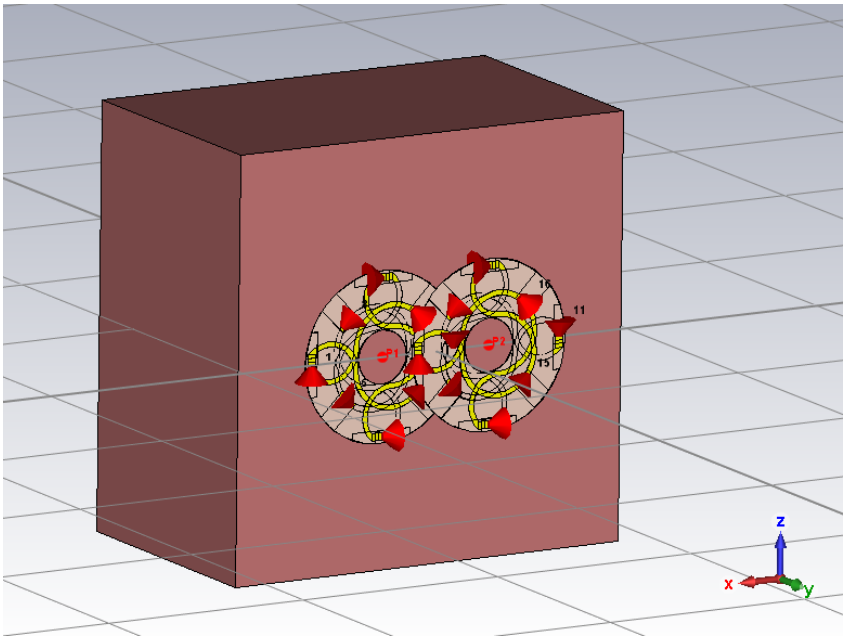


Figure 2.36: Muscle phantom setup (40 cm × 40 cm × 25 cm) for high-dielectric loading tests. P1 and P2 mark the coil centres, which were adjusted to achieve separations of 11-14 cm for geometric decoupling and homogeneity evaluation.

measurements from 11 cm to 14 cm separations, highlighted with triangular points in figure 2.35B, demonstrated a range of decoupling (-19.2 dB < S_{21} < -9.6 dB), which may be further improved through fine-tuning during subsequent engineering stages. Therefore, I believed that pushing the coils closer towards the 11 cm configuration could deliver improved B_1 homogeneity while maintaining acceptable isolation.

2.5.2 Two-Element Spacing Trade-Offs

I ran further CST simulations on the two multi-layer coil configuration placed on a muscle phantom. The phantom measured 40 cm × 40 cm × 25 cm (Figure 2.36). Coil centre positions (P1 and P2) were adjusted to achieve centre-to-centre separations of 11 cm, 12 cm, 13 cm, and 14 cm. Building on the fat phantom results which established best-case performance (Section 2.4.2), I next evaluated the two-element multi-layer array on a high-dielectric, high-load muscle phantom. This allowed me to assess heavier loading conditions. In practice, the dielectric properties of the human chest will fall between these two extremes, but testing both limits provides important design boundaries. Figure 2.37 shows the simulated B_1 fields

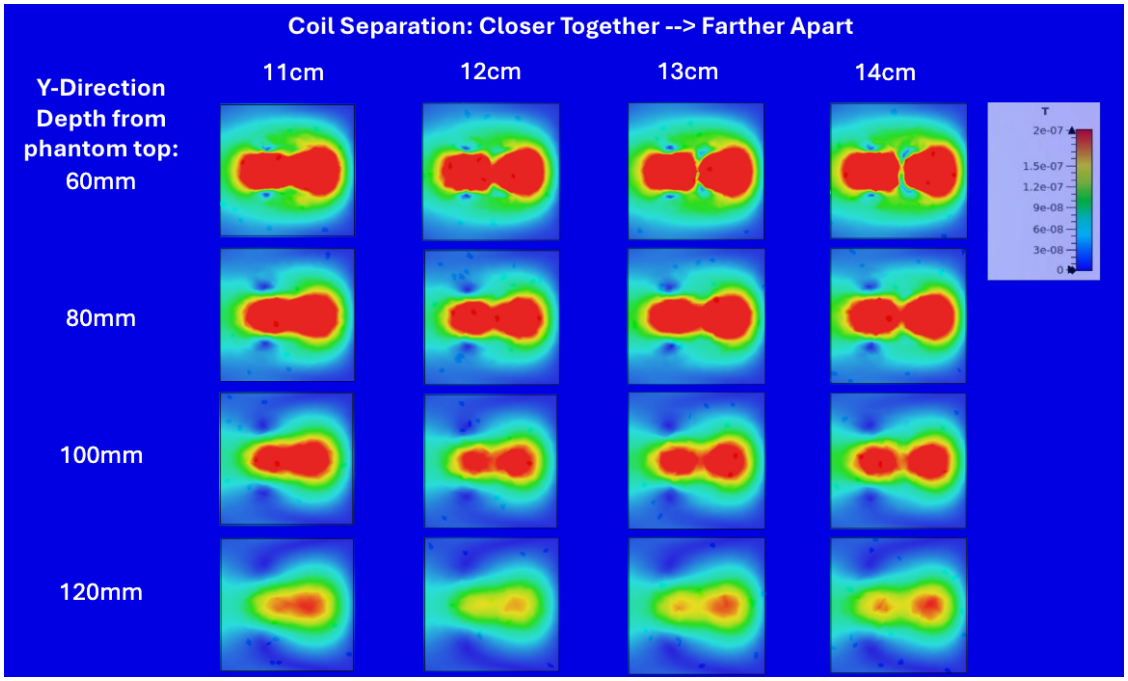


Figure 2.37: Simulated B_1^- field maps for the two-element multi-layer array on the muscle phantom at separations of 11-14 cm. Tighter spacing (11 cm) improves central field uniformity, reducing the dip between elements.

for the tested separations. Reducing the spacing increased field uniformity in the central region between the elements. This effect was particularly clear for the 11 cm configuration, which showed the most even distribution of field strength at 120mm depth, across the approximated heart region.

The S-parameters for each separation (Figures 2.38–2.41) illustrate the side of the trade-off where coil interaction increases due to coupling. Even under best-case conditions with both elements well matched and tuned (S11< -15 dB and S22< -15 dB), coupling is evident. At 11 cm (Figure 2.38), isolation is reduced compared to the wider spacings (S21 = -9.56). Increasing the separation to 12 cm (Figure 2.39) causes a marked drop in isolation performance, suggesting that this intermediate geometry is particularly susceptible to coupling (S21 = -4.89). At 13 cm (Figure 2.40), the isolation begins to recover (S21 = -10.54), and by 14 cm (Figure2.41) the array achieves its best (S21 = -24.64) isolation in this series, consistent with the geometric decoupling behaviour observed in the earlier simulations.

Quantitative metrics were calculated within an approximate heart region, defined

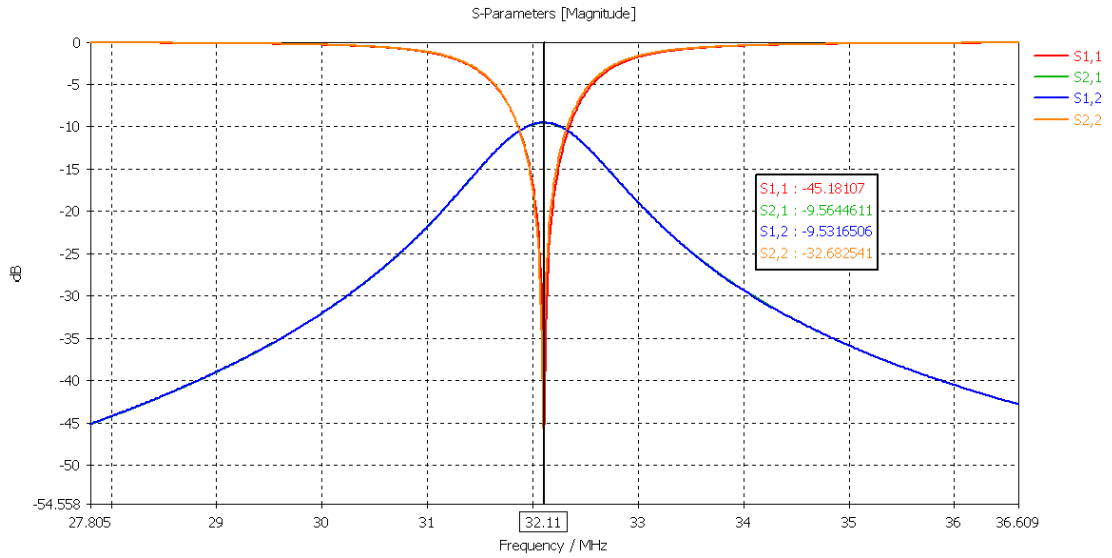


Figure 2.38: S_{21} response for the 11 cm separation. Isolation is reduced compared to wider spacings, but remains acceptable for further optimisation.

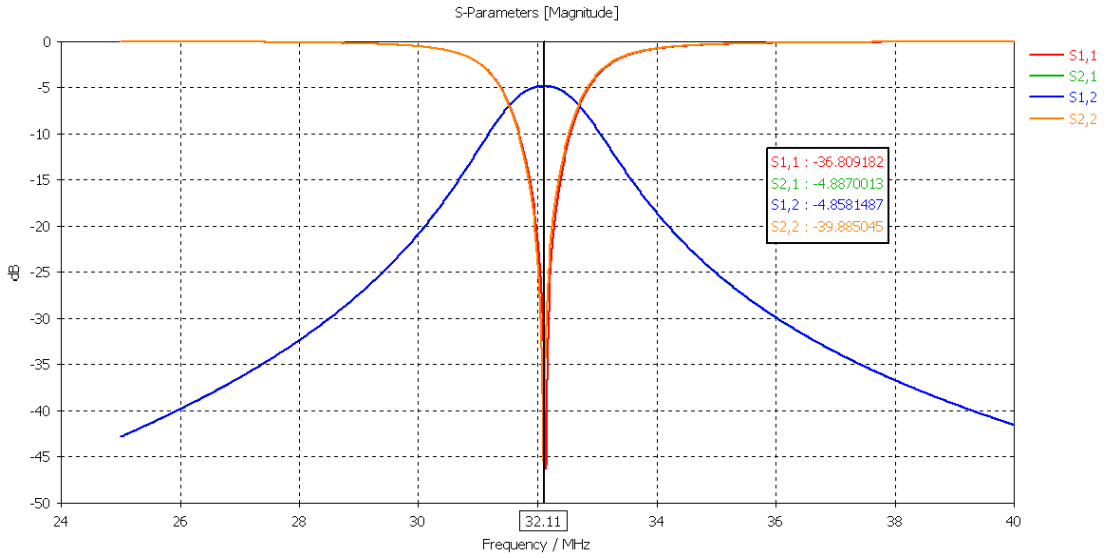


Figure 2.39: S_{21} response for the 12 cm separation. Isolation is severely reduced.

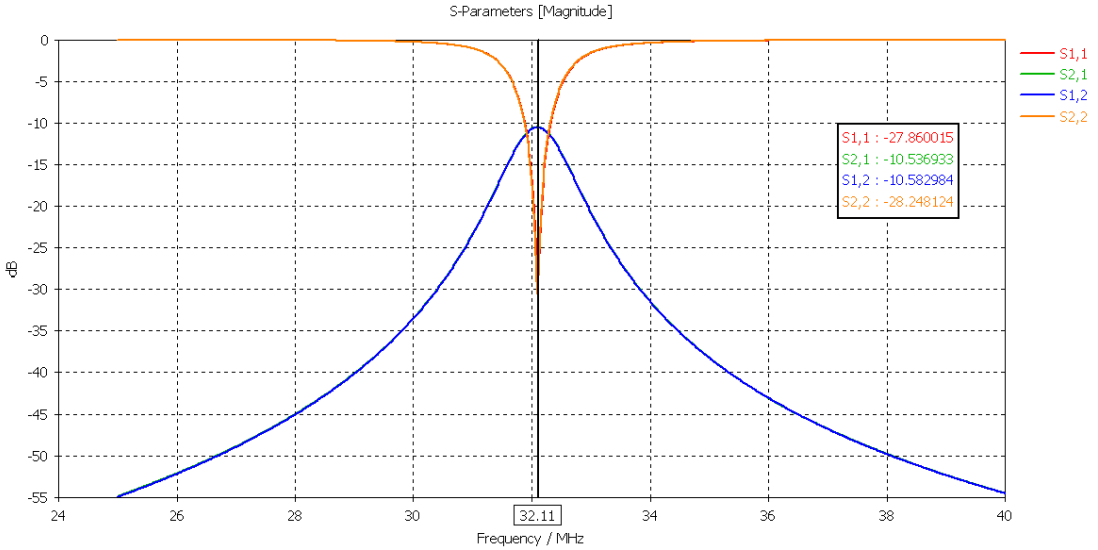


Figure 2.40: S_{21} response for the ^{13}C m separation. Isolation is starting to recover as the spacing increases.

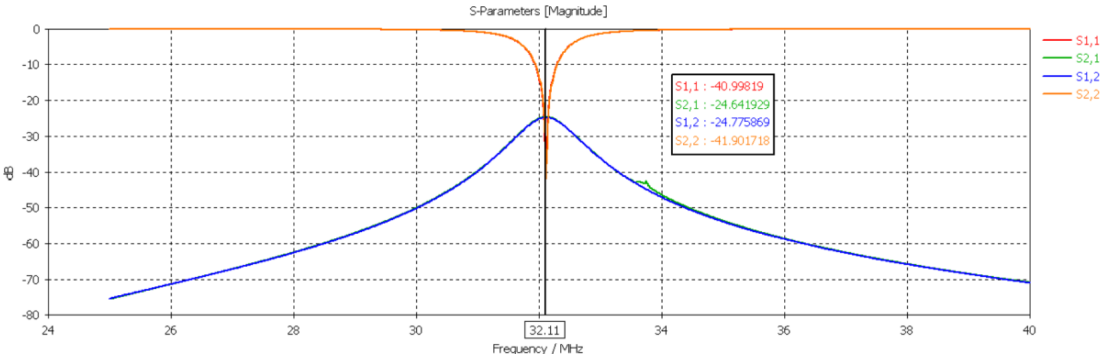


Figure 2.41: S_{21} response for the 14cm separation. This spacing provides the best isolation.

here as an 80×80×80 mm ROI centred 50 mm beneath the phantom surface. This simplified ROI represents the approximate volume occupied by the heart in the phantom, with its dimensions and depth later verified using anatomically realistic human phantom simulation (Section 2.5.4). Here, the smaller the separation, the higher the uniformity (defined as 1-CV) that is achieved, with 11cm achieving the highest mean B_1 value. Further, the violin plots in Figure 2.42 illustrate the distribution of the B_1 values. The violin plot for 11 cm separation appears block-like, with a relatively even distribution of voxel B1- values across the central range, consistent with high homogeneity. In contrast, the 12-14 cm separations are more asymmetric, each showing a pronounced bulge at a particular field strength

value, indicating a clustering of voxels around that sensitivity level, which is lower than that of the 11cm. This skewed shape illustrates reduced uniformity with a hot spot contributing most to the mean value.

Even with high loading conditions, the observation that the 11 cm configuration produces the highest mean B1- field values despite having the poorest S21 isolation highlights the trade-off between electromagnetic coupling and transmit efficiency. At this close separation, the magnetic fields from each element constructively overlap within the heart-sized ROI, boosting the measured B_1 amplitude. However, the proximity also increases mutual inductive coupling, meaning a larger fraction of the driven coil's energy is transferred into the neighbouring element rather than the load. For this project, the trade-off remains acceptable and can be mitigated with targeted matching and tuning. This result reinforces the potential of the 11 cm configuration as a strong candidate for the final design.

Table 2.1: Quantitative B_1 metrics (mean, standard deviation (SD), coefficient of variation (CV), and uniformity) for the pre-defined heart region ($80 \times 80 \times 80$ mm^3) at coil separations of 11–14 cm. The anatomical basis for the ROI is described in Section 2.5.4.

Coil Separation	Mean [T]	SD [T]	CV	Uniformity
11 cm	4.9276×10^{-7}	2.8879×10^{-7}	0.5861	0.4139
12 cm	3.8557×10^{-7}	2.3583×10^{-7}	0.6116	0.3884
13 cm	3.9738×10^{-7}	2.4735×10^{-7}	0.6225	0.3775
14 cm	3.7746×10^{-7}	2.4975×10^{-7}	0.6616	0.3384

2.5.3 Three- and Four- element Design Development

Building on these findings, I next extended the 11 cm separation configuration into three- and four-element geometries. The goal was to assess whether the improved B_1 homogeneity seen in the two-element tests could be maintained or enhanced when scaled to larger arrays, and whether the associated increase in coupling could still be managed through matching and tuning. Adding more elements changes the interaction landscape as each coil now couples not only with its immediate neighbour but also with second- and third-nearest elements, potentially amplifying the challenges observed in the two-element case. Testing these multi-element

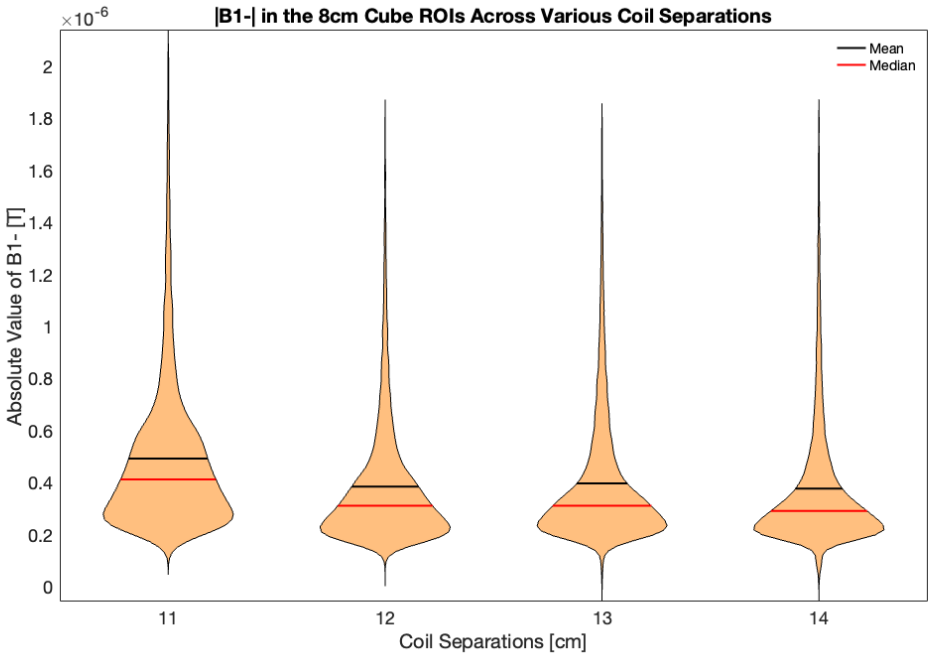


Figure 2.42: Violin plots of voxel-wise B_1 field values in the heart region (as in Figure 2.1) for separations of 11-14 cm. The 11 cm configuration shows the highest mean (black line), and the highest median (red line) and a block of high .

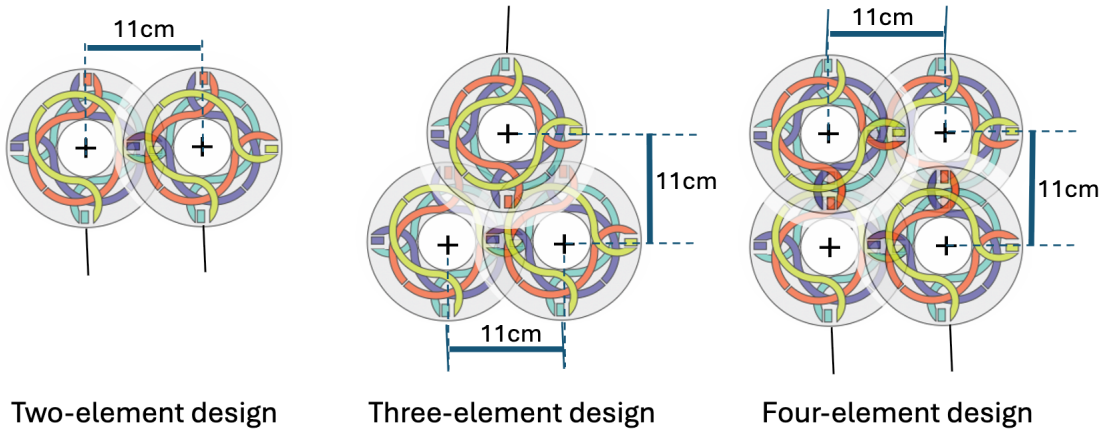


Figure 2.43: Diagram of the multi-element array designs for the multi-layer coils. The solid thin black lines show where the preamp protection and cables are attached to the coil.

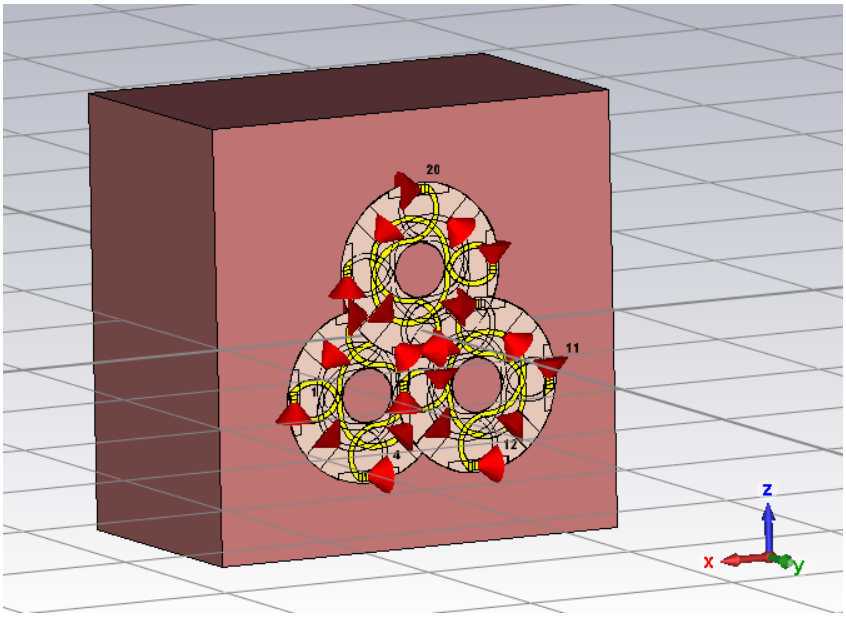


Figure 2.44: Set up of three-elements on the muscle phantom

configurations was therefore a critical step in determining whether the 11 cm spacing remains viable for the final cardiac coil design.

The EM simulations, previously optimised for the two-element multi-layer configuration, were extended to three- and four-element arrays (Figure 2.43). Simulations followed the CST workflow described in Section 2.2.5, using a tetrahedral mesh and frequency domain solver on a 40 cm × 40 cm × 25 cm muscle phantom (Figure 2.36). The heart ROI was taken as previously defined in section 2.5.2. Coil centre positions were fixed at 11 cm separation for geometric decoupling.

The three-element configuration (Figure 2.44) was arranged in a pyramid layout, with all coils well matched and tuned to the ^{13}C frequency. The resulting coupling values (S_{12} = -9.58 dB, S_{13} = -17.79 dB, S_{23} = -18.20 dB; Figure 2.45) were comparable to, or better than, the two-element case. The four-element configuration (Figure 2.46) also achieved good matching (all S_{11}, S_{22}, S_{33}, S_{44} < -15 dB) with inter-element coupling no higher than -11.40 dB (S_{34}; Figure 2.47).

Field distribution comparisons across depths (Figure 2.48) revealed that adding more elements increased the size of the homogeneous B_1 region, particularly at 120 mm depth. Quantitative analysis for the heart ROI (Figure 2.2) showed that the mean B_1^- value was highest for the two-element array and decreased with

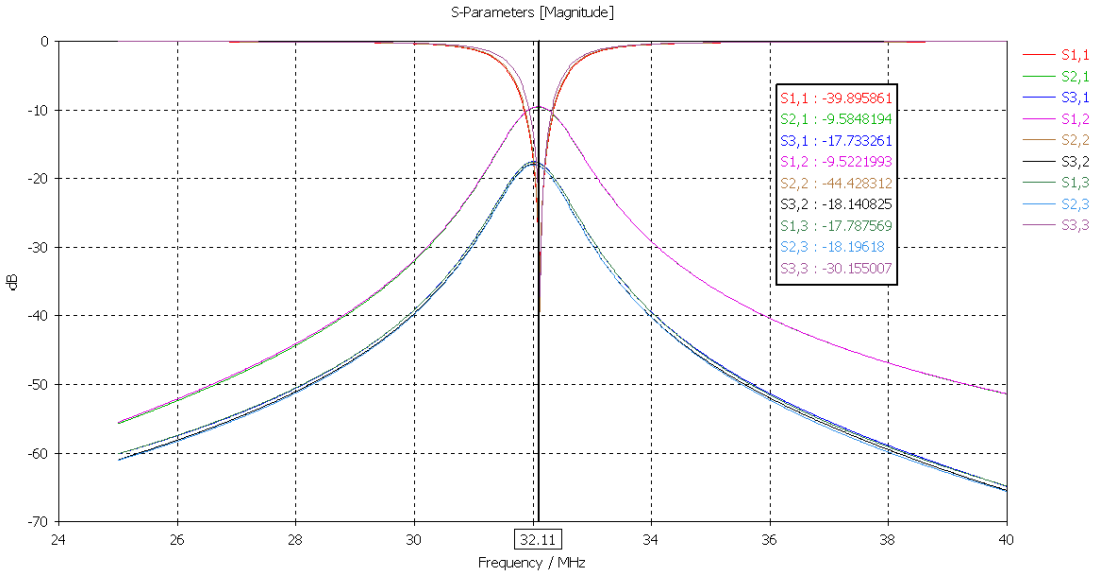


Figure 2.45: S-parameters of three-elements on the muscle phantom. All three coils are well matched and tuned to the carbon frequency and shows comparable or better values of coupling (S12 = -9.58, S13=-17.79, S23=-18.20) to those seen in the two coil array

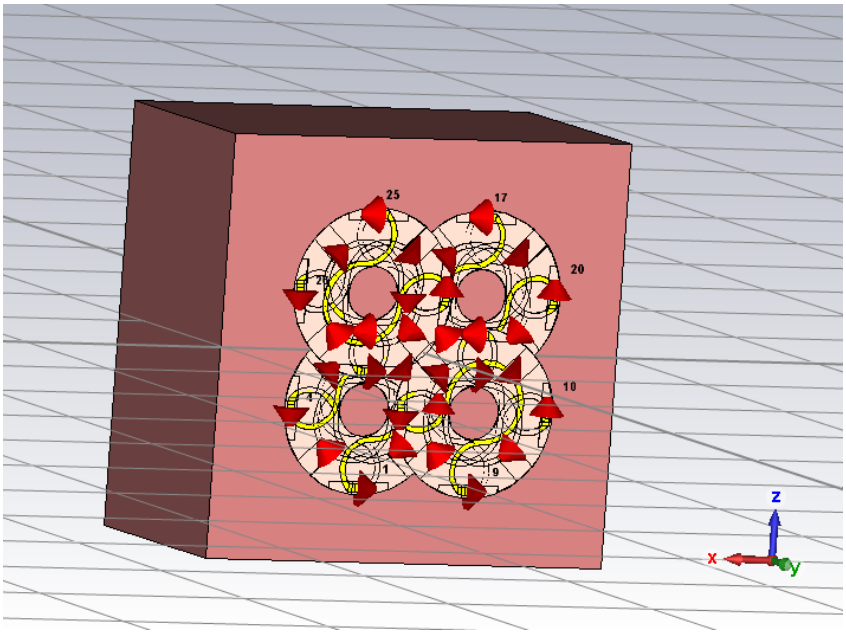


Figure 2.46: Set up of four-elements on the muscle phantom

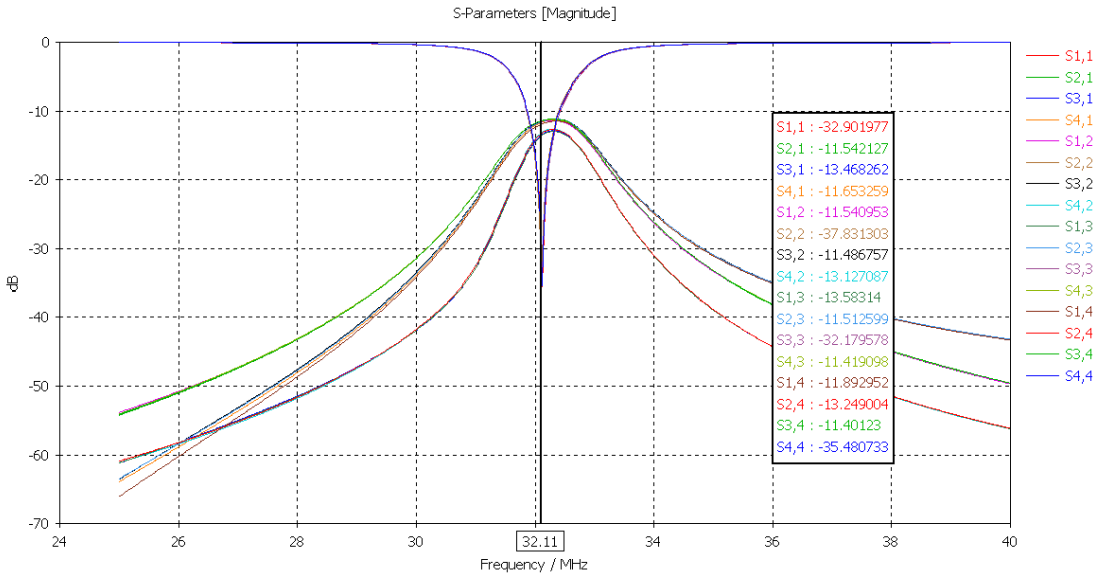


Figure 2.47: S-parameters of four-elements on the muscle phantom. All coils are well matched and tuned (<-15dB) and all inter-coil S-parameters are also all below -11.40dB, which was S34.

more elements, while uniformity (1–CV) improved with increasing coil count: the uniformity increased from 0.4139 for the two-element array to 0.4886 for three elements and 0.6616 for four elements. Violin plots (Figure 2.49) further illustrate this trend. The two-element configuration exhibited the highest mean (black line) but also a longer tail in the distribution, indicating greater voxel-to-voxel variation. In contrast, the three- and four-element arrays showed progressively more compact distributions, consistent with improved homogeneity.

Table 2.2: Quantitative B_1 metrics (mean, standard deviation (SD), coefficient of variation (CV), and uniformity) for the pre-defined heart region ($80 \times 80 \times 80$ mm^3) comparing the two-, three-, and four-element receive arrays. The anatomical basis for the ROI is described in Section 2.5.4.

Array	Mean [T]	SD [T]	CV	Uniformity
2 element	4.9276×10^{-7}	2.8879×10^{-7}	0.5861	0.4139
3 element	4.3699×10^{-7}	2.2346×10^{-7}	0.5114	0.4886
4 element	3.2619×10^{-7}	1.1040×10^{-7}	0.3384	0.6616

In summary, scaling the 11 cm configuration to three and four elements preserved acceptable coupling while improving field uniformity in the heart ROI, albeit at the expense of peak B_1 amplitude. This trade-off may be advantageous for cardiac

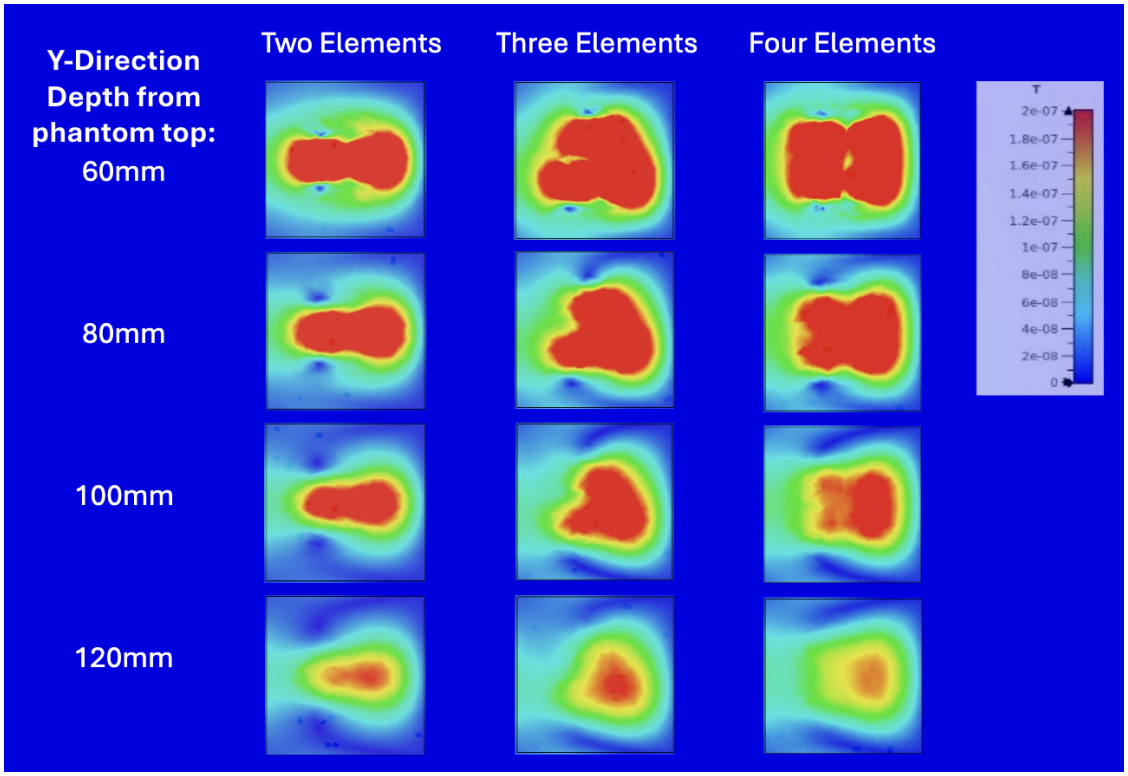


Figure 2.48: Comparison of the B_1 field strength for the two, three, and four element array configurations, across various depths. At depths of 120mm, the area of homogeneous area increases with the amount of coils added to the design.

imaging, where uniform coverage is often prioritised over maximum sensitivity.

The B_1 field maps generated from CST simulations exhibited a characteristic asymmetric smear, which I initially suspected to be related to coil combination artifacts. However, further investigation confirmed that this feature is intrinsic to circular polarisation in the z-direction, rather than a result of receive combination or meshing issues. I systematically ruled out potential sources of error, including meshing strategy (comparing hexahedral and tetrahedral elements), geometric asymmetry in the coil layout, and inaccuracies in matching and tuning. The asymmetry persisted even with a single transmit coil, confirming that the effect originates from the inherent behaviour of circularly polarised fields.

This phenomenon is consistent with findings in the literature. For example, Frass et al.[78] also report asymmetric B_1 field distributions in simulations at 3 T, even for a single loop, especially when using high-dielectric phantoms. These asymmetries arise from the right- or left-handed polarisation of the B_1 field in the presence of

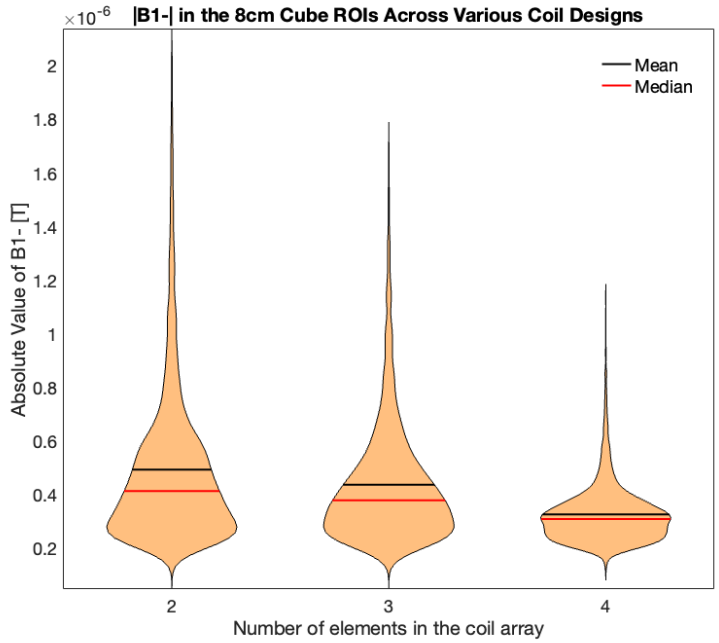


Figure 2.49: Violin plots of voxel-wise B1- values in the heart ROI for two-, three-, and four-element arrays, showing progressively more compact distributions with increasing element count.

dielectric loading, rather than from any fault in coil configuration. Nonetheless, improper matching and tuning can also distort the B_1 field, as demonstrated in Figure 2.50. It is therefore essential to carefully re-match and re-tune the coil following any simulation changes to ensure the validity of field distributions.

2.5.4 Human Phantom Simulations

To evaluate the four-element multi-layer array under realistic anatomical loading, I performed electromagnetic simulations on a high-fidelity human male phantom "Hugo" in CST Microwave Studio (Dassault Systèmes, Vélizy-Villacoublay, France). My aim was to visualise B_1 field coverage, absolute field strength, and transmit field homogeneity across the cardiac region.

At first, I considered two anatomical models, the female phantom ("Laura") and the male phantom ("Hugo") (Figure 2.51). I initially attempted simulations on the Laura model. However, the pronounced chest curvature elevated the coil elements away from the torso surface, reducing coil-to-load coupling and increasing coil-to-coil

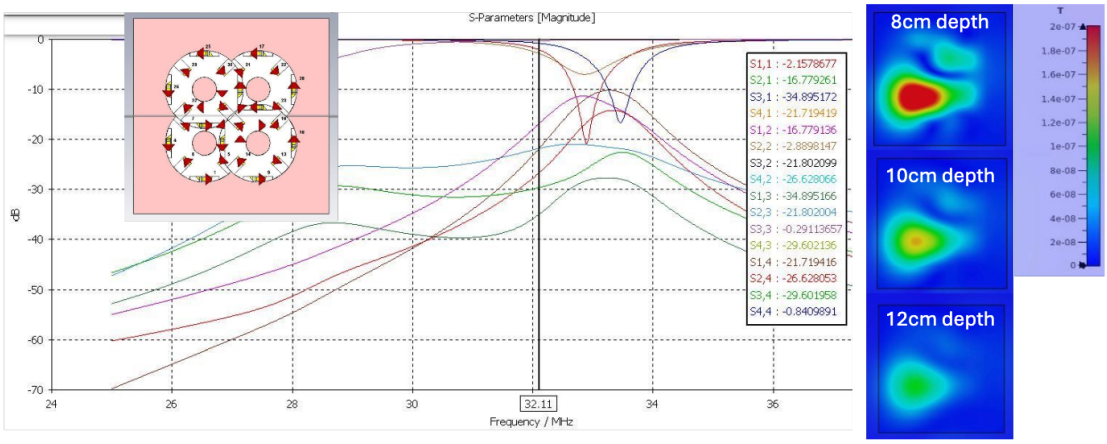


Figure 2.50: An example of bad tuning. Four-element multi-layer array with an 11cm separation has a mismatched matching and tuning shown in the graph. The B_1 field has degraded.

inductive interaction. In this configuration, small changes to tuning capacitances caused large shifts in inductance, making stable matching impractical. While physically this issue could be easily addressed by arcing the flexible PCB to conform to the chest shape, such conformal positioning is challenging to implement in CST.

For this reason, I prioritised simulations on the Hugo model, which has a flatter anterior chest and larger cross-sectional area, resulting in higher coil-to-load coupling and improved inductive stability. Although the Laura model was not used for coil simulation, both human phantoms were segmented to determine the cardiac field-of-view (FOV). As shown in Figure 2.52, the heart extended over approximately -60 mm $< y < -136$ mm in the Laura model and -53 mm $< y < -133$ mm in the Hugo model. These anatomical bounds informed the selection of B_1 visualisation slices in the subsequent analysis, focusing on the FOV slices from 60 mm to 120 mm from the anterior torso surface.

I positioned the four-element array symmetrically over the chest wall of the Hugo phantom, with each element matched and tuned to the ^{13}C frequency (32.11 MHz) (positioning shown in figure 2.51B). S-parameter analysis helped assess coil matching and tuning (S_{11}–S_{44}), and inter-element coupling ($S_{ij},\ i \neq j$), with well-matched performance defined as $S_{ii} < -15$ dB. The simulated B_1^- field maps were exported

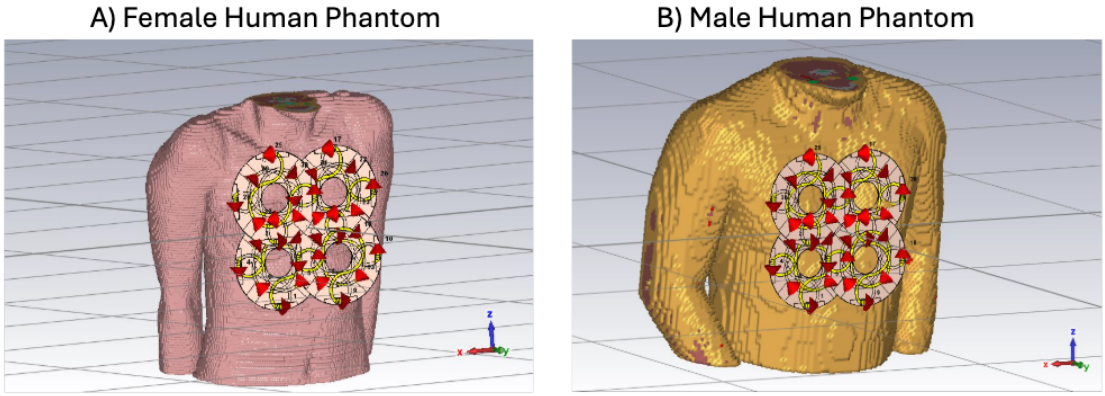


Figure 2.51: CST healthy human phantoms and coil positioning used in my simulations. A) is a female phantom model "Laura" and B) is a male phantom model "Hugo".

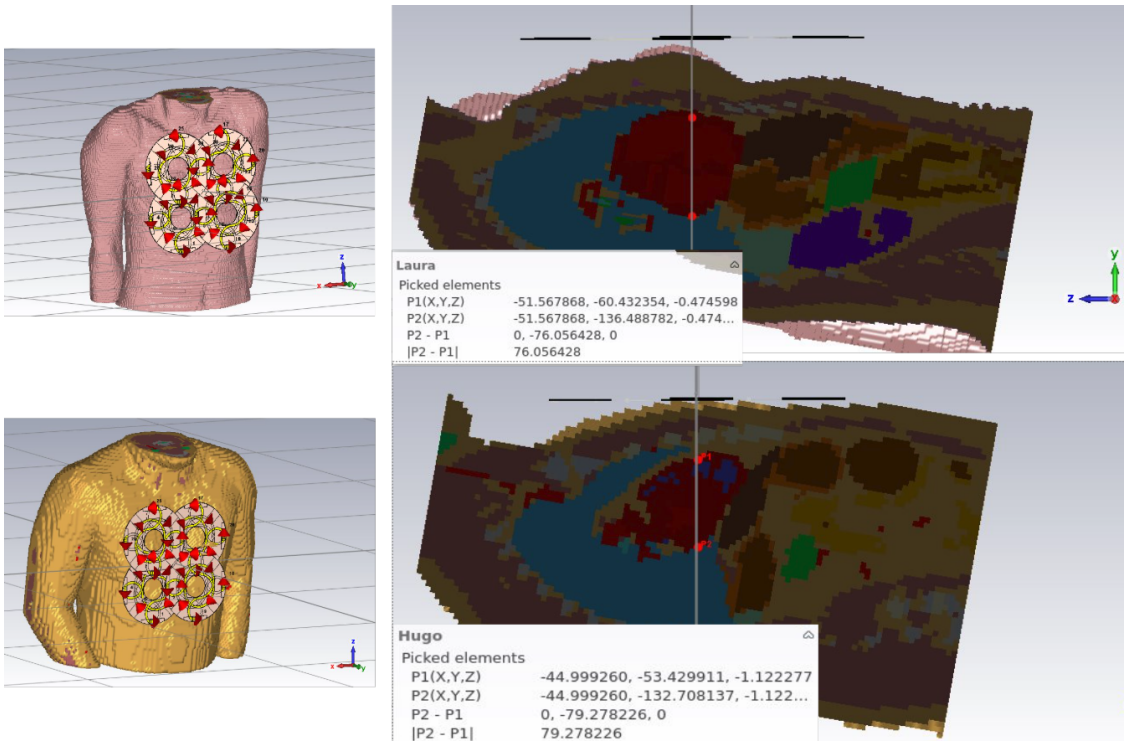


Figure 2.52: The phantom viewed in the x-direction, at the slice where the heart is thickest. The highest and lowest points of the heart in the center axis is highlighted with red dots, corresponding to the picked points shown in the opened tabs.

for magnitude visualisation across axial slices spanning the heart region, and quantitative homogeneity metrics (mean, standard deviation, coefficient of variation (CV), and uniformity = 1 − CV) were calculated within the heart ROI.

