Dynamic Nuclear Polarisation as a Probe of Metabolism in Pathophysiology

Jack Julian James Jenkins Miller

Department of Physics
Department of Physiology, Anatomy & Genetics
The Oxford Institute for Radiation Oncology
The Life Sciences' Interface Doctoral Training Centre
St. Hugh's College

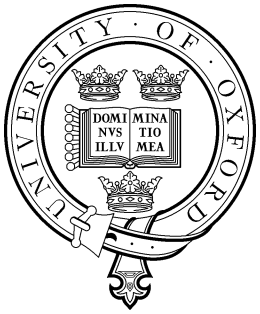

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy
Trinity Term 2015



The work described in this text would not have been possible without the financial support of the John Fell Fund, the British Heart Foundation, Cancer Research UK, GE Healthcare, the Institute of Physics, St. Hugh's College and Trinity College in the University of Oxford, the research councils BBSRC, the MRC, and, above all, EPSRC and the Life Sciences' Interface Doctoral Training Centre.

Typeset with Palatino (body; 10 pt with 20 pt leading), Garamond, Helvetica, and the LaTeX 2_ε document preparation system. This complies with the typographic requirements for DPhil Theses as stipulated by the Department of Physics.

© 2015 Jack J. Miller. This material is intended to be released, unless explicitly stated by a more proscriptive requirement elsewhere, under the CC-BY-NC-SA licence. It is requested that creators of derivative works keep the author informed of their exploits via `jack.miller@physics.org`.

To Emma,
my family (felines included),
and my late supervisor,
Prof. Harry Jones.

Things are only impossible until they are not.

— Jean-Luc Picard, When the Bough Breaks, ST-TNG

Anything less than immortality is a complete waste of time.

— Bender, Lethal Inspection, Futurama

Abstract

Dynamic Nuclear Polarisation as a Probe of Metabolism in Pathophysiology

A thesis submitted in partial fulfillment for the requirements of the degree of Doctor of Philosophy, Trinity Term 2015

Jack Julian James Jenkins Miller, St. Hugh's College

Spin-half nuclei in a magnetic field possessing a population difference in excess of that expected at thermal equilibrium are said to be hyperpolarised. Dynamic Nuclear Polarisation (DNP) is a method to generate hyperpolarised systems, pumping nuclear spins into a given state at low temperature. Hyperpolarised samples can be rapidly dissolved in a hot solvent, and maintain a highly polarised spin population that relaxes to the new thermal equilibrium over several minutes. It is possible to inject hyperpolarised molecules into an organism, sample the population difference through NMR spectroscopy, and obtain spectra with a large ($> 10^4$) increase in their signal-to-noise ratio compared to thermal equilibrium. The use of hyperpolarised [1-^{13}C]pyruvate allows for the quantification of the rate of [1-^{13}C]lactate production, increases of which are a defining characteristic of cancer; or the rate of ^{13}C-bicarbonate production, which is altered in heart disease.

This thesis describes the development of magnetisation-efficient sequences of spectrally-and-spatially selective radiofrequency pulses and rapid imaging readouts that are able to image the metabolism of hyperpolarised [1-^{13}C]pyruvate in three dimensional space, with a reconstructed resolution of $1 \times 1 \times 2$ mm^3. The technique's sensitivity to metabolism is demonstrated, and attempts are made to image cancer that has metastasised to the brain prior to the point of conventional detection through its metabolic phenotype. It is also used to image metabolism in the post ischæmic myocardium, where it reports on regions of non-viable tissue.

Additionally, hyperpolarised ^{13}C-urea is used as a 'metabolically inert' probe of perfusion in the myocardium. Flow-sensitisation gradients and a spiral readout allow ^{13}C-urea to be resolved in the myocardium. The ability to 'co-polarise' [1-^{13}C]pyruvate and ^{13}C-urea is demonstrated. This technique could simultaneously resolve perfusion deficits and metabolic changes in the acutely damaged heart, with several potential advantages over current clinical methods.

Finally, several methods for analysing data generated by hyperpolarised experiments are proposed: weighted averaging for combining spectroscopic data of varying quality; novel Bayesian approaches to quantifying metabolic rate constants of interest following the injection of hyperpolarised [1-^{13}C]pyruvate; and the Fast Padé Transformation as an improved method for the reconstruction of truncated hyperpolarised spectral data.

DNP is a technique that, following its development for particle physics experiments in the 1970s, is transitioning to clinical practice as a safe, novel molecular imaging modality. It is hoped that the techniques proposed in this text ultimately find wide utility in the assessment of cancer and heart disease, which combined are responsible for approximately 50% of deaths occurring in the developed world.

Jack J. Miller, 2015

Acknowledgements

THIS thesis is as much a work of a lifetime as it is of the last three and a half years. I am indebted to many others who have helped my work during this time. I should very much like to thank those $\sim 10^{40}$ organisms who have (evolutionarily) been responsible for my existence – but, in particular, **my family**, whose support (especially edible) has kept me going for the last 2.6×10^{-2} millenia. Words obviously cannot say 'thank you' enough. I should also like to thank those friends and educators who helped me get to the point of crossing the first $\hbar$ in this document.

Iain Johnston deserves much praise for being the person who, by combining interesting mathematics, OpenGL, and viruses, made me realise that there are many reasons for a physicist to be interested in biology. **John Chalker**, your knowledge transfer of condensed matter physics was exceptionally efficient; **Don Grainger**, thank you for showing me that experimentalists most definitely have a sense of humour. **David Freeborn**, **Alasdair Gunn**, **Sam Smart**, **Hannah & Chris Dingwall-Jones**, **Emma Taylor** and many others: you were instrumental in much more than perhaps you realise (not least for keeping me sane whilst writing up my first ever journal article, and stopping Maths, Mathematica and Matlab from being my only friends).

I must thank those who work in the Department of Physics, as they have always been exceptionally intelligent, kind, and helpful – **Mark Jones**, you are a god of e^-. **Jonathan Jones** – I am very grateful for your fantastic understanding of the quantum mechanics of NMR, and for many excellent lunches in Brasenose. **Stephen Tucker** – thank you for providing several pints of beer, light relief, administrative support and good advice. I hope to repay at least one of these items shortly.

I particularly wish to thank the fantastic support I have been lucky enough to receive from the DTC[1] during my DPhil. Through you, I have thoroughly enjoyed learning about everything from deriving spirals conduction patterns in VF with Lyapunov functions and the Bidomain equations, to the glassy funding landscape of academia. **David Gavaghan**, **Gail Preston**, **Eoin Malins**, **Francesca Summers**, **Lorraine Damerell** and **Sam Miles** – you've all been incredibly supportive, nice people who willingly share knowledge and a listening ear. Thank you. Eoin, I'm glad you like the goats: *we have both licked the helicopter.*

The Open Source Community deserve recognition for creating such wonderful tools that have benefited both my work and humanity.[2] Likewise, thanks to all the musical organisations and people who have given me so much pleasure (and food) over the last eight years; in particular *Intermezzo*, and **Mark & Sarah Jenkinson**, who have mixed music, `fsl`, and stereochemistry to great delight.

Elspeth Garman – you are a fantastic, fantastic person, and your (unrequested!) faith in me has been both touching and humbling. I've very much enjoyed teaching your maths classes, learning to share your enthusiasm for all things crystallographic, and trying to copy your unique ability to make complicated concepts clear. In the words of an (anonymous) biochemist in my year, "Elspeth's a real babe".

Cycling up a hill, **James Larkin**, your are undoubtedly the most technically minded biochemist I've ever met (and I hope the others are very grateful for it!), **Kevin Ray** – you're excellent at injecting a

[1] Yes, EPSRC, I *do* mean 'DTC' and not 'CDT'.

[2] Even if I fear that the Turing-complete combination of `sed`, `awk`, and `vim` will eventually become self aware

much-needed dose of light relief into the subterranean gloom of the RRI. After GlucoCEST, please, *please* invent Insulin Chemical Exchange Saturation Transfer. **Александр Храпичев** [*sic*], **Paul Kinchesh**, **Sean Smart**, **Irvin Teh** and **Jürgen Schneider** – thank you for many interesting discussions about the undocumented inner workings of `VnmrJ` and the console, and for sharing much useful source code. While we're up the hill, **Sébastien Serres** – it has been exceedingly entertaining to enjoy your company for the last few years. Thank you for never failing to share your mind. I hope you like Nottingham more than Ipswich.

Letting mgh reduce again, **Arne Bruyneel** – you have taught me that electronics is fun as well as frustrating, and that banging that drum about statistical understanding in biology will undoubtedly provide a lifetime of entertainment. Sharing the lab, and later an office, with you has been very enjoyable, and I will be sad for it to end: thank you for providing a constant 'nerd refuge' for me.

Carolyn Carr – as well as being a coffee drinker of the finest order, I truly have appreciated your excellent knowledge of chemistry, cell biology, people, and how I should manage my time! **Elizabeth Eggleston** and **Nicola Chubb** – without you, we would have no reagents, no plasticware, and no hope. Thank you for being a pleasant dose of (almost normal) reality! **Lisa, Latt, Luz, Brianna,** *et al.*, thank you for answering my random questions about biology; **Kieran Clarke**, thank you for prompting me to ask many more. **Lowri Cochlin** – not only are you a fellow coffee drinker, but I envy your excellent craft skills and no-nonsense approach to RF engineering. Thank you for the occasional loan of everything from a balun to a network analyser. **Esben Hansen** – it was a real privilege getting to know you (and your carrots) for the six wonderful, albeit short, months you were with us. One day, I hope to build a Linux-powered coffee machine that is half as good as your one.

Clinicians, **Rhys Evans**, **Pete Cox**, **David Holdsworth**, **Ollie Rider** and, in particular, **Andy Lewis** – you have provided (interesting!) challenges I would not have considered, contexts I was unaware of, and a (nearly) straight face when I asked where my epididymis was. **Lydia Le Page** and **Lucy Ambrose** – thank you for much experimental help, yet more context and coffee, and a good, old-fashioned chin-wag. **Vicky Ball** & **Lucia Giles** – not only have you both been bastions of support (and provided an endless supply of cake), but I have lost count of the number of times you've sneakily helped with a bit of one of my experiments whilst otherwise very occupied.

Angus Lau – you know an intimidatingly huge amount about MR, and I'm incredibly grateful that you've been able to share some of your knowledge with me over these last few years. Multiplayer `vim` with you is truly an enjoyable past-time. I have no doubt that your career-space trajectory will be very highly accelerated. You truly are, definitively, the better human being. Thank you!

Niki Sibson – it has been a pleasure working with your group for the last four years. Your encyclopædic knowledge from clinical practice to MR is both humbling and frankly gives me hope for biochemistry. I am sure that the Sibsonites shall bring many molecules to market in the forthcoming years, and wish you luck with your unbelievably intimidating Gantt charts!

Penultimately, *Dr*. **Damian Tyler** – and then *Titular Ass. Prof.* Tyler, before *Statutory Ass. Prof.* Tyler, with undoubtedly *Professor* Tyler following shortly – thank you so much for letting me spend the last few years of my life in your fantastic company, and for giving me the opportunity to do such interesting research. All it took was one seminar in the DTC, and I realised that I'd love to work on an area of *truly* interdisciplinary science (even if I expressed such a desire by babbling incoherently about the deleterious effects of deuterium whilst descending down the stairs).

Finally, **Emma Elizabeth Sadler** – thank you for not only proof-reading this entire tome, but for sharing the last five years of your life with me. I hope that we share many, many more.

$\left((x-3)\,x+\frac{9}{4}(y-3)y+(z-3)z+\frac{137}{16}\right)^3=\frac{1}{80}\left(z-\frac{3}{2}\right)^3\left(80(x-3)x+9(y-3)y+\frac{801}{4}\right),\quad x,\,y,\,z\in\mathbb{R}^+$

Research Outcomes

Peer-Reviewed Journal Articles

The bulk of the material contained in **Chapter 6** was published in May 2015 as "Robust and high resolution hyperpolarized metabolic imaging of the rat heart at 7T with 3D spectral-spatial EPI", *J. J. Miller, A. Z. Lau, I. Teh, J. E. Schneider, P. Kinchesh, S. Smart, V. Ball, N. R. Sibson, and D. J. Tyler*, in the journal Magnetic Resonance in Medicine. After one round of minor corrections, it was received with reviewer's comments "thanking [the authors] for an excellent contribution to the field", and helped the author win the Nicholas Kurti Prize for Condensed Matter Physics and an EPSRC Doctoral Prize Fellowship. The author and Angus Lau pair-programmed the sequence on the scanner, and the author reconstructed the data and wrote the first draft. Snippets of source code, advice, and intellectual input were provided from Irvin Teh, Jürgen Schneider, Paul Kinchesh and Sean Smart, and Vicky Ball provided much practical assistance with animal handling and performing dissolutions. Additionally, the method was used to make specific points relevant to the argument of "In Vivo Assessment of Cardiac Metabolism and Function in the Abdominal Aortic Banding Model of Compensated Cardiac Hypertrophy", *A-M. L. Seymour, L. Giles, V. Ball, J. J. Miller, K. Clarke, C. A. Carr, and D. J. Tyler*, published in the journal Cardiovascular Research in March 2015.

Chapter 8 describes the development of hyperpolarised ^{13}C-urea as a probe of perfusion, and was published in May 2015 as "Cardiac perfusion imaging using hyperpolarized ^{13}C-urea using flow sensitizing gradients", *A. Z. Lau, J. J. Miller, M. D. Robson, and D. J. Tyler*, in the journal Magnetic Resonance in Medicine. The author and Angus Lau pair-programmed the sequence on the scanner investigated the urea polarisation protocol, and Angus Lau reconstructed the data and wrote the first draft of the paper. The author was responsible for animal handling and performing dissolutions. The chapter is a re-written, extended and expanded form of the article that contains more experimental and theoretical detail. Additional hints are made of work that has been submitted to the journal Magnetic Resonance in Medicine as "A combined cardiac metabolism-perfusion exam using co-polarized [1-^{13}C]pyruvate and ^{13}C-urea", *A. Z. Lau, J. J. Miller, M. D. Robson, and D. J. Tyler* describing co-hyperpolarised ^{13}C-urea and [1-^{13}C]pyruvate as a simultaneous probe of perfusion and metabolism.

The section of **Chapter 9** describing the Fast Padé Transformation as a method of reconstructing accelerated MRS data is currently under review for inclusion in the journal Magnetic Resonance in Medicine. Likewise, other manuscripts based on various aspects of this thesis are currently in preparation.

In the time taken to complete the work described in this document, the author has also proposed a set of differential equations to model neurological processes occurring in adults, fitted their parameters to an experimental dataset, performed a stability analysis and wrote the article "A mathematical model of adult subventricular neurogenesis", *J. M. A. Ashbourn, J. J. Miller, V. Reumers, V. Baekelandt, and L. Geris*, published in the Journal of the Royal Society Interface in 2012. This model has subsequently been expanded to include other neurological processes, published as "Mechanisms of cell migration in the adult brain: modelling subventricular neurogenesis", *A. Van Schepdael, J. M. A. Ashbourn, R. Beard, J. J. Miller, and L. Geris*, in the journal Computational Methods in Biomechanics and Biomedical Engineering, in March 2013. Additionally, a high-resolution cardiac T_2^* mapping protocol that the author wrote

was used as a proxy for iron concentration in the mouse heart, resulting in the publication "Cardiac ferroportin regulates cellular iron homeostasis and is important for cardiac function", *S. Lakhal-Littleton, M. Wolna, C. A. Carr, J. J. J. Miller, H. C. Christian, V. Ball, A. Santos, R. Diaz, D. Biggs, R. Stillion, P. Holdship, F. Larner, D. J. Tyler, K. Clarke, B. Davies, and P. A. Robbins*, published in the journal Proceedings of the National Academy of Sciences of the USA, in February 2015. The author has additionally shown that many problems of classical Greek geometry can be solved if one is allowed to construct lines by folding the plane, in review for inclusion in the College Mathematics Journal as "Had Euclid only been to Japan: Geometric Constructions with paper Folding", *J. J. Miller and R. Norris*. For the sake of brevity, and to ensure that the work presented forms a coherent whole, none of these publications are discussed further.

Peer-Reviewed Conference Abstracts

The contents of **Chapter 5** were presented at the *4th International Dynamic Nuclear Polarisation Symposium (Helsingør, 2013):* as Hyperpolarised cardiac imaging with spectral-spatial EPI, Miller, J J., Lau, A Z., Tyler, D J. Additional subsets were presented at the *Biomedical Imaging Festival (2013, Oxford):* as Metabolic Imaging with Hyperpolarised Spectral-Spatial EPI, Miller, J J., Lau, A Z., Tyler, D J.

Chapter 6 contains material that was presented in overview to the community in *Biomedical Imaging Festival (2014, Oxford):* as Imaging the Thermodynamics of Life, Miller, J J., Lau, A Z., Tyler, D J. Additionally, a more expanded variant was presented to the *ISMRM (June 2015, Toronto)* as High resolution hyperpolarized metabolic imaging with three-dimensional spectral-spatial EPI at 7 T. Miller, J J., Lau, A Z., Tyler, D J, and a general overview of the method for an MR audience was presented to the *British Chapter of the ISMRM (2015, Stevenage)* as High resolution hyperpolarized metabolic imaging with three-dimensional spectral-spatial EPI at 7T, Miller, J J., Lau, A Z., Tyler, D J. The author was fortunate to be awarded a £250 Oral Presentation Prize at this meeting.

The work described in **Chapter 8** was also presented to the *ISMRM (2015, Toronto)* as Flow-sensitizing gradients for first-pass perfusion imaging using hyperpolarized ^{13}C-urea in the rat heart. Lau, A Z., Miller, J J., Tyler, D J.

In addition to the work described in this thesis, the author has also presented four relevant abstracts at the *ISMRM (2015, Toronto)*:

- Hyperpolarized ketone body metabolism in the *in vivo* rat heart. Lau, A Z., Miller, J J., Tyler, D J., in which the synthesis and observation of metabolism of hyperpolarised sodium [1-^{13}C]acetoacetate is observed;

- The acute effects of metformin on cardiac and hepatic metabolism: a hyperpolarized [1-^{13}C]pyruvate magnetic resonance spectroscopy study; Lewis, A L., McCullum, C., Miller, J J., Heather, L., Tyler, D J., in which the anti-diabetic drug metformin is shown to produce changes in hepatic and cardiac metabolism detectable through hyperpolarised [1-^{13}C]pyruvate spectroscopy;

- Ferroportin Regulates Cardiac Iron Homeostasis. Lakhal-Littleton, S., Miller, J J., Wolna, M., Carr, C A., Ball, V., Santos, A., Diaz, R., Biggs, D., Tyler, D J., Clarke, K., Davies, B., Robbins, P A., which details the use of relaxometry measurements to probe a novel biological system for regulating cardiac iron as related to the *PNAS* article referenced above; and

- The effects of gadolinium on the hyperpolarization of [1-^{13}C]pyruvate at 3.35 T and 5 T. Dodd, M S., Miller, J J., Tyler, D J, whereby the effect of the paramagnetic relaxation agent gadolinium on the enhancement obtainable by DNP is shown to differ substantially between the 3.35 T preclinical used for this work and the clinical 5 T polariser located in OCMR.

Public Outreach

Perhaps rather unusually, a broad overview of the work detailed in Chapters 5 and 6 was described by the author to the general public on BBC Radio Oxfordshire, broadcast in 2014. Additionally, the author has described aspects of this work to several hundred school children, as part of the 'I'm a scientist, get me out of here!' 2014 national competition; at the Abingdon Science Festival, and at open-days and University outreach events.

Contents

Abstract i

Acknowledgements ii

Contents vii

List of Figures xi

1 Overview 1

2 Physical Principles of Nuclear Magnetic Resonance and Magnetic Resonance Imaging 3
2.1 Spin and Energy . . . 3
2.1.1 Spin . . . 3
2.1.2 Zeeman Splitting . . . 5
2.2 Larmor Precession . . . 7
2.2.1 Excitation . . . 9
2.2.2 Relaxation, Detection and Chemical Shift . . . 11
2.2.3 Non-Constant Fields . . . 14
2.3 Radiofrequency pulse design . . . 17
2.3.1 Excitation k-space . . . 17
2.3.2 The Shinnar-Le Roux Algorithm . . . 18
2.3.3 Slice Selection and Excitation Bandwidth . . . 21
2.4 Magnetic Resonance Imaging . . . 22
2.4.1 Frequency encoding . . . 22
2.4.2 Phase encoding . . . 23
2.4.3 Image k-space . . . 23
2.4.4 Pulse sequences . . . 26
2.4.5 Receiver bandwidth . . . 27
2.4.6 Echo Planar Imaging (EPI) . . . 28
2.4.7 Flow and Motion . . . 31
2.5 Signal Enhancement Methods . . . 34
2.5.1 Dynamic Nuclear Polarisation . . . 34
2.5.2 The Solid Effect . . . 35
2.5.3 Thermal Mixing . . . 36
2.6 Summary . . . 38

3 Hardware: Characterisation and Development 39
3.1 Instrumentation: A detailed overview . . . 39
3.1.1 The static field . . . 40
3.1.2 Gradient coils . . . 41

3.1.3 Gradient characterisation . . . 43
3.1.4 Radiofrequency wave generation, amplification and reception . . . 45
3.1.5 Logical organisation . . . 48
3.2 Radiofrequency coils . . . 52
3.2.1 Coil characterisation . . . 53
3.2.2 Rat Butterfly surface coil . . . 54
3.2.3 Mouse Loop surface coil . . . 54
3.2.4 Quadrature surface coil . . . 56
3.2.5 Two-channel surface-receive coil & dual-tuned Birdcage volume coil . . . 56
3.2.6 Rat Brain Butterfly surface coil . . . 59
3.3 Subject handling . . . 60
3.3.1 Position, monitoring & anaesthesia . . . 60
3.3.2 Cardiac gating . . . 61
3.4 Shimming . . . 64
3.5 Dynamic Nuclear Polarisation . . . 66
3.5.1 Experimental realisation . . . 66

4 Motivations and Preliminary work 71
4.1 Metabolism . . . 71
4.2 Cardiac disease . . . 74
4.2.1 Heart failure . . . 74
4.2.2 Myocardial ischæmia, infarction, and metabolic remodelling . . . 76
4.2.3 Cardiac disease and hyperpolarised [1-^{13}C]pyruvate . . . 77
4.3 Metastatic cancer . . . 78
4.3.1 Metabolism and Cancer . . . 78
4.3.2 DNP and Cancer . . . 80
4.3.3 Pilot Study . . . 83
4.3.4 Methods . . . 85
4.3.5 Results . . . 88
4.3.6 Discussion & Conclusions . . . 91

5 Development of a Multiple-Slice Spectral-Spatial Planar 2D EPI Pulse Sequence (SpSp-EPI) 95
5.1 Introduction . . . 95
5.2 Theory of Spectral-Spatial Pulses . . . 96
5.2.1 Spectral-Spatial Excitation k-space . . . 97
5.2.2 Oscillating Gradients and $S(\mathbf{k})$ . . . 97
5.2.3 $W(\mathbf{k})$ and $B_1(t)$ . . . 99
5.3 Development & Design of a Multiple-Slice Spectral-Spatial Pulse . . . 103
5.3.1 Numerical design . . . 103
5.3.2 Pulse parameters . . . 105
5.4 Implementation & Results . . . 108
5.4.1 Implementation . . . 108
5.4.2 Results . . . 111
5.5 Discussion . . . 113
5.5.1 Centre frequency shifts . . . 113
5.5.2 Delivered flip angle . . . 114
5.5.3 Resolution requirements . . . 114
5.5.4 Evidence of distorted excitation . . . 115
5.6 Concluding remarks . . . 117

6 Development of a Three-Dimensional Spectral-Spatial Imaging Sequence 119

6.1 Overview . . . 119
6.2 Methods . . . 120
6.2.1 Spectral-Spatial Pulse Design . . . 120
6.2.2 Pulse Sequence . . . 122
6.2.3 Hardware . . . 123
6.2.4 Frequency correction . . . 124
6.2.5 Reconstruction . . . 126
6.2.6 Animal Handling and Dosing . . . 130
6.3 Validation Study Results . . . 132
6.4 Validation Study Discussion . . . 138
6.4.1 Overview . . . 138
6.4.2 Flip Angle and SNR Efficiency . . . 140
6.4.3 Scan Timing . . . 141
6.4.4 Limitations . . . 142
6.5 Myocardial Infarction . . . 143
6.5.1 Overview . . . 143
6.5.2 Results . . . 145
6.5.3 Discussion . . . 147
6.6 Conclusions . . . 149

7 Brain Metastases **151**
7.1 Methods . . . 151
7.1.1 Animals . . . 151
7.1.2 MR Imaging protocol . . . 152
7.1.3 Reconstruction and Data Analysis . . . 152
7.1.4 Statistical Analysis . . . 156
7.1.5 Histology . . . 156
7.1.6 Cell Experiments . . . 156
7.2 Results . . . 157
7.2.1 *In Vivo* Study . . . 157
7.2.2 Cell Experiments . . . 161
7.2.3 Distortion correction . . . 161
7.3 Discussion . . . 165
7.3.1 Animal model . . . 165
7.3.2 Physiological interpretation of low SNR . . . 166
7.3.3 Disease progression . . . 168
7.4 Conclusion and Future Directions . . . 168

8 Hyperpolarised Urea and Cardiac Perfusion Imaging **171**
8.1 Introduction . . . 171
8.1.1 Overview . . . 171
8.1.2 Hibernating myocardium . . . 171
8.1.3 Visualising perfusion . . . 172
8.2 Methods . . . 176
8.2.1 Pulse Sequence . . . 176
8.2.2 Phantom Study . . . 180
8.2.3 Urea Hyperpolarisation . . . 181
8.2.4 Animal Handling . . . 182
8.2.5 Data Acquisition . . . 183
8.2.6 Image Reconstruction . . . 184
8.2.7 Data Analysis . . . 185

8.3 Results . . . 185
8.3.1 Phantom Study . . . 185
8.3.2 In Vivo Study . . . 187
8.4 Discussion . . . 191
8.4.1 Importance . . . 191
8.4.2 Physiological Interpretation of Results . . . 191
8.4.3 Comparison to Contrast-Enhanced Perfusion Measurements . . . 192
8.4.4 Other Hyperpolarised Perfusion Strategies . . . 193
8.4.5 Artefacts and Limitations . . . 194
8.4.6 Alternative Methods for Reducing the Blood Pool Signal . . . 196
8.4.7 Sequence Improvements . . . 196
8.4.8 Quantitative Measurements of Absolute Perfusion . . . 197
8.5 Conclusions . . . 197

9 Dealing with DNP Data: Uncertainty, Metabolic Modelling, and Spectral Analysis 199
9.1 Spectral Peak Quantification . . . 200
9.2 Propagating Uncertainty . . . 203
9.2.1 Introduction . . . 204
9.2.2 Methods . . . 204
9.2.3 Data Analysis . . . 210
9.2.4 Results . . . 212
9.2.5 Discussion . . . 212
9.3 The Fast Padé Transform in Hyperpolarised Spectroscopy . . . 213
9.3.1 Introduction . . . 213
9.3.2 Theory & Methods . . . 214
9.3.3 Results . . . 217
9.3.4 Discussion & Conclusions . . . 219
9.4 Fitting Models to Metabolism . . . 220
9.4.1 Introduction and Overview . . . 220
9.4.2 Zierhut's Metabolic Model . . . 221
9.4.3 Hyperpolarised Image Bayesian Analysis . . . 222
9.4.4 Example results . . . 225
9.4.5 Discussion and Future Directions . . . 225

10 Conclusions, Limitations, and Future Work 229

Bibliography 235

List of Figures

2.1 Projections of angular momentum on the z-axis . . . 4
2.2 Energy levels and populations for a spin-$^1/_2$ system in a magnetic field . . . 7
2.3 Larmor precession . . . 9
2.4 A cartoon of the effect of an RF pulse . . . 10
2.5 T1 and T2* relaxation . . . 13
2.6 The simplest NMR experiment . . . 15
2.7 The effect of a gradient. . . . 16
2.8 Sinc, Gaussian and Hard RF pulses . . . 21
2.9 The effect of frequency-encoding gradients . . . 22
2.10 The effect of a prephasing gradient . . . 23
2.11 The effect of phase encoding gradients. . . . 24
2.12 Cartesian k-space trajectory. . . . 26
2.13 Example k-space trajectories . . . 27
2.14 The Gradient Echo pulse sequence . . . 28
2.15 An EPI acquisition . . . 29
2.16 An illustration of the Nyquist ghost . . . 31
2.17 Nuclear and electronic P. . . . 35
2.18 Energy levels in an isolate spin-$^1/_2$ electron-nucleon system. . . . 36
2.19 The nuclear processes that contribute to thermal mixing . . . 37

3.1 Block diagram of the scanner used. . . . 40
3.2 A cartoon of a gradient set . . . 42
3.3 A photograph of the gradient coils. . . . 42
3.4 GIRF pulse sequence . . . 44
3.5 The measured GIRF. . . . 46
3.6 RF detection overview . . . 50
3.7 The logical organisation of the 7 T scanner used. . . . 51
3.8 A surface coil . . . 52
3.9 A birdcage coil . . . 53
3.10 The characterised spatial response of several RF coils . . . 55
3.11 Bad NMR spectra acquired from a broken RF Coil, and those same spectra when fixed. . . . 57
3.12 The designed (**A**, **B**) and constructed (**C**-**E**) coil housings . . . 58
3.13 A home-built brain surface transmit/receive coil . . . 60
3.14 The designed timestamper board . . . 62
3.15 The autoshim user interface . . . 66
3.16 The structure of the OXO63 radical. . . . 67
3.17 The 'alpha system' hyperpolariser. . . . 69

4.1 An overview of the metabolic reactions inside eukaryotic cells. . . . 73

4.2 A cartoon of [1-^{13}C]pyruvate metabolism . . . 74
4.3 The processes behind PET . . . 81
4.4 The progression of brain metastases . . . 84
4.5 AMARES fits . . . 86
4.6 Kinetic turnover of metabolites following the injection of hypearpolarised [1-^{13}C]pyruvate. 87
4.7 Pre- and post-gadolinium contrast gradient echo images of brain tumours. . . . 89
4.8 Lactate production over time . . . 89
4.10 Tumour volume vs. KLacPyr . . . 90
4.9 Tumour volumes as a function of time. . . . 90
4.11 Histological immunohistochemistry stained slices of brain. . . . 91
4.12 An illustration of the insufficient spatial penetration of the coil used. . . . 93

5.1 The sampling pattern taken in spectral-spatial k-space . . . 98
5.2 $J_n(x)$, with $n = 1 \ldots 5$. . . . 99
5.3 A Gaussian-Gaussian SpSp Pulse . . . 102
5.4 Sampling bias in k-space . . . 104
5.5 The Spectral-Spatial pulse used with SpSp-EPI . . . 106
5.6 The measured properties of the SpSp pulse described above. . . . 107
5.7 Pulse sequence diagram for SpSp-EPIP . . . 109
5.8 Inital (disappointing) results from multi-slice SpSp-EPI . . . 110
5.9 Phantom experiment . . . 111
5.10 Enhanced-readout 2D multiple slice images. . . . 112
5.11 Representative Epip-SpSp images of the brain. . . . 113
5.12 Timecourse data proving errors in selective excitation . . . 116

6.1 The revised slab-excitation pulse designed for 3D spectral-spatial imaging . . . 121
6.2 The pulse sequence used. . . . 122
6.3 The flip angle scheme used with this sequence . . . 123
6.4 A comparison of image registration metrics . . . 127
6.5 Mutual Information and Image Self-Similarity as metrics . . . 128
6.6 A blip up/down 3D-EPI image of a ball . . . 129
6.7 The developed sequence imaging a syringe of pyruvate . . . 131
6.8 Mid-ventricular stills of the hyperpolarised dataset . . . 133
6.9 Hyperpolarised bicarbonate myocardial wall visibility classification . . . 133
6.10 The ratios of maximal pyruvate/bicarbonate in three different groups of animals . . . 134
6.11 Maximum intensity projections of reconstructed cardiac images with this sequence. . . . 135
6.12 Reconstructed myocardial time-courses . . . 137
6.13 The distribution of TRs . . . 137
6.14 A plot showing the SNR efficiency of 2D vs 3D hyperpolarised scans . . . 140
6.15 An illustration that cardiac gating is still required . . . 142
6.16 Ischæmia reperfusion injury . . . 144
6.17 The reason for a failed study . . . 145
6.18 Data from the MI study . . . 146
6.19 Histology and MR correlation . . . 148

7.1 A summary of the scan protocol timing . . . 153
7.2 An approximate matrix square root . . . 153
7.3 Further image registration attempts . . . 155
7.4 Representative pyruvate imaging of the brain . . . 159
7.5 Histological evidence of disease . . . 160
7.6 The timecourse of lac/pyr . . . 160

7.7 Brain lactate/pyruvate ratio images . . . 160
7.9 Cells in culture . . . 161
7.8 Longitudinal imaging of lactate production in disease . . . 162
7.10 Apparent glycolysis rates of tumour cells . . . 163
7.11 The effect of the distortion correction pipeline . . . 164

8.1 X-Ray Angiography . . . 172
8.2 CT Angiography . . . 173
8.3 Late Gadolinium Enhancement . . . 174
8.4 "4D" Flow MRI . . . 175
8.5 Example spiral trajectory . . . 177
8.6 The urea flow-sensitised sequence . . . 178
8.7 An example of golden angle sampling . . . 179
8.8 An example *in vivo* B_0 map . . . 179
8.9 Hyperpolarised ^{13}C-urea Build-Up Curve . . . 181
8.11 GIRF-predicted spiral waveforms . . . 184
8.10 The scan starting hardware . . . 185
8.12 Flow in a pipe . . . 186
8.13 Longitudinal flow in a pipe . . . 186
8.14 Proton flow-suppressed images . . . 188
8.15 Hyperpolarised ^{13}C-urea perfusing the heart at rest . . . 189
8.16 Hyperpolarised ^{13}C-urea perfusing the heart under adenosine stress . . . 190
8.17 The ratio of integrals of perfusion curves under rest and stress . . . 190
8.18 Co-polarisation build-up curve of ^{13}C-urea and [1-^{13}C]pyruvate . . . 193

9.1 The PDF of a ratio . . . 207
9.2 An overview of dealing with uncertainty . . . 211
9.3 A Padé Approximant . . . 215
9.4 A comparison of reconstruction methods of truncated data . . . 217
9.5 The FPT on truncated simulated data . . . 217
9.6 FPT reconstruction of *in vivo* spectroscopic data . . . 218
9.7 The results of Bayesian metabolic fitting . . . 226

10.1 The number of pyruvate publications as a function of time . . . 232

Chapter 1

Overview

Life is odd. It appears to create order, in a universe where disorder always increases. The set of chemical reactions that energetically power life, collectively called metabolism, are complex, but comparatively well understood. Metabolic dysregulation is a hallmark of many different pathologies, for example cancer and cardiovascular disease, which combined were responsible for $\sim 60\%$ of the deaths occurring in the United Kingdom in 2013. Specifically, a shift towards anaerobic glycolysis is a definition of cancer,[1] whereas the relative rates of fatty acid oxidation and glycolysis are altered in heart failure.[2] Understanding the mechanistic links between metabolic derangement and pathology continues to form an active area of research, but the established role of metabolism in disease is exploited clinically, primarily by Positron Emission Tomography (PET) with radioactive analogues of small metabolites. PET and related techniques have undoubtedly saved many lives, but they are subject to ultimate limitations in their resolution, in their speed of acquisition, in the radioactive dose they deliver, and in their ability to look at dynamic processes.

This thesis focusses on the development and *in vivo* validation of MRI pulse sequences that, together with dissolution Dynamic Nuclear Polarisation (DNP), form a new technique for metabolic imaging. Dissolution DNP is a technique for pumping $I \neq 0$ nuclear spins into a given state under cryogenic conditions (~ 1 K), before rapidly melting them with superheated (~ 450 K) solvent. By doing so, a "hyperpolarised" liquid is produced that instantaneously possesses the spin population difference created at low temperature, exponentially relaxing back to thermal equilibrium. The process creates a $>$10 000-fold increase in the signal-to-noise ratio of Magnetic Resonance experiments that probe this spin structure. By hyperpolarising molecules of metabolic interest, before injecting them into an organism located within a nearby Magnet Resonance Imaging (MRI) system, it is possible to follow the labelled molecule through subsequent biochemical reactions in space and across time. The main development presented in this work is a robust, three-dimensional imaging strategy that allows for the metabolism of hyperpolarised [1-^{13}C]pyruvate to be imaged with a spatial and temporal resolution of $1 \times 1 \times 2$ mm^3 and ~ 1.8 s in a large field of view. The technique is validated in the healthy rodent heart, and is applied in an attempt to detect cancer that has metastasised to the brain earlier than other methods. Additionally, novel methods of mathematical analysis for the five-dimensional data

the technique provides are presented, and the degree of myocardial perfusion in the rodent heart is quantified rapidly and directly using hyperpolarised urea.

The fundamental physics of magnetic resonance, MRI, and DNP are discussed in Chapter 2, before the hardware created and used in this work is described and characterised in Chapter 3. The need for metabolic *imaging*, as opposed to NMR spectroscopy, is emphasised by a pilot study on brain metastases in Chapter 4. A partially unsuccessful attempt at high-resolution metabolic imaging with two-dimensional spectral-spatial EPI is detailed in Chapter 5 and its shortcomings lead to the development of a 3D sequence, described in Chapter 6. An attempt was made to use the developed sequence in the early detection of cancer that has metastasised to the brain in a study detailed in Chapter 7. A successful side-project was undertaken to assess the viability of probing myocardial perfusion with hyperpolarised ^{13}C-urea, and is described in Chapter 8. Data of the form that these methods provide is highly multidimensional, and of varying quality. Chapter 9 briefly details three related techniques to interpret it. After a review of the NMR spectral peak quantification methods used in this work, a simple method to incorporate spectroscopic data of varying quality is proposed in Section 9.2. The use of the Fast Padé Transform is then investigated in Section 9.3 as a means to potentially accelerate the acquisition of hyperpolarised spectra by providing a more robust alternative to zero filling. Additionally, in order to reduce the dimensionality and improve the utility of the datasets generated by these techniques, metabolic modelling with Bayesian inference was performed to produce bounded, quantitative maps of biochemical rate constants of interest, and are detailed in Section 9.4.

Finally, the future benefit and limitations of this work are discussed in Chapter 10. DNP is a technique that is in the process of transitioning to clinical practice: it is hoped that many of the methods developed here are of wide applicability and will find routine clinical use.

Chapter 2

Physical Principles of Nuclear Magnetic Resonance and Magnetic Resonance Imaging

2.1 | Spin and Energy

2.1.1 Spin

In our universe, it is an experimental fact that fundamental particles possess an intrinsic form of angular momentum known as spin. In the Schrödinger equation formulation of non-relativistic quantum mechanics, the spin operator, $\mathbf{S}$, is defined as the generator for changes of orientation that are not accompanied by any changes of motion of the system as a whole, namely

$$\mathbf{S} = \mathbf{J} - \mathbf{L}, \tag{2.1}$$

where $\mathbf{J}$ is the operator that generates rotations of the system, and $\mathbf{L}$ is the operator that generates its displacement. Spin is constant for a given particle, always conserved, and its projections form good quantum numbers, i.e. they are time-independent eigenvalues of wavefunctions when acted on by physical Hamiltonians. The operator $S^2 = \mathbf{S} \cdot \mathbf{S}$ can be defined and shown to commute with $\mathbf{S}$. It is therefore appropriate to describe the system in terms of the eigenvalues of $\mathbf{S}$ and S^2. It can be shown that the eigenvalues of S^2 are $s(s+1)$, where s is either an integer or a half integer.[3, Ch. 6] Likewise, the eigenvalues of the projection of $\mathbf{S}$ along an axis, conventionally chosen to be the z axis, are quantised, and can only take the values $S_z = \hbar\{-s,\ -(s-1),\ \ldots,\ s-1,\ s\}$. The z axis is chosen as the quantisation axis because it is the polar axis in spherical coordinates, and states about other axes can be written in terms of S_z through the commutation relations for S_z. The quantity $2s+1$ represents the total number of states available to a spin-s particle, and is referred to as its multiplicity; for a spin-$1/2$ particle, the resulting two states are commonly referred to as "up" ($S_z = 1/2$) and "down" ($S_z = -1/2$) respectively.

Systems containing more than one particle with spin have a resulting intrinsic angular momentum that is not simply the sum of its constituents, but rather a set of possible quantised values that can be derived from considering the commutation relations of the angular momentum operators.[3, Ch. 7] For a system as complex as a nucleus, which is subject to many further physical constraints, the resulting

A	Z	N	I	Example
Even	Even	Even	0	${}^{12}_{6}\mathrm{C}_{6}$
	Odd	Odd	Integer	${}^{2}_{1}\mathrm{H}_{1}$
Odd	Even	Odd	Half-integer	${}^{13}_{6}\mathrm{C}_{7}$
	Odd	Even	Half-integer	${}^{15}_{7}\mathrm{N}_{8}$

Table 2.1: A summary of the net effect of the constraints imposed by the quantum-mechanical addition of angular momenta that aid when determining the possible nuclear spin, I, for the isotope ${}^{A}_{Z}\mathrm{X}_{N}$, where A is the atomic mass number, Z the atomic number, and N the neutron number.

total nuclear spin I is usually considered an intrinsic property of the system. However, these origins manifest themselves in the set of rules that govern nuclear spin, summarised in Table 2.1. By convention, the symbols I and its projections m are used to refer to nuclei, and S and *its* projections m to electrons. Correspondingly, a nucleus at rest has overall orbital angular momentum, J defined by $\mathbf{J} = \hbar\mathbf{I}$, with eigenvalues $J = \hbar\sqrt{I(I+1)}$ whose projection on the z axis are $J_z = m\hbar$, where m is often called the 'magnetic quantum number', and is illustrated in Fig. 2.1. Canonically, the operator I_z is then defined to have eigenvalues m; namely $I_z|\psi; m\rangle = m|\psi; m\rangle$, and it is found that m is spanned by the $2I+1$ values $-I, \ldots, I$, as above.

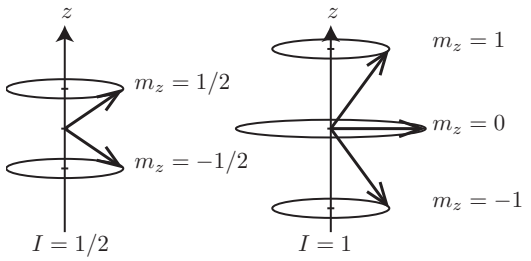

Figure 2.1: The standard vector model depiction of angular momentum, **J**, and its projection m_z on the z axis, for a spin-$1/2$ and spin-1 nucleus

An intrinsic magnetic moment $\boldsymbol{\mu}$ is associated with this intrinsic angular momentum. It can be shown that its projection on the z axis, μ_z, is of magnitude

$$\mu_z = \frac{e}{2m_p} g_N \hbar I_z \tag{2.2}$$

$$\stackrel{\text{def}}{=} \gamma \hbar I_z \tag{2.3}$$

where g_N in (2.2) is the Landé g-factor, a number specific to a given particle that famously is about two for an electron, and notoriously difficult to calculate even in that case.[1] For these reasons, the quantity γ is defined in terms of these fundamental constants

[1] Briefly, one derives $g_e = 1$ from classical (or "semi-classical") electrodynamics; $g_e = 2$ from the Lorentz invariant Klein-Gordon or Dirac equations, and one can calculate a value of $g_e \approx 2$ from local electron loop interactions with quantum electrodynamics (QED). It transpires that the value of $(g_e - 2)$ is a precision test of QED: at present, its predicted value is known with less certainty than that obtained by experiment, but the two agree to phenomenal accuracy with a relative error on the order of 10^{-12}. In QED, an order n process has a Feynman propagator with a premultiplying factor containing α^n, where the fine structure constant $\alpha \approx 1/137$ is the coupling constant that mediates the strength of each interaction occurring at a node on a Feynman diagram. High-order processes, with many virtual interactions, therefore contribute little to the overall result of the calculation as α^n is small. Correspondingly, most values of g_e are calculated up to order $n = 8$ to 10, as truncating the convergent infinite series is appropriate.[4] It is far more difficult to make accurate predictions for the value of a nuclear g factor as the coupling "constant" for quantum chromodynamics (QCD), denoted α_s, is found to be a "running coupling constant" meaning that it is a strong function of energy. At low energy one runs into the problem that α_s is effectively greater than one, and the contribution from infinite-order processes cannot be neglected. As a

(as in Eq. (2.3)), measured experimentally, and is known as the gyromagnetic (or magnetogyric) ratio, constant for a given nucleus.

2.1.2 Zeeman Splitting

If an $I \neq 0$ nucleus is placed in a magnetic field $\mathbf{B}$ then a perturbing Hamiltonian $\mathcal{H} = -\boldsymbol{\mu} \cdot \mathbf{B}$ is produced. For historical reasons, this is known as a Zeeman Hamiltonian. Taking $\mathbf{B}$ to be aligned along the z axis and of magnitude B, we see that

$$\mathcal{H} = -\hbar\gamma B I_z. \tag{2.4}$$

As $\mathcal{H}$ is small, this interaction gives rise to $2I + 1$ possible shifted energy states, of magnitude

$$E = -\gamma\hbar B_0 m, \qquad \text{where } m = -I,\ -(I-1),\ \ldots,\ I. \tag{2.5}$$

Correspondingly, transitions can occur between these states if and only if (hereafter shortened to "iff") they are of angular frequency $\hbar\omega = \Delta E$, and the matrix element for the transition is non-zero. In order to conserve the spin of an incident photon and the nucleus, the relevant matrix element is non-zero only for $m_{\text{final}} = m_{\text{initial}} \pm 1$. This selection rule therefore only permits transitions which are those adjacent in energy, irrespective of I, giving

$$\hbar\omega = \Delta E = \gamma\hbar B_0 \tag{2.6}$$

Rearranging slightly, we derive the Larmor equation, which expresses the condition for magnetic resonance:

$$\boxed{\omega = \gamma B}. \tag{2.7}$$

It should be emphasised that the Larmor equation is linear, and, despite the lack of $\hbar$, applies microscopically – we shall see that this gives rise to the great utility of NMR and MRI.

In NMR and MRI, the spins in question are nuclei with $I \neq 0$, commonly $^{1}\mathrm{H}$, $^{2}\mathrm{H}$, $^{3}\mathrm{He}$, $^{13}\mathrm{C}$, $^{15}\mathrm{N}$,$^{17}\mathrm{O}$, $^{19}\mathrm{F}$, $^{31}\mathrm{P}$ and $^{129,\,131}\mathrm{Xe}$, whereas in Electron Paramagnetic Resonance (EPR), they are unpaired electrons. It is usually the case that experiments are performed on macroscopic samples, in which many spins are considered. We therefore have to consider an ensemble of spins, in general in a superposition of their allowed states

$$|\psi\rangle = \sum_i a_i \,|\psi;\, m_i\rangle \tag{2.8}$$

consequence, it is exceptionally difficult to make an *ab initio* calculation for a hadronic g, and recent attempts at calculating g_p, g_n, and $g_{^3\mathrm{H}}$ with discreteised lattice QCD are unable to match experimental values within error, although they are within 20% of the measured values.[5]

where $|\psi_i; m\rangle$ denotes the set of i spins with quantum numbers m_i . Measurements on the system therefore act on the superposition; for example, the expectation value of I_z of the composite system is given by

$$\langle\psi|I_z|\psi\rangle = \sum_i |a_i|^2 m_i. \tag{2.9}$$

For a population of N nuclei at thermal equilibrium in a magnetic field B, and at a temperature T, thermodynamic considerations are prevalent. It is customary to assume that, in non-magnetic samples, the interactions between nuclei are weak, but that they are strong enough to allow the sample to reach thermodynamic equilibrium. If this were not the case, and spins interacted strongly with their nearest neighbours we would be considering a spin-coupled system, e.g. a ferri- or ferromagnet. Most matter in nature is not strongly spin-coupled, and we are justified in assuming that spins do not strongly interact. Correspondingly, given I states of energy E_i, the appropriate partition function is

$$Z = \sum_{i=1}^{I} e^{\frac{-E_i}{k_B T}} \tag{2.10}$$

where k_B is Boltzmann's constant. Hence, the given population of spins *at thermal equilibrium* in a state m is given by

$$N_m = N \frac{\exp\left(-\frac{E_m}{k_B T}\right)}{Z}. \tag{2.11}$$

This expression is illustrated for the common case where $I = 1/2$ in Fig. 2.2. Here, there are two states available to the nucleus: its magnetic moment can be parallel or anti-parallel to the applied field; commonly referred to as "down" and "up" respectively where the parallel orientation is of lower energy, and is called the ground state. When dealing with spin-$^1/_2$ systems, it is customary to define the thermal equilibrium polarisation P as

$$P = \frac{N_\uparrow - N_\downarrow}{N_\uparrow + N_\downarrow} = \tanh \frac{\gamma\hbar B}{2k_B T}, \tag{2.12}$$

where $N_\uparrow$ is the number of "up" spins, and $N_\downarrow$ the number of "down" spins. For systems at room temperature in a 7 T magnetic field, thermal noise dominates and P is on the order of 6×10^{-6}, indicating that six spins in a million are in excess in the ground state. A Taylor expansion of Eq. (2.12) reveals that P scales approximately linearly with B. As the total signal received in magnetic resonance experiments is proportional to γPC, where C is the concentration of the nucleus concerned, the low value of P means that NMR and MRI are usually considered inherently insensitive techniques.

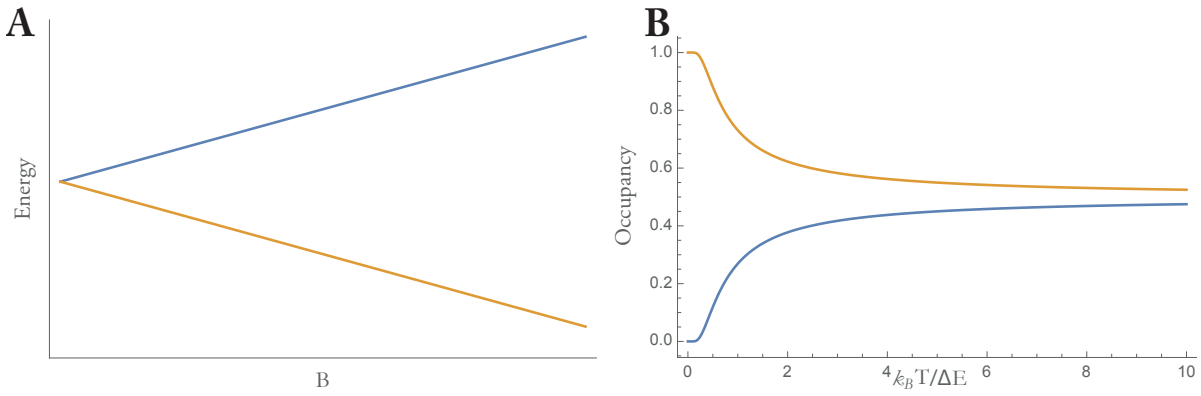

Figure 2.2: For a spin-$^1/_2$ system in a non-zero magnetic field, the two $m_z = \pm 1/2$ energy levels are split, with an energy directly proportional to the applied field, a fact referred to as 'Zeeman splitting' (**A**). The difference in energy between the lower (orange) and upper (blue) levels is $\Delta E = \hbar\omega = \hbar\gamma B$. Due to thermodynamic considerations the population in each level is nearly equal at temperatures commonly found on earth (**B**); at 7 T, $P = 0.5$ only at 2 mK.

2.2 | *Larmor Precession*

Irrespective of its origins, this small population difference of spins generates a magnetisation. In a spin-$^1/_2$ system, the magnetic moment of a spin in a given state is $\mu = \pm {}^1/_2 \hbar\gamma$, and therefore the overall magnetisation of an ensemble is therefore simply $M = {}^1/_2 N\hbar\gamma(N_\downarrow - N_\uparrow)$, where $N = N_\uparrow + N_\downarrow$, and $\mathbf{M}$ is orientated parallel to B. For a sample at thermal equilibrium at room temperature, we are justified in taking the Taylor series expansion for Eq. (2.12), and find that

$$\begin{aligned} \mathbf{M} &= N\left(\frac{\hbar\gamma}{2}\right)\tanh\left(\frac{\hbar\gamma B}{2k_BT}\right)\mathbf{k} \\ &\approx N\left(\frac{\hbar\gamma}{2}\right)\left(\frac{\hbar\gamma B}{2k_BT}\right)\mathbf{k} \qquad (2.13) \\ &\approx N\left(\frac{\hbar\gamma}{2}\right)^2\frac{B}{k_BT}\mathbf{k}. \qquad (2.14) \end{aligned}$$

We are now able to consider the behaviour of this net magnetisation in a magnetic field. The classical electrodynamic description of events is that $\mathbf{M}$ is subject to a torque given by $\mathbf{M} \times \mathbf{B}$, which is equal to the rate of change of angular momentum, $\mathbf{J}$. Classically, $\mathbf{J}$ is given by $\mathbf{M}/\gamma$. One therefore finds that

$$\boxed{\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \gamma(\mathbf{M} \times \mathbf{B})}. \qquad (2.15)$$

It is also possible to derive Eq. (2.15) from the Hamiltonian stated in Eq. (2.4) directly. Recall the commutation relations for angular momentum, namely [3, Ch. 6]

$$[I_x,\, I_y] = iI_z. \qquad (2.16)$$

Then, using Ehrenfest's theorem[2] we can calculate the time evolution of the expectation values of angular momentum, and find that

$$\begin{aligned}\frac{\mathrm{d}\langle I_x\rangle}{\mathrm{d}t} &= \frac{i}{\hbar}[\mathcal{H},\, I_x] \\ &= -\gamma B i[I_z,\, I_x] \\ &= \gamma B I_y,\end{aligned} \tag{2.17}$$

and similarly

$$\begin{aligned}\frac{\mathrm{d}\langle I_y\rangle}{\mathrm{d}t} &= -\gamma B I_x \\ \frac{\mathrm{d}\langle I_z\rangle}{\mathrm{d}t} &= 0.\end{aligned} \tag{2.18}$$

These are components of the vector equation

$$\frac{\mathrm{d}\langle \mathbf{I}\rangle}{\mathrm{d}t} = \gamma\langle \mathbf{I}\rangle \times \mathbf{B}, \tag{2.19}$$

and, since $\boldsymbol{\mu} = \gamma\hbar\mathbf{I}$, we find that the time evolution of the expectation value of the magnetic moment of the nucleon is

$$\frac{\mathrm{d}\langle \boldsymbol{\mu}\rangle}{\mathrm{d}t} = \gamma\langle \boldsymbol{\mu}\rangle \times \mathbf{B}. \tag{2.20}$$

Assuming that each spin does not interact with its neighbour, it is straightforward to show that the overall magnetisation is just the vector sum of magnetic moments, and Eq. (2.20) generalises to Eq. (2.15). It is important to note that the above makes no assumption about the behaviour of $\mathbf{B}$; namely that it still holds true if $\mathbf{B}$ is a function of time. Analogous forms of Eq. (2.15) for electronic behaviour under similar conditions can also be derived from quantum electrodynamics.[6] As the same picture evolves from classical mechanics, Schrödinger equations, and quantum electrodynamics, we can be confident that Eq. (2.15) is an appropriate description for the time evolution of magnetisation, provided that spins do not interact.

Let us now solve Eq. (2.15) in a Cartesian coordinate system xyz with unit vectors $\mathbf{ijk}$ respectively. Taking $\mathbf{B} = B_0\mathbf{k}$, with an additional subscript "$_0$" to differentiate it from other fields that we might subsequently apply, we find that

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = \gamma M_y B_0; \qquad \frac{\mathrm{d}M_y}{\mathrm{d}t} = -\gamma M_x B_0; \qquad \frac{\mathrm{d}M_z}{\mathrm{d}t} = 0 \tag{2.21}$$

[2] Ehrenfest's theorem is a statement about the evolution of expectation values. Consider a Hermitian operator Q, which is not a function of time. Then, the time-dependent Schrödinger equation $i\hbar\frac{\partial|\psi\rangle}{\partial t} = \mathcal{H}|\psi\rangle$ will tell us the time evolution of its expectation value: left-multiplying by $\langle\psi|Q$ and re-arranging, we deduce that $i\hbar\frac{\partial}{\partial t}\langle\psi|Q|\psi\rangle = \langle\psi|[Q,\,\mathcal{H}]|\psi\rangle + i\hbar\langle\psi|\frac{\partial Q}{\partial t}|\psi\rangle$. Therefore, $\frac{\partial\langle Q\rangle}{\partial t} = -i[Q,\,\mathcal{H}]/\hbar = i[\mathcal{H}, Q]/\hbar$, which is the result used.

whose solutions are

$$\begin{aligned} M_x(t) &= M_x(0)\cos(\gamma B_0 t) + M_y(0)\sin(\gamma B_0 t) \\ M_y(t) &= -M_x(0)\sin(\gamma B_0 t) + M_y(0)\cos(\gamma B_0 t) \\ M_z(t) &= M_z(0). \end{aligned} \tag{2.22}$$

This leads us naturally to define $\omega = \gamma B_0$, and realise that the solution of Eq. (2.15) is the Larmor precession of $\mathbf{M}$ at a frequency ω, schematically illustrated in Fig. 2.3. Naturally, this is exactly the same resonance condition as derived above by considering selection rules whilst conserving energy. It is worth noting that ω usually corresponds to radiofrequency photons for nuclei, and microwave for electrons. At 7 T, $2\pi\omega_{^1\mathrm{H}} \approx 300$ MHz, and $2\pi\omega_{^{13}\mathrm{C}} \approx 75.5$ MHz.

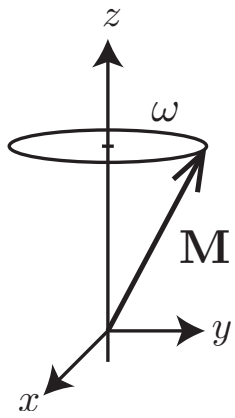

Figure 2.3: When placed in a magnetic field $B_0\mathbf{k}$, $\mathbf{M}$ precesses about z at a frequency ω.

2.2.1 *Excitation*

Let us now consider the action of applying radiofrequency photons to this system, described in terms of the B field of light, denoted $\mathbf{B_1}$. Practical, $\mathbf{B_1}$ is generated by stationary electrical devices in the experiment, referred to as RF coils. It is assumed that $\mathbf{B_1}$ lies entirely in the xy plane; this assumption can be made without loss of generality by making the substitution $B_0 \to B_0\mathbf{k} + (\mathbf{B_1})_z\mathbf{k}$ and simply redefining what the initial static field is. Consider the field applied exactly at resonance, so that $\mathbf{B_1}$ oscillates in the xy plane at ω, namely

$$\mathbf{B_1} = B_1\cos\omega t\mathbf{i} - B_1\sin\omega t\mathbf{j}. \tag{2.23}$$

Then, we find that Eq. (2.15) reduces to

$$\begin{aligned} \frac{\mathrm{d}M_x}{\mathrm{d}t} &= \gamma\left(M_y B_0 + M_x B_1\sin\omega t\right) \\ \frac{\mathrm{d}M_y}{\mathrm{d}t} &= \gamma\left(M_z B_1\cos\omega t - M_x B_0\right) \\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= \gamma\left(-M_x B_1\sin\omega t - M_y B_1\cos\omega t\right). \end{aligned} \tag{2.24}$$

Given the initial condition $\mathbf{M}(t=0) = M_0\mathbf{k}$, we solve Eq. (2.24) to find

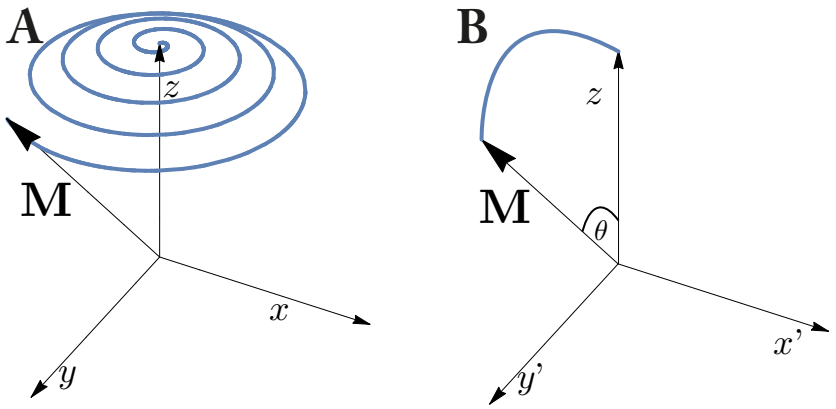

Figure 2.4: Applying RF to the system results in the precession of **M** about resultant effective field. If the RF is applied on resonance, as shown here, the net effect is precession about the rotating x' axis into the $x'y'$ plane by an angle $\theta(\tau)$, shown both in the stationary (**A**) and rotating (**B**) frames of reference.

$$\begin{aligned} M_x(t) &= M_0 \sin(\gamma B_1 t) \sin \omega t \\ M_y(t) &= M_0 \sin(\gamma B_1 t) \cos \omega t \\ M_z(t) &= M_0 \cos(\gamma B_1 t)\,. \end{aligned} \tag{2.25}$$

This therefore implies that **M** precesses simultaneously about the two fields; at ω about B_0, and at $\omega_a \stackrel{\text{def}}{=} \gamma B_1$ about B_1. This precession is illustrated in Fig. 2.4**A**. Interestingly, this result can be deduced by considering the action of a Hamiltonian $\boldsymbol{\mu} \cdot \mathbf{B}$ and finding the time evolution of $\langle \boldsymbol{\mu} \rangle$; this derivation is omitted here for the sake of brevity. As $\omega \gg \omega_a$, and B_0 is always present, it is convenient to view this process in the rotating frame of reference, in which the $x'y'$ plane rotates at ω. After some manipulation, it is revealed that this is equivalent[7, pp. 26-28] to defining an effective field during the time B_1 is "on", where the system evolves according to

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \gamma \mathbf{M} \times \mathbf{B}_{\text{eff}} \qquad \text{where } \mathbf{B}_{\text{eff}} \stackrel{\text{def}}{=} \left(B_0 - \frac{\omega_a}{\gamma} \right) \mathbf{k} + B_1 \mathbf{i}' \tag{2.26}$$

where $\mathbf{i}'$ and $\mathbf{j}'$ are unit vectors in the rotating frame.

If $\omega_a = \gamma B_0$, i.e. the RF is applied at resonance, then the effective field serves only to rotate **M** about the x' axis at an angular frequency $\omega_1 = \gamma B_1$. The most common way to perform NMR experiments is to apply a short bust of RF for a period of time τ. This serves to rotate **M** through an angle θ defined by

$$\boxed{\theta = \gamma B_1 \tau} \tag{2.27}$$

known commonly as the flip (or tip) angle, and is illustrated in Fig. 2.4**B**. For example, a $\theta = 90°$ pulse serves to rotate **M** fully into the transverse plane. Note that this description of a rotating B_1 field is equivalently expressed as the interaction of a circularly polarised field, B_1^+ with the nuclear spin system, where

$$B_1^+ \stackrel{\text{def}}{=} \frac{\mathbf{B}\cdot\mathbf{i} + i\mathbf{B}\cdot\mathbf{j}}{2} = \frac{\mathbf{B_1}_x + i\mathbf{B_1}_y}{2}, \tag{2.28}$$

if $\mathbf{B_1}$ is defined as in Eq. (2.23). B_1^+ can be shown to be the component of $\mathbf{B}$ that directly couples to the spin system, through the Pauli equation, and represents the generated field by a nearby RF coil which is responsible for the magnetic resonance phenomenon.[8] By the principle of reciprocity, the contra-rotating circularly polarised component, B_1^- defined as

$$B_1^- \stackrel{\text{def}}{=} \frac{(\mathbf{B}\cdot\mathbf{i} - i\mathbf{B}\cdot\mathbf{j})^\dagger}{2} = \frac{(\mathbf{B_1}_x - i\mathbf{B_1}_y)^\dagger}{2}, \tag{2.29}$$

corresponds to the oscillating current induced in the coil by the time varying magnetisation produced by the evolving nuclear spin system. Ideally, B_1^+ and B_1^- are the same, but electrodynamic interactions serve to violate this assumption. At high static fields the discrepancy between the two grows, and becomes a significant concern by the time $\gamma B_1 \approx 1\,\text{GHz}$ or if linear dimensions cease being small compared to ω_0.[9]

Unfortunately, the microscopic picture of the effect of an RF pulse is more complicated, as coherences of populations of states are formed that evolve in time. The most appropriate formalism to consider this problem is that of the density matrix, σ, which describes the evolution of the ensemble in terms of individual spin-up and spin-down states. It transpires that σ takes a simple form in terms of the Pauli matricies for an isolated spin-half system, namely $\sigma = \frac{1}{2}(\mathbf{I} + P\sigma_z)$ (although as the majority of spins in the sample do not participate in the radiative transition, giving rise to the identity in the above, this expression is approximately $1/2P\sigma_z$); and also that the action of a 90° pulse is essentially to form the superposition of states $\frac{i}{\sqrt{2}}(|\uparrow\rangle - |\downarrow\rangle)$.[10, pp. 38-40] Performing detailed calculations with Pauli matricies is tedious, and, for the most part unnecessary: the vector model for $\mathbf{M}$ is sufficient for all pulse sequences that do not involve coupled spins, such as multinuclear decoupling, or J-edited spectroscopy. In situations where local interactions cannot be neglected, it is necessary to consider the action of a set of RF pulses not in terms of the behaviour of $\mathbf{M}$, but rather as a set of sequential operators acting on the initial states.

2.2.2 Relaxation, Detection and Chemical Shift

2.2.2.1 Relaxation

As noted previously in the above discussion, Eq. (2.15) does not account for weak interactions between spins at all, be they spin-spin or spin-lattice interactions. It is reasonable to expect that this is a significant oversight as these processes serve to let the system maintain thermodynamic equilibrium, allowing it to exchange energy with itself and with its surroundings. It was found empirically by Bloch *et al.*[11] that the spin-lattice interaction relaxes $\mathbf{M}$ to its thermal equilibrium M_0 value exponentially with a characteristic time T_1, and that spin-spin interactions serve to reduce the components of $\mathbf{M}$ not parallel

to $\mathbf{B}$ exponentially, with a characteristic time T_2. This introduces anisotropy to Equation (2.15), resulting in the Bloch equations proper:

$$\boxed{\begin{aligned}\frac{\mathrm{d}\,M_z}{\mathrm{d}t} &= \gamma(\mathbf{M}\times\mathbf{B})_z + \frac{M_0 - M_z}{T_1}\\ \frac{\mathrm{d}\,M_x}{\mathrm{d}t} &= \gamma(\mathbf{M}\times\mathbf{B})_x - \frac{M_x}{T_2}\\ \frac{\mathrm{d}\,M_y}{\mathrm{d}t} &= \gamma(\mathbf{M}\times\mathbf{B})_y - \frac{M_y}{T_2}\end{aligned}} \tag{2.30}$$

which describe the evolution of magnetisation in a medium as a function of field. The Bloch equations (Eq. (2.30)) cannot be solved analytically with the additional relaxation terms, and analytically predicting the values of T_1 and T_2 for a given system is exceptionally difficult, as they effectively form the resultant effect of a variety of complex microscopic processes. Crucially, these values are different for spins in different environments, and therefore form the basis of contrast in traditional ^{1}H MRI.

Consequently, Eq. (2.30) are usually solved numerically, using measured or estimated values of T_1 and T_2. However, it is informative to illustrate that these additional terms still lead to Larmor precession with relaxation. In the rotating reference frame in the presence of an RF pulse, Eq. (2.30) reduce to

$$\begin{aligned}\frac{\mathrm{d}M_{x'}}{\mathrm{d}t} &= \gamma M_{y'}\left(B_0 - \frac{\omega}{\gamma}\right) - \frac{M_{x'}}{T_2}\\ \frac{\mathrm{d}M_{y'}}{\mathrm{d}t} &= \gamma\left(M_{x'}B_1 - M_{x'}\left(B_0 - \frac{\omega}{\gamma}\right)\right) - \frac{M_{y'}}{T_2}\\ \frac{\mathrm{d}M_z}{\mathrm{d}t} &= -\gamma M_{y'}B_1 - \frac{M_z - M_0}{T_1}\end{aligned} \tag{2.31}$$

and the initial condition $M_x(0) = 0$, $M_y(0) = M_0\sin\theta$, $M_z(0) = M_0\cos\theta$, stating that the RF pulse is assumed to have occurred. Solving for $\mathbf{M}$, we find that

$$\begin{aligned}M_{x'}(t) &= M_0\sin\theta\sin(\omega t)\,e^{\frac{-t}{T_2}}\\ M_{y'}(t) &= M_0\sin\theta\cos(\omega t)\,e^{\frac{-t}{T_2}}\\ M_z(t) &= M_0\left(1 - (1-\cos\theta)\,e^{\frac{-t}{T_1}}\right).\end{aligned} \tag{2.32}$$

It is informative to consider Eq. (2.32) separately. As shown in Fig. 2.5**A**, the effect of T_1 is that the longitudinal magnetisation recovers to its equilibrium value exponentially after it is perturbed, i.e.

$$M_z = M_0\left(1 - (1-\cos\theta)\,e^{-\frac{t}{T_1}}\right). \tag{2.33}$$

If the sample being probed possesses different tissues with different values of T_1, imaging a time T_I after an initial inversion pulse (i.e. $\theta = 180°$) will produce (drastically) different transverse magnetisations depending on T_1 and T_I. This is the basis of inversion recovery experiments, which are used clinically to exploit differences in T_1 between healthy and diseased tissue.

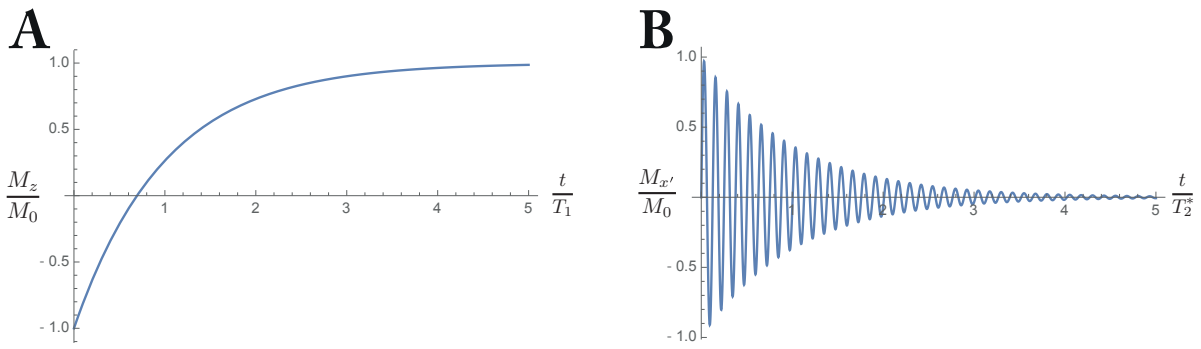

Figure 2.5: **A**: Following excitation (here by a $90°$ pulse), longitudinal magnetisation recovers exponentially with a characteristic time T_1. **B**: Similarly, following excitation (again shown here for a $90°$ pulse) the precessing magnetisation in the $x'y'$ plane exponentially decays with an envelope T_2^* (shown here for a slightly off-resonance peak).

As noted above, the transverse magnetisation is rapidly reduced by spin-spin interactions. Practically, the spins in a given environment do not experience the same magnetic field owing to local static magnetic field inhomogeneities that are caused by the complex interaction between the static field and the electromagnetic properties and geometry of the matter placed within it. These small variations can be ameliorated by a practice called shimming, described in detail in Section 3.4. The effect of these inhomogeneities is that the observed signal following a single pulse decays under an exponential envelope, with a decay constant denoted T_2^* rather than T_2 as above. The difference between T_2 and T_2^* is that T_2 is a property of the material in question, whereas T_2^* depends more on the local experimental conditions, even if those conditions are not necessarily independent of the sample in question. It is customary to define a separate relaxation time constant for these static field inhomogeneities, T_2', related by

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}. \tag{2.34}$$

2.2.2.2 *Detection & Chemical Shift*

Following excitation, the resultant magnetisation is rapidly rotating in the xy plane of the lab. It therefore induces an EMF whose magnitude is proportional to $\sqrt{M_x^2 + M_y^2}$. Experimentally, the same RF coil that generates B_1 is often used to detect this induced field, which is referred to as a free induction decay, or FID ('free' of B_1). Modern spectrometers detect the induced EMF in a phase sensitive fashion relative to a centre frequency ω_0 (c.f. Section 3.1), such that the detected signal $S(t)$ following a single broadband excitation is

$$
\begin{aligned}
M_{x'}(t) &= M_0 \sin\theta \sin(\Delta t)\, e^{\frac{-t}{T_2^*}} \\
M_{y'}(t) &= M_0 \sin\theta \cos(\Delta t)\, e^{\frac{-t}{T_2^*}} \\
\Rightarrow S(t) &\propto e^{-\frac{t}{T_2^*}}\, e^{i\Delta t}
\end{aligned} \tag{2.35}
$$

where $\Delta = \omega - \omega_0$ is the relative frequency offset for a given experiment. This is illustrated in Fig. 2.5**B**. If spin echo pulse sequences are used, the sense of precession in the $x'y'$ plane can be reversed by flipping the transverse magnetisation. Therefore, static local inhomogeneities serve to refocus the spins that they previously dephased. Correspondingly, spin echo pulse sequences predominantly have T_2 contrast. As $T_2 > T_2^*$, spin echo sequences are often used for rapid imaging sequences, discussed later.

As nuclei in different chemical environments experience slightly different values of B determined by the shielding effect of their microscopic electronic environment, they precess at different frequencies. If the radiofrequency pulse has a sufficiently broad bandwidth, it is possible to promote spins in these different chemical environments. The detected FID therefore is a superposition of the precession of the magnetisation due to each nucleus in the sample; its Fourier transform – the NMR spectrum of that sample – reveals these different frequencies of precession. This is of great utility in identifying the composition and relative concentrations of chemical species. The 'simplest' NMR experiment is depicted in Figure 2.6, where spins are excited with a single RF pulse, and the FID and hence spectrum acquired.

For most substances, the change in B that arises from electronic shielding is small, and it is common to express the relative change in field as a frequency shift relative to a specified reference. This quantity, δ, is known as the chemical shift and is usually expressed in parts per million:

$$
\delta = \frac{\omega_{\text{sample}} - \omega_{\text{reference}}}{\omega_{\text{reference}}} \cdot 10^6 \quad \text{ppm.} \tag{2.36}
$$

Most ^{1}H NMR spectra span only a few ppm, and so the overall range of frequencies that need to be sensitively detected is small.

2.2.3 Non-Constant Fields

Let us explicitly consider the effect of making $\mathbf{B}$ vary spatially. It will be shown in later sections that the spatio-temporal variation of $\mathbf{B}$ is the key requirement for encoding spatial information in magnetic resonance, and is therefore the basis of MRI.

It is customary to make small perturbations about $\mathbf{B}$ in the form of additive *linear* spatial gradients of field, $\mathbf{G} = G_x\mathbf{i} + G_y\mathbf{j} + G_z\mathbf{k}$. Ideally, $\mathbf{G}$ is a constant vector that can be switched on or off arbitrarily quickly in time. On modern systems it is the case that, to excellent accuracy, the 'gradient channels'

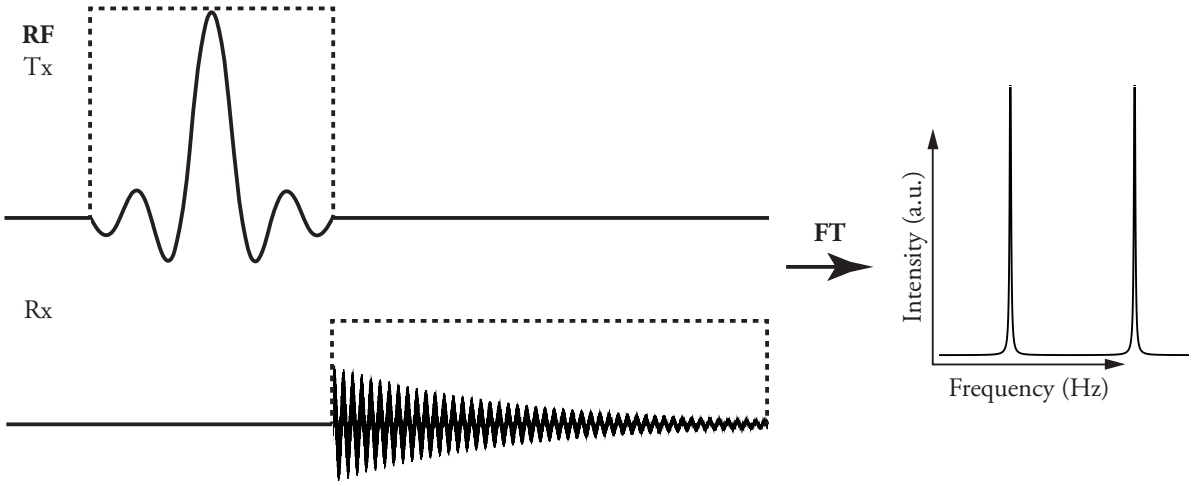

Figure 2.6: A simple 'pulse-and-acquire' NMR experiment. An RF waveform with a well-defined bandwidth is digitally synthesised, and, once a transmitter unblanking signal is received (dashed box) transmitted through the attached RF coil. After transmission, the same coil is used to receive, and a high-gain amplifier is unblanked. The resulting FID is then Fourier transformed, resulting in a spectrum of every chemical species with a resonant frequency within the transmitted bandwidth of the RF pulse.

are independent, i.e. the presence or absence of G_x has no influence on G_y. This convention has evolved arguably as it provides for a simple mathematical framework in which spatial encoding can exist. It is worth stressing, however, that it is a *convention:* the use of non-linear, non-independent perturbations to $\mathbf{B}$ generated by multiple small loops remains an active area of research, and one with many potentially exciting applications.[12–14] Moreover, historically, Raymond Damadian's precursor to MRI, Field fOcusing Nuclear mAgnetic Resonance ("FONAR") spatially varied the field over a region of interest, violating the assumptions above.[15] Nevertheless, for the sake of clarity, let us briefly illustrate the effect of spatially perturbing $\mathbf{B}$.

2.2.3.1 *Magnetic Field Gradients*

Consider a general field, $\mathbf{B}$, within a strong static field B_0 defined as

$$\mathbf{B} = B_x(x,\, y,\, z)\mathbf{i} + B_y(x,\, y,\, z)\mathbf{j} + B_0 + B_z(x,\, y,\, z)\mathbf{k} \tag{2.37}$$

where $B_x(x,\, y,\, z)$ denotes that the x-component of $\mathbf{B}$ is a function of space. Assuming that B_0 is very large, we are able to neglect B_x and B_y.[3] Secondly, suppose now that $\mathbf{B}$ has a linear spatial gradient in the $\mathbf{G}$ direction. Another way of stating this condition is to require that the field is constant within a plane perpendicular to $\mathbf{G}$, but linearly increases as we progress in the $\mathbf{G}$ direction. Vector geometry will therefore tell us that

[3]Physically, this isn't strictly possible as Maxwell requires that $\nabla \cdot \mathbf{B} = 0$ everywhere. However, B_x and B_y are far smaller than B_z, although playing a gradient as subsequently described changes the other components of $\mathbf{B}$ according to Maxwell's equations. It can be shown that the resultant x and y components of the field are weak, and scale as $1/B_0$. Correspondingly, at 7 T these 'concomitant field' effects are weak, on the order of 3 ppm, which is comparable to the homogeneity of B_0 over the volume of interest. Their effect is further reduced by the gradient correction tensor detailed in Section 3.1.3.

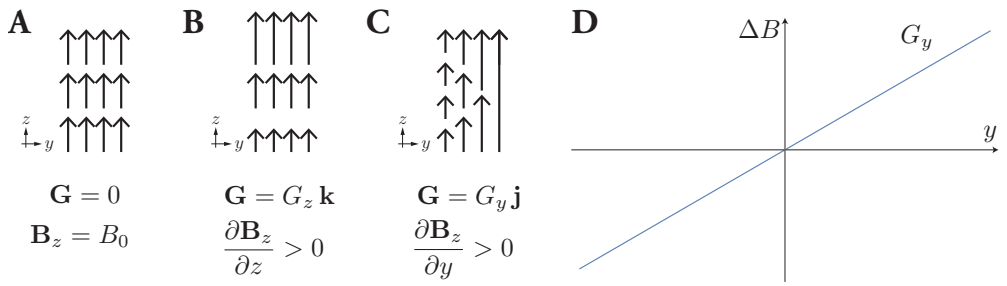

Figure 2.7: The changes in **B** following the application of a gradient **G**. **A**: Before the application of a gradient, **B** is assumed to be homogeneous and along the z axis. Applying a linear field gradient along the z or y directions (**B** and **C** respectively) then modulates the strength of the static field. To an excellent approximation, the transverse components of **B** are still zero. **D**: The effect of this situation is to provide a linear gradient of the z component of **B**, illustrated here for $G_y > 0$.

$$\nabla B_z = \mathbf{G}, \tag{2.38}$$

or, writing Eq. (2.38) in components

$$\frac{\partial B_z}{\partial x}\mathbf{i} + \frac{\partial B_z}{\partial y}\mathbf{j} + \frac{\partial B_z}{\partial z}\mathbf{k} = G_x\mathbf{i} + G_y\mathbf{j} + G_z\mathbf{k}. \tag{2.39}$$

It is straightforward to integrate the above and find that

$$\begin{aligned} B_z &= G_x x + G_y y + G_z z, \text{ or,} \\ B_z &= \mathbf{G}\cdot\mathbf{r}, \end{aligned} \tag{2.40}$$

where $\mathbf{r} = x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$. Correspondingly, the effect of playing spatial gradients is to directly alter the static field, as we find that

$$\mathbf{B} \approx B_0 + (\mathbf{G}\cdot\mathbf{r})\,\mathbf{k}. \tag{2.41}$$

This is illustrated schematically in Fig. 2.7. It is possible to consider the effect spatially of this linear gradient. Assume for the sake of clarity that $\mathbf{G} = G_x\mathbf{i}$ only. Then, the Larmor equation (Eq. (2.7)) states that $\omega = \gamma B = \gamma(B_0 + G_x x)$; i.e. there is a mapping between frequency and position along x. Likewise, assuming that the gradient is only turned on for a time τ we find that there is a 'phase dispersal' across the slice, i.e. that the phase of the precessing magnetisation vector is different in different positions. Recalling that the rate of change of phase is frequency, we can write

$$\begin{aligned} \phi(x,\,t) &= \gamma\int_0^\tau G_x(t)x\,\mathrm{d}t \\ &= \gamma G_x x\tau, \end{aligned} \tag{2.42}$$

if G_x is turned on and off instantaneously, and constant otherwise. Therefore, the presence of a linear gradient for a time τ introduces a phase shift in $\mathbf{M}$. It is customary to call G_x and similar the *gradient amplitude* (in G/cm or mT/m), and the total area $\int_0^\tau G_x(t)x\,\mathrm{d}t$ the zeroth[4] gradient moment, or just *gradient moment* for short.

The role of spatial gradients in MRI is detailed substantially in Section 2.4. However, their ability to temporarily perturb the frequency and phase of the precession of spins also finds utility in designing RF pulses. It is possible to tailor the physical region of the sample excited by playing a gradient whilst also playing RF. This suggests the possibility of designing radiofrequency pulses to perform a specific task.

2.3 | *Radiofrequency pulse design*

Given $\mathbf{B}(x,\,y,\,z,\,t)$, predicting the response of $\mathbf{M}$ is a relatively simple problem: the Bloch equations are easy to solve numerically by a variety of methods. It is therefore somewhat straightforward to numerically determine the effect of an RF pulse given an initial magnetisation, as this process is deterministic and the numerical methods required are stable. The inverse problem, however, is more difficult, as the Bloch equations are not, in general, linear and cannot be inverted analytically. Two main methods exist to aid the design of radiofrequency pulses: (1) the small flip-angle approximation, under which the effect of an RF pulse is its Fourier transform (known as the excitation k-space formalism); and (2) the Shinnar-Le Roux (SLR) algorithm, which exploits fundamental symmetries in the Bloch equations to transform the pulse design problem to that of polynomial filter design, for which exact methods exist. Both the excitation k-space formalism and the SLR algorithm have been used for designing the spectral-spatial pulses described later in this text.

2.3.1 *Excitation k-space*

Let us quantitatively illustrate the effect of playing an RF pulse whilst a magnetic gradient field is present by considering the Bloch equations (Eq. (2.30)) in the rotating frame and substituting $B_0 \mapsto B_0 + \mathbf{G}\cdot\mathbf{r}$ as per Eq. (2.41). Assuming B_1 is small, we find that the effective magnetic field acting on the spins is $B_{1,\,\text{eff}} = [B_{1x'},\,B_{1y'},\,\mathbf{G}\cdot\mathbf{r}]$, where $\mathbf{r}$ is the position vector in the rotating frame, and $\mathbf{G}$ is the applied linear gradient vector. Neglecting relaxation (i.e. assuming $T_1,\,T_2 \to \infty$), we correspondingly find that Eq. (2.31) can be expressed as

$$\begin{pmatrix} \frac{\mathrm{d}M_{x'}}{\mathrm{d}t} \\ \frac{\mathrm{d}M_{y'}}{\mathrm{d}t} \\ \frac{\mathrm{d}M_{z}}{\mathrm{d}t} \end{pmatrix} = \gamma \begin{pmatrix} 0 & \mathbf{G}\cdot\mathbf{r} & -B_{1y'} \\ -\mathbf{G}\cdot\mathbf{r} & 0 & B_{1x'} \\ B_{1y'} & -B_{1x'} & 0 \end{pmatrix} \begin{pmatrix} M_{x'} \\ M_{y'} \\ M_z \end{pmatrix} \tag{2.43}$$

[4]Higher order moments are defined as $\int_0^\tau G(t)t^n\,\mathrm{d}t$ for the nth moment; they find applications where spins are flowing, and are discussed in Section 2.4.7.

Defining the complex transverse magnetisation $M_\perp = M_{x'} + iM_{y'}$ and applied field $B_1 = B_{1x'} + iB_{1y'}$, the transverse components of Eq. (2.43) can be expressed [16] as

$$\frac{\mathrm{d}M_\perp}{\mathrm{d}t} = -i\gamma \mathbf{G}\cdot\mathbf{r}M_\perp + i\gamma B_1(t)M_z(t). \tag{2.44}$$

Given the initial condition $M_\perp(0) = 0,\ M_z(0) = M_0$, Eq. (2.44) has the solution for the magnetisation at a final time t

$$M_\perp(t) = i\gamma \int_0^t M_z(t')B_1(t')e^{-i\gamma\mathbf{r}\cdot\int_0^{t'}\mathbf{G}(s)\ \mathrm{d}s}\ \mathrm{d}t'. \tag{2.45}$$

The transverse magnetisation is therefore the inverse Fourier transform of the product of the applied RF and the remaining longitudinal transformation. In general, Eq. (2.45) is not useful practically as M_z is inherently an unknown function of time. However, the effect of a pulse is to map $M_z(0) = M_0 \mapsto M_z(t) = M_0\cos\theta$ for a flip angle θ. If one assumes that θ is small, $M_z(t) \approx M_0$. By making this assumption and defining a spatial frequency variable $\mathbf{k}(t) = -\gamma\int_0^t \mathbf{G}(s)\ \mathrm{d}s$ we find that Eq. (2.45) simplifies considerably:

$$M_\perp(t) = i\gamma M_0 \int_0^t B_1(t')e^{i\mathbf{k}(t')\cdot\mathbf{r}}\ \mathrm{d}t'. \tag{2.46}$$

We therefore see that, for small flip angles, the excited magnetisation is proportional to the Fourier transform of the transmitted RF, weighted by the velocity taken in $\mathbf{k}$. This approach is known as the *excitation k-space formalism*, and is of great utility in designing RF pulses: given a desired magnetisation response, which may be a function of space, the RF pulse that generates it is approximately its Fourier transform. In practice, the small flip angle approximation breaks down when $\theta > 90°$,[7, p. 73] as the original assumption that $M_z(t) = M_0$ becomes progressively less valid. As Fourier transforms are linear operators, practical deviations from the realm of validity for Eq. (2.46) are known as nonlinearities, and iterative methods or the SLR algorithm are used to ameliorate them.

2.3.2 *The Shinnar-Le Roux Algorithm*

When the desired flip angle is not less than 1 rad, the excitation k-space formalism becomes progressively worse, to the point where inversion pulses are essentially not playable. Whilst this is unsurprising, there are many situations in which large θ pulses are desired, and iterative approaches are unwieldy, computationally expensive, and difficult to implement. The SLR algorithm allows the inverse problem to be solved directly, and without iteration.[17] The algorithm requires the designer to specify (and make compromises between) parameters such as the excitation bandwidth, pulse duration, grid raster, and flip angle, and returns an exact pulse directly. Moreover, the SLR algorithm allows compromises between these parameters to be made explicitly, before the pulse is generated.

Very briefly, the algorithm relies upon two key assumptions: (1) that the effect of a shaped RF pulse can be broken down into a succession of short rectangular ("hard") pulses followed by periods of free precession, and that (2) each of these pulses serves to elicit a rotation of the magnetisation vector. Mathematically, three dimensional rotations can be described equivalently in two different ways known as SO(3) or SU(2). Both SO(3) and SU(2) are groups, i.e. sets with a binary operator that map any two elements of the set to form a third element, subject to various conditions (known as axioms of the group). SO(3) is the group of all rotations about the origin of a Euclidean space with properties that are "sensible", e.g. preserving distances between points and the position of the origin. Representations of SO(3) are the well-known, 3×3 real orthonormal rotation matrices that have nine real elements. As a rotation can be represented in terms of three free parameters (such as the three Euler angles χ, ψ and η), there are $9 - 3 = 6$ free parameters in the group (and therefore in these matrices). However, it is also possible to specify a rotation in a complex representation in SU(2), by a rotational operator

$$\mathbb{Q} = \begin{pmatrix} \alpha & -\beta^* \\ \beta & \alpha^* \end{pmatrix} \tag{2.47}$$

where * denotes complex conjugation and α and β are complex parameters constrained by the normalisation requirement that $|\alpha|^2 + |\beta|^2 = 1$. The matrices $\mathbb{Q}$ are unitary, that is $\mathbb{Q}^\dagger\mathbb{Q} = \mathbb{I}$ where † denotes the Hermitian conjugate. As the parameters α and β span four real numbers, representations in SU(2) contain only $4 - 3 = 1$ constraint to describe a rotation. As both SU(2) and SO(3) describe rotations, they are said to be duals, and a specific rotation can be represented with equal validity as a matrix of either form. However, the fewer number of constraints in an SU(2) representation of a rotation offers considerable mathematical simplification, as it is almost always easier to solve fewer equations with fewer constraints.

The beauty of the SLR algorithm is that N successive periods of precession and nutation are described by

$$s_{\text{final}} = \prod_{i=0}^{N} \mathbb{Q}_{i,\,\text{nutation}} \mathbb{Q}_{i,\,\text{precession}}\ s_0 \tag{2.48}$$

where s_0 is the initial magnetisation spinor (usually assumed to be $(1, 0)$, i.e. along the z axis) and multiplication occurs to the left. Inverting the Bloch equations is then reduced to the problem of inverting a large number of 2×2 Hermitian matricies, i.e. taking their conjugate transpose. For physical pulses, it can be shown that this final state is a polynomial in α and β that can be related to physical properties by expression in the variable $z = \exp(i\omega\Delta t) = \exp(i\gamma\mathbf{G} \cdot \mathbf{r}\Delta t)$, where Δt is the interval between each period of precession and nutation, defined by T/N, where T is the total pulse length, and where $\mathbf{r}$ is the position vector inside the magnet. Typically, Δt is taken to be close to the

limiting RF pulse generator resolution on the system in question. If the effect of the jth hard pulse is to produce a flip angle

$$\theta_j = \gamma |B_{1,\,j}| \Delta t, \tag{2.49}$$

and the phase of the applied RF in the rotating frame is given by

$$\phi_j = \angle(B_{1,\,j}), \tag{2.50}$$

then we find that

$$\begin{aligned} \mathbb{Q}_{j,\,\text{nutation}} = \begin{pmatrix} \cos\frac{\theta_j}{2} & ie^{-i\phi_j}\sin\frac{\theta_j}{2} \\ ie^{i\phi_j}\sin\frac{\theta_j}{2} & \cos\frac{\theta_j}{2} \end{pmatrix}, \qquad \mathbb{Q}_{j,\,\text{precession}} &= \begin{pmatrix} e^{i\frac{1}{2}\omega\Delta t} & 0 \\ 0 & e^{-i\frac{1}{2}\omega\Delta t} \end{pmatrix} \\ &= \begin{pmatrix} z^{1/2} & 0 \\ 0 & z^{-1/2} \end{pmatrix}. \end{aligned} \tag{2.51}$$

Explicitly, it then transpires that the effect of one pulse is

$$\begin{aligned} s_{j+1} &= \mathbb{Q}_{j,\,\text{nutation}}\mathbb{Q}_{j,\,\text{precession}}s_j \\ &= z^{1/2}\begin{pmatrix} \cos\frac{\theta_j}{2} & ie^{-i\phi_j}\sin\frac{\theta_j}{2} \\ ie^{i\phi_j}\sin\frac{\theta_j}{2} & \cos\frac{\theta_j}{2} \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & z^{-1} \end{pmatrix} s_j \\ &= z^{1/2}\begin{pmatrix} \cos\frac{\theta_j}{2} & iz^{-1}e^{-i\phi_j}\sin\frac{\theta_j}{2} \\ ie^{i\phi_j}\sin\frac{\theta_1}{2} & z^{-1}\cos\frac{\theta_j}{2} \end{pmatrix} s_j. \end{aligned} \tag{2.52}$$

Correspondingly, Eq. (2.48) can be evaluated subject to the condition that the initial magnetisation spinor $s_0 = (1,\,0)^\intercal$, and the final state expressed as

$$s_{\text{final}} = z^{N/2}\begin{pmatrix} A_N(z) \\ B_N(z) \end{pmatrix} \tag{2.53}$$

where $A_N(z)$ and $B_N(z)$ are two complex polynomials of order $j-1$ in z which share normalisation relations with α and β and whose coefficients are related to the amplitude and phase of the applied RF at each point in time. The problem of designing an RF pulse is therefore reduced to that of polynomial filter design subject to various (conflicting) constraints. Usually, $A_N(z)$ is constrained additionally to have roots within the unit circle, ensuring a minimum-phase design that minimises the transmitted RF power, and the phase of $B_N(z)$ is also usually chosen to be linear or minimised. A common design parameter is the dimensionless time-bandwidth product, which is a measure of selectivity for a given pulse: greater time-bandwidth product pulses are more selective, usually at the expense of pulse

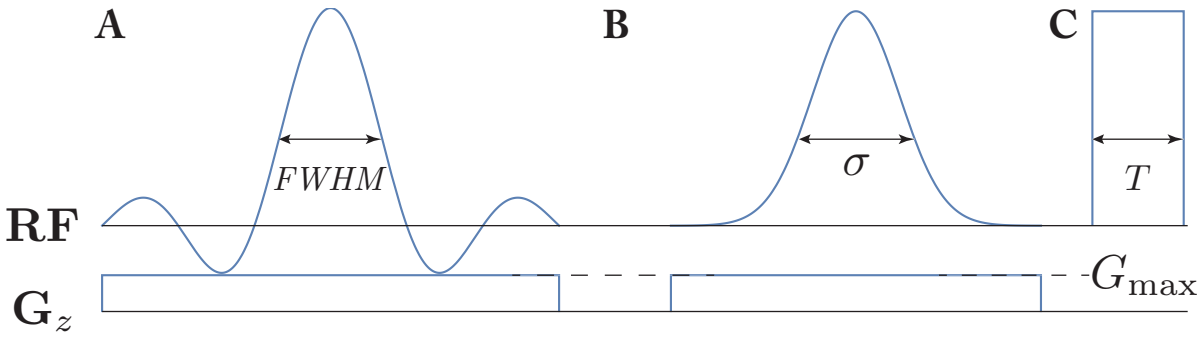

Figure 2.8: The sinc (**A**), gaussian (**B**) and hard (**C**) RF pulses are mainstays of magnetic resonance experiments, with excitation bandwidths Δf of approximately $2/\text{FWHM}$, $0.375/\sigma$ and $1.21/T$ respectively. Where a linear slice-select gradient is present, the corresponding slab thickness excited is $\Delta z = \Delta f / \gamma G_{\text{max}}$.

length. Given these large number of constraints the problem of choosing z is then overdetermined, and several exact techniques can extract the transmitted RF amplitude and phase at each point in time. The interested reader is referred to the literature for both a broader overview of the technique[7, pp. 44-58] and the original articles by Shinnar,[17] Le Roux,[18,19] and Pauly[20] for information necessary for their practical implementation.

2.3.3 *Slice Selection and Excitation Bandwidth*

It is often desirable to localise a magnetic resonance experiment to a given region in space, commonly a rectangular slab of a known thickness. This is usually achieved by the use of 'slice-selection' gradients, which only excite spins in a given plane by playing an RF pulse of a finite bandwidth $\Delta\omega$ concurrently with a gradient of strength G_{SS} along the slice axis. There is therefore a linear dependence of ω along the slice axis, and only the region between $\gamma(B + G_{\text{SS}} x_{\text{start}}) \stackrel{\text{def}}{=} \omega - \Delta\omega$ to $\gamma(B + G_{\text{SS}} x_{\text{end}}) \stackrel{\text{def}}{=} \omega + \Delta\omega$ is excited.

To tailor the exact profile of the excited slice, it is necessary to modulate the shape of the RF pulse played. For a low angle excitation, it is usually desired to excite a thin slice or thick slab of magnetisation about a given axis for subsequent imaging. For this reason, rectangular ("hard"), sinc-shaped or gaussian pulses are often used, as the resulting pulse profiles are relatively flat, and the Fourier relationships between pulse length and excitation bandwidth conceptually simple. Hard pulses are often used for unlocalised spectroscopy, as the Fourier relationship between bandwidth and pulse length means that a short pulse excites a large bandwidth and minimises de-coherence effects due to T_2^*. They are not used for slice selection due to the resulting sinc-shaped profile, and distortions arising from nonlinearity in the Bloch equations. Gaussian and sinc pulses with linear slice selection gradients are the mainstay of imaging, and are illustrated in Fig. 2.8.

Most slice-selective RF pulses (including those above) played with linear gradients generate a linear dispersion of phase across the slice. In general, the phase response of an RF pulse is simply determined

by $\angle M_{\perp}$, and it is possible to design a pulse to have a uniform (i.e. "non-spoiled") phase profile in space. However, an alternative approach is the use of slice refocussing gradients to rephase this dispersion after excitation. Usually, a gradient moment of $-1/2\times$ that played during excitation is required for refocussing. As a short-duration gradient blip can refocus this dispersion quickly, it is often more effective to use a short pulse with a large excitation bandwidth followed by a short slice-select refocussing gradient, as opposed to a longer, more complex design.

2.4 | *Magnetic Resonance Imaging*

2.4.1 *Frequency encoding*

If B is made a function of space, then the frequency of precession for a given group of (otherwise identical) spins will also vary spatially. Let $\mathbf{G}$ denote a magnetic field gradient applied in addition to $\mathbf{B}$. Then spins will precess at $\omega(\mathbf{r}) = \gamma(B_0\mathbf{k} + \mathbf{G}\cdot\mathbf{r})$. Consider a linear gradient of field along the x axis, and a single chemical species (e.g. water) present in the sample. Under these conditions, all spins at a particular value of x (known as spin isochromats) will precess at the same frequency. The recorded FID will therefore contain component frequencies from every x value occupied by the sample, and its Fourier spectrum will therefore be the projection of the density spins lying in that plane, as illustrated in Figure 2.9. This is known as *frequency encoding*, as a one-to-one mapping is made between the precessional frequency of the sample and its position along a given axis. For obvious reasons, artefacts arise if more than one chemical species is excited prior to imaging. Frequency encoding gradients are those played during the detection of the FID, and are therefore also known as 'readout gradients'.

Owing to the dispersion of phase across a slice by a readout gradient, it is necessary to prepare the phase across the excited slice to create a coherence later. This process is performed by playing a gradient

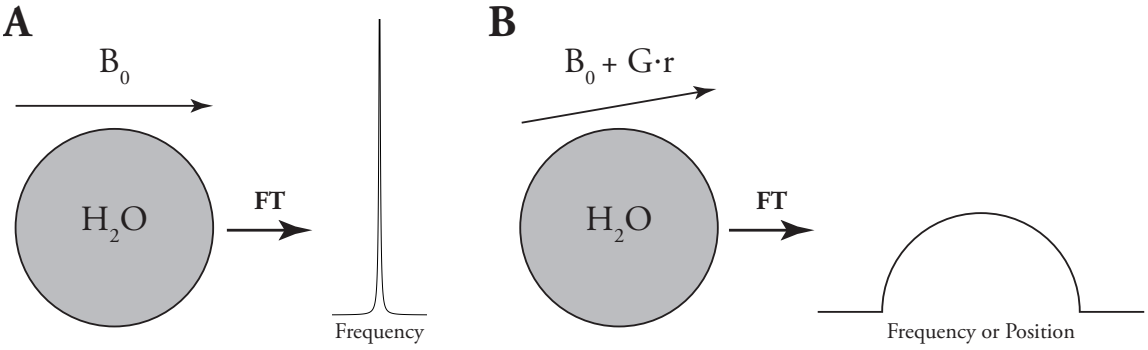

Figure 2.9: Consider an NMR experiment on a homogeneous, uniform sample of water. **A**: In the absence of a frequency-encoding gradient, the time-domain signal received after excitation consists of a single resonant frequency, namely that of the sample. **B**: Applying a frequency-encoding gradient G at the same time as the acquisition of the FID creates a mapping, whereby each frequency corresponds to a position r along the gradient axis. A range of resonant frequencies therefore is contained in the detected FID; the spectrum thus gives a projection of the imaged object.

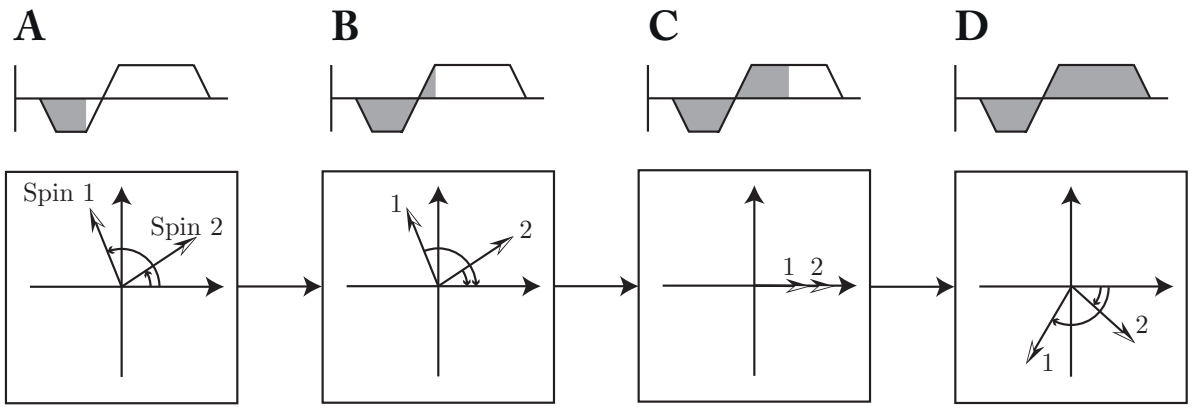

Figure 2.10: The effect of a prephasing gradient, illustrated for two spin isochromats, denoted 1 and 2. During the prephasing gradient, spins in different spatial positions accumulate phase (**A**). Under the readout gradient, this is reversed (**B**), forming an echo when isochromats are in phase in the middle of the readout lobe (**C**) before being further dephased (**D**). The effect of the gradient is therefore to ensure that the origin of k-space is in the middle of the acquisition.

moment before the readout actually occurs, known as a 'prephasing' gradient lobe, dephasing the spins sufficiently so that an echo is formed in the middle of the readout. This process can be understood as shifting the origin of image k-space in the formalism that follows (Section 2.4.3), or intuitively by considering two spin isochromats in different positions as illustrated in Fig. 2.10.

2.4.2 Phase encoding

Unfortunately, frequency encoding cannot be performed about two or more axes simultaneously as the resulting mapping is many-to-one and therefore does not have a unique inverse. Historically, projection reconstructions (such as backprojection and filtered backprojection) were used to acquire an image with many rotations of the frequency encoding axis. However, these approaches are not as SNR efficient as modern Fourier encoding schemes that are made feasible by the quadrature detection of the FID, recording both its magnitude and phase. Phase encoding creates a (usually linear) mapping between the phase of the precession magnetisation and its position, usually along an orthogonal axis to the slice selection and frequency encoding axes. Phase encoding is performed by playing a gradient moment after excitation but before readout, shifting the phase of the precessing spins. Different phase increments along an axis are obtained by changing the moment of the phase encoding gradient played. This process is illustrated schematically in Fig. 2.11.

2.4.3 Image k-space

Quantitatively, after a phase encoding gradient $G_y(t)$ of total duration T is played the total phase accumulated across the phase encoding axis y will be simply

$$\phi(y) = y\gamma \int_0^T G_y(t')\ \mathrm{d}t' \stackrel{\text{def}}{=} 2\pi k_y(T)y \tag{2.54}$$

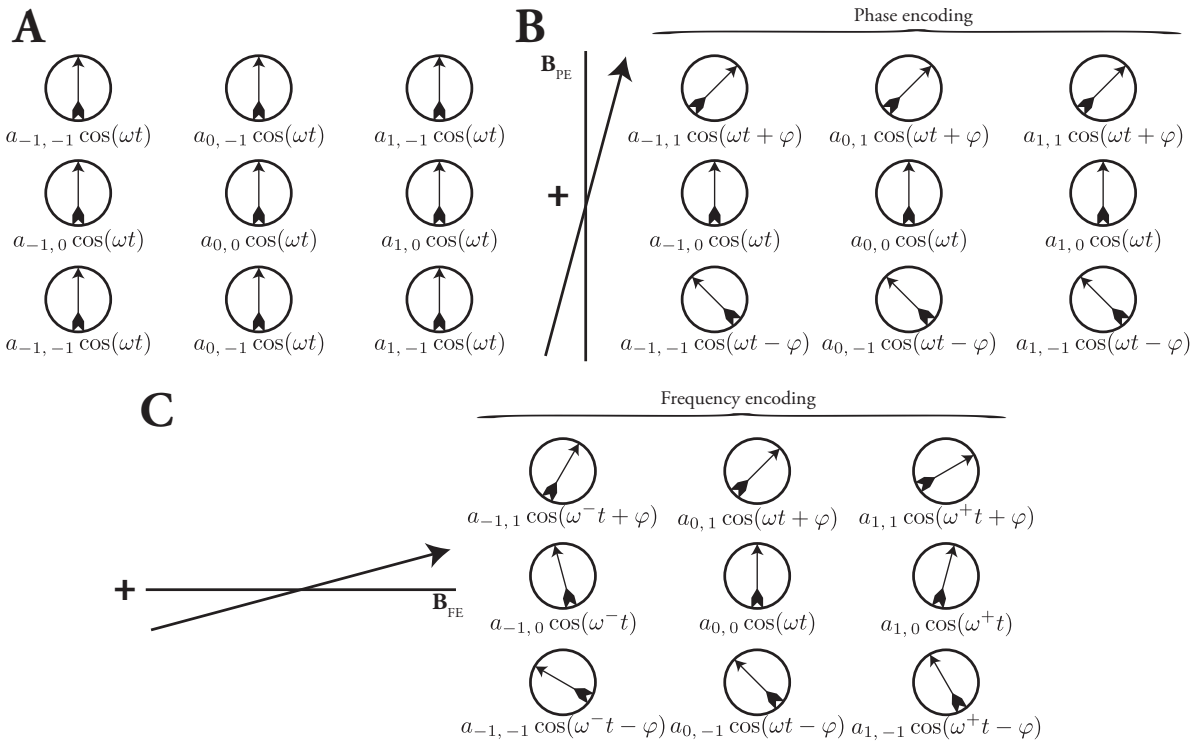

Figure 2.11: The cannonical 'clocks' diagram for phase encoding. **A**: After a 90°slice-selection pulse the magnetisation vector precesses in the xy plane, here the plane of the paper. Arbitrarily, the magnetisation is shown aligned with the y axis. Each voxel contributes a signal a_{ij} proportional to the density of excited spins in that voxel. **B**: Applying a gradient about an orthogonal axis (here the y axis) for a time τ shifts the phase of the precession, in a spatially dependent way; $\Delta\varphi = \gamma B(y)\tau$. This process encodes position in the phase of the precession, and is therefore known as phase encoding. **C**: Applying an additional gradient about the remaining orthogonal axis results in the frequency of precession shifting spatially while that gradient is applied, i.e. $\omega(x) = \gamma B(x)$. Data is therefore acquired during the presence of the frequency encoding gradient, and there exists a homomorphism between position and frequency.

where the variable $k_y(T)$ is defined as above, and has units of inverse distance. Neglecting relaxation, the signal received is the weighted vector sum of the magnetisation of every excited nuclear spins in the object imaged. Therefore, in terms of the complex transverse magnetisation, $M_\perp$, it becomes apparent that the signal received due to a phase encoding gradient is

$$s(k_y) = \int M_\perp(y) B_\perp(y) e^{-i\phi(y)}\ \mathrm{d}y \tag{2.55}$$

where $B_\perp(y)$ is the component of the receive B_1 field in the transverse plane as defined following Eq. (2.43).

Generalising Eq. (2.55) leads naturally into the concept of imaging k-space. Given a gradient vector $\mathbf{G}(t)$, the total accumulated phase in three dimensions is

$$\phi(\mathbf{r},\, t) = \gamma \int_0^t \mathbf{r} \cdot \mathbf{G}(t')\ \mathrm{d}t' \tag{2.56}$$

where all other quantities are as before. Defining

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t') \, \mathrm{d}t' \tag{2.57}$$

we find that the received signal in MR experiments is

$$S(t) = \int M_\perp(\mathbf{r}) B_1(\mathbf{r}) e^{-2\pi i \mathbf{k}(t)\cdot\mathbf{r}} \, \mathrm{d}\mathbf{r}. \tag{2.58}$$

where the integral runs over $\mathbb{R}^3$.

In Eq. (2.58), the function $\mathbf{k}(t)$ is the Fourier conjugate variable to the spatial variable $\mathbf{r}$. Correspondingly, as time evolves $S(t)$ traces out a path $\mathbf{k}(t)$ in k-space. The velocity at which k-space is scanned depends on $|\mathbf{G}|$, as, from Eq. (2.57), $|\,\mathrm{d}\mathbf{k}(t)/\,\mathrm{d}t| = \gamma|\mathbf{G}|/2\pi$. As most nuclei have positive γ (with $^{15}\mathrm{N}$ being the notable exception), k-space is scanned in the same direction as the gradient polarity, and the total distance in k spanned is determined by the gradient moment played. If every chemical species excited is the same, and B_1 spatially homogeneous, then the quantity $M_\perp(\mathbf{r})B_1(\mathbf{r})$ is proportional to the density of spins, and Eq. (2.58) then becomes an explicit Fourier transform of the spin density.

As shown in Eq. (2.57), MRI therefore directly samples a multidimensonal k-space and the experimental implementation of that sampling directly relates to the quality of the final image produced. For practical reasons, the integrals in Eqs. (2.57) and (2.58) are discretised and a number of points in k-space sampled to build up a final image. The number of points chosen to be sampled are referred to as the matrix size, commonly on the order of 128×128 for $^1\mathrm{H}$ imaging. The final image resolution, and correspondingly SNR, are determined to a large extent by the k-space trajectory taken. It can be shown that the resolution Δy in a given direction depends reciprocally on the maximum distance in k-space traversed, i.e.

$$\Delta y = \frac{1}{k_{y\,\mathrm{max}}}. \tag{2.59}$$

Similarly, the extent of the field-of-view of the image is reciprocally related to the spacing in k-space:

$$\mathrm{FOV}_y = \frac{1}{\Delta k_y}. \tag{2.60}$$

Further practical considerations of the sampling of k-space, such as receiver bandwidth and the trade-off between gradient duration and amplitude, are discussed in Section 2.4.5.

It should also be noted that Eq. (2.57) is valid in three dimensions, and that phase encoding can be performed through all of them. In this instance each acquisition at each voxel is the NMR spectrum of the sample, a process known as chemical shift imaging (CSI). As phase encoding is a slow process, CSI is usually limited to small matrix sizes.

2.4.4 Pulse sequences

The collection of RF pulses, slice-selection, phase encoding and frequency encoding gradients played together determine both the trajectory in k-space taken and the image obtained. They are referred to as a given pulse sequence. Undoubtedly, the most popular k-space trajectory is a Cartesian raster in which each line of k-space is acquired by frequency encoding whilst at a constant magnitude of the phase-encoding gradient. Steps in k_y are taken by varying the magnitude of the phase encoding gradient, and not its duration (which would otherwise introduce T_2^* weighting across k_y). All lines in the raster are parallel and separated by equal distance in k-space, and the image can be reconstructed simply by a Fourier transform.

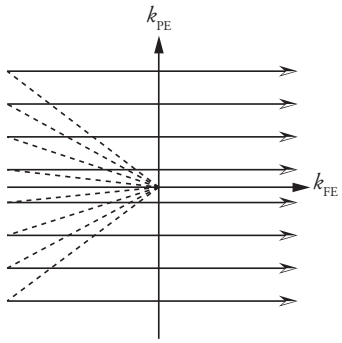

Figure 2.12: A Cartesian FLASH k-space trajectory, in which multiple lines are acquired by phase encoding. Lines acquired are shown in black, while dashed lines indicate times in "transition" between different lines through k-space.

The Cartesian trajectory is popular because artefacts arising from its imperfect implementation are far less objectionable than those from other trajectories, and it is illustrated in Fig. 2.12. In particular, the rapid imaging strategies discussed later are particularly susceptible to artefacts arising from eddy currents, local field inhomogeneities, susceptibility and magnetisation transfer effects, and other imperfections. One drawback of the Cartesian raster is that relatively long scan times result unless echo train spin refocusing is used, owing to the fact that each line requires a separate RF excitation pulse (or multiple pulses, if signal averaging is used). The use of echo trains greatly increases the sensitivity of a sequence to several artefacts, and introduces an additional weighting in the resulting image. Owing to the fact that residual magnetisation can cause an unwanted stimulated echo on successive excitations, producing an artefact in the final image, most Cartesian sequences explicitly 'spoil' the magnetisation after a line of k-space is acquired, by dephasing any remaining transverse magnetisation.

When combined with a single selective RF pulse, this Cartesian scanning is usually called Gradient echo imaging, and is illustrated in Fig. 2.14. Two important parameters are defined, the repetition rate TR between scans, and the echo time TE . Gradient echo sequences can be used to image 3D volumes in two different ways. Following the excitation of a thin slice, it is possible to subsequently excite and acquire extra slices with no time penalty whilst waiting for the longitudinal magnetisation of the previous slice to recover. This technique is called gradient echo multiple slice imaging. Alternatively, a large slab can be excited and phase encoding used to step through the third dimension. Both scenarios are useful under different circumstances. Commonly, gradient echo scans are used for localising anatomy with small flip angles ($\sim 5^\circ$) to run at with a fast repetition rate, acquiring an entire image in the order of

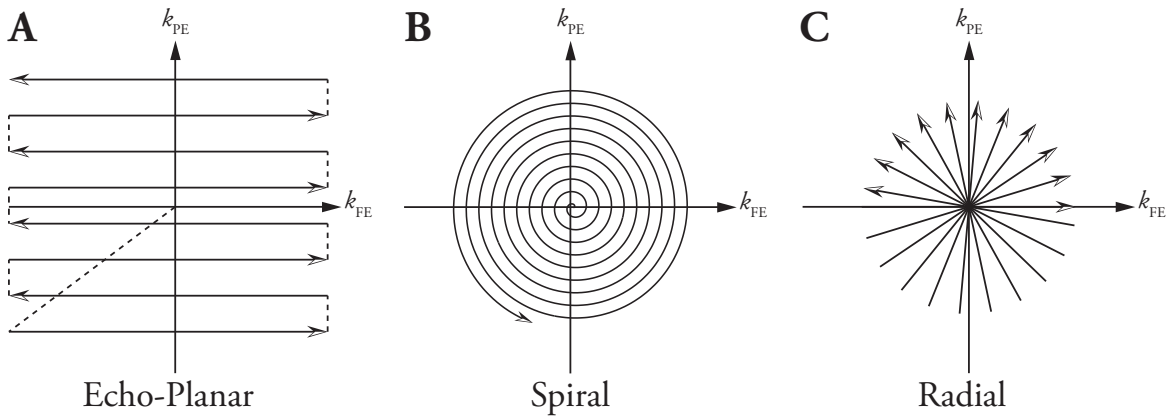

Figure 2.13: Different pulse sequences take a different trajectory through k-space; lines acquired are shown in black, while dashed lines indicate times in "transition" between different lines through k-space. Shown here are (**A**), an example EPI trajectory, where data is acquired on fast oscillating frequency encoding lobes followed by short blipped phase encodes (although note that other schemes are possible), (**B**) an example spiral trajectory, and (**C**) a radial trajectory consisting of a large number of quickly-acquired spokes.

seconds. This technique is often known as Fast-Low Angle multi SHot ('FLASH') imaging.

Numerous other k-space trajectories are commonly used. Of particular note are 'rapid' imaging sequences that traverse the whole of k-space following a single excitation ("single shot" methods), or acquire a large amount of k-space in a small number of segments. Echo planar imaging (EPI) is the rapid equivalent of Cartesian sequences, as it takes a Cartesian path through k-space in a single shot. It is also possible to traverse k-space in a polar trajectory, taking a spiral path through k_x and k_y space, or acquiring multiple radial projections through the origin. As the signal recorded decays with T_2^*, single-shot imaging techniques are challenging to implement, requiring strong gradients that rapidly oscillate to acquire data quickly.

However, the signal loss due to T_2^* usually fundamentally limits the matrix size of single-shot techniques, and also introduces artefacts in the resulting images. Three dimensional single-shot techniques, which traverse a 3D k-space following a single excitation, are correspondingly even harder to implement, and are not widely used. Several examples of k-space trajectories are illustrated in Fig. 2.13.

2.4.5 Receiver bandwidth

There are several additional parameters related to the practical implementation of a pulse sequence that have yet to be discussed. Principle amongst these is the concept of receiver bandwidth, which is the range of frequencies digitised by the ADC. A reciprocal relationship exists between bandwidth and the sampling time for each point, given by the Nyquist-Shannon sampling theorem, stating that $\omega_s \geq 2\tau$, where ω_s is the sampling bandwidth and τ the time-per-point or 'dwell time'. There is accordingly a trade-off in SNR that can be made by adjusting the sampling bandwidth, as shorter acquisitions have less signal decay due to T_2^* but the larger acquisition bandwidth increases the range of noise that is

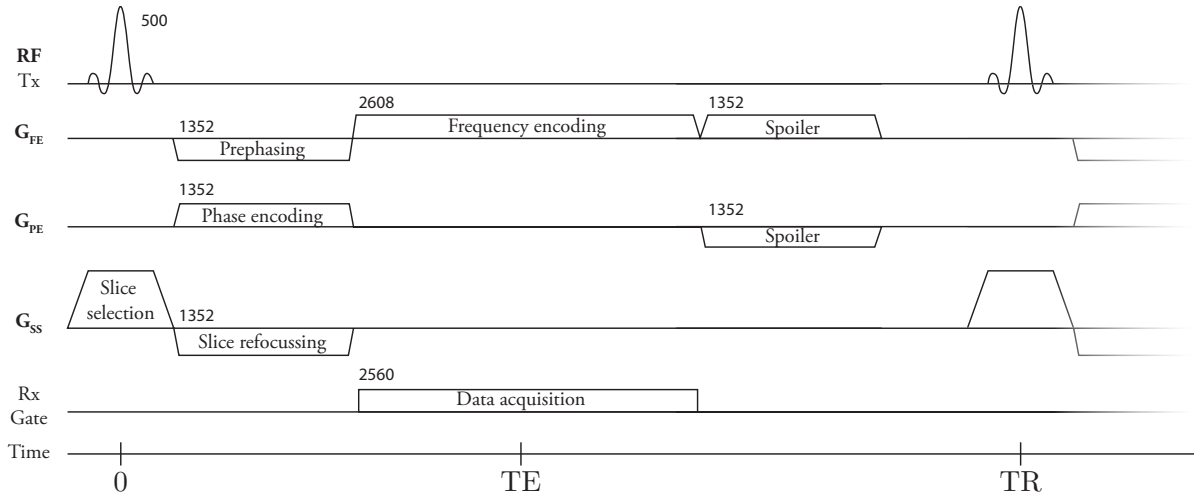

Figure 2.14: The gradient-echo imaging sequence, as implemented on our 7 T scanner. Excitation is performed, with the origin of the time axis traditionally taken to be the isodelay which for a sinc pulse is at its maximum amplitude. In order to reduce the TR , the phase encoding step and prephasing gradient occur concurrently following excitation. Data is then acquired, with the echotime TE defined as being from the RF centre to the maximum of the acquired data, here the centre of the echo. Spoilers in the transverse plane follow to dephase any remaining magnetisation, and the process repeats a time TR later.

acquired through processes known collectively as Johnson noise which scales as $\sqrt{\Delta\omega}$ (e.g., doubling the acquisition bandwidth decreases the SNR by a factor of $\sqrt{2}$). Decreasing the bandwidth below the Nyquist criterion naturally results in spectral aliasing.

Matters are further complicated by the fact that (for a given FOV) the frequency encoding gradient moment used during acquisition has to be constant, so doubling the bandwidth halves the time required for acquisition and therefore doubles the gradient amplitude required to maintain the same readout moment. As gradient fidelity scales non-linearly with amplitude, it may be found that decreasing the acquisition time by increasing receiver bandwidth causes distortions in the trajectory actually taken in k-space. Accordingly, the choice of receiver bandwidth is subject to several competing optimisation requirements associated with different processes, and the value chosen is likely to be machine specific.

2.4.6 Echo Planar Imaging (EPI)

EPI is a rapid imaging strategy in which a Cartesian path is taken to sample k space quickly, typically in a single shot following excitation. Several distinct variants of EPI exist, as a 2D k-space can be scanned in a raster by various paths. EPI is a technically demanding sequence, and these differences in implementation are accordingly of great practical importance.

2.4.6.1 Overview

A typical EPI imaging sequence is illustrated in Fig. 2.15. A large oscillating readout gradient is used to acquire individual lines of k-space in a raster, with small phase encoding blips providing movement

in k_y after each line of k_x is acquired. Note that k-space is sampled on both the positive and negative lobes of the readout gradient; the acquired data require flipping in order to re-grid them to a Cartesian raster. Two main alternatives to this strategy exist: (1) Fly-back EPI, in which the negative lobe of the readout gradient is made as large and as short as possible, and no data are acquired from under it, and (2) 'swept' acquisitions, in which the phase-encoding gradient is continually left on as opposed to being blipped. Both forms are used to reduce gradient demands or ameliorate other artefacts, but swept acquisitions traverse k-space in a zig-zag fashion and require a non-trivial regridding algorithm to map the acquired data onto a raster.

It is customary to define the echo spacing τ_e in EPI as the time between the centre of each echo (as illustrated in Fig. 2.15), i.e. the time between repeated crossings of the the $k_x = 0$ line. The effects of T_2^* decay are such that it is generally required to keep the total length of the acquisition within 2 or $3T_2^*$. For a given matrix size, this dictates the echo spacing, and accordingly the readout gradient amplitude G_RO, namely

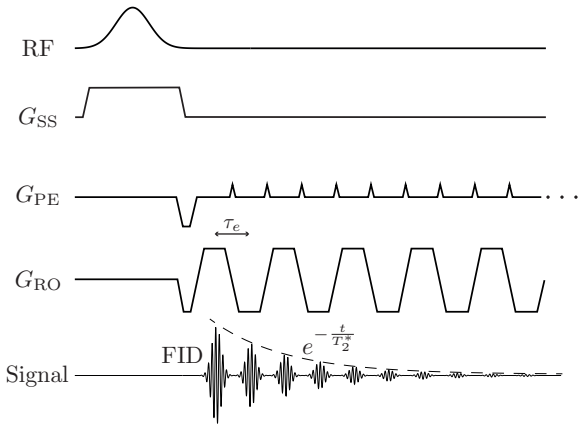

Figure 2.15: A typical EPI imaging sequence. Here a Gaussian pulse is used for excitation, and the dephasing gradients at the end of the acquisition are not shown.

$$G_\mathrm{RO} = \frac{2\pi}{\gamma L_\mathrm{RO} \tau_\mathrm{e}} \tag{2.61}$$

where L_RO is the desired field-of-view in the readout direction. As τ_e is the time between k_y samples at $k_x = 0$, it is appropriate to define an equivalent phase-encoding bandwidth

$$\Delta\omega_\mathrm{phase} = 1/\tau_\mathrm{e}. \tag{2.62}$$

As τ_e is commonly $\mathcal{O}(1\ \mathrm{ms})$, $\Delta\omega_\mathrm{phase}$ is far less than the equivalent readout bandwidth.

Owing to T_2^* decay, increasing the acquisition time hugely degrades EPI images. Consequently, it is usually the case that the oscillating readout train is made as short as possible, and is limited by the performance of the gradient set at hand. Accordingly, the readout train is usually a trapezoidal wave rather than a square wave, as the finite slew rate of the gradient set becomes a relevant limitation. In the limit as $G_\mathrm{RO} \to G_\mathrm{max}$ and $\tau_\mathrm{e} \to 2S_\mathrm{max}$ the readout train becomes triangular. This has an important consequence for sampling: when the gradient waveform is flat, samples acquired with a constant dwell time are linearly related to k_x and thus the acquisition time for each echo is

$$T_\mathrm{acq} = \frac{n_x}{2\Delta\omega} = \frac{2\pi n_x}{\gamma L_x G_\mathrm{RO}} \tag{2.63}$$

where n_x is the number of complex data-points along the readout direction (i.e. the matrix size), $\Delta\omega$ is the receiver bandwidth and L_x is the readout field-of-view. When sampling other than on the plateaux, i.e. on the ramps, however, k-space is not linearly traverse with linear temporal sampling. There exist two ways of ameliorating this problem: (1) re-gridding the nonlinear trajectory taken after acquisition, artificially spacing the edges of k_x, and (2) using non-constant sampling times, such that ramp sampling still produces a data-set that can be reconstructed by a Fast Fourier Transform. Note that T_{acq} is essentially what determines τ_e, and for these reasons a large receiver bandwidth is usually used in EPI to decrease the time of acquisition. This is associated with a penalty $\propto 1/\sqrt{\Delta\omega}$ in SNR due to Johnson noise, but, depending on the particular value of T_2^* this may be advantageous overall.

2.4.6.2 Artefacts

EPI is a pulse sequence that is particularly susceptible to magnetic field inhomogeneity effects. Firstly, a poor shim reduces T_2^*, limiting the available maximum echo time for the experiment, and hence limiting the matrix size. Secondly, variations in B_0 give rise to two other types of artefacts, signal loss artefacts and geometric distortions. Geometric distortions predominantly appear in the phase encoding direction, as $\Delta\omega_{\text{phase}}$ is (necessarily) very small. Accordingly, a small change in frequency results in an appreciable shift; if B_0 varies as $\Delta B_0(x,\, y)$, then the pixel at x, y is subject to a spatial shift Δr given by

$$\Delta r(x,\, y) = \frac{\gamma}{2\pi}\Delta B_0(x,\, y) N \tau_{\text{e}} \text{FOV}_{\text{pe}}, \tag{2.64}$$

where τ_{e} is the echo spacing in the readout train as above with N as the phase encoding matrix size. Signal loss artefacts are more complex, and arise from the variation of echo time across an image in the presence of inhomogeneities.[21] Essentially, in the presence of B_0 inhomogeneity, the behaviour of the loss artefacts become non-linear, and the direction in which k is scanned directly influences the resulting geometric distortions in the image. Accordingly, if alternate trajectories are taken in k-space, it is possible to quantify and remove these artefacts. In practice, this is commonly achieved by simply alternating the polarity of the phase encode gradient to scan k-space in a top-up or bottom-down fashion. An overdetermined estimate can then be made for the distortions, leading to their correction.[22] Moreover, this method can be used to correct for bulk position shifts arising from the incorrect setting of the transmitters' central frequency, as Δr oscillates sign on different acquisitions.[23]

Undoubtedly the most common artefact with EPI acquisitions is the so-called Nyquist (or $N/2$) ghost. EPI fundamentally relies upon a large oscillating readout gradient structure, and this gradient may produce several different errors that alternate with its polarity. For example, imperfections in gradient fidelity may effectively introduce a phase error $\Delta\phi$ between the positive and negative lobes of the readout. The signal received $S(k_x,\, k_y)$ can then be expressed as a degraded discrete Fourier transform of an 'ideal' image signal, $I(x,\, y)$, namely

$$S(k_x,\, k_y) = \begin{cases} \sum_x \sum_y I(x,\, y) \exp\left(-\frac{i2\pi x k_x}{n_x \Delta x \Delta k_x}\right) \exp\left(-\frac{i2\pi y k_y}{n_y \Delta y \Delta k_y}\right) \exp\left(-i\phi\right) & k_y \text{ even} \\ \sum_x \sum_y I(x,\, y) \exp\left(-\frac{i2\pi x k_x}{n_x \Delta x \Delta k_x}\right) \exp\left(-\frac{i2\pi y k_y}{n_y \Delta y \Delta k_y}\right) \exp\left(i\phi\right) & k_y \text{ odd} \end{cases} \tag{2.65}$$

where Δk_x, Δx etc. are the spacing between lines in either k-space or image space respectively and n_x, n_y are the matrix size. It can be shown, using the orthogonality of exponential functions, that an image I' reconstructed from S differs from I by

$$I'(x,\, y) = I(x,\, y) \cos\phi + iI\left(x,\, y - \frac{n_y}{2}\right) \sin\phi. \tag{2.66}$$

Correspondingly, this phase error splits the image into a real image correctly located and an imaginary ghost, aliased by half the field of view in the phase encoding axis, as illustrated for a circular object in Fig. 2.16. This error can be completely removed by estimating the phase error ϕ and multiplying alternate lines of k-space by $e^{\mp\phi}$ accordingly. In practice, 'navigator' echoes are acquired to improve this estimation, whereby a fast frequency-encoded projection along the centre of k-space is acquired in two directions before the start of the EPI readout proper. The phase error can then be estimated by comparing the two (nominally identical) sets of data.

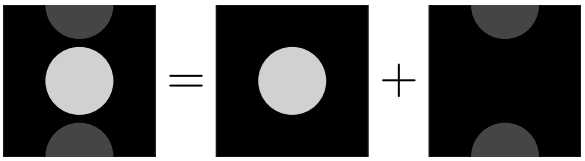

Figure 2.16: The 'Nyquist ghost' in which a ghost of the object is observed shifted by half the field of view in the phase-encoding direction. The resultant image is a sum of the real image, and imaginary ghost, corresponding to the two terms in Eq. (2.66).

In general, the effect of any error or imperfections in an acquisition can be Taylor expanded to give an oscillating phase term that is a function of space, rather than just being constant as in Eq. (2.65). Practically, "1D" phase correction is often the minimum required to reconstruct EPI images without a ghost, whereby $\phi = \phi_0 + \phi_1 x$. It has recently been shown[24] that 2D phase correction can ameliorate several other artefacts with EPI acquisitions, whereby a phase map is estimated and applied to the data. The Nyquist ghost is a well known problem with EPI, and historically led to the development of Flyback EPI, detailed above, which side-steps the issue. Alternatively, the odd and even lines of k-space can be treated as independent, yielding two (nominally identical) images with half the field of view of the full acquisition. These two images can then be added in magnitude, eliminating the ghost at the expense of the field-of-view of the acquisition.

2.4.7 Flow and Motion

The majority of MR imaging sequences used routinely to image mammals acquire data over a period of time that is comparatively long when compared to the rate of flow of blood within their circulatory

system, their heart rate, or the period of their respiratory motion. Accordingly, most MR sequences demonstrate sensitivity to flow and motion, which can lead to artefacts in many applications: for example, cardiac and respiratory motion render naïve gradient echo approaches to cardiac imaging effectively useless, with significant ghosting and signal loss artefacts present in reconstructed images. This sensitivity to motion is intrinsic, described below, and can be exploited for the quantitative imaging of blood flow or the motion of tissue.

Recall that the Larmor frequency of spins at a position $\mathbf{r}$ at time t in a region being imaged is given by

$$\omega(\mathbf{r},\, t) = \gamma B_0 + \gamma \Delta B_0 + \gamma \mathbf{r} \cdot \mathbf{G}(t). \tag{2.67}$$

If we assume that the pulse is played for a finite time $t \in [t_0,\, T]$, then the total phase accrued during this time is the integral of the above, or,

$$\phi(\mathbf{r},\, T) = \phi(\mathbf{r},\, t_0) + \gamma B_0 (T - t_0) + \gamma \int_{t_0}^{T} \mathbf{r} \cdot \mathbf{G}(t)\ \mathrm{d}t. \tag{2.68}$$

If we wish to consider the effect of flow and motion of a small voxel under the influence of the gradient $\mathbf{G}$, it is appropriate to first consider Taylor expanding its position as a function of time, obtaining $\mathbf{r} = \mathbf{r}_0 + \mathbf{v}_0 t + \frac{1}{2}\mathbf{a_0} t^2 + \ldots$ Here, $\mathbf{r}_0$ is the voxel's initial position; $\mathbf{v}_0$ its initial velocity, $\mathbf{a_0}$ initial acceleration, and so on. Performing this procedure on Eq. (2.68) we obtain

$$\begin{aligned}
\phi(\mathbf{r},\, T) &= \phi(\mathbf{r},\, t_0) + \gamma \Delta B_0 (T - t_0) + \gamma \sum_{n=0}^{\infty} \frac{1}{n!} \frac{\partial^n \mathbf{r}(\mathbf{r}_0)}{\partial t^n} \cdot \int_{t_0}^{T} (t - t_0)^n \mathbf{G}(t)\ \mathrm{d}t. && (2.69)\\
&= \phi_0 + \gamma \mathbf{r}_0 \cdot \int_{t_0}^{T} \mathbf{G}(t)\ \mathrm{d}t + \gamma \mathbf{v}_0 \cdot \int_{t_0}^{T} t\mathbf{G}(t)\ \mathrm{d}t + \frac{\gamma}{2}\mathbf{a}_0 \cdot \int_{t_0}^{T} t^2 \mathbf{G}(t)\ \mathrm{d}t + \ldots && (2.70)\\
&\overset{\text{def}}{=} \phi_0 + \gamma\big(\mathbf{M}_0 \cdot \mathbf{r}_0 + \mathbf{M}_1 \cdot \mathbf{v}_0 + \mathbf{M}_2 \cdot \mathbf{a}_0 + \ldots\big). && (2.71)
\end{aligned}$$

Here, we have defined $\mathbf{M}_i$ as the ith gradient moment, and $\phi_0 = \phi(\mathbf{r},\, t_0) + \gamma \Delta B_0 (T - t_0)$ as a combined initial phase term. Much information can obtained from Eq. (2.71): the stationary 'true phase' of the image, $\phi(\mathbf{r},\, t_0)$, can be related to the dielectric properties of the matter contained within it (yielding the fields of Electric Properties Tomography, Quantitative Susceptibility Mapping and Magnetic Resonance Electrical Impedance Tomography); the $\gamma \Delta B_0 (T - t_0)$ term allow, in the absence of motion, for the construction of maps of ΔB_0 simply by the complex voxelwise division of two images obtained at different echo times; and the use of gradients with nonzero moments allows for the reconstruction of kinematic information about the voxel in question.

Commonly, velocity information about moving spins is extracted by the use of bipolar gradients, ideally defined as

$$G(t) = \begin{cases} -G_0 & \text{for } 0 < t < \frac{T}{2} \\ +G_0 & \text{for } \frac{T}{2} \leq t < T \\ 0 & \text{otherwise.} \end{cases} \tag{2.72}$$

Such gradients have $\mathbf{M}_0 = 0$, and

$$\begin{aligned} \mathbf{M}_1 &= -\int_{t=0}^{T/2} G_0 t \ \mathrm{d}t + \int_{T/2}^{T} G_0 t \ \mathrm{d}t \\ &= \frac{G_0 T^2}{4}. \end{aligned} \tag{2.73}$$

Accordingly, the phase gained during this time for a spin labelled a moving at velocity v_0 is

$$\phi^a = \phi_0 + \gamma v_0 G_0 \left(\frac{T}{2}\right)^2 \tag{2.74}$$

$$\stackrel{\text{def}}{=} \phi_0 + \gamma v_0 \mathbf{M}_1^a + \ldots \tag{2.75}$$

If another dataset is acquired with a bipolar gradients of equal duration but differing $\mathbf{M}_1$, say $\mathbf{M}_1^b$, the resultant phase accrual ϕ^b is simply $\phi^b = \phi_0 + \gamma v_0 \mathbf{M}_1^b + \ldots$. Hence, subtracting the two separate images results in phase differences which are directly proportional to the difference in first moments and v_0, i.e.

$$\phi^a - \phi^b = \gamma v_0 (\mathbf{M}_1^a - \mathbf{M}_1^b) \tag{2.76}$$

$$\Rightarrow \Delta\phi \stackrel{\text{def}}{=} \gamma v_0 \Delta\mathbf{M}_1. \tag{2.77}$$

If v_0 is large, the resulting phase shift can be greater than 2π in magnitude. Accordingly, aliasing results, and in general it is not possible to reconstruct v_0. This leads to the definition of a maximum velocity encoding gradient, or VENC, defined as the velocity that produces a phase shift $\Delta\phi$ of π, radians, and is simply defined as

$$\text{VENC} = \frac{\pi}{\gamma \mathbf{M}_1}. \tag{2.78}$$

Conventionally, the VENC is chosen to be equal highest expected velocity in a velocity encoded scan. By iterating over the direction of $\mathbf{M}_1$, it is possible to directly map the velocity vector field of flowing liquids, which is used clinically to quantify blood flow throughout the body. A discussion of the uses of higher moments than $\mathbf{M}_1$ is outside the scope of this document.

2.5 | *Signal Enhancement Methods*

The inherent sensitivity limitations of magnetic resonance frustrate many of its potential applications. For example, *in vivo* NMR spectroscopy on nuclei other than $^1\mathrm{H}$ is exceptionally challenging, as the millimolar concentrations of compounds of interest combined with lower γ reduce the magnitude of $\mathbf{M}$ by several orders of magnitude. Correspondingly, it is traditionally necessary to perform a large amount of averaging, repeating a given experiment N times to gain a $\sqrt{N}$ increase in the SNR of the resulting data. For this reason, *in vivo* non-proton spectroscopic imaging is rarely performed as the resulting scan times are exceptionally long even at modest spatial resolution. Consequently, non-proton spectroscopy is not usually used as a clinical tool, but the surfeit of information it provides means that it finds applications, however challenging, in research.