All four elements were matched and tuned at the ^{13}C frequency, with $S_{ii} < -15$ dB (Figure 2.53). Inter-element coupling values (S_{ij}, $i \neq j$) remained acceptable,

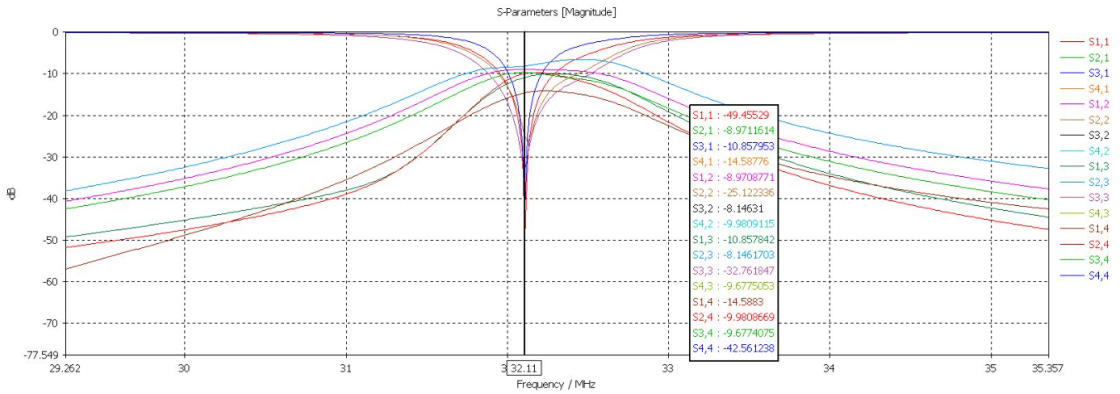


Figure 2.53: S-parameters of the four-element array design placed on the Hugo model. All coils are well matched and tuned (<-15 dB) at the carbon frequency. Inter-element coupling parameters show coupling with the highest parameter S32=-8.14631.

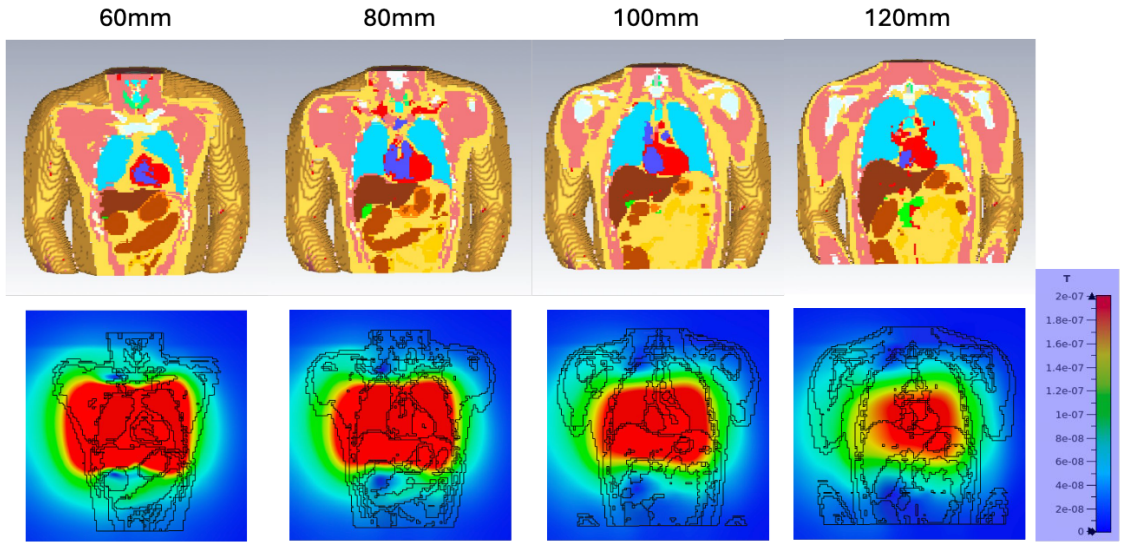


Figure 2.54: Comparison of the B_1 field strength (bottom row, overlayed on the wire diagram of the anatomy outline) of the four element array configuration on the Hugo model across various depths (anatomical slices of the phantom shown on the top row). Through all depths, the area of homogeneous B_1 field covers the heart (depected as the red and blue anatomy at the center of the

with the highest coupling measured between elements 3 and 2 ($S_{32} = -8.15$ dB).

The simulated B_1^- maps (Figure 2.54) showed consistent coverage of the cardiac region across all depths of interest (60–120 mm). The central high-sensitivity zone encompassed the full heart volume, matching the anatomical boundaries determined in Figure 2.52, with no major gaps that would compromise uniform coverage. Qualitatively, the maps confirmed that the geometric configuration from the earlier phantom studies translated well to the anatomical geometry of the Hugo model.

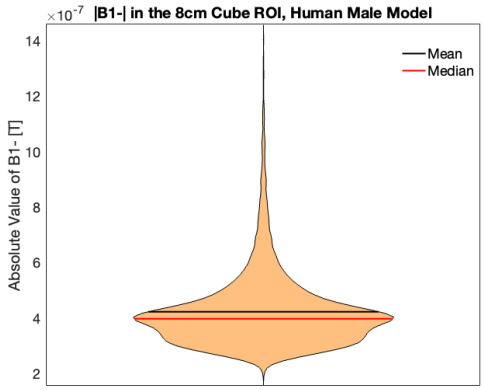


Figure 2.55: Violin plots of voxel-wise B_1 field values in the heart region (as in Figure 2.2) for the four-element array. Thick bulge at the median value with a thin tail as the absolute value of the B1- increases.

Quantitative evaluation within the heart ROI (Figure 2.2) produced a mean B_1^- value of 4.22×10^{-7} T, a standard deviation of 1.34×10^{-7} T, and a coefficient of variation (CV) of 0.3169, corresponding to a uniformity of 0.6831. The violin plot in Figure 2.55 shows a pronounced, thick bulge at the median with a short upper tail, indicating that most voxels fall within a narrow band of B_1^- values. This distribution pattern reflects high field homogeneity in the heart region, consistent with the uniformity metric.

These results confirm that the four-element, 11 cm–spaced configuration maintains strong field coverage and good homogeneity even under high-loading, anatomically realistic conditions. While some residual coupling remains, the uniform coverage achieved across the heart volume supports the viability of this configuration for the final design.

2.5.5 Discussion

The original challenge was to address the B_1^- field drop observed in the centre of the ROI between two elements. In the two-element phantom simulations, I mitigated this dip in field by bringing the coils closer together, settling on an 11 cm centre-to-centre separation. This distance offered a significant improvement

in B_1^- homogeneity without compromising field strength, even though isolation was worse than at the optimum spacing.

Scaling this 11 cm spacing to three- and four-element arrays revealed the same trade-off, where peak B_1^- amplitude decreased relative to the two-element case due to worse coupling, but uniformity improved. The three-element design expanded the field-of-view at depth, while the four-element array maximised uniformity across the heart ROI. On the Hugo human phantom, the four-element configuration maintained both B_1^- uniformity and field strength, with coupling remaining manageable ($S_{32} = -8.15$ dB). The reduced mean B_1^- in larger arrays reflects the redistribution of receive efficiency over a broader, more even coverage. From an application perspective, this is a favourable trade-off for cardiac ^{13}C imaging, where consistent coverage of the heart and surrounding chest tissue is more important than maximising absolute sensitivity at a single point. At 11 cm spacing, inter-element coupling remained within acceptable bounds across all two-, three-, and four-element configurations, demonstrating that the multi-layer design maintains effective geometric decoupling as the array scales in complexity. This consistency across configurations supports the viability of the chosen spacing as a robust foundation for the final cardiac array design.

The choice to focus on the Hugo model, rather than the Laura model, also reflects real-world constraints where varying anatomies require further development. In the CST simulations with the female phantom, the non-flat chest geometry increased the coil–coil distance from the load, amplifying inter-element coupling and making stable tuning impractical. In the Hugo model, the field maps (Figure 2.54) show that the cardiac region is consistently covered across the relevant depth range, with no central voids and minimal variation in field magnitude. This is reflected in the high uniformity value (=0.6831) and the concentrated B_1^- distribution seen in Figure 2.55. Overall, the Hugo simulations strengthen the case for the 11 cm, four-element multi-layer configuration as the baseline for the final build. The next step was to move into in-scanner validations with a phantom, and to confirm that the observed B_1 performance translates to clinical hardware.

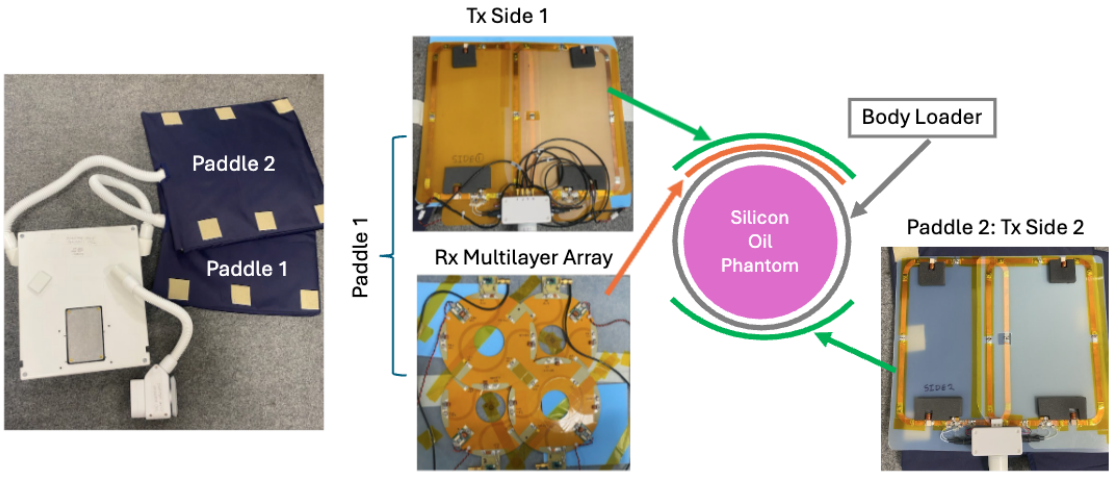


Figure 2.56: Images of the coil assembly and inner workings of the paddles. The schematic with the silicone oil phantom is the axial view of the Tx and Rx layering components in relation to the scanner set up.



Figure 2.57: Multi-layer coil set up on the GE scanner. The Tx paddles are either side of a loading phantom.

2.6 In-Bore Array Validation

This section confirms the functionality of my multi-layer array design through in-bore testing. These in-bore tests confirm successful scanner integration, validate my simulation and experimental results, and benchmark my research coil against a commercially produced chest coil.

2.6.1 Scanner Hardware Requirements

Before acquiring in-bore images, all components of the Tx coil, the Rx coil, the connections between the two, the interface with the scanner and the cabling etc.

needed to be set up for the experiment to work. My main objective was to successfully compile these components into a working coil unit for the GE scanner. In this process, I carried out the Tx coil (provided by industry) in-bore testing, the Rx coil array assembly, and the Tx Rx unit integration with the scanner. I collaborated with industrial partners to build the following components: the Tx Rx active detuning (pin diode), the scanner plug, the interface board, the Tx coil unit and the corresponding interface board connections.

Tx Coil In-Bore Testing

The ^{13}C transmit (Tx) coil was designed and manufactured by PulseTeq to industrial standards. This Tx unit consists of two paddle-shaped PCB structures, each containing two loop elements, as shown in Figure 2.56. The coil operates with a four-way power splitter (Wilkinson Splitter), delivering power to each of the four loop elements. Each loop measures 150×300 mm with a 30 mm radius fillet. Decoupling from body coil is achieved using two ^{1}H traps per element, while decoupling from the receive chain is implemented through active detuning using PIN diodes.

Testing the matching and tuning of the Tx coil was an iterative process because the measurements were dependent on the loading, which alters the resonant frequency and impedance of the coil. In the clinical setting where this coil will encounter many body types of varied loading, precise matching is impractical and transient. It is accepted that there is a practical trade-off between rigorous re-adjustment for every subject and tuning the coil under representative load conditions. Even if the Tx were to be slightly off resonance, SNR is primarily limited by the Rx chain. In this sense, it is more important to prioritise B_1^+ field uniformity and safety compliance rather than maximum power transfer efficiency.

Throughout my project, I used different phantoms as none of the available choices were the right combination of body-mimicking loading, torso sized, chest shaped, and with stable properties. With each phantom, I refined the matching and tuning to my experimental set up. In this section, I outline the Tx matching and tuning methodology for my experimental set up with the silicone oil sphere (Body Phantom Sphere, Model 2360025REV3, 10.64kg, GE Healthcare) and the GE body

loader (Body Cylinder Assembly, Model 2360037REV4, 7.29kg, GE Healthcare), which was used in the finalised experimental set up. However, at earlier points in my DPhil prior to the final phantom set up, I had my industrial supervisor act as the load within the bore. Although the human loading results are not included as the coil has progressed significantly since then, there are some learning points that would give insight to the coil unit. I noted that the breathing motion introduced variations in the measured matching values with further sensitivity to the distance between the coils and the body, as well as to the degree of bending in the flexible paddle elements. This sensitivity arises from the large loop geometry of the Tx coil, and indicated future challenges in robustness as this coil is translated into the clinic.

Returning to the phantom set up (figure 2.57), to ensure that the Tx unit was sufficiently matched and tuned for my experimental arrangement, I acquired S-parameters from the Tx, in and out of the bore with no active scanning. I ran initial tests on the two-paddle Tx coil without the combiner, with each paddle assessed independently. The active detuning circuit was powered using a 10 V, 0.3 A dual-output bench power supply (EL302RD, Aim-TTi, Huntingdon, UK). To characterise the Tx unit, I took individual measurements of each paddle both inside and outside the magnet bore, and with and without the presence of the paired paddle, all of which significantly affected the performance of the hardware. Eight possible measurement configurations were tested: the bottom paddle measured out of the bore 1) with and 2) without the presence of the top paddle, the top paddle measured out of the bore 3) with and 4) without the presence of the bottom paddle, and these four configurations repeated inside the bore. Since each paddle contains two independent loops, the left one was consistently connected to port 1 of the network analyser and the right to port 2.

First measurements from the separated channels indicated that the coil was tuned between the Siemens ^{13}C resonance frequency (30.98 MHz) and the GE resonance frequency (31.22 MHz). Consequently, the tuning was shifted to optimise for the GE scanner.

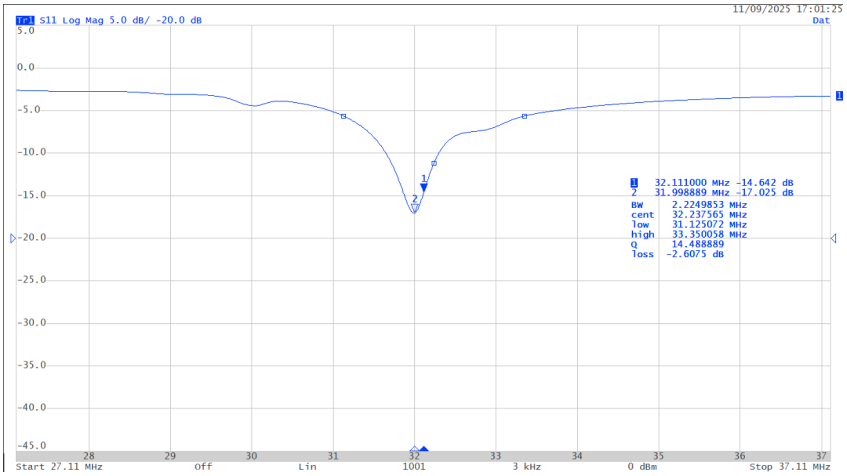


Figure 2.58: Final S11 measurement of the Tx paddles taken at the isocenter. Port 1 was connected to the Tx before the splitter, showing the reflected measure from the combined Tx line.

Further tuning and matching were refined through an iterative process: measurements were acquired in-bore with the phantom, adjustments were made on a bench outside the magnet (due to the absence of an MR-safe soldering station), and the coil was retested in-bore. In the final interaction, I tuned the Tx to the ^{13}C resonance frequency at the isocentre using a variable capacitor and by adjusting the overlap between loops. The most challenging aspect of this procedure was the observed shift in resonance frequency between out-of-bore and in-bore conditions. The process was repeated, compensating for the ≈ 0.1 MHz shift until the out-of-bore matching reached around -15 dB at 32.11 MHz. The final Tx S-parameter measurements for the fully integrated system at the isocenter are presented in figure 2.58, where the impedance match was -14.6 dB at 32.11 MHz, reaching a minimum of -17.0 dB at 32.0 MHz. These values were deemed sufficient for experimental testing, as the immediate priority was to achieve a functional integrated setup.

Rx Coil Array Assembly

I constructed the four-element receive array with an element-to-element centre separation of 11 cm and mounted on a custom PCB support. Elements 2 and 4 were positioned directly on the PCB surface, as shown in Figure 2.59A. In contrast, Elements 1 and 3 were elevated by 5 mm using a foam spacer to accommodate the protrusion of the integrated proton traps and to preserve level coil positioning, as

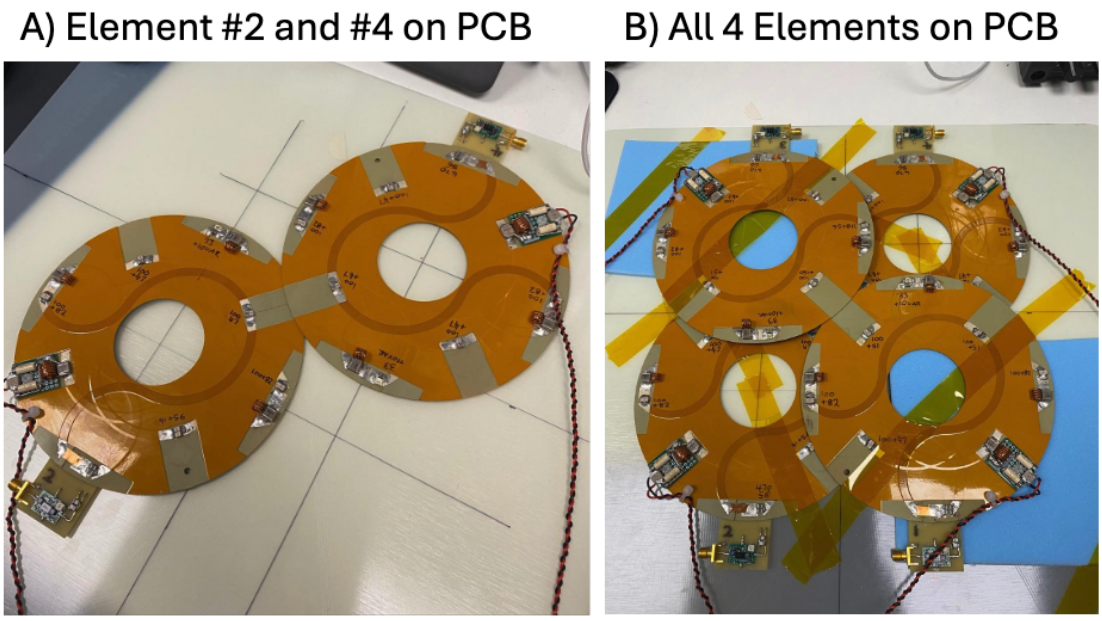


Figure 2.59: Assembly stages of the four-element Rx array on the printed circuit board (PCB). (A) Intermediate stage of construction, showing Elements 2 and 4 mounted directly on the PCB surface. (B) Final configuration, with Elements 1 and 3 elevated by a 5 mm foam spacer to accommodate the protrusion of the proton traps

illustrated in Figure 2.59B. Once the physical placement of all elements was finalised, the coaxial cables from each Rx element were soldered, routed through the PCB assembly, and connected to the transmit system in preparation for integrated in-bore testing.

Tx/Rx Unit Integration with the Scanner

The integration of the transmit and receive subsystems into a fully functional scanner-ready unit required electrical, mechanical, and design considerations. This work was undertaken in collaboration with Ian Nicholson (MRF Innovations), who served as a consultant for PulseTeq and provided engineering supervision for this phase of the project. While the Tx and Rx coils were developed and tested independently, integrating these components required a custom interface box. This interface featured a single Tx line and four Rx lines, connected to the scanner via a vendor-specific GE plug.

The on-coil transmit/receive interface box was designed to translate the scanner's RF, DC, and control signals into a phase-coherent and safety-compliant form. This

also provides an RF path regulator and circuit protection layer. It supplies the DC bias voltage required for active detuning circuits, and houses the coil identification (coil ID) hardware, enabling the scanner to recognise and configure the coil via a vendor-specific coil file. The coil file is a digital configuration that maps each Tx and Rx element to its corresponding hardware channel, embeds frequency/tuning and calibration constants, and enforces safety limits. The internal layout of the interface box, shown in Figure 2.60, illustrates the routing of RF, DC, and control lines between the scanner and the coil array.

As I assembled the system, several practical challenges arose. For example, the coil plug orientation on the interface board required reversing to achieve correct cable sequencing. Addressing such adjustments highlighted the need for design flexibility throughout this project. Another challenge was incorporating the active Tx/Rx detuning unit on the limited handle surface underneath the matching and tuning capacitors. Active decoupling can be implemented at any capacitor in the loop. After evaluating the geometry of the array, I selected the capacitors adjacent to the matching capacitors for decoupling, allowing the preamplifiers to remain on the handle surface and the decoupling circuits to be mounted on an extended handle section, as seen in figure 2.61.

I began the integration process by connecting only the Tx coil to the scanner. Upon connection, the scanner recognised the coil ID and loaded the corresponding coil file. Then I matched and tuned the Tx as described in the 'Tx Coil In-Bore Testing' section. Next, I connected the Rx array through the purple cabling with SMB endings seen in figure 2.62.

The final integrated system is summarised in the block diagram in Figure 2.63, which illustrates the routing of signals from the scanner to the Tx and Rx elements. Although I utilised four Rx lines, this configuration supports up to eight channels and can be adapted for alternative array geometries, additional receive elements, or another Rx paddle for further developments.

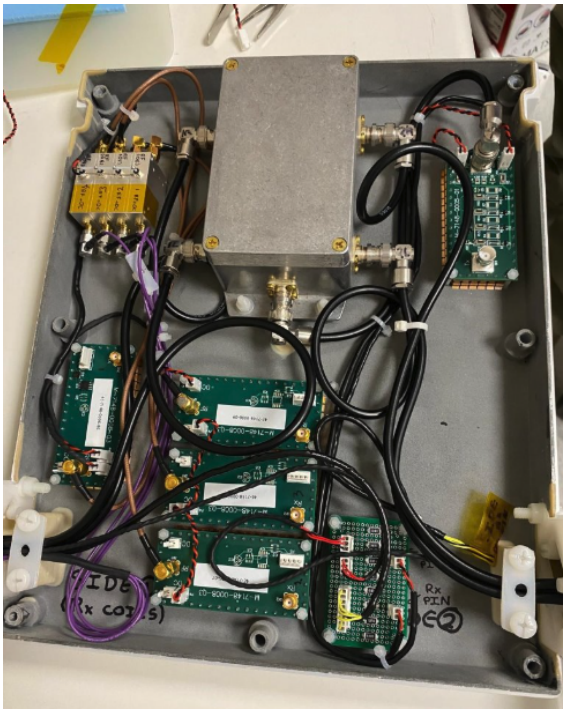


Figure 2.60: Internal view of the interface box, showing routing of RF, DC, and control signals.

2.6.2 Multi-layer Array: Proton difference map

In-bore scanner tests verified the compatibility of my multi-layer array design with the proton body coil. By taking two proton images, one using the GE body coil of the phantom set up, and another with the phantom set up plus the Tx Rx coil in the magnet, I could take a difference image. These proton difference maps are implemented as part of the quality assurance and safety testing procedure of novel coils, which is fully detailed in the following references [79, 80]. Aligning with PulseTeq's coil development specifications, I aimed for the variation of this difference image to be within 5%.

I took two fast-spin echo type proton 3-plane localisers with and without the multi-layer array set on the phantom (frequency FOV = 50cm, 20 slices of thickness = 20mm with no spacing in all 3 planes, 256 x 256 voxels, TR = 4.492 ms, and TE = 1.132 ms). I placed the two coil paddles left and right of the phantom to avoid mis-registration of the images. Figure 2.64 shows the proton signal difference maps in percentage, with and without the multi-layer array coil in place. I used

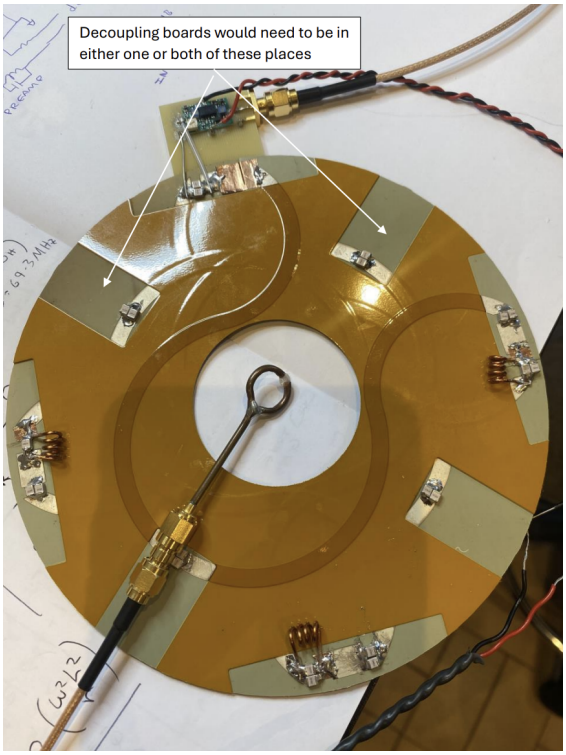


Figure 2.61: Arrows indicating potential locations for Tx/Rx decoupling boards, chosen according to the array geometry.

the scan without the coil to apply a mask around the silicon oil sphere phantom, and to calculate the percentages of the difference map. From visual inspection, there is a thin yellow line along the edge of the phantom, indicating high difference percentage, but this is due to the phantom moving very slightly inside the loader when I pulled off the coil that was strapped to the set up. There is also some coupling effect from the body coil as there is a crescent shaped artifact through the sphere at each slice depth. However, when taking the average values within a heart-sized FOV (8x8x8cm) at the centre of the sphere, the overall mean was 3.39% with a standard deviation of 0.96%. This was within the expected variation.

Increasing the amount of proton traps (currently four in the Tx and three per element in the Rx) may be required in future implementation. However, at the depths I was interested in, there seemed to be good stability for the purpose of this experiment. This minimal coupling with the proton hardware indicated that the carbon elements were protected from unwanted induced currents. Therefore the carbon excitation and reception were expected to localise to the carbon coil

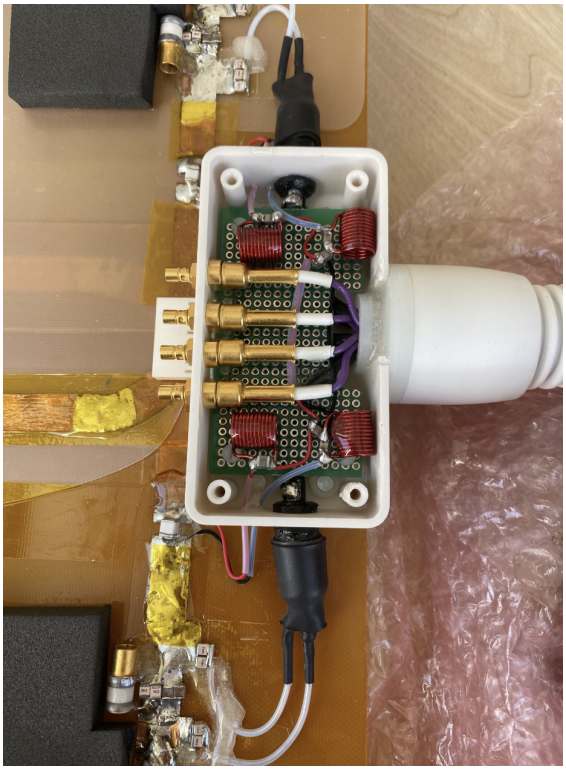


Figure 2.62: Internal wiring connections within the coil cabling unit.

elements, supporting accurate B_1 mapping without cross nucleus contamination. In the future, proton+carbon workflows would be possible without reconfiguration.

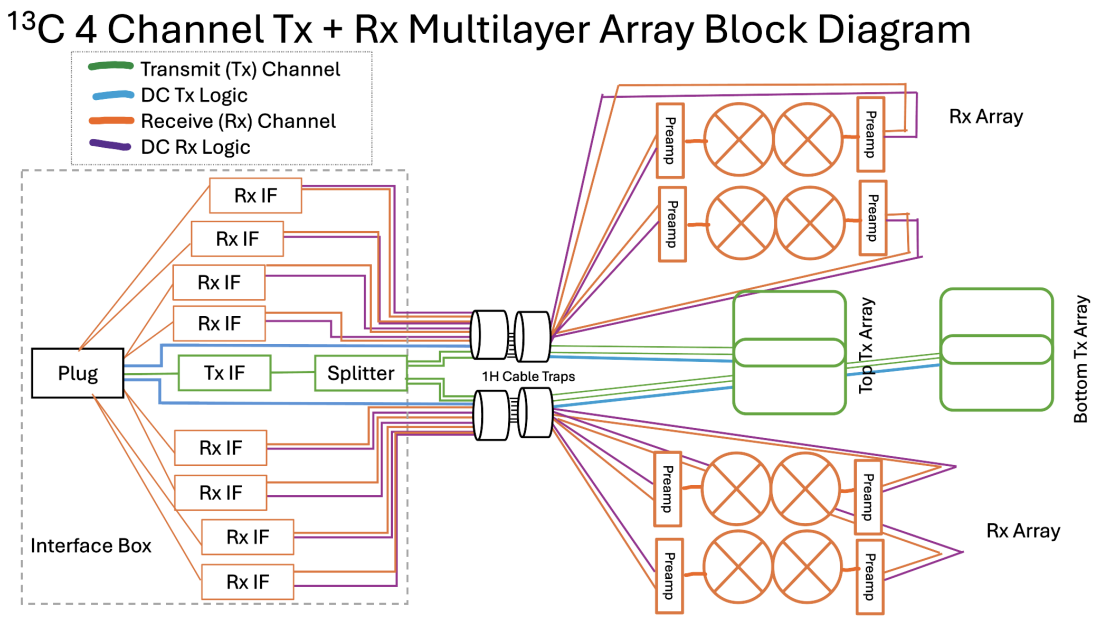


Figure 2.63: Block diagram of the multi-layer coil array configuration, showing Tx and Rx signal paths and supporting up to eight channels.

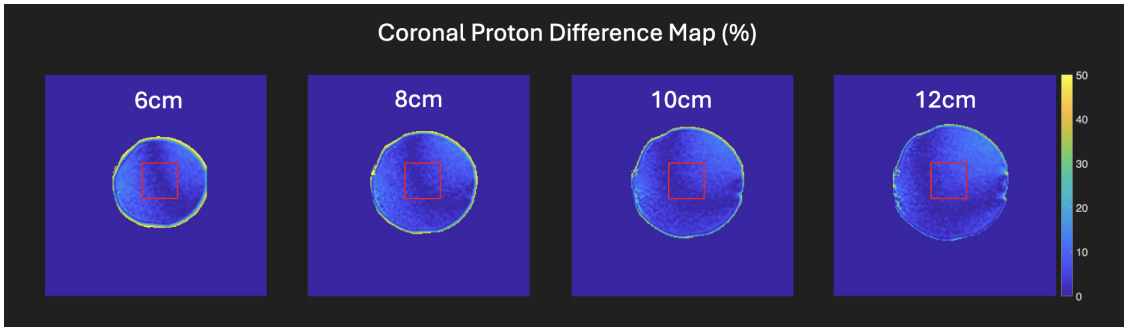


Figure 2.64: Proton difference maps across coronal slices at 6, 8, 10, and 12 cm depth. The heart ROI is delineated by red squares.

2.6.3 Multi-layer Array: B1 Field Maps

I aimed to validate the simulated receive sensitivity (B_1^-) of the multi-layer coil against experimental in-bore measurements. This simulation-to-experiment comparison was non-trivial because the CST simulations assumed electromagnetic reciprocity, where the same coil serves as the transmitter and receiver, resulting in identical B_1^+ and B_1^- profiles. In contrast, my experimental setup used a separate two-loop assembly for transmit and the multi-layer elements as receivers. Consequently, the raw magnitude data acquired by the scanner represented the product of the inhomogeneous transmit field and the local receive sensitivity. To decouple the receive profile, I acquired a Bloch-Siegert B_1 map to characterise the B_1^+ field (section 2.2.6). I then used the B_1^+ map to generate a correction map, which was applied to the magnitude image data to decouple the transmit influence and extract the receive sensitivity profile.

B_1^+ Field Characterisation for Correction Mapping

Prior to carbon imaging, I performed a proton localiser scan and a Transmit Gain (TG) calibration. This vendor-specific pre-scan measured B_1^+ efficiency under the current load and updated the transmit scaling (dB) to achieve the prescribed flip angle at the silicone oil frequency of the phantom. In instances where the maximum available transmit gain was insufficient to achieve 90°, I adjusted the nominal flip angle based on the ratio required for an effective 90° excitation.

To map the transmit field at the ^{13}C frequency, I utilised a Bloch-Siegert excitation followed by a spiral readout. This 2D multi-slice sequence acquired 16 pairs of scans (32 total acquisitions) with a ±2 kHz frequency shift relative to the resonance frequency. The scan parameters were: Coronal FOV = 40 cm, Slice Thickness = 20 mm, 4 slices (beginning 6 cm from the phantom top), TR = 7018.2 ms, and FA = 122.4°. The long TR value was based on previous measurements in our laboratory, which indicated a T_1 of approximately 1 s for the silicone oil phantom. A repetition time of 7018.2 ms greatly exceed this value, ensuring

complete longitudinal recovery between excitations and avoiding T_1 -dependent bias in the measured B_1^+ maps. The entire acquisition took a total of 9 hours 20 minutes.

The raw data consisted of 150×32 total repeats, reconstructed at a coarse resolution of 20×20 voxels ($2 \times 2\,\text{cm}$ resolution). The averages were first summed in the complex domain, after which the magnitude of the combined signal was calculated. The transmit field map, $B_1^+(x, y, z)$, was then calculated using the blosi_b1map function provided by GE Healthcare. To improve the map's computational consistency, the B_1^+ data was smoothened using cubic interpolation in the x-y plane, while z-slices were processed independently. This involved a 4×4 neighbourhood fit to a cubic polynomial to estimate the value at each spatial point.

Because my multi-layer coil utilised separate hardware for Tx and Rx, it was essential to account for spatial inhomogeneities in the Tx field. The peak RF amplitude required for 90° excitation ($B_{1,\text{desired}}$) was extracted from the mns_prescan header, which for this setup was determined to be $29.96\,\mu\text{T}$. Using a nominal flip angle (α_{nominal}) of 90°, I then calculated the actual flip angle, $\alpha(x, y, z)$, using the following relationship:

$$\alpha(x, y, z) = \alpha_{\text{nominal}} \cdot \left(\frac{B_1^+(x, y, z)}{B_{1,\text{desired}}} \right) = 90^\circ \cdot \left(\frac{B_1^+(x, y, z)}{29.96\,\mu\text{T}} \right)$$

Applying the Bloch-Siegert B_1^+ values gave me a map of the actual flip angle delivered by the transmit coil. To compensate for the effect of flip angle deviations on signal amplitude, I derived a voxel-wise correction factor based on the ratio of the sines of the nominal and actual flip angles:

$$\text{Correction Factor}(x, y, z) = \frac{\sin(90^\circ)}{\sin(\alpha(x, y, z))} = \frac{1}{\sin(\alpha(x, y, z))}$$

The resulting correction map was applied to the magnitude map to decouple the transmit and receive profiles.

Transmit Profile and Correction Map

The transmit B_1^+ maps (Figure 2.65, first row) demonstrated high homogeneity within the $8 \times 8 \times 8\,\text{cm}^3$ heart-sized FOV (uniformity = 0.87) as expected from

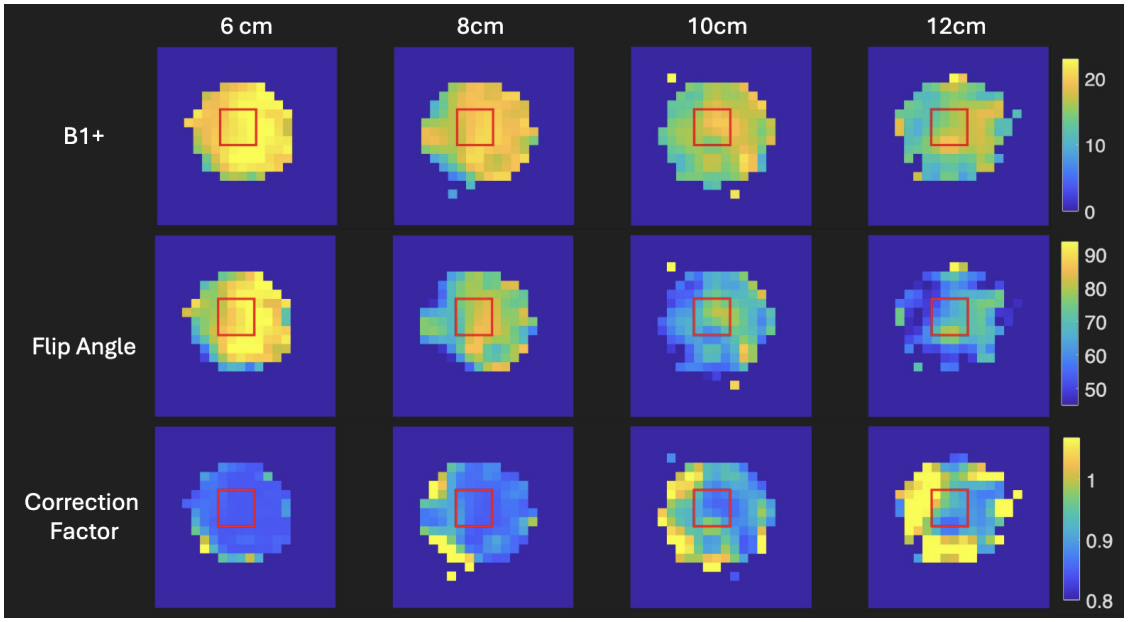


Figure 2.65: The B1+ map (uT), the calculated flip angle map (degrees) and corresponding correction factor for the transmit inhomogeneity for the multi-layer coil.

the large two loop design. Within this FOV, I measured a mean B_1^+ of 18.80 μT (Table 2.3). Flip angle maps (Figure 2.65, second row), which I derived from the measured B_1^+ field, revealed deviations from the nominal 90° excitation as the depth increased. Based on the flip angle map, the correction map was calculated and is shown in figure 2.65 in the third row. The highest correction values were primarily located outside the heart-sized FOV. Within the field of view, the corrections were modest and followed the expected depth-dependent profile of the surface coil elements.

Table 2.3: Multi-layer coil magnitude, transmit B_1^+ (μT), and corrected receive B_1^- (a.u.) field statistics measured within the heart ROI.

Heart FOV	Peak Value	Mean	Standard Dev	CV	Uniformity
Magnitude (a.u.)	1.47×10^9	7.28×10^8	2.58×10^8	0.35	0.65
Transmit B_1^+ (μT)	22.93	18.80	2.50	0.12	0.87
Corrected B_1^- (a.u.)	1.26×10^9	6.37×10^8	2.08×10^8	0.33	0.67

Magnitude Data and the Extraction of the B_1^- Receive Profile

To isolate the receive sensitivity (B_1^-) from the influence of the transmit field, I applied the correction factor derived in the previous section (Figure 2.65) to the

raw magnitude data (Figures 2.66, 2.67). The magnitude images were generated by summing the complex values of the Bloch-Siegert acquisitions, with the Rx channels combined using a Root-Sum-of-Squares (RSS) algorithm. I applied a 15% maximum signal threshold to generate masks for subsequent visualisations.

In the coil-wise magnitude images (Figure 2.66), I observed spatial variations reflecting the individual element sensitivities. Coil 1 demonstrated the highest sensitivity, with the highest signal intensity localised in the upper-left quadrant at a 6 cm depth due to its closest positioning to the phantom. As expected for surface coil designs, signal intensity across all elements attenuated over the depths of 8, 10, and 12 cm. Coil 4 showed the lowest SNR despite its physical orientation and bench-top tuning (S_{11}) being consistent with the rest of the array. The lower SNR observed for this element is likely due to a 20° deviation in the S_{21} phase, which measured 64.46° compared to approximately 85° across the remaining elements (Appendix Figure 2.89). Because the preamplifiers are operated as conventional low-noise amplifiers optimised for a 50 Ω source impedance, a phase change could shift its operating point away from the minimum noise figure condition and produce a modest but measurable SNR penalty on that channel. The elevated noise floor in channel 4 is consistent with this interpretation (Figure 2.66). The potential sources of the mismatch include capacitors that drifted slightly from their nominal value due to thermal stress during soldering or a minor difference in cable length between this channel and the others. However, the overall array performance remained within acceptable bounds, and phase consistency was identified as a parameter to further control in future iterations of this design.

The high signal feature from coil 1 was also visible in the coil-combined magnitude image at the same depth (Figure 2.67). A dip in signal intensity appeared at the centre of the 6 cm slice, which I attributed to insufficient coverage due to the physical separation between receive elements. At greater depths (8, 10, and 12 cm) the signal dropped off rapidly, consistent with the expected decay in sensitivity further from the surface coil.

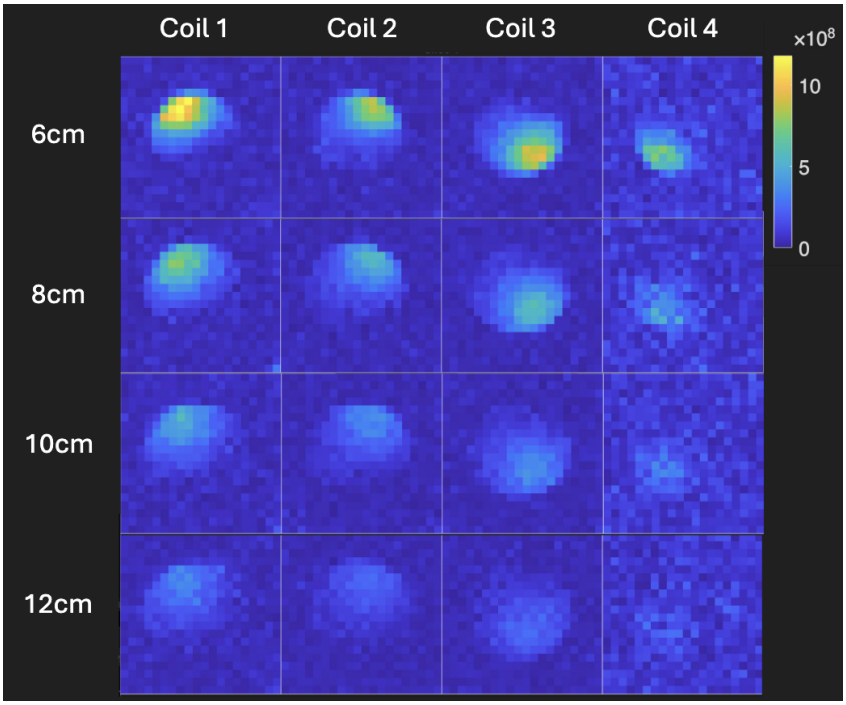


Figure 2.66: Coilwise magnitude images of arbitraty units across coronal slices at 6, 8, 10, and 12 cm depths from the top of the phantom.