These limitations all originate from the thermal equilibrium Boltzmann distribution of spins, derived and described in Section 2.1.1. It is possible to transiently generate a larger polarisation P than that found at thermal equilibrium, by pumping spins into a given state. Such systems are said to be "hyper[spin]polarised", and there are several documented methods for generating hyperpolarised substances. These include the "brute-force" method (cooling nuclei to mK temperatures before melting them); optically pumping nuclear spins using noble gasses; parahydrogen induced polarisation (PHIP) whereby a spin isomer of hydrogen is formed at low temperature before being subsequently chemically transferred to a molecule of interest; and the (Dissolution) Dynamic Nuclear Polarisation (DNP) method which is the subject of this work, and described below.

2.5.1 *Dynamic Nuclear Polarisation*

Consider a system consisting of an ensemble of nuclear spins ($I = \frac{1}{2}$) and electronic spins ($S = \frac{1}{2}$) placed in a homogeneous uniform magnetic field B at a temperature T. At thermal equilibrium we would expect that the population ratio for spins in each of their two available states to be strictly a function of temperature and B, namely Eq. (2.12) as described above.

It is possible to pump spins into non-equilibrium states by irradiating such systems with microwaves, leading to a system with a high degree of nuclear polarisation. This proceeds through a number of different nuclear mechanisms, and is collectively referred to as Dynamic Nuclear Polarisation (DNP). Such hyperpolarised spin systems generate a large ($> 10^4$) increase in the signal-to-noise ratio of magnetic resonance experiments that sample this population difference. The nuclear polarisation is instantaneously retained after a step-change in temperature, from 1 K to 37 °C (~ 310 K), letting the sample probe living organisms, before relaxing back to thermal equilibrium exponentially, with a characteristic time T_1 that depends on the molecule in question, temperature, and magnetic field.

There are two dominant mechanisms by which DNP elicits nuclear 'enhancement' (increasing P_n): "the solid effect", a historically dominant and comparatively well-understood mechanism that

was the main mechanism for the spin pumping of compounds used when DNP was investigated for applications in particle physics; and "thermal mixing", a term encompassing the thermodynamic balance of numerous other processes that can be driven to provide enhancement. For the glasses of small molecules considered in this work, the solid effect is not believed to be the dominant mechanism.

2.5.2 *The Solid Effect*

For the two-spin system described above, there are experimentally feasible regions of B and T where the majority of electronic spins will sit in the ground state, i.e. $P_e \approx 1$, and there is still no great nuclear polarisation, i.e. $P_n \approx 0$ (e.g. $B = 3\,\mathrm{T}, T = 1\,\mathrm{K}$). This is readily understood by considering a graph of P against T for electrons and a nucleus, illustrated for ^{13}C in Fig. 2.17. The dipolar spin-spin interaction permits two-quantum transitions between nuclear and electronic spin states. As shown in Fig. 2.18, this leads to four permissible transitions: first-order optical transitions between the spin states of the nucleus and electron occurring separately, with energy $\hbar\omega_\mathrm{e}$ and $\hbar\omega_\mathrm{n}$, and two-quantum dipolar interactions at $\hbar(\omega_\mathrm{e} - \omega_\mathrm{N})$ and $\hbar(\omega_\mathrm{e} + \omega_\mathrm{N})$. These transitions can occur because the dipolar interaction mixes the states of S and I into states which are not pure eigenstates of S and I. Specifically, it can be shown that the dipolar interaction

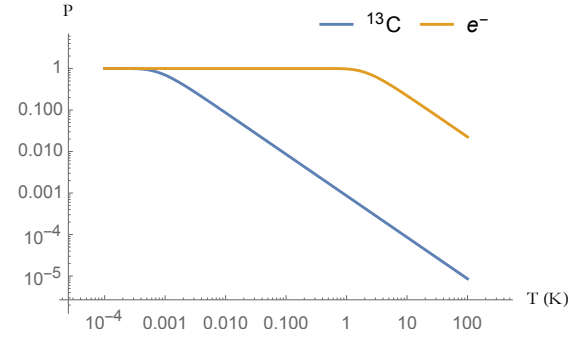

Figure 2.17: Nuclear and electronic polarisation P as a function of temperature T, illustrated for ^{13}C. The large difference in γ for electrons and nuclei leads to a large difference in P.

$$\mathcal{H}_\mathcal{D} = \frac{1}{2}\sum_{i,j}\frac{\gamma_i\gamma_j\hbar}{r_{ij}^3}\left(\mathbf{I}_i\cdot\mathbf{I}_j - 3(\mathbf{I}_i\cdot\hat{\mathbf{r}}_{ij})(\mathbf{I}_j\cdot\hat{\mathbf{r}}_{ij})\right) \tag{2.79}$$

mixes the four pure states $|a_0\rangle = \left|S = \frac{1}{2}\, I = \frac{1}{2}\right\rangle \overset{\text{def}}{=} |{+}{+}\rangle$, $|b_0\rangle = |{+}{-}\rangle$, $|c_0\rangle = |{-}{+}\rangle$ and $|d_0\rangle = |{-}{-}\rangle$ into the superpositions

$$\begin{aligned}
|a\rangle &= p|a_0\rangle - q^*|b_0\rangle \\
|b\rangle &= p|b_0\rangle + q|a_0\rangle \\
|c\rangle &= p|c_0\rangle + q^*|d_0\rangle \\
|d\rangle &= p|d_0\rangle - q|c_0\rangle.
\end{aligned} \tag{2.80}$$

The superposition coefficients p and q can be approximated through first-order perturbation theory, and it is found that $p \approx 1$, $q \ll 1$.[25] By considering transitions between these states and sources of relaxation (i.e. transitions towards thermal equilibrium) in an isolated electron-nucleus pair, it is possible to derive rate equations for the system as a whole, when irradiated with microwaves at

$\omega_e - \omega_N$. Additionally, the characteristic time for relaxation, T_1, can be estimated as a function of temperature, magnetic field strength, and polarisation. Analytic expressions for the resulting maximum nuclear polarisation can subsequently be derived, and this polarisation is found to scale linearly with temperature.

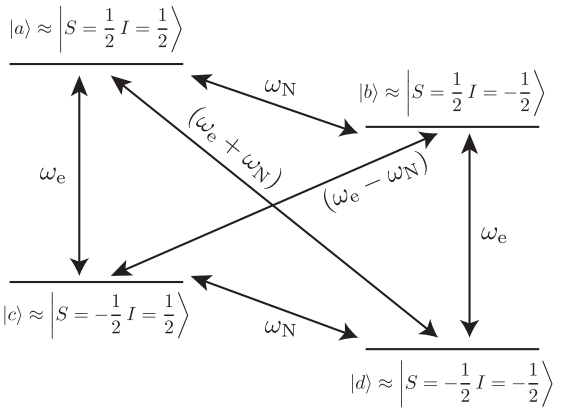

Figure 2.18: Energy levels and possible transitions in an isolated spin-$^1/_2$ electron-nucleon system. The four states $|a\rangle$ to $|d\rangle$ are approximately pure spin-states, weakly mixed by the dipolar interaction. The 'forbidden' transitions at $\omega_e \pm \omega_N$ can hence occur.

If the nuclear-electronic system is found in a glass, promoted electronic spins are typically able to relax quickly ($T_{1e} \approx 10^{-3}$ s), whereas nuclear spins at low temperature have a longer spin-lattice relaxation time (T_{1N}) on the order of hours. Under these circumstances, an appreciable fraction of nuclear spins can be pumped into one of the two eigenstates by irradiation at $\omega_e \pm \omega_N$. The process is analogous to a three-state laser: irradiation at $\omega_e - \omega_N$, say, will flip a $|c\rangle$ state to a $|b\rangle$ state. As the electronic T_1 is short, this quickly decays to the state $|d\rangle$, emitting a photon of frequency ω_e. An analogous process occurs for irradiation at $\omega_e + \omega_N$, pumping into the state $|c\rangle$ from $|d\rangle$.

If the electronic linewidth is small in comparison to the nuclear frequency (ω_N), irradiation at $\omega_e + \omega_N$ can only promote one transition, as the other is off-resonance with the driving frequency. This situation is known as the 'well-resolved solid effect', and usually requires considerable microwave power ($\mathcal{O}$(W)) due to the two-quantum transition. The NMR peak amplitude as a function of irradiation frequency, known as a 'microwave sweep', is predicted to be symmetric for a purely solid effect process, and have two Lorentzian peaks[25,26]

2.5.3 *Thermal Mixing*

Most real materials feature an ensemble of electronic spins in different environments, broadening the effective electronic linewidth. In most cases, this linewidth is not small in comparison to ω_N, and the concentration of electrons is such that spins can interact with each other. This leads to a second, complementary mechanism through which DNP can occur. Essentially, in a complex system numerous nuclear processes serve to mix the spin populations, and this mixing occurs primarily due to thermodynamic considerations. It is therefore known as thermal mixing. Owing to the fact that an inverted electronic spin population is produced by irradiation with microwaves, the system has negative temperature and therefore is driven to cool. This cooling is realised by the pumping of nuclear spins into an unfavourable entropic configuration, resulting in enhancement and increasing P_n. As the first of

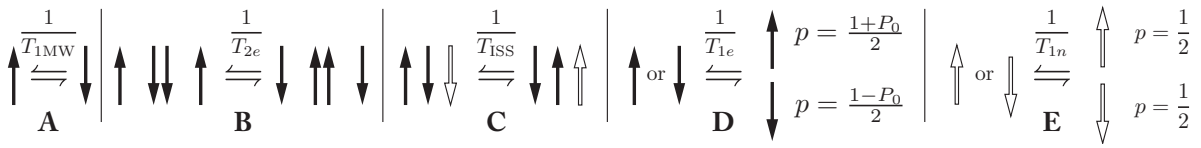

Figure 2.19: The five electronic (black arrow) and nuclear (white arrow) processes that predominantly contribute to thermal mixing. **A**: Electronic spin pumping due to MW irradiation; **B**: Spectral diffusion between electrons through the glass; **C**: the Hartmann-Hahn ISS process; **D** and **E**: Electronic and nuclear relaxation processes back towards thermal equilibrium. It is assumed that, at the temperature considered, the nuclear thermal equilibrium polarisation is still very close to zero, but that of the electrons is not; here it is denoted P_0.

these steps ('dynamic cooling') involves a single quantum transition this effect usually dominates if it can occur.[26]

Consider irradiating a sample in a magnetic field with photons close to, but distinct from, the electronic Larmor frequency, at a frequency ω, where $\Delta \stackrel{\text{def}}{=} \omega - \omega_\text{e}$ is small. It can be shown that the effective Zeeman Hamiltonian in this situation is time-independent in a reference frame rotating at the frequency of the applied irradiation. This implies that the system can be characterised in this frame, and that the concept of a 'spin temperature' in the rotating frame, T_R is well defined. Due to the broad electronic linewidth, it is possible to change the energy of the system by an amount $\hbar\Delta$, flipping a spin in a particular environment. So, if $\Delta < 0$, then such a decrease of the electron spin system's energy leads to a lowering of the system's temperature. If $\Delta > 0$, then T_R can become negative, as the spin population inverts. Sources of relaxation between the electronic spin system and the lattice couple the lattice temperature and T_R, leading to the relation

$$T_R = \frac{\Delta^2 + aD^2}{\omega_\text{e}\Delta} T \tag{2.81}$$

where T is the temperature of the lattice (i.e. cryostat), a is a model-dependent parameter on the order of 2, and D the dipolar second moment of the electronic linewidth.[25]

The second step, 'thermal mixing in the strict sense', is the coupling of this electronic spin temperature to the nuclei of interest in the laboratory frame of reference. This is a thermodynamically favoured relaxation process involving the simultaneous transition of two electron spins and one nuclear spin ($\delta m_{S1} = \pm 1$, $\delta m_{S2} = \mp 1$, $\delta m_I = \mp 1$), which proceeds through the Hartmann-Hahn mechanism, which requires a broad electronic linewidth so that the difference in energy between two electrons from the distribution is equal to the $\hbar\omega_N$. Such transitions can only occur if there is a mutual interaction between electron spins, and between an electron spin and a nuclear spin. The rate of this transition therefore increases with electron concentration. From thermodynamic considerations, one can see that this step will equalise the spin temperatures of the polarised nuclei and the electron system.[26–29]

Recent progress in understanding thermal mixing has been made by assuming that only the five nuclear processes shown in Fig. 2.19 contribute to the process.[30,31] Briefly, these are electronic spin

pumping due to the applied microwave irradiation (**A**) with a rate $1/T_{1\mathrm{MW}}$; spectral diffusion between neighbouring electrons at a rate $1/T_{1e}$ (**B**); the Hartmann-Hahn ISS process (**C** with rate $1/T_{\mathrm{ISS}}$), and longitudinal relaxation for both nuclei and electrons (**D** and **E**, relaxing at a rate $1/T_1$). Each of these rate constants can be measured or approximated analytically, and the rate equations for the system are first order in these processes, and conserving spin and energy. As reported by Serra *et. al*,[30,31] derived quantities from these equations match experimental values. However, it might be argued that this approach does not fundamentally improve understanding, as calculating rate constants *a priori* is not straightforward, and depends on parameters such as the exact geometry of the experimental apparatus and its dielectric properties. This approach is therefore likely to be most profitable if it is at least quasi-phenomenological, using experimentally determined rate constants to identify experimental regions where the maximal enhancement given by an experiment is higher.

In practice, thermal mixing remains a difficult example of mesoscopic physics, in which thermodynamic considerations drive a large disordered ensemble through a variety of processes.

2.6 | *Summary*

This chapter has presented a brief overview of the physics of nuclear spin from a non-relativistic perspective, deriving the Larmor equation by considering selection rules and the conservation of energy. The Bloch equations are then derived from both a semi-classical model, and from Schrödinger equations, and the validity of the semi-classical model for $\mathbf{M}$ discussed. NMR and MRI are introduced. Excitation is discussed in depth, and the excitation k-space formalism and SLR algorithms are derived and presented. Spatial encoding by frequency encoding and phase encoding gradients is discussed, and imaging k-space is derived. The concept of a pulse sequence is introduced, and illustrated with a slice-selective gradient echo. The Dynamic Nuclear Polarisation method of increasing signal by pumping nuclear spins into a given eigenstate is introduced and its mechanisms discussed.

Chapter 3

Hardware: Characterisation and Development

The implementation of the principles discussed above is of great practical importance. Hardware limitations and idiosyncrasies serve to dictate the usefulness of different pulse sequences, their maximum resolution, and the type and severity of any artefacts that might arise in the resulting images. All experiments described in this text were performed on a 7 T (300 MHz) Varian/Agilent Technologies[1] 'DDR' preclinical MRI scanner, which is described in detail over the following pages. To perform the studies described in subsequent chapters, it has been necessary to construct a variety of after-market components for use with the scanner, and also to undertake various aspects of repair and maintenance of the system described.

3.1 | *Instrumentation: A detailed overview*

The 7 T MRI system consists of a superconducting magnet and gradient set (described in Sections 3.1.1 and 3.1.2) connected to 'the acquisition console', a set of disparate Xilinx FPGAs and PowerPC Processors that together perform the many constituent tasks required of an MRI scanner: namely generating gradient and RF waveforms digitally, controlling shims, and receiving and digitising incident signals. A variety of RF coils plug into the acquisition console's "RF front end", and are described in detail in Section 3.2. The console's master controller ("`master1`") runs a realtime operating system (VxWorks) that is responsible for ensuring that specified tasks happen at the correct time. The console receives assembly-level instructions from, and communicates with, a "host computer" with a user-interface that is responsible for data reconstruction, image processing, study planning, pulse-sequence generation, and data storage. This host computer runs RedHat Enterprise Linux and the Varian NMR software `VnmrJ`. The logical organisation of the system is discussed in Section 3.1.5. The flow of information between the different constituent parts of the system is illustrated in Fig. 3.1 in the form of a block diagram that also summarises this overview.

[1] Varian, inc. were a manufacturer of scientific equipment, and historically developed the first commercial NMR spectrometer, in 1948. In 2010, Agilent Technologies bought Varian for $1.5 bn, before ceasing production of their MRI products (2012) and later NMR products (2014). Further details can be found in the *Journal of Magnetic Resonance* article mourning their decision.[32]

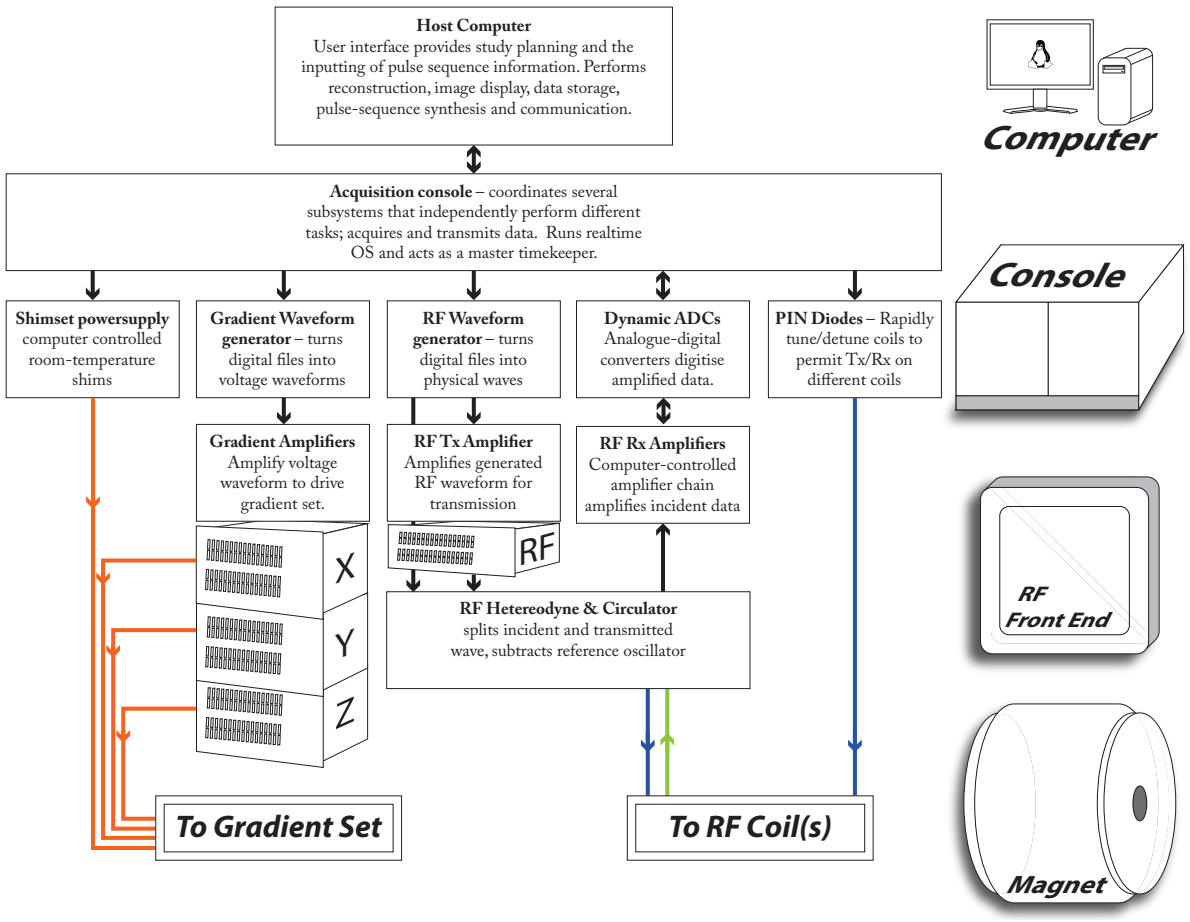

Figure 3.1: The high-level organisation of the system, together with cartoons of the relevant hardware.

3.1.1 *The static field*

The key requirement for magnetic resonance experiments is a stable, homogeneous and powerful magnetic field. This static field is provided by a 7 T superconducting Magnex[2] magnet, with an internal horizontal bore size of 155 mm. It is of an open-system 'twin can' design, with a large helium cryostat (420 L) surrounded by a smaller nitrogen cryostat (200 L) to reduce the helium boil-off rate. For engineering reasons, the 'twin cans' are actually six suspended pressure vessels, with the extra layers being radiation shields and vacuum insulation spaces. Each can be accessed independently from the central turret service port at the top, which is ordinarily sealed with a 3 mm indium wire. Exhaust helium gas is re-liquefied by the Department of Physics Central Cryomagnetics Facility. The central field of the magnet is provided by 252 A carried in superconducting NbSn wires wound through pure copper in a six-coil arrangement that is similar to a Helmholtz pair; their overall inductance is 32 H. Accordingly, the energy stored in B_0 can be estimated as $\frac{1}{2}LI^2 \approx 1$ MJ, or that produced by exploding ~ 250 g of TNT. Should any part of the superconducting B_0 coil become resistive this large

[2]Magnex Scientific, founded in 1982, were a pioneering high-field superconducting magnet design company, and the main competitor of Oxford Instruments who had a large number of historical 'first-to-market' inventions. In 2004 they were acquired by Varian, Inc, and subsequently disbanded by Agilent in 2014.

amount of stored energy is dissipated very rapidly as heat (as the resistive heating power $P = I^2R$ is instantaneously very large), vapourising the large volumes of liquid helium that surrounded the coil in a process referred to as a 'quench'. As liquid helium has a low latent heat capacity, and expands ~ 800 times when boiled, the process is explosive in nature. A graphite bursting disc and vent pipe exist to duct the large volumes of (asphyxiant and cold) helium gas out of the building and room containing the magnet. To avoid catastrophic damage to the magnet, the superconducting wires are woven through a larger copper matrix that can carry the large current resistively and dissipate it should a quench occur. A superconducting switch — an independently controllable heater wound around superconducting wire — is provided to allow for the (re-)energisation of the magnet with a high current power supply. This magnet quenched twice in the period 2012-15.

The homogeneity of the system is enhanced by the calibrated addition of small perturbing magnetic fields that vary spatially but are constant in time. This process is called shimming, and is detailed in Section 3.4. A comprehensive superconducting shim-set aims to ensure a field homogeneity of < 1 ppm within 75 mm of isocentre, the designed centre of the magnetic field and the origin of the coordinate system used within the magnet.[3] The drift rate of B over time is low, at $\ll 0.05$ ppm/hr. Superconducting shims are rarely adjusted over the lifetime of the magnet, but a room-temperature shim-set is used to ensure B_0 homogeneity that is adjusted for every experiment.

3.1.2 Gradient coils

A medium performance preclinical imaging set of gradient-generating coils, known as 'the gradient set', is located within the 155 mm bore of the cryostat, also manufactured by Magnex. It is water-cooled, of outer diameter 155 mm, inner diameter 100 mm, and it contains a room-temperature shimset along with temperature probes and water flow meters. The exact physical design of a manufacturer's gradient set is usually a "fingerprint" pattern guided by many iterations of finite element analysis and computer aided manufacture, but the approximate designs are well-known: z gradients are produced by a driven Helmholtz pair equidistant from isocentre; x and y gradients are produced by pairs of double-saddle Golay coils deformed to fit around a cylindrical former. As illustrated in Fig. 3.2, the x and y coils are identical, but rotated 90° with respect to each other; z are distinct. A distinction is made between the hardware, physical xyz axes and coils and the logical read/phase/slice axes. One key utility of MRI is that these logical axes can be orientated arbitrarily with respect to the physical axes. For example, it is common to image organs along their natural axes, such as the assessment of cardiac function by acquiring images along the short axis of the heart; or the imaging of the intervertebral disks of the spine with a coordinate system whose slice axis follows the spinal cord.

[3] Following the magnet quench in early 2015, the author was able to adjust the superconducting shims to ensure ~ 0.5 ppm homogeneity on a 75 mm diameter sphere surrounding isocentre.

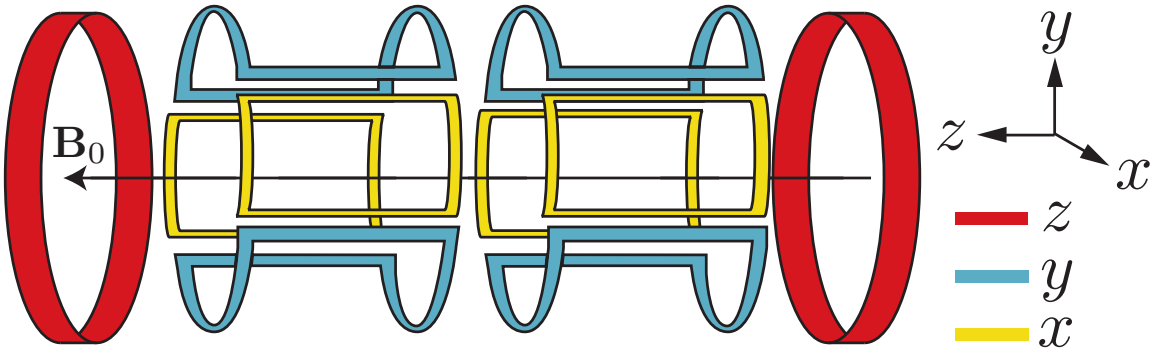

Figure 3.2: A cartoon diagram of a gradient set. The z gradient coils are shown here as a Helmholtz pair; x and y are identical but rotated $90°$ relative to each other and each consists of four saddle-coils driven independently.

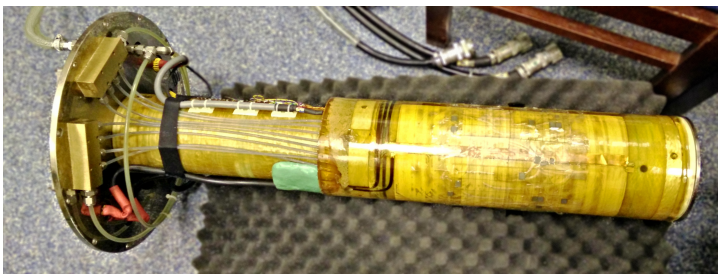

Figure 3.3: A photograph of the gradient set used, taken when it was removed following a magnet quench. The set is constructed on GRP formers and then embedded in epoxy resin, giving it its orange colour. The orientation is the same as in Fig. 3.2, with positive B_0 extending to the left. The fingerprint design of the coils can just be resolved, along with water cooling distribution channels, ferrite shims and a large quantity of Kapton tape.

The gradient set is heavy, and mechanically fixed to the cryostat bore to minimise unwanted motion. The resistance of the gradient set is usually kept as low as possible; here it is less than 1 Ω on each of the three channels. As a field is produced by driving a large current through each coil, and resistive heating is given by $P = I^2R$, the set dissipates a substantial electrical load as heat – correspondingly, cool water is recirculated continuously through the apparatus. This cooling requirement is the ultimate limitation on the maximum strength of the gradient set, as irreversible damage can occur if an acceptably low steady-state temperature is not reached. The set is made of copper deposited on a glass-reinforced plastic substrate as opposed to the alumina-ceramic composite favoured for its stiffness and thermal properties in clinical systems. The set is fused, and has a relatively low inductance (~ 0.3 mH) to minimise the time lag between V and I; the maximum rated current in the set is 15 A steady-state or up to 90 A transiently, and the maximum voltage is 120 V. There is also a limit on the maximum degree of switching experienced by the gradient set, a parameter referred to as the 'duty cycle'. Its precise definition varies by manufacturer, but conceptually it is the ratio $t_{\text{on}}/t_{\text{off}}$. On the Varian system here there are three different duty cycles, computed as 'running integrals' over different time lags. A photograph of the gradient set used is provided in Fig. 3.3.

As described in more detail in Section 3.1.5, gradient waveforms are synthesised from digital files by a custom Varian board, which consists of Xilinx Virtex 4FX FPGAs with an on-board PowerPC 405 CPU which runs the real-time operating system VXWorks. Its resolution is 4 μs, that is, each waveform is specified as a series of points on a 4 μs raster. Behaviour between these points is not specified, but it is believed to be smooth and continuous. The board outputs three gradient waveforms that lie in the range $0 - 10$ V, and can source only a few mA. The x and y axis gradient coils are driven from this output by two AE Techron 7792 amplifiers in voltage-slave mode. The maximum power output of these amplifiers is a nonlinear function of their duty cycle, and they internally consist of bipolar NPN power transistors paralleled 20 times in a bridge amplifier topology, with additional protection and linearity circuitry. Following an unfortunate fire in 2011, the z amplifier is worse than that on the x or y channels (a model 7790 AE Techron), has a more limited duty cycle, and a lower maximum current output, but is otherwise identical. Each amplifier is connected to a 415 V three phase power supply.

3.1.3 Gradient characterisation

An ideal gradient set for MRI would be arbitrarily strong in a given axis, and switch on or off instantaneously. Owing to Maxwell's equations, the limitations on the properties of matter, and the finite speed of light, this is never physically possible. Gradient sets therefore have two main design limitations: their maximum strength $G_{\max}$, and the time taken to reach this maximum strength (which is commonly expressed in units of G/cm or mT/m), known as the *rise time* (in μs). The maximum *slew rate*, $S_{\max}$ is defined as the maximum available $\frac{\mathrm{d}G}{\mathrm{d}t}$, commonly has units G/cm/ms or T/m/s, and represents the effective minimum 'switching time' for the gradient set. Our system has a current response of 2.5 mT/m/A, for a theoretical maximum steady-state gradient strength of 375 mT/m (37.5 G/cm), but this is reduced by the available amplifiers to 175 mT/m. The rise time is 200 μs, and the maximum slew rate is 875 mT/m/ms (87.5 G/cm/ms). The maximum gradient duty cycle is anisotropic across the three axes, reflecting the non-identical gradient amplifiers used.

Faraday's law states that a changing magnetic field induces a current. As gradient coils are created to temporally perturb magnetic fields, it is unsurprising that unwanted eddy currents are induced in nearby metallic structures, such as the cryostat body. These eddy currents serve to screen the field created, and exponentially die away with a time constant that depends on the material through which the current is flowing. The development of *actively shielded gradients*, in which each gradient coil is accompanied by an exterior 'screening' coil, driven in the opposite direction, reduces the magnitude of cryostat eddy currents (which have a long time constant) by two orders of magnitude or more.[33, Ch. 60] However, there are still a large number of residual eddies arising from imperfect screening in the metallic components of the scanner. To this end, modern scanners provide Eddy Current Compensation (ECC) systems, wherein the behaviour of the gradient set is linearised by distorting the waveform amplified.

The ECC system works by introducing a series of exponentially decaying current compensations with different time constants that depend on axis of the gradient played, and scale linearly with the amplitude of the played gradient.

As part of the system's installation, a tensor of correction terms is computed following measurement of the gradient set's behaviour. This tensor contains

- Six time constants and amplitudes for the three "direct" excursions (e.g. x compensation following a gradient on x, $x \rightarrow x$);
- Three time constants and amplitudes for the six "cross-term" excursions (e.g. z adjustment following a gradient on x, $x \rightarrow z$);
- Four time constants and amplitudes for the four "B_0" excursions: eddy currents can act to screen or enhance the static field everywhere, and a B_0 shim coil is pulsed to correct for this, leading to the three linear terms, e.g. B_0 compensation following a z gradient. Unfortunately, this B_0 coil itself produces eddy currents, which are compensated for with a $B_0 \rightarrow B_0$ term.

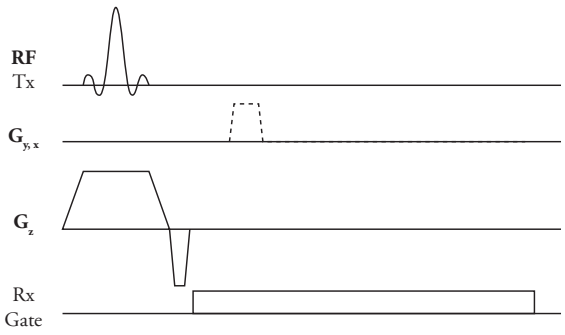

Figure 3.4: The pulse sequence used for measuring our GIRF. Slice selection occurs on an axis, before a chirp is played on another axis. The process is repeated with no chirp. The whole sequence is then repeated on other combinations of axes.

The measurement of these probably-not-independent 104 parameters is time-consuming and difficult. For this reason, these terms are measured once at G_{max}, and not usually adjusted.

Unfortunately, these eddy-current corrections are not perfect, and, moreover, the frequency response of the gradient set is usually a (non-linear) function of gradient amplitude. Correspondingly, these ECC methods ameliorate the problem of gradient non-linearities, but do not eliminate them. This is particularly troublesome if a complex path is being rapidly traversed in k-space, because a) the higher gradient demands associated with rapid imaging make it more likely that gradient non-linearities are present, and b) those non-linearities mean that there is a discrepancy between the path believed to be taken in k-space, and the actual path taken. This results in artefacts in the final image. Several methods have been proposed to correct for gradient distortions, either by directly measuring the current in the gradient coil,[34] or by using phase-sensitive magnetic resonance techniques. It was found that the nonlinearity of the gradient set used in this work was sufficient to introduce substantial artefacts in both excitation and image k-space. Consequently, the method of Duyn *et. al*[35] was used to calculate the *gradient impulse response function* (GIRF), which maps the amplitude and phase response of the gradient set as a function of the frequency of its inputs. A homogeneous

sample was placed in the magnet before the sequence shown in Fig. 3.4 was played, which consists of slice selection on a given axis followed by a chirp $G_r(t)$ with a large Fourier spectrum on another axis. Different chirps were played to cover a large spectral domain. A control measurement of the sample was then taken with just the slice-selection gradient. The process was repeated for all combinations of axes. The difference $\Delta\phi$ of the phases of the two measured MR signals is given by

$$\Delta\phi_r(t) = \int_0^t \gamma G_r(t) \cdot D_r \, \mathrm{d}t = D_r k_r(t) \tag{3.1}$$

where D_r is the distance of slice to the gradient isocentre. By measuring $\Delta\phi$ for a combination of gradient strengths the taken trajectory k_r can be derived by normalising $\Delta\phi_r(t)$ to D_r. An impulse function can then be built up by looking at the spectral response of the gradients to different stimuli, i.e. by computing

$$\mathrm{GIRF}(\nu) = \frac{1}{N} \sum_{\text{Waveforms}} (\mathcal{FT}\{G_{\text{Nominal}}(t)\}^{\dagger} \mathcal{FT}\{G_{\text{Measured}}(t)\}), \tag{3.2}$$

where $\mathcal{FT}$ represents the Fourier transform of the gradient waveform $G(t)$ either as measured or as nominally requested, ν is frequency, and N is a normalising constant. These measurements were obtained, and the GIRF was computed, both with the help of Angus Lau, and are summarised in Fig. 3.5. The GIRF is accurately able to predict the trajectory taken following a sharp gradient, and was measured at $G_{\text{max}} = 3, 6, 10, 15, 17.5$ G/cm at a variety of different slice positions (TR $= 1$ s, TE $= 11$ ms, 64 averages, eight 2 mm thick slices, 1.7 ms sinc excitation; 3.05 mm gap on a spherical phantom; 16.384 ms readout). The GIRF itself was used to correct the k-space trajectory of other experiments in this thesis, and accurately describes the limitations of the gradient set.

3.1.4 Radiofrequency wave generation, amplification and reception

In addition to gradient coils and a superconducting magnet, the other key part of an MRI system is its radiofrequency subsystem. Equipment is required to synthesise and amplify a frequency-varying shaped electromagnetic wave, transmit it into the sample, receive an induced signal, and amplify and digitise this incident wave. The Varian system features dedicated FPGA/PowerPC boards for RF waveform generation and digitisation.[4] There are two RF controllers, which synthesise waveforms for transmission on ^{1}H or nuclei other than protons respectively. It is possible to transmit on both channels at once, a feature used for multi-nuclear decoupling NMR experiments that probe spin-coupled systems. Each board is believed to drive a Programmed Test Source (PTS) waveform generator that synthesises a 1 V waveform from a digital input; the selection of transmitters is largely transparent to the user and the

[4] `172.16.0.20` or `rf1` for the ^{1}H transmitter, `172.16.0.21` or `rf2` for the X-nucleus transmitter; `172.16.0.50-53` or `ddr1-4` for the reception. A common reason for an upload to not complete successfully following an experiment is buffer overflows on `ddr1`, which can be relieved by either changing the looping structure and/or programming acquisition explicitly if not already doing so.

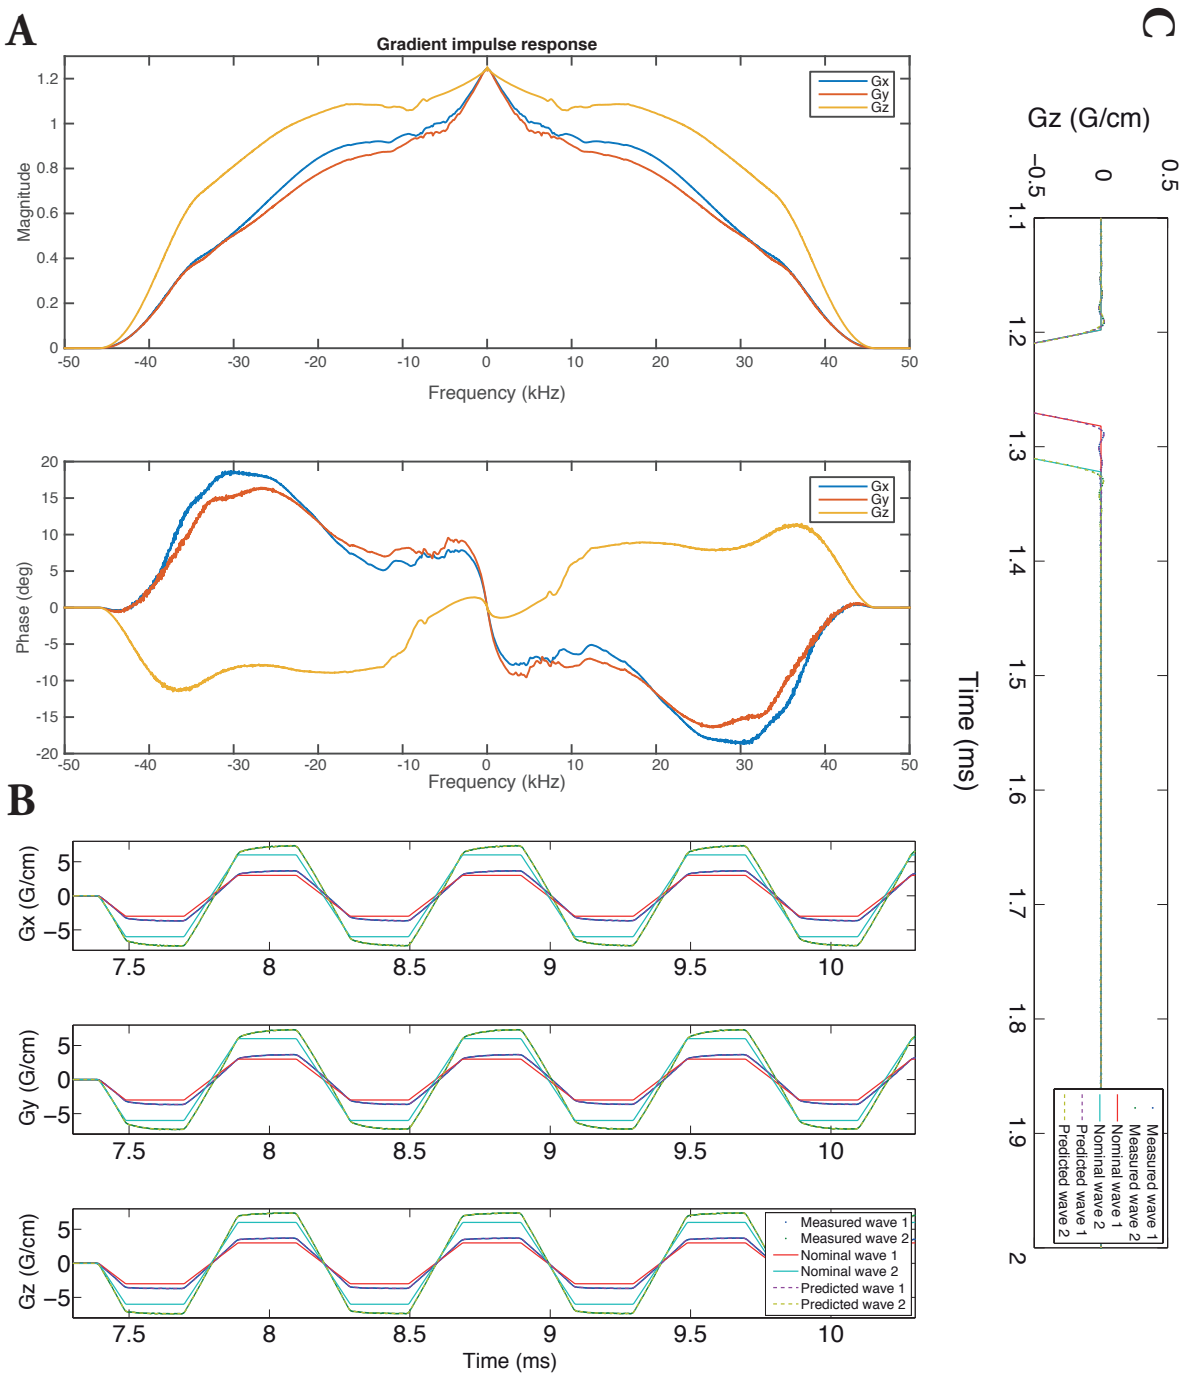

Figure 3.5: The measured response properties of the gradient set used in this work. **A**: The magnitude and phase of the (complex) GIRF, shown as a function of frequency. The response is approximately uniform for inputs < 1 kHz, whereupon distortions start. The phase lag is the greatest confound for rapid imaging strategies, as it introduces a time lag in the trajectory taken (i.e. turning $\mathbf{k}(t)$ into $\mathbf{k}(t + \delta)$. Anisotropy on all three physical axes is also visible. **B**: The effect of these gradient imperfections is to distort a rapidly oscillating gradient waveform. Illustrated here are the nominal (i.e. user-supplied trapezoidal) waveforms and their physical realisation (which are distorted). The distortion is worse for increasing gradient amplitudes. Also shown are the predictions from the GIRF of the gradient behaviour; they are superimposed on the trajectory taken. **C**: In part, this behaviour may be due to the imperfect impedance matching of the gradient set and the amplifiers. Shown here is the response following a single sharp pulse, where some degree of ringing is observed near zero. This "underdamped" behaviour may represent an impedance mismatch between the amplifier and the gradient set. The ideal, measured and GIRF-predicted trajectories are shown for waveforms of two different amplitudes and on each axis.

overall range of the system is 6-320 MHz with 0.1 Hz resolution. Digital waveform synthesis devices can either operate in a phase-continuous or frequency-continuous mode when switching frequency; it is unknown in which mode the controller operates, and this has implications for the RF pulses discussed later. The digital resolution of the generator is, at best, 200 ns, and the transmission board has > 30 MiB[5] of available waveform memory, which limits pulses to ~ 500 ms in length. A reference oscillator is also generated, whose frequency is 20 MHz above the Larmor frequency of the nucleus under investigation (or close to it), e.g. 320.2877721 MHz for ^{1}H at 7.053 T. The high resolution of the frequency generation system – typically 1 part-per-billion or better – is a key requirement for magnetic resonance experiments.

After synthesis, the RF waveform is transmitted to a 1 kW amplifier with variable gain, before being put through a variable digital-gain attenuator, controlled by the RF transmission board. The two steps of amplification/attenuation allow for a large dynamic range in the transmitted power of the RF. This is used to calibrate the flip angle delivered in a given experiment (c.f. Eq. (2.27)). Following amplification, the RF is delivered to an antenna, known as a radiofrequency coil, that transmits it to the sample. Radiofrequency coils are discussed in detail in Section 3.2: they serve to generate B_1^+, the transmitted RF, and usually also receive incident energy from the system, B_1^-. To allow multiple coils to be used, a PIN diode decoupling box is connected to the RF controller board. PIN diodes have rapid switching times, and therefore can be used to switch in or out a given coil within ~ 1 μs.

After transmission of RF, a current is induced in one or more RF coils. In comparison to the excitation voltage on the order of 1-2 kV within the coil, the received signal is on the order of µV. Hence, there is a need to separate the transmitted signal from the received signal before it can be amplified, and to ensure that the sensitive amplifier chain is not damaged by large voltages. The Varian system uses an RF circulator – a three-port passive electrical component that splits incident and transmitted waves. Traditionally, diode pairs or similar protection circuitry are also used to short the sensitive receiver out when the transmitter is on. Unfortunately, such circuitry takes a nonzero time to switch, resulting in a delay between the end of the RF transmission and the acquisition of the FID; this delay is minimised on modern systems. Following reception, the small µV signal is amplified by a low-gain ultra-low-noise GaAs FET broadband pre-amplifier and amplifier chain whose gain is computer-controlled. An oscillator driven 20 MHz *above* the centre of the receive frequency is then mixed in analogue with the amplified signal, removing the majority of the rapid Larmor precession. This signal is then digitised directly by a 24 bit 80 MHz ADC, about the 20 MHz 'intermediate frequency' window provided by the mixer. Digital signal processing techniques implemented on the fabric of the receive controller FPGAs then filter noise arising from outside the spectral window of the experiment, a process known as digital lock-in amplification, and a temporally accurate filter is able to ensure a flat spectral baseline without a 'DC offset' in the amplitude in the FID. The resulting digital dynamic range of the experiment is a

[5] All units of information in this thesis are assumed to be bits base two rather than the SI base ten; and appropriate units are used to differentiate between them, e.g. 1 MiB = 1048576 bytes; 1 MB $= 10^6$ B.

function of the receive bandwidth, e.g. 20 bit dynamic range at 10 KHz. This is usually sufficient. Four independent receivers are provided.

This process should be contrasted to traditional quadrature detection of the FID, in which a series of analogue components mix the received FID with two centre reference oscillators shifted in phase by 90° by hardware devices. These two resulting components – real and imaginary – are then analogue filtered before being digitised by separate ADCs, where the centre frequency is at 0 Hz or DC. Unfortunately, this process inherently introduces a $\sqrt{2}$ SNR penalty, along with several additional artefacts that result in an unknown DC offset to the FID signal, and also introduce a further time delay. The use of a digital receiver as soon as possible after reception of the FID is known as 'direct digital reception', from which the console takes its model name ('Varian DDR'). By capturing the NMR signal closer to the coil, one not only gets a signal with fewer distortions, but also a signal with less noise from analogue receiver stages.

Finally, some RF coils are 'quadrature coils', that is, they contain two separate elements with a fixed orientation that record M_x and M_y separately. The voltages induced on each coil are 90° out of phase with each other. The receiver hardware contains a user-accessible interface to allow direct digitisation of these two independent signals where a quadrature coil is used, improving SNR by a factor of $\sqrt{2}$. A hybrid coupler is used to direct the transmitted and incident RF appropriately. If a quadrature coil is not used, this second '$\lambda/4$' input is shorted out. The processes of RF reception and amplification for both DDR systems and the traditional quadrature detection methods are illustrated in Section 3.1.5.

3.1.5 Logical organisation

The MR system is largely modular, with a series of interconnecting components that generate different aspects of a pulse sequence. The controller board `master1` is responsible for instructing each of these subsystems. Naturally, controlling a modern imaging system by a command line that changes the stack of a real-time operating system would not make for an ideal user interface. Accordingly, a 'host computer' is responsible for 'driving' the system, sending compiled bytecode for each individual controller as a binary package (`acode`) encapsulated as serial over ethernet. This `acode` is generated by the scanner before the scan starts by a hierarchy of other processes, summarised in Fig. 3.7. The compiled `acode` can contain control statements that allow for real-time decisions to be made on the basis of the value of one of forty '`rtvars`', realtime integer variables.[6] A `C++` library on the scanner, `psg`, acts as an application-binary interface by providing functions such as `oblshapedgradient()`, which is used to play gradients, and transcoding their arguments into appropriate `acode` for the controllers. Individual pulse sequences are written in `C`, consist of a time-ordered set of such statements, and are linked against this library. They reside in `psglib` as source and are compiled as `seqfil`s in `seqlib`. The

[6] Another useful undocumented fact is that the contents of each initialised `rtvar` can be printed at runtime by `master1` if the commands `DebugLevel=-3; dmpRtVarsFlg=1` are entered. Set both symbols to zero before logging out. These commands should not be supplied into a controller other than `master1` as a buffer overflow will be produced by the large volumes of text generated.

transcoding process is not intelligent; statements similar to `for() { oblshapedgradient(); }` generate a large a number of `acode` statements and buffer overruns can occur.

The user interface for the scanner is written in a mixture of Java and XML; it also has its own `perl`-like scripting language, called Magical ("MAGnetic Instrumentation Control and Automation Language"). User-interface parameters are made available to the `C` pulse sequences via command such as `getstr()`. Magical macros are called by the user interface – for example, clicking the '`start scan`' button runs a set of image preparation macros before running the command '`go`'. These perform a variety of tasks, such as calculating inferred parameter values and ensuring that the scan parameters entered by the user are self-consistent. The sequence programming described in this thesis predominantly took place at the `psglib` level, with occasional modifications to psg and hence `acode` generation, along with Magical scripting and user-interface design.

Owing to the fact that the axis of time forms a fundamental part of magnetic resonance experiments, accurate timekeeping is important. Each controller board in the scanner issues a 'keep-alive' signal every 10 ns, which is distributed to every other board including the master controller. These signals are referred to as 'ticks' by the OS, and a report of the number elapsed for a given scan can be obtained by opening a console on `master1`; the number for each board should be constant for a given scan. Beyond checking this timing accuracy, few debugging tools are available to aid sequence development; it is not possible to attach `gdb`, use `valgrind` or other familiar tools for the purpose of debugging pulse sequences on the system. It is possible to directly intercept and decompile `acode`, either using a packet sniffer (e.g. `ettercap`) to intercept communication between the console and the host computer, or by monitoring files in `/vnmr/acqqueue/`. A rudimentary understanding of `acode` instructions can then be obtained by looking at the bytecode, e.g. with the decompiler `radare`. This tedious process is occasionally necessary, as it is unfortunately the only unambiguous way of elucidating exactly what the console is being instructed to do.

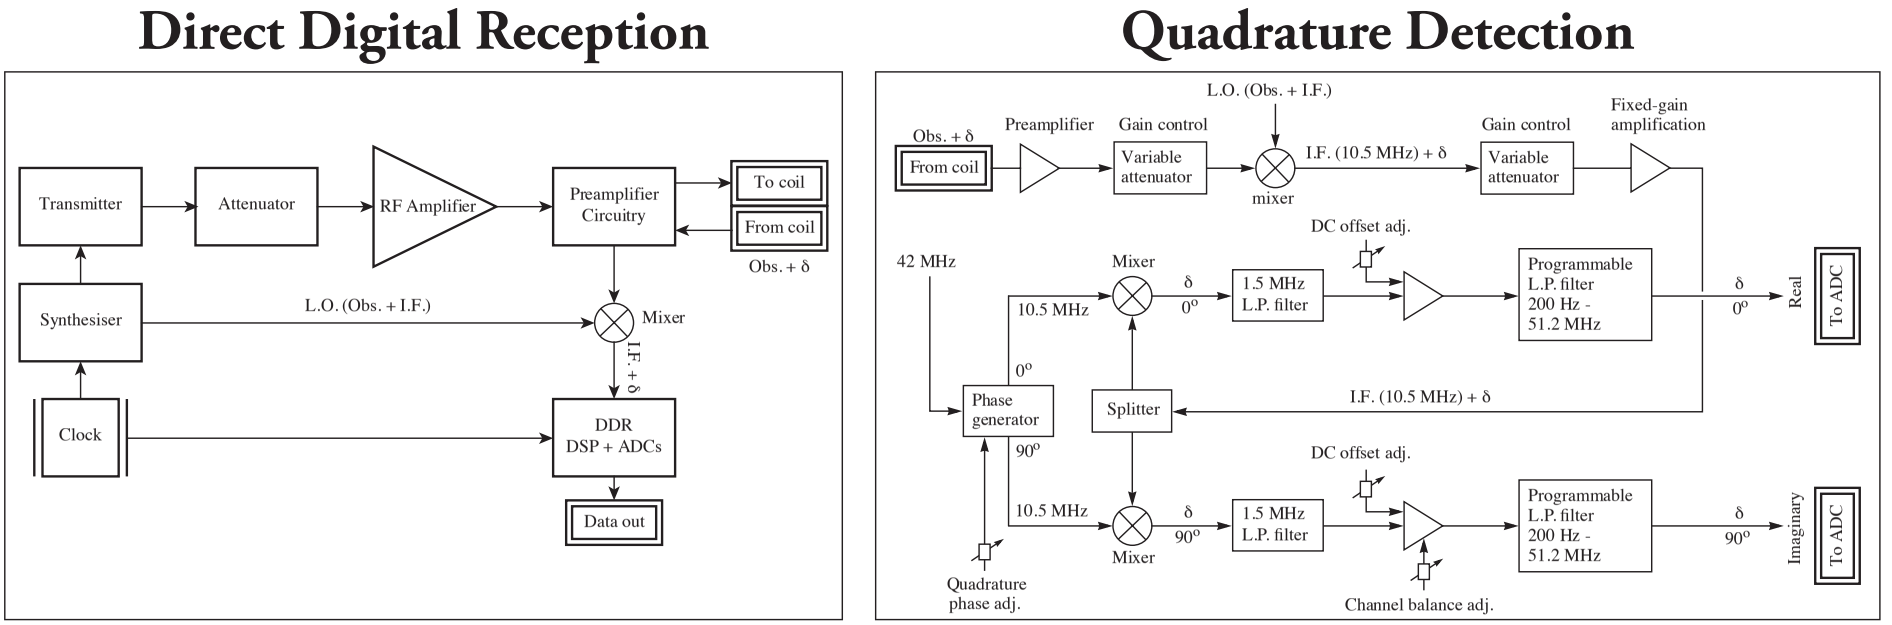

Figure 3.6: The detection of magnetic resonance signals via the DDR system used here (left) or traditional quadrature detection as implemented in the Varian Unity Inova console (right). Both block diagrams are simplified versions of the relevant parts of the receiver chain as detailed in unpublished Varian/Agilent technical documentation, privately communicated by Dr Michael Gyngell, formerly head of MR Applications Development at Agilent. In both cases the spectrometer is set to detect at a centre frequency 'Obs', subject to a small chemical shift modulation δ. Two reference oscillators of interest are created to great precision: the 'intermediate frequency' (I.F.) oscillator which is always either 20 MHz on the DDR console, or 10.5 MHz on the Unity Inova, and the 'local offset' (L.O.) frequency, which is Obs + I.F. The IF is chosen to be a static value for engineering reasons; electronic components have a nonlinear frequency response and by ensuring that the frequencies manipulated are always in a given range unexpected behaviour is avoided. In the digital design one analogue mixer is used after a GaAs FET amplification chain, resulting in a signal oscillating essentially at 20 MHz. This signal is then digitised optimally with a wide spectral and dynamic range, and time-corrected digital filtering techniques are used to remove extraneous noise and ensure the FID starts at its original, nonzero, intensity. An accurate clock together with the reference IF input are used to generate an 'in silico' heterodyne. The traditional route is more complicated. After a fixed-gain preamplifier, the signal is attenuated and mixed to obtain the IF once more. This is then split, reducing SNR, before being combined with two copies of the IF 90° out of phase. This separates out the real and imaginary components of the signal. Analogue low-pass (LP) filters remove noise, and spurious harmonics of the IF signal, with adjustable amplifiers to ensure that the two separate components are equal. The real and imaginary signals are then digitised by separate ADCs. Owing to the large number of potentially different physical devices, and at least four sensitive compensation steps (shown by the variable resistors schematically), an entirely analogue implementation of this chain is imperfect. For example, it is very likely that the filters distort the amplitude of the initial points of the FID.

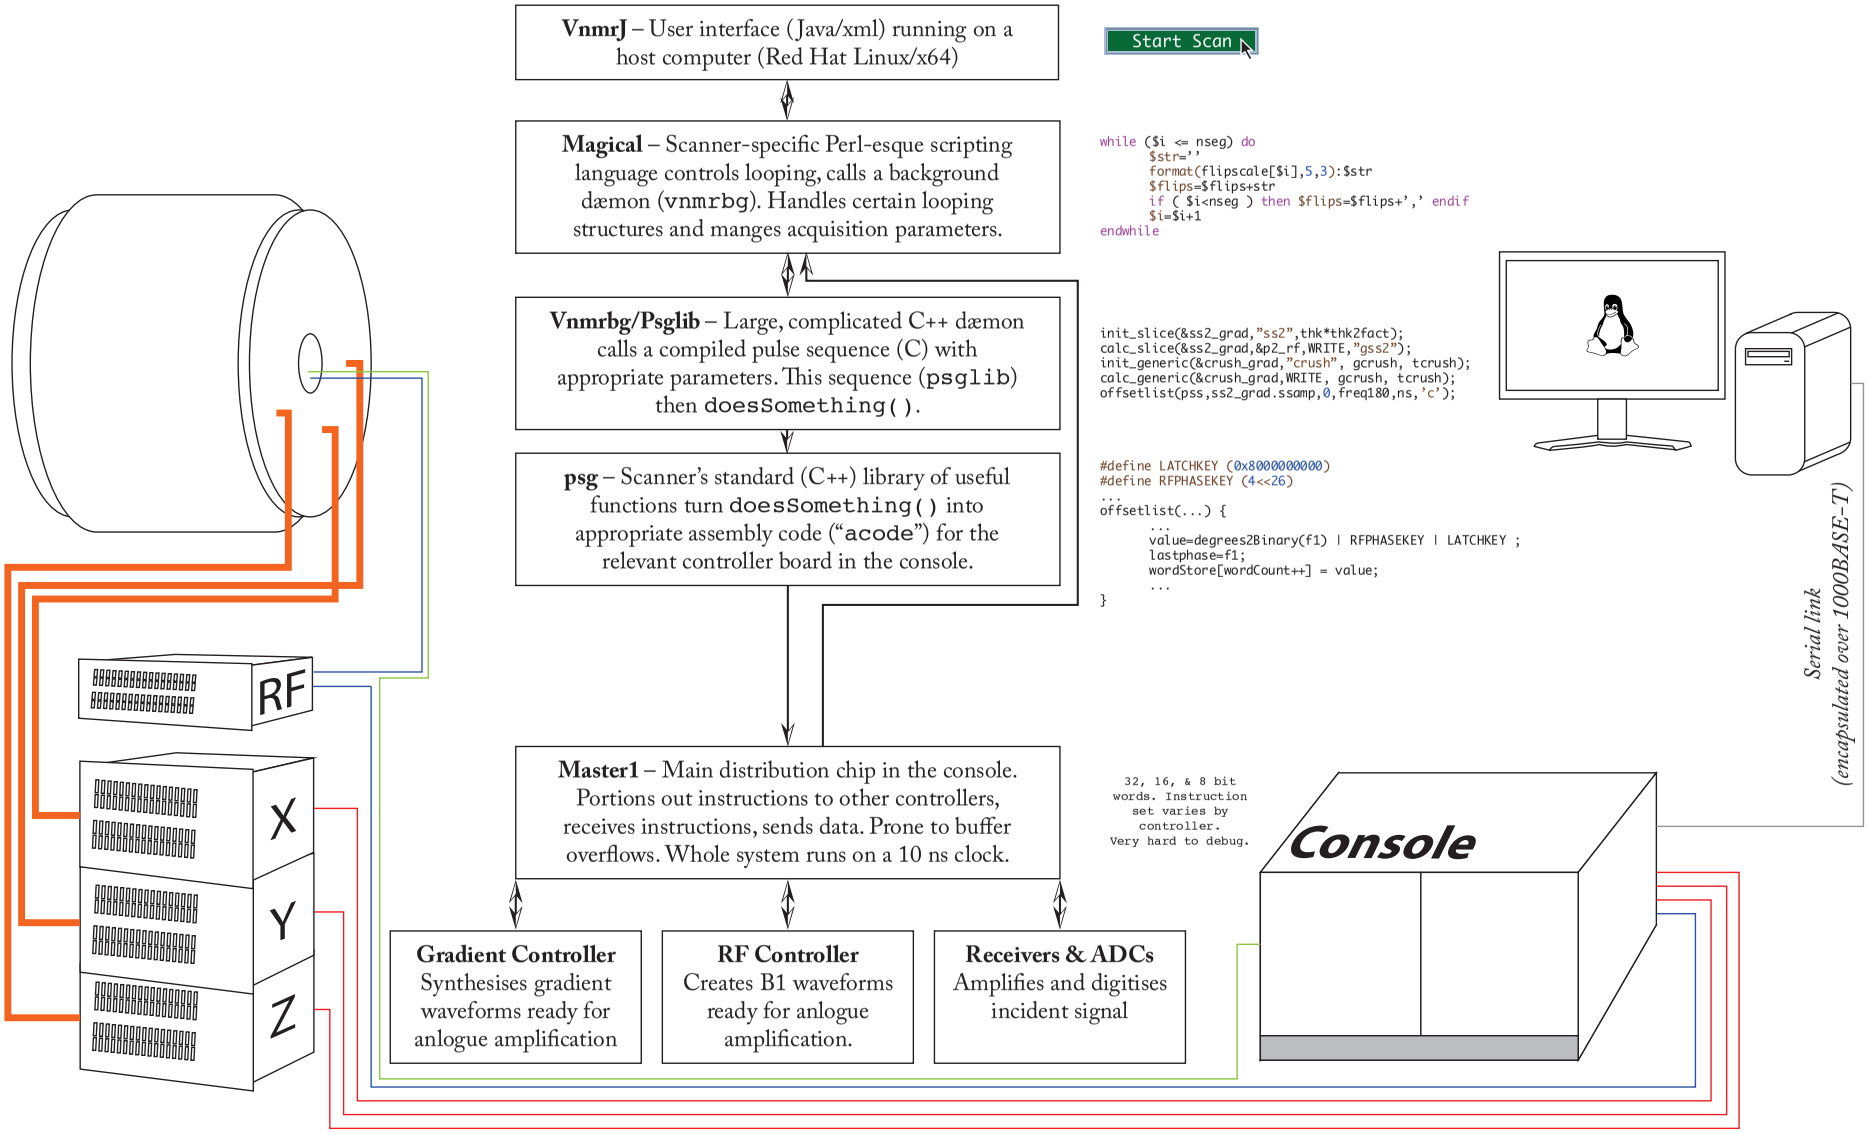

Figure 3.7: The logical organisation of `VnmrJ`. The majority of work in this thesis has taken place by programming sequence files in `C`, with occasional trips to every level of complexity specified.

3.2 | *Radiofrequency coils*

In Section 2.2.1, the effect of a radiofrequency pulse was analysed in terms of B_1, the B field it creates. There is a practical need for devices that create B_1 in order to excite spins (B_1^+) and receive resonance signals (B_1^-). Antennas that create or receive B_1 are called radiofrequency coils, even though in practice they are not solenoids. There are two main types of RF coil commonly used for imaging; volume coils – which generate $B_1^\pm$ homogeneously in a region of space within themselves – and surface coils – which, as the name suggests, sit on the surface of the object under examination.

In principle, an RF coil is a loop of wire with an LC circuit whose resonant frequency is close to the Larmor frequency of the nuclei of interest. The value L is nominally determined solely by the coil's geometry, and a 'tuning' variable capacitance C is added in series. However, as the coil is connected to an RF power amplifier whose output impedance will usually be $50\,\Omega$, it is necessary to ensure that the coil's complex impedance is also $50\,\Omega$ to maximise power transfer and avoid reflecting incident waves back into the amplifier (which may result in damage). This situation is avoided by the addition of a second capacitor in series with the coil, which allows for the matching of its impedance with that of the amplifier. This simple circuit for a surface coil is shown in Fig. 3.8.

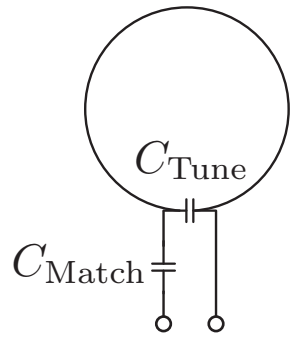

Figure 3.8: A surface coil, shown with tuning and matching capacitors.

The homogeneous region that a given RF coil generates can be determined by solving Maxwell's equations with appropriate boundary conditions reflecting the time-varying E field applied to drive the coil and also the geometry of the coil and its surroundings. For a surface coil, B_1 spatially drops off rapidly, approximately[7] following $1/r^3$. Surface coils therefore localise signal to a specific region, and are commonly used in *in vivo* NMR spectroscopy to ensure that signal arises only from a target organ of interest whilst eliminating thermal noise from the rest of the sample. This increases SNR, but at the cost of homogeneity: images acquired with surface coils display a spatially variable amplitude modulation.

The rapid spatial variation of B_1 generated by a surface coil is often undesirable. Accordingly, alternative 'volume coils' are usually used for routine ^{1}H imaging. These are constructed specifically to have a very homogeneous region of B_1 in a designed location, and have to be large enough to fit the object of study within their frame. The most common geometry of volume coil used is undoubtedly the 'birdcage' design (illustrated in Fig. 3.9), which is exceptionally homogeneous and analytically

[7] Analytically, Biot-Savart's law tells us that $(B_1)_z = \frac{\mu_0}{4\pi}\frac{2\pi R^2 I}{(z^2+R^2)^{3/2}}$ at a distance z away from the axis of a loop of radius R being driven with a current I. However, Biot-Savart's law is derived under the assumption that $\frac{\mathrm{d}\mathbf{B}}{\mathrm{d}t} = 0$. It becomes progressively invalid as the Larmor frequency (i.e. B_0) increases, yielding erroneous results by the time the wavelength and the coil geometry are comparable.

understood.[36] As the birdcage coil essentially acts as a tuned transmission line, there is a cosine variation of current around the outside of the coil, usually designed such that there is one complete standing wave around its circumference. Due to this cosine term, nodes with zero voltage occur $90°$ away from the driven part of the coil. It is then possible to introduce a second signal $90°$ out of phase ("in quadrature"), producing a circularly polarized B_1 field that couples to M_x and M_y. The use of quadrature coils improves SNR by a factor of $\sqrt{2}$.

The introduction of PIN diodes to rapidly tune or de-tune coils at run time allows for the third, obvious combination of both techniques: a volume coil can be used to transmit (where homogeneity is important) and a surface coil to receive (where SNR can be paramount). It is necessary to rapidly detune a receive-only coil when transmitting on another nearby coil to avoid unwanted electromagnetic interactions, i.e. broadcasting the RF from one coil to the other, rather than into the sample.

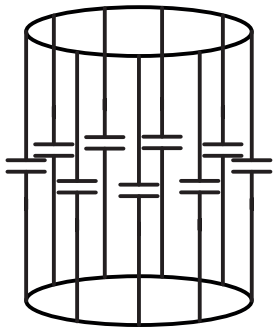

Figure 3.9: A cartoon of a birdcage volume coil.

Multiple independent RF coils can be used to parallelise the process of both excitation and reception. In both cases, the geometric information that can be inferred from the differing spatial distribution of coils can be used to increase SNR or equivalently decrease scan time, and (usually) improve homogeneity. Parallel transmission (pTx) essentially generates B_1^+ as the vector sum of the B fields of the many elements of different coils and is a relatively new development in MRI. It offers many potential avenues to overcome limitations that are increasingly relevant on high-field clinical scanners, for example, the increased RF heating of the patient at high-field and the decreased RF homogeneity of traditional equipment. Unfortunately, pTx is not available on the Varian system used in this thesis. Parallel reception (pRx) is much more established, with several well-known algorithms (such as SENSE, SMASH and GRAPPA[37, in review]) allowing for accelerated imaging by under-sampling data and using the known spatial localisation of the coils to reconstruct it. The console used in this thesis has the ability to receive from up to four channels simultaneously.

3.2.1 Coil characterisation

Like most MR facilities, we have a variety of RF coils available for differently sized samples and nuclei. Ideally, an RF coil would have zero electrical resistance, not have any intrinsic source of noise, and have a homogeneous transmit and receive B_1 field over a precisely defined region of interest. In practice, none of these conditions can ever be met exactly. Accordingly, different coils are often constructed in which differing compromises are made, and these compromises can have profound impacts on the homogeneity and sensitivity of the B_1 field generated. For the imaging of nuclei other than ^{1}H, maximising the overall received SNR is usually the most prevalent concern. Accordingly,

the transmit/receive sensitivity of each RF coil we possess was measured by acquiring gradient echo images that were axial projections along a homogeneous phantom at multiple flip angles with a TR of at least $5\mathrm{T}_1$. Field maps were then generated by the voxelwise fitting of data to the simple model $M = M_0(x, y, z)\sin\big(\alpha(x, y, z)\theta\big)$, where θ is the delivered flip angle in space, and M_0 effectively a measure of the coil's spatial response. These expressions can be related to B_1^+ and B_1^-, but from a practical point of view the ideal coil has M_0 spatially uniform (and high), and α spatially uniform (and unity). Owing to the fact that this measurement requires a large amount of scan time (i.e. 60 or more hours) on thermal equilibrium samples, it was only performed on coils of particular interest.

To facilitate the use of coils other than those that connect directly to the extant animal handling system, it was necessary to construct both mechanical equipment holders, and, in the case of the actively detuned surface/receive coil, additional RF shielding.

3.2.2 *Rat Butterfly surface coil*

This coil is a home-built double-loop surface coil constructed from 1 mm thick adhesive copper deposited on a glass-reinforced plastic (GRP) substrate. The substrate has a diameter of 65 mm, and the coil loops are 4 mm long. The coil is integrated into the animal handling system ("the cradle") and is plugged into a separate tuning and matching circuit with passive LC 0-300 MHz bandpass filters.

The coil is not driven in quadrature, nor is it actively detunable. Its sensitivity drops off sharply spatially, as expected for a surface transmit/receive device. It is, however, easy to use with the equipment originally present, and has a relatively good SNR profile: a 16 average spectrum acquired on a reference phantom (a 1.5 ml bottle of 5 M ^{13}C-urea) displays a single peak with an SNR of 175. (This sequence is used as a reference for coil comparisons; it is a hard 90° excitation with a pulse width of ~ 100 μs, TR $= 0.5$ s, 8192 Hz spectral width, 16 avgs, receiver gain of 26 dB.)