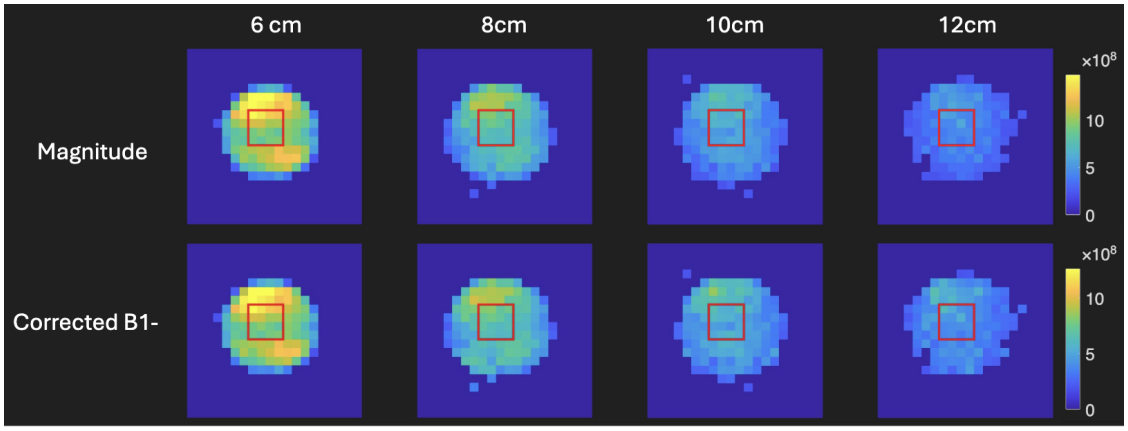
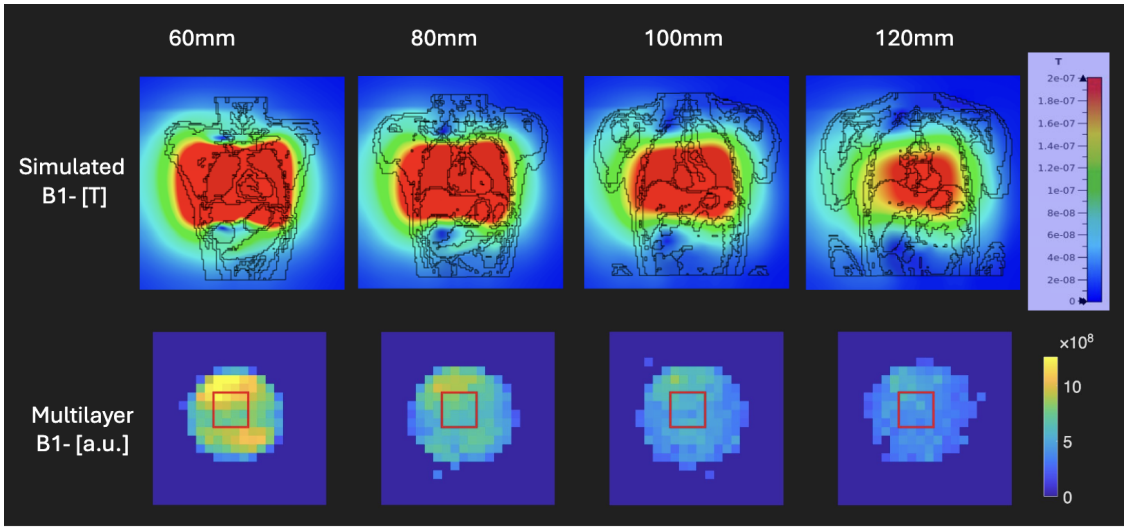


Figure 2.67: The magnitude and the corrected receive B1- images across coronal slices (6–12 cm) using the multi-layer coil.

The receive B_1^- map appeared visually similar to the uncorrected magnitude image. Within the heart ROI as previously defined, I measured peak values of 1.49×10^9 [arbitrary units] for the magnitude image, and 1.26×10^9 [a.u] for the isolated receive field (figure 2.67). The corresponding mean values were 7.28×10^8, and 6.37×10^8. The calculated uniformity $(1 - \text{CV})$ was 0.65 for the magnitude, and 0.67 for B_1^-.



Heart FOV B1-	Mean	Standard Dev	CV	Uniformity
Simulation	4.22 e^-7 T	1.34 e^-7 T	0.32	0.68
Multilayer	6.37 e^8 a.u.	2.08e^8 a.u.	0.33	0.67

Figure 2.68: Receive B_1 maps comparison between the simulated human phantom model and the in-scanner multi-layer coil experiment. Table of the B1- field statistics within the heart ROI.

Simulation vs Experimental B_1^- Receive Profile

Finally, I compared the experimental B_1^- maps to electromagnetic simulations performed on a human phantom model in CST, across the same coronal slices at 6, 8, 10, and 12 cm depth (Figure 2.68). While the simulated maps appeared more homogeneous due to the lack of experimental noise and B_0 inhomogeneities, the calculated uniformity values were nearly identical (0.68 for simulation vs. 0.67 for experimental data). This agreement validates the design hypothesis that the multi-layer geometry improves field uniformity across a cardiac-sized region of interest.

2.6.4 Benchmarking Against Commercial Carbon Coils

Next, I evaluated the performance of my multi-layer array relative to a commercially produced coil. While two commercial carbon coils were available for comparison, a two-loop chest coil developed by PulseTeq Ltd as previously described in the motivation section 2.2.1 (Figure 2.1) and a rigid chest array from Rapid Biomedical GmbH, I chose to directly compare against the Rapid coil. This decision was

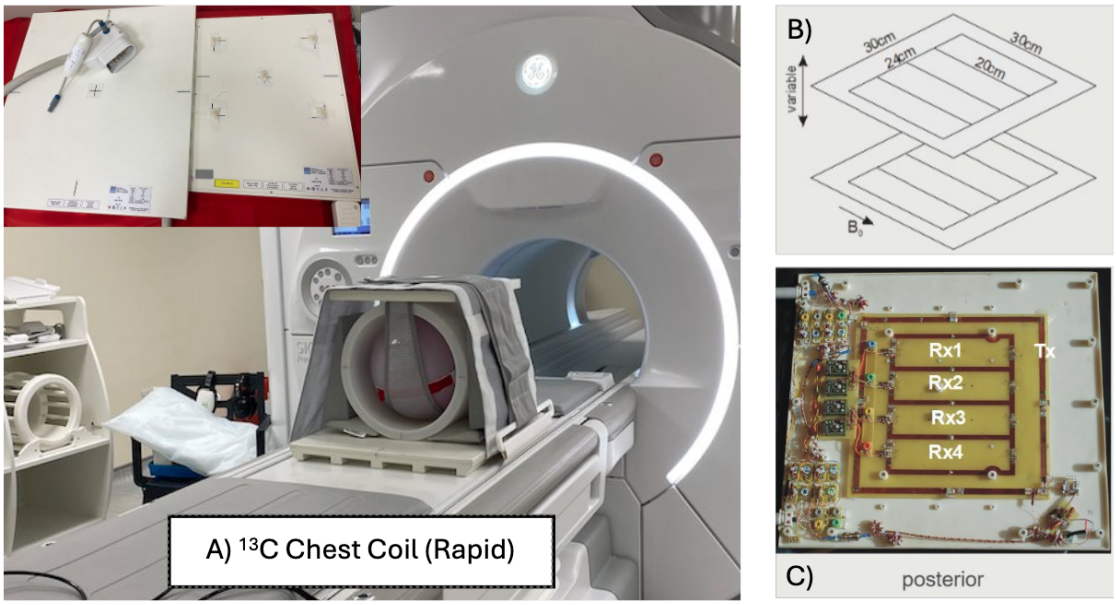


Figure 2.69: A) The set up for the imaging protocol including the coils, phantom and body loader are shown for the Chest Coil from Rapid. The coil is shown in the upper left corners. The schematics of the top and bottom slab and dimensions are shown in B. The Tx and Rx components of the posterior slab is shown in C.

based on the greater similarity in hardware configuration as both my array and the Rapid coil use separate transmit and receive components, and each has four receive channels positioned on the anterior paddle.

Figure 2.69 illustrates the external casing, internal schematic layout, and scanner setup for the Rapid chest coil. This industry-made coil includes a total of eight receive channels, with four loops embedded in both the anterior and posterior elements. The four receive loops are arranged in a ladder configuration, providing continuous coverage across the chest surface. The following results include data from all 8 elements. Data was collected with the same sequence and identical post-processing as the multi-layer results. Further, this subsection compares the corrected receive maps of the two coils, and the SNR achieved per slice within the heart ROI. Noise levels were quantified as the standard deviation of the signal within the four corners of the corrected receive maps. Although the magnitude reconstruction produces Rician-distributed background noise, this approach provided a consistent estimate for comparing the relative SNR of the multi-layer and commercial coils. These values served as the denominator for the

signal-to-noise ratio (SNR) calculations. For both the multi-layer and commercial coils, the SNR was calculated using the summed magnitude data from 32 total acquisitions. Within the resulting SNR maps, the mean and standard deviation were extracted from a heart ROI identical to that defined in Figure 2.64.

With the Rapid chest coil, I acquired a proton localiser, TG calibration, and spiral-based transmit B_1 maps, where the set up is shown in figure 2.69A. For the Rapid Chest Coil, Coronal FOV = 40cm, Slice Thickness = 20mm (AP), Slices = 4 (starting at 6cm from the top of the phantom), Freq Dim = S/I, TR = 4000, FA = 113.4, and NEX = 150 with a total acquisition time of 5 hours 20 minutes. Rx coils were handled with a RSS magnitude combine.

The Rapid Chest Coil was processed as for my multi-layer coil. From the Bloch Siegert scan, I acquired the transmit map (figure 2.70, first row), which was highly homogeneous with calculated uniformity of 0.88 and 26.24 μT for the peak value (Table 2.4). The transmit profile was then used to calculate the flip angle (figure 2.70, second row) and correction maps (figure 2.70, third row). This correction map was used to decouple the receive B_1 from the magnitude data, shown in figure 2.71. The coronal depths and slice thickness both match those of the multi-layer visualisations in figure 2.67. I then repeated the statistics (peak value, mean, std, CV, and uniformity) on the same heart ROI for the magnitude data in table 2.4.

At a 6 cm slice depth, the magnitude image appears very uniform, but the signal intensity drops off sharply in deeper slices (Figure 2.71). As expected from a large two-loop Helmholtz configuration, the transmit field is highly homogeneous, resulting in a receive B_1^- map that closely mirrors the magnitude image. In both the magnitude and receive maps, I noticed a bright region shifted toward the right-hand side. This asymmetry was due to a malfunction of the right-most receive element, which did not contribute to the measured signal during the scan. The hardware failure was confirmed after the experiment. Within the heart ROI, I measured peak values of 4.95×10^8 [a.u] for the magnitude image, and 4.59×10^8 [a.u] for the receive field in Table 2.4. Mean values were 2.79×10^8 [a.u] for the magnitude

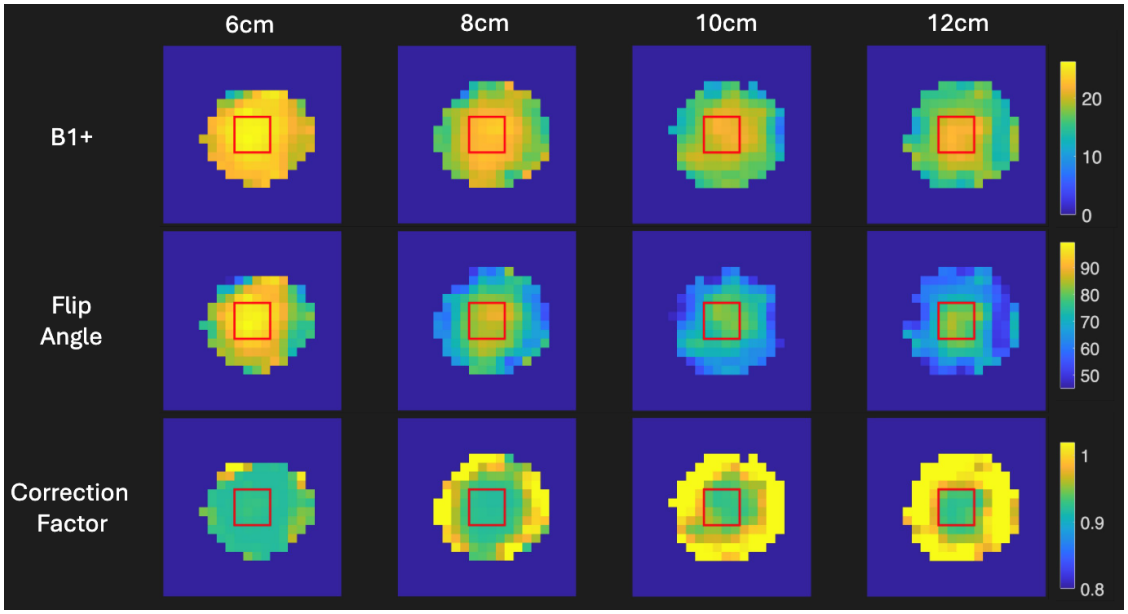


Figure 2.70: B1+ map, flip angle map, and corresponding correction factor used to decouple the transmit inhomogeneity for the Rapid coil magnitude data.

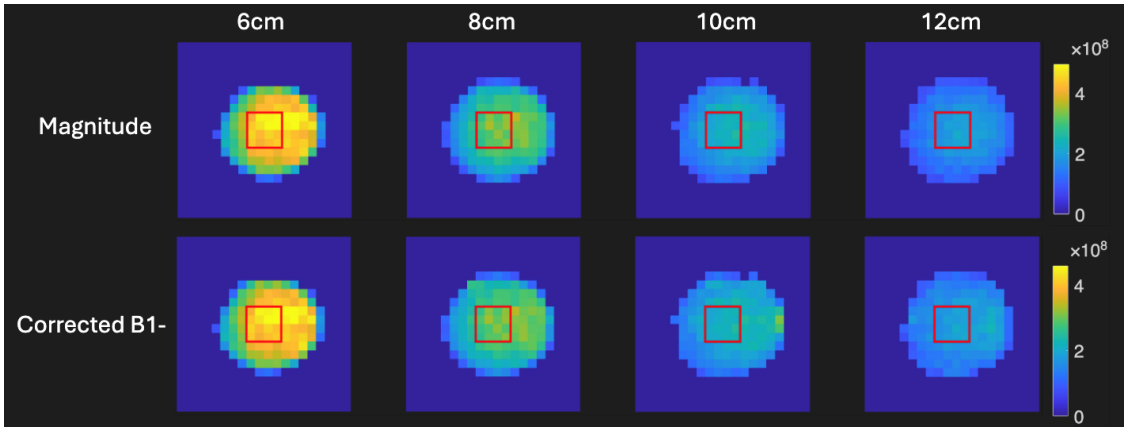


Figure 2.71: Magnitude, and the corrected receive B1- images across coronal slices (6–12 cm) for the Rapid coil.

image and 2.64×10^8[a.u] for the receive field. Calculated uniformities ($1 - \mathrm{CV}$) were 0.64 for the magnitude, and 0.65 for the B_1^- field.

When directly compared to my multi-layer array, the Rapid coil shows slightly higher B_1^+ uniformity and amplitude within the heart ROI: uniformity of 0.88 for the Rapid coil versus 0.87 for the multi-layer coil, peak B_1^+ of 26.24 versus 22.93 µT, and mean B_1^+ of 21.80 versus 18.80 µT. However, the multi-layer array achieved more than twice the receive sensitivity, with a B_1^- mean value of 6.37×10^8[a.u] compared to 2.64×10^8 [a.u.] in the Rapid coil (Table 2.3, 2.4).

Table 2.4: Rapid coil magnitude, transmit B_1^+ (μT), and corrected receive B_1^- (a.u.) field statistics measured within the heart-sized ROI.

Rapid	Peak Value	Mean	SD	CV	Uniformity
Magnitude	4.95×10^8	2.79×10^8	1.00×10^7	0.36	0.64
Transmit B_1^+ (μT)	26.24	21.80	2.92	0.12	0.88
Corrected B_1^- (a.u.)	4.59×10^8	2.64×10^8	8.99×10^7	0.35	0.65

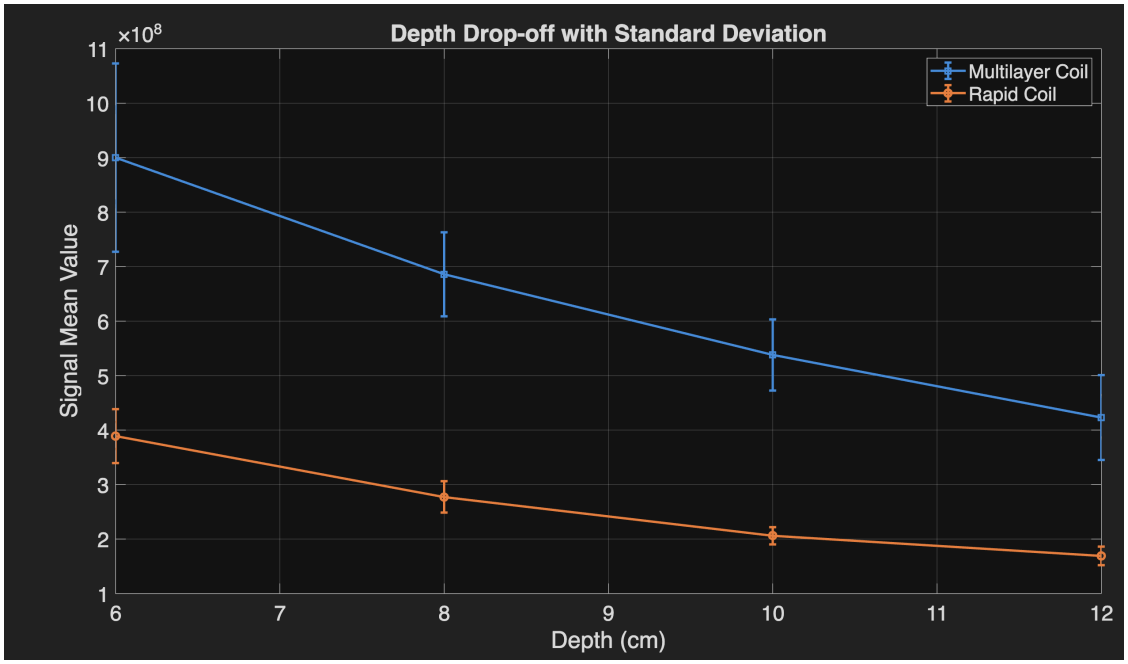


Figure 2.72: Comparison of the mean values of the corrected B_1^- signal in the heart ROI across slice depths for the multi-layer and the Rapid coil.

Table 2.5: SNR Performance Comparison. Summary of mean SNR and standard deviation ($N = 32$) for the multi-layer array and the commercial Rapid chest coil. Values are reported for individual slices (6–12 cm depth) and the aggregate 3D field of view.

Mean SNR in Heart FOV	Multilayer Array Coil	Rapid Chest Coil
Slice depth 6 cm	20.74 ± 3.99	18.56 ± 2.17
Slice depth 8 cm	15.82 ± 1.78	12.67 ± 1.19
Slice depth 10 cm	12.40 ± 1.51	9.84 ± 0.64
Slice depth 12 cm	9.76 ± 1.79	9.25 ± 1.02
All 4 slices pooled	14.68 ± 4.79	12.96 ± 3.75

These differences in receive sensitivity were particularly evident when evaluating depth-dependent performance. Figure 2.72 compares the mean signal intensity across slices for both coils. While the Rapid coil maintained a more uniform profile across the measured depths, the multi-layer array sustained significantly higher signal throughout the volume. Notably, the mean signal intensity at a 12 cm depth for the multi-layer array exceeded that of the Rapid coil at a 6 cm depth. This indicated that the multi-layer design offered superior sensitivity at depth, which is a prerequisite for cardiac imaging applications. Furthermore, the aggregate SNR within the cardiac-sized FOV was higher for the multi-layer coil (14.68 ± 4.79) than for the Rapid coil (12.96 ± 3.75) as seen in table 2.5. This was consistent with the robust receive profile of the multi-layer design, where mean signal levels within the heart ROI were more than double those of the Rapid coil across the 6–12,cm range, as shown in Figure 2.72.

My current design prioritised flexibility and depth performance, achieving strong coverage with fewer elements, but leaves room for future optimisation. In particular, I inherited the transmit hardware and did not optimise the Tx design. Starting with a more conventional single-loop Helmholtz could improve performance. Additionally, I did not simulate for flexed Rx geometry, and there is scope to optimise the spacing and positioning of elements when curved around the torso. There is further scope to create another array paddle for the anterior side of the chest to create a Helmholtz pair. These results highlight the trade-offs between uniformity, penetration, and flexibility, and suggest that with more refined Tx design and expanded Rx coverage, the multi-layer array can be further optimised for cardiac applications.

Summary and Clinical Outlook

In this chapter, I presented the development and validation of a novel multi-layer RF coil array tailored for HP ^{13}C MRI. This work addresses the limitations in sensitivity and field uniformity encountered with conventional simple loop coils, particularly in the context of cardiac imaging. I designed and simulated multi-element, multi-layer receive arrays with two-, three-, and four-channel configurations, optimising

the spatial arrangement to balance field penetration and uniformity. The four-element multi-layer array, validated via simulation, bench-top testing, and in-bore measurements, outperformed a commercial chest coil in both B_1^- homogeneity and receive signal sensitivity at depth. This receive sensitivity advantage is the direct consequence of the multi-layer design and the surface coil proximity afforded by the flexible array architecture, and translates to coverage at cardiac depths that the rigid commercial design cannot match. Together, these results establish the multi-layer array as a high-sensitivity, geometrically robust receive coil, providing the hardware foundation upon which the accelerated acquisition and spectral fitting strategies developed in the subsequent chapters of this thesis depend.

Clinical and Technical Considerations

This work presents a clinically oriented approach to hyperpolarised ^{13}C receive array design. The geometric decoupling approach streamlines assembly and tuning, removing the need for complex active decoupling circuitry. This flexible, lightweight multi-layer coil bridges the gap between bench top prototyping and scanner-integrated implementation with Tx and Rx modularity. These features make the design adaptable to a range of scanner environments, supporting future clinical integration.

With unlimited resources, I would have expanded the design to include both transmit and receive optimisation from the ground up, rather than relying on inherited transmit hardware. A custom-designed, large dual-loop Helmholtz transmitter integrated with the multi-layer Rx array could provide better flip angle control and field homogeneity. I would also have explored higher element counts with posterior coverage, implemented in flexible or segmented housings to maximise anatomical conformity.

Although the array was evaluated in its intended conformation during phantom scans, I did not explicitly simulate it in a flexed or deformed geometry. However, the observed performance suggests that body conformity strongly impacts the receive field. From a modelling standpoint, one limitation I encountered in CST was the need for manual retuning after each parameter sweep, which is an additional time sink on top of the 10-day-long processing time required to simulate on a single human model. Changing any geometric layout of the element required a re-run of the simulation. This made design-space exploration time-consuming and restricted my ability to test alternative layouts or loading conditions. Automating tuning processes or faster EM calculations could drastically accelerate optimisation in future studies, allowing more rapid convergence on clinically viable configurations. Additionally, CST does not support modelling of flexible coil bending, which would have enabled large-scale exploration with varying conformation to various anatomical shapes that might be encountered in a clinical setting.

Future simulations should explore the effect of subject-specific variability, such as torso size, body mass index (BMI), and chest geometry/anatomy, on coil performance. I would use a more diverse set of human models (e.g., male/female, obese/underweight) to fully explore my coil performance. Furthermore, understanding how element positioning and inter-element coupling vary under anatomical curvature and how bending the coil changes performance will be essential for clinical robustness across diverse patient populations.

Safety is another critical consideration. Although this project focused on receive array development, any clinical deployment will require rigorous safety evaluations, particularly with respect to Specific Absorption Rate (SAR) and thermal effects. While SAR remains the standard metric for RF safety, localised electric field (E-field) hotspots often provide a better correlate to thermal risk. To address this, future safety studies should combine electromagnetic simulations with IR thermography and in vitro heating tests using tissue-equivalent phantoms and thermal probes to ensure safe operation at clinical power levels.

Another barrier to clinical translation is the current workflow, which remains dependent on offline processing. The B_1^+ maps and associated correction factor maps used in this work were generated via offline post-processing. While this is appropriate in a research context, it is several steps removed from clinical feasibility, where image reconstruction and flip angle correction would need to occur in-line with scanner workflows, ideally producing immediately usable, corrected images. The need for manual export, custom MATLAB processing, and re-import introduces workflow complexity that would not be viable in a clinical setting. The GE research environment provided powerful tools for prototyping and validation during my coil development, but transitioning these tools to a clinical workflow will require further integration of post-processing routines and streamlined pipelines for x-nuclei calibration and reconstruction.

In terms of of the array design, the balance between the number of elements and inter-element isolation is a recurring theme. More elements improve coverage and enable parallel imaging, but they also increase the complexity of decoupling

and require more hardware components. In this work, I focused on demonstrating a proof-of-concept for a four-element array, where geometric decoupling sufficed. However, scaling the design will require more advanced strategies, particularly if posterior coverage or transmit integration is added. Additionally, there remains significant room to improve receive coil combination strategies. More sophisticated coil combination methods than the RSS used in this thesis and SNR optimisation algorithms, such as those based on noise covariance, could yield substantial gains. Exploring these techniques could form the basis of an entire additional project.

Finally, future development should explore integration with dual-tuned $^{1}H/^{13}C$ coils to enable anatomical localisation during HP ^{13}C scans, increasing the clinical interpretability of metabolite images. Furthermore, form-fitting coil housings or flexible casings will improve patient comfort and scan reproducibility. Higher channel counts may also enable parallel imaging acceleration, which is a growing need in real-time metabolic imaging. The integration of multi-channel Rx arrays to enable parallel imaging techniques is further discussed and investigated in chapter 3: Accelerating Translational HP MRI. These techniques are indispensable for accelerating data acquisition, a necessity in HP studies where imaging must be completed within a few minutes. The Rx array capability to support spatially distinct coil elements with independent receive channels forms the theoretical basis for these acceleration strategies [27, 81, 82].

Reflections on Industrial Collaboration and Research Pivot

My doctoral research was conducted within an industrial framework, a setting that demanded a dual role as both an academic researcher and a project manager. As my background prior to graduate study was in professional software engineering rather than hardware design, my first and second year necessitated a rapid, self-directed immersion into RF engineering, laboratory protocols, and hardware troubleshooting to maintain project momentum. I actively sought mentorship and aligned my technical outputs with the company's evolving priorities, which provided unique opportunities to contribute to the quality assurance and regulatory testing of industrially deployed proton and X-nucleus arrays (including ^{2}H, ^{129}Xe, and ^{31}P). I also proactively participated in business-facing activities, including product demonstrations and sales support at international conference stands. These experiences broadened my understanding of the translation from bench top prototypes to clinical-grade hardware.

While the DPhil was originally scoped as a purely hardware-oriented study, the combination of industrial engineering constraints and limited scanner access during the first year prompted a strategic pivot. During the course of this DPhil, PulseTeq Ltd. navigated significant institutional changes, including a company merger and periods of reduced engineering capacity. I chose to leverage my established expertise in software engineering to expand the research scope into post-processing, accelerated imaging, and spectral fitting. This transition allowed for a more holistic investigation of the ^{13}C imaging chain, moving from the physical interface of the coil to the computational methods required to extract metabolic data. The following chapters detail these software-driven advancements, which complement the hardware foundations established in this initial phase. In particular, the insights gained here set the stage for Chapter 3, where I investigate how parallel imaging methods, such as SENSE-accelerated chemical shift imaging, can further translate HP MRI toward clinical feasibility.

2.7 Appendix

B1 Map PulseTeq Coil Protocol:

For the PulseTeq Two-Loop coil, the spiral-based transmit B_1 map scanner protocol is as follows: Coronal FOV = 40cm, Slice Thickness = 20mm (AP), Slices = 4 (starting at 6cm from the top of the phantom), Freq Dim = S/I, TR = 8455.1, FA = 90, NEX = 150, CV3 = 3, CV4 = 32, CV14 = 4, CV19 = 16, CV20=0, CV21 = 0. A single Tx line and a single Rx line is used for this coil.

Further Rotational Geometric decoupling:

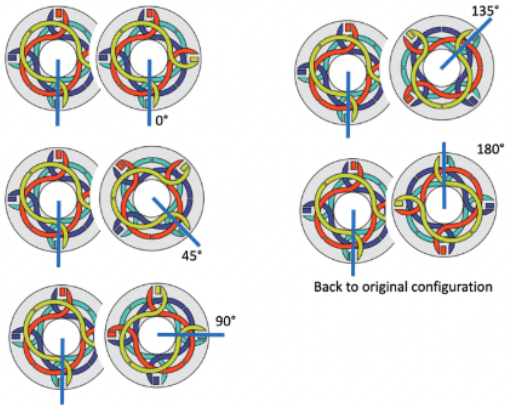


Figure 2.73: At a fixed distance (^{13}C m), the rotational difference was tested for 0°, 45°, 90°,135°, and 180° as illustrated

Bench top measurement of the two-coil setup was optimised at 130 mm distance between the centres, where the S_{21} values at 32.11 MHz went as low as –29.4 dB. Rotational decoupling was optimised at a rotation angle of 90° degrees between the coils. The testing was very sensitive to the configuration and the cable connection to the network analyser, making this a good starting point for the CST simulations.

Rotation (°)	S21 value (db) at 32.11MHz
0	-29.9
45	-14.0
90	-32.8
135	-16.1
180	-19.8

Figure 2.74: bench top S21 values of the orientations shown in figure 2.73 is quantified in this table.

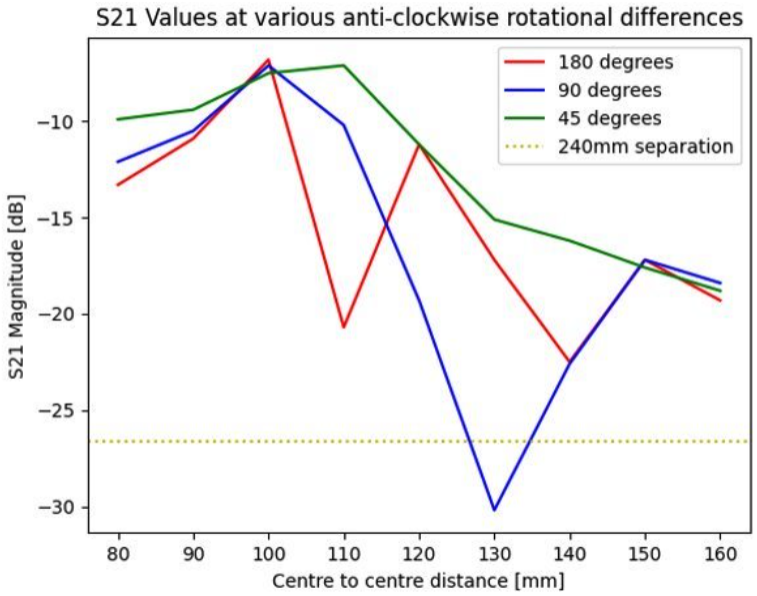


Figure 2.75: bench top S21 values of different rotational differences at a range of separation distances.

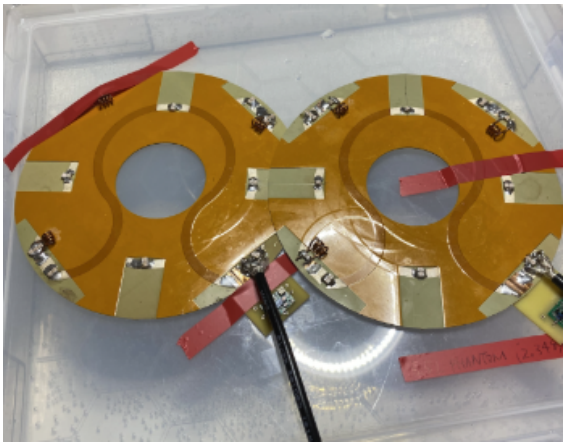


Figure 2.76: bench top experiments with both coils rotated 45 degrees to enable overlap without clashing with the proton traps.

The overlap decoupling showed expected values, but the rotational decoupling showed an unexpected difference in S_{21} values between the 0° and 180° rotations of 10.1 dB. Because the coil is symmetrical, the configuration should return after a 180° turn, so the S_{21} values was expected to be similar.

S_{21} value for both coils rotated 45 degrees: Inspired by thinking about configurations to minimise overlap space due to the proton-traps. It is surprising to see that a net 0 degrees rotation between the two 45degrees coil have low decoupling compared

Centre to centre distance (mm)	S21 value (db) at 32.11MHz
160	-16.7
150	-15.1
140	-10.4
130	-8.1
120	-7.1
110	-13.5
100	-8.9
90	-9
80	-8.4

*130-160mm allows for overlap without hitting a proton trap.

Figure 2.77: bench top S21 values from the 45 degrees rotation set up shown in figure 2.76

Distance in mm	S21 @ peak	s21 @ 32.11MHz
160	32.16, -15.8	-16
150	32.16, -14.0	-14.1
140	32.16, -14.8	-15.0
130	32.24, -21.3	-22.4
120	32.26, -31.1	-32.9
110	32.21, -8.6	-9.1
100	doublet	-6.3
90	doublet	-10.6
80	doublet	-12.4

Figure 2.78: S21 values from the offset 0 and 90 degrees set up shown in figure 2.79

to the previous rotational testing. This configuration seems to be decoupled the most when there is no overlap. Here, the coils seemed to have weak but stable decoupling between 150-160mm separation between the centres.

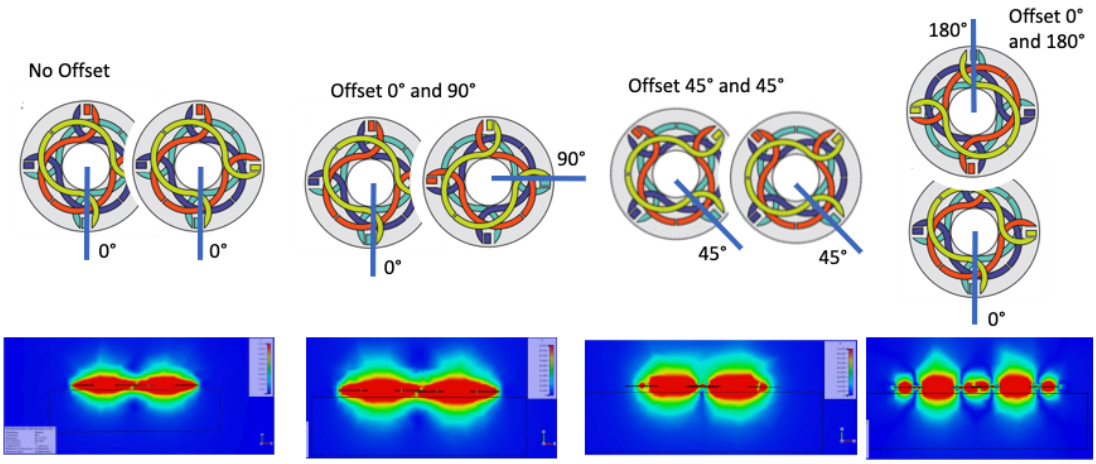


Figure 2.79: Various orientations simulated on CST.

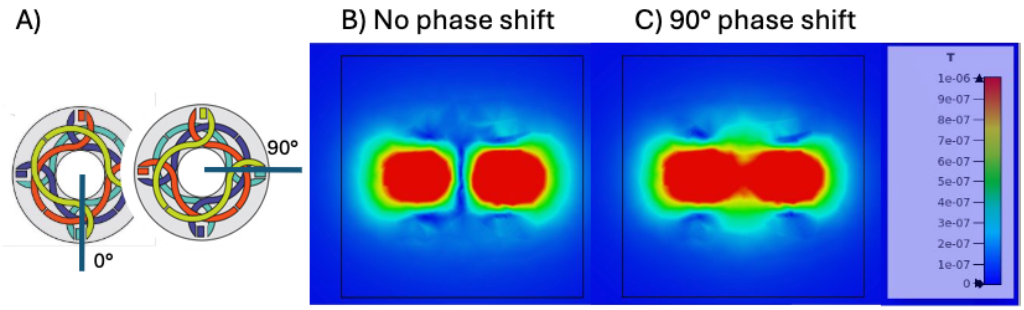


Figure 2.80: A) Set up of two elements with 14cm separation with an offset 0° and 90°. B and C show the B_1 field taken in the y-direction at 5cm depth(1e-6T scale) on a fat phantom, where B has no phase shifts and C has a 90° phase shift between the two elements. The phase shift in C looked promising in increasing the sensitivity between the elements. However, as shown in Figure 2.81, this sensitivity increase did not translate to the heart ROI at the center of the four element configuration.

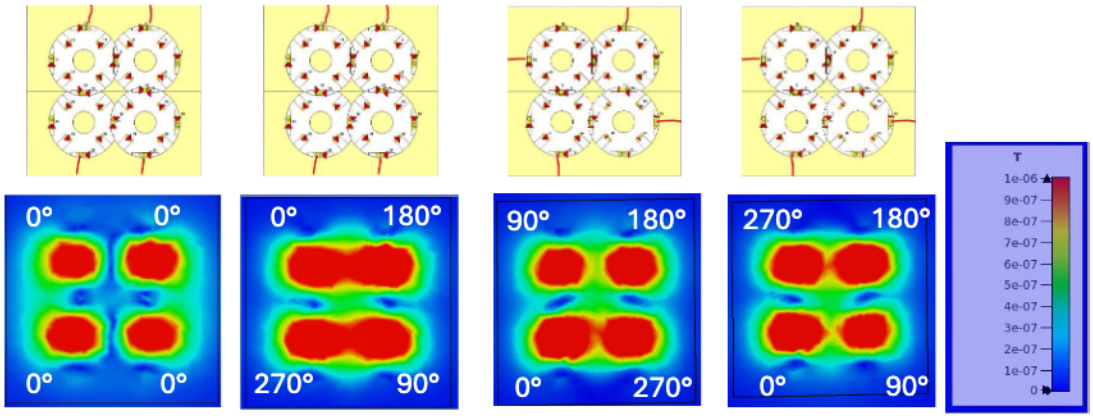


Figure 2.81: Representative cases of the 4 element configurations with various phase shifts in attempts to optimise for B_1 field increase at the center of the array. The increase in B_1 field increase between diagonal elements is essential within this design to cover the heart ROI, and requires further investigation.

Multi-layer PCB Capacitors:

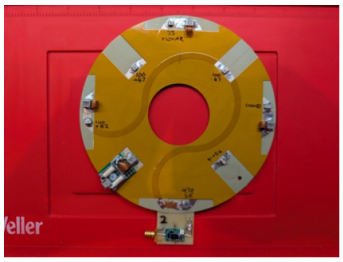


Figure 2.82: Coil 2 of the multi-layer array set up.

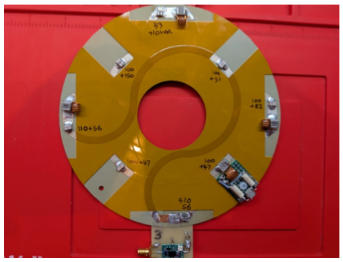


Figure 2.83: Coil 3 of the multi-layer array set up.

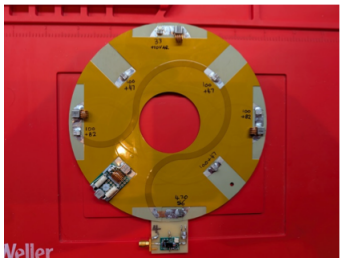


Figure 2.84: Coil 4 of the multi-layer array set up.



Figure 2.85: Coil 5 of the multi-layer array set up, fabricated as a back up element in case of hardware failure.

Rx coil elements extended S-parameters:

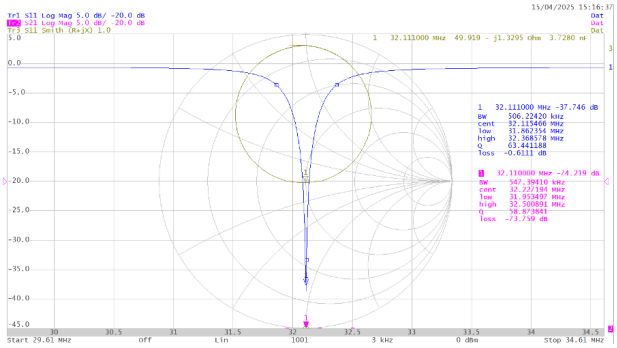


Figure 2.86: Rx coil 1 S-parameters over a chest on a representative adult male.

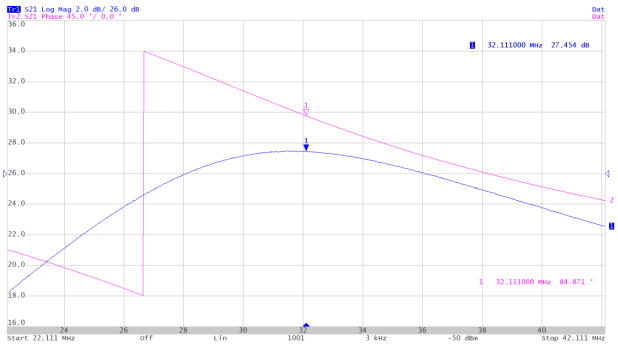


Figure 2.87: Rx coil 2 preamplifier S-parameters.

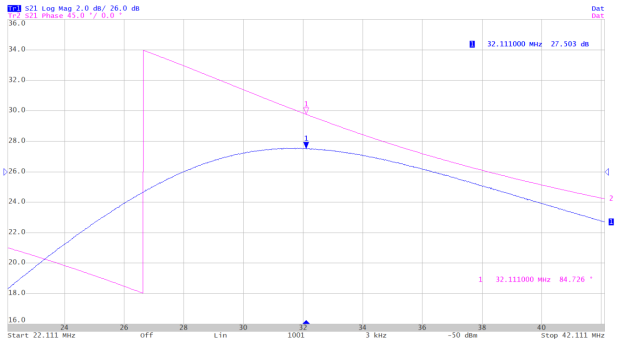


Figure 2.88: Rx coil 3 preamplifier S-parameters.

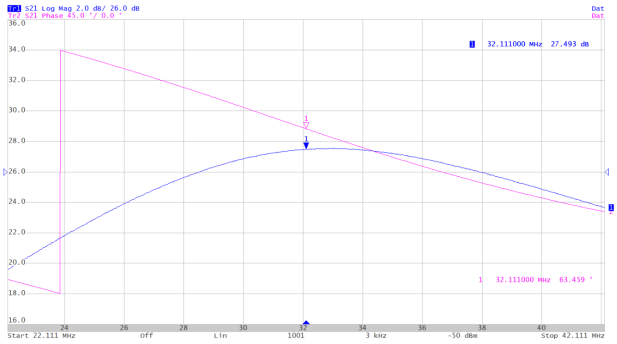


Figure 2.89: Rx coil 4 preamplifier S-parameters.

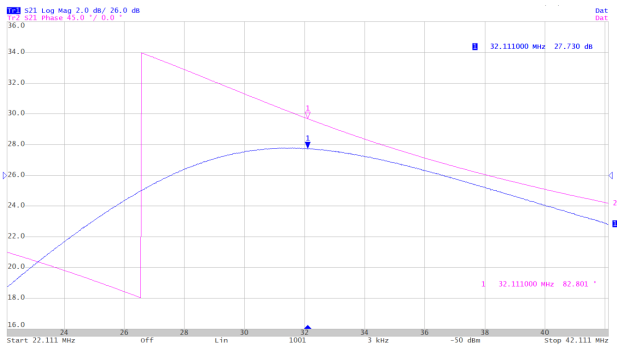


Figure 2.90: Rx coil 5 preamplifier S-parameters.

3

Accelerating Translational HP MRI

Contents

3.1 Abstract 113
3.2 Introduction 114
3.2.1 Motivation 114
3.2.2 Approaches in Fast HP MR Imaging 115
3.2.3 Parallel Imaging 116
3.2.4 Development of 2D SENSE CSI 117
3.3 Methods 122
3.3.1 SENSE Reconstruction 122
3.3.2 Phantom Study 123
3.3.3 In-Vivo Porcine Study 124
3.4 Results 127
3.4.1 Phantom Study 127
3.4.2 In-Vivo Porcine Study 128
3.5 Discussion 130
3.6 Conclusion and Future Directions 133
3.7 Appendix 135

3.1 Abstract

Hyperpolarised ^{13}C MRI is constrained by rapid signal decay, necessitating fast imaging strategies that preserve spectral and spatial integrity. Recognising the need for a method that captures this transient signal before it irreversibly decays, I

developed a robust approach that ensures clinical usability whilst accelerating a chemical shift imaging (CSI) sequence. This work investigates the application of Sensitivity Encoding (SENSE) to achieve a fourfold ($R = 4$) reduction in acquisition time. To circumvent the low natural abundance of ^{13}C, which precludes standard coil sensitivity mapping, I used a dual-tuned (^{23}Na/^{13}C) flexible receiver array to provide high-SNR sodium surrogate coil-sensitivity maps for the SENSE reconstruction.