However, as shown in Fig. 3.10**A**, the coil is not spatially homogeneous, and the received flip angle varies across the sample: the mean α is 0.69 with a standard deviation σ_α of 0.2, i.e., the delivered flip angle varies by $\sim 20\%$ across the sensitive region of the coil.

3.2.3 *Mouse Loop surface coil*

This coil is a small (∼2 cm diameter), simple loop surface coil that is usually used for extremely localised spectroscopy. After its use in an initial preliminary spectroscopic investigation (described in Section 4.3.3), it was determined that the spatial penetration depth of the coil – approximately 1 cm – was insufficient to warrant its use or further investigation.

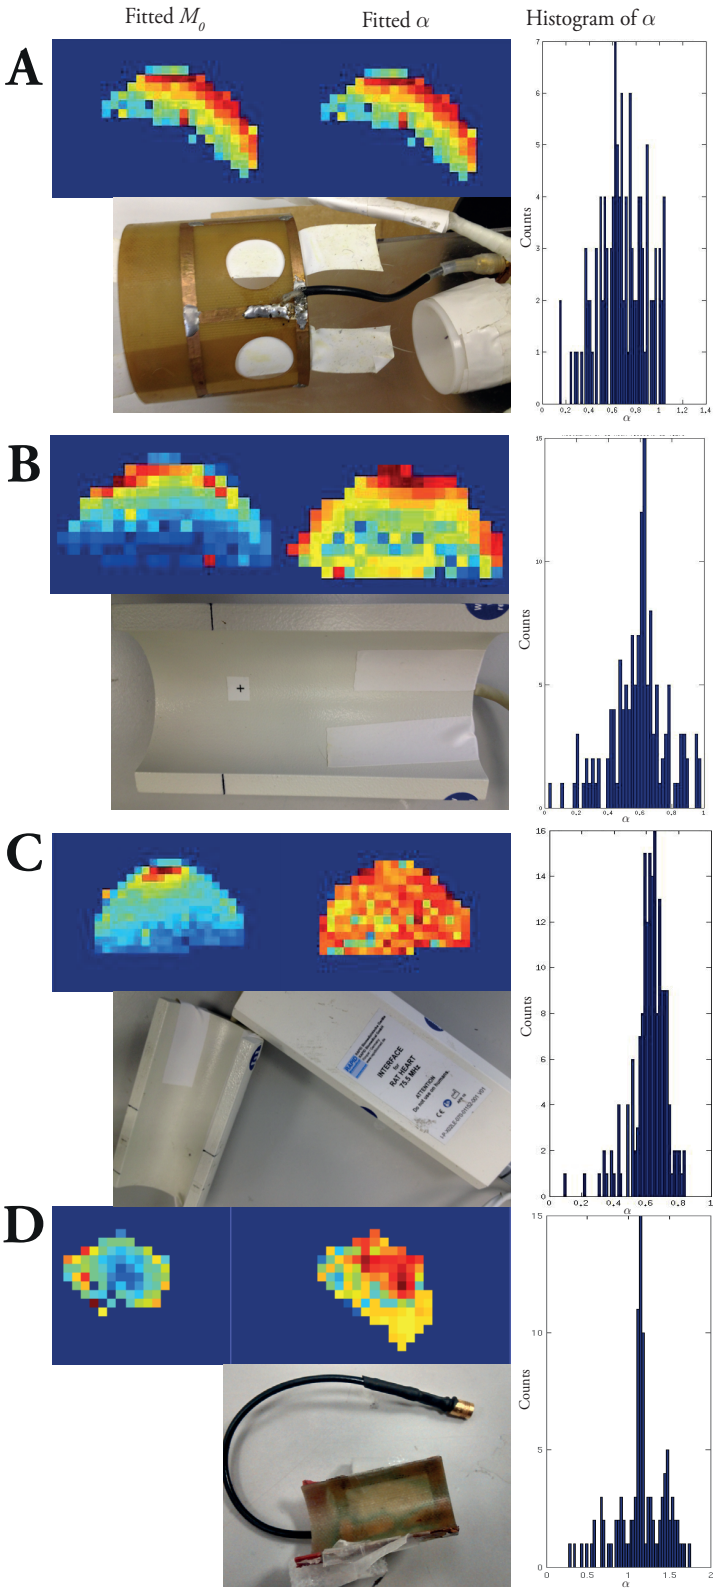

Figure 3.10: Shown here are the measured sensitivity profiles ($\approx M_0$) and flip-angle homogeneity ($\approx \alpha$) of four of the ^{13}C RF coils that we possess. A histogram of α is also shown, highlighting the variability in the flip-angle for the different instruments. A photograph of each coil is also shown for illustrative purposes. Fully relaxed gradient echo projections were acquired on homogeneous phantoms along the long axis of each coil were acquired at different flip angles, and fit to the simple model $M = M_0 \cos \alpha\theta$. From top to bottom, the coils are (**A**) the home-built rat butterfly surface coil; (**B**) the quadrature surface transmit/receive coil; (**C**) the surface coil elements of the Rapid Biomedical two-channel surface receive/volume transmit setup, shown with the original pre-amp; and (**D**) the home-built brain surface coil.

3.2.4 *Quadrature surface coil*

Rapid Biomedical GmbH (Rimpar, Germany) constructed this quadrature surface transmit/receive coil for our system in c. 2009. It contains decoupling circuitry in a nearby small Faraday caged box, and two circular loops driven in quadrature. As shown in Fig. 3.10**B**, the delivered flip angle is again spatially non-uniform ($\sigma_\alpha = 0.18$) with hotspots near the coil elements. The receive sensitivity also quickly drops off as a function of distance from the coil's RF centre. The large pre-amplifier posed significant practical problems when locating the sample and coil within the 72 mm internal diameter of the quadrature proton volume coil used for localiser imaging. As the coil did not generate a large SNR benefit over the alternatives, and presented several practical difficulties, it was not used.

3.2.5 *Two-channel surface-receive coil & dual-tuned Birdcage volume coil*

This coil, manufactured by Rapid Biomedical, consists of an actively tuned dual-tuned carbon/proton birdcage coil, and an actively detuned two-element surface receive coil. Each element is a curved 20×40 mm rectangle on a circular glass-reinforced plastic (GRP) substrate with a 67 mm outer diameter. An integrated pair of independent pre-amplifiers on a single chip is provided close to the receive elements. Active tuning/detuning is performed by PIN diodes, which (dis)connect other resonant elements to the circuit, ensuring that a) the sensitive pre-amplifiers in the receiver chain are not destroyed by the RF transmitted (at great power) by the volume coil, and b) the surface and volume coils do not interact. These PIN diodes are driven by a constant current low-noise DC power-supply with fast switching driven by TTL logic provided by the console when RF transmission and reception occurs. Correspondingly, the volume transmit/surface receive arrangement can provide the 'best of both worlds' approach to MR coil design, with the homogeneity in transmission of a volume birdcage coil, and the sensitivity and increased SNR profile of a surface coil. This is reflected in the calibration performed (Fig. 3.10**C**), with exceptionally homogeneous flip angle delivery ($\sigma_\alpha = 0.09$) and a reasonably homogeneous receive profile. Moreover, the two independent channels can be used for parallel imaging strategies, or to improve SNR.

However, this coil originally had a large ($\sim 200 \times 50 \times 20$ mm) external interface housing the pre-amplifiers and PIN decoupling circuitry. This interface was difficult to use, barely fitting within the volume coil, and requiring seven electrical connections to be made within the physically constrained space of the magnet bore (two receive channels PIN diode DC lines for the surface coil, one transmit/receive (^{1}H), one transmit only (^{13}C) and two PIN diode DC lines for the volume coil). Moreover, the response of each channel was not initially perceived to be uniform. With a urea phantom placed at the RF centre of this coil, and the coil in the isocentre of the magnet, varying results were obtained from spectroscopic studies on the phantom. As illustrated in Fig. 3.11, quite often the coil appeared to be broken, with one channel displaying far worse performance than the other; the example shown has an SNR of 106 on the

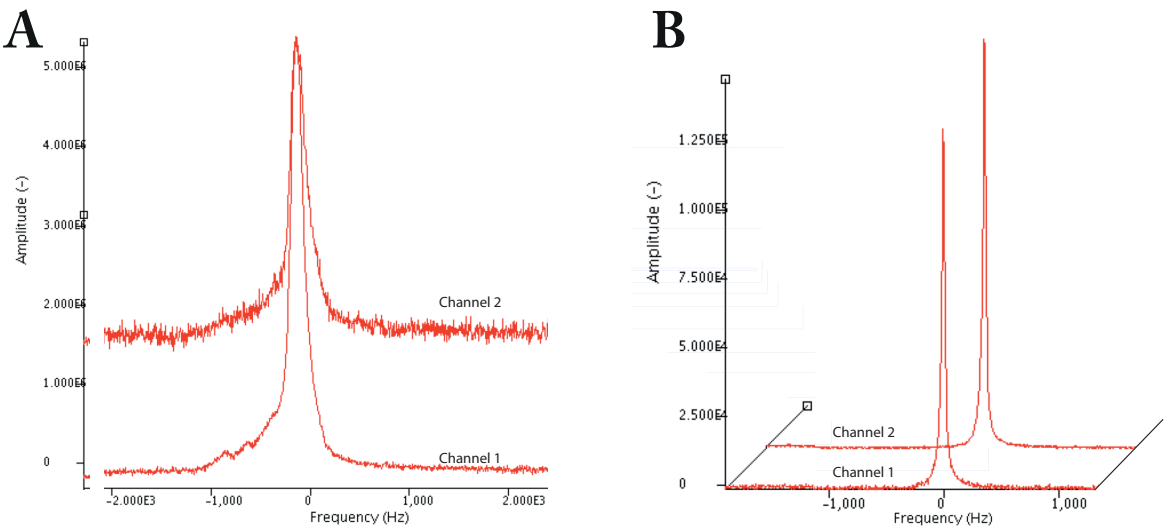

Figure 3.11: **A**: Illustrative NMR spectrum acquired from a 1.5 ml bottle of 5 M ^{13}Curea with a 90° pulse and sixteen averages. The phantom is located equidistant between the two coil elements at its designed RF centre, yet it is clear that the spectra obtained from the elements are not equal. In this instance, the SNR on channel one is an order of magnitude better than that from channel two. Which channel "was broken" was found to change between experiments, and the effect was not indicative of broken pre-amplifier circuitry; using different receivers did not ameliorate the effect (and resulted in identical spectra being obtained). Further investigation led to the hypothesis that the effect was purely down to standing wave interactions occurring in the magnet bore; spectra acquired with the hardware subsequently constructed are shown in **B**, which possess far better SNR overall, and essentially equal SNR between the channels. The decrease in linewidth is due to the fact that ameliorating these EM problems allowed for greater ease in shimming, and also because this second acquisition occurred after optimising the superconducting shims on the magnet following a quench. The arbitrary scale is not common between the two acquisitions.

first channel and 12 on the second. For the sake of illustration, the SNR when acquiring (with the same 16 average spectroscopy experiment) on the carbon volume coil alone is 10. However, repeating this experiment – and removing the coil hardware from the scanner bore – showed different results, with the SNR of the two channels varying wildly for no easily discernible cause. Under optimum conditions, which were difficult to determine, the SNR of the two channels were equal, and approximately $150 - 200$ each (resulting in a significant gain in signal over the surface coil).

With the help of Tobias Wichmann (Rapid Biomedical GMBH), the cause of the variation in the generated B_1 field was revealed to be the DC PIN diode decoupling lines introducing static B fields that interacted with the static field and the generated B_1. In particular, it was found that physically separating the diode lines from those carrying a signal greatly improved the stability of the system, as did ensuring that the outer conductor of every diode line was properly earthed. A full RF shield was then constructed by the author, out of aluminium and Delrin plastic machined by hand[8] to isolate the cabling of the volume transmit and surface receive paths at all points in the signal path. This shield also ensured that the electrical potential of the diode decoupling lines, the RF shield of each coil, the

[8]Components were cut on Colchester lathes and Bridgeport milling machines to 10 μm tolerances with safe working practices under the kind supervision of Matthew Newport in the Department of Physics research workshops.

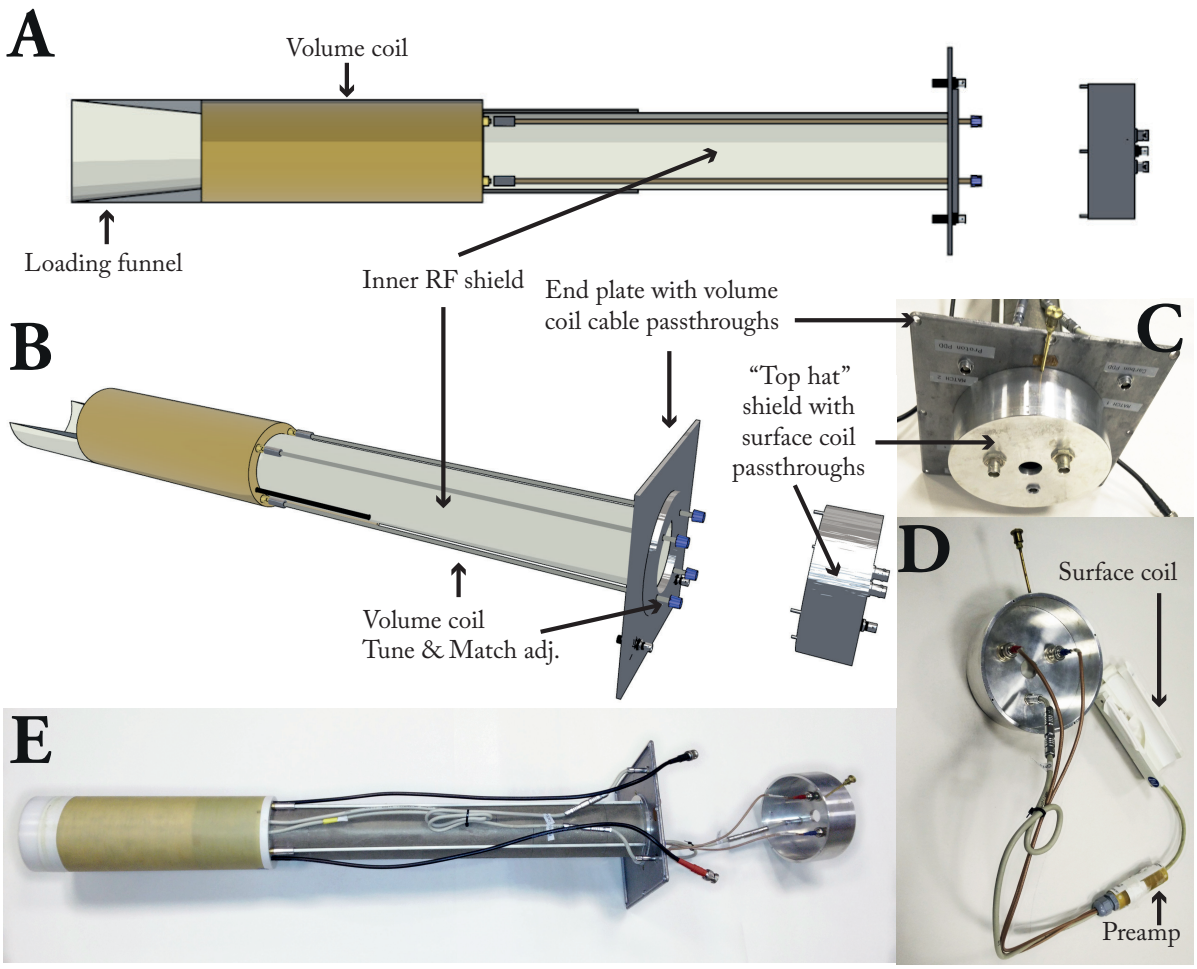

Figure 3.12: The extensions, RF shields and housing required to permit the use of the (redesigned) two-channel surface receive/volume transmit coils, as shown in top plan (**A**), and (**B**) in perspective. The sample enters in a separate cradle, which has the same internal diameter as the volume coil (72 mm). To facilitate fitting the two into each other within the bore of the magnet, an internal cone (cut from Delrin plastic) is provided. The volume coil itself has an outer diameter of 99 mm, which is close to the internal bore of the gradient set (100 mm), providing a close fit. It has four tuning and matching interfaces that need to be accessible (two each for ^{1}H and ^{13}C), so strong GRP rods are connected to the variable capacitors directly, and allow them to be manipulated from outside the magnet. These rods also mechanically stabilise the assembly. An internal RF shield (consisting of a GRP tube coated with a copper wool Faraday cage 'EMI tape') then electromagnetically isolates the pin-diode decoupling (PDD) lines and transmission lines to the volume coil from the inner bore of the volume coil. The RF shield is electrically connected to the earth of the volume coil, and connected to a mounting plate which mates with a shielding box with BNC and PDD bulkhead passthroughs. A 'top hat' further shield is then provided to isolate the *inner* bore of the coil, which fits neatly onto the end-plate. When the sample is in place, the surface coil BNC and PDD lines mate with bulkhead passthroughs on the top hat, ensuring that their earth is common. A small hole in the top hat is sufficient to pass other lines through to the sample if required. The implementation of the 'top hat' is shown in photographs **C** (where it is mated to the end-plate) and **D** (with the surface coil and its integrated preamp shown); the manufacturing tolerances are such that the whole system (**E**) matches the design (**A** & **B**) to an accuracy of at least 100 μm on plastic parts.

gradient set, and the cryostat were all at the same electrical potential, namely earth. The design of the surface receive coil was also changed to include a shorter distance from the coil elements to the pre-amplifier (improving SNR), replacing the exact pre-amplifier used with a physically far smaller GaAs FET single-chip package element, and reducing the length and path of the DC lines to minimise the magnitude of unwanted (i.e. changing transverse B fields) generated close to the receiving elements. The loading characteristics of the coil were also changed to have a very wide bandwidth, removing the need to adjust tuning and matching capacitors for each sample. Bulky in-line magnet-compatible RF chokes were also removed as the lines' outer conductors were periodically attached to a permanent earth. These changes both made the system far easier to use, with all cabling either permanently attached to the magnet, or easily slid through an isolated path in the bore before being connected to a 'top-hat' shaped RF shield.

Rendered images from the CAD system used to design the coil's housing and RF shields are shown in Fig. 3.12**A-B**, with photographs of the resulting object in **C-E**. Under the same scan conditions, the SNR is around 175 *per channel*, independent of the orientation of the equipment. This represents a significant improvement from Faraday caging the receive path, and moving the pre-amplifier closer to the elements.

3.2.6 Rat Brain Butterfly surface coil

Whilst problems with the two-channel coil were being solved, a form-fitting surface receive/transmit coil was constructed by the author for use imaging rat brains. It consists of a two-loop square birdcage design, wrapped around a GRP half-cylinder and is illustrated and characterised in Fig. 3.10**D**. The circuit design originated from a resonant LC bandpass circuit produced by a fourth-order Butterworth filter driven by a tuning and matching circuit and shown in Fig. 3.13**A**. AC analysis was performed in the open-source software QUCS[9] to obtain the scattering matrix for the system, commonly referred to as the S matrix. The element S_{11} gives the reflected power incident on the coil, which ideally is minimised at the resonant frequency. Numerical minimisation of S guided the chosen design values of C given a fixed L dependent on the geometry of the coil. Finite Element Analysis (FEA) was then performed with the commercial solver Comsol (COMSOL AB, Stockholm, Sweden) for a given value of C in an attempt to predict B_1. Unfortunately, as illustrated in Fig. 3.13**B**, the physical geometry of the surface-mount ceramic capacitors required is relatively large in comparison with the circuit, and numerically solving Maxwell's equations proved troublesome. A very large E field was predicted to be present over the primary capacitor, whose exact composition was unknown, potentially reflecting an error in the simulation. Accordingly, more detailed FEA was not performed and the device was prototyped and

[9] *Quite Universal Circuit Simulator*; `http://qucs.sourceforge.net`.

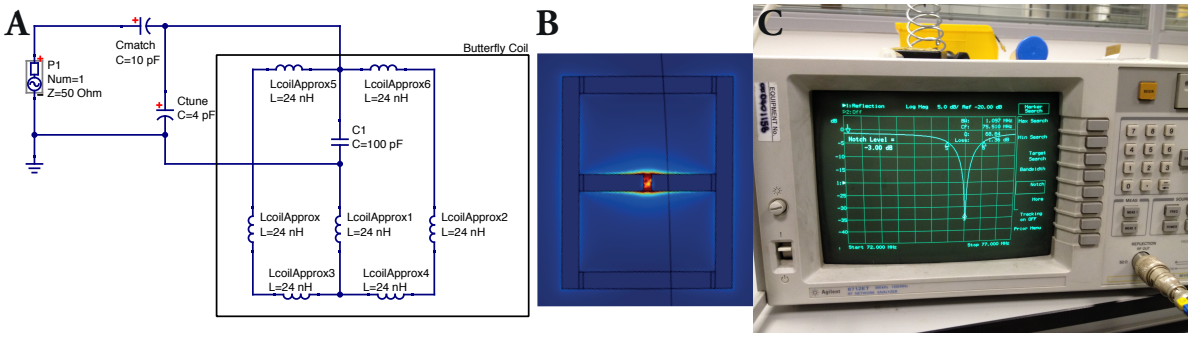

Figure 3.13: The design process of the rat brain coil. (**A**) Circuit-level simulation with QUCS lead to a resonant LC circuit with values of L and C that were physically realisable; (**B**) an attempt at finite element modelling was made to solve Maxwell's equations over the mesh of the proposed coil. However, it was found that the unknown behaviour of the ATC ceramic chip capacitor used was a significant source of error in the model, shown by the extremely large electric field across the location of the capacitor above, and it was not pursued further. Accordingly, the device was prototyped and found to have a frequency response that matched that of the design, shown on a network analyser in (**C**), and the measured values of L and C were close to those simulated.

found to match the Butterworth design.[10] The coil is highly resonant, with the dimensionless width Q (centre frequency / bandwidth) being 68, with a narrow bandwidth (~ 1 MHz) when tuned to the carbon frequency at 7 T (75.515 MHz). As shown in Fig. 3.10**D**, this coil has a spatially varying flip angle ($\sigma_\alpha = 0.3$) and sensitivity profile. For these reasons, and due to the fact the coil does not offer a significant SNR advantage, it was not developed further.

3.3 | *Subject handling*

3.3.1 *Position, monitoring & anaesthesia*

Animals scanned were anaesthetised with isoflurane to minimise their distress at being immobilised for long periods of time. Anaesthesia was induced with ~ 3.5 % isoflurane in oxygen (2 L/min) with a small quantity of nitrous oxide to provide analgesia (0.2 L/min N_2O). Anaesthesia was maintained at 1.5 % isoflurane in oxygen and nitrous oxide. Animals were placed in a cradle with a nose cone for the delivery of anaesthesia, heating air circulation, scavenging of anaesthetic gasses, and passive-filtered electrical connections for electrocardiogram (ECG) leads, which were placed subcutaneously near the triceps muscle in the rat. Home-built surface coils also interface with the cradle, and passive LC bandpass filters and tuning and matching circuits are provided for this purpose. Hyperpolarised compounds were injected into the rat via a cannula in one of the bilateral tail veins. The cradle also contains a calibrated scale that allows for the positioning of the animal at the magnet's isocentre.

All animals were housed in appropriate facilities in the University of Oxford, on a 12:12 hour light/dark cycle. Studies using animals conformed to the Home Office Guidelines of the Operation of

[10] Although the author wishes to note that some copper substrates have conductive adhesive, and other, seemingly identical and unlabelled substrates, do not.

the Animals (Scientific Procedures) Act of 1986, to the University of Oxford Institutional Guidelines, and were subject to local ethical review. All metabolic data was acquired in the fed state (i.e. between 6 am and 1 pm) unless otherwise stated.

3.3.2 Cardiac gating

It is difficult to image the beating heart. One common way of dealing with this problem is to gate the acquisition period of an imaging sequence to the heartbeat, usually the period of diastole where the heart is stillest for longest. This is usually performed by monitoring the ECG of the subject. We record an ECG from the anaesthetised rat under study by subcutaneously inserting thin needles under the skin near the triceps muscles. These needles function as electrodes, and the ECG is recorded and differentially amplified by a home-built box, constructed by Paul Cassidy in c. 2004. The box entirely consists of CMOS logic ICs with a nine-stage amplifier chain, and its purpose is to filter and differentially amplify the incident ECG signal before performing threshold-based triggering to work out when the R wave of an ECG has come in. Its output is a 4 μs long 5 V TTL pulse to the `xgate1` input of the console. This TTL represents the start of an acceptable period of imaging. The box can also amplify a respiratory signal, provided by a loop of wire that moves when the animal breathes. A large period of motion then follows, and the box keeps a second TTL line – `xgate2` – high for the estimated duration of respiration. Pulse sequences can therefore be cardiac and respiratory gated.

However, there is variability in the heart-rate of the anaesthetised rat. If this variability is not taken into account, an error will be produced in the time axis of the experiment. After searching for computational means to record the absolute time at which a trigger signal is received by the console, and being told by employees of its manufacturer that no such means existed, the decision was taken to create a hardware device to record the time at which each triggered shot of a scan occurred. Whilst it might be possible to alter the software running on the console, in either code running on VxWorks or on the fabric of the FPGAs in `master1`, the lack of accessible source code made this a daunting proposition with potentially expensive consequences should anything go wrong. Accordingly, the ECG triggering box was modified to produce 4 μs long 5 V TTL pulses at the start of both a respiratory and cardiac gating period. As illustrated in Fig. 3.14, a very simple and small board was constructed to act as an interface to the acquisition computer to these signals, with two 4013N SR latches used to hold a pin on a parallel port interface high until the acquisition computer polls it. Two Schmitt triggers are used to provide a delayed input, which when fed into the 4013N provide a measure of missed signals should the acquisition computer not poll the parallel port frequently enough. In practice, this problem never occurred as the polling software was written using the `parapin` library, originally developed for fast applications by the United States Department of Defence. A kernel mode driver was used to provide exceptionally low-latency input to the parallel port, whose timing accuracy is limited by the clock speed

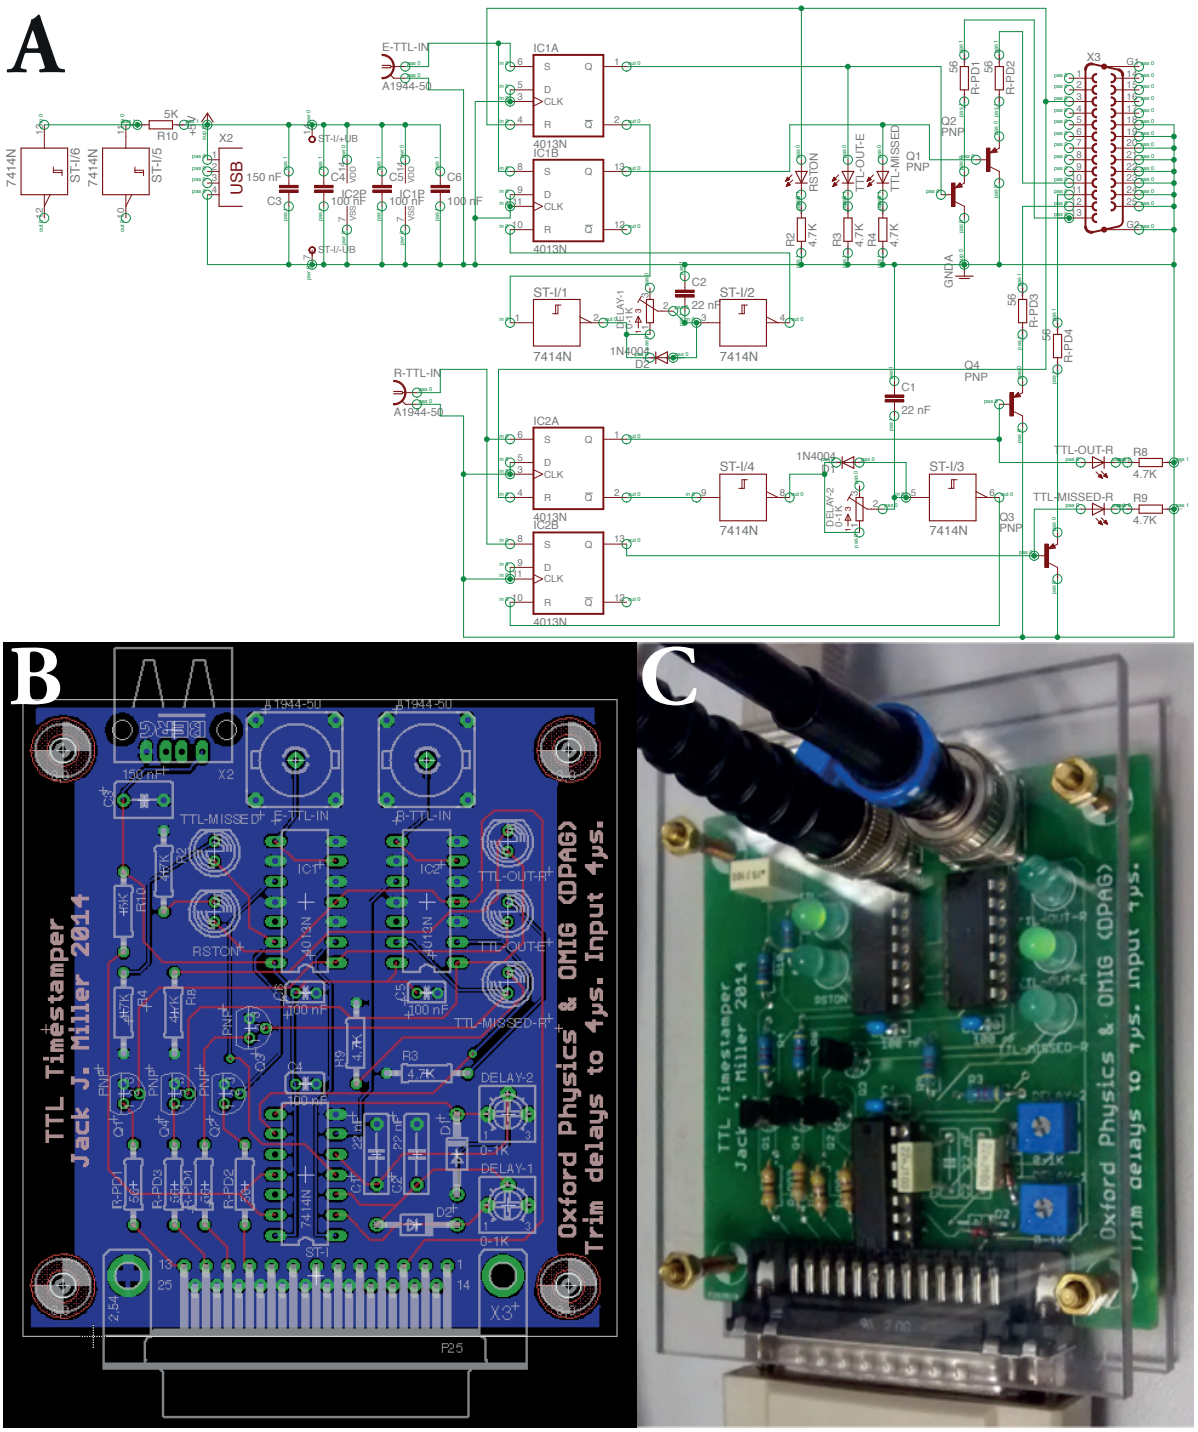

Figure 3.14: Circuit diagram (**A**), board layout (**B**) and manufactured PCB (**C**) of the 'timestamper' board, whose inputs are TTL signals from the cardio-respiratory gating box, and whose outputs are pins held high or low on a parallel port interface. Power is crudely provided by USB. LEDs inform the user if the board is being communicated with, and (high tolerance, theramlly stable, ceramic and accurate) trimming potentiometers are provided to compensate for RC delays. A crude perspex "case" was additionally constructed by the author by hand. The board functioned as designed immediately after manufacture.

of the processor and the load on the system. The process written used a specific Intel processor timing counter, `CLOCK_MONOTONIC_RAW`, to time events. It is worth noting that there is a long history of using parallel ports in this fashion as "easy" ADCs, as the driver bypasses a lot of the electrical and logical circuitry associated with using the parallel port as a printer interface. When the scan starts, a polling process is spawned as root, and is rescheduled to have the highest real-time priority on the kernel, and a high scheduling priority.[11] When a scan occurs, a signal comes in, and the board holds a pin high on the parallel interface indicating this. The re-prioritised process spawned on the acquisition computer records this, makes an accurate timestamp in a file, and resets the latch. The timing accuracy is not well defined, but in practice with a quad-core hyperthreaded Xeon processor, it is substantially less than the duration of the latch periods. Owing to the lack of a logic analyser fast enough to measure the response time, calibration experiments were performed by assuming that the console kept proper time when running ungated experiments. Correction factors were introduced according to the published switching time of the ICs used, and accurately measured LC constants at the relevant frequency in the cabling used. These showed that the constructed timestamper was accurate, and the total duration of communication was approximately 10 μs. As the scan starts on the rising edge of the input TTL, this was *post hoc* decided to be the time recorded. Subsequent tests of temporal accuracy revealed excellent results.[12] Circuit simulations were carried out in TINA Cloud, with schematic and board layout in the freeware version of EAGLE. The PCB was manually laid out and routed by the author, with much useful design advice gratefully received from Mark Jones in the Physics Electronics Workshop, before being commercially manufactured by EuroCircuits Belgium and assembled by hand in Oxford.

It might be argued that the alternative approach of putting an ADC together with a fast clock on a board might offer improved – and guaranteed – timing accuracy, independent of the activity of the host computer's CPU. However, the fast sampling rate required to accurately record the short duration of TTL pulses signalling the start of a scan is prohibitive. Alternative digital I/O devices, such as the `82C55` IC, exist, but would require significant supporting implementation circuitry, such as a microprocessor referencing inputs to a faster external clock and/or a fast FPGA. This is in contrast to the SR latches used, whose comparators have a switching time that is less than 1 ns (e.g. 200 ps at 22 °C). Moreover, encapsulating the resulting timestamp communication over USB would only add overhead and introduce the difficulty of synchronising the clock running on the board to the time of the experiment. Additionally, the exceptionally simple circuit as implemented has a bill of materials that

[11] This ensures that the process is run in nearly real-time mode, and that the computer deals with received signals promptly when they arrive.

[12] More precisely, when running ~ 2200 ungated scans with a set TR of 150 ms, the median TR as detected by the new hardware is $0.150\,000\,000\,000\,006$ s. One scan was detected at twice the TR. As mentioned above, due to the lack of an available fast logic analyser, the actual accuracy has yet to be measured conclusively; it is certainly worse than this figure may imply, as, by necessity, the returned timing accuracy is truncated to 10 ns. However, given the RR interval of the rat is ~ 150 ms, it was decided that microsecond accuracy was acceptable, and accordingly the precision of this device was not fully explored. The experiment was repeated with 2220 ungated scans with a different TR , nominally of 100 ms, a median TR of $0.100\,000\,000\,000\,009$ s was recorded.

costs ∼£10, in addition to the cost of the board (£100 for six): Xilinx FPGAs are far more expensive. In short, the implemented timestamper might be crude, application-specific, and highly non-portable, but it is effective.

3.4 | *Shimming*

It is a central requirement for effective MR experiments that B_0 is made as homogeneous as possible. The process of making B_0 more homogeneous than it would be otherwise is known as 'shimming', referring to its origins as the practice of securing small pieces of ferrous material in the B_0 coil of the main magnet. On modern systems, shimming is a three-layered process consisting of the adjustment of small physical pieces of metal and (large) superconducting adjustment currents on coils within the cryostat (both of which are adjusted when the magnet is energised and subsequently not changed), and the creation and routine adjustment of smaller correction fields by room-temperature coils present within the gradient set of the system at hand. Owing to the electrostatic interaction between sample and B_0, and the varying electric and magnetic permeabilities of different samples, shimming is a process that has to be performed routinely before each scan. For rapid imaging, B_0 inhomogeneity causes a variety of unwanted effects, discussed below.

In the absence of currents or charges, Laplace's equation for the scalar magnetic potential ξ is solved inside the magnet bore, i.e.

$$\nabla^2 \xi = 0 \tag{3.3}$$

from which the physical fields can be reconstructed as $\mathbf{H} = -\nabla\xi$ and $\mathbf{B} = \mu\mathbf{H}$. Note that the presence of materials with two different relative permeabilities will necessitate a discontinuity in $\mathbf{B}$ such that Maxwell's equations are satisfied, and no currents flow. In spherical polars, Eq. (3.3) can be solved by looking for separable solutions of the form $\xi(r,\, \theta,\, \phi) = R(r)\Theta(\theta)\Phi(\phi)$, and explicit manipulation of the Laplacian in spherical polars together with the imposition of equality constraints. This yields the general result

$$\xi(r,\, \theta,\, \phi) = \sum_{l=0}^{\infty} \sum_{m=-1}^{l} \left(A_l r^l + B_l r^{-l-1} \right) P_l^m(\cos\phi) e^{-im\theta}, \quad \{m,\, l\} \in \mathbb{Z} \tag{3.4}$$

where $P_l^m(\cos\phi)$ are the Legendre polynomials and A_l and B_l are (in general complex) constants whose values depend on the boundary conditions for the problem at hand. The angular parts of this solution are commonly joined together to form the spherical harmonic functions $Y_l^m(\theta,\, \phi) = P_l^m(\cos\phi)e^{-im\theta}$ such that the solution to Eq. (3.3) is simply a weighted sum of spherical harmonics. An equivalent way of stating this is to say that the spherical harmonics form a complete basis for Laplace's equation. Accordingly, shimming is often performed by the construction of several separate

Shorthand	Order (l)	Harmonic (m)	Functional Form
Z	1	0	z
X	1	1	x
Y	1	-1	y
$Z2$	2	0	$z^2 - (x^2 + y^2)/2$
XZ	2	1	$3xz$
YZ	2	-1	$3yz$
$X2 - Y2$	2	2	$3(x^2 - y^2)$
XY	2	-2	$6xy$

Table 3.1: List of spherical harmonics that are available as room-temperature corrections to B_0 on the used system's shimset.

coils whose generated fields are particular spherical harmonics. As spherical harmonics form an orthonormal basis set, to an excellent approximation currents present in one coil do not induce a field in any others. Correspondingly, the user of an MR system can "dial in" separate currents into each coil in the shimset to maximise the homogeneity of B_0. The system used in this work features a second-order room-temperature shimset, detailed in Table 3.1, where the maximum current through each coil is ± 2.5 A per channel. Shimset currents are quantised as 'DAC units'; 16-bit integers which take values between 32767 and -32767 corresponding to ± 2.5 A. The superconducting shimset is nearly second order, with six of the possible eight channels available for adjustment.

Field homogeneity is typically assessed by the acquisition of multi-echo B_0 mapping sequences, in which two sets of images are acquired at different echo times separated by $\Delta\tau_E$. Neglecting relaxation, the only difference between each echo should be accumulated phase in each voxel arising from changes in B_0. Accordingly, the angle of the element-wise complex division of the two images yields a phase map where the local phase change can be related to ΔB_0 by $\Delta\phi = \gamma \Delta B_0 \Delta\tau_E$. This process allows for the quantification of field homogeneity as a function of space, and is the basis of several automated shimming algorithms, through which the currents flowing in each coil are changed to elicit a more homogeneous B_0. It is worth noting that spherical harmonics are independent in 3D, but degenerate in 2D, meaning that it is fundamentally preferable to acquire 3D B_0 maps rather than thin slices. It should be noted that manually shimming a region of interest is difficult, and prone to local minima due to the fact that the higher order spherical harmonics in Table 3.1 can also contain lower order terms.

An automated shimming algorithm developed[38] by Jürgen Schneider and installed and partially calibrated by the author was used for the majority of work in this thesis; its user interface is shown in Fig. 3.15. The algorithm requires mapping behaviour of the shimset of the magnet by iterating over all possible DAC values, and calculating an impulse response function by acquiring B_0 maps of the resulting field perturbations on a homogeneous phantom. This process takes several days, but allows for the estimation of shim DAC values required to elicit a change in B_0, neglecting magnetic susceptibility effects. Shimming is then performed by acquiring (cardiac gated) segmented 3D multi-echo B_0 maps of

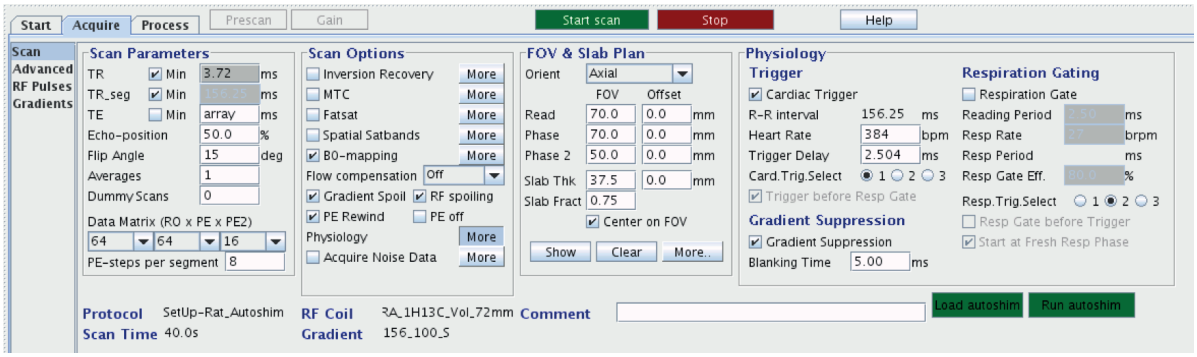

Figure 3.15: The user interface, developed by the author, for the autoshim protocol developed by Jürgen Schneider. Clicking the "run autoshim" button loads an external utility that predicts the ADC values to optimise homogeneity over a defined region of interest, and writes them to a file. These values are then loaded in with the 'load autoshim' button, and another B_0 map can be acquired to verify homogeneity.

the sample at hand, defining a region of interest and predicting the DAC values to ensure its homogeneity as quantified through a least-phase metric. This procedure is automated, and allows for the definition of a manual region of interest, further improving homogeneity at the desired location potentially at the expense of B_0 homogeneity outside that region. Moreover, it allows for some degree of compensation for the various electrodynamic effects that make cardiac shimming difficult, such as the air-tissue interface in the lungs where susceptibility effects are prevalent, and local field inhomogeneities arising from the paramagnetic nature of blood. The algorithm is regularly able to shim the rat heart to < 70 Hz homogeneity, which is in excess of that which can easily be achieved by hand. For the heart, the segmented sequence used obtained 8 phase encoding lines per RR trigger with two different TE$= 1.4, 2.2$ ms and TR$= 3.15$ ms, from a $64 \times 64 \times 16$ matrix in a $70 \times 70 \times 50$ mm^3 FOV. The flip angle was $15°$, with two averages, and a gradient suppression delay of 5 ms inserted at the end of each stack of phase encodes acquired to minimise self-gating. When scanning objects other than the heart (i.e. phantoms or the brain), the matrix size was increased to $96 \times 96 \times 32$ in a $65 \times 65 \times 65$ mm^3 FOV, cardiac gating turned off, the TR increased to 15 ms, and the flip angle decreased to $8°$.

3.5 | *Dynamic Nuclear Polarisation*

As described previously (Section 2.5), the processes that lead to Dynamic Nuclear Polarisation are comparatively well understood. The method can be used to pump an appreciable fraction of nuclear spins in a given molecule into a given state, and requires a variable-temperature superconducting cryostat with a high-power microwave source in order to do so.

3.5.1 *Experimental realisation*

In order to directly probe metabolism *in vivo* (as described later), we polarise labelled metabolites, commonly [1-^{13}C]pyruvate, with the organic radical 'OXO63' (Fig. 3.16), a carbonyl free radical suitable

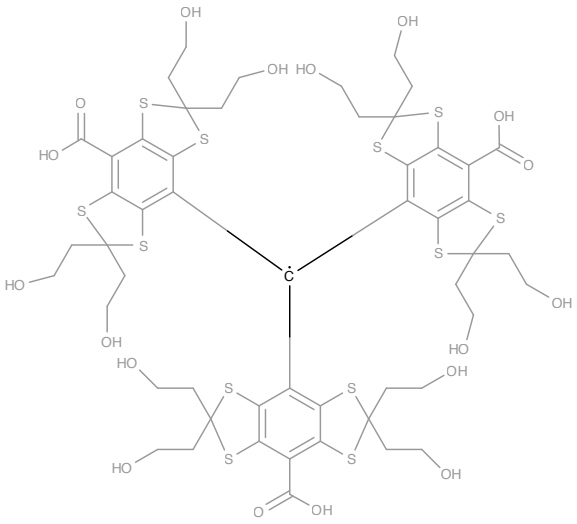

Figure 3.16: The structure of the OXO63 radical. The carbonyl radical (centre, black) provides an 'effectively free' electron to participate in the DNP process. The three (identical) trityl moieties surrounding the radical limit its chemical reactivity, eliminating toxicity *in vivo* that might otherwise occur from the injection of a reactive carbon species at any concentration. The large number of exterior hydroxyl groups on each moiety ensure that the radical is hydrophilic, and thus able to quickly enter aqueous solution. Interactions between the carbonyl radical and the 'cage' provided lead to a large (and approximately Gaussian) electronic linewidth, as demonstrated elsewhere.[39] OXO63 can be synthesised from tetra-*tert*-butylthiobenzene in house with a low yield, or commercially acquired from Oxford Instruments.[40]

for *in vivo* applications due to its low toxicity. OXO63 has a broad electronic linewidth,[41] and thus the dominant mechanism through which carbon polarisation is acquired is thermal mixing. This allows for the pumping process to occur with relatively low microwave power. The cryostat used contains an integral radiofrequency coil; it is thus possible to measure the build up of nuclear polarisation through NMR. The system is a converted superconducting 3.35 T NMR cryostat, with a variable temperature inset (VTI) that directly couples the helium reservoir to the sample chamber, which can also have its pressure varied by high vacuum pumps. The sample is contained in a small plastic (polyetheretherketone, PEEK) cylindrical cup, which is placed in a microwave and radiofrequency resonant cavity; a microwave waveguide and coaxial cable connect the cavity to the microwave and radiofrequency waveform generators. The cup is screwed onto a long hollow polytetrafluoroethylene (PTFE) tube which is stoppered when the system is evacuated. Temperature monitoring within the VTI is performed by calculating the resistivity of a set of small resistors. A separate permanent magnet and RF coil provides the ability to estimate the polarisation level of the sample post dissolution, again by NMR.

A LabView PXI controller[13] computer is responsible for electronically controlling the majority of these disparate pieces of equipment. A 4U PXI rack of bespoke electronics together with two NI-DAQ ADCs provides a custom interface to the microwave controller, VTI helium level monitoring and control

[13]PXI, and its successor PXIe, are open hardware platforms designed for industrial computing; it provides a bus interface for devices, and permits a large number of logically addressable devices that is based on the popular PCI or PCI-e protocol found on computer motherboards.

systems, and RF waveform synthesis for low-field NMR. The PXI controller (National Instruments PXI-8176) has two separate ADCs to handle waveform generation and FID digitisation. A LabView user interface sets the centre frequency of microwave irradiation, measures the temperature and pressure of the VTI, estimates the spin temperature of the sample, and controls the liquid helium level inside the VTI. W-band (94 GHz) microwaves are generated by an independent controller, an EVLA-1 which produces high power microwaves with an approximately Gaussian frequency profile with a 0.5 GHz bandwidth. The vacuum system is controlled manually, and consists of two pumping stages; rotary vane, and Roots. After the slight modification of the VTI inset to ameliorate leaks, the system can be pumped down to 10^{-2} mbar, which corresponds to a ^{4}He temperature of 0.8 K and is limited by the pumping system. As it is regularly opened to the atmosphere to load samples, there is a gradual build-up of (predominantly oxygen) ice which is ameliorated by the yearly removal and thawing of the inset. Dissolutions are performed by heating the required solvent in a pressure chamber at the end of a long stainless steel stick to 10 bar, re-pressurising the VTI to ~ 1.1 bar with helium gas, physically lifting the sample out of any liquid helium present, docking the cup to the end of the long stick, opening a valve and injecting the superheated solvent onto the frozen spin-polarised glass directly. An exit port then collects the resulting liquid in a plastic chamber, where it is sucked into a syringe before injection.

This prototype hyperpolariser is illustrated in Fig. 3.17**A** and **B**. Data showing the measured polarisation as a function of applied microwave frequency (a 'microwave sweep' curve) together with the NMR signal amplitude as a function of time of microwave irradiation (a 'build-up curve') is presented in Fig. 3.17**C** and **D**. The non-symmetric nature of the two peaks in the microwave sweep is indicative of thermal mixing, as is the slight, but significant, deviation from single-exponential behaviour in the build-up curve.[26] It is also interesting to note that the separation between $\omega_e + \omega_N$ and $\omega_e - \omega_N$ is far less than the bandwidth of the microwave source, indicating the presence of processes at equilibrium. Irrespective of the microscopic mechanisms occurring, the system can reproducibly polarise 80 mM [1-^{13}C]-pyruvic acid to 40%, measured after dissolution with NaOH. An example of the gain in the signal-to-noise ratio (SNR) of spectra acquired with the hyperpolarisation technique is shown in Fig. 3.17**E**.

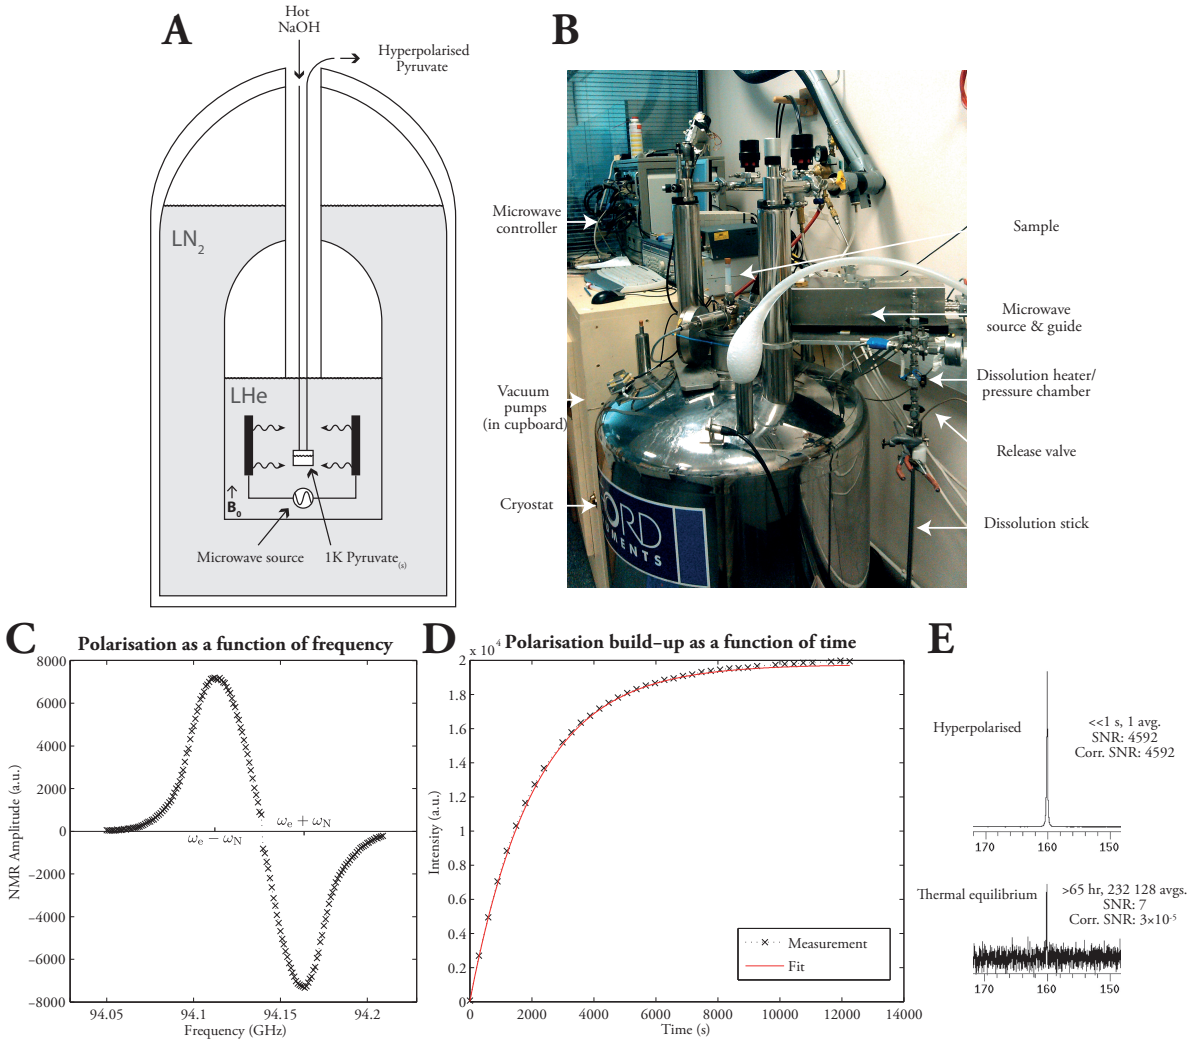

Figure 3.17: Hyperpolarisation is carried out in a 1 K cryostat with a movable sample holder to facilitate rapid dissolution and heating of the sample to 37°C. **A**: A cartoon of the process: the sample is held in a pumped liquid He cryostat at approximately 1 K, and irradiated with microwaves in a resonant cavity. When nuclear polarisation has built up to an acceptable degree (as measured by NMR in the cryostat), hot neutralising solvent at 10 bar is introduced, and the resulting polarised liquid collected for injection into an organism in a nearby MRI scanner. **B**: A labelled photo of the equipment used, described in detail elsewhere.[42] **C**: Measured NMR signal intensity ($\propto$ polarisation) as a function of applied microwave frequency with [1-^{13}C]pyruvate and OXO63. The two peaks correspond to pumping at $\omega_e - \omega_N$ and $\omega_e + \omega_N$ respectively. **D**: Measured polarisation when pumping at $\omega_e - \omega_N$. Shown is a fit to $A(1 - e^{-t/T_B})$, an approximation of the multi-exponential behaviour proposed for the build-up of polarisation in the system of copper-tutton salts.[43,44] We find that a time-constant T_B of 2000 s describes the system well; the minimum time between subsequent experiments is therefore $\sim$ 30 minutes. **E**: The net effect of the hyperpolarisation process is to procure a massive increase in the signal-to-noise ratio (SNR). As this figure by Ardenkjaer-Larsen *et al.* shows, the technique running on this equipment generates an increase in SNR on the order of 10^5.[42]

Chapter 4

Motivations and Preliminary work

4.1 | *Metabolism*

Energetically, life has to create local order. The large, complex set of chemical reactions that facilitate this process are collectively known as 'metabolism', and are subject to regulation at multiple points. As a network, metabolism has design features that ensure it is robust, error-tolerant, and scale-free.[45] The scale-free nature of metabolic networks implies that their node connectivity is described by a power-law relationship; that they are highly heterogeneous; and that their connectivity and overall topology is determined by a few highly connected nodes which link the rest of the (less connected) nodes to the system. Correspondingly, the average distance between any two metabolites is small (approximately three nodes), meaning that a desired molecule can be constructed rapidly from available ones. This suggests that organisms have evolved towards robust networks that allow the rapid and efficient response to external stimuli or internal errors. Such properties are emergent from dynamic systems that are subject to evolution, and are considered common to all cellular networks.[45–47]

Pyruvate is a metabolite that is a highly connected node in the metabolic network, as illustrated in Fig. 4.1. Pyruvate correspondingly has multiple metabolic fates: it can be reduced in the cytosol by lactate dehydrogenase (LDH) to form lactate; irreversibly decarboxylated with co-enzyme A in the pyruvate dehydrogenase complex (PDH) to produce acetyl-CoA; carboxylated by pyruvate carboxylase to form oxaloacetate; or transaminated by alanine transaminase to synthesise alanine.[48, Ch. 2] Under physiological conditions, pyruvate, whose $pK_a = 2.39$, is almost entirely dissociated into its anionic (charged) form. Accordingly, the diffusion of the molecule through lipid bilayers is not possible and it is instead actively transported by members of the monocarboxylate transporter (MCT) family of integral membrane proteins into the cell. Under physiological conditions, the majority of pyruvate used by the cell is generated endogenously from glucose by glycolysis, and from the oxidation of lactate. Under aerobic conditions, pyruvate and its flux through PDH act as the entry point for carbohydrate metabolism into the Krebs cycle, which then oxidises the carbon transferred. Reflecting its role as a gateway to the Krebs cycle, PDH is spatially localised to the inner mitochondrial membrane, and pyruvate is actively transported across the outer mitochondrial membrane by a family of transporters,

the mitochondrial pyruvate carriers (MPCs). Carbon dioxide is cleaved from pyruvate by PDH, and is in rapid equilibrium with bicarbonate (HCO_3^-), catalysed by carbonic anhydrases in the cytosol. In the presence of fatty acid oxidation, also known as β-oxidation, acetyl-CoA is produced independently of pyruvate, and PDH flux is inhibited accordingly. The kinase pyruvate dehydrogenase kinase (PDK) also serves to inhibit PDH flux, and is itself is highly regulated. For example, high concentrations of ADP and pyruvate inhibit PDK, increasing the flux through PDH. If sufficient oxygen is not available, pyruvate can undergo reduction in the cytosol by LDH, mediated by the relative concentrations in the cytosol of NADH and NAD^+. Pyruvate Carboxylase (PC) does not ordinarily significantly contribute to pyruvate metabolism in mammals as the reaction pyruvate $\rightarrow$ oxaloacetate is energetically disfavoured, requiring ATP. The reverse reaction, oxaloacetate $\rightarrow$ pyruvate, is one of the key gluconeogenic pathways that allow the synthesis of glucose from other precursors.[49–51]

The use of hyperpolarised pyruvate allows for these reactions to be measured *in vivo*, as the downstream metabolic products of pyruvate can be followed by changes in chemical shift. If labelled in the first position, hyperpolarised [1-^{13}C]pyruvate forms an effective probe of metabolism, as the visible multiple fates of pyruvate report on the rate of β-oxidation, glycolysis, and amino acid synthesis in the cell. An illustration of the labelled products visible with hyperpolarised [1-^{13}C]pyruvate is shown in Fig. 4.2; decreases in glycolysis consequently serve to reduce the visible flux to lactate. An increased rate of fatty acid β-oxidation serve to increase the concentration of acetyl-CoA entering the TCA cycle, inhibiting PDH and decreasing the visible flux of [1-^{13}C]pyruvate to ^{13}C-carbon dioxide and ^{13}C-bicarbonate. These fluxes, therefore, act as the (continuous) 'logical not' of fatty acid oxidation.

As might be expected for a complex network subject to regulation at numerous different points, metabolic dysregulation is implicated in disparate diseases. This work primarily deals with cardiac disease and cancer, and a brief discussion of the main metabolic changes occurring in each disorder follows.

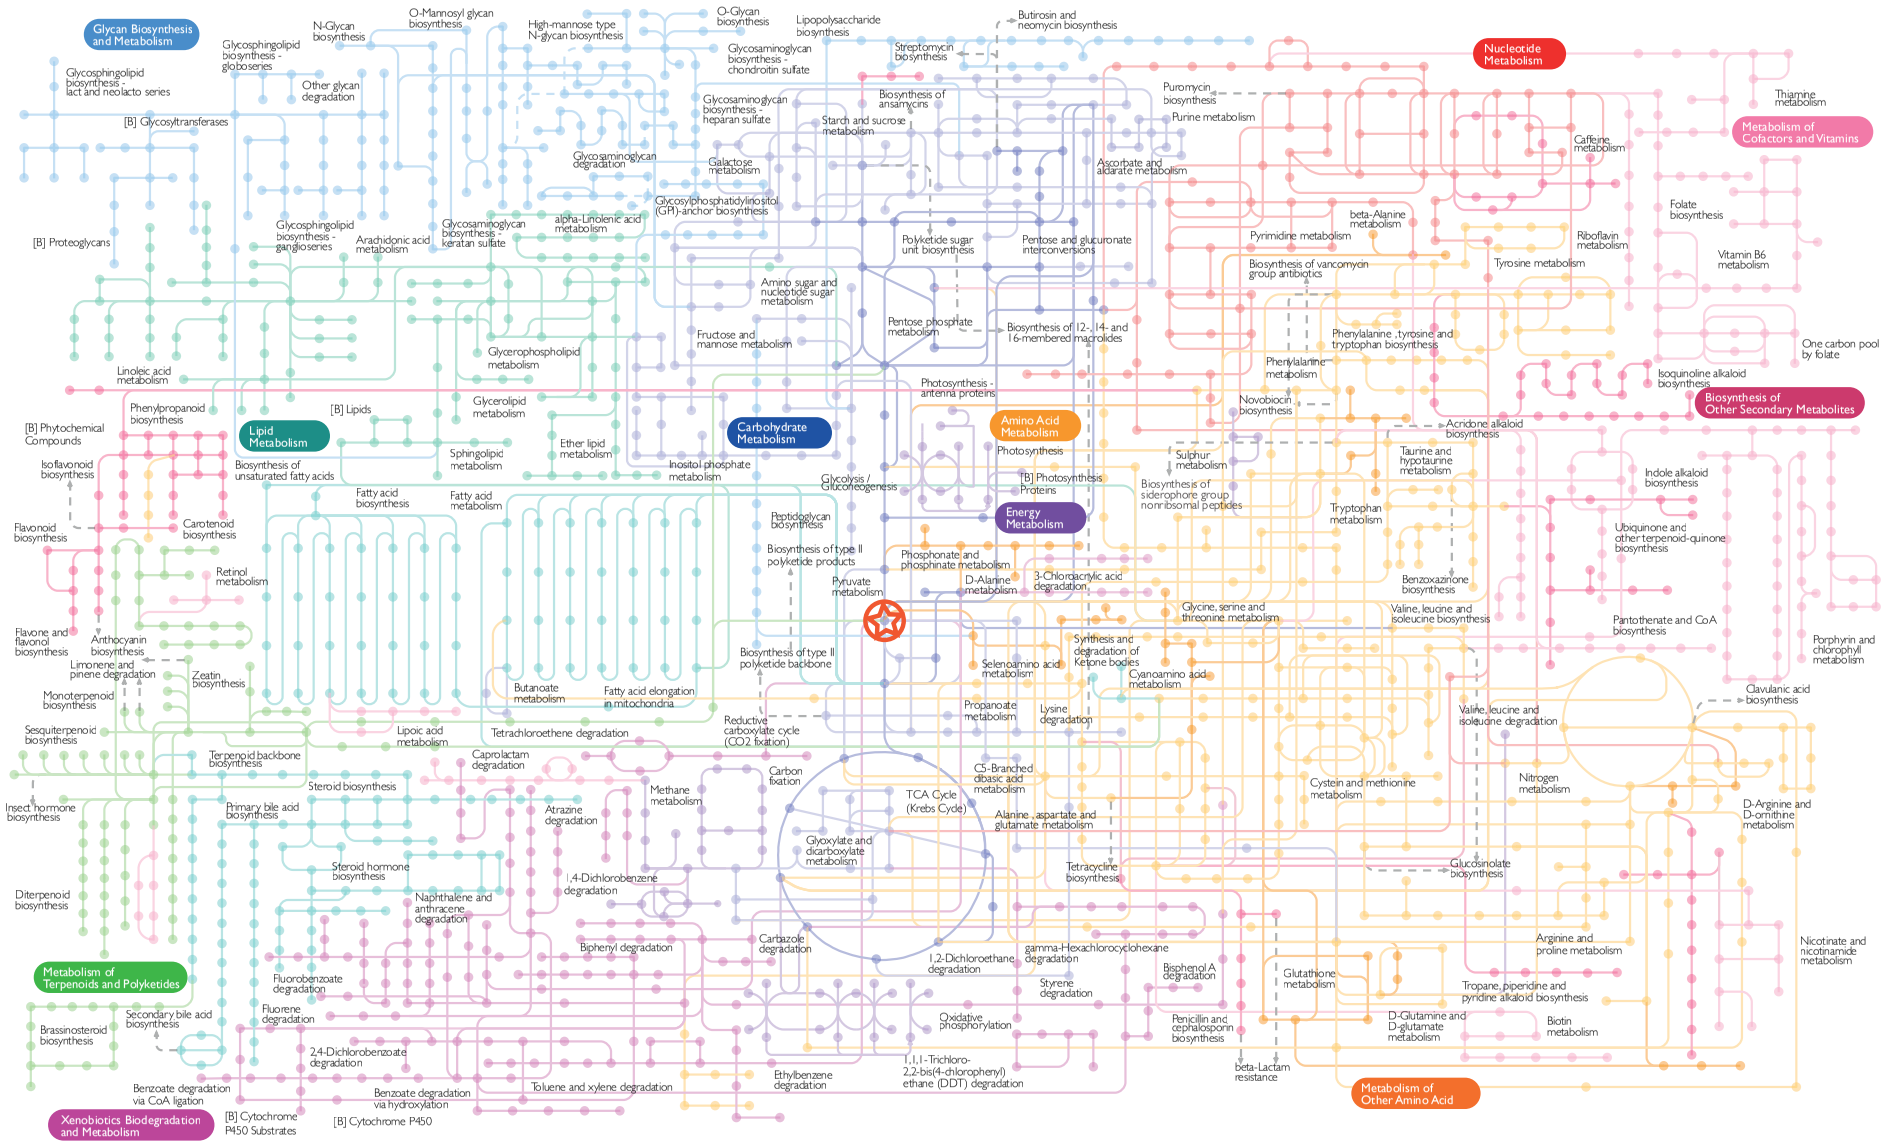

Figure 4.1: A brief overview of the chemical reactions that energetically power life, reconstructed from reactions documented in the KEGG database.[52] The location of pyruvate is highlighted (with a starred circle) and the names of several commonly discussed pathways are included. The complexity and connectedness of the metabolic network is apparent. Such datasets are frequently used for attempts at understanding the control of metabolism as a large, interacting set of chemical species with (mostly) known chemical kinetics. DNP aims to provide *in vivo* values of these rate constants for larger organisms, which are otherwise known to rapidly change.

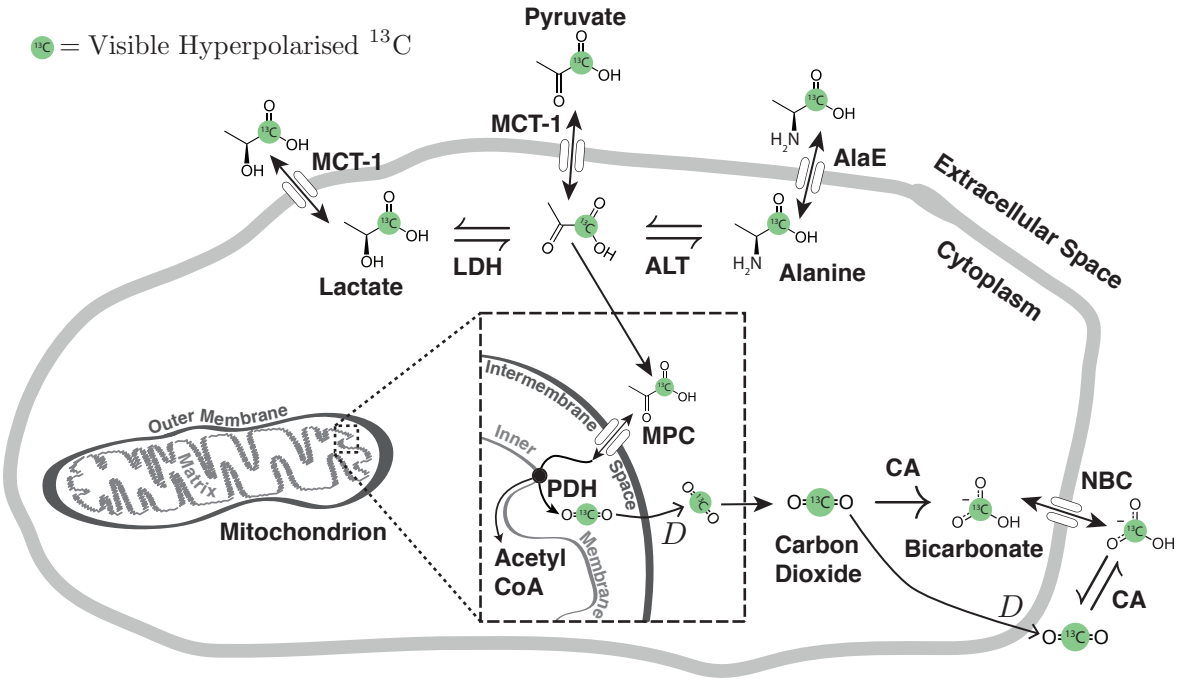

Figure 4.2: A cartoon of the possible metabolic fates of hyperpolarised [1-^{13}C]pyruvate in the cell. Pyruvate enters the cell through the monocarboxylate transporter MCT-1 and can then either be transaminated to form alanine by ALT; removed from the cytosol by protein synthesis or by the alanine exporter AlaE; reduced by LDH to form lactate, which is then extracellularly exported by MCT-1; or transported to the intermembrane space in the mitochondrion by the mitochondrial pyruvate carrier (MPC) before being decarboxylated by PDH. The resulting labelled $^{13}CO_2$ is able to diffuse across lipid membranes before equilibrating with bicarbonate in the cytosol, catalysed by carbonic anhydrases (CA) prior to its export by the sodium bicarbonate carrier (NBC). As acid-base balance is vital to biological systems, most species also contain CAs in their circulatory systems.