I validated this approach through retrospective undersampling of phantom data and a porcine study (n=3), reducing the acquisition time from 36 s to 9 s per time point. The experiments demonstrated that two-dimensional SENSE acceleration can be successfully implemented in CSI to reconstruct anatomical metabolite maps and correct for aliasing artefacts, with g-factor analysis confirming manageable noise amplification in the regions of interest. By combining a vendor-supplied CSI sequence with hardware-enabled acceleration, this method enhances the clinical feasibility of HP ^{13}C imaging and enables wider application in research and diagnostics without the need for complex, site-specific sequences.

3.2 Introduction

3.2.1 Motivation

Hyperpolarised ^{13}C MRI is a clinically translational, non-invasive, technique used to detect and localize low-signal metabolites *in vivo* through injections of ^{13}C-labelled endogenous substrates [56]. This technique offers a unique opportunity to study metabolic processes *in vivo*, promising to provide key information to further our understanding of early metabolic changes in pathology. The fundamental constraints of HP carbon-13 MRI, including the non-renewable magnetisation, T1 decay, and the resulting narrow acquisition window, are established in Section 1.2.4. Imaging time is therefore scarce, and consequently motivate the fast imaging approaches reviewed in the following section.

3.2.2 Approaches in Fast HP MR Imaging

Achieving higher temporal resolution in hyperpolarised ^{13}C MRI involves balancing the trade-offs between accelerating imaging, maintaining a high signal-to-noise ratio (SNR), and achieving sufficient spatial and spectral resolution. These technical challenges, together with the complexity of existing acquisition methods and specialised hardware requirements, have limited the widespread clinical adoption of HP ^{13}C MRI, confining its use largely to specialist centres. To address the needs of carbon-13 HP MRI, various temporal acceleration strategies have emerged, each with their own advantages and limitations. In current literature, the high-level approaches in temporal acceleration for hyperpolarised MR include: acquiring multiple k-space points per TR, under sampling during k-space traversal, decreasing TR, and harnessing spectral or spatial prior knowledge during data reconstruction. To enhance the benefits of these various methods, approaches may be combined, which is covered more extensively in the following references [41, 83].

One method is spectral spatial imaging (SPSP), where a metabolite-specific RF pulse excites a specific frequency and a fast imaging readout encodes the spatial information[84]. Techniques such as single-shot echo-planar imaging [85] or spiral readouts [27] are used with this approach. SPSP is RF-efficient but requires prior spectral knowledge of the metabolites expected to be observed in the experiment. Incorrect setting of the centre frequency can lead to failure in exciting the appropriate resonance, resulting in a wasted experiment [86].

Another strategy is fast spectroscopic imaging, which uses oscillating echo-planar readout gradients to acquire spatial and spectral information simultaneously during a single readout. By sampling multiple lines of k-space within each excitation, these methods substantially increase acquisition efficiency and reduce the total scan time. Methods such as Echo-Planar Spectroscopic Imaging (EPSI), spiral, and radial acquisitions [86–89] fall under this category. These techniques can achieve high temporal resolution, although this requires a trade off between temporal resolution, spectral bandwidth, and spatial resolution within the limits of the available gradient hardware [86].

When the spectral profile is well defined, multi-echo balanced steady state free precession (bSSFP) can be used. In this method, data is acquired at echo times centred between the RF pulses, i.e., at $2 \times \text{TE} = \text{TR}$ [90]. Multi-echo bSSFP has been performed on rat kidneys achieving good spatial resolution[91].

Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL) is a fast chemical shift encoding method. It combines the asymmetrical acquisition of echoes with a decomposition algorithm[92]. While IDEAL provides high temporal resolution, it can suffer from bleed-through artefacts [92].

Phase-encoded CSI was selected as the acquisition framework for this work due to its broad spectral bandwidth and insensitivity to field inhomogeneities, as described in Section 1.2.2. The following section outlines how CSI can be accelerated using parallel imaging.

3.2.3 Parallel Imaging

Parallel imaging has become a standard approach for accelerating imaging. As established in Section 1.2.2, reducing k-space sampling density below the Nyquist limit constrains the FOV and introduces spatial aliasing. Parallel imaging exploits the distinct spatial sensitivity profiles of receive array elements to correct for this aliasing, enabling a reduction in the number of phase-encoding steps and therefore a shorter acquisition time [93, 94]. When applied to CSI, this approach preserves the spectral resolution inherent to the CSI method and is particularly well-suited to array coils that provide distinct sensitivity profiles across the field-of-view.

The two established frameworks for aliasing correction are Sensitivity Encoding (SENSE) [69] and Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [70]. GRAPPA operates in k-space by estimating missing data points from neighbouring acquired samples using interpolation weights derived from an autocalibration signal, before image reconstruction. SENSE performs the reconstruction in the image domain by explicitly incorporating coil sensitivity maps to mathematically unfold the aliased signals. Although more advanced strategies such as non-Cartesian sampling or compressed sensing offer higher

acceleration factors, they require specialised sequences or computationally intensive iterative reconstructions that are difficult to translate into a standard clinical workflow [19, 83, 95–97]. In contrast, the Cartesian SENSE approach allows for acceleration using standard, manufacturer-supplied CSI sequences while preserving image resolution through a straightforward geometric reconstruction, providing a robust path to increasing temporal resolution in hyperpolarised ^{13}C imaging without sacrificing spectral integrity.

A key requirement for accurate SENSE reconstruction is the availability of coil sensitivity maps, typically acquired in a separate pre-scan. In the case of ^{13}C imaging, however, the low natural abundance of ^{13}C makes such pre-scans impractical, and the high cost of the hyperpolarised injection, currently exceeding £4000 per dose, precludes the use of a dedicated sensitivity injection for calibration [95]. Multiple approaches to estimating a coil sensitivity profile have been proposed: auto-calibration [58, 94], obtained from the centre of a fully sampled k-space; calibrationless reconstruction [98], which iteratively solves a matrix completion problem; and hardware approaches [27], which pre-acquire a sensitivity map using a surrogate nucleus.

To address this limitation in the present work, I implemented a cross-nuclear approach using naturally abundant ^{23}Na. Given its similar Larmor frequency and sufficient thermal signal, sodium sensitivity maps can be acquired without additional contrast agents or sequence modifications, and have been shown to provide accurate ^{13}C sensitivity profile estimates at 3T [27]. This enabled the application of SENSE-accelerated CSI achieving up to fourfold acceleration while maintaining both spatial and spectral resolution. I validated this method in phantom and porcine kidney studies, demonstrating effective aliasing correction and establishing its feasibility for dynamic hyperpolarised metabolic imaging.

3.2.4 Development of 2D SENSE CSI

The aim of this work was to accelerate CSI by limiting FOV and using the sodium-carbon coil mentioned above to increase temporal resolution whilst retaining spatial

and spectral resolution. Building on the hardware developments, the aforementioned coil was expected to correct for coil weighted effects, which is a characteristic of the flexible coil, and for aliasing artefacts, which is a characteristic of accelerated parallel imaging. In accelerated ^{13}C CSI, these aliased images contaminate the metabolic maps. For clarity, there is a distinction between the spatial dimensions of the acquisition and the dimensionality of the Cartesian k-space acceleration. While a 2D CSI sequence was used to encode the spatial domain (x, y), the 1D/2D SENSE acceleration refers specifically to the undersampling pattern in either one or both k-space dimensions (k_x and k_y), independent of the fact that the resulting data is reconstructed into a two-dimensional image space.

As of February 2026, the most recent Fidall package has phased out their SENSE reconstruction algorithm, and requires a custom code base to reconstruct accelerated data. At the time of this project in 2024, GE Healthcare supported proprietary SENSE reconstruction in the Fidall package, but it was still not compatible for my specific requirements of retrospective undersampling on the phantoms. The Fidall SENSE reconstruction pipeline automatically extracts acquisition parameters from the GE raw data header, providing a robust framework for reconstructing scanner-acquired datasets. However, because the reconstruction relies on this metadata, retrospectively undersampling the k-space also required the corresponding acquisition metadata to be updated. Re-engineering this metadata within the GE header structure proved non-trivial. Consequently, I began with the fundamental principles of SENSE reconstruction to a 2D Cartesian implementation. This development was informed by the parallel imaging tutorial by Chiew [99], which provided the mathematical basis for unaliasing overlapping signals in the image domain.

The undersampled imaging problem can be modelled as $\mathbf{k} = \Phi F S \mathbf{m}$, where $\mathbf{k}$ is the k-space data, $\mathbf{m}$ is the target magnetization vector, S is the coil sensitivity matrix, F is the Fourier Transform, and Φ is the k-space sampling matrix [99]. If k-space is fully sampled, Φ is the identity matrix. When solving for $\mathbf{m}$ in the least squares sense, representing this inverse explicitly is very costly computationally and memory-wise.

Conveniently, the use of a Cartesian undersampling pattern allows for a simplification. When the sampling matrix Φ systematically skips every R^{th} line of k-space, the resulting aliasing follows a regular pattern, where each voxel in the reduced field-of-view (FOV) represents a superposition of exactly R source voxels from the full FOV. This allows the reconstruction to be decomposed into a set of independent, localized sub-problems. For a given coil l, the acquired image $C_l(x)$ is the product of the complex coil sensitivity $S_l(x)$ and the magnetisation $M(x)$, which is the final, unaliased image:

$$C_l(x) = S_l(x) \cdot M(x) \tag{3.1}$$

When it is not possible to acquire the ^{13}C sensitivity maps S_l directly, surrogate coil maps are required, such as the sodium sensitivity maps.

In the 1D k-space acceleration case, this is expressed as a matrix equation where the vector of coil signals at (x, y) is the product of a sensitivity matrix and the source signals as shown in figure 3.1. For an acceleration factor of $R = 2$, this system consists of N_c equations and two unknowns: the magnetization at the primary location and the magnetization at the location shifted by half the FOV. In other words, within the reduced Field of View (FOV), each aliased image is a superposition of the true image plus copies shifted by a specific distance. The unaliasing process then becomes an exercise in finding the best estimate $\hat{m}$ in a least-squares sense. By leveraging the spatial diversity of the receiver array, where each coil provides a unique spatial signal, the algorithm can mathematically separate the overlapping signal. This inversion outputs a single, coil-combined image that eliminates the coil-specific weighting while restoring the correct anatomical localization of the signal.

In the 2D k-space acceleration implementation, I extended this framework to account for simultaneous shifts in both the x and y directions. For a total acceleration of $R = 4$ ($R_x = 2, R_y = 2$), the sensitivity matrix expands to include four columns to account for the four spatial superpositions: (x, y), $(x + N_x/2, y)$, $(x, y + N_y/2)$, and $(x + N_x/2, y + N_y/2)$. The unknown magnetization $\mathbf{m}$ is then

$$\begin{bmatrix} c^1_{x,y} \\ c^2_{x,y} \\ \vdots \\ c^{N_c}_{x,y} \end{bmatrix} = \begin{bmatrix} S^1_{x,y} & S^1_{x+N_x/2,y} \\ S^2_{x,y} & S^2_{x+N_x/2,y} \\ \vdots & \vdots \\ S^{N_c}_{x,y} & S^{N_c}_{x+N_x/2,y} \end{bmatrix} \begin{bmatrix} m_{x,y} \\ m_{x+N_x/2,y} \end{bmatrix}$$

Figure 3.1: Equation from reference [99] of an R=2, 1D acceleration in the N_x direction

recovered in a least-squares sense using the localized matrix inversion:

$$\hat{m} = (S^H \Psi^{-1} S)^{-1} S^H \Psi^{-1} \mathbf{a} \tag{3.2}$$

By solving these independent sub-problems voxel-by-voxel, the algorithm leverages the spatial diversity of the receive array to separate the overlapping signals and restore the original image.

To evaluate the noise penalties inherent in this acceleration, I performed a geometry factor (g-factor) analysis. The g-factor is a quantitative measure of the noise amplification produced by the coil hardware and the mathematical unaliasing process [69]. It represents the degree to which the coil sensitivity profiles are distinct, where a high g-factor indicates that the overlapping signals are difficult to separate, leading to localized noise. To empirically derive these maps, I employed a Monte Carlo noise simulation approach, following the pseudo-replica method defined by Robson et al. [100]. The practical implementation was based on a MATLAB tutorial made freely available by Hansen [101], which I adapted for the 2D SENSE reconstruction requirements of this study. The total signal-to-noise ratio (SNR) in a SENSE-accelerated image is governed by the following relationship:

$$\mathrm{SNR}_{SENSE} = \frac{\mathrm{SNR}_{full}}{g\sqrt{R}} \tag{3.3}$$

where R is the acceleration factor. While the $\sqrt{R}$ term accounts for the loss of SNR due to the reduced total sampling time, the g-factor accounts for the additional noise amplification caused by the non-orthogonal nature of the coil sensitivities. I prioritized the g-factor analysis over standard SNR metrics for this validation.

While SNR is a common measure of image quality, it is not a robust metric for evaluating the performance of a parallel imaging reconstruction across different HP injections. The SNR of a fully sampled reference image (SNR_{full}) is heavily influenced by experimental factors independent of the reconstruction method, such as fluctuations in hyperpolarisation, variations in bolus timing, and physiological motion. Because these factors change between experiments and even between individual injections, SNR does not isolate the performance of the SENSE algorithm itself. In contrast, the *g*-factor provides a direct, hardware-centric quantification of how much noise is amplified by the parallel imaging geometry alone. By mapping the *g*-factor across the phantom and *in vivo* data, I was able to evaluate the viability of a fourfold acceleration independent of the experimental factors that change the signal and noise levels.

The following sections in this chapter cover the phantom and animal studies I carried out to investigate the SENSE CSI implementation for carbon-13 MRI. To validate the reconstruction, I first employed retrospective undersampling. By taking a fully sampled k-space acquisition on the scanner from a head phantom and artificially removing phase and frequency encoding steps, I could directly compare the SENSE-reconstructed images against a known ground truth. These results are detailed in the Phantom Study section 3.3.2 of this chapter. Following successful bench-testing on the phantom model, I transitioned to an *in vivo* experiment through a collaboration with Aarhus University, Denmark. This study, as detailed in the *in vivo* Porcine Study section 3.3.3, involved the acquisition of porcine kidney images following the injections of hyperpolarised ^{13}C pyruvate, one with the fully sampled acquisition and another with the undersampled protocol. A critical hardware innovation in this study was the use of a flexible dual-tuned (^{23}Na/^{13}C) coil as described in the previous subsection. By acquiring a high-signal sodium map, I was able to generate a surrogate sensitivity profile for the ^{13}C reconstruction.

My research represents one of the first comprehensive investigations of SENSE acceleration for HP ^{13}C 2D CSI. The findings contribute to the development of faster, more efficient HP ^{13}C MRI techniques, offering a promising solution for metabolic

imaging across diverse clinical scenarios. By addressing these key challenges and developing a more accessible technique, this work aims to accelerate the translation of HP ^{13}C MRI into clinical practice, thereby expanding its utility in metabolic imaging.

3.3 Methods

All images were acquired using a GE 3T scanner (GE Healthcare, Waukesha, WI, USA). I performed ^{13}C imaging with a ^{13}C-tuned clamshell volume coil (Rapid Biomedical) for transmission and reception via a flexible eight-channel ^{23}Na/^{13}C receive array (fabricated by Sanchez-Heredia in reference [27]), shown in Figure 3.2B. Free induction decay (FID) CSI, whereby the pulse sequence starts with a wide-bandwidth RF pulse, followed by phase encoding, and FID sampling, was implemented in the phantom and porcine study[9]. All computation was carried out in MATLAB (R2023b, Mathworks), and code is accessible through the laboratory GitHub page.

3.3.1 SENSE Reconstruction

I developed a 2D Cartesian SENSE reconstruction in MATLAB, designed to handle the requirements of my hyperpolarised ^{13}C dataset. The algorithm operates by iterating through the undersampled image coordinates and calculating the corresponding indices for the source voxels that have aliased into the current location. I first implemented a preprocessing step where a binary mask was generated from the combined ^{23}Na signal. This mask ensures that the reconstruction inverted the sensitivity matrix in regions where the coil profiles were well-defined.

For this study, I implemented a simple pseudo-inverse (pinv) to solve $\mathbf{c} = S\mathbf{m}$ with a Generalized Least Squares (GLS) framework. By incorporating a noise covariance matrix, Ψ, the reconstruction operator can be expanded to $\mathbf{u} = (S^H\Psi^{-1}S)^{-1}S^H\Psi^{-1}$. While Ψ was treated as an identity matrix for unaliasing in this work, which is equivalent to the standard pseudo-inverse, the code has the capability for coil-weighting in future works. By applying this inverse to the vector of aliased coil signals, the algorithm reassigns intensity to its correct spatial coordinate, and

unfolds the image into a single, coil-combined result. This implementation treated the x and y acceleration directions as separable, allowing the 2D problem to be solved efficiently through sequential 1D-by-1D operations.

3.3.2 Phantom Study

To calibrate for the sodium and carbon centre frequency and transmit gain, I used a Bloch-Siegert method[71] on a Specific Anthropomorphic Mannequin (SAM) phantom filled with ethylene glycol doped with 17g/L of NaCl (291mmoL/L ^{23}Na concentration)[27].

To assess the coil profile of the ^{23}Na/^{13}C receive coil (Figure 3.2B) on a head-shaped phantom, I pre-acquired signal at the sodium frequency using a single-slice CSI sequence (FOV = 240 mm × 240 mm, NEX = 150, FA = 90°, TR = 40 ms, receiver bandwidth = 5 kHz, acquisition matrix size = 24 × 24, slice thickness = 15 mm), ensuring the FOV was larger than the phantom. As illustrated in Figure 3.2C, I wrapped the array coil around the head phantom, though it was too short to cover the entire circumference.

I collected a second single-slice CSI at the carbon frequency using the exact same setup as the sodium acquisition, maintaining the coil and phantom geometry. This CSI (FOV = 240 mm × 240 mm, NEX = 8, FA = 50°, TR = 29.3 ms, bandwidth = 50 kHz, matrix size = 24 × 24, slice thickness = 40 mm) utilized an identical FOV to the sodium acquisition.

During post-processing, I generated coil-wise sodium sensitivity maps by normalising the sodium images to the peak signal across all coil elements, smoothing the signal, and masking out data beyond both the receive coil and the phantom. To produce the carbon images, I extracted the peak intensity from each voxel spectrum from one of the ethylene glycol resonances. I then created an undersampled image with an acceleration factor of R = 4 by retrospectively reducing the fully sampled carbon CSI k-space data, filling every other row and column in the matrix with zeros. I reconstructed this undersampled image using a 2D SENSE algorithm in

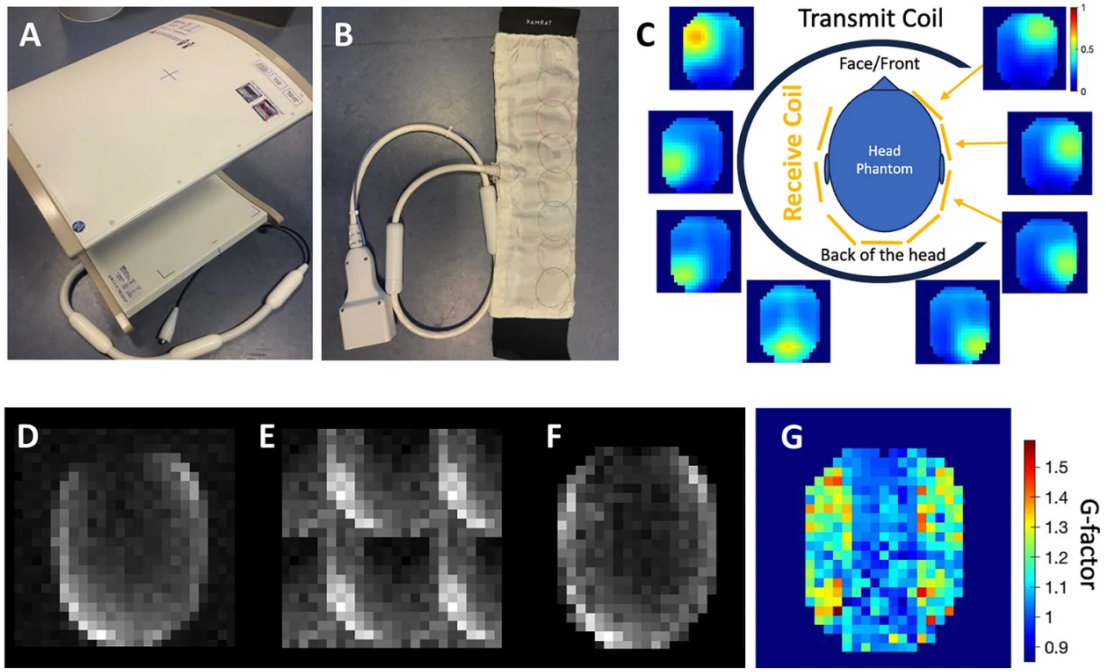


Figure 3.2: (A) ^{13}C clamshell transmit coil (Rapid Biomedical, Germany) (B) an eight-channel array, flexible 23Na/^{13}C receive coil (fabricated by Juan D. sanchezheredia2022) (C) a phantom experiment set up is illustrated, where the head phantom sits inside the clamshell transmit coil and the head is wrapped with the receive coil. The location of the receive elements are depicted as yellow lines. Coil-wise sodium sensitivity images surround the illustration. (D) The fully sampled ^{13}C phantom image. (E) The retrospectively undersampled ^{13}C phantom image. (F) The undersampled SENSE reconstructed image. (G) The g-factor map of R = 4, where acquisition is accelerated by R = 2 in both the frequency and phase directions.

MATLAB (R2023b, Mathworks), applying the sodium sensitivity maps to correct for the acceleration artefacts.

Finally, to generate the R = 4 g-factor map, I reconstructed 100 noise images at R = 1 and another 100 accelerated noise images at R = 4 using the 2D SENSE algorithm. I mapped the ratio of the standard deviations of these R = 1 and R = 4 reconstructions per voxel to produce the final g-factor map. Images were normalised to the highest signal in each image.

3.3.3 In-Vivo Porcine Study

Figure 3.3 provides an overview of the experimental workflow implemented in this study. First, [1-^{13}C]pyruvate was hyperpolarised using dynamic nuclear polarisation (DNP), increasing the available ^{13}C signal by more than four orders of

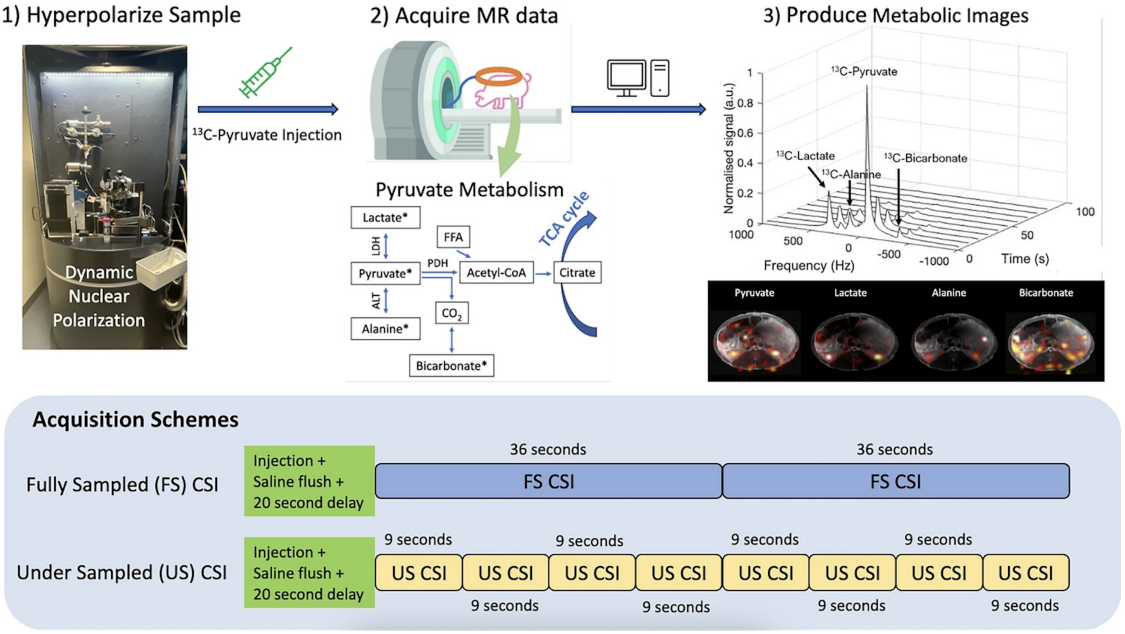


Figure 3.3: Overview of the HP ^{13}C experimental work flow (steps 1–3) carried out in this study. The illustrations of the two CSI acquisition schemes are both FID CSIs taken consecutively following an injection, saline flush, and a 20 second delay. The fully sampled (FS) CSI (36 seconds per acquisition) and the undersampled (US) CSI (9 s per acquisition) are collected after separate HP ^{13}C injections about 90min apart. To visualize the four-fold acceleration, the schemes are illustrated directly on top of each other, where four consecutive US CSIs are taken within a single FS CSI.

magnitude. The hyperpolarised substrate was then injected intravenously and its downstream metabolism was monitored using either a fully sampled or accelerated CSI acquisition. Finally, the acquired data were reconstructed using the sodium-guided SENSE framework described in the previous section to generate metabolic images for quantitative analysis. The following sections describe each stage of this workflow in detail.

Animal handling and clinical procedures were performed by a licensed collaborator at Aarhus University, Denmark in accordance with the license and ethical review granted by the Danish animal inspectorate regulations (2019-15-0201-00298), and were reported in accordance with the ARRIVE guidelines. Three female Danish domestic pigs were sedated with a zoletil mix (0.1 mL/kg IM; 25 mg/mL tiletamine + 25 mg/mL zolazepam) and intravenous (i.v.) access was obtained through an ear venflon. Following a 3 mL propofol bolus administered through the i.v., the animals were intubated. I was responsible for the coordination of the experimental timing and

scanning, while the collaborator maintained anaesthesia throughout the study using a combination of propofol (10 mL/h), midazolam (10 mL/h), and fentanyl (4 mL/h).

I positioned the porcine subject within the MRI bore as shown in the experimental setup in Figure 3.4A. To visualise the torso anatomy, I acquired anatomical reference images (Figure 3.4B) using the body coil with an axial T_2-weighted fast spin echo sequence (FOV = 360 mm × 360 mm, TR = 4300 ms, echo time = 121 ms, matrix = 512 × 512, slice thickness = 10 mm), allowing me to identify the kidney regions of interest. For sensitivity mapping, I acquired a single-slice sodium CSI (FOV = 240 mm × 240 mm, NEX = 16, TR = 38.5 ms, receiver bandwidth = 5 kHz, matrix = 24 × 24, slice thickness = 10 mm).

The [1-^{13}C]pyruvate injections using a DNP system (SpinAligner, Polarize, Denmark) were prepared on site, and I supervised two injections of approximately 30 mL (250 mM) per subject. The bolus was administered intravenously at approximately 5 mL/s followed by a 20 mL saline flush, with imaging commencing 20 seconds after the injection was complete. I separated the two injections by approximately 1.5 hours. After each injection, I acquired either a fully sampled or undersampled ^{13}C CSI. I achieved undersampling by limiting both the FOV and the matrix size in each direction (fully sampled FOV = 240 mm × 240 mm, matrix = 24 × 24; undersampled FOV = 120 mm × 120 mm, matrix = 12 × 12), which preserved spatial resolution while reducing the temporal resolution from 36 s to 9 s. Three pigs were imaged in total.

In my post-processing workflow, I extracted metabolite signals from the spectra using the MATLAB trapz function to calculate the approximate integral via the trapezoidal method. For the dynamic spectra, I summed the magnitude signals within the two kidney ROIs per time point and normalized the plots to the highest pyruvate values. My image reconstruction utilized a 2D SENSE algorithm implemented in MATLAB (R2023b), where I used the pre-acquired sodium sensitivity maps to correct the carbon images. Critically, I carried out spectral processing first, followed by the SENSE reconstruction to unalias the

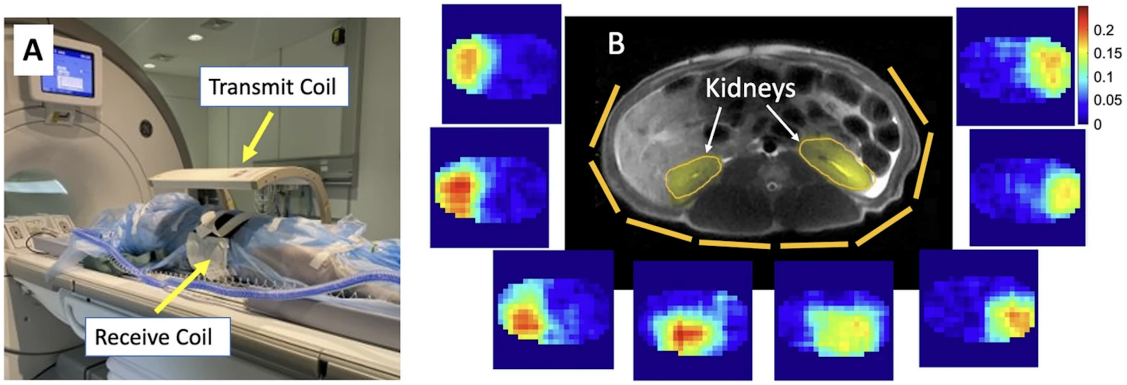


Figure 3.4: (A) A porcine experiment set up with a supine pig inside the clamshell transmit coil and the torso wrapped with the flexible receive coil. (B) Axial proton image of the porcine torso shows the location of the kidneys. The yellow lines visualize the flexible receive coil element positions and coil-wise sodium sensitivity images surround the proton image.

spatial data. I generated signal and g-factor maps using the same methodology established in my phantom study.

3.4 Results

3.4.1 Phantom Study

The coil hardware, experimental set up, ^{23}Na sensitivity maps for SENSE reconstruction of undersampled ^{13}C images, and results acquired from a Specific Anthropomorphic Mannequin (SAM) phantom is summarised in Figure 3.2. I illustrated the experimental setup consisting of a ^{13}C clamshell transmit coil (Figure 3.2A), a ^{23}Na/^{13}C receive coil (Figure 3.2B), and a head phantom (Figure 3.2C), surrounded by the localised ^{23}Na sensitivities of the eight receiver elements. The sensitivity dropped rapidly away from the coil elements. Because the loop coil was too short to wrap entirely around the phantom, I noted a blind spot in the sensitivity profile at the apex where no elements were present, as illustrated in Figure 3.2C. Although the phantom was technically too large for the receive coil, I found that this did not impede the reconstruction of images for this phantom or the subsequent *in vivo* study. In the conventional, fully sampled CSI (Figure 3.2D), I showed the ground truth signals acquired at the carbon frequency. The signal decreased toward the centre of the phantom, and reflected the "blind spot"

at the top of the arc. In the retrospectively undersampled CSI (Figure 3.2E), I demonstrated the wrap-around artefact characteristic of accelerated parallel imaging, where the anatomy was replicated and overlapped onto the original structure. I then performed a SENSE reconstruction on the undersampled image (Figure 3.2F), correcting the aliasing through the use of the sodium sensitivity maps. Through this process, I recovered the oval shape of the phantom, the opening at the top, and the gradual signal decrease toward the centre (Figure 3.2F). I demonstrated the resulting g-factor map in Figure 3.2G, where the yellow-red regions on the upper left and right-hand sides highlighted the noise amplification inherent to the coil hardware. Because this sodium-sensitivity SENSE reconstruction (Figure 3.2F) recovered an image highly similar to the fully sampled CSI (Figure 3.2D), I subsequently implemented this approach *in vivo*.

3.4.2 In-Vivo Porcine Study

Surrounding the anatomy in Figure 3.4B, are the sodium sensitivity profiles I acquired for each receiver element. I evaluated the performance of the SENSE reconstruction by comparing pyruvate maps generated via conventional fully sampled CSI against those from accelerated CSI methods. Using the kidney ROIs defined in the axial T_2-weighted proton image (Figure 3.5A), I observed high-intensity pyruvate signal corresponding to the kidneys in the fully sampled acquisition (Figure 3.5B). By reducing the FOV in the undersampled acquisition, I induced the expected aliasing artefacts shown in Figure 3.5C. I then applied SENSE reconstruction using the ^{23}Na sensitivity maps, which successfully recovered the signal in the kidneys (Figure 3.5D). The resulting unaliased image demonstrated high structural similarity to the fully sampled reference, confirming the removal of folding artefacts. Finally, I quantified the coil performance using a g-factor map (Figure 3.5E), where I found the highest values localised at the top of the torso, safely distal to the kidneys.

I tracked the metabolic dynamics in the kidneys of a separate porcine subject, as illustrated in Figure 3.6. In both the fully sampled (Figure 3.6A) and accelerated (Figure 3.6B) acquisitions, I found the spectral peaks were well-separated, allowing

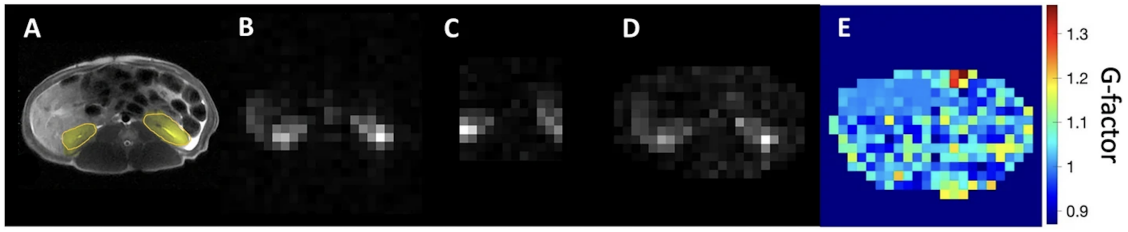


Figure 3.5: (A) Axial T2 proton image of the pig anatomy with the kidneys highlighted in yellow. (B) [^{13}C -1] pyruvate signal from fully sampled CSI taken over 36 seconds. (C) [^{13}C -1] Pyruvate signal from undersampled CSI taken over 9 seconds. (D) [^{13}C -1] pyruvate SENSE Reconstructed CSI based on the undersampled CSI and the sodium sensitivity image. (E) The g-factor map of R = 4, where acquisition is accelerated by R = 2 in both the frequency and phase directions.

me to capture high-quality metabolic data. By implementing the accelerated CSI acquisition, I achieved four-times higher temporal resolution, resulting in a more detailed decay time course compared to the standard acquisition (Figure 3.6A vs B).

I utilised ^{13}C imaging to demonstrate the acceleration achieved through SENSE reconstruction of metabolite ratio distributions in healthy porcine kidneys. In the top row of Figure 3.7, I generated ratio maps for lactate-to-pyruvate, alanine-to-pyruvate, and bicarbonate-to-pyruvate using data acquired over 36 s. In the second row, I illustrated the aliasing caused by the reduced FOV and the corresponding acceleration down to 9 s per scan. My SENSE-reconstructed images (Figure 3.7, bottom row) successfully recovered the folding effect at the kidneys, exhibiting ratio distributions highly consistent with the fully sampled maps.

In Figure 3.1, I produced the individual mean and standard deviation of these metabolic ratios per porcine subject (numbered "Pig 1", "Pig 2", and "Pig 3"), summing scans to a total acquisition time of 72 seconds. While I omitted the fully sampled CSI for Pig 3 due to a technical issue during acquisition, I calculated the mean values for Pig 1 and Pig 2 as: lactate/pyruvate (0.44 ± 0.11), alanine/pyruvate (0.19 ± 0.04), and bicarbonate/pyruvate (0.09 ± 0.01). Across all three subjects, my SENSE-reconstructed maps yielded ratios of: lactate/pyruvate (0.31 ± 0.02), alanine/pyruvate (0.23 ± 0.03), and bicarbonate/pyruvate (0.13 ± 0.02). Furthermore, I determined the coefficients of variation for these reconstructed ratios within the kidneys to be 6.6% for lactate, 13% for alanine, and 15% for bicarbonate.

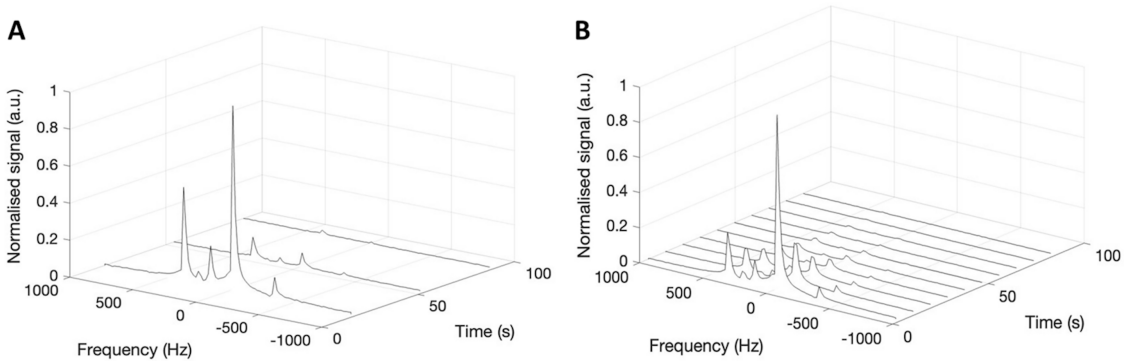


Figure 3.6: (A) Dynamic ^{13}C spectra of the fully sampled CSI in the kidneys of a pig. (B) Dynamic ^{13}C spectra of the undersampled CSI prior to the SENSE reconstruction in the kidneys of the same pig.

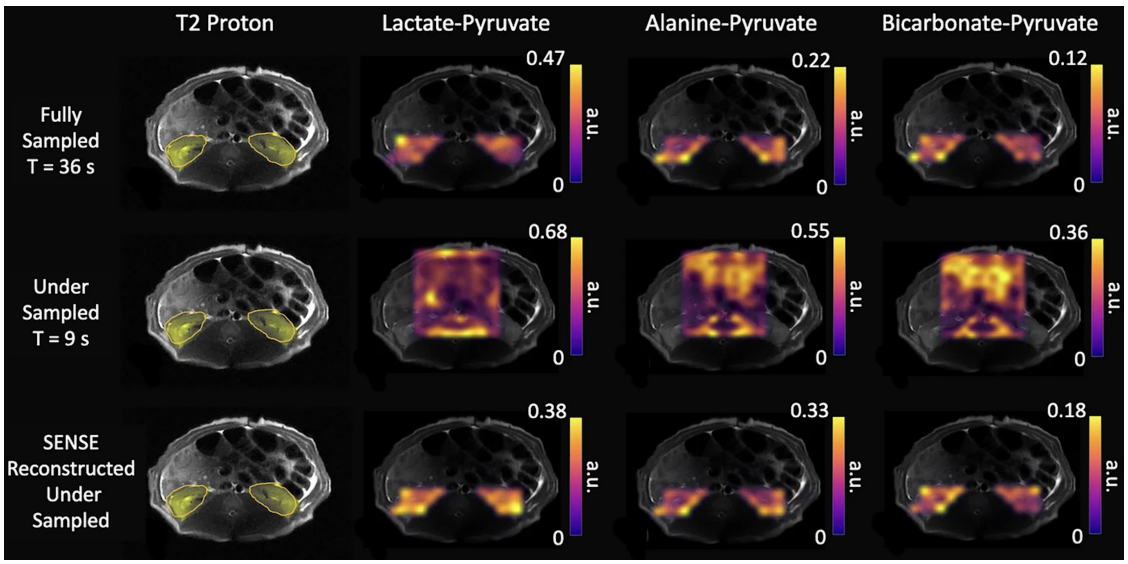


Figure 3.7: Ratio maps of lactate-to-pyruvate, alanine-to-pyruvate, and bicarbonate-to-pyruvate were calculated to show metabolism in the kidneys in the single scan fully sampled image, the undersampled image, and the SENSE reconstructed undersampled image of a healthy representative porcine subject. The axial T2 images have also been shown to indicate the kidney locations in yellow ROIs.

3.5 Discussion

Although SENSE-CSI is not the fastest acquisition approach currently available, it combines a well-established, easy-to-use, and robust sequence with hardware-enabled acceleration that provides a balance of spectral, temporal, and spatial resolution necessary for 13-carbon HP MRI, characterised by fast decaying signals with dynamic metabolic conversion. For example, fast imaging using SPSP RF pulses allows for very efficient scan times that are much faster than that implemented in this study.

Table 3.1: Mean ± SD metabolite-to-pyruvate ratios measured in the kidneys for each porcine subject. Fully sampled data were unavailable for Pig 3 due to a technical issue that prevented image acquisition.

Metabolite	**Pig 1 Full**	**Pig 1 SENSE**	**Pig 2 Full**	**Pig 2 SENSE**	**Pig 3 SENSE**
Lac:pyr	0.36 ± 0.07	0.29 ± 0.04	0.52 ± 0.16	0.32 ± 0.10	0.33 ± 0.03
Ala:pyr	0.16 ± 0.04	0.21 ± 0.05	0.22 ± 0.08	0.21 ± 0.06	0.26 ± 0.05
Bic:pyr	0.08 ± 0.02	0.11 ± 0.03	0.10 ± 0.02	0.15 ± 0.04	0.13 ± 0.02

However, SPSP approaches rely on a known spectrum with a well-separated profile to compensate for the low spectral resolution making the technique vulnerable to centre frequency shifts or inhomogeneous fields. IDEAL has also been used with spiral trajectories, EPI and other k space under sampling methods to create very fast scans but suffers from bleed through artefacts. Compared to the method used in this study, imaging with shortened TR, such as bSSFP, is much faster at data acquisition but often has low spectral resolution. Fast spectroscopic methods have good spectral resolution, although inferior to CSI, and harness frequency-encodes to reduce scan time such as spiral/radial k-space trajectories and EPSI. While offering very short scan times, fast spectroscopic imaging is limited by the hardware gradient slew rate in order to obtain maximum acceleration. Non-cartesian k space trajectories also produce artefacts that requires further complex post-processing to recover an interpretable image.

My work demonstrated the feasibility of accelerating ^{13}C MRI by combining a standard 2D CSI sequence with SENSE reconstruction, providing a robust pathway for alias-corrected metabolic imaging. I successfully applied geometric corrections for coil sensitivity and folding artefacts in both phantom and *in vivo* porcine models. This acceleration method proved effective for obtaining metabolic images with high temporal resolution while maintaining spatial and spectral integrity, particularly in clinical scenarios where ^{13}C sensitivity profiles cannot be pre-acquired.

A key strength of this approach was demonstrated *in vivo*, where I achieved a four-fold temporal acceleration by reducing the FOV. Importantly, the coil hardware and FOV parameters utilized in this study are directly translatable to clinical settings, such as human metabolic brain imaging, without modification. By

employing a dual-tuned (^{23}Na/^{13}C) receive coil, I eliminated the need for additional hyperpolarised injections or image co-registration. The resulting 9-second temporal resolution provided detailed decay curves essential for kinetic modelling. In cases with unknown metabolite spectra or limited HP MRI specialist knowledge, this robust acceleration method ensures comprehensive capture of metabolic signals, mitigating the risk of missing spectral data.

In my analysis of the porcine subjects, I found that SENSE-reconstructed metabolic ratio maps were reproducible, with coefficients of variation within 15% across the three subjects. I observed a slight decrease in the mean lactate-to-pyruvate ratio in the undersampled data compared to the fully sampled reference. However, given the small sample size and the fact that ratios for both fully sampled and undersampled data remained within one standard deviation of each other, I could not determine if this was a systematic hardware difference or a result of biological variation, such as anaesthetic timing or the physiological response to sequential injections.

The primary limitation of this approach is the reliance on specialised multi-element receiver hardware, which in this study had the maximum acceleration factor of four. In addition, the proposed method requires a dual-tuned (^{23}Na/^{13}C) RF system to obtain the sodium sensitivity maps used for SENSE reconstruction. Most hyperpolarised MRI research centres are equipped with dedicated multi-nuclear RF hardware, but the benefits of this approach come at the cost of requiring a specialised dual-tuned RF system. Wider clinical adoption would therefore require similar multi-nuclear hardware to be routinely available. Although other accelerated ^{13}C imaging methods are better at achieving very high temporal resolution, such as SPSP, CSI is readily available on most major MRI systems, robust to B0 fluctuations, and easy to use. To prioritize temporal resolution, I constrained the spectral bandwidth and sampling points, resulting in a coarse spectral resolution of 1.22 ppm, and 12.1 ppm in one acquisition. Under these conditions, I determined that complex spectral fitting algorithms were unnecessary; instead, I utilized trapezoidal integration (`trapz` in MATLAB), which provided effective quantification for the well-separated, high-SNR

peaks observed. Future work integrating advanced spectral fitting approaches, as discussed in Chapter 4, could further improve precision in lower-SNR scenarios.