4.2 | *Cardiac disease*

4.2.1 *Heart failure*

The failing heart, by definition, is unable to supply the body with sufficient blood to permit function with unaltered physiology.[53] Heart failure (HF) has several clinically distinct ætiologies that share either some degree of myocardial injury, or an increase in the expected myocardial workload. These causes can be abrupt and sudden, such as a myocardial infarction, or gradual and progressive, such as systemic pressure overloading due to hypertension or stenosis. HF is the end-point of many cardiac or cardiovascular diseases.[54] With a median survival time of five years, and limited options for clinical management, the prognosis for HF is poor.[55] It is now generally accepted that changes in myocardial metabolism are implicated in HF, although the nature of the metabolic alterations in question is contentious.[2] At the end stages of the disease, there is a profound shift towards increased glucose oxidation as opposed to fatty acid oxidation. Decreased fatty acid oxidation is present at both early and late stages of the disease, together with maladapted mitochondrial morphology, function and biogenesis, and an increase in anaplerotic reactions that channel carbon back towards the Krebs cycle. These alterations are concomitant with a decrease in ATP availability in the failing heart, but the causal

link between metabolic alteration and HF remains uncertain.[56]

As the phenotype of the failing heart is reduced contractile function, the traditional management of the disease consisted of pro-contractile (positively inotropic) drugs such as dobutamine that would increase the force of cardiac muscle contraction and, it was believed, increase ejection fraction. However, these drugs increase energy usage in the failing heart, which is now expected to worsen cardiac performance in the long term. They fail to have a long-term benefit in patients, but maintain a role in certain specific situations and during cardiac surgery.[57] Current chronic HF treatments focus on reducing the load on the cardiovascular system, either by surgically increasing the diameter of compromised blood vessels; by mechanically unloading the left ventricle of the heart; or by pharmacologically activating various aspects of the neurohormonal systems. Treatments such as beta-receptor blockade,[58,59] ACE inhibitors[60,61] or angiotensin II blockers[62,63] are energy-sparing, and may counter-intuitively improve the prognosis of HF by decreasing contractility. The efficacy of these compounds in current clinical use motivates the development of novel pharmacoactive agents to directly target myocardial metabolism, and additionally hints at a causal link between metabolism and HF.

Fatty acid metabolism requires more oxygen per mole of ATP generated than glucose oxidation, and correspondingly results in an increase in lactate production in situations of reduced oxygen availability. As the leading causes of HF are vascular disease and myocardial infarction, some reduction in oxygen availability is at least transiently likely.[55] Under these conditions, the continued production of lactic acid disrupts the acid-base balance of the cell and leads to tissue acidosis, which leads to a decrease in myocyte contractility.[64] Accordingly, several partial inhibitors of fatty acid oxidation ("pFOX inhibitors") have been developed and show promise in ameliorating HF.[65] It is known that the agents etomoxir, oxfenicine and perhexiline are mitochondrial carnitine palmitoyl transferase I inhibitors, and therefore block mitochondrial fatty acid uptake and subsequent oxidation.[66] Unfortunately, etomoxir failed phase-II clinical trials due to excessive toxicity;[67] but clinical trials of oxfenicine[68] and perhexiline are ongoing.[69,70] Other pFOX inhibitors have been developed, such as the piparazine derivatives trimetazidine and ranolazine, but a mechanistic understanding of their action is lacking as they were initially developed as anti-anginal agents. As others have commented,[71] the successful management of HF by decreasing fatty acid utilisation is counter-intuitive, as end-stage HF is associated with a profound decrease in fatty acid oxidation.

The modulation of metabolic substrate utilisation for the treatment of heart failure forms a promising avenue for the future management of the disease. However, the efficacy of agents already under investigation is not conclusive, and their mechanistic role of action is unknown. The safety of the agents discussed above is uncertain, and the utility of shifting substrate selection on a population level unconfirmed. Therefore, there is a need both for further preclinical investigation into the mechanistic role of metabolism in HF, and for techniques to non-invasively assess cardiac metabolism *in vivo*, ideally

in human clinical trials.[71]

4.2.2 *Myocardial ischæmia, infarction, and metabolic remodelling*

Ischæmia is defined as the pathological reduction of blood flow to tissue,[72] and is a condition in which the metabolism of pyruvate is altered. In the heart, myocardial ischæmia commonly originates from an obstruction in the coronary arteries, and the resulting hypoxic state rapidly causes a switch towards anaerobic glycolysis to maintain ATP synthesis.[73] This is associated with a subsequent increase in the rate of glycogen breakdown and glucose uptake.[74] This increased rate of glycolysis results in the production of lactate and NADH, leading to acidosis and a decrease in myocyte contracility. In practice, patients undergoing some form of myocardial ischæmic event experience a 'wavefront' of injury, with the hypoxic region starting in the endocardium (interior surface of the myocardium) and progressing towards the epicardium (external).[75,76] The functional response of the myocardium to ischæmia is complex, and depends on the length and severity of the ischæmic event. Prior to cell death due to ischæmia, the myocardium can be 'stunned', referring to recoverable contractile dysfunction after the return of normal blood flow; or 'hibernating', akinetic or with impaired contractile function during ischæmia.[64,77] If the ischæmic insult is sufficiently large, the resulting acidosis is great enough to cause necrotic cell death, resulting in a myocardial infarction (MI, commonly known as a "heart attack"). Myocardial infarction is a common cause of subsequent heart failure.[78] It is currently difficult to distinguish between viable, hibernating myocardium and tissue that is dead; accordingly, most patients believed to have acute MI are the subject of interventional revascularisation therapies that attempt to reperfuse and rescue the ischæmic myocardium. Interventional reperfusion remains the front-line treatment for acute MI,[79] and it is effective: reperfusion following acute MI is associated with lower rates of early and late death following the MI, and fewer hospital readmissions for subsequent ischæmia or HF.[80]

Reperfusion following ischæmia triggers an acute set of metabolic and inflammatory responses.[81] It is these inflammatory responses that ordinarily lead to wound healing and repair following an injury. However, in the infarcted myocardium the influx of macrophages, neutrophils and lymphocytes into the tissue can be disadvantageous. Dysregulated inflammatory responses or metabolic changes that lead to further injury are termed 'reperfusion injuries', and have been shown in preclinical models of disease to worsen outcome.[82,83] In humans, sustained inflammation following MI is strongly correlated with adverse outcome and reduced survival,[84] yet attempts at improving patient outcomes through the use of anti-inflammatory drugs have not been successful.[85–87] Therefore, there is a need to be able to resolve this 'reperfusion conflict', and understand under which circumstances reperfusion is deleterious. The metabolic signature of inflammation – increased regional glycolysis arising from an increased density of highly glycolytic T-cells – suggests that DNP with hyperpolarised [1-^{13}C]pyruvate might provide an

approach to probe the response of the heart to injury.

If the MI is survived, a fibrotic scar forms over a period of weeks that is believed to lessen the probability of cardiac rupture. This scar is non-contractile, and correspondingly decreases the functional performance of the organ by reducing the ejection fraction of the relevant chamber. It is this change in functional performance that drives the development of HF from MI: the effective workload of the remaining tissue is increased following a decrease in its useful stroke volume. This period of (mal-)adaptation is known as myocardial remodelling, and occurs over a comparatively long timescale following an MI. Increasing cardiac workload is a necessary response to the inability of the damaged heart to adequately perfuse the body following MI, and is achieved by cytokine and neurohormonal signalling, resulting in increased heart rate and increased force of myocardial contraction. Such changes provide short term support by maintaining cardiac output, but ultimately can drive the descent into heart failure following MI.[88]

4.2.3 Cardiac disease and hyperpolarised [1-^{13}C]pyruvate

As outlined above, metabolism is of interest in cardiac pathology. Changes in the metabolic phenotype could act as a direct measure of the ischæmic tissue in MI, and report on metabolic remodelling and HF over greater periods of time as the disease progresses. Such changes are of interest for basic research into the disease, as it is likely that any successful therapy for HF will have a metabolically modulating effect. As metabolic alterations in HF have been known for decades (starting with the 1939 observation of decreased creatine content in the failing heart by Herrmann and Decherd[89]), there is a considerable body of evidence showing the utility for metabolic research into the disease. Accordingly, cardiac metabolism was one of the first areas of study following the development of DNP, with Merrit *et. al* showing in the perfused rat heart that the production of visible bicarbonate following hyperpolarised [1-^{13}C]pyruvate injection is obtained exclusively by flux through pyruvate dehydrogenase.[90] Further work in the perfused heart by Schroeder *et. al*[91] revealed that hyperpolarised [2-^{13}C]pyruvate could act as an effective probe of Krebs cycle intermediates, and notes a substantial increase in visible lactate in the ischæmic myocardium. Additionally, hyperpolarised [1-^{13}C]pyruvate has been shown to discriminate between acute and chronic MI by metabolic phenotype,[92] and to be sensitive to the substrate switching occurring in both hyperthyroidism[93] and during cardiac hypertrophy as probed in the spontaneously hypertensive rat.[94] Likewise, a significant decrease in the visible bicarbonate/pyruvate ratio is found in the failing heart, with dilated cardiomyopathy being associated with a reduction of a third in the bicarbonate/pyruvate ratio in pigs.[95]

These studies have fundamentally demonstrated that hyperpolarised [1-^{13}C]pyruvate spectroscopy is able to effectively probe metabolism in heart disease, and is sensitive to the known differences in metabolism between disease states. As the field has developed, interest has shifted towards the spatial

imaging of hyperpolarised [1-^{13}C]pyruvate and its downstream metabolic products in disease. The utility of such approaches to image myocardial ischæmia has been verified in large animal models of MI, with myocardial lactate production during coronary occlusion providing a direct visualisation of ischæmia.[96] As discussed in Chapter 5, imaging hyperpolarised [1-^{13}C]pyruvate and its downstream products is technically challenging. However, current clinical techniques used in the management of patients with MI rely on indirect measures of ischæmia (such as perfusion abnormalities or changes in wall motion during stress), and it is therefore hoped that the direct visual assessment of myocardial ischæmia in the form of 'lactate imaging' will be able to aid guided interventional therapies, such as revascularisation. As revascularisation only forms an effective therapy if targeted to the region of significant ischæmic burden, it is hoped that hyperpolarised [1-^{13}C]pyruvate imaging will substantially improve patient outcome.[97]

In summary, a substantial body of evidence exists that demonstrates the utility of using hyperpolarised endogenous metabolites, specifically [1-^{13}C]pyruvate, to probe metabolism in cardiac disease. Since its introduction in 2003, dissolution-DNP has progressed from being a new technique validated with known biochemistry in *ex vivo* situations, to a tool that can garner information about the metabolic state of an organism *in vivo* and act as a diagnostic aid. In a cardiac setting, DNP has been firmly established as a sensitive tool that is of interest both clinically and for basic science. As noted by Rider and Tyler, there is a need for methods to image hyperpolarised metabolites at high spatial and temporal resolution in order to effectively define the viable myocardium in situations of cardiac ischæmia.[97]

4.3 | *Metastatic cancer*

4.3.1 Metabolism and Cancer

In addition to cardiovascular disease, metabolic dysregulation is also a hallmark of cancer.[1] It has been known since 1931 that cancer cells metabolise differently from non-malignant cells, primarily regenerating ATP via a high rate of non-oxidative anaerobic glycolysis, as opposed to a low rate of glycolysis followed by the complete oxidation of pyruvate.[1,98] This greatly increased rate of glycolysis, known as the Warburg effect, results in lactate excretion from the cells lowering extracellular pH, and killing neighbouring healthy tissue. A number of chemotherapeutic agents therefore attempt to block the glycolytic pathway.[99] It is current clinical practice to assess the metabolism of tumours using ^{18}F-fluorodeoxyglucose Positron Emission Tomography (^{18}F-FDG PET), in which the radiolabelled glucose analogue ^{18}F-FDG is administered to the patient. ^{18}F-FDG is phosphorylated by hexokinase, but cannot be isomerised by phosphoglucose isomerase, and hence accumulates in tissues with high metabolic activity, such as malignant tumours, the heart, kidneys, and the brain. The ^{18}F nucleus then undergoes β^+ decay with a half life of 109 minutes, emitting a positron with a maximal kinetic energy of 634 keV[100] and an electron neutrino (Fig. 4.3**A** and **B**). This positron then annihilates (Fig. 4.3**D**), and

the resulting back-to-back pair of 511 keV photons are detected by a 'gamma camera', typically either a scintillation crystal backed by photomultiplier tubes or an array of doped semiconductors with a small bandgap that produce spectroscopic as well as positional information (Fig. 4.3**E**). Techniques such as filtered backprojection can then reconstruct the likely origin of the positron, producing a map of glucose uptake (Fig. 4.3**F**). Owing to the large radiation dose received by the patient (typically $\sim$15 – 30 mSv, approximately equivalent to 3000 chest X-rays) with a concomitant increased risk of cancer,[101] PET is clinically limited to oncology settings, with several common conditions (such as youth and pregnancy) often serving as relevant contraindications. Tumour metabolism is inferred by assuming excessive glucose uptake is synonymous with glycolysis.

PET is a mature technique that is used extensively in the clinic, reflecting the utility of metabolic imaging in diseases where metabolism is changed. However, PET features several key limitations, which have driven research into DNP as an alternative metabolic probe in oncology. DNP has three key advantages over PET as a method to acquire metabolically relevant information *in vivo*, namely:

4.3.1.1 Safety

The fact that ^{13}C is a stable nucleus with a measured half-life greater than the estimated age of the universe represents a significant safety advantage compared to PET. Patients need not be subjected to a dose of ionising radiation, and consequently the technique is more likely to find use outside of oncology. In humans, DNP with hyperpolarised [1-^{13}C]pyruvate is safe.[102]

4.3.1.2 Limiting resolution

PET is fundamentally a low spatial and temporal resolution technique. The radioactive labels used in PET necessarily need to be present at low concentrations *in vivo* to avoid excessive doses of radiation to the patient. This, combined with the stochastic nature of radioactive decay (and in particular the relatively long half-life of the ^{18}F nucleus) serve to decrease the resolution of the final images, and increase the time taken for the scan to complete. Additionally, the required conservation of four-momentum means that the probability for a high-energy positron annihilating is very small as the available phase-space for the decay to occur is constrained. Hence, the annihilation event only occurs once the positron has a comparable energy to nearby electrons. Thus, the ultimate limiting spatial resolution of PET is determined by the random walk undertaken by the emitted positron as it loses energy due to electrostatic interactions with the surrounding matter (Fig. 4.3**C**), fundamentally limiting the maximum possible spatial resolution to ~ 1 mm^3. In practice, this resolution is further degraded by detector physics (typically to ~ 5 mm[103, p. 110]) and the inevitable small motion of the patient during the half-hour long scan times.[104] The fundamental limiting spatial resolution for imaging hyperpolarised compounds is not as clearly defined, being determined by the available M_0 and the efficiency of the hardware. The imaging strategies proposed in this thesis function with spatial resolution comparable to the limiting

resolution of PET. The dynamic nature and (transiently) high SNR of the technique means that it is able to operate at superior spatiotemporal resolution than PET, if compromises are made in temporal and spectral resolution.

4.3.1.3 Metabolic versatility

A given PET scan can only probe one particular part of a metabolic pathway, and does so 'statically'. By utilising a radio-labelled *analogue* of a particular molecule of interest that is itself not then further metabolised, PET necessitates often complex radiochemistry and provides only a snapshot of the concentration of the radio label after a long period of time. PET is unable to identify uniquely the molecule containing the nucleus that the emitted positron originally came from, and in contrast to magnetic resonance techniques, is not therefore chemically selective. In reality, metabolic reactions form a system in dynamic equilibrium where fluxes are rapidly regulated, and the flux observed through the nth step in a pathway with a molecule designed to accumulate in the $n + 1^{\text{th}}$ step may not be the same as that flux with its native metabolite. PET, therefore, trades specificity for metabolic versatility. Developing different tracers is difficult as each molecular analogue needs to be designed and individually approved as a pharmaceutical agent. There is no expected commonality to their safety profile. Accordingly, despite the large number of reported tracers in the literature only a small number are approved for use in humans, and ^{18}FDG remains the only molecule routinely used in clinical practice.[107] In contrast, DNP allows endogenous metabolites to be MR-visible, sidestepping both the radioactive dose and radiochemistry requirement. The different chemical shift of different metabolites also ensures that multiple steps in a metabolic pathway can be observed (if metabolism is sufficiently rapid compared to T_1), and there is no need to use analogues that are accumulated rather than metabolised. Additionally, the fact that most endogenous metabolites have well-characterised toxicological data and measured physiological ranges *in vivo* is likely to shorten the route to regulatory approval for novel 'DNP tracers'.

4.3.2 DNP and Cancer

The majority of studies using DNP in an oncology setting published to date have focussed on the [1-^{13}C]pyruvate $\rightarrow$ [1-^{13}C]lactate reaction specifically.[108] Hyperpolarised [1-^{13}C]pyruvate has been shown to form a sensitive probe to tumour metabolism, with an increase in the apparent rate constant for lactate production reflecting the underlying shift to glycolysis. Oncology formed one of the earliest applications of the technique since its inception in 2003, with Golman reporting a doubling of the visible [1-^{13}C]lactate intensity present in implanted P22 tumours in rat in 2006.[109] More recently, hyperpolarised [1-^{13}C]pyruvate has been shown to form an effective therapy response marker in oncology, with a decrease in the visible [1-^{13}C]lactate signal correlated with a positive response to

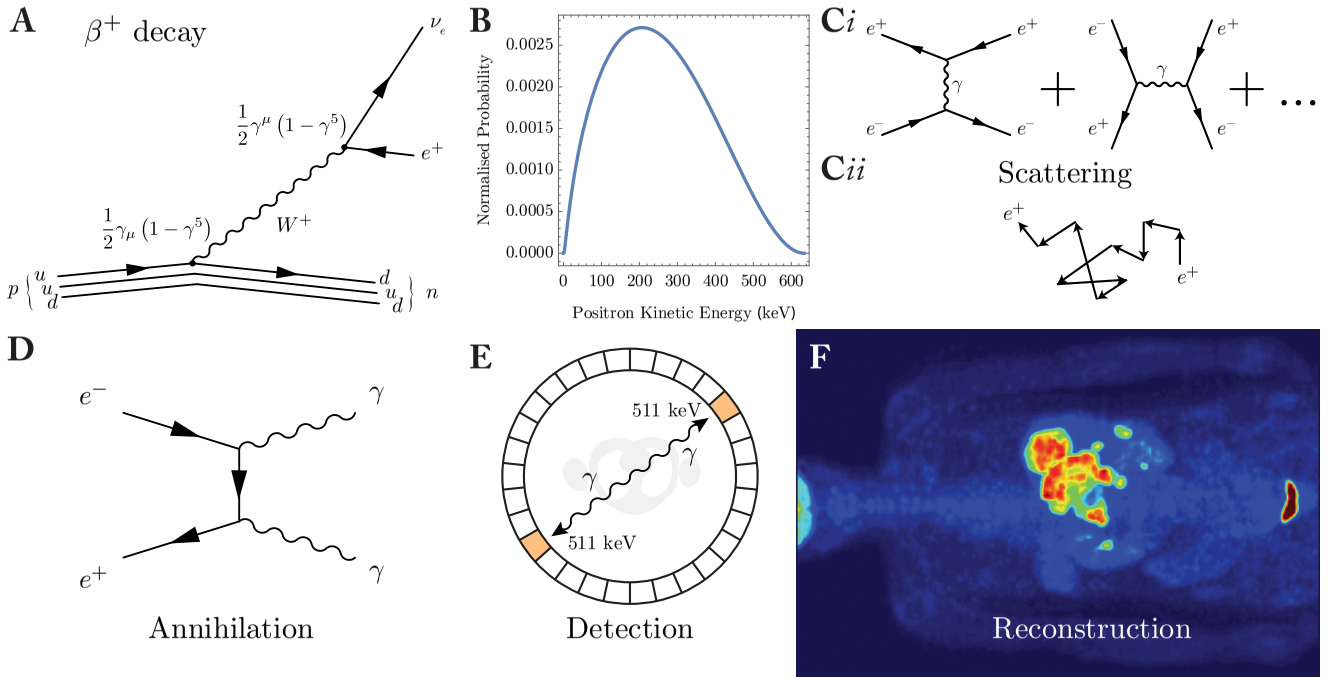

Figure 4.3: Positron Emission Tomography. **A**: A radioactive nucleus, commonly ^{18}F, spontaneously undergoes β^+ decay. Shown here is the process of β^+ decay, including the Feynman propagators for the weak force. The presence of two particles as decay products, combined with the vector-axial form of the weak interaction, means that the energy of the positron is not uniquely determined, and a distribution is observed whose shape is shown in **B**. It is well known that this distribution can be calculated by Fermi's Golden Rule, which states that the transition amplitude is proportional to a matrix element multiplied by the density of (final) states.[105] From the Feynman rules for electroweak interactions, the relevant matrix element for this decay is proportional to $\int \left[\psi_n^\dagger(\mathbf{x})\gamma_\mu \frac{1}{2}\left(1-\gamma^5\right)\psi_p(\mathbf{x})\right]\left[\psi_\nu^\dagger(\mathbf{x})\gamma^\mu \frac{1}{2}\left(1-\gamma^5\right)\psi_e(\mathbf{x})\right]\,\mathrm{d}^4\mathbf{x}$, where $\psi_i(\mathbf{x})$ is the wavefunction for particle i, and the γ^i are the Dirac matrices. The density of states for this situation can be shown to be $p^2(\sqrt{S}-E_{ke})^2$, where p is the positron's momentum, $\sqrt{S}$ the total energy available for the decay, and E_{ke} the kinetic energy of the positron. This density is further multiplied by a Fermi function to account for Fermi-Dirac statistics and the Coulomb repulsion of the positron from the nucleus. It is then possible to obtain the energy spectrum as it is simply proportional to the mod-square of the transition amplitude when plotted as a function of energy. The energy spectrum of the emitted positron for ^{18}F as calculated by the author is plotted in **B**. Once emitted, the positron then interacts electrostatically with matter, leading order processes of which are illustrated in **C***i*. The net effect of these interactions is that the positron undergoes a random walk, losing energy by ionising the surrounding matter (**C***ii*). Finally, when the positron's kinetic energy is sufficiently low, the amplitude for e^+e^- annihilation becomes large. Annihilation (**D**) produces two back-to-back photons, which may be 'residually' Lorentz boosted, meaning that in practice they are emitted within $\sim 180 \pm 0.6°$ to each other.[106] These annihilation photons are detected (**E**), and reconstruction methods such as filtered backprojection used to compute the emitter density within the patient, as shown in **F** (public domain image courtesy of Jens Langer, Technische Universität Dresden).

chemotherapy.[110] This decrease was thought to be due to the cytotoxic chemotherapy lysing cells and decreasing the concentration of NAD(H) required to catalyse the reaction. As this chemotherapy was first administered 24 hours prior to the second scan, the technique provides the ability to monitor the response to therapy on a more rapid timescale than other techniques.[110] Separate studies have shown that ^{18}FDG-PET correlates with the apparent label exchange of [1-^{13}C]pyruvate with [1-^{13}C]lactate, and that DNP compares favourably with ^{18}FDG.[111] As well as the excessive production of lactate due to the Warburg effect, alanine production following [1-^{13}C]pyruvate injection is of interest in an oncology setting, as alanine usage reports on biosynthetic processes within the cell,[112] and this 'extra information' is likely to be of utility in future imaging strategies with pyruvate. For these reasons, the first human phase-I trial of hyperpolarised pyruvate occurred in patients with biopsy-proven prostate cancer, where the technique was well tolerated, able to confirm the previous diagnosis and, in one case, detect a cancerous lesion that was not detected by conventional imaging strategies (i.e. contrast changes on T_2 weighted images).[102]

Tumours in organs with ordinarily high glucose turnover – such as the brain – are not visible to ^{18}FDG-PET.[113] Hyperpolarised [1-^{13}C]pyruvate, therefore, forms an attractive alternative to ^{18}FDG in the brain.[108] Many cancer patients develop secondary metastases, and those located in the brain are only detectable when sufficiently large as to be difficult to treat. This difficulty arises because neurons in the brain exist in an isolated, controlled environment that ordinarily serves to shield them from opportunistic infection, but in oncology this makes the detection and management of metastatic tumours exceptionally challenging. In the brain, endothelial cells lining the walls of the micro-vasculature are physically adjoined to each other by specialised 'tight-junction' proteins, and form an effective barrier between the brain and the blood. This cellular barrier does not allow large macromolecules or microscopic objects into the brain, and limits the diffusion of smaller molecules across it.[114, in review.] The majority of required substances in the brain – such as glucose and other metabolites – are actively transported across this blood-brain barrier (BBB). Some pathogenic organisms can physically force their way through these tight junctions, a process referred to as extravasation. As illustrated in Fig. 4.4, cancer cells circulating in the blood stream can extravasate across the BBB, growing between the endothelium and the smooth muscle cells.[115] The BBB then acts to protect the metastasis, shielding it both from chemotherapy and the immune system whilst letting it consume glucose, which is, in part, responsible for the exceptionally poor prognosis of patients with cancer metastasised to the brain. Only when metastases grow sufficiently large to deform neighbouring anatomy, and start secreting angiogenic factors are they detectable by MRI or other techniques. This effect occurs because vessel growth stimulated by the secretion of these angiogenic factors is disorganised, and the vessels grown lack functional tight junctions between endothelial cells, resulting in the localised breakdown of the BBB. The current clinical method of detecting metastases is to administer an MRI contrast agent, chelated

gadolinium. As the agent can only leave the vasculature in this angiogenic region of disease, or where the structure of the BBB is compromised, hyperintensities on a T_1-weighted image indicate a breakdown of the BBB at sites of large metastases. Unfortunately, by this point the brain is likely to contain a large number of other, undetected metastases. Additionally, these large metastases are aggressive, and difficult to treat with systemic therapy by the heterogeneously permeable nature of the BBB. Prognosis is accordingly very poor. A number of studies have been proposed to transiently circumvent the BBB for the delivery of chemotherapeutics to brain tumours. Connell *et al.* have demonstrated that the pleotropic, pro-inflammatory cytokine Tumour Necrosis Factor α (TNF-α) is able to temporarily disrupt the tight junctions between endothelial cells when administered intravenously, and that disruption enhances the uptake of clinical relevant chemotherapies and gadolinium.[116] Such methods hold the promise of being translatable agents that act to improve the delivery of chemotherapy and imaging contrast agents to micro metastases in the brain at a time when treatment may be more effective.

There is a small, but growing body of work where hyperpolarisation has been used to probe neurovascular metabolism. Both [1-^{13}C]pyruvate and [2-^{13}C]pyruvate[118,119] have been shown to cross the BBB sufficiently rapidly to allow spectroscopic information to be acquired, but with a relatively low signal-to-noise ratio compared to other tissues. The more lipophilic pyruvate analogue, ethyl-[1-^{13}C]pyruvate is taken across the BBB much more rapidly, but acts as a mild anaesthetic whose biological role and subsequent metabolism is not fully understood.[120] Hyperpolarised succinate is reported to not be transported sufficiently rapidly to be of utility beyond that of perfusion imaging, with no metabolic products being spectroscopically observed following injection.[121] Studies that have used [1-^{13}C]pyruvate in neuro-oncology report a significant increase in the level of lactate production in large tumours, commonly glioblastoma,[122–124] but the coarse spatial resolution of the sequences used (e.g. $5 \times 5 \times 10$ mm^3)[118, Fig. 3] limits their applicability to situations where the tumour is already clearly visible on anatomical images, and usually has a volume comparable to that of nearby healthy tissue.

4.3.3 Pilot Study

4.3.3.1 Overview

As the feasibility of using DNP to detect cancer that had metastasised to the brain was uncertain, a pilot study was performed (in 2012) to determine whether the uptake of pyruvate across the intact blood-brain barrier was sufficient to allow detection of metabolism by hyperpolarised [1-^{13}C]pyruvate spectroscopy. This study also formed an attempt to differentiate between healthy and diseased regions of the brain through differential lactate production. As different metabolic changes are associated with different cancer phenotypes (and ultimately prognosis),[125] understanding metabolism is still of clinical value even when the metastases present are too large to be effectively treated, as changes in metabolism are of value for monitoring the response to chemotherapy. This study used eight male BD-IX rats with

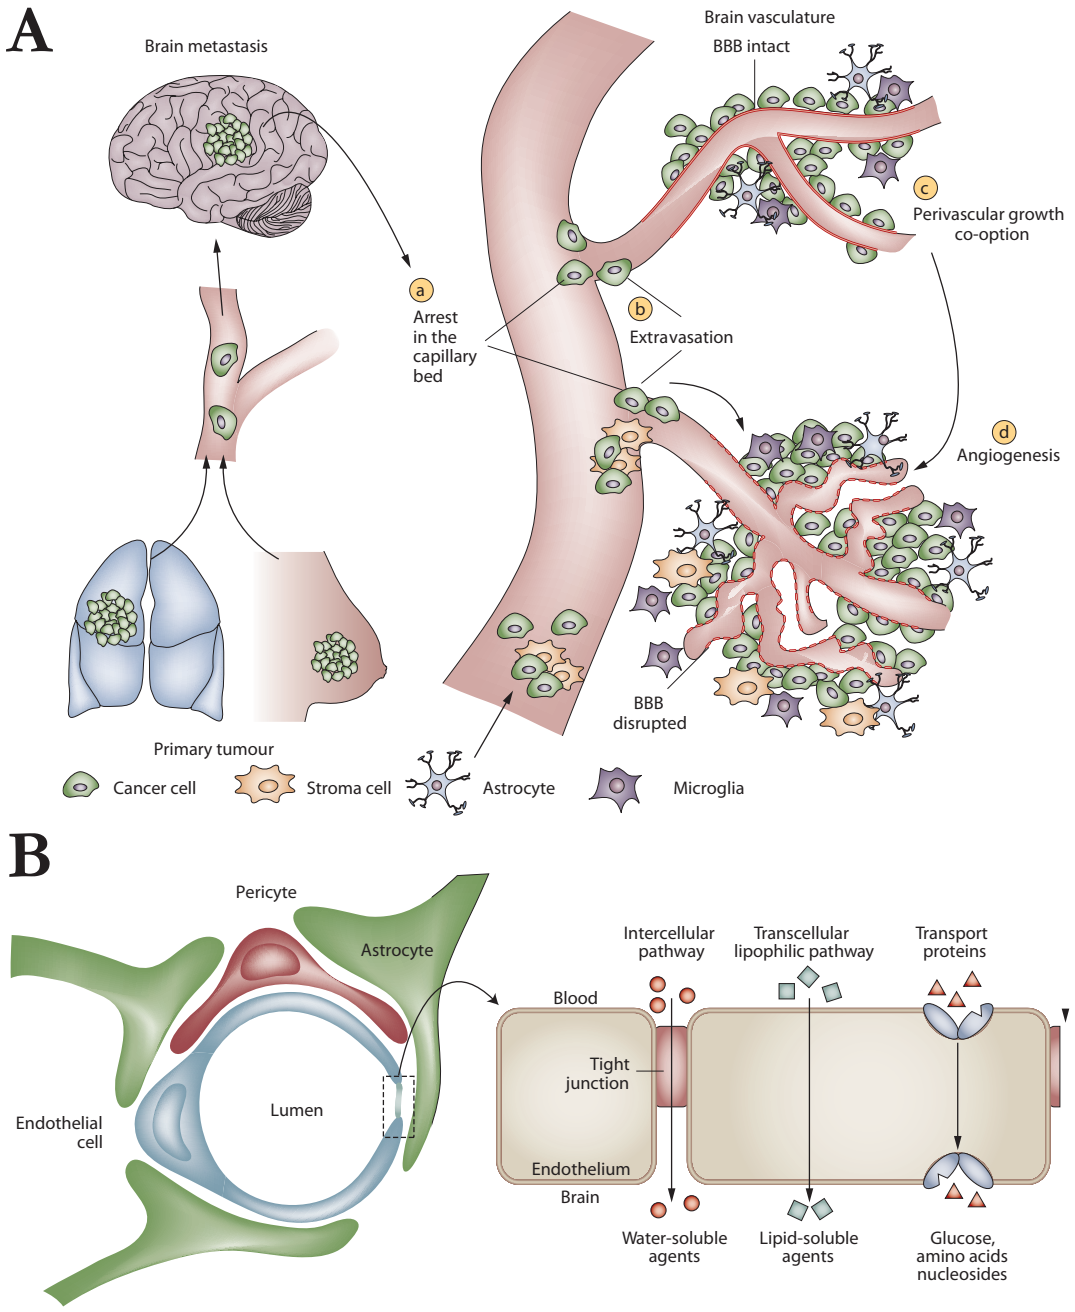

Figure 4.4: **A**: The progression of brain metastases. Tumour cells circulating from a primary tumour can adhere to the endothelium within capillaries (a), force their way into the perivascular niche (b), and then proliferate (c). Only when tumours start becoming sufficiently large as to be angiogenic (d) and 'leaky' will contrast agents, such as gadolinium, reveal their presence in T_1-weighted imaging. **B**: The structure of a capillary and the blood-brain-barrier. The blood flowing through the body of the capillary (lumen) is separated from the neural environment therein by tight junctions, proteinaceous junctions that only permit the diffusion of certain small molecules. Everything else required by the brain – including glucose and pyruvate – is actively transported. Images adapted from Eichler *et al.* [117]

small tumours induced in one hemisphere of the brain and attempted to spectroscopically differentiate between the two hemispheres on a metabolic basis, with the unaffected hemisphere providing a control. TNF$-\alpha$ was administered to a subset of the animals in an attempt to increase the uptake of pyruvate across the BBB. After injection of hyperpolarised [1-^{13}C]pyruvate, slice-selective spectra were acquired interleaved from each half of the brain, and differences in the apparent rate of lactate production were examined.

4.3.4 Methods

4.3.4.1 Animals

Male Berlin-Druckrey IX (BD-IX) rats ($n = 8$) were stereotactically intrastriatially injected in the left striatum by Sébastien Serres (+1.00 mm anterior, +3.50 mm lateral of Bregma, -3.50 mm depth) under 3% isoflurane anaesthesia with $10\,000$ ENU1564 cells, metastatic tumour cells originating from a primary mammary carcinoma in the female BD-IX rat induced by the organic carcinogen N-ethyl-N-nitrosourea (ENU).[126] Animals were housed in accordance with institutional guidelines, allowed food and water *ad libitum* and imaged one, eight, fifteen and twenty-two days post injection. Anaesthesia was induced by 2% isoflurane in 90% O_2, 10% N_2O, and ECG electrodes subcutaneously implanted for subsequent monitoring. Anaesthetised animals were imaged lying supine on the 2 cm 'mouse' loop surface ^{13}CRF coil within the 7 T preclinical MRI system described previously in Chapter 3. Separate spectra were obtained from each hemisphere of the animal's brain: the hemisphere contralateral to the injection site was used as a natural control. Lyophilised recombinant rat TNFα (Brain Laboratories Inc., IL, USA) was reconstituted in 250 μL sterile saline and administered as an IV bolus to each animal. [1-^{13}C]pyruvate spectra were obtained prior to intravenous TNF injection, and two hours post injection. TNF was not administered at the first time-point as it was believed that any tumour present would be exceptionally small, and additionally that the immune inflammatory response to surgery would conflict with the administration of the pro-inflammatory cytokine TNF.

4.3.4.2 MR protocol

After the acquisition of initial Gradient Echo scout images to check the co-localisation of the ^{13}C coil and the brain, slice-selective spectra were obtained every second within two slices covering each hemisphere of the brain. Slices were separated by a 1 mm gap, in an attempt to avoid the middle cerebral artery and to prevent spectral cross-contamination between the slices. Aside from the gap, the profile of the signal obtained was dominated by the field profile of the coil. One hundred and twenty spectra were obtained (60 per hemisphere, TR $= 1$ s, excitation flip angle 5° with a five-lobe sinc excitation pulse, 10 mm thick, sweep width 6 kHz, centre frequency approximately that of [1-^{13}C]pyruvate). After hyperpolarised spectroscopy, animals were imaged with axial 2D gradient and spin echo sequences (11 slices, 1.5 mm

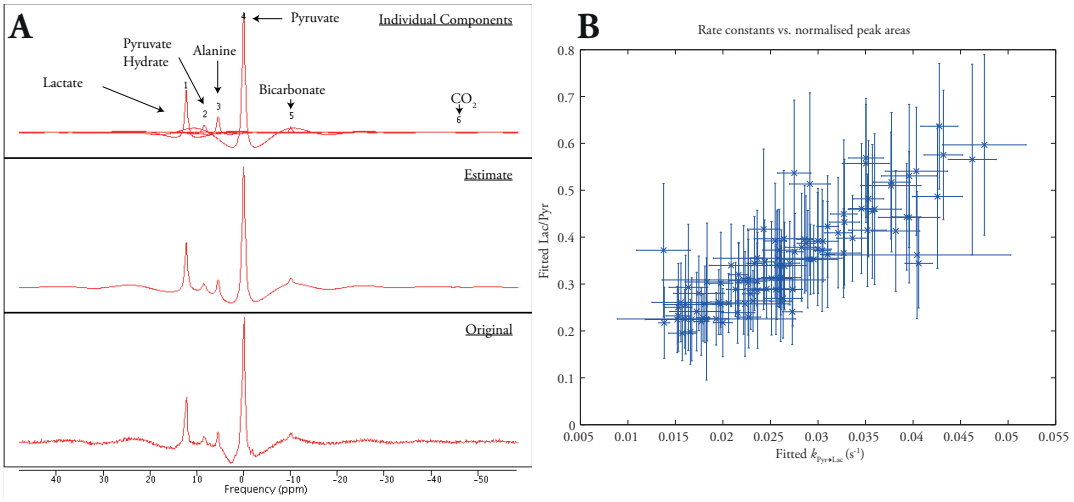

Figure 4.5: **A**: The expected six peaks were fitted to the acquired data. The AMARES algorithm is fits expected curves in the FID's time domain, using prior knowledge of the location of the metabolite spectrum. The resulting fit has a nearly flat residual, aside from deviations caused by local field inhomogeneities. **B**: It is then possible to fit the acquired spectra to a kinetic model of metabolism, the returned rate constants of which strongly correlate with metabolite-to-pyruvate ratios, shown here for $k_{\text{Pyruvate}} \rightarrow k_{\text{Lactate}}$.

thick, centred on the injection site; gradient echo: TR 120 ms, TE 4.68 ms, Matrix 192×192, 90° flip angle; spin echo: TR 500 ms, TE 20 ms, 4 averages, Matrix 128×128). Gadolinium (1.35 mmol/kg Gd-DPTA Chelate as DOTAREM®; Guerbet Laboratories, Villepinte, France) was then administered intravenously, and the same gradient echo and spin echo scans were repeated five minutes later. TNF (60 μg/kg) was delivered intravenously, and each animal allowed to recover for two hours before being reassessed using the above protocol.

4.3.4.3 *Analysis*

The acquired ^{13}C spectra were analysed and quantified using the AMARES peak fitting algorithm in the jMRUI spectral analyser,[127] which was successfully ported to several operating systems by the author (further details in Section 9.1.0.2). As illustrated in Fig. 4.5**A**, the expected NMR spectrum of downstream metabolites (pyruvate, lactate, alanine, bicarbonate and carbon dioxide) was fitted to the time-series spectral data obtained, and the height of each signal component recorded at each time-point. The algorithm returns the overall magnitude of each contribution to the signal, together with an estimate, the Cramér-Rao bound on the variance, of the error on each peak. As the acquired spectra varied in their signal-to-noise ratio as a consequence of the varying polarisation achieved for each experiment, this bound cannot be neglected. The time-integrated spectrum for each hemisphere was also quantified, which possesses greater SNR than the individual time-point spectra. As different experiments result in different degrees of polarisation, we normalised these time-integrated spectra for each experiment by simply dividing each peak in the time-integrated spectra by the magnitude of the corresponding pyruvate peak, obtaining a metabolite-to-pyruvate ratio.

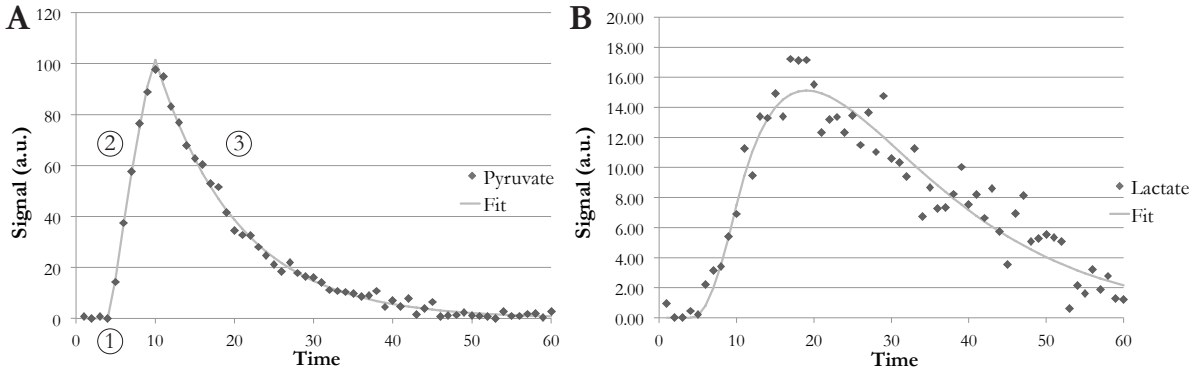

Figure 4.6: Fits to the turnover of metabolites in a representative spectrum as a function of time: shown are both the delivery of pyruvate (**A**) and the arrival of a lactate signal (**B**) following its metabolism. The piecewise continuous model describes three processes: the arrival of pyruvate into the voxel at a time t_a (**A**1), the continued increase in signal as the bolus continues to arrive (**A**2), and the subsequent decay of the signal as the hyperpolarised pyruvate is both metabolised and returns to thermal equilibrium (**A**3). Further information follows in Chapter 9.

Time series data were further quantified by fitting the peak amplitudes as a function of time to a simple kinetic model that described the biochemical interconversion of pyruvate to secondary metabolites. The model takes into account the arrival of the pyruvate signal, its decay, and the subsequent arrival and decay of all other metabolites, and is discussed fully in Section 9.4; particularly Eqs. (9.29) and (9.30). As illustrated in Fig. 4.6, the model provides an appropriate description of the fundamental processes undertaken by the hyperpolarised metabolite, and yields metabolic rate constants of interest, such as $k_{\mathrm{Pyr}\rightarrow\mathrm{Lac}}$. This method is highly positively correlated with the metabolite-to-pyruvate ratio ($R^2 > 0.99$, correlation plot shown in Fig. 4.5**B**), but has the advantage that it can statistically discriminate between noise and signal, and provides bounded estimates for the rate constants fitted. Correspondingly, the uncertainty resulting from performing this fitting is less than that obtained by taking the ratio of the two (uncertain) metabolite amplitudes. A detailed discussion on the propagation of uncertainty throughout the process of spectral fitting follows in Chapter 9.

The difference between pre- and post-contrast images (illustrated in Fig. 4.7) was used to provide a surrogate for the volume of the tumour present. After subtraction, the region of enhancement was manually segmented in ImageJ.

4.3.4.4 Immunohistochemistry

The expression of TNF-α type-I and -II receptors was confirmed by immunohistochemistry. Sections were stained with rabbit anti-TNF-α-I, or rabbit anti-TNF-α-II as primary antibodies using a standard protocol, and then counterstained with Cresyl Violet ((9-dimethylamino-10-methyl-benzo[a]phenoxazin-5-ylidene)ammonium chloride; a nuclei stain).[128] Animals were euthanised by an intraperitoneal injection of pentobarbital (500 mg), and a left thoracotomy performed. The right atrium was transected, and the left ventricle perfused initially with heparinised saline to exsanguinate the animal. A fixative

perfusate was then injected, consisting of paraformaldehyde (3.33 M), gluteraldehyde (5.2 mM), lysine HCl (200 mM), sodium periodate (10 mM), and di-sodium hydrogen orthosphosphate dihydrate (100 mM). Brains were excised, soaked in the fixative perfusate overnight, and then soaked in 30% sucrose for several days until dehydrated for cryoprotection. After being frozen at -40° in isopentane and embedded in Tissue-Tek, 10 μm coronal sections were cut within approximately 5 mm of the tumour. Slides were then soaked for 30 s in acidified water (0.1% ethanoic acid in water) followed by approximately one minute in Cresyl violet, before being dehydrated through a gradient of ethanol and water (75%, 90% and 100% ethanol for 5 minutes each). Slides were then washed in xylene, and coverslips were attached with DPX mountant (Sigma-Aldritch). For antibody staining, slides were washed in PBS for five minutes before being immersed in 250 methanol + 2.5 ml H_2O_2 for 20 minutes to quench exogenous activity. Slides were then washed in PBS again (10 mins, two washes) and blocked with 10% horse serum for an hour. The primary antibody was incubated overnight, before slides were washed in PBS, and incubated with the (biotinylated) secondary antibody for 45 minutes. Sections were subsequently washed and stained with a commercial 'ABC reaction' kit (Vector Laboratories, Peterborough, UK), before being incubated with the organic reagent DAB for 30 seconds.

4.3.4.5 In vitro *metabolic quantification*

To confirm the *in vivo* measurements of relative lactate and pyruvate levels via DNP, the ENU cell line was cultured and assayed for its pyruvate and lactate content. A highly glycolytic human breast carcinoma cell line – the MDA cell line – served as a positive control. Cells were grown in 90% DMEM with 10% fœtal bovine serum and a penicillin/streptomycin antibiotic admixture. Lactate was quantified spectrophotometrically on a Pentra 'Clinical Chemistry' analyser;[129] pyruvate quantification proceeded by spectrophotometrically measuring the reduction of NADH by lactate dehydrogenase.[130, p. 2127] The extracellular lactate concentration was measured as a function of time for both cell lines, in an attempt to quantify the rate of lactate efflux.

4.3.5 Results

Tumour presence and development was routinely assessed by the subtraction of images taken before and after the injection of gadolinium; if BBB breakdown had occured, gadolinium accumulation would result in a region of hyperintensity present in the T_1-weighted images. This breakdown was present after two ($n = 3$) or three ($n = 5$) weeks; Fig. 4.7 shows a representative image pre- and post-gadolinium together with the difference between the two images. The approximately exponential growth of tumour volume against time is shown in Fig. 4.9.

As shown in Fig. 4.8, there were no significant differences between the apparent rate of lactate production between the control and injected hemispheres, and TNF administration did not have any

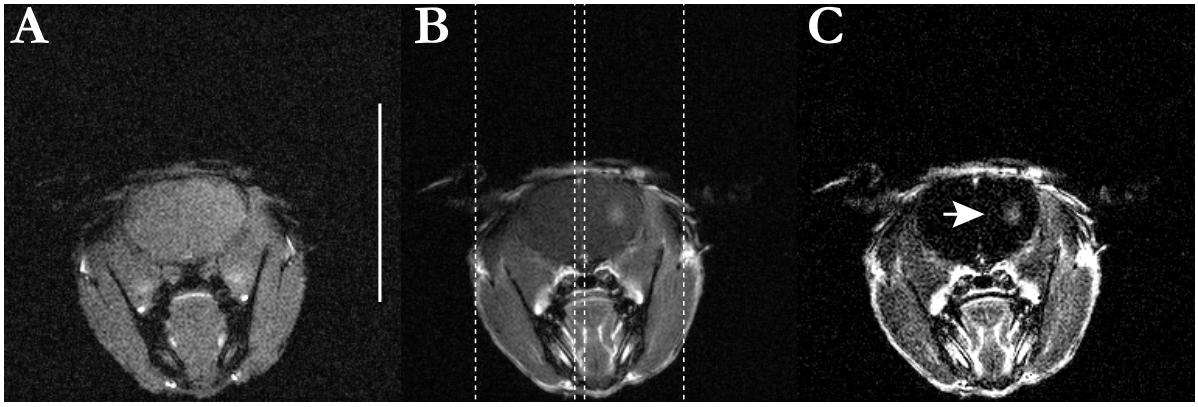

Figure 4.7: Axial GEMS slices showing the brain before (**A**) the injection of a gadolinium contrast agent; after the injection of gadolinium (**B**) an increased MR signal is obtained from gadolinium 'leaking' through the BBB. **C**: The difference between the two images clearly shows the tumour within the striatum of the left hemisphere (arrow). This animal had the largest tumour and the earliest significant breakdown of the BBB, visible two weeks post injection. Also shown in **B** is the location of the voxels chosen for the acquisition of slice-selective spectra. Scale bar 20 mm.

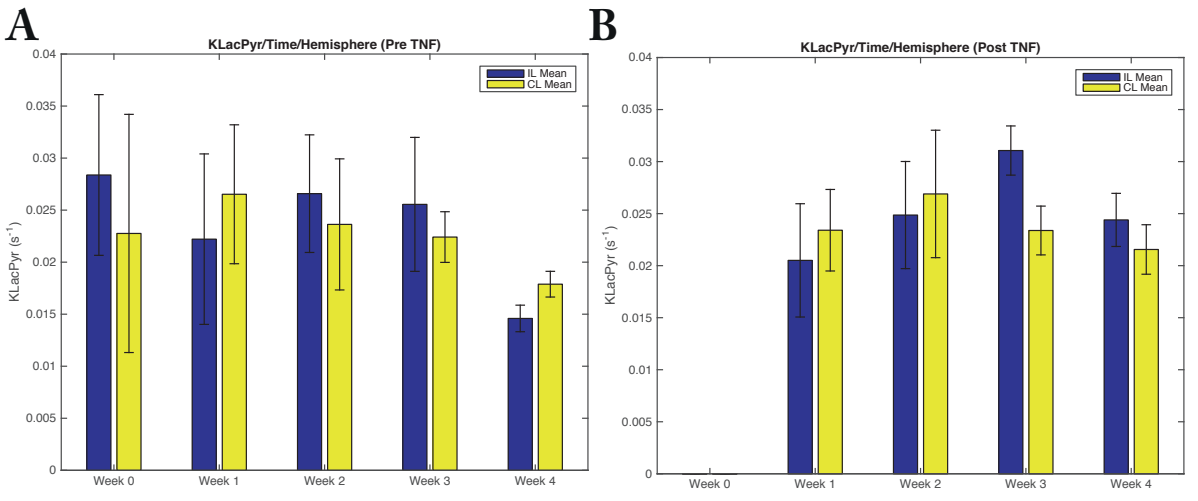

Figure 4.8: Average rates of lactate label incorporation shown at each timepoint, both without (**A**) and with (**B**) prior administration of TNF. Data are shown as weighted mean $\pm$ weighted standard deviation for both the affected (ipsilateral, IL) and control (contralateral, CL) hemispheres. There are no significant differences between groups ($P > 0.05$ between all groups by either Student's t-test with Bonferroni's correction for multiple comparisons, or the methods outlined in Section 9.2.2.2).

visible effect. The spectra were of varying signal-to-noise ratio, but there was no significant trends for changes in the delivered visible pyruvate (i.e. the raw peak integral) over time. There was no correlation with visible tumour volume and lactate production, as measured by the fitted rate constant (Fig. 4.10).

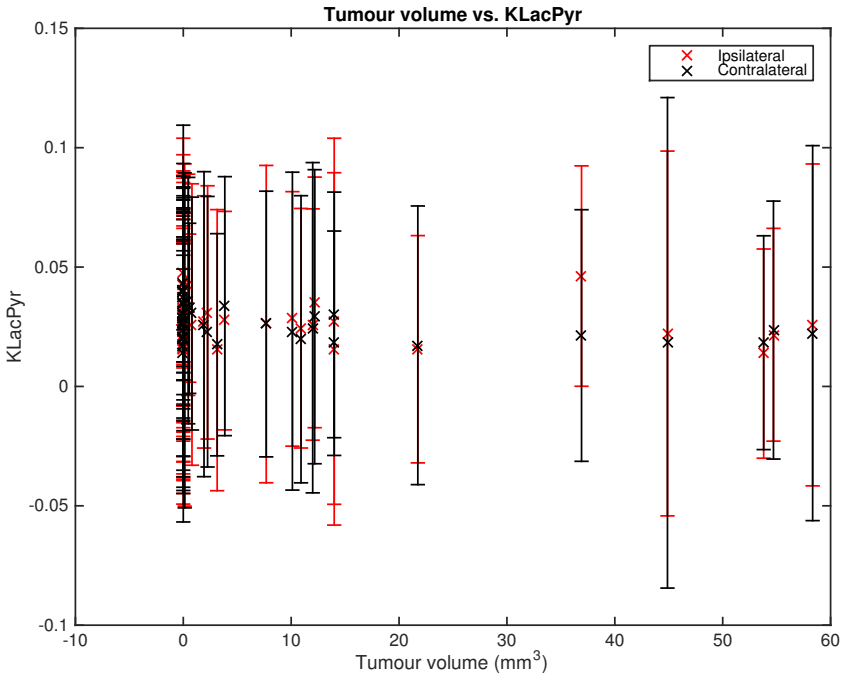

Figure 4.10: Measured apparent rate constant as a function of tumour volume, for both the affected (ipsilateral) and control (contralateral) hemispheres. There is no correlation with apparent lactate production and tumour volume.

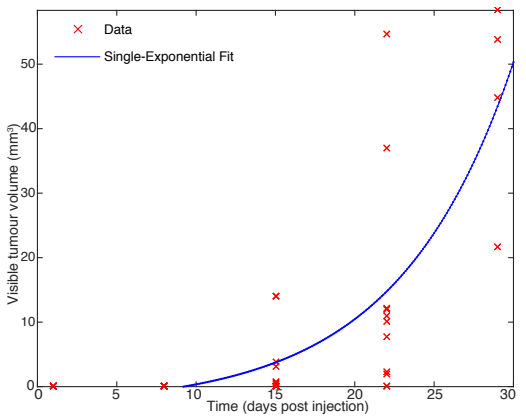

Figure 4.9: Tumour volumes as a function of time; tumours were only visible at either 15 or 22 days post injection, as determined by gadolinium enhanced MRI. A single exponential fit is also shown in blue.

The effect of TNF administration was unclear. TNF did not consistently change detected metabolite levels at any point (c.f. Fig. 4.8), nor did TNF result in a statistically significant increase in the measured tumour volumes as measured by the gadolinium signal visible to MRI. This finding contrasts to previous reports on the increased uptake of both therapeutic antibodies and gadolinium by brain metastases, when preceded by the administration of TNF.[116]

Histology confirmed the presence of tumours in each animal. As illustrated in Fig. 4.11, the tumour maintained a metastatic morphology aside from a large necrotic central region. Immunohistochemistry showed that TNFα receptor type-I (TNFαR-I) was mostly expressed on glial cells, tumour-associated macrophages, and neutrophils. It was also seen in the neurons around the tumour. TNFα receptor type-II (TNFαR-II) was mostly expressed on neurons,

microglia and leukocytes, which can extravasate across the blood-brain barrier. Some macrophages seemed to become trapped in the endothelium, which might imply that the endothelium expresses TNFαR-II. It appeared to be the case that TNFαR-II and extravasated tumour cells were co-localised (c.f. Fig. 4.11**C**), although it may be necessary to use the higher achievable in-plane resolution of confocal immunofluorescence microscopy to confirm this hypothesis. It was not possible to quantify the density of the receptors within the tumour in comparison to the physiologically healthy brain.

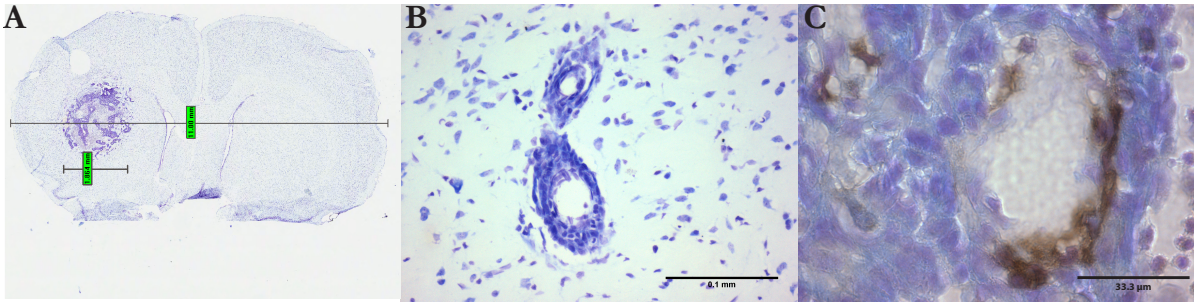

Figure 4.11: Histological confirmation of the metastatic nature of the tumour. **A**: Cresyl violet stained brain, illustrating metastatic foci in the left hemisphere. Scales shown are 11.00 mm (the approximate width of the brain), and 1.864 mm (the width of the metastatic region). **B**: Despite the centre of the tumour being necrotic, the tumour exhibited growth and extravasation along capillaries as observed in human brain metastases. Here we see two colonies growing in the perivascular niche, surrounding each capillary. (Cresyl violet; scale 0.1 mm) **C**: Immunohistochemistry revealed the apparent co-localisation of TNFαR-II and the metastatic colony. Staining clearly reveals expression of the receptor on the inside of the capillary, above the level in the unaffected part of the brain. (Cresyl violet and anti-TNFαR-II staining; scale bar 33.3 μm).

4.3.6 Discussion & Conclusions

This study reports the novel *in vivo* observation of metabolism in a model of brain metastases. It is clear that the technique was not able to discriminate between healthy and diseased tissue on the basis of metabolic phenotype. However, it appears that the action of monocarboxylic acid transporters provided sufficient transport of hyperpolarised pyruvate across the BBB for the observation of all downstream metabolic products. The fact that there was no correlation between observed tumour volume and apparent lactate production is surprising, as it contradicts the known glycolytic nature of this cell line *in vitro*, and the well-known metabolic phenotype of cancerous cells in general.

The response to TNF injection was uncertain. No significant changes were observed in the metabolism of hyperpolarised [1-^{13}C]pyruvate following its injection. Based on the previous work of others, it was expected that TNF may be able to systemically permeabilise the BBB, and it was hoped that the up-regulated expression of the receptor in the region of disease might lead to selective permeabilisation near the tumour. If permeabilisation had occurred, one may expect to see an increase in the apparent rate of lactate production after the injection of TNF compared to that without it. However, it was not possible to readily quantify the histological data obtained robustly to confirm this hypothesis. Therefore, it is not known whether (*i*) TNF has no biologically relevant effect in this experimental model

(a possibility, as earlier work was performed on mice) (*ii*) the administered dose was insufficient to induce BBB permeabilisation, or (*iii*) TNF did have an effect on the uptake of pyruvate, but the inability of the spectroscopic experiment to distinguish between diseased and healthy brain meant that any such effects were not distinguishable.

Accordingly, it was decided to investigate the origin of this lack of asymmetry in apparent lactate production between cancerous and healthy brain. At this point, in 2012, it was decided to measure the field map of the RF coil used for the study, the $^{13}\mathrm{C}$ mouse surface coil. This coil is a loop approximately 2 cm in radius, and it was therefore assumed previously that, as can be shown from Biot-Savart's law for a single axial loop of radius R, the B_1 field would extend on-axis as

$$\begin{aligned} B_{\text{Single loop}} &= \frac{\mu_0}{4\pi}\frac{2\pi R^2 I}{(z^2+R^2)^{\frac{3}{2}}} \\ &\approx \frac{\mu_0}{2}\left(\frac{I}{R} - \mathcal{O}\left(\frac{z^2}{R^3}\right) + \dots\right) \end{aligned} \tag{4.1}$$

and therefore be approximately constant for a small distance $\propto R$, i.e. would be able to approximately probe metabolism up to a depth in tissue on the order of 2 cm. After measuring the field map of the coil, it was found that this assumption was decidedly incorrect, and that in reality the signal acquired was spatially localised to the superficial region of the brain (shown in Fig. 4.12**A**). This is likely due to the difference between an idealised loop of wire that has zero radius and is small in comparison to the infinite wavelength of the DC current flowing through it, and the physical coil used in the study with a $^{13}\mathrm{C}$ Larmor frequency of 75.5 MHz.

To predict the effects of this coil, a simulation was performed in which anatomical images were segmented, the tumour voxels were assigned an increase in 'apparent lactate signal' compared to non-tumour voxels, and the co-registered fieldmap multiplied on the resulting 'estimated data'. Tumours were dissected *ex vivo*, and the lactate concentration measured in the body of the tumour, the surrounding tissue, and the contralateral hemisphere in order to determine the relative increase in each. These measurements were obtained by automated enzyme-linked spectrophotometric assays, performed on an 'AB Pentra 400' clinical biochemistry analyser (HORIBA Medical ABX SAS, Montpellier, France). The amount of lactate in the thawed tumour lysate was approximately 2.5$\times$ that of healthy tissue. A model was then generated using the measured field map of the RF coil used for the study (Fig. 4.12**A**), whereby a prediction for the total 'signal' received for each hemisphere was stochastically estimated. Reference images of the size of the tumours were segmented, and each voxel of healthy tissue was assumed to contribute a normally distributed amount of 'signal' with $\mu = 1$, $\sigma^2 = 1/4$. This signal was then multiplied by the normalised sensitivity of the RF coil at that point. Tumour voxels were assigned a distribution of mean 2.5 and unit variance, producing the 'image' shown in Fig. 4.12**B**. After spatial integration over the two slices (mirroring the spectroscopy experiment), it was found that the

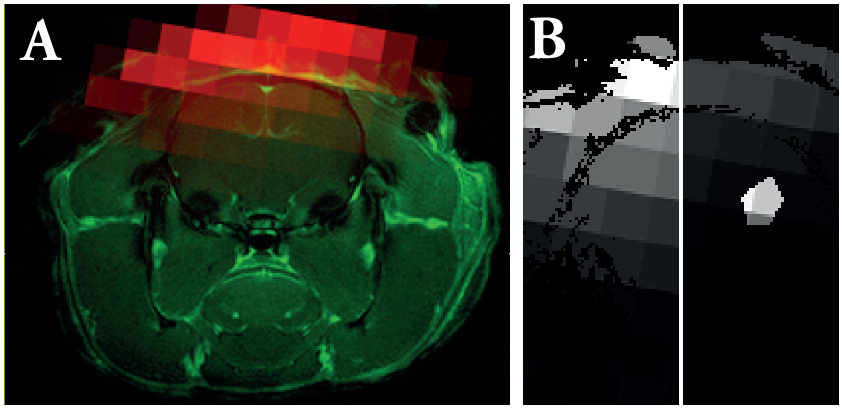

Figure 4.12: **A**: Measured field map of the ^{13}C mouse loop coil overlaid on the proton rat image. The spatial penetration of the coil used was insufficient to adequately probe the tumour, located in the striatum. **B**: An illustration of the simulation performed as described in the text; the brain was segmented into two hemispheres, and the tumour (bright region on the right) assigned a greater value of 'lactate' production than the surrounding tissue. When multiplied by the field map, the resulting image is as shown. This heralds promise for the utility of hyperpolarised imaging strategies.

average differences in signal between the two hemispheres were small, and unlikely to be resolved experimentally. However, it was realised that the difference within a voxel is large enough in magnitude to allow the tumour to be imaged by its metabolic phenotype. The measurements were dominated by noise, and it was considered highly likely that the glycolytic nature of the tumour was hidden amongst thermal noise, inter-experiment variability, and the contribution from non-tumour voxels. Nevertheless, this exercise clearly indicated that the highly glycolytic nature of tumour voxels would render them visible to any imaging strategy that was truly able to image glycolysis. As disease is spatially localised this method would serve as a useful diagnostic tool.

Fundamentally, the imaging of metabolism is likely to provide more useful information for quantifying disease and to be of clinical utility. This study indicated that pyruvate metabolism could be visible spectroscopically in the brain, but confirmed the seemingly obvious statement that imaging needs to be used in order to adequately discriminate between healthy and diseased tissue. As pyruvate metabolism is changed in many prevalent and currently poorly managed diseases, the spatial imaging of hyperpolarised metabolites is likely to be broadly applicable and of wide utility. For example, the ability to image lactate production following the injection of hyperpolarised [1-^{13}C]pyruvate could be applied in both an oncology setting, detecting the tumours discussed above, and also to detect regions of metabolic change in the heart following infarction.

For the reasons above, subsequent work was directed at the aim of creating a high-resolution metabolic imaging sequence that could track the metabolism of hyperpolarised [1-^{13}C]pyruvate *in vivo*. This development is described in subsequent chapters. Additionally, the significant role played by the RF coil used in determining sensitivity lead to the characterisation and subsequent development of appropriate ^{13}C RF coils that were suitable for hyperpolarised imaging, described in Section 3.2.

Chapter 5

Development of a Multiple-Slice Spectral-Spatial Planar 2D EPI Pulse Sequence (SpSp-EPI)

5.1 | *Introduction*

Radiofrequency pulses that sample the enhanced magnetisation generated by DNP necessarily destroy it. The second law of thermodynamics ensures that spins excited by any RF pulse (where $\theta \bmod 2\pi \neq 0$) will relax to thermal equilibrium rather than to their initial hyperpolarised state. As additionally, in general, $T_2 < 2T_1$ (and often $T_2 \ll T_1$)[131, §11.10] any potential rapid imaging readout for a hyperpolarised compound has an inherently limited temporal window in which to function. These facts serve to make hyperpolarised imaging particularly challenging. In addition to minimising the overall RF energy deposited, any hyperpolarised imaging scheme has to be temporally rapid, as the hyperpolarised compound of interest returns to thermal equilibrium. For the case of [1-^{13}C]pyruvate, and other compounds which are actively metabolised, *in vivo* imaging strategies have to be designed to cope with the chemical shift evolution of substrates over time. In a cardiac setting this challenge is even greater for two key reasons: (1) the heart moves, and this motion is difficult to retrospectively correct for as the observed signal is inherently a function of time, independent of cardiac motion; and (2) the paramagnetic nature of blood makes single-shot imaging techniques more challenging to apply owing to B_0 inhomogeneities, a decrease in T_2^*, and artefacts arising from flowing blood.

The simplest method to obtain both spectral and spatial information following the injection of [1-^{13}C]pyruvate is conventional Chemical Shift Imaging (CSI), in which either no frequency encoding is used (and each voxel is obtained entirely by phase encoding), or by the use of multiple echo times to reconstruct spectral information *post hoc*. CSI is flexible in the sense that no prior spectral information is required for efficient imaging, but the cost of this flexibility is increased acquisition time. Owing to the need for rapid imaging outlined above, this acquisition time constraint represents a significant challenge, and has lead to various techniques to accelerate acquisitions, such as compressed sensing[132] or parallel imaging.[133] IDEAL Spiral CSI represents a highly refined technique for the chemical shift imaging of hyperpolarised [1-^{13}C]pyruvate, in which a rapid spiral k-space trajectory is traversed with echo time shifting responsible for spectroscopic reconstruction.[134] IDEAL Spiral CSI effectively undersamples

the spectral domain, with echo times chosen *a priori* based on the expected location of resonances in the ^{13}C spectrum following the injection of [1-^{13}C]pyruvate. It is necessary to interleave the acquisition of a spectrum in order to guide the reconstruction, and for [1-^{13}C]pyruvate imaging this necessitates the acquisition of seven FIDs per image, further limiting temporal resolution. However, it remains infeasible to achieve whole-heart coverage within a 20 second cardiac-gated acquisition period using conventional 3D phase-encoded spectroscopic acquisitions, even with the introduction of rapid imaging readouts. Following injection of hyperpolarised [1-^{13}C]pyruvate IDEAL spiral CSI is able to obtain an image within seven ~ 150 ms TRs, which makes multi-slice coverage of the heart a realistic possibility. However, IDEAL is prone to several artefacts such as spectral-spatial ringing, whereby a large signal at one frequency can produce spurious signals in another, and requires at least seven broadband excitation pulses to be played per set of reconstructed images. Furthermore, as every metabolite is excited in this technique before sampling, an effective limit is placed on the flip angle that can be used with this technique as the depletion of pyruvate magnetisation would end the production of visible downstream metabolites, effectively ending the experiment. Additionally, the multiple echo times result in varying degrees of T_2^* decay before the readout. This, combined with the sensitivity of spiral readouts to off-resonance effects, makes the reconstruction of IDEAL CSI data a non-trivial problem.

This chapter describes the development and brief characterisation of an alternative method for spectroscopic imaging, with spectrally and spatially selective excitation provided by a 2D RF pulse, followed by a rapid-imaging EPI readout. This approach allows for the excitation of multiple slabs of magnetisation with the delivery of independent flip angles to different metabolites. Such strategies can be considered more magnetisation efficient, as the large 'reservoir' of hyperpolarised [1-^{13}C]pyruvate can be minimally sampled, whilst the continually produced labelled metabolites of interest, namely [1-^{13}C]lactate and ^{13}C-bicarbonate, can be probed with larger flip angles. This allows the acquisition of images with approximately equal SNR in all metabolites, despite the fact that the apparent visible concentration of labelled metabolites downstream of [1-^{13}C]pyruvate is far lower than that of [1-^{13}C]pyruvate itself.

5.2 | *Theory of Spectral-Spatial Pulses*

Spectral-Spatial (SpSp) RF pulses are, as implied by the name, selective in both chemical shift and space. Originally introduced for fat-water imaging in the 1990s by Meyer *et al.*,[135] they form an attractive alternative for hyperpolarised imaging as the spectral complexity possible *in vivo* following the injection of a hyperpolarised metabolite is inherently confined. All spectral-spatial pulses consist of multiple RF 'sub-pulses' that are played under a broad envelope, together with an oscillating bipolar slice-selection gradient.[7, p. 153] The effect of the pulse is best understood in the excitation k-space, which is derived (in Section 2.3.1) under the assumption that θ is small. This analytic approach is informative, and a brief

sketch of the action of such a pulse follows.

5.2.1 *Spectral-Spatial Excitation k-space*

Recall that the effect of an RF pulse is approximately its Fourier transform, i.e.

$$M_\perp = i\gamma M_0 \int_0^T B_1(t') e^{i\mathbf{k}(t')\cdot\mathbf{r}} \, \mathrm{d}t'. \tag{5.1}$$

where $\mathbf{k}(t) = -\gamma \int_0^t \mathbf{G}(s) \, \mathrm{d}s$ is the spatial frequency variable, the pulse runs from a time 0 to T, and all other symbols have their usual meaning. Spectral-spatial pulses are multidimensional, and exist in a $k_\mathbf{r}$, k_ω (spatial, spectral) k-space. For a 2D pulse it is most convenient to define k-space axes corresponding to z and ω as

$$k_z(t) = -R \int_t^T G_z(s) \, \mathrm{d}s, \qquad \text{and} \qquad k_\omega(t) = t - T \tag{5.2}$$

where the constant R depends on the design problem at hand; for 2D spatial-spatial excitation, R would be γ, and k_z would have units of reciprocal distance. Here, the SI units of k_z and k_ω are different, and we chose R such that k_z has the same units as k_ω, namely those of reciprocal frequency, i.e. seconds. The numerical choice of R is hence arbitrary, as it essentially represents the choice of axes in k_z, k_ω space.