In the current implementation, the SENSE reconstruction utilized an identity matrix for the noise covariance term (Ψ), effectively treating each of the eight receiver elements with equal weighting. In practice, however, the individual loops of the flexible sodium-carbon array exhibit varying sensitivity profiles depending on their proximity to the anatomy, resulting in non-uniform SNR across the elements. By incorporating a measured noise covariance matrix into the GLS framework, future iterations of this work could implement optimal coil weighting. Such a refinement would suppress contributions from noisier, more distant elements while prioritizing signal from those with higher local sensitivity. This would be particularly beneficial for managing the noise floor in low-SNR metabolic maps, such as bicarbonate, where the geometry factor penalties are most acutely felt.

Overall, the combination of accelerated CSI and SENSE reconstruction provided ^{13}C images with sufficient spatial and temporal resolution to observe renal metabolism through a computationally inexpensive reconstruction. By further optimizing parameters such as flip angle, bandwidth and spectral points, there is scope to further enhance the signal acquired from a finite hyperpolarised pool, depending on the resolution requirements of the study. This novel integration of sodium sensitivity profiles with accelerated CSI offers a promising, robust solution for metabolic imaging across diverse clinical scenarios.

3.6 Conclusion and Future Directions

I have demonstrated that utilising ^{23}Na sensitivity profiles to inform SENSE-reconstructed CSI successfully produces alias-corrected metabolic images in both phantom and *in vivo* porcine kidney studies. By recovering unaliased images from accelerated acquisitions, I have shown that SENSE-CSI is a robust, and easy-to-implement method for reducing the scan time of hyperpolarised experiments.

The clinical significance of this framework lies in its efficiency and reliability. Because hyperpolarised ^{13}C signal is a finite and costly resource, the ability to pre-acquire sensitivity maps using the high-signal, endogenous sodium pool eliminates the need for calibration injections. This brings the technique closer to the throughput required for clinical implementation, where patient and MRI time are at a premium. Furthermore, by building this acceleration on a standard CSI, I have provided a translation ready technique for clinical sites, unlike the more complex spectral-spatial readouts that can fail due to frequency shifts or poor shimming.

By achieving a four-fold temporal acceleration to 9 seconds per scan, this work enables the detailed kinetic modelling necessary to distinguish healthy from pathological metabolism in real-time. Future research will focus on integrating a measured noise covariance matrix into the reconstruction to optimise coil weighting and explore advanced spectral fitting models to enhance precision. Further work into optimising k-space trajectories may also enhance acceleration. Ultimately, this framework offers a promising, robust solution that balances the high-performance demands of hyperpolarised MRI with the practical needs of clinical diagnostic settings.

3.7 Appendix

To verify the mathematical robustness of the SENSE unaliasing algorithm, I performed a series of controlled tests on a standard imaging phantom before proceeding to *in vivo* studies. This allowed me to isolate the performance of the 2D Cartesian reconstruction from the physiological variables of a live subject. Before the head phantom, I utilized a 10.64 kg silicone oil sphere (Body Phantom Sphere, Model 2360025REV3, GE Healthcare) as a calibration subject with the proton 21-channel head-neck coil from GE Healthcare. I used proton imaging and took the peak amplitude as signal. Reference datasets provided by Chiew [99] were also used to validate the code against established parallel imaging benchmarks (figure 3.11).

To simulate acceleration, I performed retrospective undersampling by manipulating the fully sampled k-space data. To achieve an acceleration factor of $R = 2$, I systematically deleted every other line of the k-space matrix and replaced them with zeros. This manipulation was initially evaluated in one dimension (figure 3.8). I simulated phase direction acceleration by skipping lines along the vertical axis, resulting in vertical aliasing (top-down folding). For frequency direction acceleration, I skipped columns along the horizontal axis, resulting in horizontal aliasing (left-right folding). I then extended the manipulation to 2D by undersampling both the x and y directions (figure 3.9). The core logic of the reconstruction relies on the fact that phase and frequency encoding directions are orthogonal. Because of this orthogonality, the aliasing problem can be treated as separable. I solved the SENSE equations for each dimension independently, effectively unfolding the image in a sequential 1D-by-1D process. This improved the computational efficiency of the MATLAB implementation and ensured the stability of the matrix inversion. The resulting undersampled images were reconstructed using my custom SENSE algorithm and evaluated using g-factor maps (figure 3.10). By testing on the silicone sphere with high SNR proton imaging, I confirmed that the algorithm could successfully recover the phantom's geometry. The g-maps provided a baseline quantification of the noise amplification inherent to the 8-element coil array geometry.

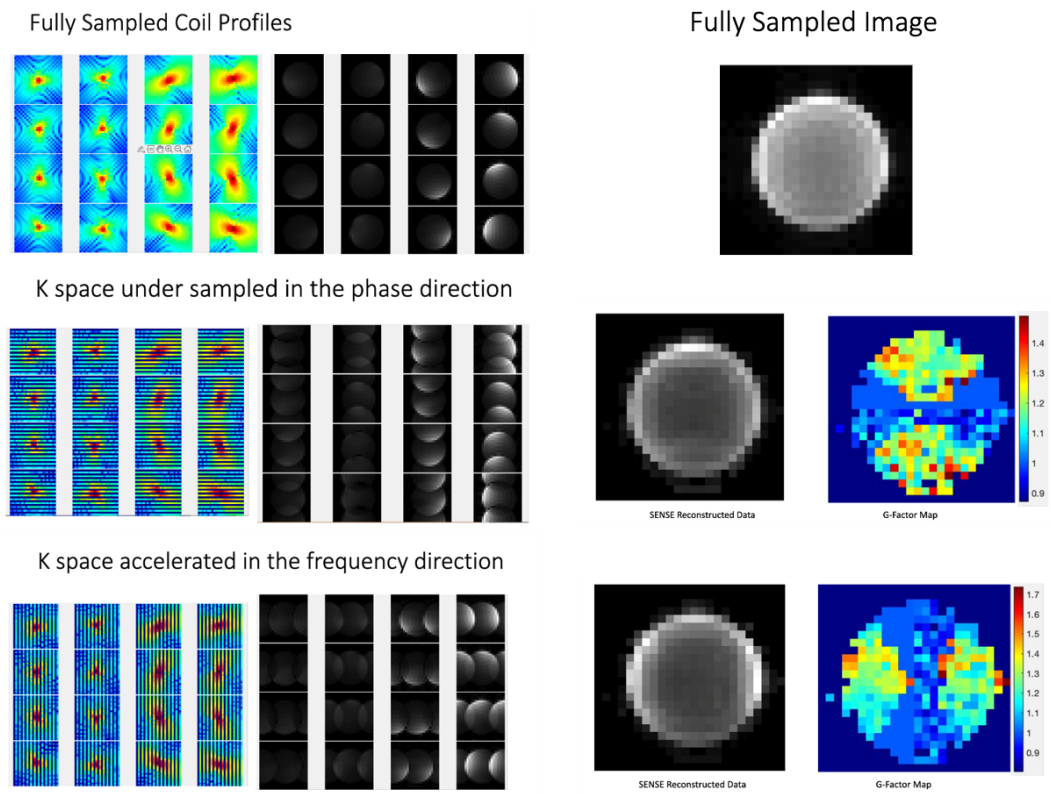


Figure 3.8: Acceleration in one direction, and the reconstructed images of the sphere phantom.

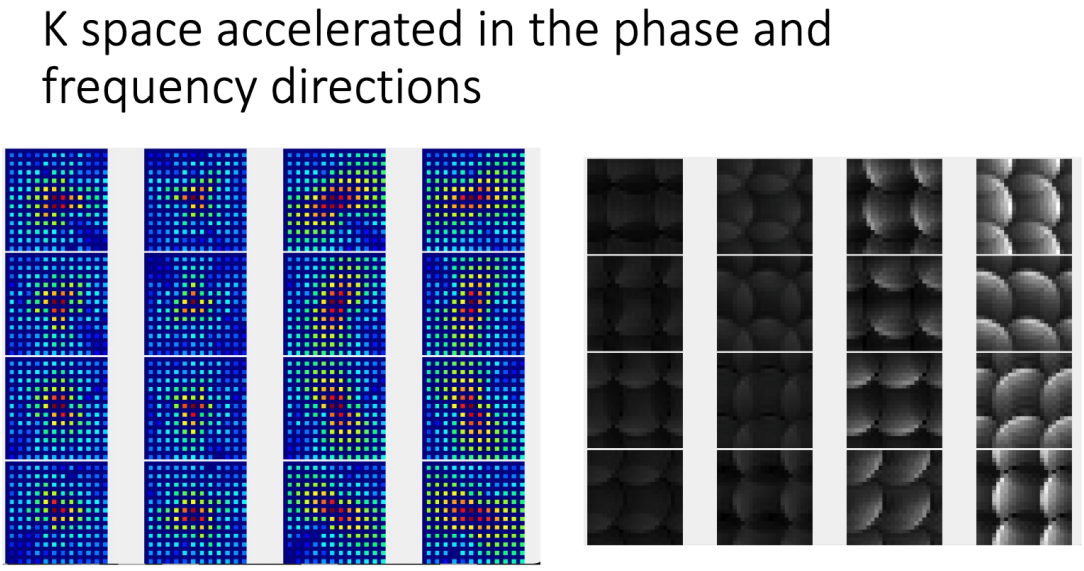


Figure 3.9: Acceleration in two directions, and the accelerated images of the sphere phantom.

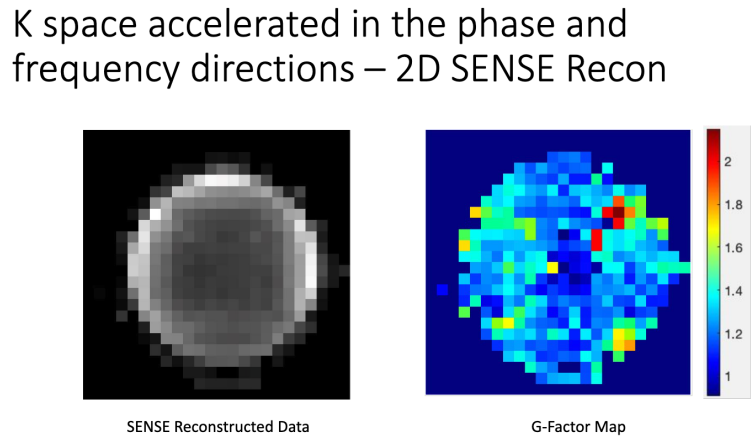


Figure 3.10: SENSE reconstructed image of the 2D accelerated acquisition for the sphere phantom, and its g-factor map.

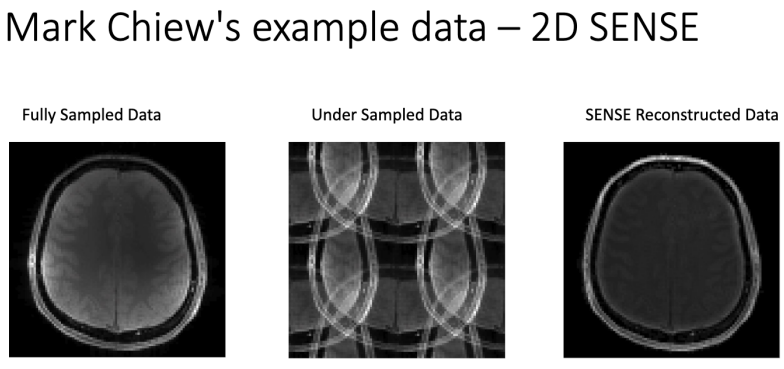


Figure 3.11: Proton head phantom from Mark Chiew's example data reconstructed through the acceleration algorithm.

4

Denoising HP MR Spectroscopy using Spectral Fitting

Contents

4.1 Abstract 139
4.2 Introduction 140
4.2.1 Motivation 140
4.2.2 Spectral Fitting in Spectroscopy 141
4.2.3 Kinetic Modelling and Metabolic Rate Constants 142
4.2.4 Spectral Challenges in HP MRI 142
4.2.5 Current Approaches in Spectral Fitting 146
4.2.6 Proposed Dynamic Fitting Technique 150
4.2.7 Chapter Objectives 151
4.3 Methods 152
4.3.1 Data Acquisition and Study Cohorts 152
4.3.2 Spectral Fitting and Kinetic Modelling Implementation 153
4.3.3 Development of an In-house Graphical User Interface . . 154
4.3.4 Peak Fit Model Comparisons 157
4.4 Results 165
4.4.1 Rat Heart Study 165
4.4.2 Mouse Heart Study 169
4.4.3 Human Brain 171
4.4.4 Human Liver 174
4.5 Discussion 180
4.6 Conclusion and Future Directions 183

4.1 Abstract

Robust quantification of metabolites in hyperpolarised (HP) ^{13}C MRI is often hindered by rapid signal decay and low signal-to-noise ratio (SNR) for downstream metabolites like bicarbonate in time-resolved datasets. The challenges in fitting independent time points in the dynamic spectra include fitting unstable noise and the overlap from the dominant pyruvate signal. In the work described in this chapter, I have developed a 'dynamic fit' algorithm that leverages high-SNR information obtained from a summed spectrum as a-priori knowledge to constrain the fitting of individual low-SNR time points.

Using an OXSA-based MATLAB workflow, I validated this approach across rat heart ($n = 7$), mouse heart ($n = 10$), human brain ($n = 1$), and human liver ($n = 1$) datasets. The dynamic fit consistently outperformed standard independent fitting, reducing the mean Cramér-Rao Lower Bound (CRLB) by up to 78% and improving kinetic model Root Mean Square Error (RMSE) by up to 76%. In a mouse model of high-fat diet (HFD) induced obesity, the improvement from the dynamic fit revealed significant physiological differences in bicarbonate kinetics ($k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$, $p = 0.03$; $k_{\mathrm{x\,bic}}$, $p = 0.03$) that were otherwise obscured by noise in the simple fit ($p = 0.28$ and $p = 0.22$, respectively). This improvement in spectral quantification enabled a 79% reduction in the required sample size, from 28 to 6 subjects, to achieve 95% statistical power for the $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ endpoint. These results demonstrate that integrating a-priori information from summed spectra provides a robust framework for metabolic quantification, enhancing the statistical power of HP ^{13}C studies and supporting future translation into clinical diagnostic settings as demonstrated by a first step feasibility in single healthy human volunteer examples.

4.2 Introduction

4.2.1 Motivation

A key objective of hyperpolarised (HP) ^{13}C MRI is the robust quantification of metabolic rate constants, which serve as non-invasive biomarkers for real time metabolism to inform upon early diagnosis, prognosis, and treatment response. In an idealised spectroscopy experiment, the acquired Free Induction Decay (FID) would yield a spectrum of well-separated, high-SNR Lorentzian peaks. This is because an ideal FID follows a first-order exponential decay (e^{-t/T_2^*}), the Fourier Transform of which is mathematically defined as a Lorentzian function in the frequency domain.

Within this spectrum, the area under the curve (AUC) is a transient signal representing a combination of metabolite concentration, the excitation flip angle, and transverse relaxation (T_2^*). The complexity of this quantification arises because these measurements evolve constantly due to longitudinal relaxation (T_1) and the dynamics of the hyperpolarised bolus. Furthermore, HP ^{13}C data is characterised by a rapid signal-to-noise ratio (SNR) decrease over a narrow temporal window. This decay, coupled with the low signal of downstream metabolites following a [1-^{13}C]pyruvate injection, often leaves the dynamic signal buried within baseline noise, complex peak shapes, and the spectral overlap of the dominant pyruvate peak.

Conventional analytical methods frequently address these SNR limitations by summing all temporal data into a single spectrum. While this "summed spectrum" approach is effective for identifying metabolite locations and improving SNR for static visualisation, it inherently sacrifices the time-resolved information essential for characterising metabolic kinetics. Conversely, when temporal resolution is prioritised, the resulting individual spectra are often too noisy for reliable mathematical fitting. Current techniques often attempt to denoise or tweak these low-SNR frames in isolation, ignoring the high-quality global information that is readily available from the total acquired dataset.

I propose that the high-SNR information contained within the summed spectra should be leveraged as a-priori knowledge to constrain the fitting of the low-SNR

dynamic time points. By assuming that fundamental peak characteristics, such as peak shape, chemical shift, and phase, remain relatively constant throughout the acquisition, the summed spectra can be used to define precise fitting boundaries and initial conditions. While *in vivo* factors such as B_0 drift and physiological motion can introduce temporal phase and lineshape fluctuations, treating these parameters as stationary constraints allows the fitting algorithm to more effectively distinguish between metabolic signals and noise, as the noise lacks the consistent lineshape of the global signal. The work described in this chapter focuses on integrating this a-priori information into the model, providing a robust quantification of metabolic changes in time-resolved datasets without the loss of kinetic detail.

4.2.2 Spectral Fitting in Spectroscopy

Spectral fitting is the computational process required to extract quantitative metabolic information from the raw MR signal. While the underlying physics of the chemical shift and the Free Induction Decay (FID) provide the basis for signal generation, the objective of fitting is to determine the area under the curve (AUC) for individual metabolites, which is proportional to the concentration of the hyperpolarised nuclei. Established gold-standard methods, such as LCModel and AMARES further discussed in section 4.2.5, provide fitting approaches that are extensively utilised in conventional proton MRS and have been widely adapted for other X-nuclei. When applied to static MRS experiments, traditional spectral fitting algorithms typically assume steady-state conditions, fixed metabolite line shapes, and relatively stable baseline environments.

In contrast, hyperpolarised ^{13}C spectra present unique challenges for quantification. Because the magnetisation is non-renewable and subject to constant T_1 decay and RF depletion, the signal-to-noise ratio (SNR) varies significantly throughout a dynamic acquisition. Furthermore, spectral line shapes may fluctuate due to physiological motion or B_0 field instabilities over the two-minute imaging window. These temporal variations must therefore be accommodated by any fitting model, including approaches that incorporate prior spectral information. Consequently,

conventional fitting tools often struggle with accuracy when applied to the low-SNR regions of these transient datasets.

4.2.3 Kinetic Modelling and Metabolic Rate Constants

Kinetic modelling provides the mathematical link between the measured time-course of hyperpolarised signals and the underlying biological activity. In HP ^{13}C MRI, the distribution and conversion of the tracer are typically described using a compartmental model. The most common framework is the "Two-Site Exchange Model", which assumes that metabolites exist in distinct, well-mixed "pools" (e.g., a pyruvate pool and a lactate pool) and move between them via first-order kinetics.

The evolution of the magnetisation for a metabolite pool, such as lactate (M_L), can be described by the following differential equation:

$$\frac{dM_L(t)}{dt} = k_{PL}M_P(t) - (k_{LP} + R_{1,L} + \mathrm{RF}_{loss,L})M_L(t) \tag{4.1}$$

In this expression, $M_P(t)$ is the magnetisation of the parent pyruvate pool, k_{LP} represents the reverse reaction rate, $R_{1,L}$ is the longitudinal relaxation rate ($1/T_1$), and $\mathrm{RF}_{loss,L}$ accounts for the depletion of the signal due to the application of RF flip angles.

A significant variable in this model is the Arterial Input Function (AIF), or the bolus shape. The AIF describes the delivery of the hyperpolarised pyruvate to the organ of interest following injection. Because the conversion to downstream metabolites is dependent on the concentration of the available substrate, the shape and timing of this arrival must be accurately modelled to avoid bias in the calculation of k_{PL} and k_{PB}.

4.2.4 Spectral Challenges in HP MRI

Proton vs HP Carbon Spectra

NMR theory states that the size and shape of a spectral peak predominantly depends on 5 factors: 1) the concentration of the measured nuclei, 2) the T1 and T2 times

of the metabolite, 3) the magnetic field homogeneity, 4) the presence of hidden or overlapping peaks and 5) whether the line is a singlet or a multiplet [102].

The challenges with spectral fitting differ between proton and HP carbon MRS. Proton MRS, specifically in the brain, is characterised by a high density of metabolites, often resulting in significant peak overlap and complexity in the baseline [103]. A major challenge here is that water and fat signals are much more abundant than the metabolites of interest and can easily obscure lower-concentration metabolites [103]. Furthermore, because the metabolites sit within a narrow 10 ppm window, the lack of dispersion means that scalar coupling (J-coupling) interactions between protons on the same molecule lead to complex, overlapping multiplet patterns [104]. While ^{1}H-MRS is susceptible to B_0 and B_1 drifts over long acquisition times, its primary advantage is a relatively high thermal SNR and a steady-state signal that allows for extensive temporal averaging.

In contrast, hyperpolarised ^{13}C MRI presents a different set of challenges. The spectra are relatively sparse, with well-separated chemical shifts and a limited number of expected metabolites (typically pyruvate, pyruvate-hydrate, lactate, alanine and bicarbonate following the injection of hyperpolarised [1-13C]pyruvate) [105]. However, any alterations in the FID can change the idealistic Lorentzian peak in the spectra (more details in the following Dynamic Peak Morphology subsection 4.2.4). In particular, the signal from the pyruvate injection bolus is orders of magnitude higher than that of the downstream products. Similar to water in proton MRS, the broadened base of this dominant resonance ("pyruvate wings") can overlap with and obscure neighbouring low-SNR metabolite peaks. In hyperpolarised carbon-13 MRI, this challenge is compounded by the rapidly changing metabolite amplitudes throughout the dynamic acquisition. There is also the issue of the irreversible loss of signal due to T_1 relaxation and RF excitations [106]. This makes acquiring data a battle against time to get high-resolution information before the signal reaches the noise floor. Because the HP signal decays so rapidly and is non-renewable, these smaller peaks are frequently lost to the noise. Additionally, the baseline is

further complicated by slice-selection gradients and refocusing requirements, which introduce phase rolls that must be accounted for during quantification.

Dynamic Peak Morphology

Given the challenges of the HP ^{13}C spectra, analysing its time course brings further complexities. As data is collected across a time window, the peak shape and phase slightly fluctuates due to flow, B_0 drift, and physiological motion. In addition, rapid FID decay causes the Lorentzian peak to broaden. This is particularly impactful in the fitting quantification when this leads to a broad pyruvate linewidth. The broad wings from the base of the pyruvate peak artificially change the amplitude of adjacent metabolic peaks.

Following the injection of hyperpolarised [1-13C]pyruvate, the dominant signal in the time course is the injected pyruvate. The assumption when making kinetic models for the temporal evolution of the pyruvate signal is that the measured amplitude represents the total available pool for enzymatic conversion. However, signal originates from distinct environments: intra-luminal (blood vessel), interstitial (extracellular), and intracellular (inside the cell at the metabolic site). Following the intravenous injection, the pyruvate bolus must first transit the vascular system, permeate the tissue vasculature, and cross the cellular membrane (via Monocarboxylate Transporters, MCTs) before it can be metabolised [107, 108]. During the initial phase of bolus arrival, the signal is dominated by intravascular pyruvate that has not yet entered the cell. As pyruvate subsequently distributes between the vascular, extracellular and intracellular compartments, signals arising from these distinct physiological environments contribute to the observed pyruvate resonance, producing a more complex peak shape.

In terms of organ-specific challenges, the blood flow within the heart creates significant partial volume effects. The initial bolus of pyruvate first arrives in the high-volume heart chambers (right and left ventricles), producing a massive signal that is dominated by the blood pool. Then as pyruvate is transported from the blood into the heart muscle via the coronary arteries, there is pyruvate available at the

metabolic site. Intracellular conversion to lactate (via LDH), alanine (via ALT) and bicarbonate (via PDH) occurs within the myocardium [109]. Therefore, significant partial volume effects between the chamber blood and the thin myocardial wall are observed in cardiac metabolic imaging. The complexity of multiple pyruvate signals from different compartments means a complex peak shape. Despite challenges to quantification, human HP ^{13}C cardiac studies have found clinical insights into cardiac hypertrophy, ischaemia, and heart failure[110–112].

In the brain, the metabolic profile expects to see pyruvate, pyruvate-hydrate, lactate, and bicarbonate. Alanine signal typically originates from the surrounding scalp muscles, and is often omitted from analysis [37, 113]. Once the pyruvate is injected, the delivery rate is limited by the blood-brain barrier and requires some time for delivery to the site of metabolic conversion[114]. Some studies of pathology/physiology in the human brain performed using HP MRI include imaging of tumours, traumatic brain injury and the effects of ageing [115–120].

The liver is expected to produce all five metabolites: pyruvate, pyruvate-hydrate (PH), alanine, lactate and bicarbonate, although PH is often excluded from analysis as it is not metabolically active [121, 122]. However, the proximity to large vessels (vena cava, portal vein) causes significant partial volume effects with intra-luminal pyruvate[123]. Due to the location in the torso, the liver is susceptible to B_0 fluctuations and spatial blurring due to the breathing cycle, which complicates the stability of the peak shape. Early HP ^{13}C hepatic metabolism work in animals and humans has shown promise to address clinical needs such as metabolic dysfunction-associated steatotic liver disease (MASLD), cirrhosis, diabetes and liver injury [121, 123–127].

Baseline Distortions

In thermal ^{1}H MRS, baseline distortion from macromolecules is a well-characterised phenomenon. In HP ^{13}C MRI, most baseline issues are handled on a study by study basis, where techniques can range from a background-free assumption[128], drawing a straight line between the integration interval[88], subtracting a spline across

the spectra[129–132], to packages that have inbuilt baseline capabilities[133–135]. This topic is complex because many factors, such as imperfections in the gradient waveforms, Gibbs ringing (truncation artifacts from finite k-space sampling), and B_0 field inhomogeneities can all alter the spectra. Furthermore, in HP MRI, sometimes it is useful to use acquisition schemes such as slice selection. The difficulty with slice selection is that the delay between excitation and acquisition (length of the start time) causes a first-order linear phase shift in the FID. If uncorrected, these phase errors lead to twisted peak shapes (mixing dispersive and absorptive components), which ruins the validity of peak integration and Lorentzian fitting. Accurate baseline modelling therefore remains a source of quantification error in HP carbon-13 MRI and continues to be an active area of research.

4.2.5 Current Approaches in Spectral Fitting

While spectral fitting for proton (^{1}H) spectroscopy is well-established in clinical research settings, the ecosystem for HP MRI and X-nuclei is more fragmented, consisting a variety of legacy and modern software toolboxes.

There are broadly two mathematical approaches to spectral fitting. One is the frequency based methods, where the FID is Fourier transformed and the resulting spectral peak shapes are fitted. In the frequency-based category, the HP MRI community have used adaptations of the LCModel software (written in Fortran by Provencher [130]) [136, 137]. The LCModel takes a basis-set, i.e. the perfect metabolite shape simulated in a lab, and matches a weighted linear combination of the basis set to the data. Standard clinical research tools, such as Osprey [138], are primarily optimised for frequency-domain linear combination modelling of steady-state metabolites in the brain. Similarly, FSL-MRS [139] provides a modern Bayesian framework that is highly effective for ^{1}H datasets but requires significant customisation of basis sets for ^{13}C applications. INSPECTOR [140] is another frequency-domain approach written in MATLAB that offers sophisticated density matrix simulations to generate these basis sets. To account for real-world variations, these packages also incorporate a flexible baseline component to handle

background signals and then applies phase and frequency shifts to the entire basis set simultaneously.

While the frequency-based strategy is highly effective for complex proton (^{1}H) spectra where many metabolites overlap, it can be rigid when applied to the dynamic nature of hyperpolarised ^{13}C experiments. Because these algorithms rely on static and ideal signal templates, they may struggle to accommodate the rapid lineshape distortions and SNR transitions that occur during a metabolic acquisition. For decades, LCModel was the expensive commercial industry standard, though it has recently transitioned to an open-source model.

Another spectral fitting approach is time-domain based algorithms, where the FID is modelled as a sum of exponentially decaying sinusoids. In the AMARES framework, the time-domain signal $y(t)$ is modelled as a sum of K exponentially decaying sinusoids [141, 142]:

$$y(t) = \sum_{k=1}^{K} a_k e^{i\phi_k} e^{-d_k t} e^{-\sigma_k^2 t^2} e^{i2\pi f_k t} + \epsilon(t) \tag{4.2}$$

In this formulation, a_k and ϕ_k represent the amplitude and phase of the k-th peak, respectively, while f_k denotes the frequency shift and $\epsilon(t)$ represents complex Gaussian white noise. The signal decay is characterised by the Lorentzian damping constant d_k, representing T_2 relaxation and homogeneous broadening, and the Gaussian damping constant σ_k, which accounts for inhomogeneous broadening due to B_0 variations. By incorporating both terms, the model accommodates the Voigt lineshape, the mathematical convolution of Lorentzian and Gaussian profiles, which more accurately reflects the complex peak morphology typical of *in vivo* ^{13}C spectra. These parameters are iteratively optimised to match the raw FID. Subsequently, a Fourier transform is applied to the resulting fit for spectral visualisation and residual analysis.

Both LCModel and AMARES use an iterative least-squares procedure (Levenberg-Marquardt) to optimise the residual between the data and the predicted parameters [141].

Because LCModel-type software fits on an ideal-world basis-set, it handles data with absorptive peaks (aligned phases and symmetric shapes) very well, but cannot handle data that has phase errors, which is frequently seen in hyperpolarised experiments due to the high flow of pyruvate, B_1 inhomogeneity, or timing delays. In contrast, AMARES parameterises the phase (ϕ_k) and frequency (f_k) as free variables, allowing the algorithm to automatically calculate the optimal value for every individual peak. This flexibility makes AMARES robust to the zero-order and first-order phase shifts common in rapid spectral-spatial acquisitions [143]. Furthermore, HP ^{13}C spectra are typically sparse with a small number of well-separated peaks (pyruvate, lactate, alanine, bicarbonate). The complex unmixing power of LCModel is computationally unnecessary for such sparse data, whereas the parametric approach of AMARES is computationally efficient and easily constrained for these specific metabolites.

FID-driven AMARES implementations remain the standard for carbon imaging, widely utilised via jMRUI (Java)[134] and OXSA (MATLAB) [141], and the more recently published python counterpart: pyAMARES (Python)[143]. jMRUI is a proprietary software that is freely distributed but has been slow with creating releases, most recently released in 2023, and prior to that in 2017. In the software engineering world, where releases are continuously published, jMRUI is regarded more as a legacy tool that requires custom plug-ins for x-nuclei applications. Because it is proprietary, it cannot be easily coded and changes are subject to upgrade releases. jMRUI is available to download for Microsoft Windows and GNU/Linux systems. The VARPRO algorithm [144] is the predecessor of OXSA, and allows for the incorporation of prior knowledge, which is essential for stabilising fits as the signal-to-noise ratio (SNR) declines, but this software is more mathematically oriented than for clinical usage. Clinical software like MRSpecLab[145] and SIVIC [133] prioritises DICOM-to-PACs compliant pipelines that utilise integrated area methods for metabolite quantification, but it requires use of external plugins or complementary software for rigorous spectral analysis. VesPa [135] is a mix of both frequency and time domain, which uses a graphical interface for designing

Software	Language	Open Source?	Workflow Ease 13C	Primary Algorithm	Original Reference
OXSA	MATLAB	Yes	High (Dynamic loops)	AMARES	Purvis et al. (2017)
jMRUI	Java/C++	No	High (but no customisation)	AMARES	Stefan et al. (2009)
pyAMARES	Python	Yes	High (Script-based)	AMARES	C. Wang et al. (2023)
LCModel	Fortran	Yes	Low (Requires basis set)	Linear Combination	Provencher (1993)
FSL-MRS	Python	Yes	Medium (needs config)	Linear Combination	Anderson et al. (2021)
Osprey	MATLAB	Yes	Medium (optimised for 1H)	Linear Combination	Oeltzschner et al. (2020)
TARQUIN	C++	Yes	Low (Basis set needed)	Least-Squares	Wilson et al. (2011)
SIVIC	C++	Yes	High (best for clinic workflow)	Peak Integration	Crane et al. (2013)
VesPA	Python	Yes	Medium	Separable Nonlinear Least Squares	Soher et al. (2011)
INSPECTOR	MATLAB	Yes	Medium	Linear Combination	B. Juchem (2015)
VARPRO	Fortran/C	Yes	Low (Academic code)	Variable Projection	Van der Veen (1988)
FID-A	MATLAB	Yes	Pre-processing only	N/A (Toolkit)	Simpson et al. (2017)
MRSpecLab	Python	Yes	Medium (great for preprocessing)	Linear Combination	Xiao et al. (2025)
SpectrAI	Python	Yes	Low (requires training, but great GUI)	Deep Learning (CNN/RNN)	Horgan et al. (2021)
NNFit	C/MATLAB	Yes	Low (requires extensive training)	Artificial Neural Network (ANN)	Giuffrida et al. (2025)

Figure 4.1: Comprehensive comparison of spectral fitting software packages for X-nuclei spectroscopy.

custom RF pulses combined with spectral simulations from the PyGAMMA library, where all code is open-source on their GitHub [135].

Recent developments have introduced Deep Learning (DL) as a third way to address the computational bottleneck of iterative fitting in dynamic time-series data. Frameworks such as NeuralNetwork Fitting (NNFit) [146] and SpectrAI [147] utilise neural networks trained on synthetic datasets to perform quantification in milliseconds. NNFit is a self-supervised deep learning model trained on simulated spectra with varying noise and linewidths developed by Wang et al.[146] NNFit can quantify metabolite concentrations orders of magnitude faster than the iterative least-squares models (such as LCmodel or AMARES)[117]. While promising for real-time metabolic mapping, these methods currently rely on training data generated by classical algorithms (AMARES or LCModel) and inherits the biases of their supervisors. Table 4.1 summarises the comparisons of these X-nuclei packages, updated as of early 2026. I chose OXSA as the primary analysis platform for the work described in this chapter because it offers a native MATLAB implementation of the AMARES time-domain algorithm. For hyperpolarised ^{13}C, fitting in the time domain is more appropriate than frequency-domain approaches as it replaces rigid, idealised templates with a flexible parametric model better suited to sparse spectra.

Because the slice-selective RF pulses and gradient rephasing commonly used in HP MR protocols often lead to truncated FIDs and significant first-point phase errors, time-domain fitting allows these issues to be handled by the parametrisation of phase and frequency. In order to fit the spectra, OXSA requires a prior knowledge (PK) file that defines starting values and bounds to the parameters. The open-source nature of OXSA was also a major advantage, as it allowed me to directly modify the Prior Knowledge (PK) constraints. By integrating OXSA into a custom MATLAB pipeline, I was able to process dynamic datasets in a scripted, repeatable way. This ensured that the kinetic rate constants were derived from a consistent and mathematically transparent routine, rather than a "black-box" automated fit in the deep learning frameworks.

4.2.6 Proposed Dynamic Fitting Technique

Despite the advances in optimising spectral fitting, the existing toolboxes treat individual time points independently, failing to leverage the high-SNR information available in a summed spectrum to understand the metabolic time course. This gap in the existing software landscape gives way to the development of my dynamic fitting pipeline, which aims to improve quantification robustness by algorithmically linking the summed reference peak characteristics to the time-resolved analysis.

In this section, I introduce the novel concept of using a high-SNR summed spectrum as a reference for fitting individual time points. I build upon the existing AMARES-based OXSA workflow by updating the prior knowledge using high-SNR, summed spectra. My proposed workflow is illustrated in figure 4.2. I implemented the dynamic fitting pipeline by first summing the time-resolved data to generate a high-SNR reference spectrum. By fitting this summed spectrum, I obtained precise peak parameters, specifically the centre frequencies, phases, and linewidths, which I used as the basis for the prior knowledge in the subsequent individual time point analysis. These values provided well-informed initial estimates and enabled tight constraints during the non-linear least-squares fitting of the individual low-SNR time points. This approach assumes that these spectral parameters remain sufficiently

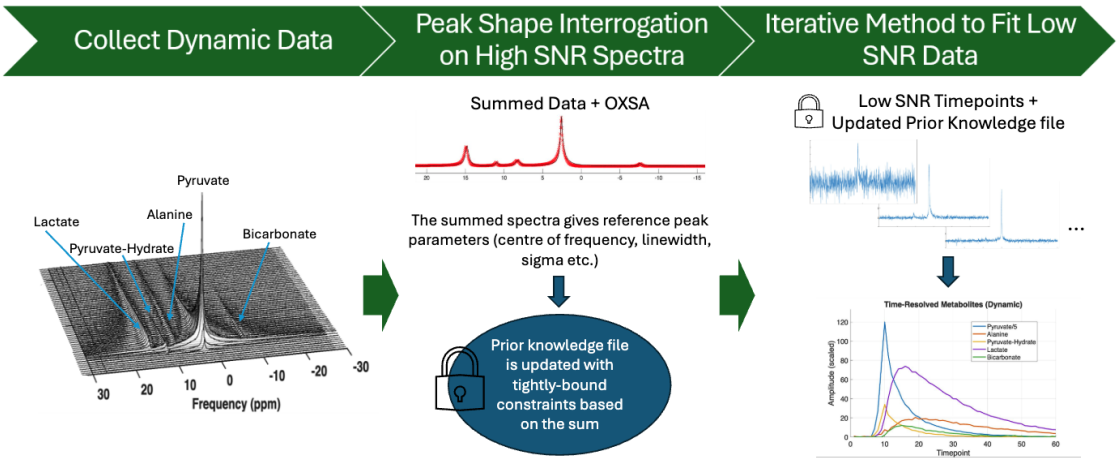


Figure 4.2: Overview of the iterative fitting pipeline for dynamic ^{13}C MRSI: Time-resolved data is first summed to produce a high-SNR reference spectrum, providing precise peak parameters for use as prior knowledge. These values serve as tight constraints and well-informed initial points during the subsequent non-linear least squares fitting of individual, low-SNR time points. This iterative workflow ensures robust spectral quantification and improves the accuracy of the resulting time-resolved metabolic profiles.

stable throughout the dynamic acquisition. Consequently, while the constraints suppress random fitting fluctuations caused by noise and improve robustness in low-SNR spectra, they may also reduce sensitivity to genuine physiological changes in spectral parameters arising from processes such as motion or B_0 drift. Therefore, the use of a global high-SNR reference represents a trade-off between fitting stability and the ability to capture true temporal variations in spectral line shape.

4.2.7 Chapter Objectives

The objectives of this chapter were: (1) To develop a 'dynamic fit' algorithm that leverages peak-shape information from high-SNR summed spectra to constrain time-resolved fitting; (2) To validate this method against the standard 'simple fit' across a range of SNR regimes, including a high-SNR rat cardiac dataset, a low-SNR mouse dataset from a high-fat diet (HFD) fed model of obesity, and preliminary human clinical data by comparing theoretical lower bounds of parameter variance for each metabolite; and (3) To quantify the downstream impact of this spectral fitting strategy on kinetic modelling accuracy and its statistical power to differentiate physiological states.

4.3 Methods

4.3.1 Data Acquisition and Study Cohorts

I analysed seven rat hearts, ten mouse hearts, one human brain and one human liver through retrospective spectral fitting. While I did not perform the original data acquisition, these datasets were available through the laboratory database.

Rat Heart: To evaluate the performance of my dynamic spectral fitting algorithm, I processed *in vivo* hyperpolarised ^{13}C MRS data from a pre-established rat model of cardiac hypertrophy. The control rats ($n = 7$) from the abdominal aortic banding (AAB) - Control cohort from the study by Seymour et al. [148] provided the initial case study for comparing the precision and stability of the dynamic fit against the standard simple fit, as these spectra provide relatively high SNR with well-defined metabolic peaks. Rat heart spectra were acquired using a 7.0 T Agilent magnet with a Varian console following an established protocol [148]. Briefly, seven male Sprague Dawley rats were injected with a 1 mL bolus of 80 mM hyperpolarised [1-^{13}C]pyruvate over 10 seconds via a tail vein. ECG-gated pulse-acquire spectra was acquired for 1 minute using a 20 mm butterfly surface coil positioned over the myocardium.

Mouse Heart: To extend the case study, I analysed a lower SNR, unpublished dataset. This data set was a cohort of control ($n = 5$) and high-fat diet (HFD) fed mice ($n = 5$). These animals were maintained on either standard chow or a high-fat diet as carried out by Anderson[149]. Briefly, male C57BL/6JOlaHsd mice were fed either a HFD or a control diet for 16 weeks, commencing at 8 weeks of age. The HFD was a lard-based formulation providing 60% energy from fat, 35% from protein, and 5% from carbohydrates (LBS Serving Biotechnology Ltd, 829197). The control cohort received standard chow (Teklad, Envigo) consisting of 13% fat, 20% protein, and 67% carbohydrates by energy. Both diets were provided ad libitum throughout the duration of the study. Imaging was performed on a 7 T Agilent magnet with a Varian console. Each mouse received a 0.15 mL bolus followed by a 0.05 mL flush (total volume 0.2 mL) of 80 mM hyperpolarised [1-^{13}C]pyruvate via a tail vein

cannulation following the protocol proposed by Dodd et al. [150]. Sixty ECG-gated carbon-13 pulse-acquire cardiac spectra was acquired over approximately 60 seconds using a ^{13}C surface coil (10 mm radius).

Human Brain: As a proof of concept for cross-species and cross-organ validation, I investigated a human brain multi-slice MRS dataset. The healthy human brain data was acquired with slice selective spectroscopy (FOV = 240 mm × 240 mm, flip angle (FA) = 10°, repetition time (TR) = 4 s, receiver bandwidth = 4 kHz, spectral points = 1024, slice thickness = 15 mm) on a 3T GE Premier scanner (GE HealthCare, WI, USA). A 5 mL per second injection of approximately 250 mM hyperpolarised [1-^{13}C] pyruvate at 0.4 mL per kg was delivered as described in Rider et al.[46].

Human Liver: To test the performance of my dynamic fit at very low SNR regimes, I implemented this workflow on a healthy human liver MRS dataset. The liver data was acquired from a healthy volunteer with slice selective spectroscopy (FOV = 400 mm × 400 mm, flip angle (FA) = 90°, repetition time (TR) = 1 s, receiver bandwidth = 3.5 kHz, spectral points = 1024, slice thickness = 30 mm) on a 3T GE Premier scanner (GE HealthCare, WI, USA). The injection protocol is identical to that of the brain data[46].

These data sets helped assess the versatility of the dynamic fit across different physiological environments and signal characteristics.

4.3.2 Spectral Fitting and Kinetic Modelling Implementation

I performed all spectral fitting, post-processing, and visualisations in MATLAB (R2025a). While the core fitting logic was provided by the open-source Oxford Spectroscopy Analysis (OXSA) toolbox, I helped develop a custom graphical user interface (GUI) to manage the data processing. This application allowed for the interactive identification of the start time (i.e., the receiver dead time between RF excitation and signal acquisition), zero-order phase, and centre frequency shift in a few clicks by updating the spectral display in real time. This tool was necessary for

the pre-processing workflow as it provided a method for rapid, intuitive parameter adjustment that was not available in the standard toolbox.

4.3.3 Development of an In-house Graphical User Interface

To streamline the analysis of the complex multi-dimensional datasets acquired in this work, I helped develop an in-house MATLAB-based GUI. Access to the laboratory source code repository (Oxford Molecular Imaging GitHub repository [151]) is restricted due to GE Healthcare IP and unpublished collaborative work in the wider code base, however, access to the workflows specific to this thesis can be granted to the examiners upon request for review. While the OXSA framework provides a robust engine for spectral fitting, I found that an interactive platform was necessary to bridge the gap between raw data acquisition and batch quantification. This tool allows for the rapid visualisation of spectra with only a few clicks, enabling me to verify data quality and optimise fitting parameters without writing custom code for every individual dataset. Once the initial parameters are set, the GUI automates the OXSA fitting process across the entire temporal or spatial array, visualising each fit as a figure to ensure consistency.

The sparse and highly dynamic nature of hyperpolarised ^{13}C spectra presents unique challenges that this GUI was specifically designed to address. A primary hurdle in these experiments is the presence of significant zero-order and first-order phase shifts, which are common in rapid spectral-spatial acquisitions or when hardware timing delays occur. The interface includes interactive sliders that allow the user to adjust these phase values and visualise the corrected spectrum in real time. This is particularly helpful for correcting first-order phase shifts introduced by the delay between RF excitation and signal acquisition, which can distort the spectral lineshape and reduce fitting accuracy. Correcting this phase ensures that the spectrum is properly aligned before fitting with AMARES. Furthermore, the GUI is essential for refining the Prior Knowledge (PK) files. It allows for a visual check of the chemical shift accuracy for each peak and adjustments to the frequency shift (F_0) to account for B_0 inhomogeneities. By centring the spectrum and checking

the metabolite locations against the PK file, I could ensure that the AMARES algorithm was initialised correctly.