Under this formalism, the effect of an RF pulse can be expressed as

$$M_\perp(\mathbf{r}) = i\gamma M_0(\mathbf{r}) \int_\mathbf{K} W(\mathbf{k}) S(\mathbf{k}) e^{i\mathbf{r}\cdot\mathbf{k}} \, \mathrm{d}\mathbf{k}, \tag{5.3}$$

$$\text{where } W(\mathbf{k}(t)) = \frac{B_1(t)}{|\gamma \mathbf{G}(t)|} \overset{\text{def}}{=} \frac{B_1(t)}{|\dot{\mathbf{k}}|}, \tag{5.4}$$

$$\text{and } S(\mathbf{k}) = \int_0^T \delta\big(\mathbf{k}(t) - \mathbf{k}\big) |\dot{\mathbf{k}}| \, \mathrm{d}t. \tag{5.5}$$

Here, the function $W(\mathbf{k})$ is a weighting function in multidimensional k-space, and $S(\mathbf{k})$ represents a sampling grid. The factor $|\dot{\mathbf{k}}|$ is essentially a normalising constant related to the velocity in which k-space is scanned. The effect of the pulse is therefore just the Fourier transform of the product $W(\mathbf{k})S(\mathbf{k})$.[16]

5.2.2 *Oscillating Gradients and $S(\mathbf{k})$*

As the spectral k trajectory is constrained to move linearly with time, the gradient structure has to be oscillated to adequately sample k_z, k_ω space. In general, any arbitrary waveform $G(t)$ could be used, but the effect of a spectral-spatial pulse can be analytically derived if $G_z(t)$ is taken to be explicitly sinusoidal, i.e.

$$G_z(t) = G \cos\big(\Omega(t - T)\big). \tag{5.6}$$

It is then straightforward to find the corresponding trajectory taken in k-space as

$$k_z = \frac{RG}{\Omega}\sin\big(\Omega(t-T)\big) \equiv \frac{RG}{\Omega}\sin(\Omega k_\omega), \text{ where } -T \le k_\omega \le 0. \tag{5.7}$$

Note the difference of sign compared to image k-space; as a consequence of the minus sign in Eq. (5.2) this oscillatory trajectory ends, rather than starts, at the origin. As $\mathbf{k}$ is defined, we proceed to calculate W and S. With this choice of G, $S(\mathbf{k})$ is given by

$$S(\mathbf{k}) = \sqrt{R^2G^2\cos^2\big(\Omega k_\omega\big)+1}\,\delta\Big(k_z - \frac{RG}{\Omega}\sin\Omega k_\omega\Big) \tag{5.8}$$

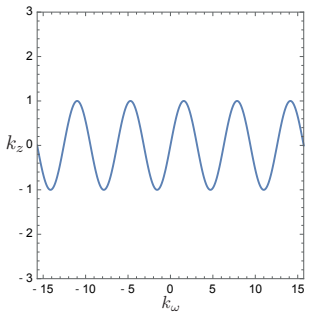

Figure 5.1: The function $S(\mathbf{k})$ defined in Eq. (5.8).

which takes the form of a sinusoidal 'line delta', as sketched in the contour plot Fig. 5.1 for the case $RG = \Omega = 1$, where white represents zero and blue is undefined. Analytically, the domain of $S(\mathbf{k})$ is the unbounded $k_z, k_\omega \in \mathbb{R}^2$ plane, although any physical realisation of such pulses will take place in a bounded subset of this space. As $S(\mathbf{k})$ is clearly periodic (with period $2\pi/\Omega$), it can be expressed as a Fourier series with k_z as a parameter, i.e.

$$S(k_z, k_\omega) = \sum_{n=-\infty}^{\infty} C_n(k_z)e^{-ink_\omega\Omega}, \tag{5.9}$$

and where each series coefficient $C_n(k_z)$ can be found by evaluating the integral

$$C_n(k_z) = \frac{\Omega}{2\pi}\int_{-\pi/\Omega}^{\pi/\Omega}\sqrt{R^2G^2\cos^2\big(\Omega k_\omega\big)+1}\,\delta\Big(k_z - \frac{RG}{\Omega}\sin\Omega k_\omega\Big)e^{in\Omega k_\omega}\,\mathrm{d}k_\omega. \tag{5.10}$$

By expressing the argument of the delta function in the form $k_\omega - f(k_z)$ this integral can be computed by the defining properties of the Dirac delta. Unfortunately, as $\arcsin$ is a multivalued function, the required rearrangement is not uniquely defined, and actually performing the integration Eq. (5.10) is somewhat involved. It can be shown[135] that the resulting coefficients have the form

$$C_n(k_z) = \begin{cases} 2\beta(k_z)\ \mathrm{rect}\left(\frac{\Omega k_z}{2RG}\right)\cos\left(n\arcsin\left(\frac{\Omega k_z}{RG}\right)\right), & \text{for } n \text{ even}, \\ 2i\beta(k_z)\ \mathrm{rect}\left(\frac{\Omega k_z}{2RG}\right)\sin\left(n\arccos\left(\frac{\Omega k_z}{RG}\right)\right), & \text{for } n \text{ odd}. \end{cases} \tag{5.11}$$

where $\mathrm{rect}(x)$ is the rectangle function defined as $\mathrm{rect}(|x| < 1/2) = 1$; $\mathrm{rect}(|x| \ge 1/2) = 0$, and $\beta(k_z)$ is a rational function of $\Omega k_z/RG$. Finally, the Fourier transform of $S(\mathbf{k})$ can be computed as

$$\mathcal{FT}^{-1}\{C_n(k_z)\} = K_1 \frac{J_1\left(\sqrt{1+\frac{1}{(RG)^2}}\frac{RGz}{\Omega}\right)}{2\sqrt{1+\frac{1}{(RG)^2}}\frac{RGz}{\Omega}} \otimes J_n\left(\frac{RGz}{\Omega}\right) \tag{5.12}$$

$$\stackrel{\text{def}}{=} K_1 \frac{J_1(Az')}{2Az'} \otimes J_n(z') \tag{5.13}$$

where $\otimes$ denotes convolution, $\mathcal{FT}^{-1}$ the inverse Fourier transformation, $J_n(z)$ are Bessel functions of the first kind of order n (sketched for the sake of completeness in Fig. 5.2), K_1 is an arbitrary constant, z is the distance from isocentre along the slice direction, and we have defined

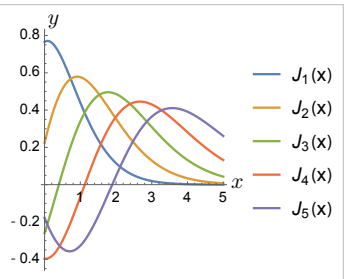

Figure 5.2: A sketch of the first five Bessel functions $J_n(x)$.

$$A = \sqrt{1+\frac{1}{(RG)^2}} \tag{5.14}$$

$$z' = \frac{RGz}{\Omega} \tag{5.15}$$

for the sake of notational convenience. Accordingly, given the series representation for $S(\mathbf{k})$ (Eq. (5.9)), it is clear that the inverse FT of $S(\mathbf{k})$ is given by

$$\begin{aligned}\mathcal{FT}^{-1}\{S(k_\omega, k_z)\} &= \sum_{n=-\infty}^{\infty} \mathcal{FT}^{-1}\{C_n(k_z)\}\delta(\omega - n\Omega) \\ &= K_1 \frac{J_1(Az')}{2Az'} \otimes \sum_{n=-\infty}^{\infty} J_n(z')\delta(\omega - n\Omega).\end{aligned} \tag{5.16}$$

In other words, the transform of $S(\mathbf{k})$, the sampling density function, is a weighted comb of Dirac deltas parallel to the k_z axis and separated by Ω, the gradient oscillation frequency. The weighting is proportional to a product of Bessel functions which arise from several integrals, convolved with the Dirac comb. Ultimately, this structure is responsible for the z, ω pattern of excitation that spectral-spatial pulses provide. Here, the form of $G(t)$ has been assumed to be sinusoidally varying, although in practice trapezoidal shapes are used. Such shapes necessitate the piecewise evaluation of Eq. (5.10) given a piecewise-defined trapezoidal $G(t)$. The effect is to change the form of the coefficients $C_n(k_z)$ above.

5.2.3 $W(\mathbf{k})$ and $B_1(t)$

If the form of $W(\mathbf{k})$ is taken to be separable, i.e. $W(k_z, k_\omega) \equiv W(k_z)W(k_\omega)$, then the final Fourier transform necessary to compute the excited magnetisation is separable, and can be expressed through the convolution theorem in terms of known quantities. With this further approximation, the form of the excited magnetisation is

$$M_\perp(z,\omega) = 2\pi i\gamma M_0(z,\omega)\,\mathcal{FT}^{-1}\{W(k_\omega)\}\delta(k_z) \otimes_{2\mathrm{D}} \sum_{n=-\infty}^{\infty} \mathcal{FT}^{-1}\{C_n(k_z)W(k_z)\}\delta(\omega - n\Omega)$$
$$\propto \mathcal{FT}^{-1}\{W(k_\omega)\}\delta(k_z) \otimes_{2\mathrm{D}} \mathcal{FT}^{-1}\{W(k_z)\} \otimes \frac{J_1(Az')}{2Az'} \otimes \sum_{n=-\infty}^{\infty} \Big(J_n(z')\delta(\omega - n\Omega)\Big), \quad (5.17)$$

where $\otimes_{2\mathrm{D}}$ denotes 2D convolution. As this expression consists of several convolutions of different terms, some properties of this expression are apparent upon detailed inspection. Each Dirac delta function serves to pick out appropriate periodic "islands" of magnetisation in both z and ω space following convolution (e.g. when $\omega = n\Omega$), and these islands have the form of smoothed Bessel functions (i.e. J_n) in the z direction. In the spectral direction, the profile of the pulse is given by the inverse Fourier transform of $W(k_\omega)$. This implies that the usual Fourier relationships between the length of the pulse and its spectral properties will still apply, but that the excited magnetisation will spatially and spectrally vary as the outer product of the spectral and spatial W functions. Each island of excited magnetisation will have a finite spectral width, commonly called its passband, determined by the exact evaluation of Eq. (5.17). Similarly, the period between the islands is commonly known as the stopband as no excitation occurs for ω in this region. Note that no assumptions about $B_1(t)$ have been made, only that a pulse is played concomitantly with an oscillating gradient along one axis, and that the frequency of that oscillation (Ω) determines the spacing between islands.

Several properties of spectral-spatial pulses can now be derived from Eq. (5.17). If, arbitrarily, we choose R such that RG is large, then

$$\begin{aligned} |\dot{\mathbf{k}}(t)| &= \sqrt{\dot{k}_z^2 + \dot{k}_\omega^2} \\ &= \sqrt{(RG(t))^2 + 1} \\ &= \sqrt{R^2G^2\cos^2\big(\Omega(t-T)\big) + 1} \\ &\to RG\cos\big(\Omega(t-T)\big) \text{ as } \frac{1}{RG} \to 0, \end{aligned} \quad (5.18)$$

and the function $W(k)$ takes the form

$$W(t) \propto \frac{1}{\cos\big(\Omega(t-T)\big)} B_1(t). \quad (5.19)$$

The simplest illustrative choice of $B_1(t)$ is that of several Gaussian sub-pulses under a separate Gaussian envelope, i.e.

$$B_1(t) = B_{1\,\max} \overbrace{\exp\left[-\pi\left(\frac{\sin\big(\Omega(t-T)\big)}{U}\right)^2\right]}^{\text{Sub-pulses}} \overbrace{\exp\left[-\pi\left(\frac{t-\frac{T}{2}}{V}\right)^2\right]}^{\text{Envelope}} \overbrace{\cos\big(\Omega(t-T)\big)}^{\text{Gradient centres}}, \tag{5.20}$$

where $U > V > 0$.

This is sketched, together with the underlying oscillating gradient chain, in Fig. 5.3**A**. One can hence compute W as

$$W(t) \propto \exp\left[-\pi\left(\frac{\sin\big(\Omega(t-T)\big)}{U}\right)^2\right] \exp\left[-\pi\left(\frac{t-\frac{T}{2}}{V}\right)^2\right] \tag{5.21}$$

and thus Eq. (5.17) can be evaluated directly. Analytically, this calculation is protracted, but the relevant terms from the envelope can be simplified by use of the convolution theorem:

$$\begin{aligned} &\ldots \mathcal{FT}^{-1}\left\{\exp\left[-\pi\left(\frac{\sin\big(\Omega(t-T)\big)}{U}\right)^2\right] \exp\left[-\pi\left(\frac{t-\frac{T}{2}}{V}\right)^2\right]\right\} \ldots \\ \equiv &\ldots \mathcal{FT}^{-1}\underbrace{\left\{\exp\left[-\pi\left(\frac{\sin\big(\Omega(t-T)\big)}{U}\right)^2\right]\right\}}_{\text{Frequency}} \otimes_{\text{2D}} \mathcal{FT}^{-1}\underbrace{\left\{\exp\left[-\pi\left(\frac{t-\frac{T}{2}}{V}\right)^2\right]\right\}}_{\text{Space}} \ldots \end{aligned} \tag{5.22}$$

which can reveal the behaviour of the different parts of the pulse above. A numerical solution of the Bloch equations for a particular choice of U, V, and Ω is shown in Fig. 5.3. By explicitly evaluating these (inverse) Fourier transforms analytically, it can be shown that:

- As the Fourier transform of a Gaussian is another Gaussian of reciprocal width, the overall length of the pulse is proportional to the width of each passband of excited magnetisation, i.e. the FWHM of the islands discussed above scale as
$$\Delta\omega \propto \frac{1}{V}. \tag{5.23}$$

- Likewise, the gap between each sub-pulse is determined by the gradient oscillation frequency Ω, and this determines the spacing of a comb of Dirac delta functions in the resulting magnetisation excitation pattern. Accordingly, the stopband, the spacing between the islands of excited magnetisation is directly proportional to Ω;

- The spatial profile of excitation is approximately the Fourier transform of each sub-pulse. For the Gaussian example given;
$$\Delta z \propto \frac{1}{V}, \tag{5.24}$$
where Δz is the width of the (Gaussian) slice excited.

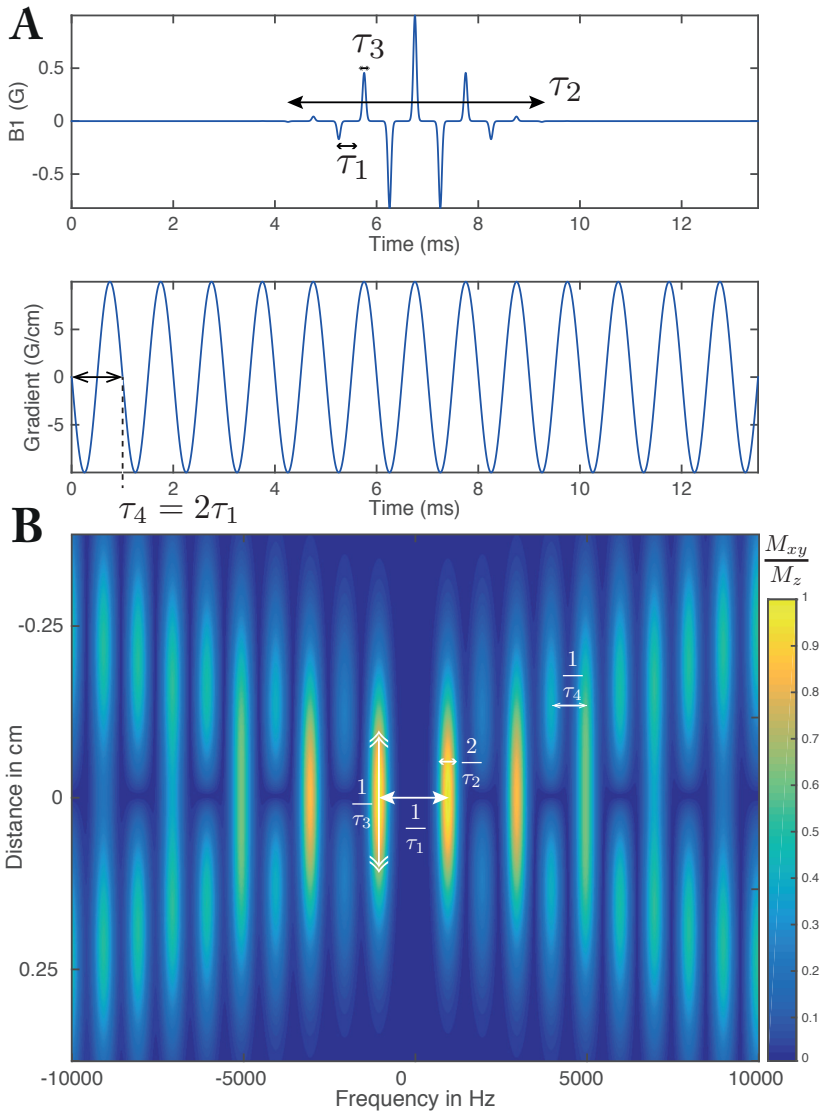

Figure 5.3: A pulse of the type described in Eq. (5.20), (**A**) and corresponding $M_\perp$ (**B**). Here, Ω is 1000 Hz, the number of sub-pulses is 27, G is 10 G/cm, $U = 0.5$, and $V = 0.002$. The corresponding time-per-sub-pulse is 0.5 ms, and the total pulse length is 13.5 ms. Several Fourier relationships are apparent: the time between sub-pulses τ_1 determines the centre gap between the first two 'islands' of excited magnetisation, known as the stopband which is of width $1/\tau_1$. The envelope of the pulse, τ_2 determines the width of the pass-band, the first islands of excited magnetisation of spectral width $2/\tau_2$. The spatial profile is approximately determined by that of each sub-pulse; which for the Gaussian example here is inversely proportional to their width τ_3, when convolved with a Bessel function J_n for the nth sub-lobe. The repetitive islands of magnetisation are excited at a nodal pattern determined by the spacing of the gradient chain, i.e. $\tau_4 = 2\tau_1$. The centre points of these islands form a cross with lines defined by $\pm 1/(\gamma G)$.

That many Fourier relationships exist between the design parameters of spectral-spatial pulse and its effects is arguably unsurprising given the existence of excitation k-space. The preceding text has served to demonstrate the origins of such relationships, and provides a quantitative explanation for the appearance of the plot of $M_\perp$ as shown in Fig. 5.3**B**.

However, it might also be argued that the above analysis is at the limit of analytical tractability and utility. It is time consuming and difficult to directly evaluate expressions such as Eq. (5.17), for both skilled humans and computer algebra systems such as Mathematica. For these reasons, a numerical approach is usually adopted in which the Fourier operators are discretised, analogously to techniques in physical chemistry that discretise Hilbert space to estimate observables in a quantum mechanical system of interest. Specifically, if F_{ij} is a $N \times N$ matrix given by

$$F_{jk} = \exp\left(2\pi i \frac{\left(j - 1 - \frac{N}{2}\right)(k-1)}{N}\right), \tag{5.25}$$

then it can be shown that s and S, column vectors $\in \mathbb{C}$ of length N, are approximate Fourier transforms of each other if

$$S = Fs.$$

As F is unitary, $F^\dagger = F^{-1}$, which provides a significant computational benefit if N is large.

5.3 | *Development & Design of a Multiple-Slice Spectral-Spatial Pulse*

5.3.1 *Numerical design*

The numerical approach first proposed by Cunningham *et al.*[136] was used and adapted to design a spectral-spatial pulse for hyperpolarised imaging, and is summarised below. The majority of the code used was freely obtained from John Pauly's `rf_tools` Matlab package, with modifications by Angus Lau, Chuck Cunningham, Peder Larsen, the author and others.[1]

Initially, design parameters for the spectral and spatial profiles of the pulse were set, such as its time-bandwidth product in both dimensions according to the expected excitation profile of the pulse (as discussed subsequently). Design parameters for the gradient waveform, such as maximum strength and slew-rate, the number of sub-pulses, their duration, the raster spacing, and hence the overall length of the pulse were also provided. A filter design method then generated β_i, the coefficients in the SLR polynomials (defined in Eq. (2.53)) that specify the overall pulse B_1. An oscillating trapezoidal gradient train was then prescribed matching the supplied information. At this point, the pulse was nearly fully specified. However, two corrections are then required: (1) the RF amplitude of each sub-pulse was

[1] See `http://rsl.stanford.edu/research/software.html` for further information.

'VERSEd', i.e. modified in order to correct for changing gradient amplitude underneath it to allow for excitation in minimum-time; and (2) the irregular sampling of k_z space was corrected for with zipper correction, described below.

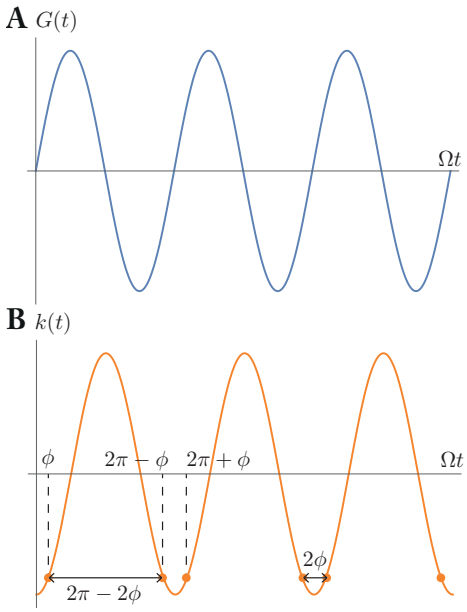

Figure 5.4: The density of samples in k-space is equally spaced at $k_z = 0$, but 'staggered' at some offset ϕ away. This is ultimately responsible for the islands of magnetisation shown in Fig. 5.3.

In the analytical discussion above, the sinusoidal slice-select gradient $G(t)$ gives rise to a series of islands of magnetisation separated by $1/\tau_4$ in frequency space, as represented by the Dirac comb in Eq. (5.17). This is undesirable, as it implies that the only way to have a sparser excitation is to decrease τ_4, increasing the gradient oscillation frequency, which comes at the expense of the minimum slice thickness available and may incur hardware penalties or reach hardware limitations. An alternative approach is to compensate for these unwanted islands by changing the k-space weighting function W to explicitly minimise them. An equivalent explanation for the origin of these islands is that the oscillatory gradient effectively introduces a sampling bias in k_z space; as shown in Fig. 5.4, at $k_z = 0$ the time between samples is equally spaced. However, at some small time offset $\phi \in \mathbb{R} \subset [0, \frac{\pi}{2}]$, samples on $k_z \neq 0$ are staggered, with alternate spacing between subsequent samples of the same point. These samples in k_z are taken at a time ϕ, $2\pi - \phi$, $2\pi + \phi$, ..., $2\pi n + \phi$, where $n \in \mathbb{Z}$. Accordingly, there are two spacings between samples, one large ($2\pi - 2\phi$) and one small (2ϕ). Owing to the periodicity of this sampling $\bmod 2\pi$, it is appropriate to think of this staggering as a sample delay Δ that alternates between successive samples k, giving a delay of $(-1)^k\Delta$. Accordingly, the appropriate discretised Fourier operator describing the delay (i.e. Eq. (5.25)) becomes modified to take into account Δ on a particular line, giving

$$G_{jk} = \exp\left(2\pi i \frac{\left(j - 1 - \frac{N}{2}\right)\left(k - 1 + (-1)^k\Delta\right)}{N}\right). \tag{5.26}$$

The B_1 that is hence required in each point of the non-uniformly sampled line of excitation k-space can therefore be found by solving the weighted least squares problem

$$\begin{aligned} GB_1 &= H \\ \Rightarrow WGB_1 &= WH \\ \Rightarrow B_1 &= (G^\dagger WG)^{-1}G^\dagger WH \end{aligned} \tag{5.27}$$

where H is the desired spectral response, B_1 is a column vector specifying the amplitude of $B_1(t)$ at point $1 \ldots N$, and W a diagonal matrix of weights that specify the important parts of the magnetisation response to fit to. The form $(G^{\dagger}WG)$ hence acts as a pre-conditioner to improve the numerical stability of this problem. For this work, W is taken to be 1 in the central part of the spectral stopband, and 0.01 elsewhere, indicating a desired flatness in the stopband. Equation (5.27) was then solved numerically by a least-squares method and the column vectors containing B_1 and the gradient waveform returned. The net effect of Eq. (5.27) is essentially a non-linear warping in excitation k-space that corrects for unequal sampling in k_z by slightly modifying the RF waveform. This then effectively removes every other island of magnetisation, doubling the stopband without sacrificing time. The effect of the pulse can then be determined by the traditional numerical solution of the Bloch equations or by an appropriate SLR transformation. The numerical aspects of the design process take a few minutes on a modern computer (with an Intel Core i7-3770, 3.40 GHz 8-core CPU, and 16 GiB RAM running Linux/3.13.0-33 generic) with appropriate Matlab MEX files.

5.3.2 Pulse parameters

A 17 sub-lobe, 6.8 ms long spectral-spatial pulse was designed with the algorithm detailed above, with the intention of using it in a multiple-slice EPI rapid imaging sequence. Its spectral pass-bandwidth was 400 Hz (FWHM), which was believed to provide a degree of immunity against B_0 inhomogeneity. The smallest achievable slice thickness was 1 mm/4 mm (^{1}H/^{13}C) with $G_{\max} = 17.5$ G/cm on the 7 T system described previously. The chosen time-bandwidth products were three in both the spectral and spatial domain, with a $4\,\mu s$ raster on both the RF and gradient waveforms. This grid resolution matches that of the gradient controller on the Varian system, although the RF controller is capable of using a much finer grid. The pulse is shown in Fig. 5.5**A**, with the spectral-spatial excitation pattern in **B**, cross-sections of which are shown in **C** and **D**.

The pulse was initially implemented with a gradient echo readout for verification purposes, where both the slice-select and RF were played on the basis of MATLAB-generated text files produced by the above algorithm. Its frequency and slice profiles were measured, and these are plotted in Fig. 5.6**A** and **B**. Some Gibbs ringing was observed in the slice profile, which is believed to be a consequence of the finite temporal resolution of the pulses played. It was initially found that a discrepancy arose from the incomplete 'zig-zag correction' described above, leading to the spurious excitation marked by the arrow in Fig. 5.6**A**. This artefact has been reported previously, and is believed to be due to a timing shift between the RF and oscillating gradient waveform.[137] Given the characterised behaviour of the gradient set, such errors are not surprising: the non-zero phase waveforms in Fig. 3.5, particularly on z, serve to introduce a propagation delay between the gradient synthesis and amplification at high frequencies. For this pulse, exciting a 4 mm thick slab of magnetisation requires an oscillating square

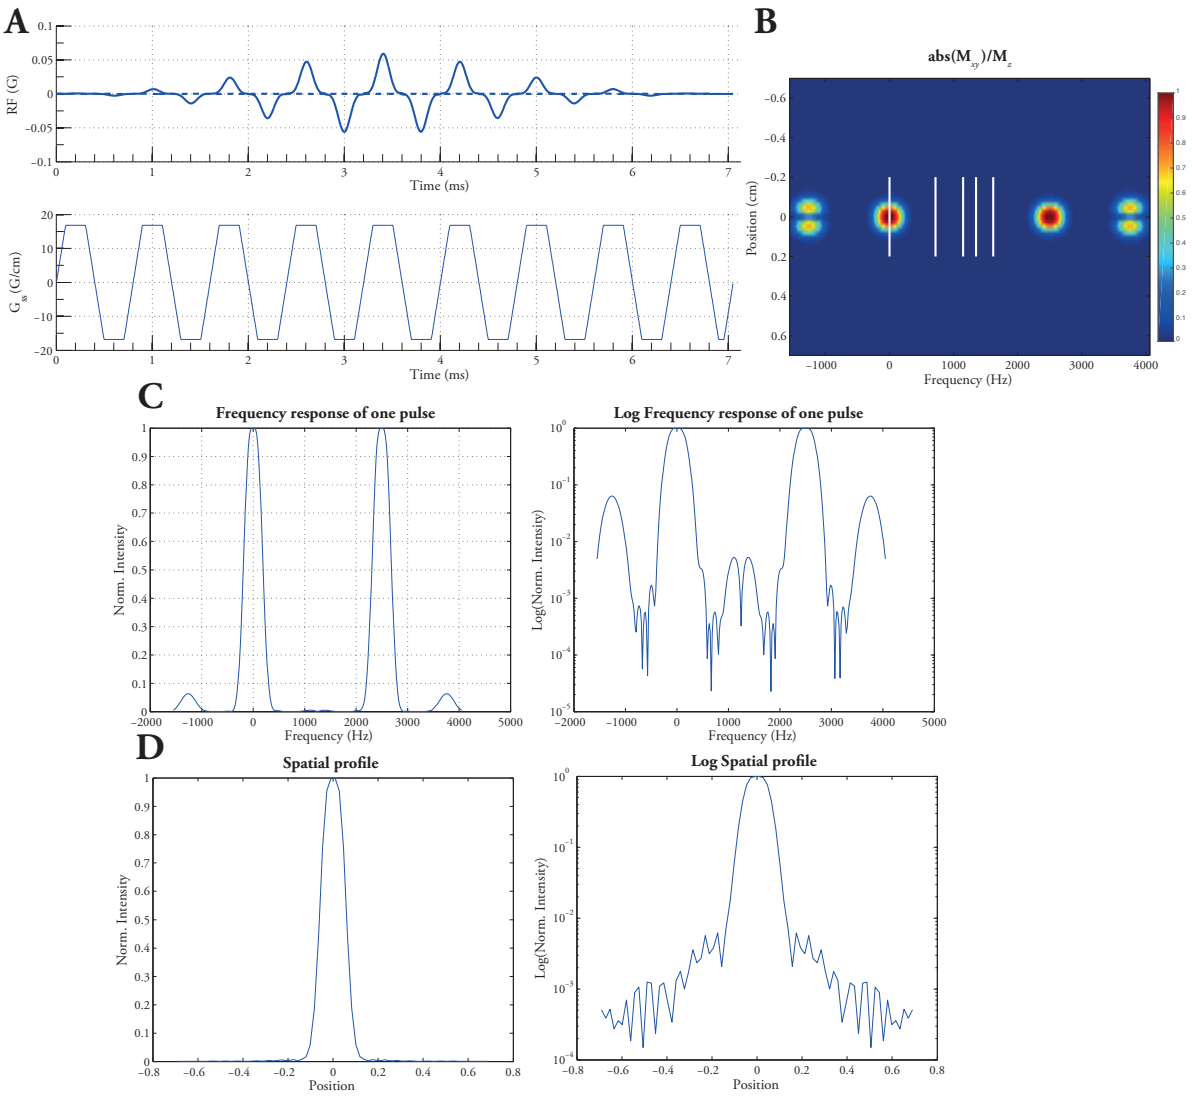

Figure 5.5: **A**: The spectral-spatial pulse, a superposition of 17 sub-lobes that are approximately Gaussian pulses played under an approximately Gaussian envelope. **B**: Solving the Bloch equations for the pulse reveals the spectral-spatial profile of the pulse. The vertical white lines shown correspond to the peaks in the [1-^{13}C]pyruvate spectrum following injection: from left to right, bicarbonate, pyruvate, pyruvate hydrate, lactate and alanine. **C:** The spectral selectivity of the pulse is clearly emphasised in the semi-log graph. **D**: The spatial excitation of the pulse is narrow, and uniform (shown at $G = 17.5$ G/cm).

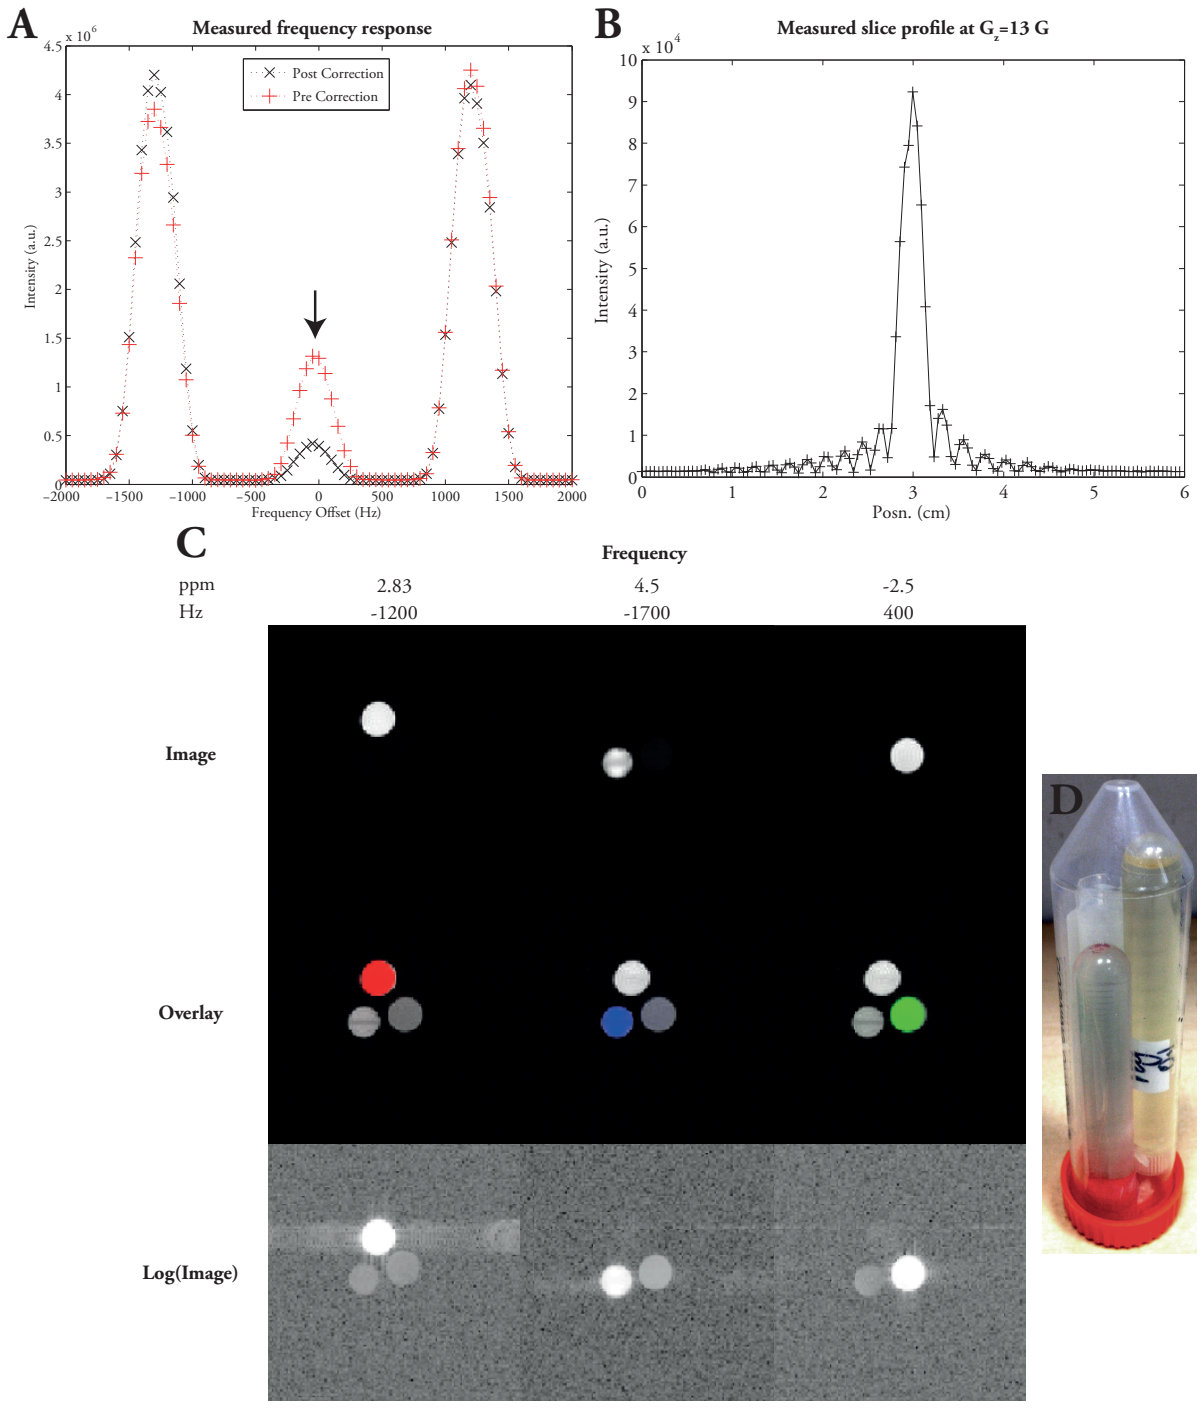

Figure 5.6: Measured properties of the spectral-spatial pulse when implemented with a Gradient Echo readout. **A**: By varying the transmitter frequency whilst imaging a homogeneous sample, the spectral domain of the pulse can be scanned and its excitation profile measured, predicted in Figure 5.5C. Basic eddy current compensation improved performance at the point shown (arrow). The measured slice profile at relatively high gradient strengths is, again, not as ideal as that predicted. **C**: Nevertheless, images produced by the gradient echo spectral-spatial pulse sequence of a chemical shift phantom, consisting of three tubes of different substances (acetone, water and vegetable oil; shown in **D**) were encouraging.

wave of period $400\,\mu$s running at 17.5 G/cm.

To ameliorate this pattern of spurious excitation, corrective delays were introduced between the RF and gradient waveforms, and their value set by minimising the excitation at 0 Hz. This process improved the central selectivity of the pulses by approximately a factor of two (c.f. Fig. 5.6**A**). However, this extinction coefficient was approximately an order of magnitude higher than that predicted in the design of the pulse, where the ratio of excitation close to 0 Hz to that at the maximum was predicted to be ~ 300 (c.f. Fig. 5.5**C**). This pattern of spurious excitation was reproduced when the Bloch equations were solved with a deliberate time offset between the RF and gradient.

SpSp-Gradient Echo images of a chemical shift phantom (a 50 ml Falcon tube containing three smaller tubes of water, acetone and vegetable oil, chosen because of their similar chemical shifts to pyruvate, bicarbonate and lactate) showed clear spectral separation in each acquired channel between the different compounds. Upon examining the natural logarithm of the image, a slight degree of cross-contamination was observed. This was initially believed to reflect both the wide dynamic range of the receiver and the small spurious excitation at 0 Hz. These results were seen as encouraging, and verified the function of the pulse designed above.

5.4 | *Implementation & Results*

5.4.1 *Implementation*

In order to facilitate rapid imaging, this pulse was incorporated into an EPI readout trajectory. In order to use a vendor-supplied EPI module, it was necessary to generate the correct slice-select gradient waveform that corresponded to the SpSp RF pulse on the fly, as opposed to reading in a pre-existing text file. Additionally, it was necessary to modify `psg`, the scanner's standard library, to permit the excitation of slices away from isocentre. Ordinarily, the use of a constant linear slice-select gradient means that a spatial translation linearly corresponds to a shift in frequency, $\Delta\omega = \gamma G \Delta z$. The use of an oscillating gradient in a spectral-spatial pulse serves to make this frequency offset a function of time: $\Delta\omega(t) = \gamma G(t) \Delta z$, where $G(t)$ may vary between slices.[137] Unfortunately, hardware limitations meant that the transmitter frequency cannot be changed on a microsecond basis, but that it was possible to vary the phase of the transmitted RF on a 4 μs raster. Accordingly, an extension to `psg` was written to solve the differential equation $\frac{\mathrm{d}\phi(t)}{\mathrm{d}t} = \gamma G(t) \Delta z$ at runtime, and generate RF transmitter `acode` (binary words) that corresponded to a particular phase increment.[2] Additionally, it was believed that this approach may be limited by the quantised nature of the phase changes: the Nyquist sampling theorem implies that spatial aliasing may result if $2 \times \Delta\omega > 4\,\mu$s. As it was not known whether the RF transmitter

[2]This process was technically challenging, and the author was obliged to visit Agilent in Yarnton in 2013 with source code samples to demonstrate that such low-level modifications would not damage the hardware, and was the appropriate way to solve this problem.

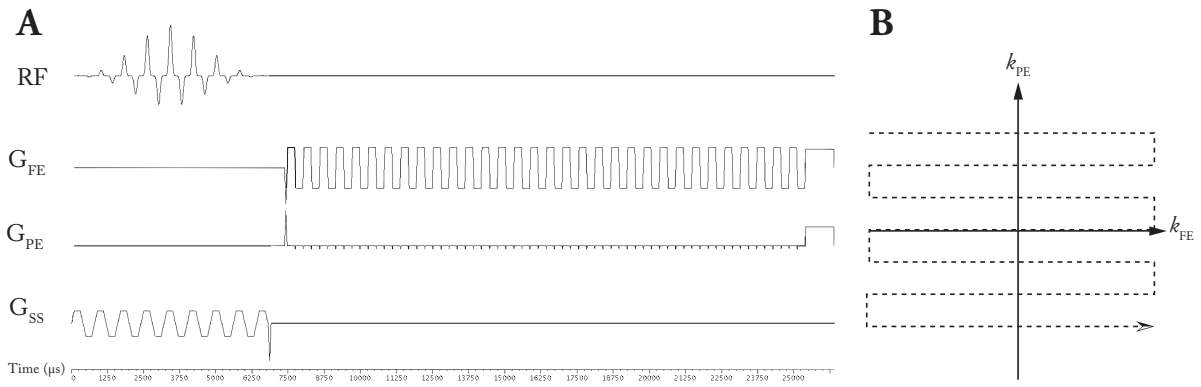

Figure 5.7: **A**: The designed spectral-spatial pulse followed with a Cartesian EPI readout sequence, as implemented. **B**: The corresponding k-space trajectory. An image can be acquired in < 30 ms, with a single RF excitation.

operated in a phase continuous fashion with or without phase memory, it was therefore decided to avoid moving slices more than ~ 2 cm from isocentre, which would limit any deleterious effects from the quantised nature of these phase shifts.

The sequence as it was initially implemented is illustrated in Fig. 5.7. As EPI had not been run on this particular system before, it was necessary to define scan parameters such as the field of view, receiver bandwidth, the equivalent phase encoding bandwidth, overall readout length, and matrix size. Initial attempts were limited by the inappropriate (default) setting of the receiver bandwidth, which, when combined with the default priority method of gradient moment calculation limited the smallest FOV to be 150×150 mm^2 with a 64×32 matrix for a resolution of $2.4 \times 4.8 \times 4$ mm^3. The receiver bandwidth was 250 kHz, readout gradient strength 15.57 G/cm, TR 500 ms per metabolite, TE 12 ms, and no navigators were used. Phase blips were 120 μs, of magnitude 1.04 G/cm, and the total duration of each combined readout/PE lobe was 0.46 ms. Ramp sampling was not performed. Metabolite acquisition was performed in the order

$$ {}_1X^1_1,\ {}_1X^1_2,\ \ldots,\ {}_1X^1_n,\ {}_2X^2_1,\ \ldots,\ {}_2X^2_n,\ {}_3X^3_1,\ \ldots,\ {}_3X^3_N,\ {}_1X^4_1,\ \ldots \tag{5.28} $$

where ${}_\alpha X^j_i$ denotes imaging metabolite α at time-point j for the slice $i = 1 \ldots n$. Owing to the limited cardiac window for imaging, and the comparatively large maximum slice thickness, it was decided to avoid attempts at multiple-slice imaging in the rat heart, i.e. $n = 1$. The surface butterfly coil was used, with a nominal flip angle of $\theta = 15^\circ$ for pyruvate and $\theta = 90^\circ$ for other metabolites.

An initial attempt was made to image rodent cardiac metabolism with this sequence as a proof of principle. These initial results from this readout were unsatisfactory, and are illustrated in Fig. 5.8. It was quickly noted that the resolution requirement for meaningful metabolic imaging was not realisable with this EPI readout. Additionally, the Nyquist ghost, discussed in Section 2.4.6.2, proved to be troublesome, and data was reconstructed (in Matlab) by taking alternate lines of k-space, reconstructing an image for

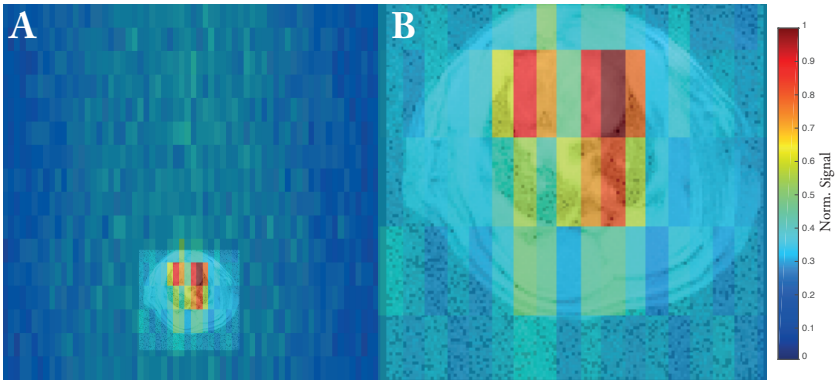

Figure 5.8: Unsatisfactory, low-resolution initial pyruvate image of the rodent heart acquired with the sequence without ramp-sampling, using the default EPI readout and limited by gradient duty cycle (**A**). A crop to the animal is shown in **B**.

each set of alternate lines, and adding the two resulting images in magnitude. This process halved the (excessively large) FOV to ameliorate the ghost.

It was also noticed that images acquired in each of the three metabolite channels did not necessarily co-localise with either each other or their expected location on the underlying anatomy. It was quickly realised that these artefacts would arise from errors in setting the central receiver and transmitter frequencies respectively, as an unknown shift in frequency corresponds to an unknown shift in image space. This raised two issues: (1) it was realised that an accurate method to find the central frequency of [1-^{13}C]pyruvate before injection was required, in order to minimise bulk shifts in the resulting stack of images; and (2) B_0 homogeneity was particularly important to prevent partial spatial excitation of a metabolite, or differential shifts in image space.

These issues were initially tackled by the implementation of an accurate automated shimming algorithm (detailed in Section 3.4), and by the inclusion of a bottle of 5 M ^{13}C-urea next to the subject scanned, which acted as a central frequency reference (and a positive control that the RF coil installed was correctly connected). A simple two-echo B_0 mapping sequence was implemented, which allowed for the measurement of the frequency difference between the urea reference and the region of interest, allowing the central transmitter frequency to be set correctly,[3] as $\omega_{\text{transmit}} = \omega_{\text{urea}} \pm \Delta B_0 \gamma_{^{13}\text{C}} / \gamma_{^{1}\text{H}}$, where ω_{urea} is the difference in chemical shift between urea and pyruvate at 7 T (approximately 300 Hz in the heart) and ΔB_0 the variation in B_0 measured.

To improve resolution, the EPI readout kernel was rewritten with ramp sampling achieved by a variable sampling time method (which also decreased the echo time and gradient duty cycle), and the gradient readout calculation algorithm was improved. With a readout receiver bandwidth of 125 kHz, and a readout maximum gradient amplitude of 8.75 G/cm, the minimum achievable FOV

[3]The actual implementation of this on the Varian system is somewhat convoluted; `tof` is the relevant transmitter offset parameter, a double which has the units of Hertz and the sign convention of ppm. Zero corresponds to the transmitter frequency being equal to that of the corresponding lock frequency for Deuterium; `tof` is therefore usually on the order of a few kHz, where negative values increase the transmitter frequency.

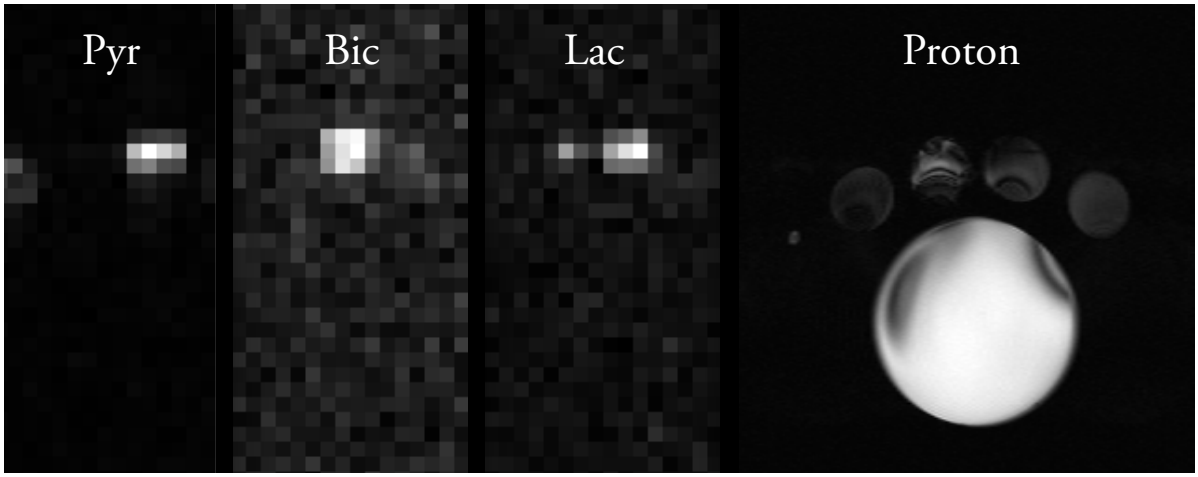

Figure 5.9: Single time-point images of pyruvate, bicarbonate and lactate produced *in situ* in the chemical shift phantom following injection of hyperpolarised [1-^{13}C]pyruvate. The relatively high concentration of different chemical species and small size of the syringes imaged made conventional proton gradient echo imaging difficult, as illustrated on the right.

was $60 \times 60 \times 4$ mm^3. The sequence was additionally modified to include two linear navigators for estimation of the linear phase term $e^{i\phi}$ responsible for the Nyquist ghost. The phase encoding part of the EPI kernel had a similar equivalent bandwidth (2.5 kHz) with a maximum gradient of 3.53 G/cm, a TE of 14 ms, and other parameters unchanged.

5.4.2 Results

As a proof-of-principle of the updated sequence, a chemical shift phantom was constructed and imaged. The phantom consisted of a series of 5 ml syringes (outer diameter ~ 1 cm) filled with either water, lactate dehydrogenase (LDH), or hydrogen peroxide, next to urea as a frequency reference. The first syringe was empty to allow for the generation of a [1-^{13}C]pyruvate image. The second chamber contained 10% hydrogen peroxide solution at pH$= 8.8$ and buffered by TRIS-K$^+$, which rapidly evolves ^{13}C-bicarbonate when mixed with hyperpolarised [1-^{13}C]pyruvate. The third chamber contained 1 kUnit of LDH and 10mM NADH, to generate [1-^{13}C]lactate when mixed with [1-^{13}C]pyruvate (all reagents purchased from Sigma Aldritch). The syringes were connected to a three-way vein infusion kit outside the magnet, and could be individually (re-)charged with their respective reagents, or all simultaneously injected with pyruvate. A larger tube containing water was also added for localisation. The phantom allowed the more rapid optimisation of EPI readout parameters than using live animals, and an example image obtained (with sequence parameters described above) is shown in Fig. 5.9. The FOV has been halved in the reconstruction of these images to remove the Nyquist ghost, leading to some aliasing.

As shown in Fig. 5.9, this readout was less prone to artefacts, shortened temporally, and was able to resolve metabolism in the the phantom discussed above. Accordingly, the same sequence was used

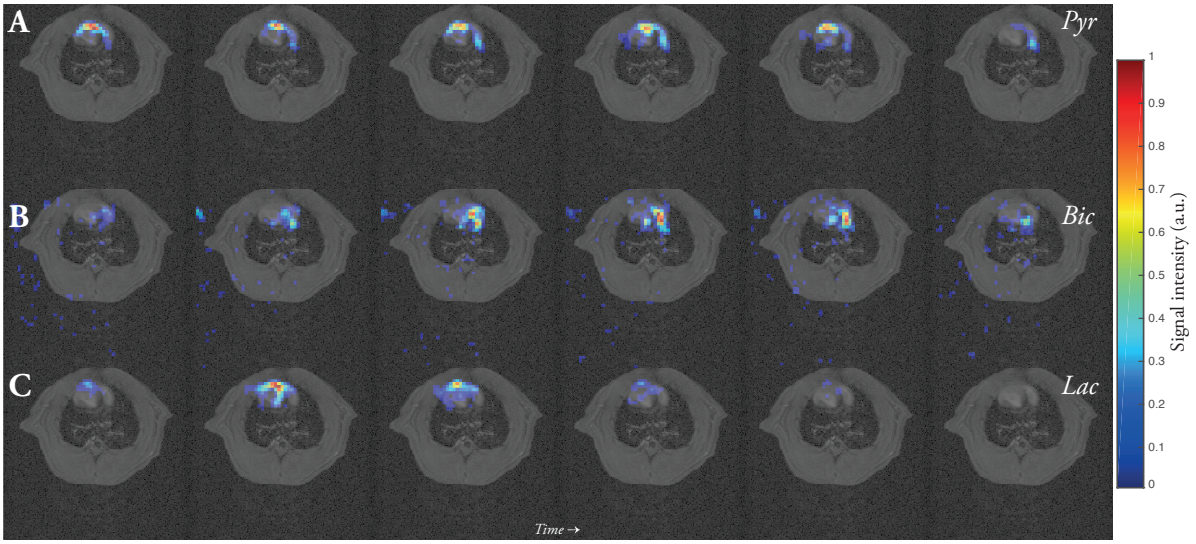

Figure 5.10: Single slice images of [1-^{13}C]pyruvate (**A**), ^{13}C-bicarbonate (**B**) and [1-^{13}C]lactate (**C**) following the injection of hyperpolarised pyruvate. Time increases from left to right. Pyruvate is visible in the ventricles of the heart; lactate and bicarbonate in the myocardium. Both the bias in these images towards the front of the chest, and the fact that they appear somewhat distorted are likely due to to the profile of the rat butterfly surface coil used in this experiment.

to acquire data from the *in vivo* rat heart, as shown in Fig. 5.10. The slice thickness was deliberately increased to improve SNR, leading to a scan with a resolution of $2 \times 2 \times 8\,\text{mm}^3$. Pyruvate could be observed perfusing the ventricles of the heart, and lactate and bicarbonate seen to approximately co-localise with the myocardium. However, it was subsequently realised that the images obtained were still subject to significant distortions, and that they were predominantly biased towards the ventral chest wall of the animal. It was believed that the limited spatial coverage of the surface coil used was responsible for these shortcomings, as the spatially varying flip angle and limited region of homogeneity would predominantly cause a drop off and non-linear distortions that would be progressively worse with any in-plane flow. This fact drove the development of the hardware necessary to use the volume transmit/two-channel surface receive array coil combination, which features superior transmit homogeneity and receive sensitivity, but, prior to its redesign, was not able to be used reliably due to the EM interactions discussed in Section 3.2.5.

This sequence was additionally used to image metabolism in the brains of healthy rats, with the home-made ^{13}C brain coil. Scan parameters were as above, but without cardiac gating. Temporally summed single-slice images acquired with the ^{13}C brain coil are shown in Fig. 5.11, and are believed to show [1-^{13}C]pyruvate in the middle cerebral artery (MCA), and lactate in the brain and MCA. The brain coil used featured greater spatial homogeneity in flip angle than the rat butterfly coil, but was difficult to position accurately on the rat's head within the tight spatial confines of the magnet bore.

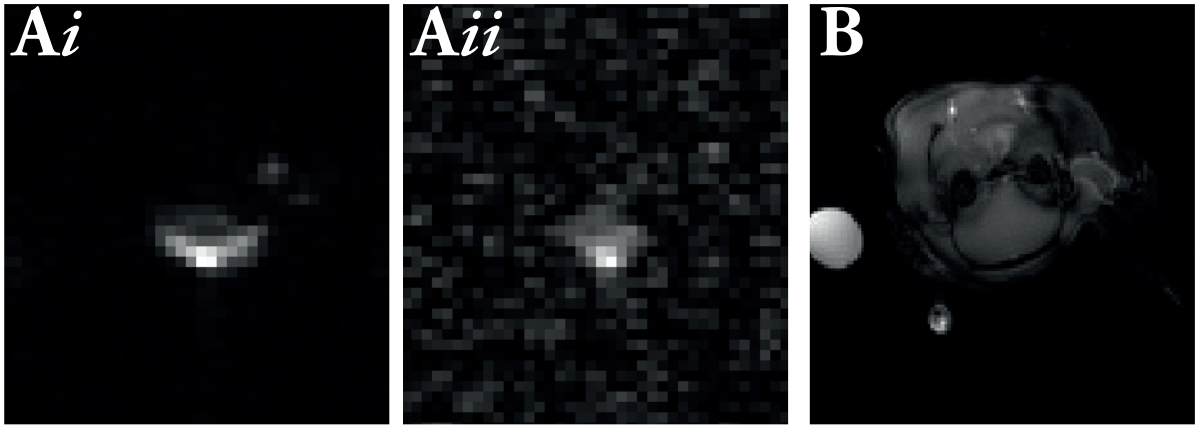

Figure 5.11: Temporally summed single-slice images of both (**Ai**) hyperpolarised [1-^{13}C]pyruvate and (**Aii**) [1-^{13}C]lactate acquired with the sequence discussed previously. **B**: A proton overlay of the same FOV and slice thickness but with a higher matrix size shows that the pyruvate signal is believed to originate primarily from the middle cerebral artery, in the centre of the brain (and the larger Circle of Willis, not shown). A 1.5 ml bottle of ^{13}C urea was placed next to the brain as a frequency reference, and is visible on the left of the image. The coil used contains a positioning phantom, a blown glass bead filled with gadolinium-doped saline, visible in the lower left of the image.

5.5 | Discussion

It was found that the imaging sequence developed was not particularly robust with either coil: the success or failure of an experiment to be reconstructible into artefact free images was approximately independent of the success of previous experiments; and approximately 50% of experiments failed. This lack of robustness reduced the utility of the sequence to such an extent that it was not considered ethical to subsequently demonstrate its utility with animal models of disease. Neglecting both operator error and hardware failure,[4] four main reasons are proposed for this lack of robustness:

1. Centre frequency shifts and B_0 inhomogeneity;
2. Flip angle uncertainty and variability;
3. Insufficient z-resolution; and
4. Gradient infidelity.

5.5.1 Centre frequency shifts

The key premise of spectral-spatial techniques for metabolic imaging is that the ^{13}C spectrum obtained after injection of [1-^{13}C]pyruvate is known exactly *a priori*, i.e. the frequency of the metabolites of interest are known before the scan starts. Therefore, the spectrometer has to be provided with an accurate centre ^{13}C frequency for the forthcoming scan. The relatively wide spectral bandwidth of the excitation (400 Hz) means that any minor uncertainty in this value will still produce excitation, and it

[4] Both of which were unfortunately not infrequent.

was initially assumed that one spectroscopy study on the target organ of interest would provide an adequate frequency reference for future work. Unfortunately, it was quickly realised that this was not the case: routine shimming changes B_0 sufficiently to vary the centre [1-^{13}C]pyruvate frequency by more than the bandwidth of the excitation pulse, resulting in no excitation of [1-^{13}C]pyruvate and its downstream products, or, depending on the magnitude of the error, the appearance of [1-^{13}C]pyruvate in the other metabolite channels. As described above, a bottle of ^{13}C-urea was included next to the organ of interest, and used as a frequency reference. B_0 maps were also routinely acquired to correct for the frequency difference in space between the reference and the target organ of interest. It was noted that the apparent B_0-compensated chemical shift of injected hyperpolarised [1-^{13}C]pyruvate relative to urea was different between the heart and the brain. This fact necessitated occasional 'frequency scout' spectroscopy experiments before an imaging sequence could be run, which might further degrade the utility of the sequence proposed.

Owing to the small phase encoding equivalent bandwidth (c.f. Eq. (2.62)) used in the above EPI readout ($\mathcal{O}(2500)$ Hz or ~ 40 Hz/pixel), any residual error in the centre transmitter frequency causes a spatial shift in the image along the phase encoding axis. Unless the exact value of the frequency difference between the centre transmitter frequency and the actual frequency of the injected [1-^{13}C]pyruvate can be obtained, this error would prevent the accurate co-registration of the hyperpolarised data and anatomical images.

5.5.2 *Delivered flip angle*

As hyperpolarised magnetisation is not renewable, it was quickly realised that the delivered flip angle was crucially important for hyperpolarised experiments. It was believed that apparently different kinetic curves might result from a nominally identical experiment if the flip angle delivered differed from that requested. Where a decaying metabolite of interest is spatially constrained within one region – for example, ^{13}C-bicarbonate within the myocardium – then any flip angle variation over that region will produce apparent regional differences in kinetics that are erroneous. This fact is believed to be responsible for some of the distortions present in Fig. 5.10, and drove the development of appropriate hardware to use the volume transmit/surface receive array coil, which featured far better homogeneity. Additionally, it was realised that the volume transmit homogeneity of this hardware would allow for a flip angle calibration to be performed on the urea frequency reference present with every experiment.

5.5.3 *Resolution requirements*

The rodent heart is not large: it measures approximately 10 mm from base to apex, and is approximately 10 mm in diameter. The left ventricular wall comprises the greatest density of tissue, and is approximately, 4 mm thick at diastole. The right ventricle is far smaller, with a wall thickness on the order of 1 mm, and

is otherwise crescent shaped with principal lengths of 10 and 3 mm. The atria are typically not resolved on preclinical proton CINE MR scans, but occupy the anterior ~ 4 mm of the length of the heart. It is in this context that the minimum achievable slice thickness of 4 mm of the sequence proposed above should be considered: it represents a projection along almost the entire left ventricle, and therefore is comparatively coarse. Additionally, multi-slice imaging of the rodent heart proves challenging due to the large duty-cycle requirement that this sequence imposes.

Conversely, the rodent brain is comparatively long at ~ 20 mm compared to the slice thickness used here. Multi-slice imaging approaches with different numbers of slices in different organs would produce kinetic data that would not be directly comparable, as the different magnetisation sampling patterns would contribute directly to the apparent metabolite timecourses shown. In particular, attempting to perform multi-slice imaging in the heart could lead to significantly different results depending on the order of slice acquisition and the direction of blood flow: if pyruvate was imaged early in one slice before flowing to another, that second slice would have an apparently reduced level of perfusion or, potentially, of visible metabolite production from pyruvate. There is therefore a need for a sequence which possesses improved slice resolution with improved coverage, and ideally that sequence would sample magnetisation uniformly in a large volume of space.

5.5.4 Evidence of distorted excitation

The known chemical kinetics of pyruvate metabolism imply that the temporal profile of pyruvate images and downstream metabolite images should always differ. Specifically, no downstream metabolite should be observed instantaneously at the moment pyruvate is injected (or before), and the nonzero time taken for the enzymatic metabolism of pyruvate strongly implies that a delay should be present between the time of maximum pyruvate signal and maximum downstream metabolite signal. These conditions were satisfied for the majority of data acquired with this imaging sequence. However, a significant minority of imaging datasets contained time-courses that violated these requirements, indicating an artefact attributable to the imaging sequence. Figure 5.12 illustrates example pyruvate and lactate time-courses, reconstructed from a bright voxel corresponding to the MCA in two separate brain imaging experiments. One of these experiments was successful, the other not. For the 'good' data set (Fig. 5.12**A**) a difference of approximately 15 frames (TRs) was observed between the point of maximum pyruvate and lactate production in the MCA, corresponding to a delay of approximately 1.5 s between the two maxima. The SNR of each time course was relatively high, and pyruvate was observable for several frames over the three minute acquisition period. Figure 5.12**B** shows a 'bad' data set, where the SNR of both timecourses (and images) was low, and, crucially, no difference in time was observed between the point of maximum pyruvate intensity and maximum lactate intensity. The signal was not visible for more than ~ 5 frames.

It was believed that the non-linear behaviour of the gradient set (c.f. Fig. 3.5**A**) used could be

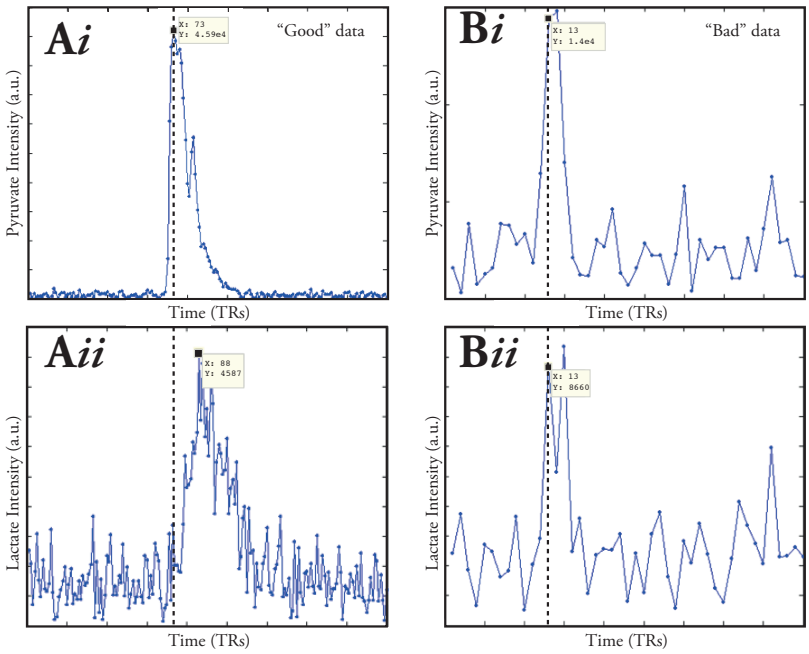

Figure 5.12: Two example pyruvate (***i***) and lactate (***ii***) timecourses obtained from a single bright voxel (believed to be the MCA) in separate brain imaging experiments. One dataset, **A**, is 'good', in the sense that the reconstructed pyruvate and lactate timecourses have a shape consistent with that of hyperpolarised metabolism, and displays a clear temporal separation between the point of maximal pyruvate signal (frame 73; dashed line) and maximal lactate signal (frame 88). The other timecourse, **B**, is 'bad' in the sense that this temporal separation is not present, and both maxima are coincident in frame 13 (dashed line). The x-axis differs between **A** and **B**.

responsible for these artefacts: as previously shown in Fig. 5.6**A**, eddy current effects create spurious excitation at 0 Hz that degrades the spectral profile of the pulse. It was believed that, depending on the dielectric loading of the magnet (i.e. the mass of the rat contained within it, and the position of the stainless steel IV cannula needle within the rat), the eddy current profile could be slightly altered. This alteration could change the effective value required for the corrective delay introduced previously, and hence lead to a degradation in the spectral profile of the pulse. The observed voxel time-courses are explained by the imperfect off-resonance excitation of pyruvate with both the pyruvate and lactate channels, or, say, pyruvate hydrate. This proposed sample-dependent degradation of gradient performance is well known on human MR systems, and has lead to the development of field cameras to independently measure the imaging k-space trajectories actually taken by fast imaging sequences. As the zipper correction method taken in designing the pulse for excitation requires a relatively high level of gradient fidelity, the known poor performance of the gradient set remains a plausible, though unproven, candidate explanation for the incorrect timecourses shown in Fig. 5.12.

As it is not possible to easily measure either the gradient set performance, nor to optimise carbon excitation profile at 0 Hz with a rat in the magnet, the presence of these artefacts, however infrequently, was concerning.

5.6 | Concluding remarks

This chapter has documented the *de novo* development of a spectral-spatial imaging sequence at 7 T, and provides brief demonstrations of its potential utility for imaging metabolism in the rodent heart and brain. Four key limitations with the sequence as proposed were identified: (1) the need to identify the correct centre frequency and correct for shifts in image-space arising from errors in this value; (2) the need for homogeneous B_1^+ and a method for within experiment flip angle calibration (which did not exist when this work was performed); (3) a desire for resolution in excess of that which could be achieved in order to provide a capability for meaningful whole organ metabolic imaging; and (4) the presence of excitation artefacts indicates a potential gradient fidelity issue. Consequently, this sequence was not used for any larger scale pre-clinical studies or demonstrations of efficacy, as the sequence was not considered robust enough for routine use.

Instead, methods of overcoming these limitations were investigated, leading ultimately to the extension of the above work with the adoption of a slightly modified 3D 'flyback' spectral-spatial excitation pulse, described subsequently in Chapter 6. Nevertheless, spectral-spatial excitation has been shown to effectively allow for the rapid, magnetisation-efficient imaging of hyperpolarised [1-^{13}C]pyruvate *in vivo*, in small animals at 7 T.

Chapter 6

Development of a Three-Dimensional Spectral-Spatial Imaging Sequence

6.1 | *Overview*

As discussed in the preceding chapter, spectral-spatial excitation enables the magnetisation-efficient imaging of the metabolism of hyperpolarised [1-^{13}C]pyruvate following its injection into living systems. However, the two-dimensional multiple slice sequence developed suffered from several key limitations in resolution and spatial coverage, and additionally was prone to multiple artefacts arising from either frequency shifts and B_0 inhomogeneity, or poor gradient fidelity.