The GUI implementation follows a direct workflow: raw FIDs are loaded into the app, where interactive sliders are used to adjust the zero-order phase, frequency shifts (F_0) and start-time delays. After selecting the relevant nucleus (^{13}C or ^{129}Xe) and navigating to specific spatial or temporal voxels, I assign a Prior Knowledge (PK) file to define the metabolite constraints. The user can then choose to fit a single summed spectrum or automate the analysis across the entire dynamic array, with all results and fitted spectra automatically exported to structured .mat files for downstream kinetic modelling.

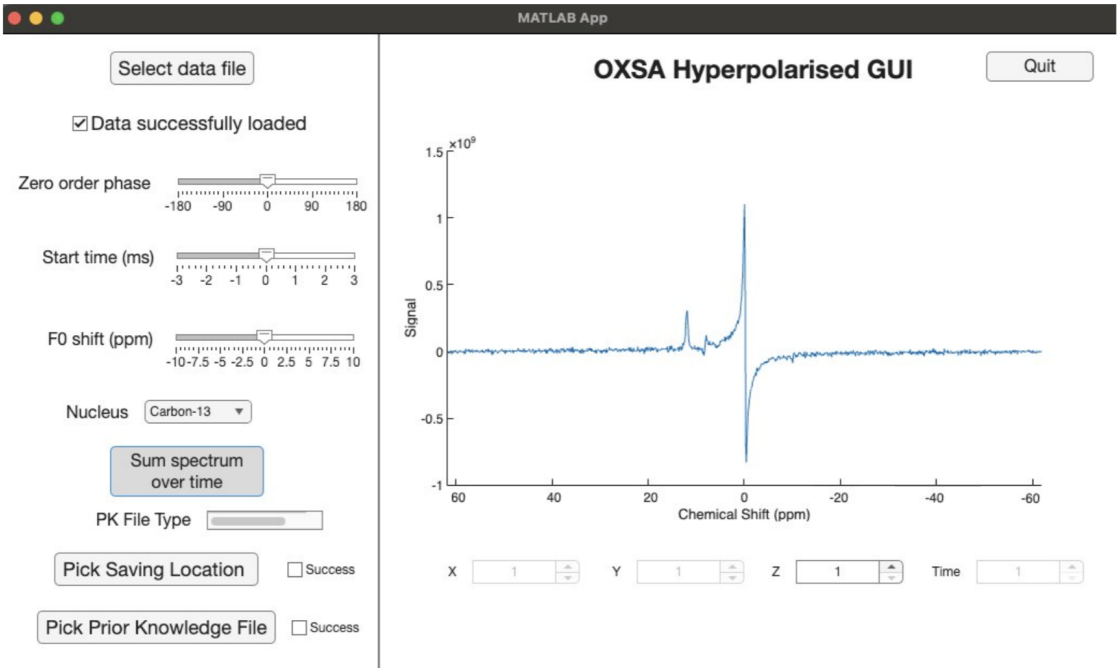


Figure 4.3: MATLAB Graphical User Interface (GUI) developed to interface with the OXSA spectral fitting engine. The left control panel enables interactive pre-processing, including sliders for zero-order phase and start-time (first-order phase) correction. The interface supports the selection of Carbon-13 or Xenon-129 nuclei and allows the user to navigate through 5D datasets (X, Y, Z, and Time) using specialised spinners. A dedicated workflow permits the fitting of either a summed spectrum over the entire time-course or the automated batch fitting of dynamic time-series data using a specified Prior Knowledge (PK) file. Real-time visualisation of the resulting fit and residuals ensures that the AMARES algorithm is correctly initialised, with all quantitative results exported to structured .mat files for downstream kinetic analysis.

For OXSA to accurately fit the FID, the software requires a prior knowledge

file. In my workflow, I created two prior knowledge (PK) files: 1) a baseline PK file that was manually created and was standard protocol in the OXSA-based analysis, and 2) an updated PK file that was algorithmically produced based on high-SNR data. The PK file consisted of bounds and initial values for all peaks: chemical shift, amplitude, linewidth, phase and sigma, as defined in the AMARES framework outlined in subsection 4.2.5. These parameters directly constrain the variables within the AMARES model. The amplitude (a_k) signifies the scaling factor or total area under the peak, which is proportional to metabolite concentration, while the phase (ϕ_k) represents the initial phase offset at $t = 0$. The linewidth (d_k) defines the Lorentzian damping constant, where the frequency-domain Full Width at Half Maximum (FWHM) is related by d_k/π, representing natural T_2 decay. Sigma (σ_k) represents the Gaussian damping constant, accounting for the quadratic decay component in the time domain caused by inhomogeneous broadening (T_2^*) that rounds the peak centre. Finally, the chemical shift (f_k) denotes the resonance frequency relative to the transmitter frequency, allowing for the identification of specific molecular species based on their ppm values.

My post-processing workflow began with visualising the summed data through the custom GUI to identify the zero-order phase, start time and any centre frequency offset, which I kept constant throughout the dataset. Once I identified the correct values, the summed spectra showed multiple distinct peaks correctly phased. I then identified the chemical shift location in ppm for all ^{13}C metabolites (pyruvate, pyruvate-hydrate, lactate, alanine and bicarbonate). In some cases, specific metabolites were absent, such as bicarbonate in the HFD mice or alanine in the human brain. For the HFD mice, I selected a representative control mouse with a bicarbonate peak to provide an approximate location. For the human brain, I omitted the quantification of alanine based on literature indicating this metabolite is not typically detected in healthy brain tissue in similar MRS studies [113].

4.3.4 Peak Fit Model Comparisons

To characterise the optimal fitting strategy for the large pyruvate peak, I implemented a two-pass subtraction pipeline designed to isolate the complex pyruvate lineshape from downstream metabolic signals on the rat data. In the first pass, the pyruvate peak was evaluated against five candidate models, ranging from a standard single Voigt to a dual-component "Two-Voigt" approach. After subtracting the primary pyruvate fit, the remaining metabolites were fitted to the residual using phase-locked single Voigt peaks.

The five candidate models evaluated within the OXSA prior knowledge framework included 1) a standard single voigt (SingleVoigt) laboratory template ($LW_{init} = 20, \sigma = 3$), 2) a generic Dual-Voigt model (VoigtGeneric) consisting of two identical Voigt functions and 3) a Dual-Lorentzian model (TwoLorentz) utilising a summation of two Lorentzian functions. Additionally, I tested a 4) Lorentzian + Voigt model (LorentzGauss) and 5) a hybrid Dual-Component model (VoigtLG), which was specifically designed to capture asymmetric broadening using one high-amplitude Lorentzian-like peak ($LW_{init} = 20, \sigma = 1$) paired with a lower-amplitude, Gaussian-broadened base ($LW_{init} = 10, \sigma = 30$).

Preliminary optimisation of the pyruvate lineshape was performed across all cohorts to ensure the fitting accurately captured the metabolic peaks. Based on the minimisation of the median residual norm, a hybrid Dual-Component model (VoigtLG) was selected for the rat and mouse heart cohorts (Figures 4.4 and 4.5). For the human brain, 9 slices were independently fitted for the pyruvate model selection as shown in figure 4.6, where the generic two Voigt model is omitted because the relative residual norms in all slices scored higher than RN > 25. In the brain, the lowest residual norm was obtained with the single-Voigt model. In the liver dataset, the highest SNR slice was chosen for pyruvate model selection. The residual norm of the single-Voigt model was 0.03, the TwoLorentz model was 0.04, the VoigtLG was 0.04, the LorentzGauss was 0.08, and the VoigtGeneric was 14.85. For the liver data, a laboratory standard single-Voigt model was therefore used for the subsequent analysis. The specific model parameters (linewidth and

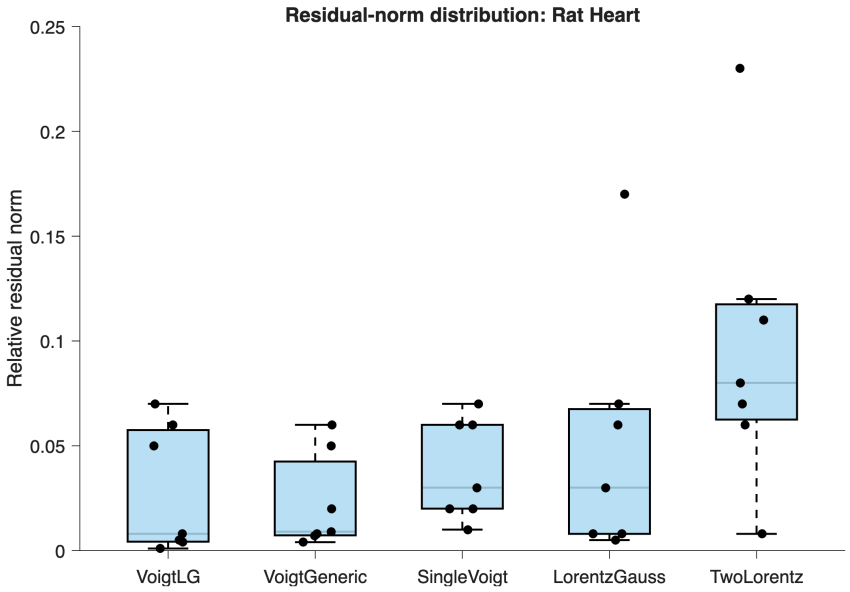


Figure 4.4: Residual norms distribution for the pyruvate model selection of the rat heart cohort. Comparison of fitting residuals for five lineshape configurations. The Two-Voigt model with Lorentzian and Gaussian broadening initialisation was identified as the best-performing model (left to right, lowest to highest median residual norm).

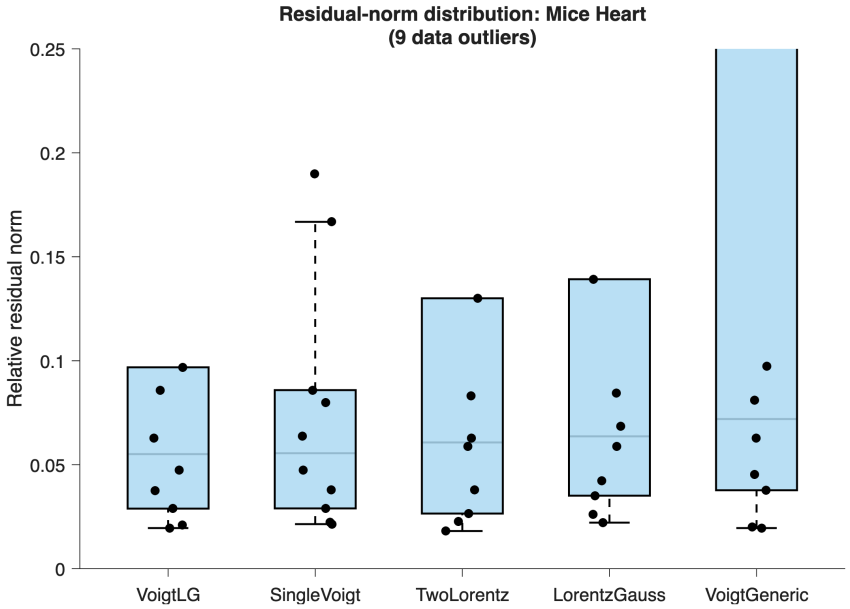


Figure 4.5: Residual norms distribution for the pyruvate model selection of the mice heart cohort. From left to right is the lowest to highest median residual norm.

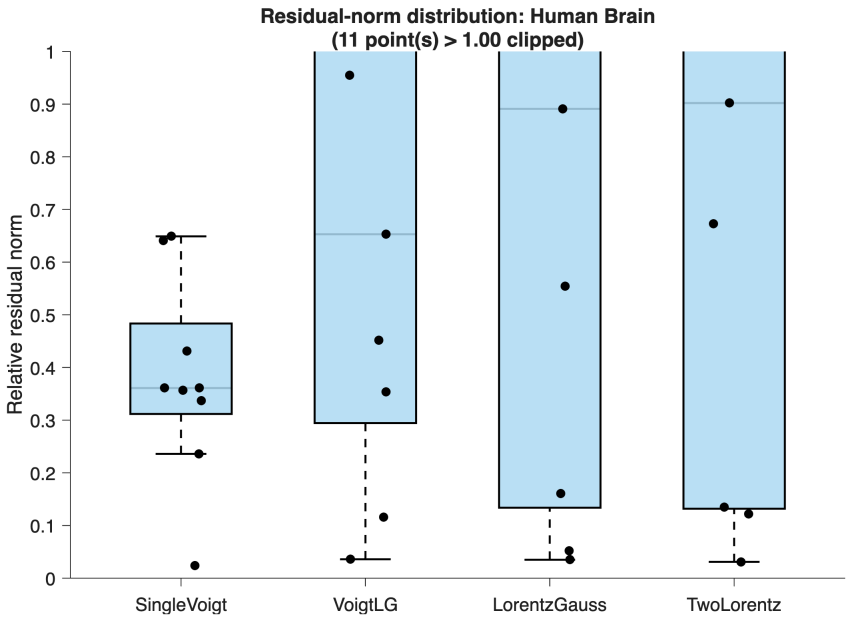


Figure 4.6: Residual norms distribution for the pyruvate model selection of the mice heart cohort. From left to right is the lowest to highest median residual norm. Points above 1.0 were omitted as outliers.

sigma) were defined within the OXSA prior knowledge framework and the optimal configuration for each cohort was selected based on the minimisation of the residual norm. I conducted this selection process prior to the dynamic quantification described in the following sections.

After defining the pyruvate peak model, I ran a pre-process summed spectra fit on a representative spectra with all non-pyruvate metabolites set to values consistent with the laboratory template (initial linewidth=20 with bounds = [5, inf] , sigma=3 with bounds = [0,50]). The non-pyruvate metabolite linewidths were averaged, rounded to the nearest 10, and assigned in the baseline PK file (rat data: linewidth = 30 and sigma = 3, mouse data: linewidth = 40 and sigma = 3, and human data: linewidth = 10 and sigma =3). I adjusted the linewidth bounds to [5, inf] for the rat heart, consistent with the standardised processing templates used within the laboratory for OXSA-based analysis, [2, 80] for the mouse heart, following the established constraints used by Schroeder et al. in previous cardiac studies [152], and [5, 80] for the human brain to adapt the mouse data linewidth to the inter-species differences and noise levels. For all datasets, I implemented sigma bounds of [0, 50],

following the laboratory's standard processing OXSA template. Amplitude bounds were set from zero to the maximum observed signal height amplitude in the summed spectra during the initial GUI visualisation, prior to the fitting routine. I locked all metabolite phases to that of the first pyruvate peak. I defined this set of bounds and initial values as the baseline PK file for my workflow. At this stage, the baseline method (the 'simple fit') and my proposed novel method (the 'dynamic fit') diverge.

Simple Fit (Baseline Method)

In the simple method, I applied the baseline prior knowledge described above to fit every time point independently. I report the metabolite amplitude extracted from the OXSA fit as the signal value in the results section. In the rat and mouse data, the two pyruvate peaks were summed to obtain the total pyruvate amplitude. I refer to the resulting metabolite time courses as the 'simple fit' data.

Dynamic Fit (Novel Method)

In the dynamic method, I iterate on the baseline PK file with high-SNR spectral information from the summed data as a reference. Assuming that metabolite peaks retain consistent characteristics (centre frequency and phase) while amplitudes decay, I leveraged the peak shape information from the high-SNR summed spectra to establish well-informed and very restricted boundaries for the low-SNR time points in the dynamic data. This approach adds meaningful information to the fit, as noise lacks a consistent peak shape and can be filtered by enforcing tight boundaries based on expected peak behaviour.

By fitting the high-SNR summed data, I obtained the metabolite centre frequencies, maximum amplitudes, phases, linewidths, and sigmas. I used these values to algorithmically generate a PK file, using the fit values from the summed data as initial values. I enforced much stricter boundaries than the baseline method, as I assumed large variations from these shapes represented noise. I set chemical shift boundaries to ± 0.5 ppm around the initial value. The amplitude was set between zero and the maximum metabolite peak amplitude as the upper boundary. I bounded the linewidth and sigma to ± 0.1 s^{-1} around their initial values. The

phase was locked to the same global value across all metabolites, with a ±10 degree boundary to account for potential phase drifts due to physical or physiological factors. In the human brain case, I created a slice-wise PK files based on the summed spectrum from the corresponding slice. In the human liver case, I located the highest SNR slice and created an updated PK file data to be used for all slices.

To prevent the algorithm from fitting the noise floor when a metabolite was absent from the summed spectrum, I implemented a validation step. If a linewidth of the peak (excluding pyruvate) deviated by more than three times the average of the other metabolites, the prior knowledge file was overwritten with the average linewidth and sigma values of the other peaks. This assumes that the peak shapes for alanine, lactate, pyruvate-hydrate, and bicarbonate should be approximately similar. I used this updated PK file to fit every time point. I report the extracted amplitude as the signal for each metabolite, with the pyruvate signal representing the sum of the two peaks in the rat and mouse data. These time courses are presented as the 'dynamic fit' data.

Fitting Error Metrics

I quantified the precision of the spectral fits using the CRLB, as implemented in the OXSA toolbox [141]. The CRLB provides a theoretical lower bound on the variance of any unbiased estimator; in the context of HP MRI, it is typically expressed as a percentage of the estimated amplitude.

$$\sigma_m \geq \sqrt{(F^{-1})_{mm}} \tag{4.3}$$

This metric is derived from the Fisher Information Matrix (F) and depends on the noise variance, which is estimated from the tail of the Free Induction Decay (FID), and the sensitivity of the model to parameter changes [142, 153]. A low CRLB indicates that the metabolite signal was sufficiently high relative to the noise to allow for confident quantification. This statistical metric is essential for filtering out unreliable voxels in dynamic datasets where the signal eventually

decays below the noise floor. All uncertainties are calculated in line with the principles of error propagation [154].

The OXSA fitting process yielded time-resolved CRLB values, which I visualised as shaded regions in the amplitude-versus-time plots. To compare the two fitting strategies, I calculated the difference between the mean CRLB value per subject for the simple and dynamic fits; then normalised to the simple fit mean and averaged to produce the values I report in Table 4.1. In the case of the human brain data, I performed this calculation on a slice-by-slice basis. Once the time-resolved metabolic amplitudes were established, I used these data as the primary inputs for kinetic modelling to determine the metabolic exchange rates *in vivo*.

Kinetic Model

I adopted the kinetic modelling established by Schroeder et al. [152], which was adapted to include a delay term between pyruvate bolus arrival and metabolic conversion, which is frequently seen in cardiac metabolic flux, but has also worked on the brain and liver data. I implemented the model as a two-stage fit using non-linear least squares regression (via the `lsqcurvefit` function in MATLAB) to estimate conversion rates while accounting for bolus arrival and decay.

In the first stage, I modelled the pyruvate magnetisation, $M_P(t)$, using a piece-wise function to characterise the injection phase and subsequent metabolic decay. The fitted parameters included the pyruvate decay rate constant (k_{pyr}), the injection rate ($Rate_{inj}$), the bolus arrival time (T_{arr}), and the time at which injection ends (T_{end}). The residual magnetisation at the end of the injection was $M_{P,end}$. The input function is defined as:

$$M_P(t) = \begin{cases} 0 & t < T_{arr} \\ \frac{Rate_{inj}}{k_{pyr}}(1 - e^{-k_{pyr}(t-T_{arr})}) & T_{arr} \leq t \leq T_{end} \\ M_{P,end}e^{-k_{pyr}(t-T_{end})} & t > T_{end} \end{cases} \tag{4.4}$$

Following the characterisation of the pyruvate input, I fitted the downstream metabolites (lactate, alanine, and bicarbonate) in the second stage to estimate specific conversion rates ($k_{pyr\rightarrow x}$). I omitted pyruvate-hydrate as it has no known

conversion enzyme and is considered enzymatically inactive [152]. This stage incorporates a temporal delay (T_{delay}) to account for the transit time between pyruvate arrival and metabolic conversion. This defines an effective time $t_{eff} = t - T_{delay}$.

The downstream magnetisation $M_x(t)$ is described by the conversion rate from pyruvate ($k_{pyr\to x}$), the metabolite-specific decay rate (k_x), and the residual magnetisation ($M_{x,end}$). During bolus arrival, the magnetisation is determined by the balance between the constant injection rate ($Rate_{inj}$) and the simultaneous decay of pyruvate (k_{pyr}) and the metabolite (k_x). After the injection ends, the magnetisation is described by the decay of the remaining pyruvate pool, where $M_{P,end}$ is the pyruvate magnetisation at T_{end}. For continuity, the value at the boundary ($M_{x,cont}$) is evaluated at $t_{eff} = T_{end}$:

$$M_x(t) = \begin{cases} 0 & t_{eff} < T_{arr} \\ \frac{k_{pyr\to x}\cdot Rate_{inj}}{k_{pyr}-k_x}\left[\frac{1-e^{-k_x(t_{eff}-T_{arr})}}{k_x} - \frac{1-e^{-k_{pyr}(t_{eff}-T_{arr})}}{k_{pyr}}\right] & T_{arr} \le t_{eff} < T_{end} \\ \frac{M_{P,end}\cdot k_{pyr\to x}}{k_{pyr}-k_x}\left(e^{-k_x(t_{eff}-T_{end})} - e^{-k_{pyr}(t_{eff}-T_{end})}\right) & \\ +M_{x,cont}e^{-k_x(t_{eff}-T_{end})} & t_{eff} \ge T_{end} \end{cases} \tag{4.5}$$

I performed parameter estimation using `lsqcurvefit` in MATLAB, with the pyruvate parameters (k_{pyr}, $Rate_{inj}$, T_{arr}, T_{end}, $M_{P,end}$) fixed as inputs for each metabolite fit.

I applied the model separately to the amplitudes derived from the simple and dynamic spectral fits. I assessed the goodness-of-fit via the Root Mean Square Error (RMSE). Because the original data is identical between the simple and dynamic fits, I was able to directly compare the RMSE values. I calculated the RMSE as:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{4.6}$$

I performed all calculations in MATLAB. I omitted data points prior to the pyruvate start time from the mean calculation. To produce the kinetic fit improvement values in Table 4.2, I subtracted the dynamic fit RMSE from the simple fit RMSE, normalised this to the simple fit value, and averaged this

percentage across the cohort (or slices, in the human data sets). Based on the prior knowledge, metabolites that do not exist in the summed spectra were omitted from kinetic analysis.

I calculated the mean and standard error (SEM = $\frac{\sigma}{\sqrt{n}}$) for all kinetic parameters across the two mouse cohorts (control and HFD). I then implemented a student's t-test of unpaired groups via the `ttest2` function in MATLAB. To calculate the sample size based on a specific kinetic parameter as an endpoint, I took the mean and standard error of the control and the mean of the HFD to determine the minimum sample size for adequate study power, where the probability of type I error was set to 0.05 and the probability of type II error was set to 0.95 (power = 95%). The equation for the sample size calculation is below based on reference [155]:

$$n_1 = \frac{(\sigma_1^2 + \sigma_2^2/\kappa)(z_{1-\alpha/2} + z_{1-\beta})^2}{\Delta^2}$$

where the corresponding sample size for the second group is defined by $n_2 = \kappa \cdot n_1$. The parameters used in this calculation are defined as: $\Delta = |\mu_2 - \mu_1|$, the absolute difference between the two cohort means. σ_1 and σ_2 are the variance for the mean of group 1 and the mean of group 2, respectively. n_1, n_2 is the required sample size for group 1 and group 2. α is the probability of a Type I error (set at 0.05) and β is the probability of a Type II error (set at 0.95). z is the critical Z-value for the given significance level (α) or power ($1 - \beta$). κ is the ratio of the sample size of group 2 to group 1.

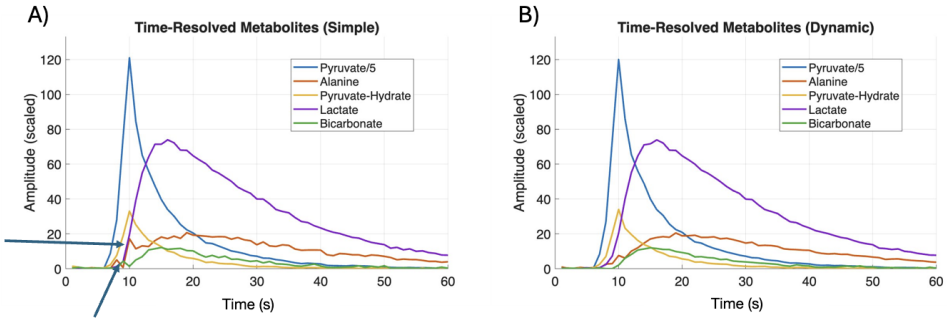


Figure 4.7: Metabolite time-course signals for rat heart AAB02 using (A) simple fit and (B) dynamic fit methods.

4.4 Results

Table 4.1: Relative CRLB reduction (%), reported as mean ± SD, across rat heart, mouse heart, human brain, and human liver datasets. Percentages represent the reduction in mean CRLB achieved using dynamic fitting relative to simple fitting, normalised to the simple mean CRLB.

Relative CRLB Reduction	Pyruvate	PH	Lactate	Alanine	Bicarbonate
Rat hearts ($N = 7$)	$67\% \pm 42\%$	$15\% \pm 19\%$	$6\% \pm 6\%$	$15\% \pm 7\%$	$6\% \pm 7\%$
Mouse hearts ($N = 10$)	$100\% \pm 0\%$	$57\% \pm 4\%$	$37\% \pm 8\%$	$43\% \pm 14\%$	$42\% \pm 25\%$
Human brain (slice 9)	$68\% \pm 22\%$	$78\% \pm 13\%$	$51\% \pm 16\%$	N/A	$66\% \pm 14\%$
Human liver (slice 2)	$44\% \pm 6\%$	$20\% \pm 3\%$	$48\% \pm 19\%$	$20\% \pm 17\%$	$70\% \pm 10\%$

Table 4.2: Improvement in kinetic modelling, reported as the mean percentage decrease in RMSE (± SD) when comparing the dynamic fit with the simple fit across rat heart, mouse heart, human brain, and human liver datasets.

Kinetic Fit Improvement	Pyruvate	Lactate	Alanine	Bicarbonate
Rat hearts ($N = 7$)	$3\% \pm 7\%$	$21\% \pm 13\%$	$36\% \pm 13\%$	$32\% \pm 14\%$
Mouse hearts ($N = 10$)	$23\% \pm 18\%$	$32\% \pm 24\%$	$53\% \pm 18\%$	$67\% \pm 12\%$
Human brain (Slice 9)	$2\% \pm 12\%$	$41\% \pm 11\%$	N/A	$71\% \pm 11\%$
Human liver (Slice 2)	$29\% \pm 25\%$	$31\% \pm 14\%$	$30\% \pm 11\%$	$76\% \pm 2\%$

4.4.1 Rat Heart Study

Across all seven rats, the dynamic fit reduced the mean CRLB for all metabolites compared to the simple fit. The relative mean CRLB reductions were: pyruvate (67% ± 42%), pyruvate-hydrate (15% ± 19%), lactate (6% ± 6%), alanine (15% ± 7%) and bicarbonate (6% ± 7%) (Table 4.1). This improvement is visualised as qualitative process improvements in figures 4.7- 4.10.

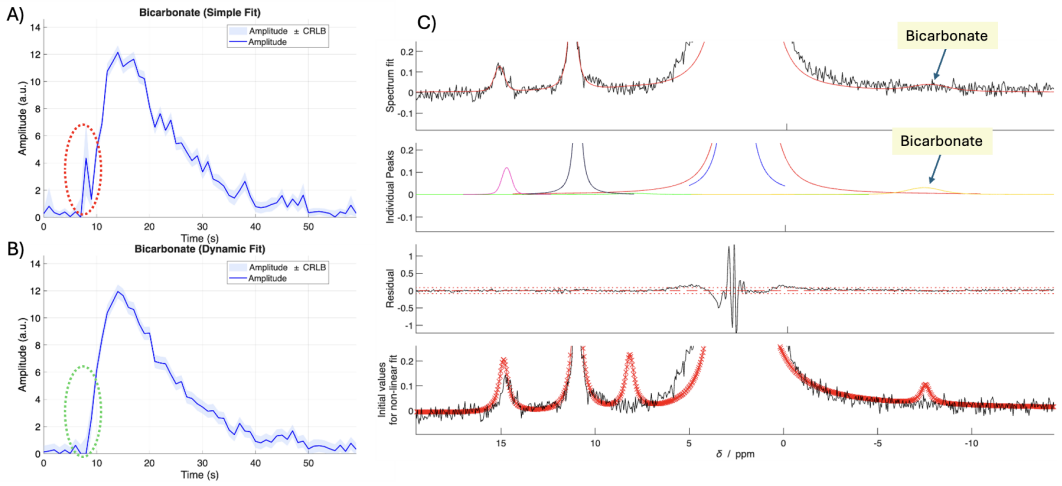


Figure 4.8: Bicarbonate time-resolved signal (AAB02) demonstrating an unphysical local maximum at around 9 s in the simple fit (A). The spectral fit at this time point (C) shows signal assignment to the pyruvate wings and noise. This artifact is resolved in the dynamic fit (B, red dotted circle).

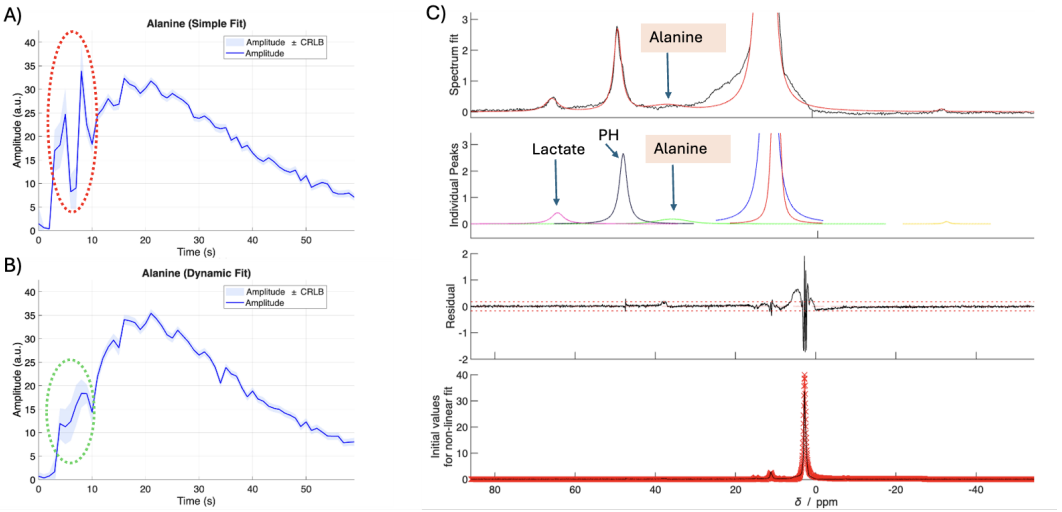


Figure 4.9: Alanine time-resolved signal (AAB08) showing unphysical local maxima in the simple fit (A). Spectral analysis at 8 s (C) indicates alanine signal assignment to noise and the pyruvate wings. These artifacts are corrected in the dynamic fit (B).

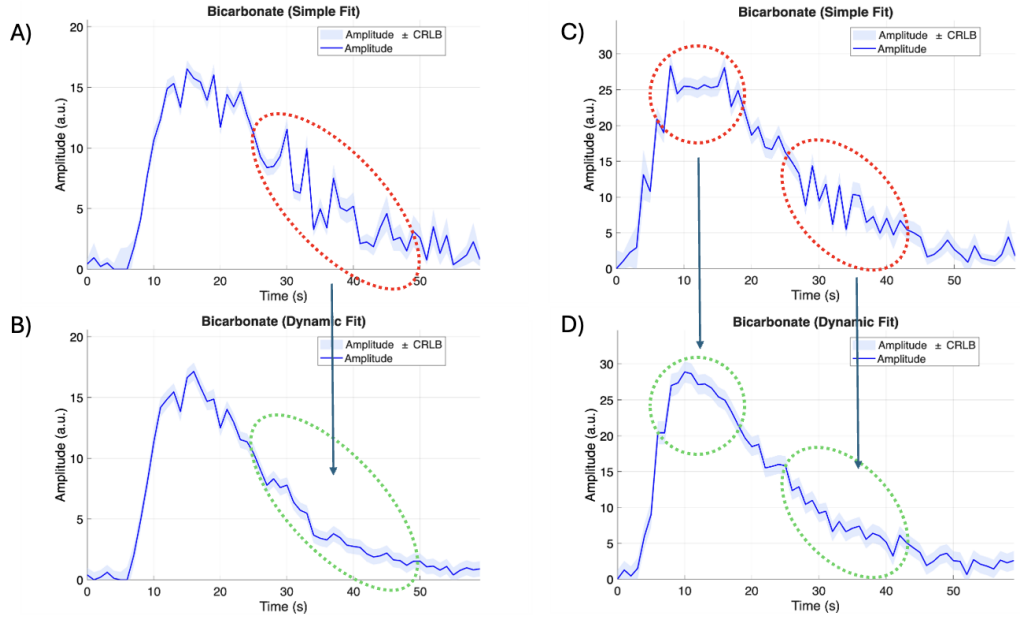


Figure 4.10: Comparison of bicarbonate signal stability in two representative rat hearts (AAB04: A,B; AAB05: C,D). dynamic fitting (B,D) reduces temporal noise compared to simple fitting (A,C), more closely reflecting expected metabolic kinetics.

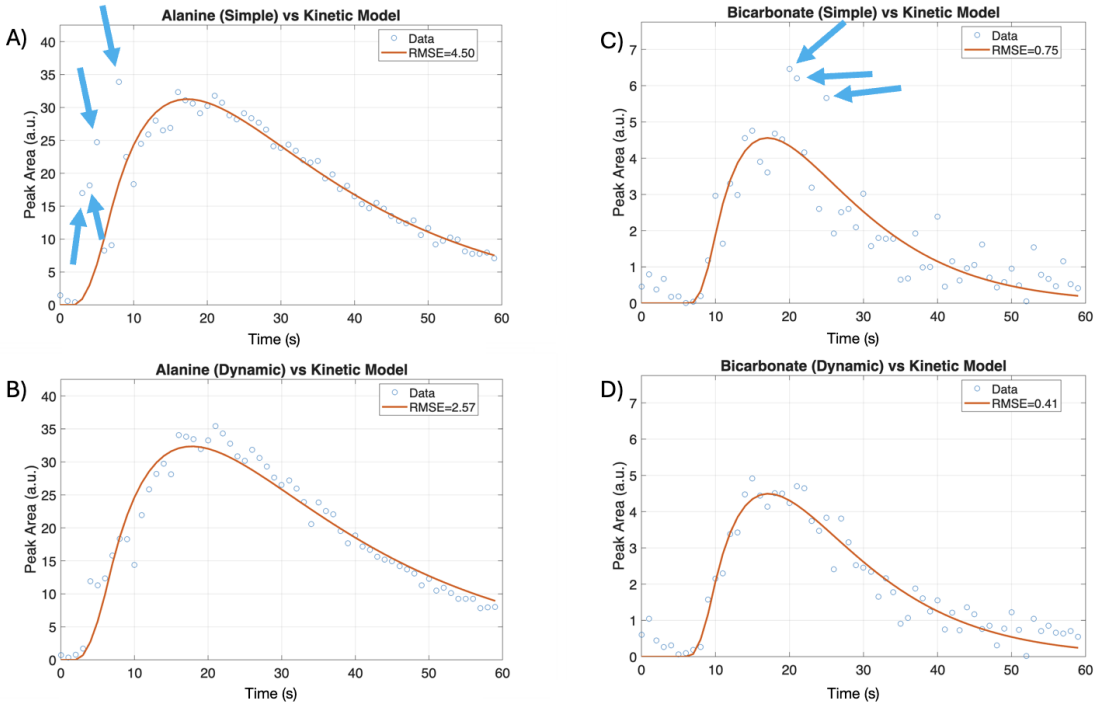


Figure 4.11: Kinetic model fits for alanine (AAB08; A,B) and bicarbonate (AAB03; C,D). dynamic fitting (B,D) results in higher agreement with the kinetic model compared to simple fitting (A,C).

In figure 4.7, the time courses of the metabolites based on a simple fit (A) and dynamic fit (B) are shown where the blue arrows point out elevated signals from alanine and bicarbonate at the same time as pyruvate reaches its maximum amplitude (t=9,10). Physiologically, the signals from downstream metabolites lag behind the pyruvate appearance because of the need for pyruvate uptake and for the enzymatic systems, Lactate Dehydrogenase (LDH) for lactate, Alanine Transaminase (ALT) for alanine and Pyruvate Dehydrogenase (PDH) for bicarbonate, to process the influx of substrate [156]. This means that the simultaneous pyruvate, alanine, and bicarbonate signals are all over fitting to the same large pyruvate signal. The dynamic fit corrects for this, eliminating the non-physiological local maxima corresponding to the pyruvate peak amplitude. A closer inspection of bicarbonate simple fit data (figure 4.8A) at the pyruvate maximum time point reveals that the simple fit erroneously assigns signal from the pyruvate wings and noise to the bicarbonate peak (Figure 4.8C). This produces an artificial local maximum in the bicarbonate time course at 9 seconds, which is eliminated by the dynamic fit (figure 4.8B). I have included a second example with alanine in figure 4.9. In figure 4.9A, the simple fit shows an initial local maxima around 8 seconds, and figure 4.9C is the simple fit at this local max time point. The Spectrum fit shown at the top of Figure 4.9C is fitting the alanine peak to the pyruvate wings and noise. This misfitting is reduced in figure 4.9B by implementing the dynamic fit process. The dynamic fit further demonstrated improved temporal stability in later acquisition times of low signal metabolites like bicarbonate, where the dynamic fit suppressed the signal fluctuations of the simple fit, which better reflects the smooth metabolic kinetic models shown in two representative rat hearts in figure 4.10.

This process improvement in the fitting of the raw data in turn increased the kinetic modelling accuracy by eliminating many outliers and sharpening the metabolic decay curve. In quantitative terms, the kinetic models (simple vs dynamic fit) saw RMSE decreases in the dynamic fits. Dynamic fitting decreased the mean RMSE for pyruvate (3% $\pm$ 7%), alanine (36% $\pm$ 13%), lactate (21% $\pm$ 13%) and bicarbonate (32% $\pm$ 14%) (table 4.2). To illustrate this kinetic modelling

improvement, two different metabolites from two representative rats are shown in figure4.11. In figure4.11A and B, the alanine RMSE is reduced from 4.50 to 2.50, a 44% reduction in RMSE relative to the simple fit. The outliers to the kinetic model in the simple fit are highlighted by blue arrows. These are data points that are erroneously fitting to the pyruvate peak in the earlier time points, and are eliminated in the dynamic fit. The bicarbonate example in figure 4.11C and D also shows a RMSE decrease from 0.75 to 0.42, a 43% improvement. This example shows the reduction of outliers and the tightening of data points along the kinetic model curve, which both contribute to the lower RMSE value in the dynamic fit. These results demonstrate that improved spectral fitting directly translates to higher fit stability in the metabolic quantification and more robust kinetic modelling.

4.4.2 Mouse Heart Study

Compared to the simple fitting approach, the dynamic fit yielded a substantial reduction in the mean CRLBs across all metabolites in the mouse heart study. Specifically, relative CRLB reductions were observed for pyruvate-hydrate ($57\% \pm 4\%$), lactate ($37\% \pm 8\%$), alanine ($43\% \pm 14\%$), and bicarbonate ($42\% \pm 25\%$) (Table 4.1). While the mean CRLB improvement for pyruvate is reported as $100\% \pm 0\%$, this figure represents a significant six-fold decrease in absolute CRLB values when transitioning from the simple to the dynamic fit.

Figures 4.12 and 4.13 compare the spectral fitting performance where the dynamic fit corrects the pyruvate wings and noise effects seen in the simple fit. Figure 4.12 presents metabolite time courses from a representative high fat diet mouse. Lactate, alanine and bicarbonate all show artificial peaks around 49 seconds in the simple fit (figure 4.12A,C,E), which is corrected in the dynamic fit as highlighted with the green circles (figure 4.12B,D,F). The results are similar in the control subject in figure 4.13. The simple fit produces signal maxima that are not physiologically possible (figure 4.13A,C,E), which are resolved by the dynamic fit procedure (figure 4.13B,D,F). Dynamic fitting prevents the assignment of pyruvate peak-wing signal or noise to downstream metabolites, which occurs in the simple fit

due to low SNR. Compared to the simple fit, the dynamic fit data is smoother across time, which improves the metabolite dynamics quantification for the next step.

In all mouse heart data, the improved metabolite quantification translated to lower RMSE values in the kinetic modelling for pyruvate (23% ± 18%), alanine (53% ± 18%), lactate (32% ± 24%) and bicarbonate (67% ± 12%) (table 4.2). In figure 4.14 A and B, the RMSE is reduced from 0.13 to 0.07, a 46% difference to the lactate kinetic model. The outliers to the kinetic model in the simple fit (figure 4.14 A) are pointed out by blue arrows. These are data erroneously fitting to the pyruvate peak. In the last third of the acquisition, the simple fit produced scattered noise around the baseline (figure 4.14 A). Both the erroneous peaks and the noise fluctuations are eliminated in the dynamic fit (figure 4.14 B). The alanine example (figure 4.14 C,D) illustrates the RMSE reduction from 0.16 to 0.06, a 63% improvement. Again, the data points are much tighter to the kinetic model in the dynamic fit. The (figure 4.14 E,F) bicarbonate example shows the RMSE reduction from 0.15 to 0.06, a 60% difference and similarly illustrates a better fit to the kinetic model.

Table 4.3: Population mean±SEM values for the kinetic parameters estimated using the simple and dynamic fitting methods.

Parameter	**Control**	**HFD**
Simple fit		
$k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ ($p = 0.21$)	0.026 ± 0.019	0.0008 ± 0.0001
$k_{\mathrm{x\,lac}}$ ($p = \mathbf{0.03}$)	0.0711 ± 0.0040	0.0579 ± 0.0030
$k_{\mathrm{x\,bic}}$ ($p = 0.28$)	1.1328 ± 0.9747	0.0047 ± 0.0015
Dynamic fit		
$k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ ($p = \mathbf{0.03}$)	0.0014 ± 0.0003	0.0005 ± 0.00001
$k_{\mathrm{x\,lac}}$ ($p = \mathbf{0.04}$)	0.0710 ± 0.0045	0.0595 ± 0.0019
$k_{\mathrm{x\,bic}}$ ($p = \mathbf{0.03}$)	0.0362 ± 0.0112	0.0070 ± 0.0014

The dynamic fit decreased the p-values between the control and HFD populations for two of the kinetic parameters. Figure 4.15 summarises the population differences for each parameter, where Figure 4.15A shows the simple fit and Figure 4.15B shows the dynamic fit. Of the metabolic parameters, only $k_{\mathrm{x\,lac}}$ showed a significant difference between the control and HFD groups using the simple fit, whereas $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$, $k_{\mathrm{x\,lac}}$, and $k_{\mathrm{x\,bic}}$ were all significantly different using the dynamic fit. The population

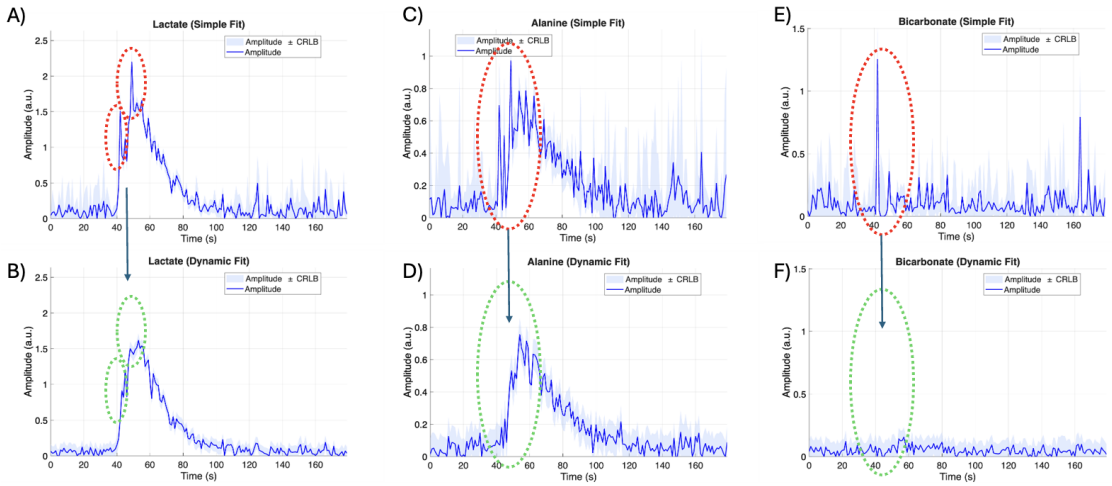


Figure 4.12: Comparison of simple and dynamic fit time courses for lactate (A,B), alanine (C,D), and bicarbonate (E,F) in a high-fat diet (HFD09) mouse. Erroneous peaks resulting from pyruvate wings overlap and noise are circled in red (A,C,E); the resolution of these unphysical maxima via dynamic fitting is circled in blue (B,D,F).

mean±SEM values and corresponding statistical comparisons are summarised in Table 4.3. Notably, dynamic fitting identified significant differences in both $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ and $k_{\mathrm{x\,bic}}$, which were not detected using the simple fitting approach.