At the same time as this work was being undertaken, other researchers became aware that single or multiple slice spectral-spatial excitation pulses form an attractive method for hyperpolarised imaging. Correspondingly, their utility for metabolic imaging has been shown in the pig heart[138] and in mice[139], and for saturation-recovery rate mapping in rat tumour models.[140] The inverse relationship between RF pulse length and excitation bandwidth necessitates that all such proposed spectral-spatial pulses are relatively long. As noted by Sigfridsson *et al.*,[141] and as discovered in the preceding chapter, this puts strenuous requirements on both the maximum strength and slew rate of the gradient set used for imaging. At high field, these requirements effectively preclude the excitation of thin slices. To pursue higher resolution imaging, Sigfridsson *et al.* proposed a different multi-band approach, in which a temporally short spectral-spatial pulse was used to excite multiple metabolites with different flip angles, and a multi-echo readout chain was then used for multiple metabolite imaging. The shorter pulse allowed for the excitation of a slice whose thickness would not have been achievable with conventional pass/stopband single-excitation pulses. However, as noted by Sigfridsson, this scheme fails if B_0 is not homogeneous, and requires accurate knowledge of the centre frequency of the arriving pyruvate bolus. On account of the oscillatory excitation profile, a small error in this frequency translates into a large error in the flip angle that each metabolite experiences. Pass/stop-band spectral-spatial pulses are more tolerant of off-resonance effects and B_0 inhomogeneity, which improves their utility *in vivo*. However, both approaches become SNR limited as the slice thickness decreases.

This chapter describes the development of a centric ordered fully three-dimensional pass/stopband

spectral-spatial EPI sequence that enables the imaging of hyperpolarised metabolites with a spatial resolution of $2 \times 2 \times 3.8$ mm^3, and a (duty-cycle limited) maximum temporal resolution of 780 ms per 3D image. This sequence is an extension of that described in Chapter 5, and overcomes its limitations. To minimise the effects of gradient infidelity, a fly-back spectral-spatial pulse has been developed and used, whereby RF is only transmitted on the positive lobes of the underlying slice selection gradient. The pulse excites a comparatively large slab (45.5 mm) of magnetisation, with subsequent phase encoding in the z-direction trading time for z-resolution. An EPI readout train then follows, with a phase encoding direction that flips on alternate acquisitions to allow for the *post hoc* correction of frequency shifts and, theoretically,[142] flow and susceptibility artefacts in the resulting cardiac images.

After verifying the technique by imaging phantoms, its sensitivity to metabolism was demonstrated by imaging rodents whose metabolic flux through pyruvate dehydrogenase (PDH) had been modulated by being fed, fasted, and dosed with the pyruvate dehydrogenase kinase (PDK) inhibitor dichloroacetate (DCA). A significant difference in apparent PDH flux was observed in these three circumstances, in accordance with the known biochemistry. The sequence was prospectively cardiac gated and timestamped by a custom hardware device, which prevented an error in the time axis arising from variability in the R-R interval of the anaesthetised rat. This gating reduces the temporal resolution from 780 ms to ~ 1.8 s per 3D image, although the resulting TR is not constant. The proposed sequence is believed to represent a good compromise between spectral, spatial and temporal resolution, and, as demonstrated, is able to effectively image metabolic dynamics in the rodent heart.

6.2 | *Methods*

6.2.1 *Spectral-Spatial Pulse Design*

As detailed in Fig. 6.1**A**, a 9 ms long pulse was designed under the RF excitation k-space formalism with zipper correction as described previously in Chapter 5. The time-bandwidth product used was 3 in the spectral domain and 9 in the spatial domain, and the RF amplitude was set to zero on negative gradient lobes to produce a fly-back pulse. This design excites a large slab of magnetisation with minimal ($< 1\%$) spectral ripple, and avoids problems of spurious excitation at 0 Hz by having a designed immunity to the k-space sample staggering described in Section 5.3. The pulse was slew-rate (S_{max}) limited rather than G_{max} limited. This reflects the prediction from the GIRF (c.f. Section 3.1.3) that excitation waveforms played at S_{max} rather than at G_{max} would be reproduced with significantly better fidelity than those running at G_{max}. The pulse had a sublobe duration of 246 μs and 37 sublobes, corresponding to an excitation bandwidth of 240 Hz, and a stop-band of $\sim$2 kHz. As shown, this was sufficient at 7 T to avoid every resonance in the spectrum of hyperpolarised [1-^{13}C]pyruvate, denoted as white vertical lines in Fig. 6.1**B**. The excitation bandwidth was sufficiently large (3 ppm) to provide relative immunity to transmitting off-resonance and B_0 inhomogeneity, improving *in vivo* utility. The minimum excitation

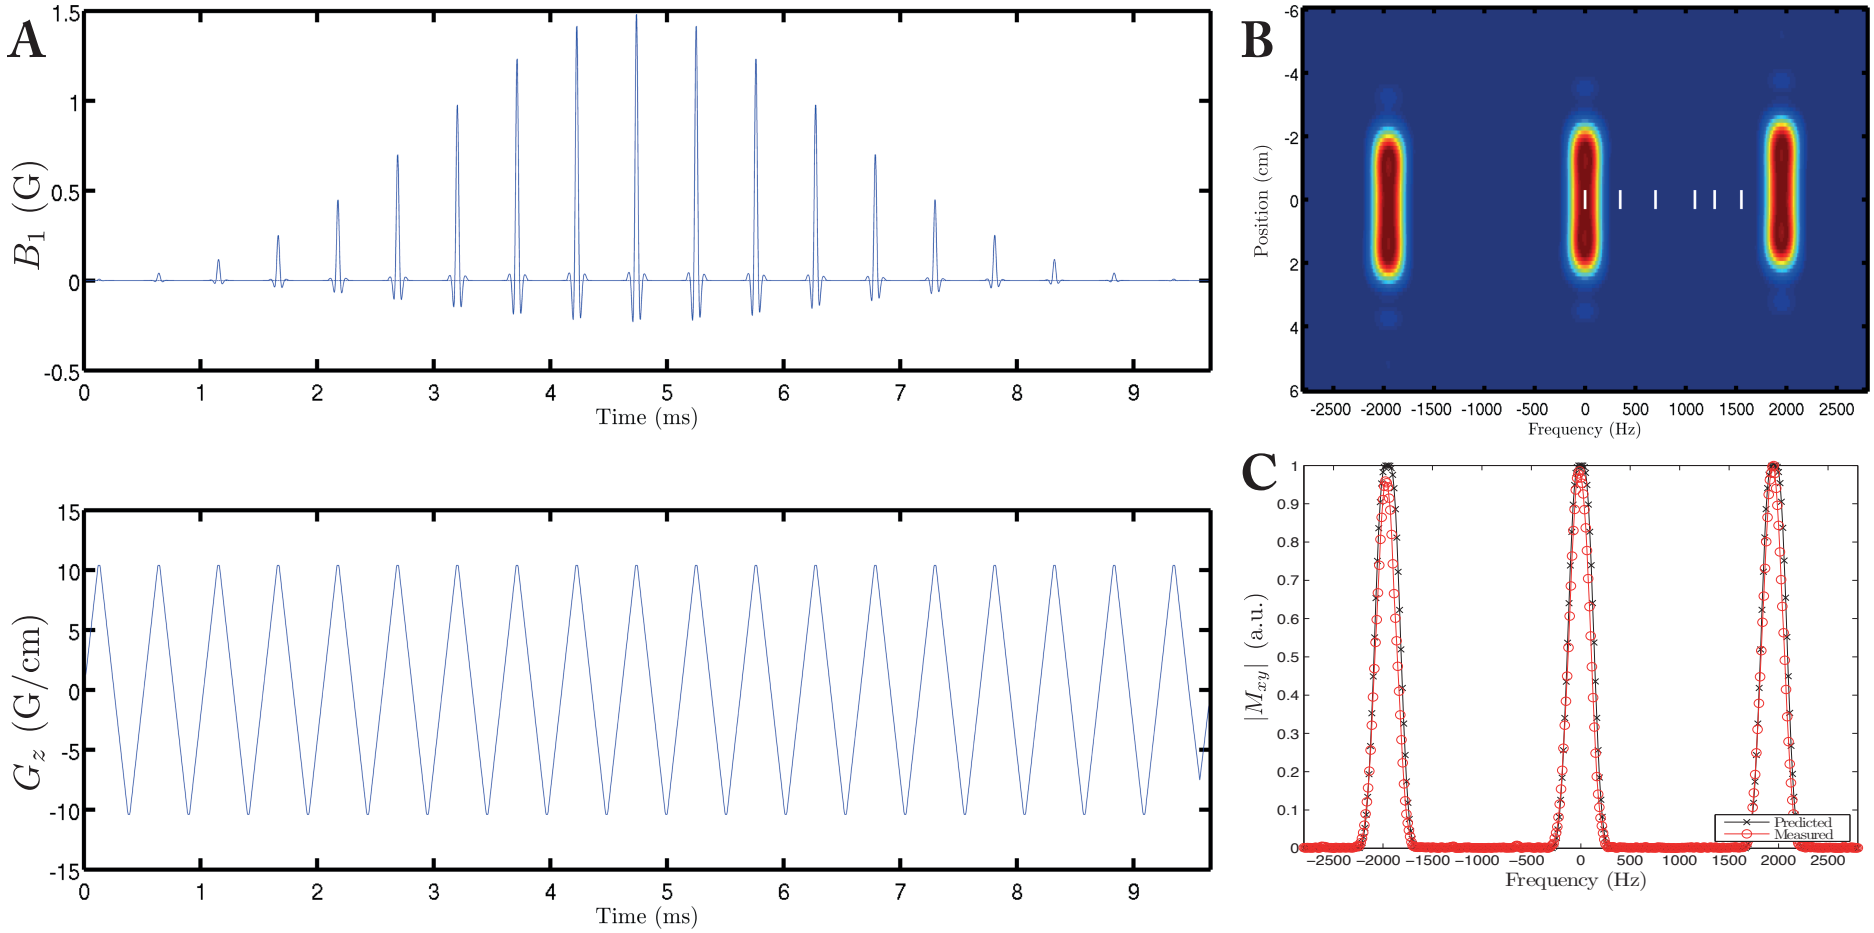

Figure 6.1: The spectral-spatial pulse designed for use in this study (**A**) running at $S_{\max} = 87.5$ G/cm/ms has an excitation bandwidth of ~ 250 Hz (FWHM, **B**) and a stop-band of 2 kHz. The location of the visible resonances following injection of [1-^{13}C]pyruvate are shown as white bars, and are, in order, bicarbonate, urea (a thermal equilibrium sample of which was used as a frequency reference), pyruvate, alanine, pyruvate hydrate, lactate. The measured frequency profile of the pulse agreed with that predicted (**C**).

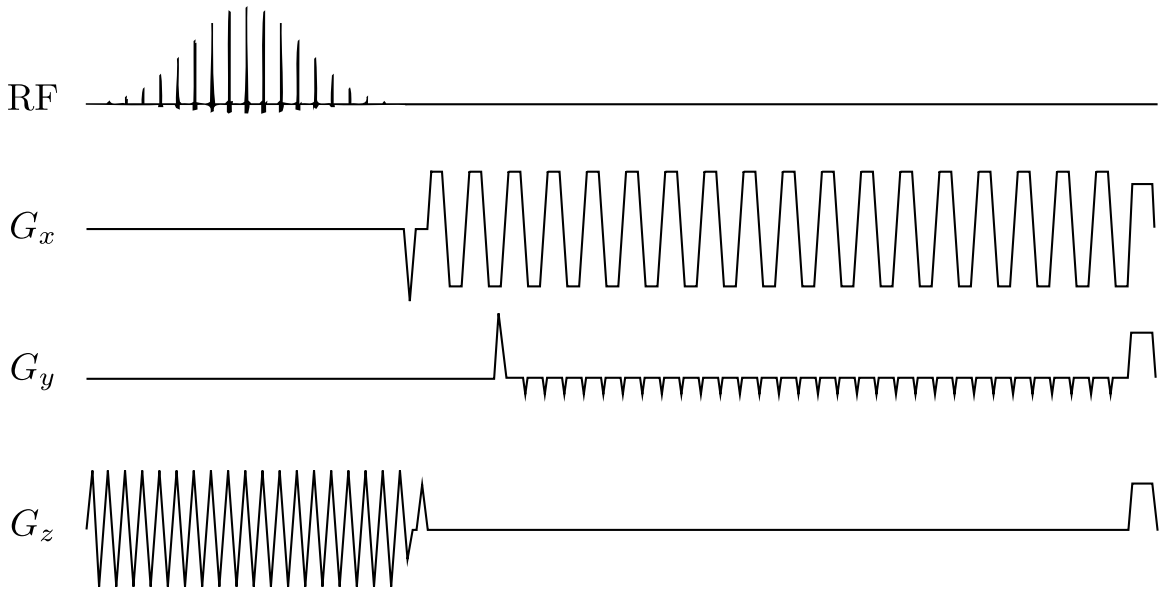

Figure 6.2: The pulse sequence used in this study, consisting of the spectral-spatial pulse discussed previously followed by z-ordered phase encoding and a Cartesian xy EPI readout with two navigators.

slab thickness was obtained when operating at $S_{\max}$, which corresponded to $G_z = 10$ G/cm and a slab thickness of 45.5 mm. The spectral excitation profile of the pulse was measured on a spherical phantom under the designed conditions: the implementation matched the design (Fig. 6.1**C**). This condition represents a significant improvement on the corresponding multi-slice approach (c.f. Fig. 5.6**A**).

6.2.2 Pulse Sequence

After slab excitation with the above spectral-spatial pulse, phase encoding with centric-ordered shots was then used to trade temporal resolution for z-resolution. An EPI readout with three navigators followed in plane (Fig. 6.2; EPI readout parameters were the same as the revised version described in Section 5.4.). The matrix size was $32 \times 32 \times 12$ with a FOV of $64 \times 64 \times 45.5$ mm^3, for an acquisition resolution of $2 \times 2 \times 3.8$ mm^3; TE $= 16.34$ ms; minimum TR $= 70$ ms (duty cycle limited; sequence length $= 31.48$ ms). As illustrated in Fig. 6.3, excitation was performed sequentially across metabolites in the order

$$ {}_1X_1^1,\ {}_1X_2^1,\ \ldots,\ {}_1X_N^1,\ {}_2X_1^2,\ \ldots,\ {}_2X_N^2,\ {}_3X_1^3,\ \ldots,\ {}_3X_N^3,\ {}_1X_1^4,\ \ldots \tag{6.1} $$

where ${}_\alpha X_i^j$ denotes time-point j and phase encoding step $i = 1 \ldots N$ for metabolite ${}_\alpha X$. Here, $\{{}_1X,\ {}_2X,\ {}_3X\}$ = [Pyr, Bic, Lac], and $N = 12$. The sequence was ECG-gated, and the polarity of the in-plane phase encoding (y) axis was alternated on successive acquisitions of the same metabolite. The flip angle per shot was chosen to be $\theta = 3^\circ$ for $X =$ Pyr, and $\theta = 20^\circ$ for $X \neq$ Pyr, such that the overall flip angle for each 3D image is $\vartheta = \arccos(\cos^N \theta) = 17^\circ$ for $X =$ Pyr and 61° for $X \neq$ Pyr.

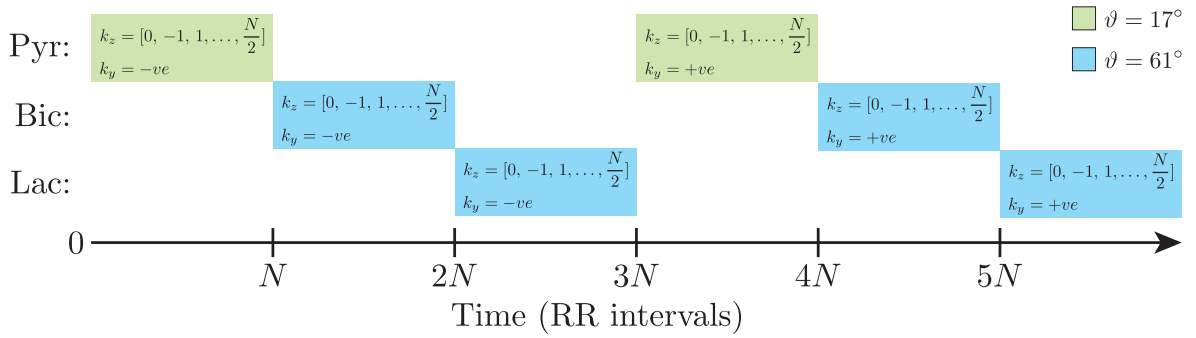

Figure 6.3: A cartoon of the acquisition strategy. Each imaging block consists of N phase encoding steps in the z direction, which are ordered centrically. One phase encoding block was obtained per R-R interval. Three metabolites are acquired in sequential order: pyruvate (Pyr), bicarbonate (Bic) and lactate (Lac). The overall flip angle for each 3D image was ϑ. The in-plane blipped EPI phase encoding axis, k_y, was alternated on successive acquisitions of the same metabolite.

Three dimensional, cardiac-gated volume shimming was performed using the previously mapped behaviour of the magnet's shim set and a spherical harmonic cardiac auto-shimming algorithm as described previously in Section 3.4. Multi-echo cardiac gated gradient echo B_0 maps were acquired (48×48 matrix; 70×70 mm FOV; 1.2 mm slice; TR $= 100$ ms; TE $= 3.06, 3.56$ ms, $T_{read} = 1$ ms; total scan time 60 s) to verify homogeneity over the whole volume of acquisition. As before, the ^{13}C centre frequency $\omega_{transmit}$ was set by identifying the frequency of the urea phantom ω_{urea}, and compensating for any small measured B_0 variation between the urea phantom and the heart ΔB_0 as $\omega_{transmit} = \omega_{urea} + \Delta B_0 \gamma_{^{13}C}/\gamma_{^{1}H}$.

Anatomical images were acquired with a segmented gradient echo CINE sequence described previously,[143] with the same slice thickness and FOV as the hyperpolarised data (TR $= 4.6$, TE $= 1.41$ ms, 28 frames, FA $= 25°$).

6.2.3 *Hardware*

Experiments were performed on the Varian 7 T system described in Chapter 3, using the actively detuned transmit/surface receive setup consisting of a 72 mm dual-tuned $^{1}H/^{13}C$ proton/carbon birdcage volume coil with a 40 mm two-channel ^{13}C surface receive array with an integrated preamp. As documented in Section 3.2.5, the home-made combined RF shield and animal handling system was required to ameliorate undesired electrostatic standing wave interactions occurring inside the magnet bore. The volume coil ^{13}C flip-angle calibration was performed on a cylindrical reference phantom consisting of 500 mL of 1 M natural abundance urea with thermal equilibrium slice-selective spectroscopy. It was found that the position and width of the resonant frequency of the volume coil did not change appreciably between loading with the phantom and with a rat, indicating that the delivered carbon flip angle was as calibrated. This was periodically monitored by performing a flip angle calibration on the 5 M urea reference, and assuming that B_1^+ was homogeneous. As discussed in

Section 3.2.5 and Fig. 3.10, this B_1^+ homogeneity of the volume coil was assessed by acquiring a multiple flip angle B_1 map on the urea phantom with TR at least $5T_1$. The spatial variation of flip angle was less than 9%, which was believed to be dominated by noise in the acquired data.

To allow for the reconstruction of an accurate time axis, it was necessary to create a simple hardware device (documented in Section 3.3.2) that recorded when each ECG-gated scan occurred to a resolution of at least 10 μs. This high temporal resolution ensured that the error on recording the time at which each shot occurs could be neglected. Scans were gated to the R-wave. Respiratory gating was not performed, as no artefacts were observed that were believed to be due to respiratory motion. The period of respiratory motion (c. ~ 1.5 – 2 s in the anaesthetised rat) was comparable to the duration of each acquisition.

6.2.4 Frequency correction

6.2.4.1 Introduction

As described in Section 5.5.1, the small phase encoding equivalent bandwidth of an EPI readout train serves to make EPI readouts very susceptible to artefacts arising from off-resonance effects. In particular, a small uncertainty in the centre frequency of the excited metabolite is highly probable in hyperpolarised experiments, as the centre frequency is not known until a dissolution is performed and the hyperpolarised compound injected. The resulting small spatial shifts in image space would preclude the co-registration of hyperpolarised images and anatomy unless a method for estimating this frequency offset could be determined. If B_0 is not homogeneous across the image, these spatial shifts become a function of space, resulting in non-linear distortions and artefacts in the image.

The fact that these issues are common to EPI readouts has inspired several methods to correct for them in both image[144] and k-space.[145] Arguably the simplest and most popular method to correct for EPI distortions is that proposed by Andersson *et al.*[142] in which the direction of the phase encoding axis is alternated between successive acquisitions. If B_0 is homogeneous, but the centre transmitter frequency is set incorrectly, the resulting pair of nominally identical images will be shifted by the same amount in different directions. Co-registering the same images (by shifting them in image space until they overlap) allows for not only the estimation of the error in transmitter frequency, but also the use of the same data with this correction.

If B_0 is not homogeneous, then the signal observed over time $S(t)$ varies as

$$S(t) \propto \iint_{x,\,y} \rho(x,\,y) e^{i\gamma(\Delta B(x,y)+g_f(x,y,t)+g_p(x,y,t))}\,\mathrm{d}x\,\mathrm{d}y, \tag{6.2}$$

which can be expressed as a linear equation

$$\mathbf{S} = \mathbf{A}\rho, \tag{6.3}$$

where ρ is an image, g_f and g_p are the time-integrals of the frequency and phase-encoding gradients G_x and G_y, usually chosen such that $\mathbf{A}$ is the Fourier matrix if ΔB is zero everywhere. The key insight of Andersson *et al.* is to factor Eq. (6.3) as the product of an explicitly 'on resonance' Fourier matrix $\mathbf{F}$ and $\mathbf{A}$:

$$\mathbf{s} = \mathbf{F}^\dagger \mathbf{A}\rho. \tag{6.4}$$

Accordingly, the 'true' undistorted image could be found, if ΔB is known, by

$$\begin{aligned}\hat{\rho} &= \left(\mathbf{F}^\dagger \mathbf{A}\right)^{-1} \mathbf{s} \\ &\stackrel{\text{def}}{=} \mathbf{K}^{-1}\mathbf{s},\end{aligned} \tag{6.5}$$

resulting in an undistorted image $\hat{\rho}$. However, even given an independent measurement of ΔB, the inverse of $\mathbf{K}$ is not guaranteed to exist, and $\mathbf{K}$ may be very poorly conditioned. In particular, if large gradients in the ΔB map coincide with non-zero spatial gradients in image space, the condition number of $\mathbf{K}$ is effectively infinite and the inversion problem can not be solved.

However, if the direction of k-space traversal is alternated, two separate $\mathbf{K}$ are sampled, say, $\mathbf{K}_+$ and $\mathbf{K}_-$ for two different datasets $\mathbf{s}_+$ and $\mathbf{s}_-$. As the 'true' image $\hat{\rho}$ should be independent of the direction in which k-space is traversed, one can write down that

$$\begin{pmatrix}\mathbf{s}_+ \\ \mathbf{s}_-\end{pmatrix} = \begin{pmatrix}\mathbf{K}_+ \\ \mathbf{K}_-\end{pmatrix} \hat{\rho} \tag{6.6}$$

$$\text{And hence } \hat{\rho} = \left(\left(\mathbf{K}_+^\top \mathbf{K}_-^\top\right) \begin{pmatrix}\mathbf{K}_+ \\ \mathbf{K}_-\end{pmatrix} \right)^{-1} \left(\mathbf{K}_+^\top \mathbf{K}_-^\top\right) \begin{pmatrix}\mathbf{s}_+ \\ \mathbf{s}_-\end{pmatrix}. \tag{6.7}$$

The properties of $\mathbf{K}$ ensure that this inversion problem is well defined, and it is possible to calculate a (pseudo-)inverse to satisfy Eq. (6.7). Accordingly, the non-linear compression and expansion resulting from an imperfect EPI acquisition can be corrected for *post hoc* if such data are acquired. This approach is routinely used in clinical fast imaging readouts, and is implemented in the FMRIB Software Library `fsl` by the module called `topup`.[146,147]

6.2.4.2 *Application to hyperpolarised* ^{13}C

It is difficult to apply the preceding method directly to hyperpolarised imaging, as we do not expect $\mathbf{s}_+$ and $\mathbf{s}_-$ to reconstruct the same image $\hat{\rho}$ because $\hat{\rho}$ inherently changes as a function of time. Directly

applying `topup` to multidimensional hyperpolarised datasets therefore does not appropriately correct for distortions as discussed above; instead, it rather generates a large set of additional ones. As realised by Cunningham *et al.*, it is substantially easier to use image domain metrics to try and correct for bulk frequency shifts alone.[23] Alternating sets of images can therefore be aligned by gross features which change slowly over time, through finding the stationary points of appropriate metrics of image similarity, such as image mutual information (as proposed by Cunningham), image self-similarity (as proposed by the author) or simpler metrics, e.g. minimising the least-squares difference between the 'blip up' (s+) and 'blip down' (s−) image stacks. A value for the frequency offset can then be estimated by co-registering each slice. The value derived should be consistent with that estimated from B_0 maps and knowledge of the centre transmitter frequency.

6.2.5 *Reconstruction*

Data were regridded appropriately in Matlab, and linear phase correction performed in-plane on the EPI readout. Alternate lines of k-space were then Fourier transformed independently, and added in magnitude in image space; halving the FOV and eliminating the ghost. This process was performed to avoid edge cases where Nyquist ghost reconstruction failed; all data were treated equally. The noise decorrelation matrix for the two channel array was calculated following the method of Kellmen and McVeigh,[148] and the data 'pre-whitened' to ensure that the noise in each channel was independent. Images were combined and coil sensitivities estimated by the method of McKenzie.[149] The source code used to perform this noise equalisation and parallel imaging reconstruction was obtained freely from the 2013 ISMRM Educational Weekends Sunrise Parallel Imaging session.[1]

The image alignment method of Cunningham *et al.*[23] was implemented. Blip-reversed EPI data were spatially co-localised by both a mutual information maximisation method, manually, and by the image self-similarity metric discussed below.

Mutual information (MI) is a metric commonly used in information theory to quantify the mutual dependence of two sets. It is defined as the sum of image entropy in both images, minus the joint entropy of the two images, i.e. as

$$\begin{aligned}\mathrm{MI}(x,\,y) = -\Bigg(&\int_{\text{Image 1}} p(x)\log_2 p(x)\;\mathrm{d}x + \int_{\text{Image 2}} p(y)\log_2 p(y)\;\mathrm{d}y \\ &- \iint_{\text{Image 1 and 2}} p(x,\,y)\log_2 p(x,\,y)\;\mathrm{d}x\,\mathrm{d}y\Bigg),\end{aligned} \tag{6.8}$$

where $p(x)$ is the probability associated of having an intensity value x for a given image, i.e. the histogram count for that intensity level in a given image. The joint entropy term $\iint p(x,y)\log_2 p(x,\,y)\;\mathrm{d}x\,\mathrm{d}y$ is therefore the joint histogram between the two slices being registered. This provides a measure of

[1] And is available online at `http://hansenms.github.io/sunrise/sunrise2013/hansen_exercise/ismrm_hansen_practical_session.html`

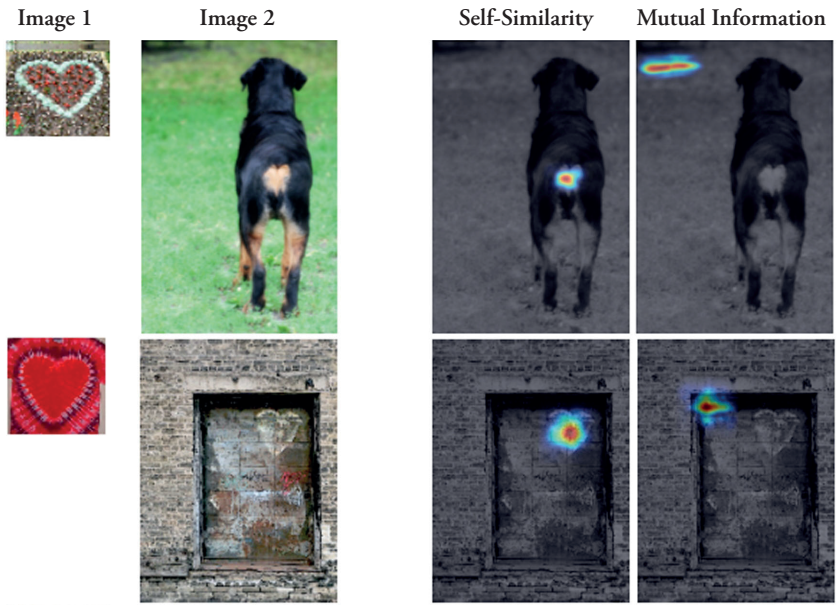

Figure 6.4: A comparison of self similarity and mutual information as image co-registration metrics. Given a target input – here, a cartoon shape of a heart – target densities are shown overlaid in other images where heart shapes are believed to lie. Self-similarity is shown to outperform mutual information. Figure adapted from Shechtman and Iraini.[150]

overall feature density that is penalised if the joint image is significantly 'busier' (having a greater feature density) than each image individually. This metric is believed to not necessarily be as appropriate if the two images being registered are substantially different as the term $\int\int p(x,y)\log_2 p(x,\,y)\;\mathrm{d}x\,\mathrm{d}y$ will independently change.

In contrast, image self-similarity is a statistical score based on maximising generalised correlations between images independent of features such as colour or local changes. Self-similarity is a more generalised metric, and is able to accurately co-register images that do *not* share common image features (such as textures or edges), but *do* share a similar geometric layout of features. It is defined as the binned log-polar transform of S_q, where

$$S_q = \exp\left(-D_q(x,y)/\max(\mathrm{var}_{\mathrm{noise}},\,\mathrm{var}_{\mathrm{auto}}(q))\right) \tag{6.9}$$

represents a normalised correlation surface between the image at point q and the reference image. In Eq. (6.9), D_q is the weighted sum-of-squares difference between the reference image and that being registered at a location q; $\mathrm{var}_{\mathrm{noise}}$ is a constant estimating the 'acceptable' noise level in the image; and $\mathrm{var}_{\mathrm{auto}}(q)$ is the maximal variance of all patches within a small region of q. In this fashion, sharp edges are more tolerable to pattern variations than smooth regions. Self-similarity represents a powerful generalised difference metric, and is, for example, able to accurately co-register the shape of a heart in several different forms, as illustrated in Fig. 6.4.[150] This versatility comes at a cost, however: self-similarity, as an iterated local metric, is approximately two orders of magnitude more computationally expensive than mutual information. Here, the self-similarity implementation provided by Chatfield *et*

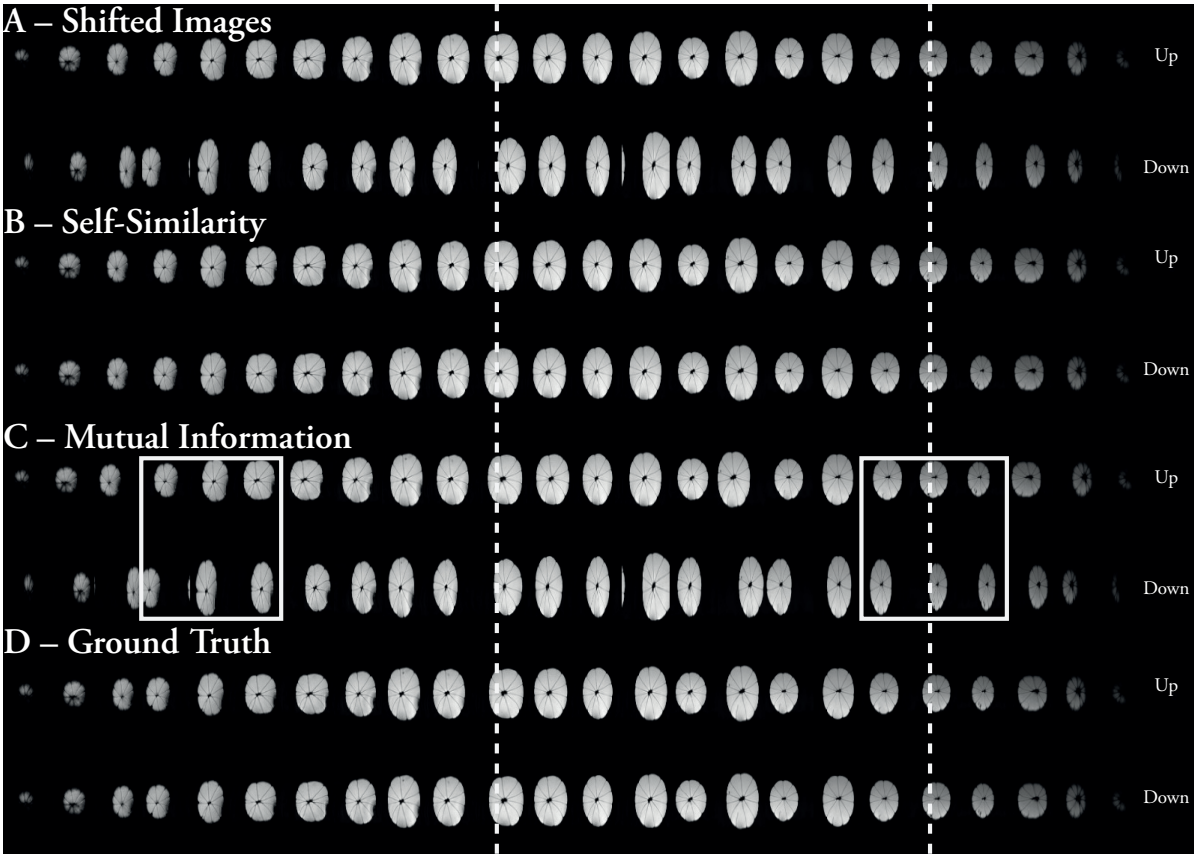

Figure 6.5: A comparison of the co-registration metrics used to estimate off-resonance shifts in the acquired blip up ("up") and blip down ("down") EPI datasets. **A**: This computational experiment used distorted, shifted images of an object (a lemon, with z concatenated from left to right) which were co-registered using either the self-similarity metric (**B**), image mutual information (**C**) or a number of other, less-successful methods. Compared to the ground truth (**D**), image self-similarity was able to successfully estimate the correct feature shift, even for distorted images. Two vertical lines highlight shifting features between the different methods. Regions where mutual information has converged to the incorrect extremum are highlighted with a box.

al.[151] was used.[2]

For co-registering distorted EPI data, the mutual information metric for both blip-up and blip-down images was computed after pre-filtered each image with either a (2D) Gaussian or a Chebyshev kernel, blurring them to improve the co-registration. It was, however, found that the mutual information metric was not always the most appropriate measure of image alignment. This error is illustrated graphically in the phantom experiment summarised in Fig. 6.5, in which images of an object (a lemon) that were distorted to mimic EPI tearing artefacts were shifted in the phase-encoding axis by a random amount (distributed as $p \sim N(0, 15)$ pixels on each slice), producing the 'blip up', 'blip down' dataset shown in Fig. 6.5**A**. The co-registration algorithm written by the author based on that proposed by Cunningham *et al.* was then used to register the images and estimate the frequency shift corresponding to p. A number of metrics were evaluated, including image self similarity maximisation (Fig. 6.5**B**) and mutual information maximisation (Fig. 6.5**C**). It was found that, compared to the ground truth reference

[2]This implementation is available from `http://www.robots.ox.ac.uk/~vgg/software/SelfSimilarity/`.

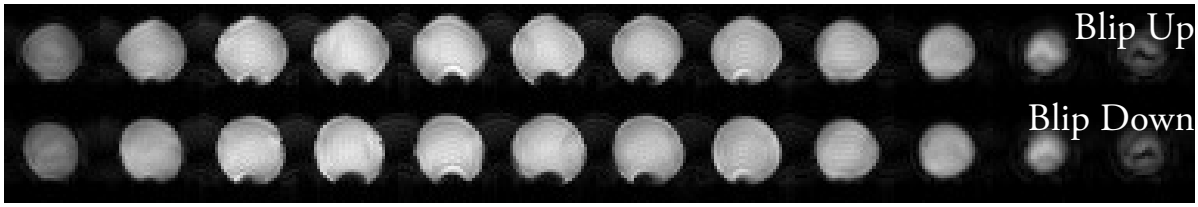

Figure 6.6: A blip up/down reconstructed 3D Spectral-Spatial EPI image stack of a ball of water. z increases from left to right.

(Fig. 6.5**D**), the image self similarity metric proposed by Shechtman and Irani[150] was successfully able to co-register dissimilar stacks of distorted images with greater accuracy than mutual information. This behaviour is reported elsewhere.[150]

When applied to $^{1}\mathrm{H}$ images, this algorithm was successfully able to co-register blip up/down 3D-EPI images of a phantom (a plastic ball containing water), as shown in Fig. 6.6. This also demonstrates the Nyquist ghost removal algorithm used to reconstruct the data. T_2^* artefacts are visible, in particular around a small bubble of air (reconstructed here at the bottom of each image), and at the last slice position. This proton image is shown without zero-filling.

If the registration failed, estimates for the requisite frequency shift were provided manually. Attempts at directly computing $\mathbf{K}$ as defined in Section 6.2.4 unfortunately were not successful. It was believed that this failure originated from the comparatively low SNR of hyperpolarised images, and their motion over time. Future work will aim to measure $\mathbf{K}$ from fat-saturated proton EPI images, and then apply the resulting correction matrix to acquired hyperpolarised data.

The resulting five-dimensional image stacks were zero-filled by a factor of two in all spatial directions, producing an image stack with a final in-plane resolution of $1 \times 1 \times 1.9$ mm^3. To confirm the technique, ~ 4 ml hyperpolarised [1-^{13}C]pyruvate was injected into a 50 mL syringe (outer diameter ~ 30 mm) containing 5 ml water (to enable shimming) and then pyruvate images acquired with a TR of 100 ms without cardiac gating. The sequence was able to resolve the arrival and decay of pyruvate within the syringe, and the temporal sum of the resulting images are shown in Section 6.2.6**Aii** together with a proton overlay in Section 6.2.6**Ai**. The effect of the surface receive coil is apparent, with poor spatial penetration into the back of the comparatively large syringe. It is also possible to resolve the small pipe used to inject pyruvate into the syringe at the extremities of z during the initial time points.

The visible time-course of pyruvate (shown in Section 6.2.6**B** as the total magnitude of the image as a function of time) matches that expected for decay back to thermal equilibrium and with a small degree of RF depletion. Pyruvate is consequently resolved for a several minutes. The image stack, concatenated in both z and t is shown in Section 6.2.6**C**, and represents a baseline for subsequent imaging experiments.

6.2.6 *Animal Handling and Dosing*

To demonstrate the *in vivo* sensitivity to metabolism, the hearts of healthy rodents were scanned in three different metabolic states: fasted, fed, and fed and given the PDK inhibitor DCA. All experiments were performed in accordance with relevant UK legislation (with personal, project and institutional licences granted under the Animals (Scientific Proceedures) Act 1986), and were subject to local ethical review and an independent cost-benefit analysis.

Approximately 40 mg of [1-^{13}C]pyruvate doped with 15 mM OX063 trityl radical and 3 μL of a 1 in 50 dilution of Dotarem (Guerbet Laboratories Ltd) gadolinium chelate was polarised and dissoluted in a prototype polariser as described previously in Section 3.5. The estimated liquid state polarisation was $\sim$32 % after dissolution, and two millilitres of the resulting 80 mM pyruvate solution were injected manually over a 20 s period through a previously placed tail vein cannula in anaesthetised male Wistar rats (2% isoflurane in 90% O_2, 10% N_2O; $n = 4$; body weight 410 ± 70 g). Pyruvate was injected approximately 10 s after dissolution, and its pH was measured after each dissolution to ensure correct neutralisation in all experiments. The scan was started prior to the dissolution of hyperpolarised pyruvate.

Four healthy rats were scanned on three separate occasions with at least two days as a recovery period between scans. Animals were imaged three times: (1) in the fed state after being allowed food and water *ad libitum*; (2) after being fasted for at least eight hours prior to imaging; and (3) following injection with the PDK inhibitor, DCA, in the fed state. Rats were dosed with DCA at approximately 30 mg/kg as previously proposed,[152] 30 minutes before hyperpolarised imaging. To minimise the total volume of injected agents, DCA was administered as a 100 μL bolus of neutralised 10% DCA solution at pH 7.2-7.8. Data were then reconstructed as above, and overlaid on anatomic images for reference. A manually segmented ROI was drawn in the lumen (pyruvate) or myocardium (bicarbonate, lactate) and the ratio of maximum bicarbonate to maximum pyruvate computed after subtracting the mean noise as measured from the last ten time points of the experiment.

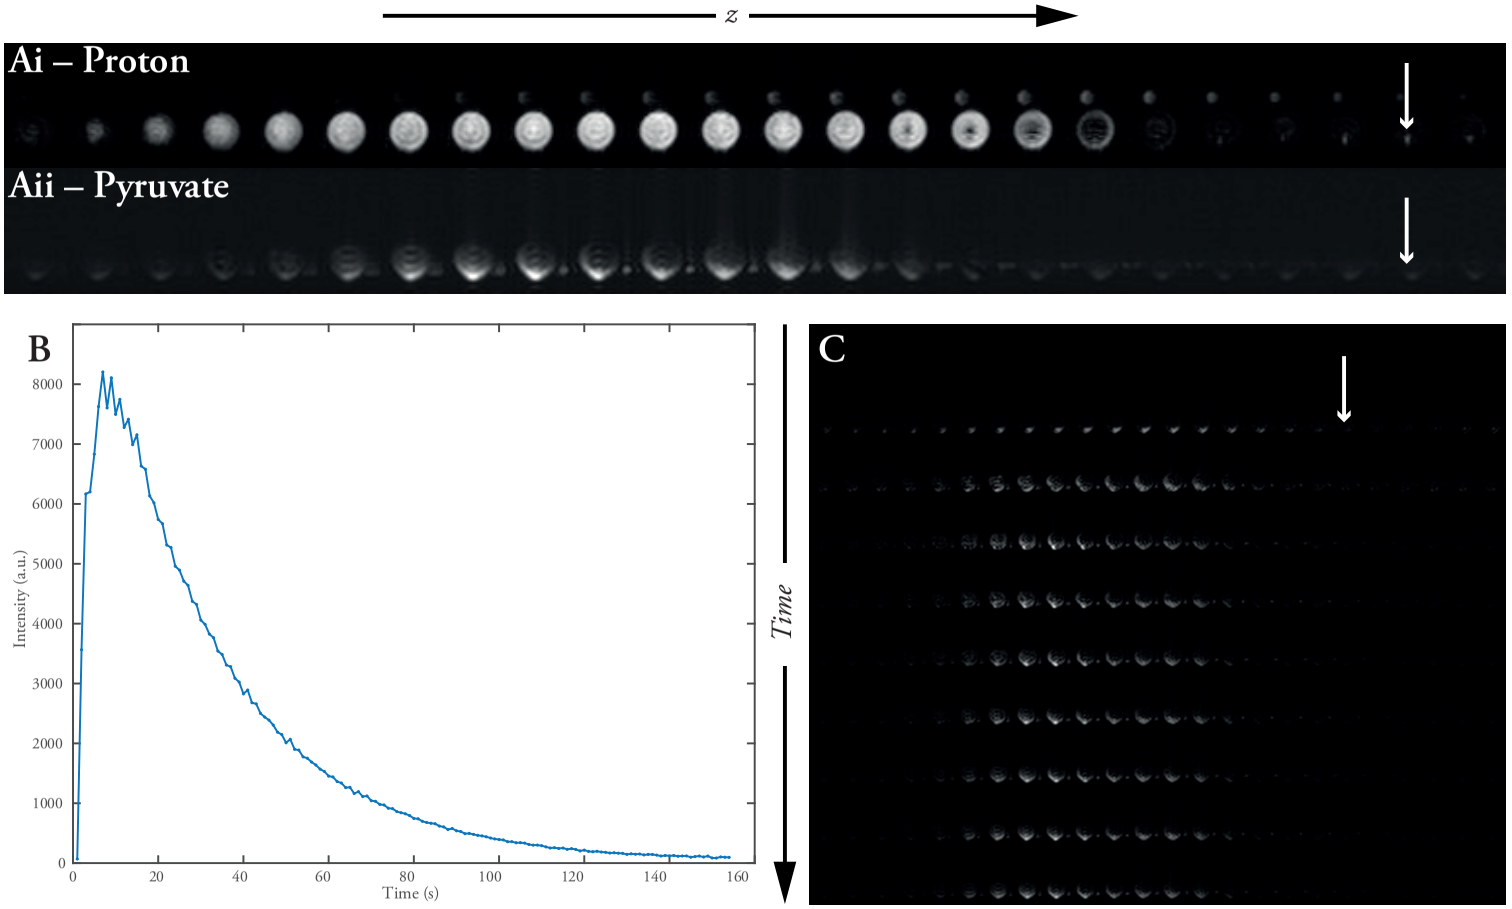

Figure 6.7: The developed sequence imaging hyperpolarised [1-^{13}C]pyruvate in a tube of water. **Ai** Proton overlay of the tube of water, for comparison with **Aii**, the hyperpolarised images. The profile of the surface receive coil used dominates the image obtained. **B** However, the efficient imaging strategy allows pyruvate to be resolved for several minutes following injection. This data is shown concatenated in **C**. The small pipe into the syringe is shown on ^{1}H and ^{13}C images by a white arrow.

6.3 | *Validation Study Results*

The proposed sequence was able to resolve the metabolism of pre-polarised [1-^{13}C]pyruvate in the in vivo rat heart with a reconstructed spatio-temporal resolution of $1 \times 1 \times 1.9$ mm^3 and ~ 1.8 s. The three-dimensional sequence reports on hyperpolarised metabolism in a $64 \times 64 \times 45.5$ mm^3 FOV, which was sufficient to encompass the whole heart and the liver. The large field of view was advantageous as we obtained information from multiple organs at the same point in time. The data were five-dimensional and an example maximum intensity projection (summed over time) showing the acquisition volume is illustrated in Fig. 6.11.

Alternating the EPI blip direction successfully allowed for the accurate correction of bulk frequency shifts, and for the co-registered data to be overlaid with anatomical CINE images, illustrating myocardial metabolism (*cf*. Fig. 6.8). The video time-courses of this data are available online as supporting information with Miller *et al*.[153] The B_1 homogeneity arising from the volume transmit/surface receive hardware used was sufficient to ensure a uniform excitation across the volume of interest, ensuring that the flip angle delivered was spatially constant. This homogeneity should be contrasted to the spatially non-uniform flip angle arising from a surface transmit/receive coil. However, due to the sensitivity profile of the surface receive coil used, it was not always possible to resolve segments of the myocardium distal from the chest wall. Across the twelve scans, the distal wall was completely resolved twice, partially resolved four times, and not resolved six times. An illustrative example of each of these three gross classifications is provided in Fig. 6.9. The variation between these three situations was believed to be dominated by inter-experiment variability in the relative placement of the animal and the coil used.

The temporal resolution of the technique was sufficient to resolve metabolic dynamics following injection of hyperpolarised [1-^{13}C]pyruvate in a single voxel. Concomitant with the known anatomy and physiology, pyruvate was observed perfusing the great vessels at early times, before arriving in the right and then left ventricles. Myocardial bicarbonate signal and lactate production signal subsequently followed. Summing the data acquired in a manually segmented ROI over the heart allowed for the reconstruction of metabolic dynamic time-courses (Fig. 6.12), as expected. The SNR of the acquired data was such that metabolic time-courses could be reconstructed in any cardiac voxel, if required.

A distribution of acquisition TRs was observed (Fig. 6.12**A**), with a large peak at the R-R interval and a small number of shots at $2\times$ R-R interval, reflecting a number of missed triggers that we conjecture are due to respiratory motion. It was noted that the TR of the experiment was not constant, and that the distribution of 3D image TRs was well described by a gamma distribution (Fig. 6.12). The mean of this distribution was *not* equal to the number of z phase encodes, N, multiplied by the R-R interval. The observed coefficient of variation of TR /image was $40.4 \pm 11.4\%$ (distribution $\sigma/\mu \pm$ s.e.m. over 16 scans), reflecting a significant variability in the R-R interval of the anaesthetised rat. This variability vanished when a signal generator was used to provide an input comparable to an ECG into the cardiac

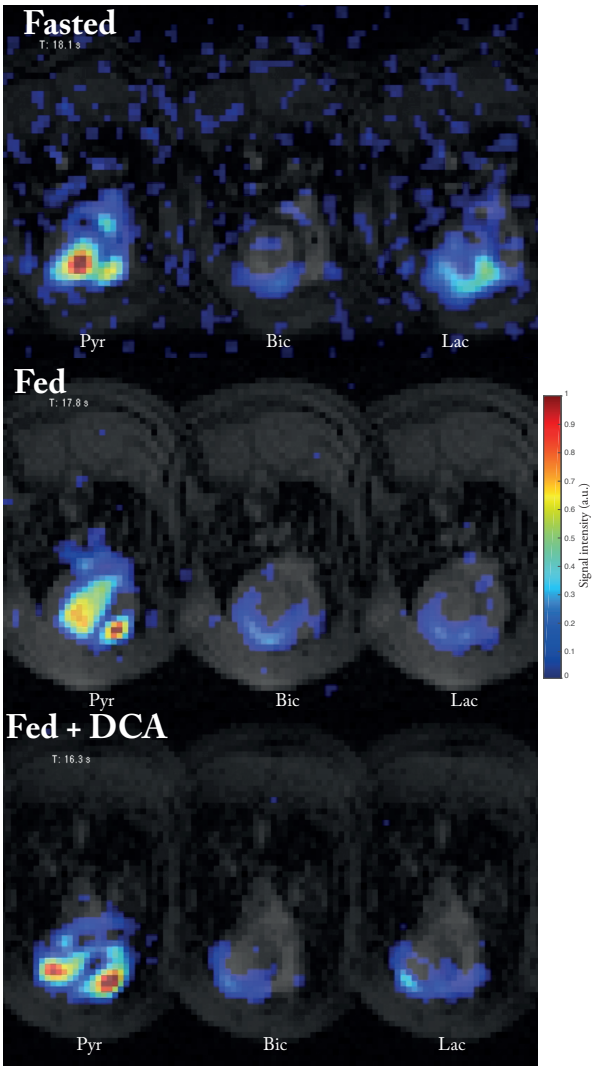

Figure 6.8: Representative mid-ventricular single-slice single time-point data overlaid on anatomical reference images. Shown here are the three metabolites (pyruvate, bicarbonate and lactate) at the timepoint of maximum pyruvate signal, in the three groups of rats. The colour axis for each set of three images is scaled relative to the maximum pyruvate. These data are stills from Supporting Videos S1-3 as provided by Miller *et al.* [153]

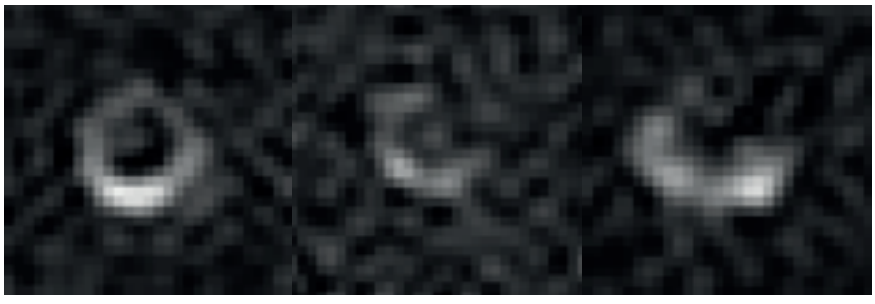

Figure 6.9: Illustrative examples of single-timepoint single-slice hyperpolarised ^{13}C-bicarbonate data, together with the corresponding classification for the degree of resolution of the distal wall of the myocardium. Across the twelve rats scanned in this validation study, the distal wall was completely resolved twice, partially resolved four times, and not resolved six times.

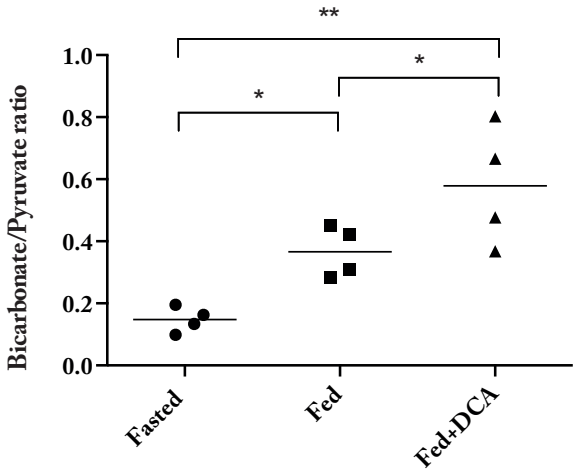

Figure 6.10: The ratio of maximum bicarbonate/pyruvate was significantly different across the three groups of animals scanned, reflecting the known inhibition of PDK by DCA. ($*$ denotes $P < 0.05$; $**$ $P < 10^{-5}$.)

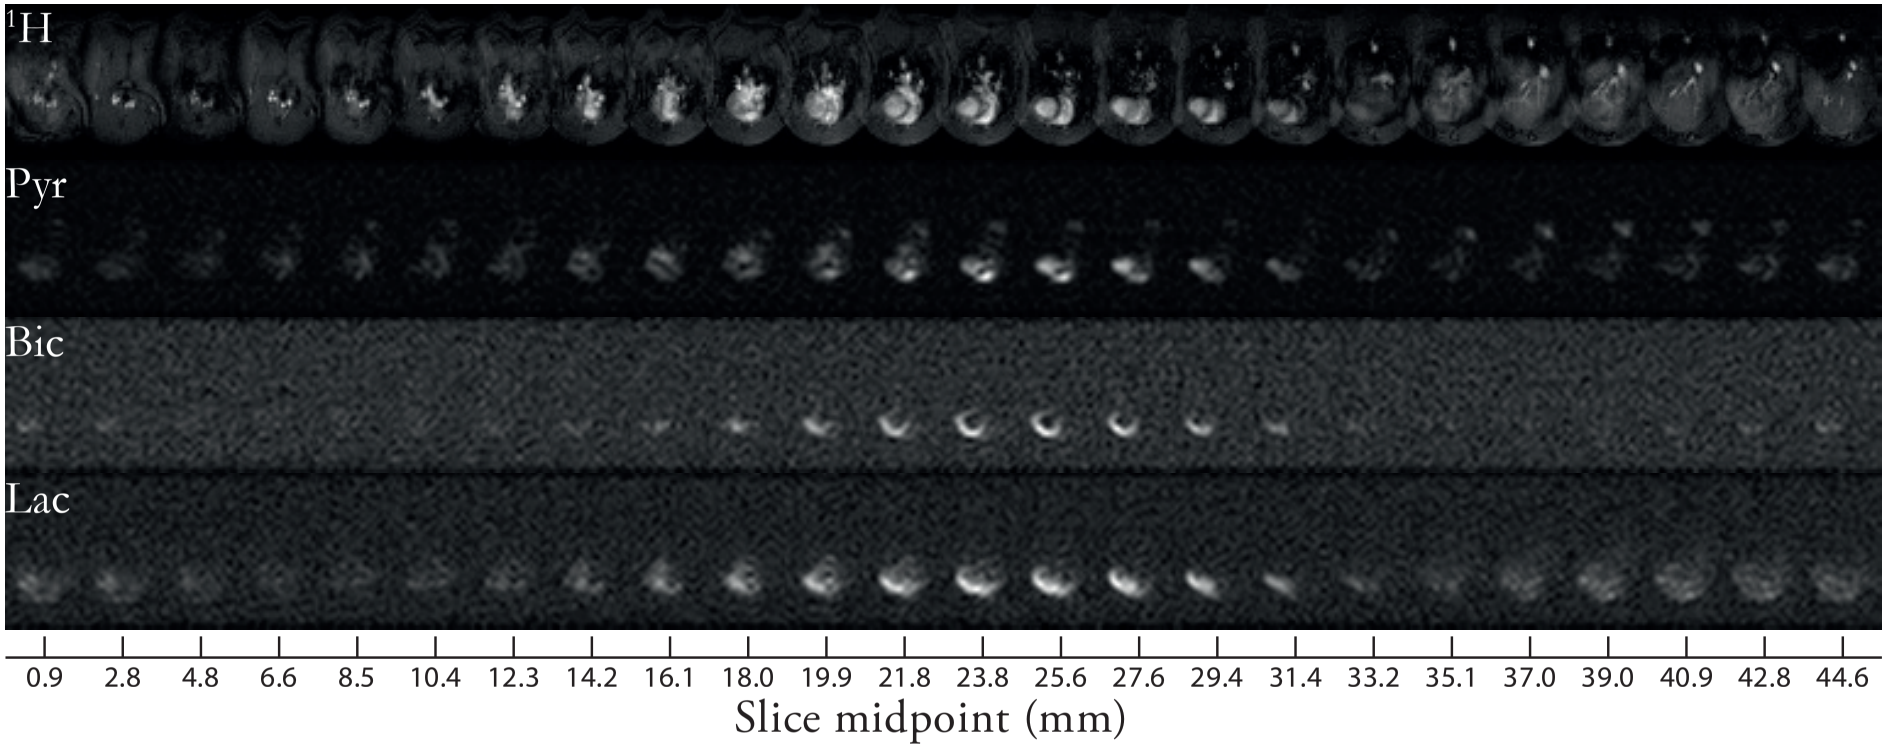

Figure 6.11: Representative axial maximum intensity projections of data acquired from animals fed and dosed with DCA. Images shown concatenated with z decreasing from left to right; the reference CINE images (top) span the neck (top left), heart (top middle) and anterior lobes of the liver (top right). Pyruvate was resolved in the ventricles of the heart, and additionally in the inferior vena cava. The production of hyperpolarised bicarbonate occurred predominantly in the myocardium, and in this instance it was not possible to resolve metabolism in segments distal from the chest wall. Hyperpolarised [1-^{13}C]lactate was also detected after production in the myocardium, and, to a far lesser extent, the liver.

triggering hardware used, demonstrating that this effect is biological in origin and not an experimental artefact. The correct time axis is fundamentally important in hyperpolarised imaging, as it is desirable to create quantitative maps of metabolic rate constants of interest.

Reflecting the known modulation of PDK by both DCA and fasting, significant differences in the ratio of the summed myocardial maximum bicarbonate to maximum pyruvate were observed across the three groups of animals scanned (Fig. 6.10) where Student's t-test $P = 0.03$ between fasted and fed, and fed and DCA; $P = 0.0008$ between fasted and DCA with Bonferroni's *post hoc* correction for multiple comparisons. The administration of DCA serves to approximately double this ratio compared to the fed animals, and quadruple it compared to fasted animals (DCA to fed: 1.6 ± 0.3, DCA to fasted: 3.9 ± 2.8; values ratio means $\pm$ s.e.m. propagated in quadrature). The approximate doubling of PDH flux between fed animals and those dosed with DCA is the same as that reported in the directly analogous experiment of simple un-localised $^{13}\mathrm{C}$ spectroscopy following hyperpolarised [1-$^{13}\mathrm{C}$]pyruvate injection as reported by Atherton *et al*.[152]

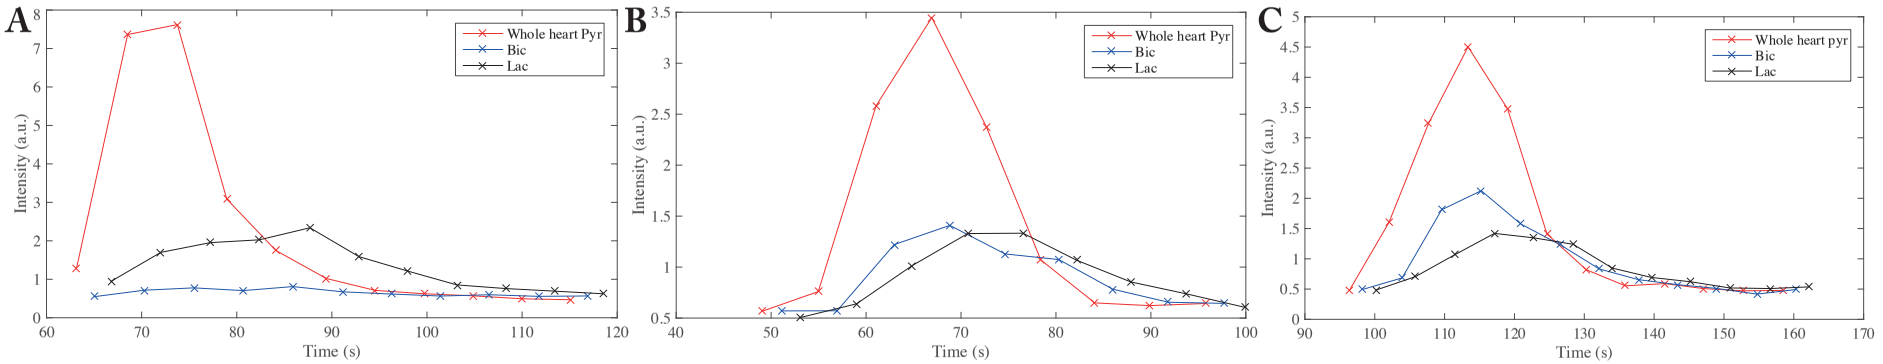

Figure 6.12: Representative single-slice whole-heart time-courses acquired from fasted animals (**A**), fed animals (**B**), and animals fed and administered DCA (**C**).

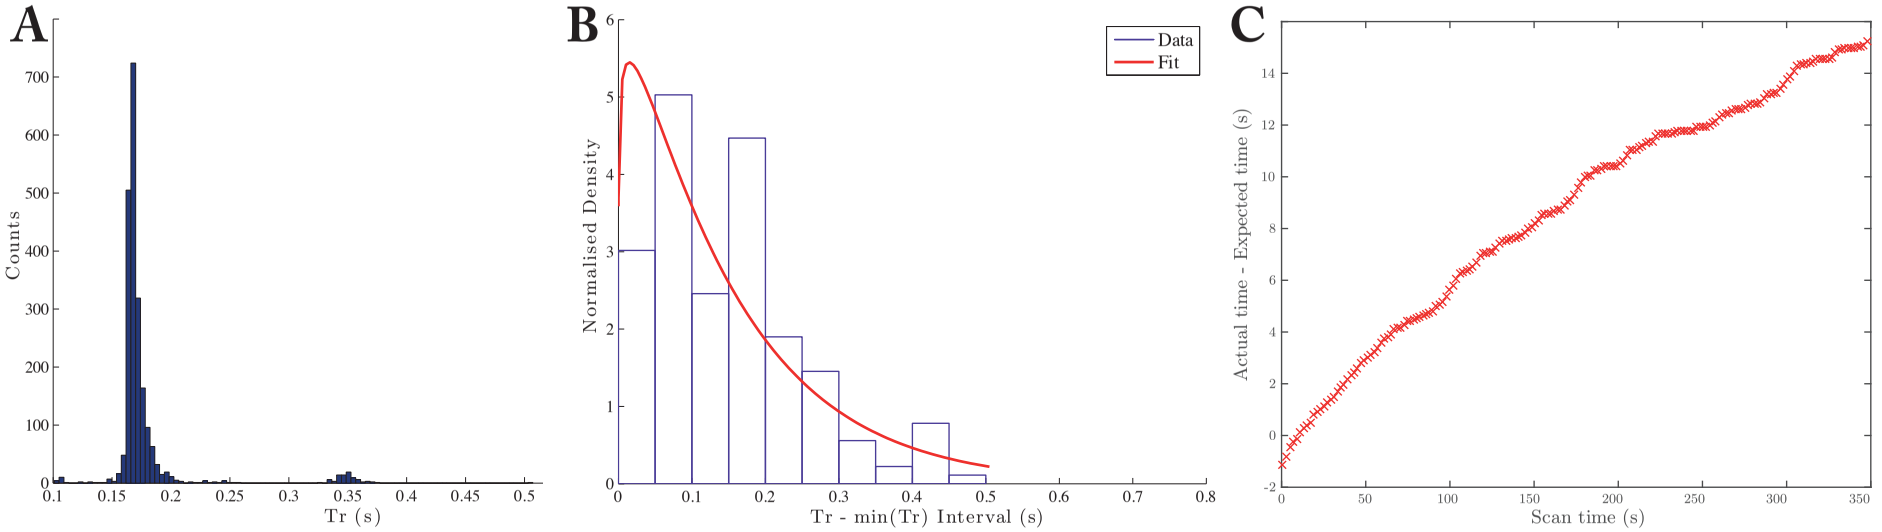

Figure 6.13: **A**: The distributions of gated shot TR was distributed around the R-R interval (here $\approx$ 170 ms), with a small number of missed triggers (at $\approx$ 350 ms). **B**: As expected for a stochastic waiting process, the subsequent distribution of image TR was described by a gamma distribution (fit), with a large spread in values (min/max TR /image = 1.966 s /2.418 s). **C**: The error between the expected ($N\times$ R-R interval) and measured scan times becomes progressively worse as time increases.

6.4 | *Validation Study Discussion*

6.4.1 Overview

As demonstrated above, the novel proposed sequence enables the spectroscopic imaging of hyperpolarised [1-^{13}C]pyruvate and its downstream metabolites in the *in vivo* rat heart. The proposed sequence can be considered robust in the sense that (1) the fly-back 3D spectral-spatial excitation reduces gradient demands, and was played with better fidelity than pulses that excite a thin slab, (2) the excitation bandwidth of the pulse was relatively large (3 ppm), minimising sensitivity to B_0 inhomogeneity, and (3) bulk frequency shifts on the order of 30 Hz that arose from uncertainty of the pyruvate frequency *a priori* could be corrected *post hoc* with the alternating EPI readout and the co-registration algorithm described in Section 6.2.4. The timestamping hardware developed in Section 3.3.2 and used here provided accurate data about the acquisition time of each image. This information was used to compensate for the variability of the heartbeat of the anaesthetised rat.

Cardiac small animal imaging is challenging, as the required resolution and rapid heart rate increases both gradient hardware and SNR demands. Nevertheless, it is desirable to image rodents as they form the predominant animal model of choice for cardiovascular research, and unlike large animals, mice (and increasingly rats) can be readily genetically manipulated to further probe the genetic basis of disease,[154] and form the species of choice to aid in the development of novel pharmaceutical agents, including new hyperpolarised probes of metabolism. Therefore, whilst small animal imaging is challenging, it is of considerable utility owing to the large number of different conditions that could be investigated with the technique. It is also a central tenet of ethical animal experimentation to minimise the degree of discomfort experienced by using a less sentient model species if one is available.[155]

Additionally, small animal hyperpolarised cardiac MR can be considered a 'worst-case' scenario for human imaging, as the required spatial and temporal resolution are in excess of those required for human use. As the total magnetisation available for hyperpolarised imaging is independent of the scanner hardware (instead depending on the polariser used and the quantity of visible carbon injected), the ultimate SNR limited resolution for a given sequence and a given quantity of visible carbon will also be largely independent of subject and hardware. For the reasons outlined above, it is therefore desirable to develop the best possible imaging sequence for capturing hyperpolarised metabolism in the rodent heart, even though scaling arguments predict that this is more challenging than imaging larger animals, including humans.

The proposed sequence is a good compromise between spectral, spatial and temporal resolution. As shown, we were able to spatially resolve metabolism within the myocardium and liver with temporal resolution sufficient to capture metabolic dynamics of interest. The fly-back spectral-spatial excitation effectively represents the limit of spectral undersampling and reduces demands on gradient

hardware. Metabolism was recorded in multiple organs simultaneously, with an acquisition resolution of $2 \times 2 \times 3.8$ mm^3 in a comparatively large FOV that ensures whole heart coverage. Moreover, the large excitation slab reduced flow artefacts that can occur due to pyruvate in the blood flowing into (and out of) a thin excitation slab for 2D multi-slice imaging approaches. In-plane flow ghosting artefacts were not observed, reflecting the comparatively short readout train as EPI readouts are known to be particularly prone to these distortions.[156] The resolution of this sequence compares favourably to other approaches, such as multi-band,[141] IDEAL CSI,[134] high resolution accelerated spiral CSI,[157] spectral-spatial excitation followed by a multi-slice spiral readout,[138] or k-t undersampling following spectral-spatial excitation with an EPI readout,[158] for which the reported resolutions are $2 \times 2 \times 4$ mm^3, $4.4 \times 4.4 \times 10$ mm^3, $2.7 \times 2.7 \times 10$ mm^3, $8.8 \times 8.8 \times 10$ mm^3 and $1 \times 1 \times 8$ mm^3 respectively. Whilst these sequences have been developed on a variety of hardware with different limitations, and utilised on a variety of different animal models of disease, they all fundamentally encounter the SNR limitations of pyruvate M_0 generated in a 3 T pre-clinical hyperpolariser.

Accordingly, an argument can therefore be made that the ultimate achiveable resolution of hyperpolarised pyruvate imaging depends on the quantity of visible carbon injected, which is approximately constant, and the compromises made in the design of sequences to image it. Accordingly, it is believed that the proposed sequence represents an acceptable compromise between spatial, spectral and temporal resolution that will be generally applicable to those wishing to image hyperpolarised metabolism in any organism. The scanner used is a common pre-clinical system, and is far from state-of-the-art. It was found that the fly-back pulse running at S_{max} produced a waveform whose biggest distortion was a constant time offset, which could be corrected by the introduction of a small delay between the RF relative to the gradient waveform, and whose effect otherwise was to slightly degrade the slice profile. This should be contrasted to non-fly-back spectral-spatial pulses, where infidelities cause spurious excitation at 0 Hz, producing unwanted mixed metabolic excitation. The duration of the fly-back pulse is chosen to have a narrow spectral passband and a large stopband, i.e. to avoid exciting pyruvate hydrate and lactate at the same time. The pulse proposed here is accordingly more tolerant of gradient fidelity issues. Changing the performance of the gradient set serves to adjust the excitation slab thickness. Reducing the slab thickness increases resolution at the cost of gradient requirements and SNR; increasing it reduces either the spatial or temporal resolution with associated SNR and gradient requirement benefits.

As expected, the sensitivity to metabolism of this hyperpolarised imaging strategy was confirmed by modulating the metabolic state of the rodents imaged, whereupon significant differences in PDH flux were observed between animals that were fasted, fed, and fed and given the PDK inhibitor DCA. The fact that the relative increase in apparent PDH flux was the same as when the same experiment was performed spectroscopically validates the technique.