Consequently, the increased precision of the dynamic model led to a substantial reduction in the sample size required to reach statistical significance. For a study where $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ is the primary endpoint ($\alpha = 0.05$, power = 95%), the required cohort size decreased from 28 subjects using a simple fit to just 6 subjects with the dynamic model. Likewise, for $k_{\mathrm{x\,bic}}$, the required sample size was reduced from 38 to 8 subjects. The dynamic fitting provides a more robust framework for detecting metabolic changes in the HFD populations with substantially fewer subjects.

4.4.3 Human Brain

I further evaluated the performance of my dynamic fitting algorithm using a multi-slice human brain dataset from a healthy volunteer. As shown in the representative summed spectrum (Figure 4.16), I identified distinct peaks for pyruvate, pyruvate-hydrate, lactate and bicarbonate; I did not detect alanine, which is consistent with reported metabolic profiles for healthy brain tissue in hyperpolarised studies. By implementing the dynamic fit, I achieved a substantial reduction in the mean relative

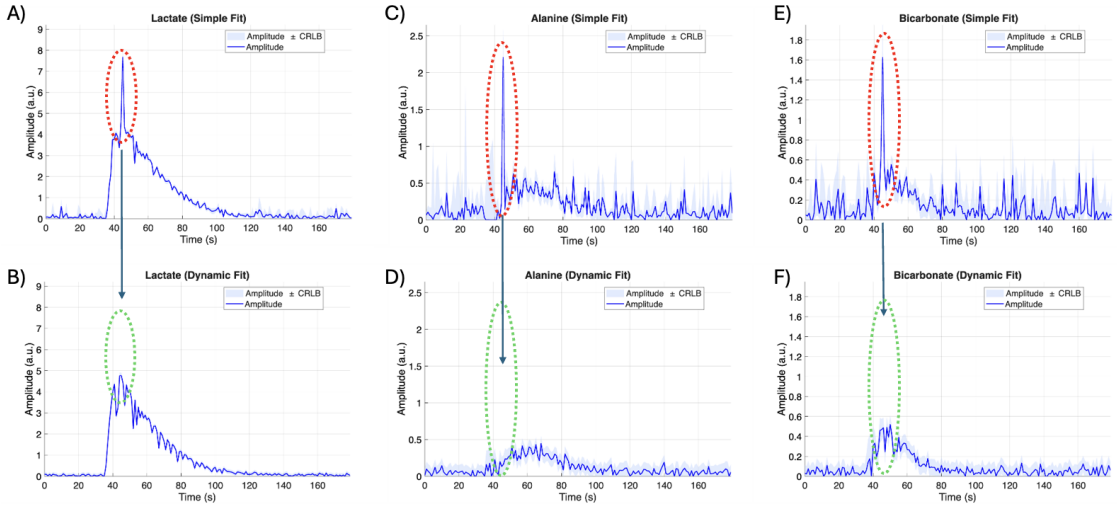


Figure 4.13: Comparison of simple and dynamic fit time courses for lactate (A,B), alanine (C,D), and bicarbonate (E,F) in a control mouse (HFD02). Artificial local maxima in the simple fit are circled in red (A,C,E) and corrected in the dynamic fit (B,D,F).

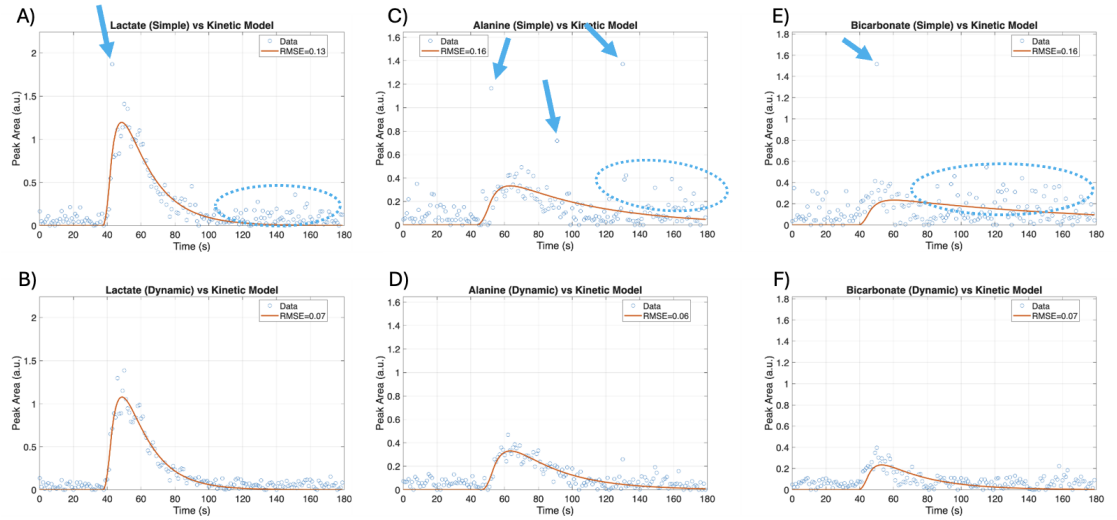


Figure 4.14: Kinetic modelling results for lactate from HFD07 (A,B), alanine from HFD06 (C,D), and bicarbonate from HFD01 (E,F) using simple and dynamic spectral fits. Blue arrows and circles indicate outliers in the simple fit that increase RMSE. These are eliminated in the dynamic fit, resulting in higher agreement with the kinetic model.

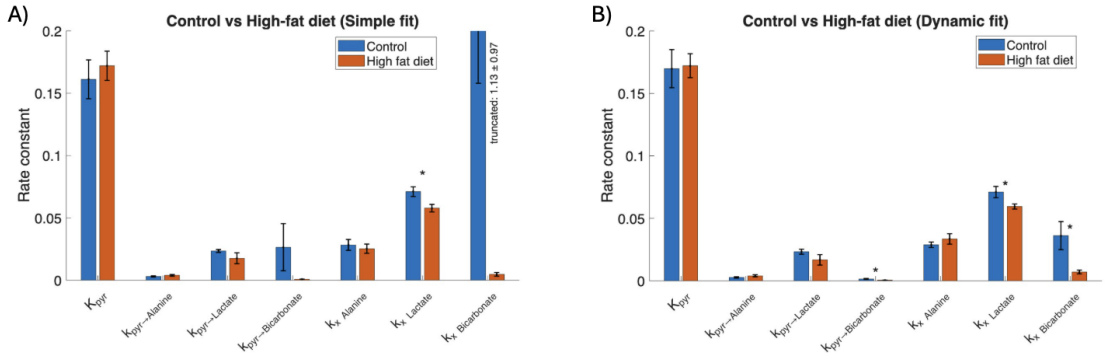


Figure 4.15: Comparison of kinetic parameters between control and high-fat diet populations. Results are shown for (A) the simple fit and (B) the dynamic fit methods. Asterisks denote statistical significance via t-test ($p < 0.05$).

CRLB across all observed metabolites compared to the simple fit: pyruvate (68% ± 22%), pyruvate-hydrate (78% ± 13%), lactate (51% ± 16%) and bicarbonate (66% ± 14%) (Table 4.1).

The qualitative improvements in the metabolite time courses are illustrated in Figure 4.17. In the simple fit (Figure 4.17A), the bicarbonate signal shows a local maxima simultaneous with the pyruvate bolus maximum ($t = 5$). At this point, the simple fit erroneously assigned signal from the pyruvate wings to the bicarbonate frequency. Through the dynamic fit (Figure 4.17B), the bicarbonate signal peaked at approximately $t = 8$, which more accurately reflects the temporal delay required for pyruvate uptake and metabolic conversion. Additionally, I found that the dynamic fit effectively filtered the noise fluctuations that were prevalent in the later timepoints of the acquisition in the simple fit.

I compared specific spectral fitting performance for pyruvate-hydrate, lactate and bicarbonate in Figure 4.18. The red indicators in the simple fit (Figures 4.18A, C, E) highlight the erroneous signal assignments during the initial bolus arrival (Figure 4.18E) and the increased variance during the decay phase (Figures 4.18C,E). The dynamic fit reduced fluctuations (Figures 4.18B, D, F, highlighted in green), where the use of strictly constrained peak parameters ensured the fitting remained robust even as the SNR declined.

Finally, I assessed the impact of these fitting improvements on the downstream kinetic modelling across the brain slices. The dynamic fit reduced the mean RMSE

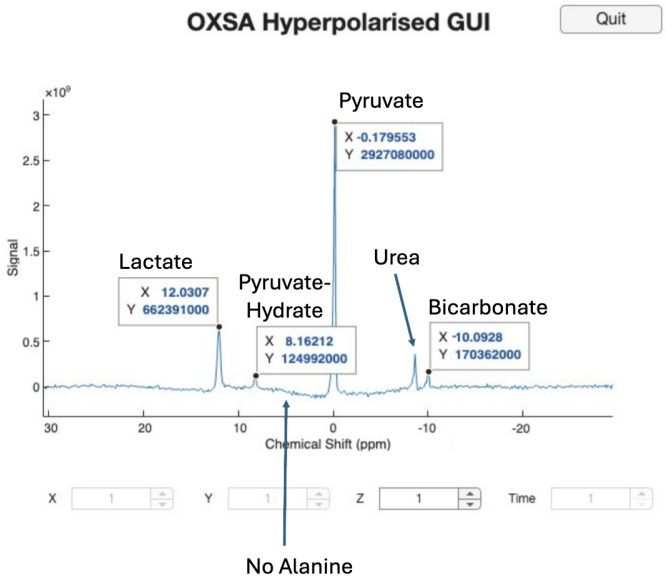


Figure 4.16: Representative summed slice showing pyruvate, lactate, pyruvate-hydrate, and bicarbonate, but no alanine. The urea peak is from a reference vial attached to the head coil.

for pyruvate (2% ± 12%), lactate (41% ± 11%) and bicarbonate (71% ± 11%) (Table 4.2). Figure 4.19 presents representative kinetic fits for bicarbonate from three different slices. In the simple fit (Figures 4.19A, C, E), blue markers identify outliers, particularly during the initial time points when the pyruvate bolus arrives, that deviate from the kinetic model and inflate the error metrics. In particular, figure 4.19A peaks simultaneously with pyruvate and the bicarbonate kinetic model mimics that of pyruvate. The kinetic model completely changes to a lower amplitude with a smooth curve delaying from the bolus, which is more realistic, in figure 4.19B. This increased stability is critical for the reliable estimation of conversion rates in future clinical translation.

4.4.4 Human Liver

Finally, I applied the dynamic fitting method to human liver data acquired from a healthy volunteer. Three slices were acquired, but the first slice was rejected from analysis due to low signal in the summed prior knowledge spectra as the FOV captures only a partial liver. Similar to the previous datasets, the dynamic

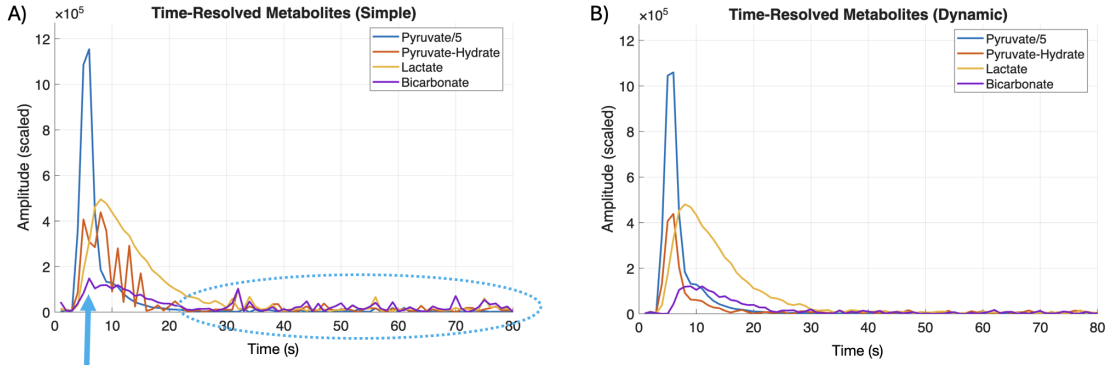


Figure 4.17: Metabolite time-course signals for brain slice 4 using (A) simple fit and (B) dynamic fit methods.

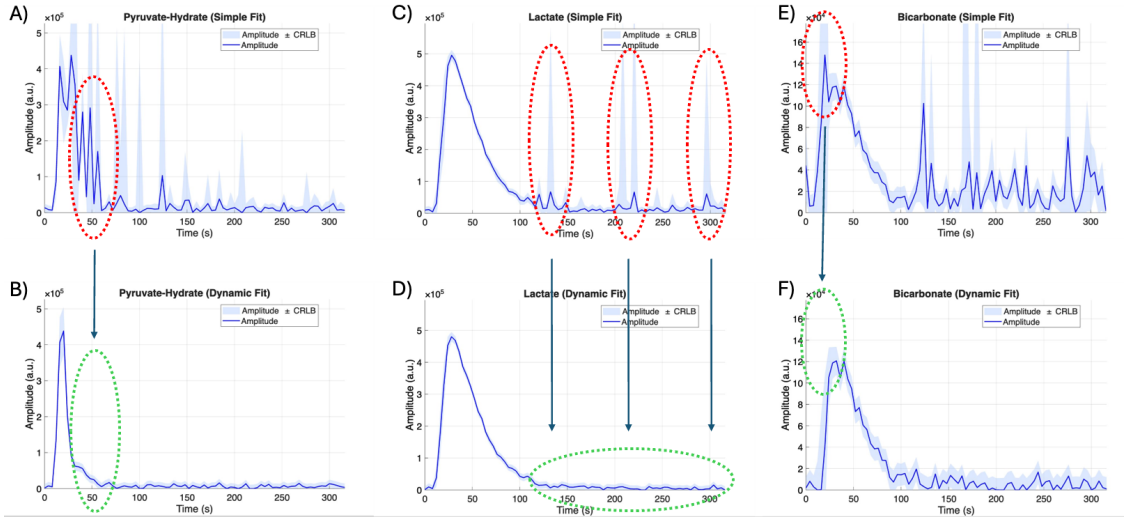


Figure 4.18: Comparison of simple and dynamic fit time courses for pyruvate-hydrate (A,B), lactate (C,D), and bicarbonate (E,F) in the healthy human brain (slice 4). Erroneous peaks in the simple fit due to pyruvate wings overlap and noise are circled in red (A,C,E); the corrections via dynamic fitting is circled in green (B,D,F).

fit consistently improved the precision of all the metabolic estimates. I observed mean relative CRLB reductions for pyruvate (44% ± 6%), pyruvate-hydrate (20% ± 3%), lactate (48% ± 19%), alanine (20% ± 17%) and bicarbonate (70% ± 10%) compared to the simple fit (Table 4.1).

Figure 4.20 shows the summed spectra and resulting metabolite time courses for the two liver slices. In slice 1, I detected signals for pyruvate, pyruvate-hydrate, lactate, alanine and bicarbonate but did not detect a bicarbonate peak in slice 2 (Figure 4.20A,D). I found that the simple fit produced substantial fluctuations that were inconsistent with the summed data in figure 4.20A,B. For example in

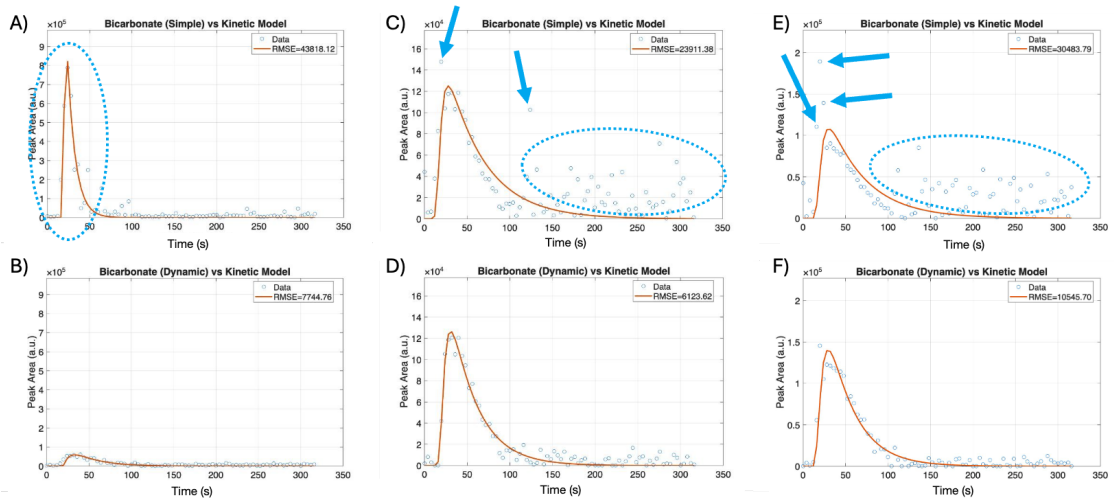


Figure 4.19: Kinetic modelling results for bicarbonate from slice 1 (A,B), from slice 4 (C,D), and from slice 5 (E,F) using simple and dynamic spectral fits. Blue indicators highlight outliers in the simple fit that increase RMSE (A,C,E). The dynamic fit (B,D,F), corrects for the outliers, and produces a more stable time course that aligns closely with the kinetic model.

figure 4.20B, the bicarbonate reaches higher values across t=[21 25], than those in the pyruvate-hydrate, lactate and alanine time courses, which is improbable for a low signal downstream metabolite. The dynamic fit in figure 4.20C addresses these discrepancies, yielding more physiologically plausible and smoother curves in all metabolites. In slice 2, the time courses from the simple fit are very noisy (Figure 4.20E), and it is hard to see the kinetics. Figure 4.20F substantially improves the signal definition. Notably, the bicarbonate is correctly reduced to the noise floor figure 4.20F.

To understand the alanine and bicarbonate behaviour in figure 4.20C, I further examined the performance of the dynamic fit during the peak of the pyruvate bolus (Figure 4.21). At time points $t = 22$–24, I was able to clearly resolve the alanine and bicarbonate peaks from the raw data, as shown in the individual spectral fits (Figures 4.21B–D). Figures 4.21B–D shows the heavy influence of the pyruvate wings on the alanine and bicarbonate peaks. The alanine peak is visible, but the peak amplitude is artificially enhanced by the pyruvate wings distorting the baseline.

In figure 4.22, the time-resolved signals for lactate, alanine and bicarbonate in slice 1 showed several erroneous peaks and local spikes in the simple fit results

(indicated by red circles in figure 4.22A,C,E). The dynamic fit eliminated these artifacts caused by noise and overlap with the pyruvate wings (figure 4.22B,D,F). This correction was particularly important for alanine and bicarbonate, which exhibit lower SNR and are sensitive to fitting errors during the high-intensity pyruvate bolus.

Finally, I evaluated the impact of the dynamic fit on the kinetic modelling of the liver metabolites. I achieved mean RMSE reductions for pyruvate (29% ± 25%), lactate (31% ± 14%), alanine (30% ± 11%), and bicarbonate (76% ± 2%) (Table 4.2). As illustrated in figure 4.23A,C,E, the simple fit resulted in several outliers (blue markers) that deviated from the expected kinetic decay, particularly where the pyruvate bolus arrives. The dynamic fit corrected these points, resulting in a significantly tighter alignment with the kinetic model (figure 4.23B,D,F). As shown in Figure 4.23B and F, the maximum signal intensities observed in the time course more closely reflect the low amplitudes expected in the liver. This is consistent with prior knowledge derived from the summed spectra, which identifies alanine and bicarbonate as the two lowest-intensity metabolites. These results demonstrate that the dynamic fitting approach provides a robust framework for quantifying liver metabolism where signal instability is a known challenge.

Overall, the transition from simple to dynamic spectral fitting consistently improved the reliability of metabolic quantification across rat, mouse and human cohorts. By suppressing non-physical signal assignments caused by pyruvate wings overlap and low SNR, the dynamic fit reduced relative CRLB percentages (ranging up to 70% reduction in low signal metabolites, figure 4.1) and lowered kinetic model RMSE (up to 76% improvement in figure 4.2). This increased precision directly enhanced the statistical power of the mouse heart study, enabling the detection of significant physiological differences in $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ and $k_{\mathrm{x\,bic}}$ between control and high-fat diet populations that were otherwise obscured by fitting noise. The dynamic fit produced significant p-values between the mouse populations for $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ ($p = 0.03$) and $k_{\mathrm{x\,bic}}$ ($p = 0.03$), which were non-significant in the simple fit ($p = 0.28$ and $p = 0.22$, respectively) (figure 4.15). These results underscore the advantages of

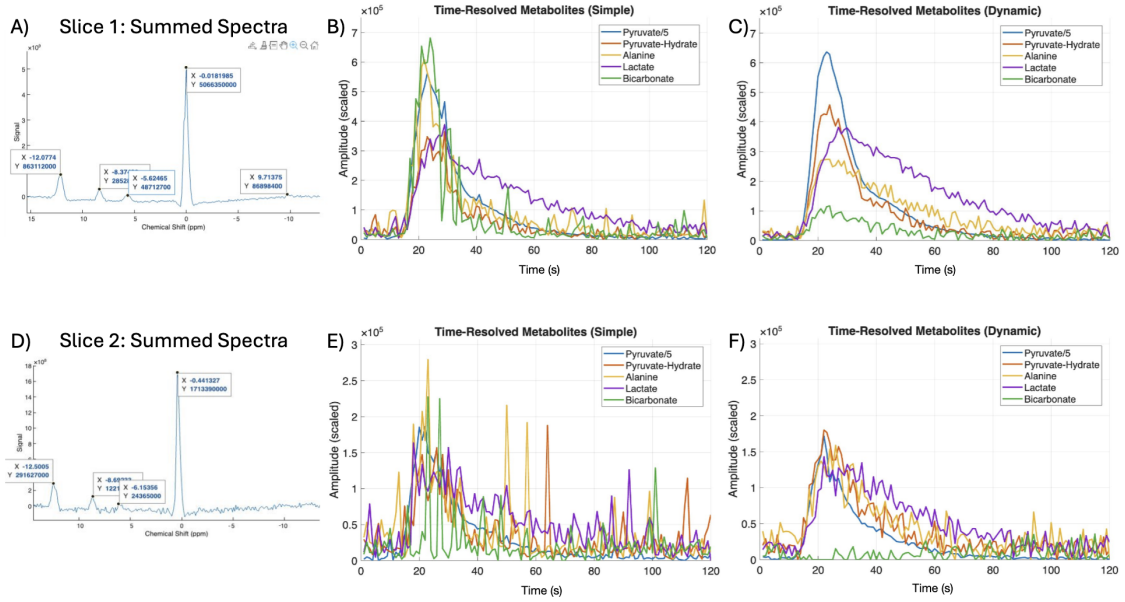


Figure 4.20: Comparison of summed spectra and metabolite time courses for two human liver slices. The top row (A-C) shows the results for slice 1, including the summed spectra (A) and a comparison between the simple (B) and dynamic (C) fit methods. The bottom row (D-F) illustrates the same comparison for slice 2.

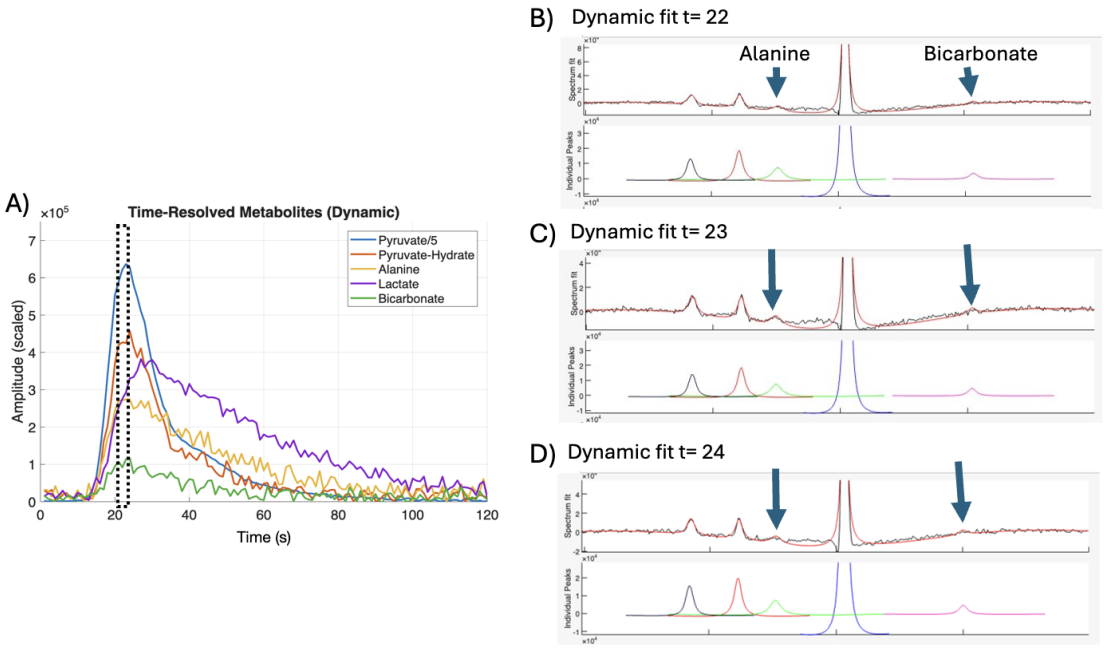


Figure 4.21: Dynamic fit results for individual time frames for liver slice 1. (A) The dynamic fit metabolite time course, with the black dotted rectangle highlighting the time points of the pyruvate bolus peak (t=22-24). (B-D) Individual spectral fits at these time points show the fitted alanine and bicarbonate peaks (arrows) using the OXSA dynamic fit. The red lines represent the OXSA fit, the black lines represent the raw data, and the individual fitted metabolic peaks are plotted below.

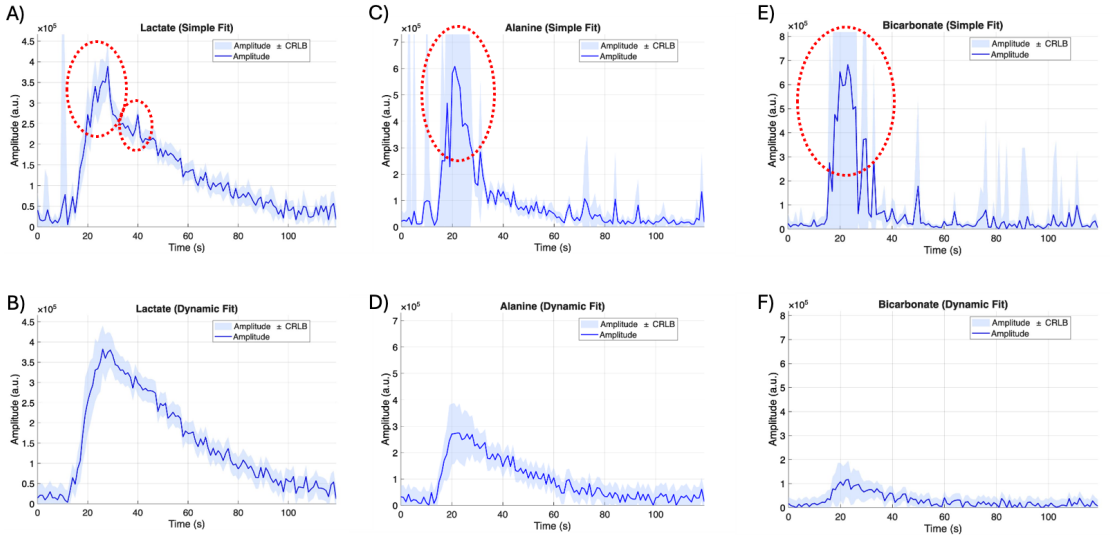


Figure 4.22: Comparison of simple and dynamic fit time courses for lactate (A,B), alanine (C,D), and bicarbonate (E,F) in the healthy human liver (slice 1). Erroneous peaks in the simple fit due to pyruvate wings overlap and noise are circled in red (A,C,E).

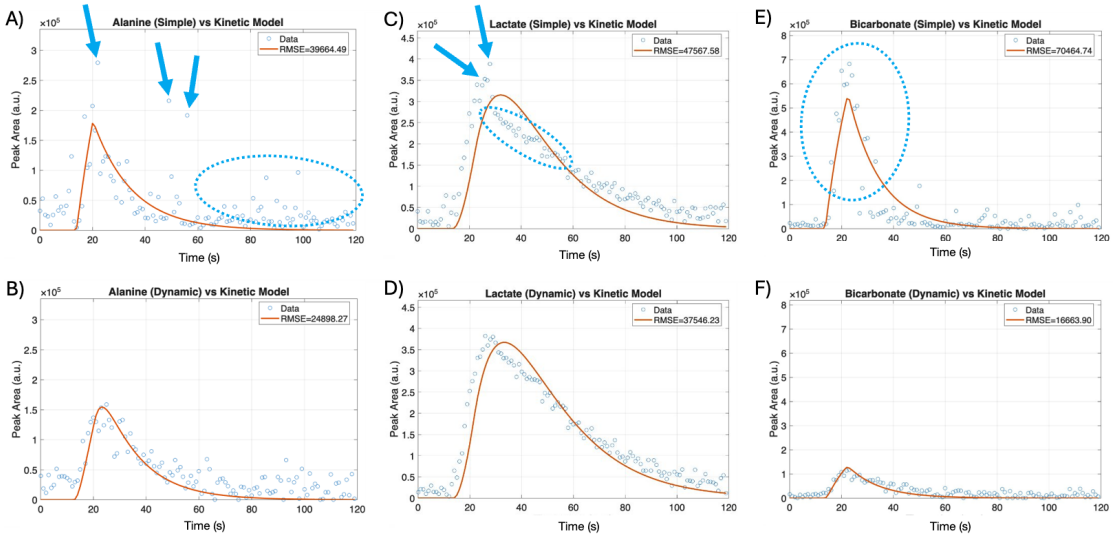


Figure 4.23: Kinetic modelling results for alanine from slice 2 (A,B), lactate from slice 1 (C,D), and bicarbonate from slice 1 (E,F) using simple and dynamic spectral fits. Blue indicators highlight outliers in the simple fit that increase RMSE (A,C,E). The dynamic fit (B,D,F), corrects for the outliers, eliminates pyruvate wings effects, and produces a more stable time course that aligns closely with the kinetic model.

incorporating well-informed constraints taken from high SNR data into the spectral fitting process for robust hyperpolarised ^{13}C kinetic modelling.

4.5 Discussion

The dynamic fitting method effectively decouples metabolite signals from noise by enforcing consistent peak shapes derived from a high-SNR summed prior knowledge reference. Standard OXSA fitting allows linewidth and frequency parameters within a wide boundary to vary freely at each time point. In the low-SNR regimes typical of hyperpolarised ^{13}C experiments, this freedom often leads the algorithm to erroneously assign noise or pyruvate wings overlap to downstream metabolites. I found that by constraining these parameters, I could prevent the unphysical local maxima often observed during the initial bolus arrival, which otherwise skew the arrival slope and peak amplitude of the kinetic model.

In downstream metabolites, the substantial decrease in the relative CRLB reduction reaching up to 78% demonstrated that the dynamic fit provided a "pre-denoised" input for the kinetic fitting process, which showed mean RSME improvements between 21% and 76%. By narrowing the confidence intervals (CRLB) and removing spectral artifacts, I ensured that the estimated rate constants reflect metabolic conversion rather than fitting instability. All datasets (rat, mouse and human) reflected this CRLB improvement through implementing the dynamic fit, but in the liver dataset, most strikingly illustrated in figure 4.20, the simple fit produced a bicarbonate curve that mirrored the pyruvate bolus, whereas the dynamic fit revealed a smoother metabolic decay curve in slice 1 and a flat zero amplitude in slice 2, which were both consistent with expected physiology.

A key outcome of this work is the decrease in sample size I achieved by implementing the dynamic fit in the mouse heart study. I found that $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ and $k_{\mathrm{x\,bic}}$ were significantly different between control and high-fat diet populations when using the dynamic fit ($p<0.05$). In the simple fit, these physiological differences were obscured by fitting noise. Consequently, the sample size calculation for these specific kinetic parameter endpoints pointed to a decrease in the minimum number

of subjects for adequate study power (α = 0.05,β=0.05) from 28 to 6 mice for $k_{\mathrm{pyr}\rightarrow\mathrm{bic}}$ and 38 to 8 mice for $k_{\mathrm{x\,bic}}$. This substantial reduction in animal numbers directly aligns with the 3Rs (Replacement, Reduction, and Refinement) of scientific animal research and contributes to the University of Oxford's formal commitment "to promoting knowledge of the moral and legal requirements for animal research; and to promoting a culture of care in all aspects of research" [157]. Furthermore, at the University, scan time costs approximately £600 per 3 mice. In decreased scan time alone, this improvement could save thousands of pounds, excluding cost savings in fewer injections, consumables, personnel time, husbandry, and veterinary oversight. The decrease in sample size thereby converts to significantly more cost-effective experiments and fewer animals sacrificed.

Beyond the financial and ethical improvements, these parameters could offer further physiological insight. While enzymatic conversion rates (e.g., $k_{\mathrm{pyr}\rightarrow\mathrm{lac}}$) are often the focus of metabolic imaging, the decay rates (k_{x}), which represent the clearing of the metabolite from the system, also provide critical metabolic information. These decay rates are influenced by T_1 relaxation and metabolic flux, which reflect tissue characteristics, micro-environmental factors and intracellular transport. However, decay rates remain understudied. The ability to reliably quantify these parameters opens new avenues for investigating additional molecular environment factors, potentially serving as a new class of metabolic biomarkers.

Despite these improvements, there is much more scope to overcome the limitations to this approach, especially the sensitivity to the quality of the reference: the summed data. If the summed spectrum itself is corrupted by significant motion artifacts, B_0 drift over time, or poor shimming, these errors will be propagated as fixed constraints (phase, frequency shift, and linewidth) into every individual time point. However, in my datasets the high SNR of the summed data generally mitigates this.

Furthermore, the complex, non-Lorentzian peak shape of pyruvate remains a challenge. A two-Voigt model was required for the rat and mouse heart data, whereas a single Voigt model was sufficient for the human brain and liver. This is consistent

with the known differences in pyruvate signal composition across organs. In the heart, the large intravascular blood volume and rapid bolus transit through the cardiac chambers produce a composite pyruvate signal from multiple compartments with distinct linewidths, as described in the Dynamic Peak Morphology subsection 4.2.4. In contrast, in the brain and liver, the vascular contribution to the voxel signal is comparatively smaller, and the pyruvate lineshape is better approximated by a single Voigt function. Whilst this interpretation is consistent with the literature on cardiac partial volume effects [46], further investigation is needed to confirm these effects alongside the potential of complex peak shape as a result of B_0 inhomogeneity, flow effects or the specific RF pulse profile. It is also important to mention that the dynamic fit method improves the stability of the total pyruvate signal, but it cannot separate intra-luminal blood and intracellular tissue contributions. Consequently, the calculated rate constants must be viewed as apparent rates representing a mixture of perfusion, membrane transport and enzymatic conversion.

I also assume peak shape constancy over time. Factors such as changes in local pH (affecting chemical shift) or oxygenation levels (affecting T_2^* and linewidth) could cause subtle variations in peak morphology over the course of the experiment. By enforcing strict ± 0.1 boundaries on linewidth and ± 0.5 ppm on frequency, the dynamic fit may force the model to ignore genuine physiological fluctuations in favour of stability.

My dynamic fitting framework also assumes that peak shape remains constant over time. Factors such as changes in local pH, which may affect chemical shift, or oxygenation, which may influence T_2^* and linewidth, could produce subtle variations in peak morphology during the experiment. By enforcing strict boundaries of ± 0.1 on linewidth and ± 0.5 ppm on frequency, the dynamic fit may favour stability at the expense of genuine physiological variation. Future work could address this limitation by introducing adaptive constraints or by modelling frequency and linewidth as smoothly varying temporal parameters. A rolling-window approach, in which the prior spectral information for each time point is derived from a summed spectrum containing the preceding (X) time points and the following (X) time points,

rather than from data summed over the entire acquisition, could be an avenue of investigation. This would provide a locally representative a priori spectrum that can adapt to gradual changes in peak position and linewidth while retaining sufficient signal-to-noise ratio for robust parameter estimation. The window width could be selected according to the temporal resolution and signal-to-noise ratio of the dataset, although excessively wide windows may smooth genuine temporal changes, whereas narrow windows may provide less stable prior information.

The kinetic modelling can also be optimised to find the best fit given the metabolic data points. I noted that when I tweak my initial variables, the resulting kinetic parameters changed slightly. This highlights an underlying mathematical problem of local minima in non-linear least-squares optimisation. This becomes a very complex optimisation problem beyond the scope of this thesis, but the stability provided by the dynamic fitting process significantly reduces the RMSE in these estimates compared to standard methods.

4.6 Conclusion and Future Directions

In this chapter, I developed and implemented a dynamic fitting algorithm designed to improve the reliability of hyperpolarised ^{13}C spectral quantification. My primary findings demonstrate that leveraging high-SNR summed spectra to constrain time-resolved fitting consistently reduces parameter variance and improves kinetic modelling accuracy across species and organs. Specifically, I achieved a significant reduction in mean relative CRLB (up to 78%) and a substantial decrease in kinetic model RMSE (up to 76%) for all metabolites. In the mouse heart cohort, I found that this improvement enabled the detection of significant physiological differences in $k_{\mathrm{pyr}\to\mathrm{bic}}$ ($p = 0.03$) and $k_{\mathrm{x\,bic}}$ ($p = 0.03$) between control and high-fat diet populations, differences that were obscured by noise when using the standard simple fit. Such differences meant a 79% decrease in the necessary cohort size, decreasing the financial and ethical burden of future studies.

Implementing the dynamic fit procedure allowed me to resolve several non-physiological signal fluctuations such as downstream metabolites appearing to peak

simultaneously with the pyruvate bolus, which were observed across all datasets. While the fits are not all perfect, particularly in low-signal regions like liver slice 2, the dynamic approach significantly outperformed the simple fit alternatives. In these low-SNR cases, the dynamic fit produced far more plausible kinetic curves, even if they did not perfectly align with the expected physiological time delays for lactate and alanine. Ultimately, these results demonstrate that the dynamic fitting routine provides a more robust representation of metabolic arrival and decay; by enforcing temporal consistency, it prevents the algorithm from erroneously assigning the 'pyruvate wings' overlap or baseline noise to low-signal metabolite peaks.

While the dynamic fit effectively addresses several challenges in spectral quantification, it also highlights areas for future development such as peak morphology investigations. My use of a two-Voigt peak model for pyruvate served as a pragmatic solution to handle the complex peak morphology observed *in vivo*; however, the physical origins of this shape, whether due to flow effects, hardware inhomogeneities, or compartmentalisation, remain to be fully characterised. I identified that the pyruvate signal is a composite of intra-luminal blood and intracellular tissue, and future work should focus on whether a more physically grounded model can separate these components to improve the accuracy of the conversion rates.

I also observed that kinetic modelling remains sensitive to initial parameter conditions and local minima. Implementing a more global optimisation strategy, or a Bayesian approach to kinetic fitting, could further enhance the robustness of the calculated rate constants. Finally, the metabolite-specific decay rates (k_x) proved to be a sensitive marker in the HFD mice study. Investigating the physiological drivers of these clearance rates, such as intracellular transport and T_1 relaxation changes, could establish k_x as a valuable new biomarker for disease progression and treatment response.

Whilst significant physiological differences were only detected in the mouse heart cohort, the consistent improvement in spectral quantification across all species and organs in this project gives me confidence that the dynamic fit will generalise to future human studies. The healthy volunteer data in the brain and liver presented

here is limited in that it only demonstrates feasibility of this technique. With more time and resources, larger human cohorts, including subjects with pathology, are the next steps to show the clinical applications for this work.

Overall, the transition from independent frame-by-frame fitting to a constrained dynamic approach represents a step toward robust work flows required in HP carbon-13 experiments. By ensuring that metabolic quantification is driven by physiological signals rather than fitting noise, I have developed a method that significantly enhances the ability to investigate *in vivo* metabolic activity.

5 Conclusion

Contents

5.1 Coil Hardware Design 186
5.2 Accelerating Translational HP MRI 190
5.3 Denoising HP MRI Using Spectral Fitting 193
5.4 The Future of HP MRI 195

In this DPhil, I addressed some of the key translational limitations of HP carbon-13 MRI, which hamper the ability to extract reliable metabolic information within the short, non-renewable signal lifetime, while maintaining spatial and spectral resolution. Across Chapters 2–4, I focused on three components of the HP imaging chain: (i) improving receive sensitivity and coverage through RF coil design (Chapter 2), (ii) reducing acquisition time using a robust, scanner-available spectroscopic imaging sequence accelerated with parallel imaging (Chapter 3), and (iii) improving quantification robustness using spectral fitting methods tailored to low-SNR, dynamic HP data (Chapter 4).

5.1 Coil Hardware Design

In Chapter 2, I developed and characterised a multi-layer receive array aimed at improving sensitivity and coverage for chest HP ^{13}C applications, where signal

at depth and spatial coverage were limiting factors. I systematically evaluated element separation to balance geometric decoupling against field sensitivity and uniformity, using electromagnetic simulations supported by bench top S-parameter measurements. In the two-element multi-layer configuration, the simulations identified a wide decoupling range (approximately 11-16 cm) with an optimum overlap of 14 cm for decoupling. Benchtop measurements of the two-element multi-layer configuration closely followed simulated predictions, exhibiting a similarly broad decoupling range. While simulations suggested an optimal overlap of 14 cm, the experimental optimum was found to be 13 cm. However, when I evaluated receive field behaviour with depth, I found that an 11 cm separation provided the most homogeneous B1- behaviour across depths of interest, indicating that the separation that maximises decoupling is not necessarily the separation that optimises imaging performance in the target region.

I then extended the approach to three- and four-element configurations with the 11cm separation to prioritise homogeneous B1- behaviour and quantified performance within a heart-sized ROI. The three-element configuration provided the highest average B1- magnitude (i.e., highest sensitivity), whereas the four-element configuration provided the lowest coefficient of variation (i.e., best uniformity), establishing a clear design trade off between sensitivity and homogeneity that can be selected according to the clinical question and anatomical target. To pursue the best uniformity, I built and tested the four-element array inside the bore of a 3 T clinical scanner. Through a phase-based Bloch-Siegert B_1^+ mapping protocol, the receive sensitivity was decoupled from transmit inhomogeneities, revealing that the multi-layer array achieved more than twice the mean sensitivity of a commercial rigid chest array within the cardiac field-of-view. Specifically, the array achieved a mean SNR of 14.68 ± 4.79, outperforming the commercial reference (12.96 ± 3.75) and demonstrating superior depth penetration, where mean signal intensity at 12 cm depth exceeded the mean signal performance of the commercial coil at 6 cm.

Reflecting on this project, my steepest learning curve stemmed from pivoting from my background in software engineering to the nuanced, hands-on, and more

unpredictable world of RF hardware. Physically building a delicate and intricate system highlighted the patience and attention to detail required for hardware research. For example, each two-millimetre-wide capacitor must be mounted with the minimal amount of soldering material. Across the seven elements I constructed, including two practice elements, this required the successful placement of 150 capacitors, with every tweak requiring a cycle of desoldering and resoldering.

Transitioning from the bench to in-bore experimentation, I quickly learned that hardware intuition is a hard-won skill, such as finding hairline fractures in the solder joints that led to an open circuits, stopping the scanner from operating any pre-scan operations, effectively halting a planned scan day. These hardware failures were compounded by the vague language of the MRI system's error messages. The scanner errors rarely provided a direct explanation as to why a specific coil was failing or why a certain scan would stop operating. This means a rigorous, manual troubleshooting where systematically verifying every cable, connector, and capacitor became the only path forward. While frustrating, this exercise afforded me a much deeper understanding of the entire signal chain.

I was incredibly fortunate to have the mentorship of Tony Zhou, Ian Nicholson, and Damian Tyler, whose experienced guidance kept the project moving in the right direction. The process of engineering a functional MRI coil provided me with a robust foundation in RF engineering and testing. I believe this skill set is one of the most valuable outcomes of my DPhil. It future-proofs my expertise, providing a physical, engineering-based counterweight to my software background in an era where digital tasks are increasingly replaced by AI. Ultimately, seeing the first successful acquisition from my multilayer array remains the most memorable, and rewarding milestone of my doctoral research.

As I reflect on the time I spent troubleshooting during my simulation phase of my project, I realise that the manual burden of clicking through the clunky CST user interface could have been mitigated through better computational strategies. If I were to approach this project again, my primary objective would be to automate the simulation-tuning loop. While the time spent on hands-on engineering was

an invaluable learning experience, the repetitive nature of setting up simulations, waiting for runs to complete often days at a time, and manually tuning the system only to repeat the entire cycle was a significant time sink.

While I used CST to leverage the inherited domain knowledge of the laboratory, my experience with the software was often frustrating. The software felt opaque and the workflow was excessively manual. The licensing restrictions were a constant bottleneck, preventing me from accessing the university's GPU cluster and forcing me into depending on a local supercomputer that had limited availability, which meant competing for run time.

In hindsight, I would have preferred to invest those months into trialling an alternative CAD platform like Sim4Life (ZMT Zurich MedTech AG). Its native Python API would have allowed me to leverage my software engineering background to replace manual iterations with scripted optimisations. This would have enabled me to automatically and systematically sweep centre-to-centre distances and allowed for curved conformations to the natural and diverse shapes of the human torso, which was too time consuming to manually set up on CST. More importantly, it would have allowed for seamless integration with the University-available GPU resources, bypassing licensing hurdles and eliminating downtime whilst waiting for the local super computer to become available.

Given this compute upgrade, there are multiple pathways I could have directed this hardware research. I could have shifted the focus toward a more curvature-based design, accommodating the diverse BMIs and anatomical variations seen in a clinical population, or more rigorously optimised the four-element geometry through granular decoupling sweeps on a human model. However, moving beyond geometric refinement and into clinical translation necessitates a safety characterisation such as a comprehensive heating profile.

While these hardware refinements represent a clear engineering path forward, they exist within a rapidly evolving context for the field. Throughout the course of this DPhil, my attendance at the International Society for Magnetic Resonance in Medicine (ISMRM) conferences provided a broader perspective on the trajectory

of the field. My personal thoughts, based on observations over these four years, are that the volume of presentations and posters dedicated to fundamental RF coil engineering appears to be decreasing. Within RF engineering, the trend seems to be moving towards extreme channel counts, ultra-high field strengths (7T and above), and complex dual- or triple-tuned innovations. This shift in the field, combined with the increasing prevalence of abstracts on smarter and often AI-driven reconstruction, prompted my own strategic pivot. I realised that while hardware innovation remains the front line of the signal detection chain, the significant gains in clinical utility may now lie in the computational frameworks for acceleration and spectral fitting. Consequently, my projects moved from the physical interface of the coil to the software-driven methodologies that enable high-resolution metabolic mapping.

5.2 Accelerating Translational HP MRI

In Chapter 3, I addressed the temporal inefficiency of conventional two-dimensional free-induction decay chemical shift imaging (2D CSI) by implementing parallel imaging acceleration using SENSitivity Encoding (SENSE), with the specific aim of improving temporal sampling of a robust and widely available spectroscopic sequence. A key limitation for SENSE in carbon-13 is that coil sensitivity maps cannot be reliably acquired at the carbon frequency *in vivo* because of low natural abundance and HP time constraints. I therefore used naturally abundant sodium sensitivity maps as a surrogate for carbon-13 coil sensitivities, acquired with a ready-made dual-tuned receive array, and used these maps for SENSE reconstruction of undersampled carbon-13 CSI. I demonstrated, in both phantom experiments and *in vivo* porcine kidney studies, that sodium-based sensitivity maps enabled reliable unfolding of reduced-FOV aliasing and supported a four-fold temporal acceleration for carbon-13 imaging. In the porcine experiments, this produced repeatable metabolite ratio mapping across subjects, with coefficients of variation for SENSE-reconstructed ratios within 15% across kidney regions of interest. I also showed that the increased temporal resolution produced a more detailed arrival

and decay time course, which is directly relevant for subsequent kinetic modelling and for denoising strategies that harness temporal structure.

Due to the refurbishment of the MRI facilities at the John Radcliffe Hospital in Oxford, this project necessitated an international collaboration with the Christoffer Laustsen group who has a large animal facility at Aarhus University, Denmark. This collaboration served as a significant learning point in research management. Organising large-animal experimental procedures and designing protocols for an external facility required communication and logistics that exceeded previous phantom-based work. I was fortunate to have the meticulous supervision of James Grist, who travelled with me to Denmark and guided the initial MRI sessions in Aarhus. This preparation provided my technical foundations to independently manage these *in vivo* scans as I carried out my porcine experiments on the subsequent trip. On the day of my experiments, I had medics preparing and anaesthetising pigs from 6AM, while technical specialists prepared my pyruvate dissolution from the early hours of the morning. Once I had my scanning protocol ready, I had someone run between the hyperpolariser room and the scanner to inject the £4000 HP [1-^{13}C] pyruvate into my pigs, all while the lab professor and a group of master students stood behind me watching the experiment. Unlike the solitary and iterative nature of building hardware, large-animal experiments were a high-pressure team effort where months of preparation culminated in a 60-second window of signal. Reflecting on this phase of the project, this experimental chapter was the most exhilarating, as it transitioned the research from the laboratory bench to the complex environment of a large-animal facility in Aarhus.

The experiments were the exciting part of this chapter, and the technical execution was more routine but not without its frustrations. A hurdle was the requirement to transition my coding language from Python, which is my preference, to MATLAB. As this part of the research preceded the current era of widespread generative AI coding assistants, this was a process of relearning syntax by manually debugging code and navigating documentation through trial and error. This developmental phase was further supported by the Transfer of Status viva, where

the critical feedback provided by the examiners, Aaron Hess and Ferenc Mozes, was instrumental in refining the experimental objectives. These inputs directly facilitated the peer-reviewed publication of these findings. Navigating the peer review process itself was a good exercise in scientific communication, teaching me to present findings with the logic and conciseness required for high-impact research, a standard I have endeavoured to maintain throughout this thesis.

Over the last four years, the landscape of hyperpolarised ^{13}C MRI has shifted toward multi-centre studies, marked by the release of several key consensus papers such as: "Consensus recommendations for hyperpolarised [1-^{13}C]pyruvate MRI multi-center human studies". I had the privilege of contributing to this dialogue at the Hyperpolarized Magnetic Resonance Technology Center (HMRTC) conference at UCSF in 2024, where I presented my work on accelerated SENSE CSI. My work in SENSE CSI directly addresses these emerging guidelines by providing a practical solution to the "calibration bottleneck" mentioned in the consensus paper. My *in vivo* experiments demonstrated that naturally abundant sodium sensitivity maps can reliably support 4x acceleration, without the need for expensive and time-consuming carbon-13 calibration scans. This approach aligns with the consensus drive for surrogate-based workflows, ensuring that high-resolution metabolic mapping can be performed consistently across different hardware platforms.

Crucially, by implementing these accelerations on a robust and widely available CSI sequence, this work significantly lowers the barrier for entry for other research sites. Rather than requiring specialised pulse sequences that are difficult to port between systems or may be under intellectual property protection, accelerating CSI allows multiple sites to access an identical scanning sequence on standard clinical hardware. This democratisation of high-performance acquisition is essential for the field's move toward large-scale, multi-site clinical trials.

While the primary objective of the SENSE study was to achieve four-fold acceleration, my experiments successfully provided a more detailed view of the metabolic time course than was previously possible. In the porcine studies, the acquisition commenced after the initial bolus arrival, meaning the pyruvate signal

rise was not captured. The absence of this initial kinetic data complicated the model-fitting process and highlighted a critical requirement: high-speed acquisition must be paired with both precise timing and robust analysis to fully characterise metabolic data. This insight led directly to the work presented in Chapter 4, where I developed an iterative spectral fitting and kinetic modelling framework. By bridging the gap between rapid imaging and quantitative analysis, this progression establishes a more comprehensive framework for metabolic characterisation.

5.3 Denoising HP MRI Using Spectral Fitting

In Chapter 4, I developed an analysis framework based on the requirement that robust quantification is essential for clinical utility, particularly when rapid signal decay and low SNR leave downstream metabolites near the noise floor. The motivation for this project is that metabolic signals are frequently obscured by baseline noise and spectral overlap from the dominant pyruvate peak. To address these limitations, I developed a 'dynamic fit' framework that leverages high-SNR information from summed spectra as prior knowledge to constrain the fitting of individual, low-SNR time points. Unlike conventional independent fitting methods that treat each temporal frame in isolation, this approach enforces physical consistency across the time course by using global peak characteristics to stabilise the quantification of transient signals, effectively preventing unphysical signal assignments.

Validated across rat, mouse, and human datasets, the dynamic fit consistently outperformed standard independent fitting, reducing the mean Cramér-Rao Lower Bound (CRLB) by up to 78% and improving kinetic model Root Mean Square Error (RMSE) by up to 76%. Most significantly, in a mouse model of obesity, the increased precision of the dynamic fit revealed significant physiological differences in both the pyruvate-to-bicarbonate conversion rate ($k_{pyr \to bic}$, $p = 0.03$) and the bicarbonate decay rate (k_{x_bic}, $p = 0.03$), which were otherwise obscured by noise in the simple fit ($p = 0.28$ and $p = 0.22$, respectively). This improved sensitivity enabled a 79% reduction in the required sample size, from 28 to 6 subjects, to achieve 95%

statistical power for the $k_{pyr \rightarrow bic}$ endpoint, demonstrating a robust framework for enhancing the statistical rigour and cost-effectiveness of HP ^{13}C clinical studies.

I particularly enjoyed the breadth of data I was able to process through this project. Applying the dynamic fit across three species: rat, mouse, and human, and three distinct organs: the heart, brain, and liver, provided an invaluable exposure to the diversity of *in vivo* ^{13}C spectra. Each organ presented unique challenges, from the pyruvate wings dominating the cardiac spectra to the low-SNR peaks seen in the human brain and liver. Helping to develop a custom Graphical User Interface (GUI) to manage these fits was also a satisfying engineering task. This project was purely computational, allowing me to focus entirely on post-processing development. Coupled with this, the use of AI-assisted coding further accelerated my development.

My solution to spectral fitting was intentionally simple and ran contrary to the current trend of implementing complex machine learning solutions. As a former software engineer at an AI company, I understood the power of automation, but I prioritised a physics-based, transparent model. Unlike supervised learning models that require vast, pre-existing training datasets, the dynamic fit is self-referencing. This ensures that the results are mathematically interpretable, a prerequisite for clinical confidence, and avoids the black box nature of neural networks. While I suspect there are many aspects of innovation in HP MRI that can benefit from ML, I hope future researchers focus on the specific physical constraints of the problem rather than the allure of novel AI architectures.

While this framework was implemented within the MATLAB-based OXSA environment to leverage existing laboratory tools, the underlying logic is platform-agnostic. My transition to AI-assisted coding during this phase highlighted how the next generation of MR tools must be built for interoperability, moving toward the open-source Python ecosystems that currently drive global innovation.

Reflecting on this phase of the research, the most profound impact was the direct link between software precision and the ethical conduct of science. The 79% reduction in the required sample size for the mouse obesity study illustrates how better analysis can have a significant ethical and financial impact. This

work demonstrates that the software-driven developments are as impactful as the hardware improvements toward the many steps needed for successful translation of hyperpolarised ^{13}C MRI.

5.4 The Future of HP MRI

Taken together, the new aspect I bring to the hyperpolarised community is an integrated, practical pathway that connects hardware, image acquisition, and software quantification innovations. While Chapter 2 improved the signal at the source and Chapter 3 accelerated its capture, Chapter 4 ensured that the resulting data could be translated into reliable physiological parameters. On the hardware side, I provide a multi-layer receive array design strategy with experimentally grounded optimisation, and a quantified sensitivity–uniformity trade-off across array sizes that is directly relevant to cardiac HP imaging. On the acquisition side, I provide a technically simple and reproducible method for accelerating a vendor-available CSI sequence using SENSE with surrogate-nucleus sodium sensitivity mapping, which addresses a common barrier to parallel imaging in multi-nuclear HP studies. On the post-processing side, I introduce a spectral fitting framework that harnesses known temporal structures to quantify metabolic kinetics with high statistical power in an animal model. Clinically, the developments in this thesis increase the usable signal that a HP study yields, and therefore provides interpretable metabolic information in the time available, and they reduce dependence on highly specialised, centre-specific sequences and workflows.

Ultimately, the technical developments presented in this thesis form a cohesive imaging chain designed to move hyperpolarised ^{13}C MRI from a specialised research tool toward a robust clinical diagnostic. By addressing the sensitivity bottleneck with multi-layer coil geometries, the acceleration bottleneck with SENSE CSI mapping, and the quantification bottleneck with dynamic spectral fitting, I have demonstrated that the most significant improvements in HP MRI are found at the intersection of physical engineering and computational rigour. This journey has mirrored my own professional evolution, from a software engineer navigating the unpredictable

physics of RF hardware to a researcher prioritising transparent, physics-informed analysis in an era of AI-driven automation. As the field moves toward multi-centre human trials, the success of metabolic imaging will depend on integrated, accessible workflows from hardware through to post-processing. By bridging these domains, this work provides a framework for the next generation of metabolic research, ensuring that the signal of hyperpolarised ^{13}C can be captured, accelerated, and quantified to its fullest potential in the pursuit of improved patient care.

References

[1] Jan H. Ardenkjaer-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* (18 2003), pp. 10158–10163.

[2] Klaes Golman et al. “Metabolic imaging by hyperpolarized 13C magnetic resonance imaging for in vivo tumor diagnosis”. In: *Cancer research* 66.22 (2006), pp. 10855–10860.

[3] Sam E Day et al. “Detecting tumor response to treatment using hyperpolarized 13C magnetic resonance imaging and spectroscopy”. In: *Nature medicine* 13.11 (2007), pp. 1382–1387.

[4] Steen Hylgaard Joergensen et al. “Detection of increased pyruvate dehydrogenase flux in the human heart during adenosine stress test using hyperpolarized [1-13C] pyruvate cardiovascular magnetic resonance imaging”. In: *Journal of Cardiovascular Magnetic Resonance* 24.1 (2022), p. 34.

[5] Andrew JM Lewis, Damian J Tyler, and Oliver Rider. “Clinical cardiovascular applications of hyperpolarized magnetic resonance”. In: *Cardiovascular Drugs and Therapy* 34.2 (2020), pp. 231–240.

[6] Sven Månsson et al. “13C imaging—a new diagnostic platform”. In: *European radiology* 16.1 (2006), pp. 57–67.

[7] Panayiotis Nikolaou, Boyd M Goodson, and Eduard Y Chekmenev. “NMR hyperpolarization techniques for biomedicine”. In: *Chemistry–A European Journal* 21.8 (2015), pp. 3156–3166.

[8] FA Jenkins. “Nuclear Spins of the Carbon Isotopes”. In: *Physical Review* 74.4 (1948), p. 355.

[9] Robin De Graaf. *In vivo NMR Spectroscopy : principles and techniques*. Wiley Online Library, 2007, pp. 512–540.

[10] Ray Hashemi, William Bradley, and Christopher Lisanti. *MRI: the basics*. Second. Lippincott Williams and Wilkins, 2003, pp. 27–56.

[11] Felix Bloch. “Nuclear Induction”. In: *Physical Review* 70 (1946), pp. 460–474.

[12] Robert W Brown et al. *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons, 2014.

[13] Paul C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242 (1973), pp. 190–191.

[14] A.N. Garroway, P.K. Grannell, and P. Mansfield. “Image formation in NMR by a selective irradiative process”. In: *Journal of Physics C: Solid State Physics* 7 (1974), pp. L457–L462.

[15] W.A. Edelstein et al. “Spin warp NMR imaging and applications to human whole-body imaging”. In: *Physics in Medicine and Biology* 25.4 (1980), pp. 751–756.

[16] Donald B. Twieg. “The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods”. In: *Medical Physics* 10.5 (1983), pp. 610–621.

[17] Dong Liang et al. “Accelerating SENSE using compressed sensing”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 62.6 (2009), pp. 1574–1584.

[18] Jian-Xiong Wang et al. “Accelerated chemical shift imaging of hyperpolarized 13C metabolites”. In: *Magnetic resonance in medicine* 76.4 (2016), pp. 1033–1038.

[19] Dirk Mayer et al. “In vivo application of sub-second spiral chemical shift imaging (CSI) to hyperpolarized 13C metabolic imaging: Comparison with phase-encoded CSI”. In: *Journal of Magnetic Resonance* 204 (2 June 2010), pp. 340–345.

[20] Bernhard Gruber et al. “RF coils: A practical guide for nonphysicists”. In: *Journal of Magnetic Resonance Imaging* 48 (3 Sept. 2018), pp. 590–604.

[21] Daniel Hernandez and Kyoung-Nam Kim. “A Review on the RF Coil Designs and Trends for Ultra High Field Magnetic Resonance Imaging”. In: *Investigative Magnetic Resonance Imaging* 24 (3 2020), p. 95.

[22] Giulio Giovannetti et al. “Radio Frequency Coils for Hyperpolarized 13C Magnetic Resonance Experiments with a 3T MR Clinical Scanner: Experience from a Cardiovascular Lab”. In: *Electronics* 10.4 (2021), p. 366.

[23] Giulio Giovannetti et al. “Hyperpolarized ^{13}C MRS surface coil: Design and signal-to-noise ratio estimation”. In: *Medical Physics* 37.10 (2010), pp. 5361–5369.

[24] Thomas William Redpath. “Signal-to-noise ratio in MRI.” In: *The British journal of radiology* 71.847 (1998), pp. 704–707.

[25] DI Hoult. “The NMR receiver: a description and analysis of design”. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 12.1 (1978), pp. 41–77.

[26] Cecil E Hayes, Neil Hattes, and Peter B Roemer. “Volume imaging with MR phased arrays”. In: *Magnetic resonance in medicine* 18.2 (1991), pp. 309–319.

[27] Juan D. Sanchez-Heredia et al. “RF coil design for accurate parallel imaging on 13C MRSI using 23Na sensitivity profiles”. In: *Magnetic Resonance in Medicine* 88.3 (2022), pp. 1391–1405.

[28] Steven M Wright. “Full-wave analysis of planar radiofrequency coils and coil arrays with assumed current distribution”. In: *Concepts in Magnetic Resonance: An Educational Journal* 15.1 (2002), pp. 2–14.

[29] J Thomas Vaughan and John R Griffiths. *RF coils for MRI*. John Wiley and Sons, 2012.

[30] Hiroyuki Fujita et al. “RF surface receive array coils: the art of an LC circuit”. In: *Journal of magnetic resonance imaging* 38.1 (2013), pp. 12–25.

[31] Arne Reykowski, Charles Saylor, and G Randy Duensing. “Do we need preamplifier decoupling”. In: *Proc 19th Annual Meeting ISMRM, Montreal*. 2011.

[32] Ceylan Celil Guclu et al. “A method for preamplifier-decoupling improvement in quadrature phased-array coils”. In: *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 19.2 (2004), pp. 255–258.

[33] KM Huber et al. “A preamplifier for 7T MRI with extended dynamic range and integrated cable trap”. In: *Proc. Int. Soc. Mag. Reson. Med.* Vol. 16. 2008.

[34] Scott B King et al. “The effects of ultra low input impedance preamplifiers on phased array coil design”. In: *Proceedings of the 10th Annual Meeting of ISMRM, Honolulu, Hawaii*. 2002, p. 330.

[35] Arne Reykowski, Steven M Wright, and Jay R Porter. “Design of matching networks for low noise preamplifiers”. In: *Magnetic resonance in medicine* 33.6 (1995), pp. 848–852.

[36] Otward M Mueller et al. *Preamplifier circuit for magnetic resonance system*. US Patent 5,545,999. 1996.

[37] Peder EZ Larson et al. “Current methods for hyperpolarized [1-13C] pyruvate MRI human studies”. In: *Magnetic resonance in medicine* 91.6 (2024), pp. 2204–2228.

[38] J. Wolber et al. “Generating highly polarized nuclear spins in solution using dynamic nuclear polarization”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 526.1-2 (2004), pp. 173–181.

[39] Thomas C Eisenschmid et al. “Para hydrogen induced polarization in hydrogenation reactions”. In: *Journal of the American Chemical Society* 109.26 (1987), pp. 8089–8091.

[40] Ralph W Adams et al. “Reversible interactions with para-hydrogen enhance NMR sensitivity by polarization transfer”. In: *Science* 323.5922 (2009), pp. 1708–1711.

[41] Michal Fiedorowicz et al. “Hyperpolarized 13C tracers: Technical advancements and perspectives for clinical applications”. In: *Biocybernetics and Biomedical Engineering* 41 (4 Oct. 2021), pp. 1466–1485.

[42] Ralph E Hurd et al. “Hyperpolarized 13C metabolic imaging using dissolution dynamic nuclear polarization”. In: *Journal of Magnetic Resonance Imaging* 36.6 (2012), pp. 1314–1328.

[43] Jan Henrik Ardenkjær-Larsen et al. “Cryogen-Free Dissolution Dynamic Nuclear Polarization Polarizer Operating at 3.35 T, 6.70 T, and 10.1 T”. In: *Magnetic Resonance in Medicine* 81.3 (2019), pp. 2184–2194.

[44] Ying-Chieh Lai et al. “Hyperpolarized carbon-13 magnetic resonance imaging: technical considerations and clinical applications”. In: *Korean Journal of Radiology* 25.5 (2024), 459–472).

[45] Otto Warburg. “The metabolism of carcinoma cells”. In: *The Journal of Cancer Research* 9.1 (1925), pp. 148–163.

[46] 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.

[47] Marie A Schroeder et al. “Real-time assessment of Krebs cycle metabolism using hyperpolarized 13C magnetic resonance spectroscopy”. In: *The FASEB journal* 23.8 (2009), pp. 2529–2538.

[48] Marie A Schroeder et al. “In vivo assessment of pyruvate dehydrogenase flux in the heart using hyperpolarized carbon-13 magnetic resonance”. In: *Proceedings of the National Academy of Sciences* 105.33 (2008), pp. 12051–12056.

[49] Yan Xing et al. “Optimal variable flip angle schemes for dynamic acquisition of exchanging hyperpolarized substrates”. In: *Journal of magnetic resonance* 234 (2013), pp. 75–81.

[50] Peder Eric Zufall Larson and Jeremy W Gordon. “Hyperpolarized metabolic mri—acquisition, reconstruction, and analysis methods”. In: *Metabolites* 11.6 (2021), p. 386.

[51] Mark J Albers et al. “Hyperpolarized 13C lactate, pyruvate, and alanine: noninvasive biomarkers for prostate cancer detection and grading”. In: *Cancer research* 68.20 (2008), pp. 8607–8615.

[52] DJ Tyler et al. “Application of hyperpolarized magnetic resonance in the study of cardiac metabolism”. In: *Applied Magnetic Resonance* 34.3 (2008), pp. 523–531.

[53] Matthew G Vander Heiden, Lewis C Cantley, and Craig B Thompson. “Understanding the Warburg effect: the metabolic requirements of cell proliferation”. In: *science* 324.5930 (2009), pp. 1029–1033.

[54] Kristin L Granlund et al. “Hyperpolarized MRI of human prostate cancer reveals increased lactate with tumor grade driven by monocarboxylate transporter 1”. In: *Cell Metabolism* 31.1 (2020), pp. 105–114.

[55] Andreas Seraphim et al. “Quantitative cardiac MRI”. In: *Journal of Magnetic Resonance Imaging* 51.3 (2020), pp. 693–711.

[56] John Kurhanewicz et al. “Hyperpolarized 13C MRI: Path to Clinical Translation in Oncology”. In: *Neoplasia* 21 (1 Jan. 2019), pp. 1–16.

[57] William Dominguez-Viqueira et al. “Intensity correction for multichannel hyperpolarized 13C imaging of the heart”. In: *Magnetic resonance in medicine* 75.2 (2016), pp. 859–865.

[58] Michael A Ohliger et al. “55Mn based fiducial markers for rapid and automated RF coil localization for hyperpolarized 13C MRI”. In: *Magnetic resonance in medicine* 85.1 (2021), pp. 518–530.

[59] Boris Keil et al. “A 64-channel 3T array coil for accelerated brain MRI”. In: *Magnetic resonance in medicine* 70.1 (2013), pp. 248–258.

[60] Tony Zhou et al. “Multilayer Radiofrequency Coils: A Novel Surface Coil Design for Improved B1+ Efficiency and Signal-to-Noise”. In: *Proceedings of the 28th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*. ISMRM. 2020, p. 4028.

[61] JX Wang et al. “An 8-channel coil array for small animal 13C MR imaging”. In: *Proceedings of the 18th Annual Meeting of ISMRM*. 2010, p. 1489.

[62] Zihan Zhu et al. “Coil combination methods for multi-channel hyperpolarized 13C imaging data from human studies”. In: *Journal of Magnetic Resonance* 301 (2019), pp. 73–79.

[63] Azma Mareyam et al. “31-Channel Brain Array for Hyperpolarized 13C Imaging at 3T”. In: *Proceedings of the 25th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*. ISMRM. 2017, p. 1225.

[64] Juan Diego Sanchez-Heredia et al. “Multi-site benchmarking of clinical 13C RF coils at 3T”. In: *Journal of Magnetic Resonance* 318 (2020), p. 106798.

[65] M Robin Bendall, Alan Connelly, and Jamie M McKendry. “Elimination of coupling between cylindrical transmit coils and surface-receive coils for in vivo NMR”. In: *Magnetic resonance in medicine* 3.1 (1986), pp. 157–163.

[66] Peter B Roemer et al. “The NMR phased array”. In: *Magnetic resonance in medicine* 16.2 (1990), pp. 192–225.

[67] Boris Keil et al. “Size-optimized 32-channel brain arrays for 3 T pediatric imaging”. In: *Magnetic Resonance in Medicine* 66.6 (2011), pp. 1777–1787.

[68] Marija Nikolovski. “Detailed modeling of the human body in motion to investigate the electromagnetic influence of fields in a realistic environment”. In: (2018).

[69] Klaas P Pruessmann et al. “SENSE: sensitivity encoding for fast MRI”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 42.5 (1999), pp. 952–962.

[70] Mark A Griswold et al. “Generalized autocalibrating partially parallel acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine* 47.6 (2002), pp. 1202–1210.

[71] Rolf F Schulte et al. “Transmit gain calibration for nonproton MR using the Bloch–Siegert shift”. In: *NMR in Biomedicine* 24.9 (2011), pp. 1068–1072.

[72] Laura I Sacolick et al. “B1 mapping by Bloch-Siegert shift”. In: *Magnetic resonance in medicine* 63.5 (2010), pp. 1315–1322.

[73] Martin Meyerspeer et al. “An improved trap design for decoupling multinuclear RF coils”. In: *Magnetic resonance in medicine* 72.2 (2014), pp. 584–590.

[74] Wantcom Inc. *WMA3RA 0.03–1.0 GHz Low Noise Amplifier Datasheet*. Rev. R1. Accessed: May 22, 2024. Wantcom Inc. Chanhassen, MN, USA. URL: `https://www.wantcominc.com/DataSheets/WMA/WMA3RA-R1.pdf`.

[75] Danilo Brizi et al. “Design of distributed spiral resonators for the decoupling of MRI double-tuned RF coils”. In: *IEEE Transactions on Biomedical Engineering* 67.10 (2020), pp. 2806–2816.

[76] Pascal Spincemaille et al. “Optimal coil array design: the two-coil case”. In: *Magnetic resonance imaging* 25.5 (2007), pp. 671–677.

[77] Ray F Lee, Randy O Giaquinto, and Christopher J Hardy. “Coupling and decoupling theory and its application to the MRI phased array”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 48.1 (2002), pp. 203–213.

[78] Roberta Frass-Kriegl et al. “Multi-loop radio frequency coil elements for magnetic resonance imaging: theory, simulation, and experimental investigation”. In: *Frontiers in Physics* 7 (2020), p. 237.

[79] Nicola De Zanche et al. “ISMRM best practices for safety testing of experimental RF hardware”. In: *ISMRM RF Hardware Recommendations* (2022).

[80] Jens Hoffmann et al. “Safety testing and operational procedures for self-developed radiofrequency coils”. In: *NMR in Biomedicine* 29.9 (2016), pp. 1131–1144.

[81] Yesu Feng et al. “Development and testing of hyperpolarized 13C MR calibrationless parallel imaging”. In: *Journal of Magnetic Resonance* 262 (2016), pp. 1–7.

[82] Ayaka Shinozaki et al. “Enabling SENSE accelerated 2D CSI for hyperpolarized carbon-13 imaging”. In: *Scientific Reports* 14.1 (2024), p. 20591.

[83] Wolfgang Bogner, Ricardo Otazo, and Anke Henning. “Accelerated MR spectroscopic imaging—a review of current and emerging techniques”. In: *NMR in Biomedicine* 34.5 (2021), e4314.

[84] Craig H Meyer et al. “Simultaneous spatial and spectral selective excitation”. In: *Magnetic resonance in medicine* 15.2 (1990), pp. 287–304.

[85] Jeremy W Gordon, Daniel B Vigneron, and Peder EZ Larson. “Development of a symmetric echo planar imaging framework for clinical translation of rapid dynamic hyperpolarized 13C imaging”. In: *Magnetic resonance in medicine* 77.2 (2017), pp. 826–832.

[86] Jeremy W Gordon et al. “Fast imaging for hyperpolarized MR metabolic imaging”. In: *Journal of Magnetic Resonance Imaging* 53.3 (2021), pp. 686–702.

[87] Y-F Yen et al. “Imaging considerations for in vivo 13C metabolic mapping using hyperpolarized 13C-pyruvate”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 62.1 (2009), pp. 1–10.

[88] Dirk Mayer et al. “Fast metabolic imaging of systems with sparse spectra: application for hyperpolarized 13C imaging”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 56.4 (2006), pp. 932–937.

[89] Marc S Ramirez et al. “Radial spectroscopic MRI of hyperpolarized [1-13C] pyruvate at 7 tesla”. In: *Magnetic resonance in medicine* 72.4 (2014), pp. 986–995.

[90] Oliver Bieri and Klaus Scheffler. “Fundamentals of balanced steady state free precession MRI”. In: *Journal of Magnetic Resonance Imaging* 38.1 (2013), pp. 2–11.

[91] Christian Østergaard Mariager et al. “Multi-echo balanced steady state imaging of hyperpolarized [1-13C]-pyruvate at 9.4 T”. In: *ISMRM 2017 Proceedings* (2017).

[92] Scott B Reeder et al. “Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): application with fast spin-echo imaging”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 54.3 (2005), pp. 636–644.

[93] Jesse Hamilton, Dominique Franson, and Nicole Seiberlich. “Recent advances in parallel imaging for MRI”. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 101 (Aug. 2017), pp. 71–95.

[94] Arjun Arunachalam et al. “Accelerated spectroscopic imaging of hyperpolarized C-13 pyruvate using SENSE parallel imaging”. In: *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In vivo* 22.8 (2009), pp. 867–873.

[95] James Tropp et al. “Multi-channel metabolic imaging, with SENSE reconstruction, of hyperpolarized [1-13C] pyruvate in a live rat at 3.0 tesla on a clinical MR scanner”. In: *Journal of Magnetic Resonance* 208 (1 Jan. 2011), pp. 171–177.

[96] Rie B. Olin et al. “Three-dimensional accelerated acquisition for hyperpolarized 13 C MR with blipped stack-of-spirals and conjugate-gradient SENSE”. In: *Magnetic resonance in medicine* 84 (2 Aug. 2020), pp. 519–534.

[97] Simon Hu et al. “Compressed sensing for resolution enhancement of hyperpolarized 13C flyback 3D-MRSI”. In: *Journal of magnetic resonance* 192.2 (2008), pp. 258–264.

[98] Jeremy W. Gordon et al. “3D hyperpolarized C-13 EPI with calibrationless parallel imaging”. In: *Journal of Magnetic Resonance* 289 (Apr. 2018), pp. 92–99.

[99] Mark Chiew. *Teaching*. `https://mchiew.github.io/Teaching.html`. Accessed: 2025-02-02. 2026.

[100] Philip M Robson et al. “Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 60.4 (2008), pp. 895–907.

[101] Michael S. Hansen. *SUNRISE 2013: Parallel Imaging Tutorial*. `http://hansenms.github.io/sunrise/sunrise2013/`. Accessed: 14 January 2024. 2013.

[102] Ibrahim Karaman, Rui Climaco Pinto, and Gonçalo Graça. “Metabolomics data preprocessing: from raw data to features for statistical analysis”. In: *Comprehensive analytical chemistry*. Vol. 82. Elsevier, 2018, pp. 197–225.

[103] Martin Wilson et al. “Methodological consensus on clinical proton MRS of the brain: Review and recommendations”. In: *Magnetic resonance in medicine* 82.2 (2019), pp. 527–550.

[104] James Keeler. *Understanding NMR Spectroscopy*. 2nd. Chichester, West Sussex: John Wiley and Sons, 2010.

[105] Geoffrey J Topping et al. “Hyperpolarized 13C spectroscopy with simple slice-and-frequency-selective excitation”. In: *Biomedicines* 9.2 (2021), p. 121.

[106] GA Nagana Gowda and Daniel Raftery. “Overview of NMR spectroscopy-based metabolomics: opportunities and challenges”. In: *NMR-Based Metabolomics: Methods and Protocols* (2019), pp. 3–14.

[107] Yi Rao et al. “Hyperpolarized [1-13C] pyruvate-to-[1-13C] lactate conversion is rate-limited by monocarboxylate transporter-1 in the plasma membrane”. In: *Proceedings of the National Academy of Sciences* 117.36 (2020), pp. 22378–22389.

[108] Yoichi Takakusagi et al. “Hyperpolarized [1-13C] pyruvate NMR spectroscopy reveals transition of tumor energy metabolism in microscale multicellular spheroids”. In: *Scientific Reports* 15.1 (2025), p. 19303.

[109] Oliver J Rider and Damian J Tyler. "Clinical implications of cardiac hyperpolarized magnetic resonance imaging". In: *Journal of Cardiovascular Magnetic Resonance* 15.1 (2013), p. 93.

[110] Helen J Atherton et al. "Role of pyruvate dehydrogenase inhibition in the development of hypertrophy in the hyperthyroid rat heart: a combined magnetic resonance imaging and hyperpolarized magnetic resonance spectroscopy study". In: *Circulation* 123.22 (2011), pp. 2552–2561.

[111] Yafang Xu et al. "Hyperpolarized 13C magnetic resonance imaging can detect metabolic changes characteristic of penumbra in ischemic stroke". In: *Tomography* 3.2 (2017), p. 67.

[112] Michael S Dodd et al. "Impaired in vivo mitochondrial Krebs cycle activity after myocardial infarction assessed using hyperpolarized magnetic resonance spectroscopy". In: *Circulation: Cardiovascular Imaging* 7.6 (2014), pp. 895–904.

[113] James T. Grist et al. "Quantifying normal human brain metabolism using hyperpolarized [1–13C]pyruvate and magnetic resonance imaging". In: *NeuroImage* 189 (Apr. 2019), pp. 171–179.

[114] Jack J Miller et al. "13C pyruvate transport across the blood-brain barrier in preclinical hyperpolarised MRI". In: *Scientific reports* 8.1 (2018), p. 15082.

[115] Ilwoo Park et al. "Development of methods and feasibility of using hyperpolarized carbon-13 imaging data for evaluating brain metabolism in patient studies". In: *Magnetic resonance in medicine* 80.3 (2018), pp. 864–873.

[116] Vesselin Z Miloushev et al. "Metabolic imaging of the human brain with hyperpolarized 13C pyruvate demonstrates 13C lactate production in brain tumor patients". In: *Cancer research* 78.14 (2018), pp. 3755–3760.

[117] Alexander S Giuffrida et al. "An Accelerated Spectroscopic MRI Metabolite Quantification Based on a Deep Learning Method for Radiation Therapy Planning in Brain Tumor Patients". In: *Cancers* 17.3 (2025), p. 423.

[118] Fulvio Zaccagna et al. "Imaging glioblastoma metabolism by using hyperpolarized [1-13C] pyruvate demonstrates heterogeneity in lactate labeling: a proof of principle study". In: *Radiology: Imaging Cancer* 4.4 (2022), e210076.

[119] Edward P Hackett et al. "Imaging acute metabolic changes in patients with mild traumatic brain injury using hyperpolarized [1-13C] pyruvate". In: *Iscience* 23.12 (2020).

[120] Biranavan Uthayakumar et al. "Age-associated change in pyruvate metabolism investigated with hyperpolarized 13C-MRI of the human brain". In: *Human Brain Mapping* 44.10 (2023), pp. 4052–4063.

[121] Zheng Ye et al. "Hyperpolarized carbon 13 MRI in liver diseases: Recent advances and future opportunities". In: *Liver International* 42.5 (2022), pp. 973–983.

[122] Uffe Kjærgaard et al. "Hyperpolarized [1-13C] pyruvate as a possible diagnostic tool in liver disease". In: *Physiological reports* 6.23 (2018), e13943.

[123] Philip M Lee et al. "Whole-abdomen metabolic imaging of healthy volunteers using hyperpolarized [1-13C] pyruvate MRI". In: *Journal of Magnetic Resonance Imaging* 56.6 (2022), pp. 1792–1806.

[124] Matthew E Merritt et al. “Flux through hepatic pyruvate carboxylase and phosphoenolpyruvate carboxykinase detected by hyperpolarized 13C magnetic resonance”. In: *Proceedings of the National Academy of Sciences* 108.47 (2011), pp. 19084–19089.

[125] Chung-Man Moon et al. “Metabolic biomarkers for non-alcoholic fatty liver disease induced by high-fat diet: In vivo magnetic resonance spectroscopy of hyperpolarized [1-13C] pyruvate”. In: *Biochemical and biophysical research communications* 482.1 (2017), pp. 112–119.

[126] Sonal Josan et al. “Assessing inflammatory liver injury in an acute CCl4 model using dynamic 3D metabolic imaging of hyperpolarized [1-13C] pyruvate”. In: *NMR in Biomedicine* 28.12 (2015), pp. 1671–1677.

[127] Cornelius von Morze et al. “Non-invasive detection of divergent metabolic signals in insulin deficiency vs. insulin resistance in vivo”. In: *Scientific reports* 8.1 (2018), p. 2088.

[128] Klaes Golman, René in ‘t Zandt, and Mikkel Thaning. “Real-time metabolic imaging”. In: *Proceedings of the National Academy of Sciences* 103.30 (2006), pp. 11270–11275.

[129] Sarah J Nelson and Truman R Brown. “The accuracy of quantification from 1D NMR spectra using the PIQABLE algorithm”. In: *Journal of Magnetic Resonance (1969)* 84.1 (1989), pp. 95–109.

[130] Stephen W. Provencher. “Estimation of metabolite concentrations from localized in vivo proton NMR spectra”. In: *Magnetic Resonance in Medicine* 30 (6 Dec. 1993), pp. 672–679.

[131] Martin Wilson et al. “A constrained least-squares approach to the automated quantitation of in vivo 1H magnetic resonance spectroscopy data”. In: *Magnetic Resonance in Medicine* 65 (1 Jan. 2011), pp. 1–12.

[132] Robin Simpson et al. “Advanced processing and simulation of MRS data using the FID appliance (FID-A)—an open source, MATLAB-based toolkit”. In: *Magnetic resonance in medicine* 77.1 (2017), pp. 23–33.

[133] S. J. Nelson, E. Ozturk-Isik, and D. B. Vigneron. “SIVIC: Spectroscopic Imaging Visualization and Computing”. In: *Proceedings of SPIE* 7964 (2011), 79640P.

[134] DDCF Stefan et al. “Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package”. In: *Measurement Science and Technology* 20.10 (2009), p. 104035.

[135] Brian J Soher et al. “VeSPA: integrated applications for RF pulse design, spectral simulation and MRS data analysis”. In: *Magnetic resonance in medicine* 90.3 (2023), pp. 823–838.

[136] Jun Chen et al. “Preoperative imaging of glioblastoma patients using hyperpolarized 13C pyruvate: potential role in clinical decision making”. In: *Neuro-Oncology Advances* 3.1 (2021), vdab092.

[137] Donghyun Hong et al. “Acquisition and quantification pipeline for in vivo hyperpolarized 13C MR spectroscopy”. In: *Magnetic resonance in medicine* 87.4 (2022), pp. 1673–1687.

[138] Georg Oeltzschner et al. “Osprey: Open-source software for Magnetic Resonance Spectroscopy data analysis”. In: *NMR in Biomedicine* 34.12 (2021), e4426.

[139] William T Clarke, Charlotte J Stagg, and Saad Jbabdi. “FSL-MRS: An end-to-end spectroscopy analysis package”. In: *Magnetic resonance in medicine* 85.6 (2021), pp. 2950–2964.

[140] Martin Gajdošík et al. “INSPECTOR: free software for magnetic resonance spectroscopy data inspection, processing, simulation and analysis”. In: *Scientific reports* 11.1 (2021), p. 2094.

[141] Lucian A.B. Purvis et al. “OXSA: An open-source magnetic resonance spectroscopy analysis toolbox in MATLAB”. In: *PLOS ONE* 12 (9 Sept. 2017), e0185356.

[142] 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 (1 Nov. 1997), pp. 35–43.

[143] Jia Xu et al. “PyAMARES, an Open-Source Python Library for Fitting Magnetic Resonance Spectroscopy Data”. In: *Diagnostics* 14.23 (2024), p. 2668.

[144] J. W. C. van der Veen et al. “Accurate quantification of in vivo ^{31}P NMR signals using the variable projection method and prior knowledge”. In: *Magnetic Resonance in Medicine* 6.1 (1988), pp. 92–98.

[145] Ying Xiao et al. “A graphical pipeline platform for MRS data processing and analysis: MRSpecLAB”. In: *Frontiers in neuroimaging* 4 (2025), p. 1610658.

[146] Zuojun Wang et al. “Using a deep learning prior for accelerating hyperpolarized 13C MRSI on synthetic cancer datasets”. In: *Magnetic Resonance in Medicine* 92.3 (2024), pp. 945–955.

[147] Conor C Horgan and Mads S Bergholt. “Spectrai: A deep learning framework for spectral data”. In: *arXiv preprint arXiv:2108.07595* (2021).

[148] Anne-Marie L Seymour et al. “In vivo assessment of cardiac metabolism and function in the abdominal aortic banding model of compensated cardiac hypertrophy”. In: *Cardiovascular research* 106.2 (2015), pp. 249–260.

[149] S. Anderson. “Application of magnetic resonance imaging and spectroscopy to the assessment of cardiovascular disease”. PhD thesis. University of Oxford, 2023.

[150] Michael S Dodd et al. “In vivo mouse cardiac hyperpolarized magnetic resonance spectroscopy”. In: *Journal of Cardiovascular Magnetic Resonance* 15.1 (2013), p. 19.

[151] Oxford Molecular Imaging Group. *Oxford Molecular Imaging Framework*. Private repository. Access restricted due to GE Healthcare IP and unpublished work; access granted to examiners upon request. 2026.

[152] Marie A. Schroeder et al. “Hyperpolarized magnetic resonance: A novel technique for the in vivo assessment of cardiovascular disease”. In: *Circulation* 124 (14 Oct. 2011), pp. 1580–1594.

[153] S Cavassila et al. “Cramér–Rao bounds: an evaluation tool for quantitation”. In: *NMR in Biomedicine: An International Journal Devoted to the Development and Application of Magnetic Resonance In Vivo* 14.4 (2001), pp. 278–283.

[154] David M. Harrison. *Error Analysis in Experimental Physical Science: Propagation of Errors*. Department of Physics, University of Toronto. Accessed: 2025-05-22. 2011. URL: https://faraday.physics.utoronto.ca/PVB/Harrison/ErrorAnalysis/Propagation.html.

[155] Bernard Rosner. "Fundamentals of Biostatistics (The 7th edition)". In: *Boston, MA: Brooks/Cole* (2011).

[156] Matthew L Zierhut et al. "Kinetic modeling of hyperpolarized 13C1-pyruvate metabolism in normal rats and TRAMP mice". In: *Journal of magnetic resonance* 202.1 (2010), pp. 85–92.

[157] University of Oxford. *University Policy on the Use of Animals in Scientific Research*. Accessed: 2024-05-22. 2022. URL: https://www.ox.ac.uk/news-and-events/animal-research/university-policy-on-the-use-of-animals-in-scientific-research.