6.4.2 *Flip Angle and SNR Efficiency*

Three-dimensional hyperpolarised imaging strategies need to be designed to take into account the effect of the depletion of the longitudinal magnetisation by each phase encode on the resulting image. Neglecting diffusion and flow, the relative SNR efficiency of 3D pulses for hyperpolarised imaging is well documented following earlier work by Wild *et al.* on hyperpolarised gas imaging,[159,160] and can be approximated by the ratio

$$\frac{\mathrm{SNR_{2D}}}{\mathrm{SNR_{3D}}} = \frac{\sin\vartheta_{\mathrm{2D}}}{\sqrt{N}\sin\theta_{\mathrm{3D}}} \tag{6.10}$$

where N is the number of z phase encodes. The derivation for Eq. (6.10) proceeds analogously to that for the Ernst angle; as documented by Wild,[159] the remaining longitudinal magnetisation after shot N in a phase-encoded 3D scan is $M_{zN} = M_0(\cos\theta)^N e^{\frac{-N\mathrm{T_r}}{\mathrm{T_1}}}$. Consequently, the transverse magnetisation excited at shot N is $M_{\perp N} = (\cos\theta)^{N-1} e^{\frac{-(N-1)\mathrm{T_r}}{\mathrm{T_1}}} \sin\theta$. The image we acquire is a sum of N phase encoded shots, where the shot noise scales as $\sqrt{N}$. If the SNR is assumed to be dominated by the slab at $k_z = 0$, the result Eq. (6.10) follows. Interestingly, if one assumes that all of k-space contributes equally to SNR, Eq. (6.10) is modified by an additional multiplicative term α, which is the sum of a geometric series of scan parameter terms, and T_1. Depending on these parameters, α varies between approximately 0 and 5, but its average value, integrated over $\theta \in [0, \frac{\pi}{2}]$ for the TR, N, and T_1 used here, is close to one. Equation (6.10) is illustrated graphically in Fig. 6.14. Here $N = 12$, $\vartheta_{\mathrm{2D}} = 17°$ or $61°$, and $\theta_{\mathrm{3D}} = 3°$ or $20°$. This predicts an approximately equal SNR efficiency for pyruvate imaging (2D/3D ratio$= 0.993$), and a $\sim 25\%$ SNR benefit for metabolites other than pyruvate (ratio$= 0.743$). This SNR improvement from exciting a large volume comes at a cost, as noted by Wild *et al.*, as the effect of the depletion of the longitudinal magnetisation is to apply a filter in k_z that serves to blur the image through the z direction. The size of this effect is related to $1 - (\cos\theta)^N$, the acquisition scheme, and the ratio T_r/T_1 that together determine the Full Width Half Max (FWHM) of the filter. The quantity $1 - (\cos\theta)^N$ here is small (~ 0.02) for pyruvate and $\approx 1/2$ for metabolites other than pyruvate.

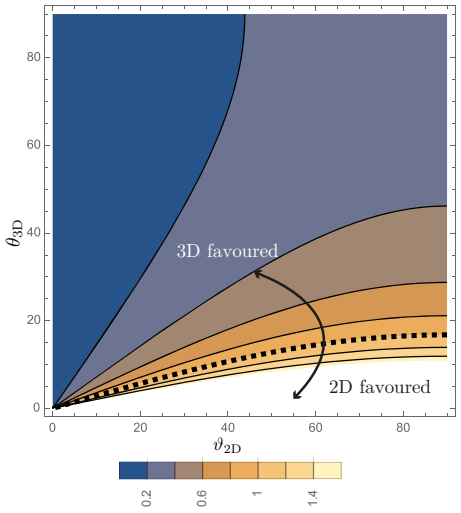

Figure 6.14: A contour plot of Eq. (6.10). The dashed line marks unity.

To quantify the effect of this proposed acquisition scheme, the k_z filter $H(z)$ as defined by Wild *et al.* was numerically calculated for the 3D sequence described here (i.e. with $N = 12$ centric ordered shots), and Fourier transformed to find the point spread function. The resulting k_z filter had a FWHM of 1.012 for pyruvate, and 1.176 for metabolites other than pyruvate. The resulting sub-voxel blurring in the

slice direction has to be considered in the context of the $\sim 25\%$ gain in SNR resulting from the three dimensional approach.

6.4.3 Scan Timing

The error in gated scan timing, here defined as the difference between actual and 'expected' ($N\times$ R-R interval) scan time, became progressively worse over time (Fig. 6.12**C**). The mean of this distribution was not equal to $N\times$ R-R interval; and it was broad: the coefficient of variation of TR /image was $40.4 \pm 11.4\%$ (distribution $\sigma/\mu \pm$ s.e.m. over 16 scans). As shown in Fig. 6.12, the observed distribution of TRs were well-described by the two-parameter Gamma distribution

$$f(x;k,\theta) = \frac{x^{k-1}e^{-\frac{x}{\theta}}}{\theta^k \Gamma(k)} \quad \text{for } x > 0 \text{ and } k, \theta > 0 \tag{6.11}$$

where $\Gamma(k)$ is the gamma function $\Gamma(k) = \int_0^\infty x^{k-1}e^{-x}\,\mathrm{d}x$, and k and θ are the shape and scale parameters that were fit to the acquired data. As the ECG can be considered a stochastic time-series that is a sum of underlying fundamental processes,[161,162] the scan counts are positive semidefinite, and the distribution of increment times – i.e. TRs – are Gamma distributed, the number of gated scans at a given time can be considered a Gamma process, a type of generalised Poisson process.[163] This fact implies that dynamic techniques which neglect this information will have a skewed time axis that becomes progressively distorted as time increases, despite the relatively small distribution in the R-R interval itself. Computed 'system independent' quantities of interest such as rate constants or perfusion values, which fundamentally depend on an accurate time axis, are therefore liable to be biased if data is not time-stamped accurately. Note that the majority of pre-clinical scanners (i.e. Varian/Agilent and Bruker) do not record the exact time at which a gated shot occurs. Future work might be able to exploit the distribution of repetition times in this acquisition with sequences inspired by stochastic NMR.[164]

Owing to the fact that z phase encoding is centric ordered, it is proposed that the effective timepoint for each 3D volume acquired is defined as the echo time of the first shot, when the centre slab of z k-space is acquired. If this definition were to be contentious, the use of an accurate timestamping device allows for the *post hoc* reconstruction of the time axis as required; here, the axis was defined as above. The acquisition strategy spans several orders of magnitude in time, with the pulse length ($= 31.48$ ms) short in comparison to the cardiac cycle (~ 150 ms) and respiratory motion ($\sim 1.5 - 2$ s). In the cardiac data acquired, no artefacts were observed that could be ascribed to respiratory motion, but a separate acquisition showed that cardiac triggering was still required to avoid motion artefacts in the myocardium. This control acquisition is illustrated in Fig. 6.15.

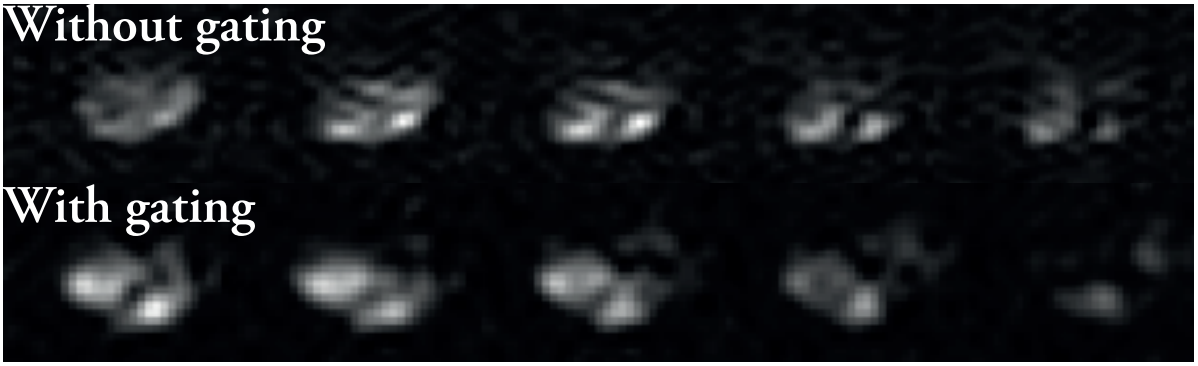

Figure 6.15: Single timepoint [1-^{13}C]pyruvate images of the heart acquired with the developed sequence. Even with the short readout and similarity of the nominal TR of the experiment to the heart rate of the anaesthetised rat, cardiac gating is required to reconstruct distortion-free images.

6.4.4 Limitations

There are several potential limitations to this 3D approach for hyperpolarised imaging. Firstly, the SNR benefit and the width of the k_z filter arising from the 3D approach is highly dependent on the actual flip angle delivered, and therefore requires good B_1^+ homogeneity. In this study, a volume transmit coil was used that was homogeneous over the region of interest. It is likely that this acquisition strategy will be challenging with a surface transmit/receive coil, where the flip angle delivered is inherently a function of space. One would expect the resulting images to have a significant degree of interslice blurring arising from the large flip angle close to the surface of the coil. However, if these effects can be ameliorated by ensuring homogeneous B_1^+ (as they have been here), the 3D approach can fundamentally deliver better SNR.

Secondly, hyperpolarised imaging sequences that use phase encoding effectively apply a blurring filter in the phase encoding direction, here z, due to the depletion by each shot of the initial magnetisation. This effect is proportional to $1 - \cos(\theta)^N$, which grows quickly with θ and N. However, the SNR benefit grows only as $\sqrt{N}\sin\theta$, which is approximately linear at small θ, and increases slowly with N. There are therefore several regions of θ and N where this approach is inferior to 2D imaging with multiple slices, e.g. as $\theta \to 90°$, where the slice resolution approaches the FOV, and at very low flip angles and small N, where the 2D approach has better SNR. As demonstrated, we have chosen values of θ and N such that this effect was ameliorated, and a substantial SNR benefit was obtained at the cost of a modest trade-off in slice resolution.

Thirdly, one would expect that the high slew-rate requirements of the fly-back pulse would introduce nonlinear distortions that serve to degrade the excitation profile of the pulse. However, this has not proved to be the case (*c.f.* Fig. 6.1**C**), despite the prediction from the measured gradient impulse response function that this pulse would occur in a region where gradient performance was non-linear.

Fourthly, the 3D sequence is relatively robust to flow artefacts, as the large excitation slab thickness ensures that excited spins will be read out irrespective of their motion. However, we have not quantified

the effect of flow over the 12 R-R intervals over which each volume is acquired, and it is likely that different parts of k_z correspond to flowing spins in different positions, which could result in a flow artefact. Two types of artefact are expected to arise from flow: ghosting in the phase encoding direction arising from in-plane motion during the EPI readout;[156] and a linear shift in z arising from through-plane motion during the 3D readout. This last effect arises as the k_z gradient blips can serve to act as velocity encoding gradients, creating a phase ramp on spins through z. In the proposed sequence, the largest blip is 328 μs long, and 7.51 G in amplitude. This corresponds to a large velocity encoding moment (c.f. Eq. (2.78)); spins flowing at 23 m/s would be shifted by one voxel in z. This artefact is correspondingly negligible in the rat heart.

It is also appropriate to consider the effect of metabolism on the resulting images as the $^{13}\mathrm{C}$ label is transferred to downstream metabolites during the acquisition period of the 3D volume. This effect is analogous to modulating the T_1 of the compound imaged, as the net effect of the (dis)appearance of a small quantity of hyperpolarised metabolite in bulk will be analogous to the increase or decrease in T_1 for that compound. Accordingly, the effect on the image is a slight blurring or sharpening through k_z, as the non-renewable magnetisation is depleted or slightly replenished between phase encodes. No through-plane blurring was observed, indicating that this effect is small.

Finally, the first-order phase correction method used did not always remove the Nyquist ghost arising from the EPI readout train used. Owing to the large acquisition FOV of the technique, it was possible to recombine odd and even lines separately in every case, halving the FOV and eliminating the ghost. However, in general it is possible to perform a 2D phase correction to remove the distortion.[24] Unfortunately, the dynamic nature of hyperpolarised imaging means that many standard algorithms will not work 'out-of-the box', and instead a manual in-plane phase map has to be created and applied for each acquisition, estimated from high SNR images that cover as large a volume as possible, including both ventricles and the myocardium. This is a potentially challenging and time-consuming step that might drive a desire to perform fly-back EPI, or other rapid imaging readout trajectories.

6.5 | Myocardial Infarction

6.5.1 Overview

As an extension of the above work, a study was designed to image [1-$^{13}\mathrm{C}$]pyruvate metabolism following a period of myocardial ischæmia. As outlined in Section 4.2.2, it is currently difficult to distinguish between viable, hibernating myocardium and tissue that is dead following a myocardial infarction (MI). Imaging the rates of glycolysis and fatty acid oxidation may provide the ability to discriminate between infarcted, viable tissue – which would be highly glycolytic – and necrotic tissue that would progress to fibrotic scar formation; characterised by a 'hole' in nearby regions of healthy metabolism. Additionally, the novel high-resolution imaging of metabolism in the whole heart following MI might reveal new

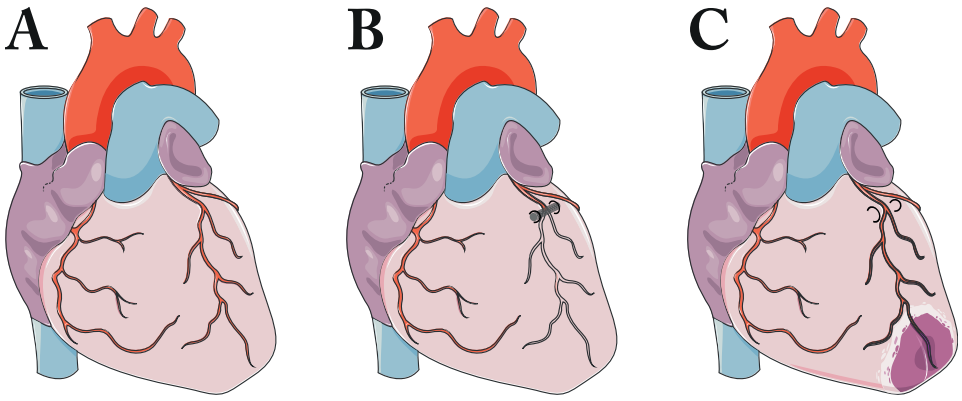

Figure 6.16: The process by which the ischæmia reperfusion injury model of MI was induced. The heart was exposed (**A**), and the left anterior descending coronary artery occluded by a needle (**B**). After removal of the occlusion, a reproducible infarct in the front wall left ventricle could be produced.

information about the metabolic remodelling occurring in the peri-infarct region, which may be of utility in understanding the descent into heart failure following MI.

Accordingly, three male Wistar rats were used as a model of MI, and subject to a 45 minute period of localised myocardial ischæmia followed by reperfusion. As illustrated in Fig. 6.16, animals were anaesthetised with 2.5% isoflurane in O_2, before a left anterolateral thoracotomy was performed (by Carolyn Carr, Andrew Lewis and Vicky Ball) and the heart exposed (c.f. Fig. 6.16**A**). The left anterior descending (LAD) coronary artery was ligated approximately 2 mm from the origin, using a 5-0 proline suture tied around a blunted 18 gauge needle that was placed over the artery (c.f. Fig. 6.16**B**), and the chest closed. After 45 minutes, the chest was reopened, and the suture removed, leading to a reproducible infarct (Fig. 6.16**C**). Sufficient peri- and post-operative analgesia was provided. As before, all experiments were performed in accordance with relevant UK legislation (with personal, project and institutional licences granted under the Animals (Scientific Procedures) Act 1986), and were subject to local ethical review and an independent cost-benefit analysis.

It was originally planned to use the developed sequence to scan each animal at three, seven, fourteen, and twenty one days post infarction, and also six to eight weeks following MI, to provide an accurate characterisation of the metabolic changes occurring in the infarcted heart as a function of time. Unfortunately, two days after surgery, the magnet quenched, as shown in Fig. 6.17. The reason for the quench, which occurred at approximately 1800 hours when the magnet was not being used, was not known, and could not be determined. Owing to the presence of (oxygen) ice within the cryostat, and the degree of difficulty in ramping it back to field, it was not possible to subsequently image at these chosen time points.

Several months later, the animals infarcted were imaged with the developed sequence, and traditional gradient echo CINE to assess cardiac function. Hearts were then excised under terminal anaesthesia, fixed in formaldehyde, embedded in Tissue-Tek and sectioned by Andrew Lewis. Sections of heart were

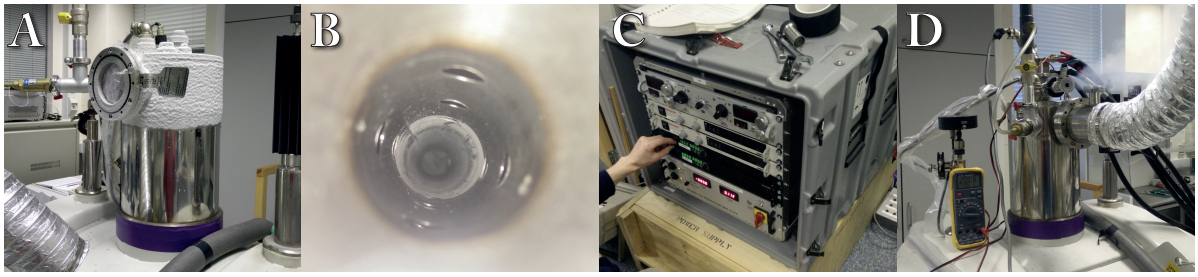

Figure 6.17: The aftermath of a superconducting quench, which prevented the longitudinal study discussed above from running as planned. **A** The overpressure created by the rapid boiling of the large volume of liquid helium surrounding the superconducting magnet was sufficient to burst the bursting disc (as designed), which allowed the venting of helium gas from the inner cryostat to the emergency quench pipe (the bursting disc was temporarily replaced by clingfilm in this photograph). Additionally, it was also sufficient to rip and blow off the quench pipe, break two windows, damage the floor, and relocate several light items within the magnet room (which, presumably, was not as designed). **B**: Owing to the presence of ice within the voltage sense probe connection within the service turret on the top of the magnet, it was difficult to monitor the magnet adequately and to try and elucidate the cause of the quench. **C**, **D**: Re-energisation of the magnet with a large (500 A max) constant current supply. Large, vacuum insulated and helium-cooled conductors are directly inserted onto terminals within the cryostat body, with an exceptionally high impedance voltmeter allowing for the measurement of the potential difference between the magnet terminals (2.003 V). The magnet was re-energised approximately two months after it quenched.

then stained with Masson's trichrome, a three-colour histological stain that colours intact myocytes red, collagen blue, and cell nuclei black. The whole heart sections were then scanned by $20\times$ optical microscopy on a Leica Aperio slide scanner "digital histology" station. No attempt was made at deformable co-registration between the histological data and MR images.

6.5.2 Results

Of the three animals subject to MI, only one was observed to show evidence of cardiac dysfunction by CINE MRI, in the form of an akinetic region in the anterior wall of the myocardium, consistent with LAD occlusion. This correlated with a decrease in the measurable ejection fraction from the CINE stack. The other two animals were apparently normal upon CINE, and did not show evidence of MI or myocardial dysfunction.

As shown in Fig. 6.18, the hyperpolarised imaging data acquired from these animals was variable, and the apparent degree of polarisation for the third animal was substantially lower (for an unknown reason) than that of the other two. Nevertheless, as before, it was possible to adequately resolve metabolism in the rodent heart. Interestingly, regional deficits were observed in the amount of ^{13}C-bicarbonate present within the myocardium, in the anterior myocardial wall, when compared to that obtained from healthy animals (c.f. Fig. 6.11). As illustrated in Section 6.5.2, these defects approximately co-localised with the histological ground truth. It was confirmed by histology that one animal had extensive anterior damage to the myocardial wall that extended substantially along the long axis of the heart, as evidenced by the large areas of fibrosis and thinning, consistent with the ^{1}H MR findings. Likewise, the infarct region in

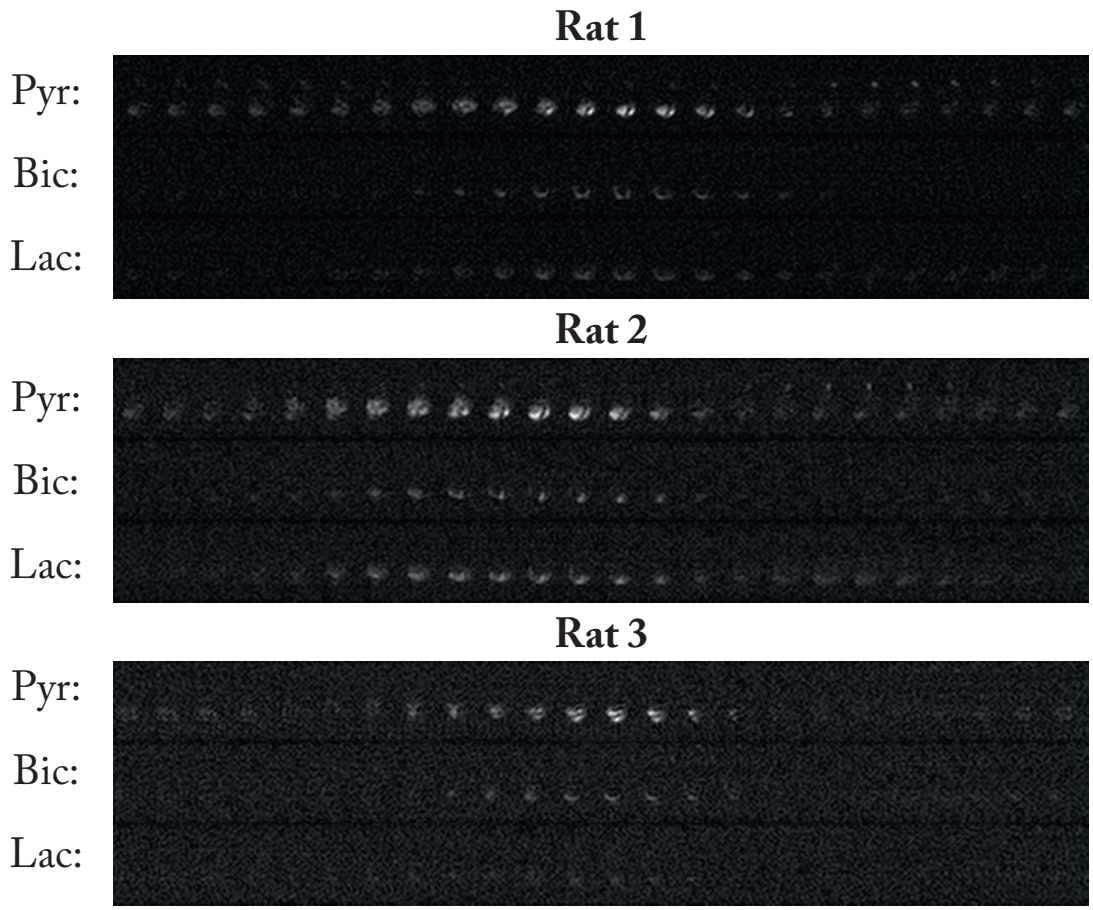

Figure 6.18: The temporal sum of the data acquired from the hyperpolarised [1-^{13}C]pyruvate imaging of the three infarcted rodents. In each image z increases (along the rostral/caudal axis) from left to right. Regional defects in ^{13}C-bicarbonate production are apparent in the infarcted animals, which are shown in greater detail in Section 6.5.2.

the other two animals was substantially smaller, and did not penetrate the full depth of the myocardium. This was consistent with the varying size of 'the hole' in the expected crescent of ^{13}C-bicarbonate production, with the first animal showing substantial disruption in the myocardial wall, the second animal a one-to-two voxel defect, and the third animal apparently normal.

6.5.3 Discussion

The ischæmia-reperfusion injury rodent model of MI is well documented and believed to be a physiologically meaningful parallel to the type of injury suffered in human patients.[165] It is therefore unsurprising that the degree and location of the injury sustained is variable: aside from the technical challenges of performing cardiac surgery on rodents, small variation in the location and duration of the suture used to occlude the LAD produce a large effect on the viability of the downstream tissue. Importantly, the hyperpolarised [1-^{13}C]pyruvate imaging sequence developed was able to provide information consistent within this variability, and, if interpreted diagnostically, would lead to a prediction of varying degrees of infarction for each of the three animals. Whilst the study as planned was unable to be completed due to hardware failure, the 'very chronic' infarction imaged, displayed metabolic shifts that were consistent with those believed to occur in the damaged heart.

This study represents a promising pilot that demonstrates the feasibility of using three dimensional spectral-spatial EPI to image hyperpolarised [1-^{13}C]pyruvate metabolism in the rodent heart. Future work will examine that time-course of metabolic derangement, from infarct to heart failure. By understanding the compensatory period of metabolic maladaption, it is hoped that targeted modulators of metabolism could be applied to ameliorate the rate of progression from MI to first compensated heart failure, and then acute decompensated heart failure.

There are at least two challenges highlighted by this study that should be overcome before progressing to larger scale studies of the response of the heart to infarction. First, it is apparent that the sensitivity profile of the coil used could be improved: any variability in visible ^{13}C-bicarbonate reception is highly deleterious when attempting to visualise a localised defect in a ring of ^{13}C-bicarbonate production corresponding to the left ventricle. This effect was mitigated by choosing to ligate the LAD coronary artery in an attempt to localise the resulting infarct to the front wall of the myocardium. The construction of a four-channel surface receive array, designed to have uniform spatial coverage of the whole heart, is currently ongoing.

Second, before undertaking a study similar to the above to answer pertinent biological questions, it is necessary to consider the method of quantification of these images. As discussed at length in Chapter 9, this is a nontrivial problem as the variability between [1-^{13}C]pyruvate dissolutions, between the degree of infarction of each animal, and between the response of each animal cannot be neglected, or de-convolved. Automated, sophisticated methods based upon solving constrained systems of differential

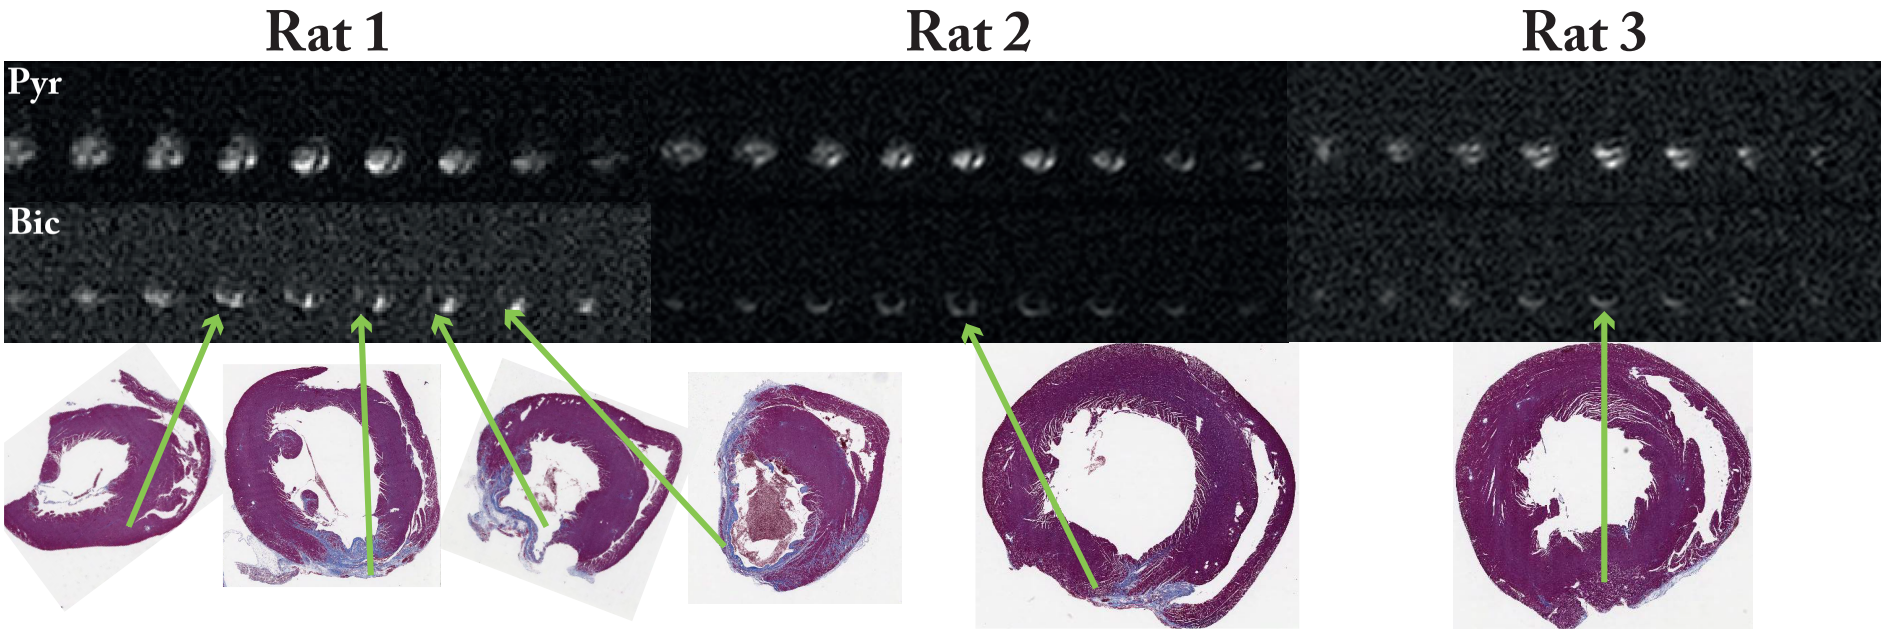

Figure 6.19: An expanded view of the summed [1-^{13}C]pyruvate and ^{13}C-bicarbonate images acquired around the heart for the three infarcted animals. Regional defects in ^{13}C-bicarbonate production are visible, and correspond with histological slices acquired from the same region, as illustrated below. The infarct Rat 1 acquired was substantially larger than that of the second or third animals, and extended from the top of the ventricle (leftmost histological slice) below the chordæ tendinæ, along the length of the papillary muscles and became substantially worse towards the apex. This corresponded with a larger visible drop-out in apparent bicarbonate production.

equations may be ultimately the most appropriate way of quantifying the high-dimensional datasets DNP imaging sessions provide, but also the most time-consuming way to quantify these images. Simpler methods, such as physician-guided myocardial segmentation followed by the computation of metabolite ratios such as [1-^{13}C]lactate/[1-^{13}C]pyruvate may be substantially quicker and easier to use, but less robust to noise and subject to greater variability. These issues should be investigated in future work to allow strong, statistically sound conclusions to be obtained from the imaging of a highly heterogeneous population of rodent infarcts.

6.6 | Conclusions

This chapter has described the development of a novel ECG-gated three-dimensional pulse sequence for the cardiac spectroscopic imaging of hyperpolarised [1-^{13}C]pyruvate and its downstream metabolites *in vivo*. The sensitivity of the developed sequence to metabolism has been confirmed systemically by manipulating the metabolic state of the rodents imaged, and locally by scanning rodents subject to myocardial infarction. The 3D spectral-spatial excitation is relatively immune to B_0 inhomogeneity, and does not place unphysical demands on gradient hardware in order to image at high spatiotemporal resolution. The sequence is robust, and is suitable for routine use.

A simple argument is presented to support the hypothesis that three-dimensional hyperpolarised imaging is more SNR efficient than multislice approaches, at the cost of a slight blurring in the slice direction. Additionally, it has further been shown that cardiac time-stamping is important when imaging dynamic processes because of the stochastic nature of cardiac gating, a fact that is often overlooked. This work allows the spectrally efficient metabolic imaging of cardiac metabolism in small animals at high spatio-temporal resolution. This is of utility for pre-clinical research into both cardiovascular disease and more broadly diseases where metabolism is dysregulated.

Chapter 7

Brain Metastases

The successful development of the high-resolution, three-dimensional metabolic imaging sequence described in Chapter 6 led to renewed interest in attempting to diagnose cancer that has metastasised to the brain by its metabolic phenotype. As outlined previously (in Section 4.3), hyperpolarised [1-^{13}C]pyruvate imaging is uniquely placed to potentially differentiate between diseased and healthy tissue based on the average rate of glycolysis occurring within a given voxel in the brain. Therefore, it was decided to undertake a modified form of the study described in Section 4.3.3 using the developed spectral-spatial imaging sequence to image [1-^{13}C]pyruvate metabolism in the brain of tumour-bearing rats. It was hoped that this would represent a significant improvement over slice-selective spectroscopy, as the resolution of the developed imaging sequence was believed sufficient to distinguish between different anatomical regions in the rat brain. In an attempt to model a more diffuse, metastatic tumour, the protocol for intracerebral injections was slightly modified, introducing a smaller number of cells in a larger volume. Additionally, owing to the lack of availability of the BD-IX rat in the UK, Sprague Dawley rats were used instead. This difference in strain was not believed to present a problem, as the immunological compatibility of the cell line used and the Sprague Dawley rat had been independently confirmed by other researchers (Sébastien Serres and James Larkin).

7.1 | Methods

7.1.1 Animals

Twelve female Sprague Dawley rats were divided randomly into two equal groups, and subjected to intracerebral stereotactic injection of either a total of $4\,000$ ENU 1564 cells or an equivalent volume ($0.5\,\mu L \times 4$) of sterile phosphate buffered saline as a control. Surgery was performed by Sébastien Serres, and the injections were stereotactically delivered to four separate locations ($0.5\,\mu L$/site) close to the cortex through two small burr holes: 3 mm lateral of Bregma, at depths of 1.5 and 3.5 mm and ± 1 mm anterior/posterior of Bregma. Appropriate postoperative analgesia was provided, and animals were monitored daily. As before, all experiments were performed in accordance with relevant UK legislation (with personal, project and institutional licences granted under the Animals (Scientific Procedures) Act

1986), and were subject to local ethical review and an independent cost-benefit analysis. The author was blinded as to which rats were tumour-bearing or control.

7.1.2 MR Imaging protocol

Following the induction of brain metastases, rats were anaesthetised and scanned at 1, 8, 15, 22, 29, 36 and 43 days post intracerebral injection, on the Varian/Agilent 7 T preclinical system described previously, with the dual tuned volume transmit/surface receive RF coil setup (c.f. Section 3.2.5). The imaging protocol consisted of 3D gradient echo localisers ($100 \times 100 \times 100$ mm^3 FOV, $128 \times 128 \times 128$ matrix, TR$= 3.15$ ms, TE$= 1.56$ ms, 20° flip) followed by the automatic shimming algorithm described in Section 3.4. A two-echo multi-slice B_0 map was then acquired with the same FOV as the developed 3D spectral-spatial hyperpolarised imaging sequence (i.e. $64 \times 64 \times 45.5$ mm^3) at higher in-plane resolution (matrix size 128×128) and with a 20° flip angle. The transmit power for the ^{13}C volume coil was then calibrated on a bottle of 5 M thermally polarised ^{13}C-urea that was included next to the animal on every scan. This calibration was performed by non-localised spectroscopy (1 ms hard pulse, 16 averages) transmitting on the volume coil and receiving on one channel of the surface array. The transmit power was arrayed through five values chosen *a priori*, a $M_0 \sin a\theta$ curve fitted, and the appropriate value extracted. The transmit power required was found to be consistent between experiments (with 50.1 ± 0.6 dB and a 1 ms pulse generating $\theta = 90^\circ$), validating the assumption of constancy made in Chapter 6. The urea phantom also acts as a positive control to ensure the correct operation of the transmit/receive circuitry.

After ^{13}C power calibration, each rat was scanned with the hyperpolarised spectral-spatial imaging sequence as described in Chapter 6, but without cardiac gating. The sequence was started prior to the injection of 2 mL of hyperpolarised [1-^{13}C]pyruvate over ~20 s. Following this scan, high-resolution gradient echo (GE) and fast spin echo (FSE) multi-slice images were acquired pre- and post gadolinium contrast (GE: TR$= 120$ ms, TE$= 4.68$ ms, $\theta = 90^\circ$, 40×40 mm^2 FOV, 192×192 matrix, 8 averages, 11×1.5 mm thick slices. FSE: TR$= 1300$ ms, TE$= 9.6$ ms per segment, 2 segments, gauss RF excitation ($\theta = 90^\circ$) and refocussing ($\theta = 180^\circ$) pulses; 40×40 mm^2 FOV, 128×128 matrix, 2 averages; 11×1.5 mm thick slices). Additionally, a 3D gradient echo scan was acquired with the same FOV as the hyperpolarised scan for subsequent ease of image overlay (TR$= 10$ ms, TE$= 1.96$ ms, $128 \times 128 \times 128$ matrix, $64 \times 64 \times 45.5$ mm^3 FOV, 1 average, 20° flip angle).

The time taken for each step in this per-rat protocol is summarised in Fig. 7.1.

7.1.3 Reconstruction and Data Analysis

Hyperpolarised data were regridded in Matlab and reconstructed as described in Section 6.2.5. Owing to the reduced SNR present in brain images, it was necessary to temporally sum the data in all cases,

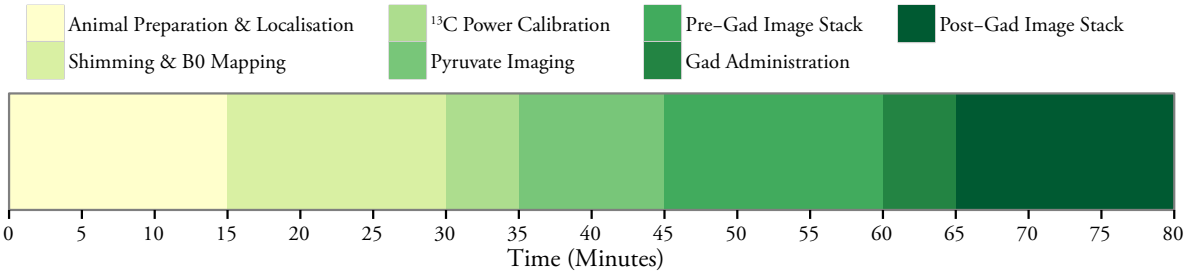

Figure 7.1: A summary of the scan protocol used to image each animal.

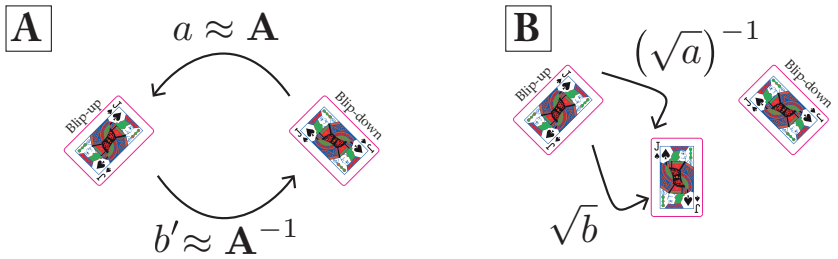

Figure 7.2: A cartoon illustration of the distortion correction procedure applied here. **A**: Affine registration of the blip-up and blip-down datasets produce two separate deformation matrices, a and b, which are experimental approximations of idealised deformation **A**.

sacrificing temporal information to reduce noise. Data were only summed from a temporal window coinciding with visible signal in the pyruvate channel, in order to maximise apparent SNR. [1]

Additionally, it was noted that B_0 homogeneity over the region of the rat head was poor compared to that achieved in the heart. In order to minimise spatial shifts experienced by different parts of the anatomy, the `fsl` module `topup` was used to distortion correct blip-up and blip-down data.[22,146,147] In order to use `topup`, it was necessary to compile a local version that was multinuclear compatible and did not contain certain memory-management related issues that became apparent in the process of reading in multinuclear data.[2] Additionally, regridded data were exported from Matlab with appropriate metadata as `nifti` files for `fsl` registration. It was also necessary to create a new stack of default registration parameters and subsampling methods to produce acceptable co-registrations with `topup`. Blip-up and blip-down data were temporally summed, before being sent to `topup` with the invocation options `-subsamp=4,2,1 -fwhm=8,4,0 -miter=15,15,40 -numprec=double -lambda=50,20,2 -ssqlambda=1 -estmov=0 -scale=0 -regmod=membrane_energy`, corresponding to multi-level image domain subsampling, blurring with a Gaussian kernel and a membrane-energy based regularisation kernel under the assumption that the anaesthetised rat had not moved during the scan.

Additionally, the blip-up and blip-down datasets were co-registered by `flirt`, the `fsl` image registration tool, called as part of `topup`'s post-processing `eddy_correct` tool. This step automates

[1] This approach is noted as not necessarily being SNR optimal, but was required due to the comparatively low signal subsequently acquired.

[2] These fixes were subsequently provided to the developers.

the alignment of the datasets, and essentially finds an affine translation between the two.[3] Additional regularisation parameters were imposed on `flirt` to ensure that the affine registration was 'reasonable', i.e., the resulting 4×4 matrix (consisting of rotations, translations, shears, and scaling parameters due to estimated distortions) was close to the identity matrix. By co-registering first blip-down data onto blip-up data, generating a deformation matrix a, and then blip-up data onto blip-down data, producing a deformation matrix b', two different approximate B_0 correcting affine transformations were obtained. In the absence of noise and non-linearities, $a^{-1} = b'$, and $b'^{-1} = a$. As these matrices are defined in different image coordinate systems, it is convenient to pick one arbitrarily and define $b = b'^{-1}$, say, such that a and b should be identical. In the presence of noise, this is not the case, and the two differ slightly. These two returned matrices a and b were then read into Matlab, and an approximate 4×4 matrix square root computed under the assumption that any residual deformation was linear, producing a mapping "between" the two opposite distortions, as illustrated in Fig. 7.2.

It is worth noting that, if (perhaps due to noise) each individual deformation matrix is not diagonalisable then no matrix square root exists and, in general, it is impossible to provide an adequate numerical approximation of one. This problem goes both ways: in general, any arbitrary matrix $\mathbf{A}$ may have any integer number of roots $k \in [0, \infty]$. Fortunately, the affine deformation matrices generated by `fsl` are Hermitian, and have eigenvalues located in the upper half plane, and, hence, are positive-semidefinite.[166] Accordingly, they possess exactly one positive-semidefinite square root, which is that used. However, owing to the presence of experimental noise, the two constructed 'unwarped' images are not identical, and are subject to distortions. The construct

$$\mathbf{M} = \sqrt{\left(\sqrt{b}\right)^{-1}\left(\sqrt{a}\right)^{-1}} \tag{7.1}$$

acts as an interpolator between the two affine transformation matricies, and is proposed as an effective way of dealing with noise. Indeed, if $a = b$, Eq. (7.1) reduces to $\sqrt{a^{-1}}$, i.e. the ideal mapping discussed above (with the inverse serving to change the definition of the original mapping).

As an illustration, this process is demonstrated in Fig. 7.3 for a 4×4 affine transformation matrix obtained by applying the `topup/fsl` algorithm to a representative hyperpolarised dataset. Rather than being used to correct distortions in this dataset, Fig. 7.3 shows a reference 3D image (a T_1 weighted brain image) that was initially distorted appropriately to mimic EPI data. The 'naïve' approach of correcting by $\sqrt{a}$ or $(\sqrt{b})^{-1}$ as appropriate produces two sets of candidate images that are similar to the ground truth, but the mean combination is closer still. This approach was used to co-register temporally summed stacks of blip-up and blip-down hyperpolarised 3D spectral-spatial EPI data.

After calculation of an appropriate distortion correction field across temporally summed data, these fields could retrospectively be applied to the time-series data, should they have sufficient SNR to support

[3] Affine transformations are any transformation that preserves collinearity between points and ratios of distances; i.e. the superset of rotations, scales, shears, and translations.

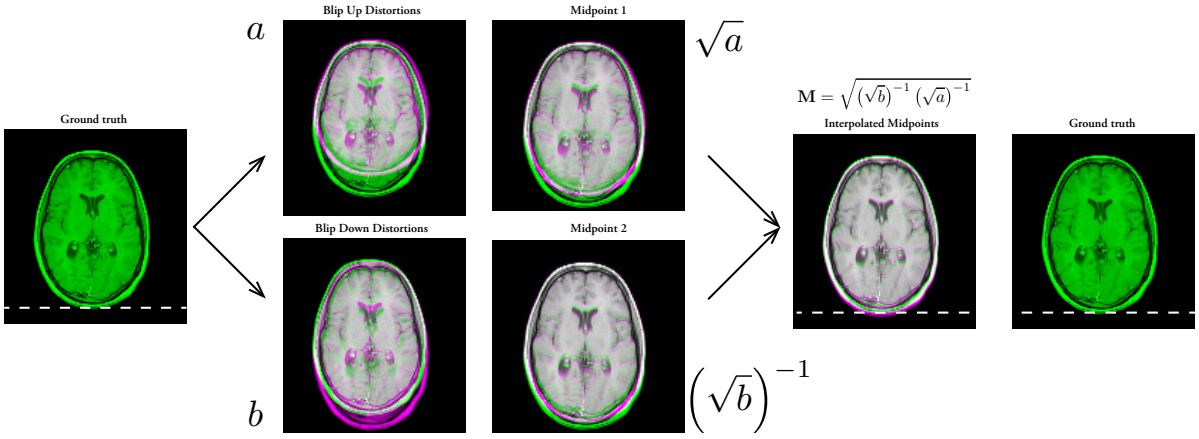

Figure 7.3: An illustration of the co-registration steps used to apply the output of `topup` and `flirt` to hyperpolarised data sets. In this computational experiment, a ground truth image (always shown in green) was subject to distortions estimated by `topup` and `flirt` to apply to hyperpolarised data. These are different between blip-up and blip-down EPI data, leading to the images labelled a and b which are overlaid in magenta on the green ground truth: co-localisation is therefore clearly visible as a white image. The naïve estimation of the ground truth, $\sqrt{a}$ and $(\sqrt{b})^{-1}$, certainly approximates the ground truth, but the estimator $\mathbf{M}$ is closer. The white dotted line highlights one residual area of difference between the ground truth and the corrected image.

this process. Finally, it is worth noting that, as an entirely image-domain correction step, `topup` is unable to function in a low SNR environment.

Images were quantified by segmenting out a region corresponding to the approximate location of the injection site (namely the superficial half of the ipsilateral hemisphere) on the proton images, and then computing the mean lactate to pyruvate ratio within that region. As the gain between pyruvate and lactate channels is identical, and the acquired RF noise is white, it is expected that the lactate to pyruvate ratio would be one in the absence of signal, and less than one in the presence of signal. Measurement of a control image stack in the absence of hyperpolarised signal revealed that the expected value of the ratio is 1, with a standard deviation of approximately 0.08. Therefore, the effect of noise in the ROI masks was ameliorated by thresholding the obtained lactate-to-pyruvate ratio images to remove voxels greater than 0.8 in magnitude. Whilst this approach places an upper limit on the detectable sensitivity of the technique, it allows for tolerance should the manually constructed ROIs be located marginally outside the brain region, which would otherwise introduce noise and variability. Other methods used to more robustly quantify DNP data using statistical methods that can discriminate between signal and noise are described in Section 9.4, but they require sufficient SNR for time-resolved imaging to be effective. The lactate to pyruvate ratio is a simple method of accounting for the substantial polarisation variability that is present, and is reported to be both robust and highly correlated with other, more complex quantification methods.[167]

7.1.4 Statistical Analysis

In order to adequately deal with data missing due to technical failures or the of death of animals, and allow for the paired nature of the longitudinal comparison, a linear mixed effects model was constructed and fit to the data using the `nlme` package[168] and `R`.[169] Linear mixed effects models can be considered as a generalised form of repeated measures ANOVA, but are more versatile in the sense that they can can be constructed from any arbitrary mix of categorical and continuous data, are able to ascribe variability to different sources, and are able to function in the presence of missing data.[170] Another advantage over traditional repeated measures ANOVA is that the within-group errors are allowed to be correlated and/or have unequal variances, meaning that there is no requirement to test for heteroscedasticity. Additionally, such models are amenable to Bayesian methods for maximum likelihood estimation,[4] such as Markov Chain Monte Carlo algorithms, which are guaranteed to (eventually) converge on the correct posterior form of the likelihood function.[171] This numerical approach is often required as, in general, it is not possible to find the likelihood function of generalised linear mixed effects models in a closed form.

The identity of each animal was considered as a random effect in a linear mixed effects model, with time and the disease status of each animal considered fixed continuous and categorical variables respectively. Disease status – namely whether a rat belonged to the 'cancer' or 'control' group – was determined by the size of the tumour ultimately grown. This model can be solved in a closed form, and was evaluated as `anova(lme())` in `R`. Missing data were omitted from the calculation of the model, though typically they may be imputed if required. In order to verify that the mixed effects model constructed was appropriate for describing the data, control models were constructed where either status or time were not considered fixed effects.

7.1.5 Histology

Sébastien Serres sacrificed rats six weeks post injection and performed histological sectioning and immunohistochemical staining as per the protocol discussed in Section 4.3.4.4. Slides were then scanned as before, by Sébastien Serres.

7.1.6 Cell Experiments

As an *in vitro* control to directly measure the apparent flux through LDH, the ENU 1564 cell line used in the *in vivo* experiments along with three other types of cancer cells (U251, a glycolytic glioblastoma cell line; HCT116, a human colon cancer cell line; and W8CTCs, derived from circulating tumour cells of a mouse mammary carcinoma) were grown in culture. Normal human astrocytes were provided as a control. All cells except from HCT116 were a gift to the author by Sébastien Serres; the HCT116 data

[4] The likelihood is defined as the probability of observing the data given parameters in a model, and its maximisation is usually the central objective of Frequentist statistics.

had already been acquired by Michael Dodd. Dulbecco's Modified Eagle Medium with 10% (v/v) FBS and 1% (w/v) glutamine was used as culture medium. Additionally, 1% penicillin and streptomycin solution (w/v) was added to minimise the probability of spontaneous infection. After thawing, cells were grown in a T-25 flask and incubated at 37 °C until 80% confluent, before being split to a T-175 flask with a 0.05% (v/v) solution of the non-specific protease Trypsin. Cells were pelleted by centrifugation at $700 \times g$, and re-suspended in 1 mL complete medium at 37 °C in a warmed (1 cm diameter) NMR tube.

The warmed NMR tube was then transported in a bottle of warm water to a heat bath next to 500 MHz Bruker Avance spectrometer, which was positioned adjacent to a 'Hypersense' (commercially available) DNP hyperpolariser. After shimming on a tube containing an equal volume of water, the cell suspensions were placed in the spectrometer. The dissolution was performed and the evolved liquid diluted threefold into a buffer containing 300 µM K^+-EDTA, before 1 mL was injected directly into the cell suspension. Simple pulse-acquire spectroscopy (hard pulse, 15° flip, TR= 1 s) was then used to quantify the rate of lactate production following the injection of hyperpolarised [1-^{13}C]pyruvate.

After the dissolution, the number of cells present in the NMR tube was estimated by sampling 12.5 µL volumes from the total volume, and mixing with the diazo dye Trypan Blue to differentially stain live and dead cells. These samples were then automatically counted by an Invitrogen 'Countess' hæmocytometer, which images the cells in a regular grid and performs image domain segmentation over a large known volume to identify live cells (those without intracellular Trypan Blue) and dead cells (with the dye).

Owing to insufficient SNR, spectra were temporally summed prior to quantification with the AMARES algorithm in jMRUI as described in Section 4.3.4.3. The ratio of total [1-^{13}C]lactate to [1-^{13}C]pyruvate per live cell was then computed, with uncertainty from every source propagated in quadrature.

7.2 | Results

7.2.1 In Vivo *Study*

Unfortunately, of the twelve animals used in this study, only one had developed a tumour visible by MRI by 43 days after injection of metastatic cancer cells. At this point, animals were sacrificed by Sébastien Serres. Additionally, one rat unexpectedly died under anæsthesia during the imaging protocol 15 days after injection. It was believed that this death was a complication of isoflurane anæsthesia, as no other plausible explanation for sudden cardiac arrest could be found.

The imaging sequence developed in Chapter 6 was robust, and functioned without failure for the duration of the study. However, on three occasions (out of 84), the tail vein was not cannulated properly, and the experiment was not able to be repeated owing to Home Office licence restrictions. As shown in Fig. 7.4, it was found that [1-^{13}C]pyruvate could be observed perfusing the brain and major vessels

of the head and neck, and was seen to co-localise with the expected position of the Circle of Willis, the common carotid arteries, and the jugular vein. This perfusion pattern was seen across all animals, and is consistent with the known anatomy of the head and neck of the rat. The fact that pyruvate perfusion was observed in the vasculature near the ventral surface of the rat, spatially separated from the surface receive coil used in the study, indicates that the sensitivity profile of the coil was not limiting the reception of hyperpolarised signal from deep areas of the brain.

In control animals, [1-^{13}C]lactate was observed in the superficial regions of the brain, again corresponding to the location of the Circle of Willis. Even summed across time, the overall SNR of the detected [1-^{13}C]lactate signal was poor, with a range estimated (by the ratio of the brightest voxel to those outside the animal) at ~ 1.6 to ~ 5. The time course of disease for the single animal that developed a tumour is shown in Fig. 7.8. It was just possible to resolve localised increases in apparent lactate production in the region surrounding the injection site after surgery (Fig. 7.8**A**, first row, shown magnified in **C**), but the scan on day eight possessed low SNR, and it was not readily apparent that any discernible structures were present. In this animal, [1-^{13}C]lactate was observed in the approximate region of the tumour at approximately the same time it became visible with contrast enhancement (shown in Fig. 7.8**B**), namely 15 days post injection. By day 29, a clear enhancement in the region of disease was present both in the gadolinium enhanced ^{1}H images and also by an increase in the detectable amount of [1-^{13}C]lactate.

As shown in Fig. 7.5**A**, the presence of gross carcinoma was confirmed by histology in the rat with the tumour visible by MR. For all other rats, no evidence of tumour was found. A slight injury was noted at the injection site (Fig. 7.5**B**), but no nearby cells were stained by the DAB immunohistochemistry protocol. It is believed that a small number of nearby slightly dark cells surrounding the injection site correspond to erythrocytes in the region of the injection site. Such cells are believed to reflect an artefact of the perfusion-fixation procedure, which ordinarily washes out blood from the tissue. An apparent localised increase in cell density was also observed in the region of the injection. This is consistent with the recruitment of inflammatory cells in the region following injection.

A decreasing trend was observed in the apparent lactate-to-pyruvate ratio over time, shown in Fig. 7.6. It is believed that this decreasing trend in control animals represents the inflammatory response to injection. Time was a significant factor in determining the apparent lactate-to-pyruvate ratio in the linear mixed effects model used, and the single rat with disease could be considered an anomaly (Time F value 47.086, $p < 0.0001$; Status F value 4.556, $p = 0.0586$). The statistical model used was appropriate for describing the data (Akaike Information Criterion (AIC) $= -162.82$ and Bayesian Information Criterion (BIC) $= -151.65$ compared to control AIC $= -127.00$, BIC $= -120.29$; likelihood ratio test $p = 2.25 \times 10^{-9}$), with normally distributed residuals. The data were highly variable, with intra-animal variation approximately twenty times higher than that between animals. This is consistent with the

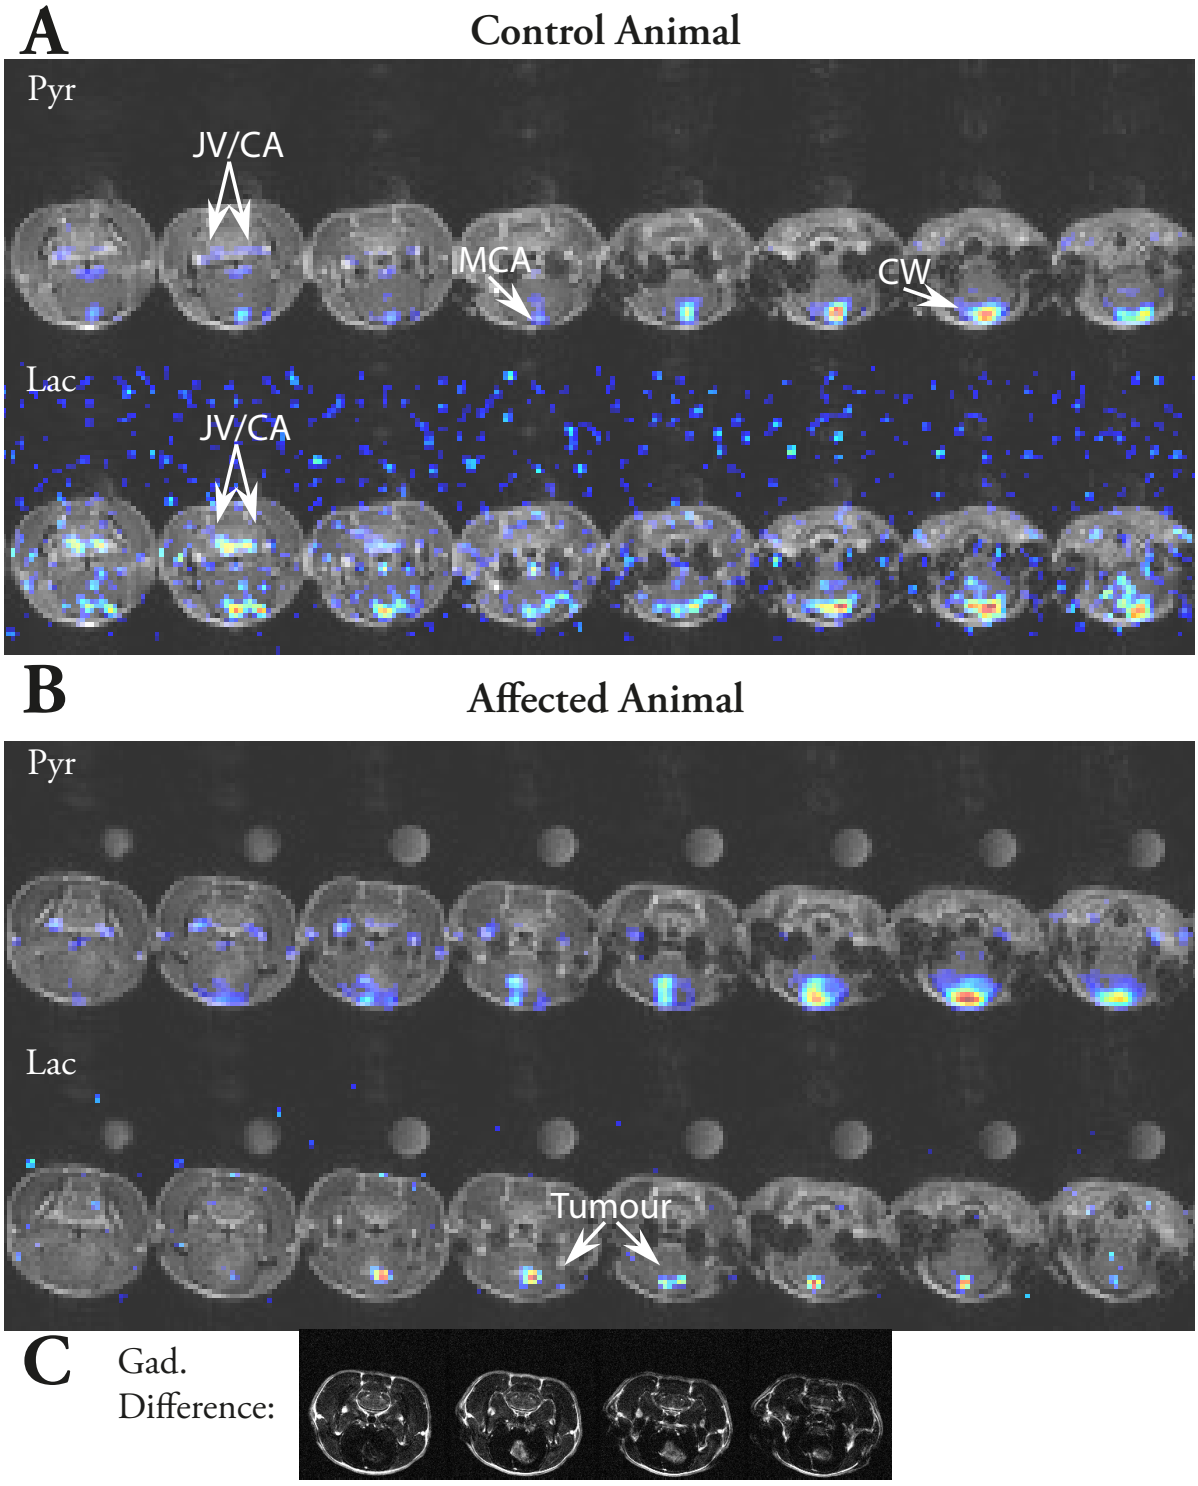

Figure 7.4: Representative data acquired from this study. **A**: Temporally summed [1-^{13}C]pyruvate and [1-^{13}C]lactate volumes shown concatenated across space, from the level of the oropharynx to the approximate location of the clavicle in a control animal. Pyruvate was seen perfusing in the approximate location of the Circle of Willis (CW) and the middle cerebral artery (MCA), along with the common carotid arteries and the jugular veins (JV/CA). The observed lactate signal is close to the noise floor of the experiment, but is again approximately seen in the location of the Circle of Willis. The colour axes in these plots are scaled such that pixel intensities $< 70\%$ or 30% of maximum (for pyruvate and lactate respectively) are transparent; red represents the maximum intensity in each metabolite. **B**: As above, but for the animal that developed a large tumour, shown here 36 days after injection. Lactate was observed in the body of the tumour. The colour axis is as for A. **C**: Subtracted pre-and-post gadolinium contrast high resolution proton images in the tumour-bearing region. The metastatic region is clearly visible. Note that the slice spacing is reduced compared to the hyperpolarised images.

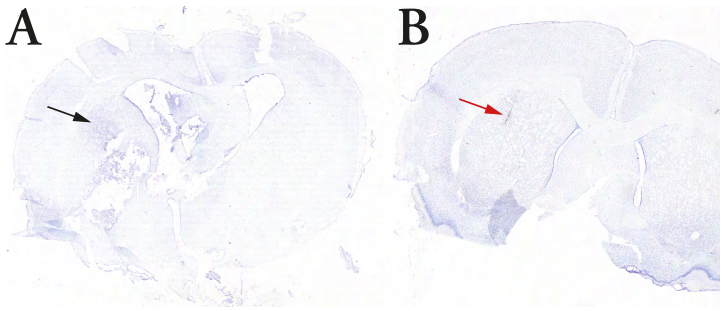

Figure 7.5: Histological slices of example affected (**A**) and unaffected (**B**) animals. **A**: The one rat which developed disease visible to MR was found histologically to possess a large, necrotic tumour (black arrow). The example slice shown in **B** was believed to be injected with 4 000 ENU tumour cells, but no evidence of this is shown. A slight injury is apparent at the injection site. The darker cells present are believed to be erythrocytes, and a slight increase in neuronal cell density is also apparent. It is believed that the erythrocytes visible are an artefact of the perfusion-fixation procedure, and the increase in apparent cell density is consistent with an inflammatory response to injection (red arrow).

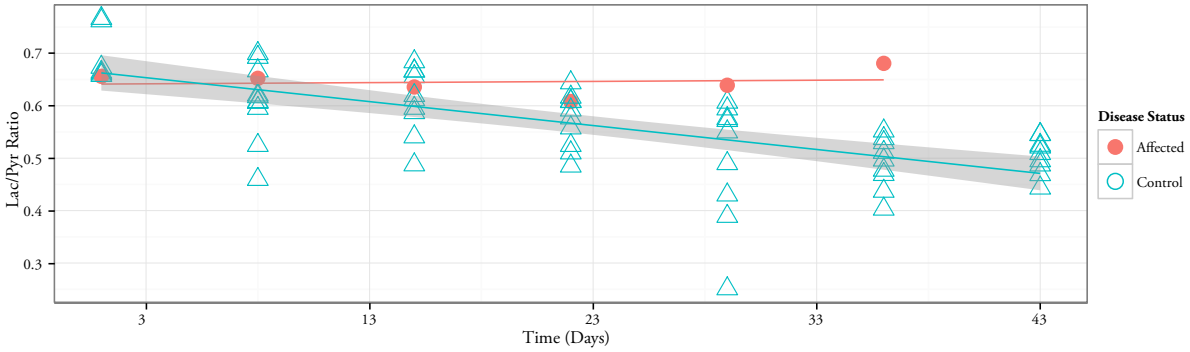

Figure 7.6: The timecourse of the apparent lactate-to-pyruvate ratio; the rat with disease is shown in red. A decreasing trend for all except the affected animal is apparent, and a linear model fit to both populations is shown, with uncertainty shaded in grey. It is believed that this decreasing trend in control animals represents the inflammatory response to injection.

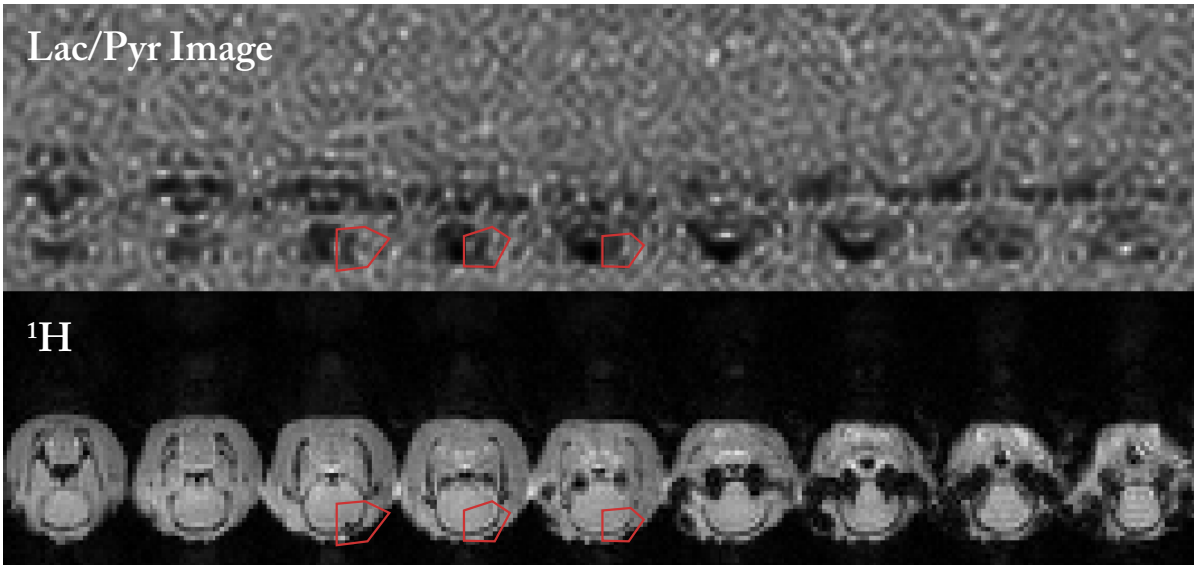

Figure 7.7: Illustrative lactate-to-pyruvate ratio image obtained from a control animal. The expectation value of this image in the presence of noise is one; regions of pyruvate perfusion therefore appear black. The poor SNR of the acquired data is apparent. The red polygons represent the segmentation ROIs drawn in the region of the injection, as can be identified from ^{1}H images.

variability arising from measurement error as opposed to biological variation between animals. The SNR in acquired lactate-to-pyruvate ratio maps was poor, as illustrated (with a segmentation ROI) in Fig. 7.7. This poor signal ensured that the the apparent lactate-to-pyruvate ratio was not greater than one in any voxel present across every dataset acquired in the regions where signal was observed. This fact validates the normalisation strategy used in computing the average ratio within an ROI.

7.2.2 *Cell Experiments*

With the exception of the normal human astrocyte control, it was possible to culture three different cancer cell lines *in vitro*, and grow the $\sim 10^6$ cells required for hyperpolarised experiments. All normal human astrocyte samples given to the author died immediately after thawing, and were not viable. It was noted that the U251, ENU 1564 and CTC cells proliferated readily in culture, as illustrated in Fig. 7.9.

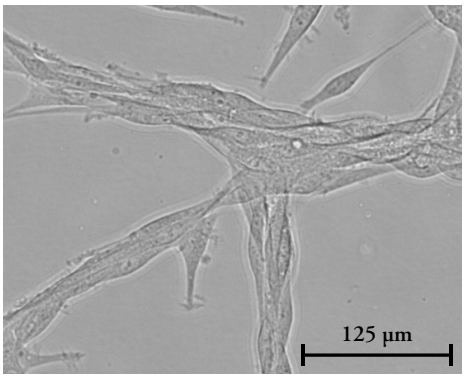

Figure 7.9: ENU 1564 cells grown in culture.

Following the injection of hyperpolarised [1-^{13}C]pyruvate onto living cells inside the spectrometer, it was possible to observe the small rate of production of visible [1-^{13}C]lactate, along with [1-^{13}C]pyruvate and [1-^{13}C]pyruvate hydrate formed in the dissolution buffer. Figure 7.10 shows the apparent lactate/pyruvate ratio per cell for the four different cell lines studied, with errors propagated in quadrature. Across the four different cancer cell lines, the lactate/pyruvate ratio per cell was consistent, with a median value of 1.6×10^{-9} per cell, and an inter-quartile range of 1×10^{-9}, despite the fact that the lines studied were from a variety of species and primary organs. The ENU 1564 cell line used in the *in vivo* experiments were not as glycolytic as the U251 glioblastoma line, which are consistently reported as being highly glycolytic.[172]

It was possible to resolve metabolism with hyperpolarised [1-^{13}C]pyruvate with as few as $\sim 10^6$ cells via spectroscopy with a non-selective 10° flip in this experiment. An estimate of the cell density in a tumour was obtained via manual cell counting in histology images (provided kindly by James Larkin); with ENU 1564 cells, a rough estimate is $\sim$2.8 $\times$ 10^6 cells/mm^3, including a rough correction factor for the non-uniform cell density within a tumour. This tentatively indicates that, if coil effects are neglected, it may be just possible to observe the lactate production of a $\sim$1 mm^3 tumour.

7.2.3 *Distortion correction*

The distortion correction additions to the reconstruction pipeline discussed in Section 7.1.3 were able to adequately remove residual distortions in the hyperpolarised image stacks believed to originate

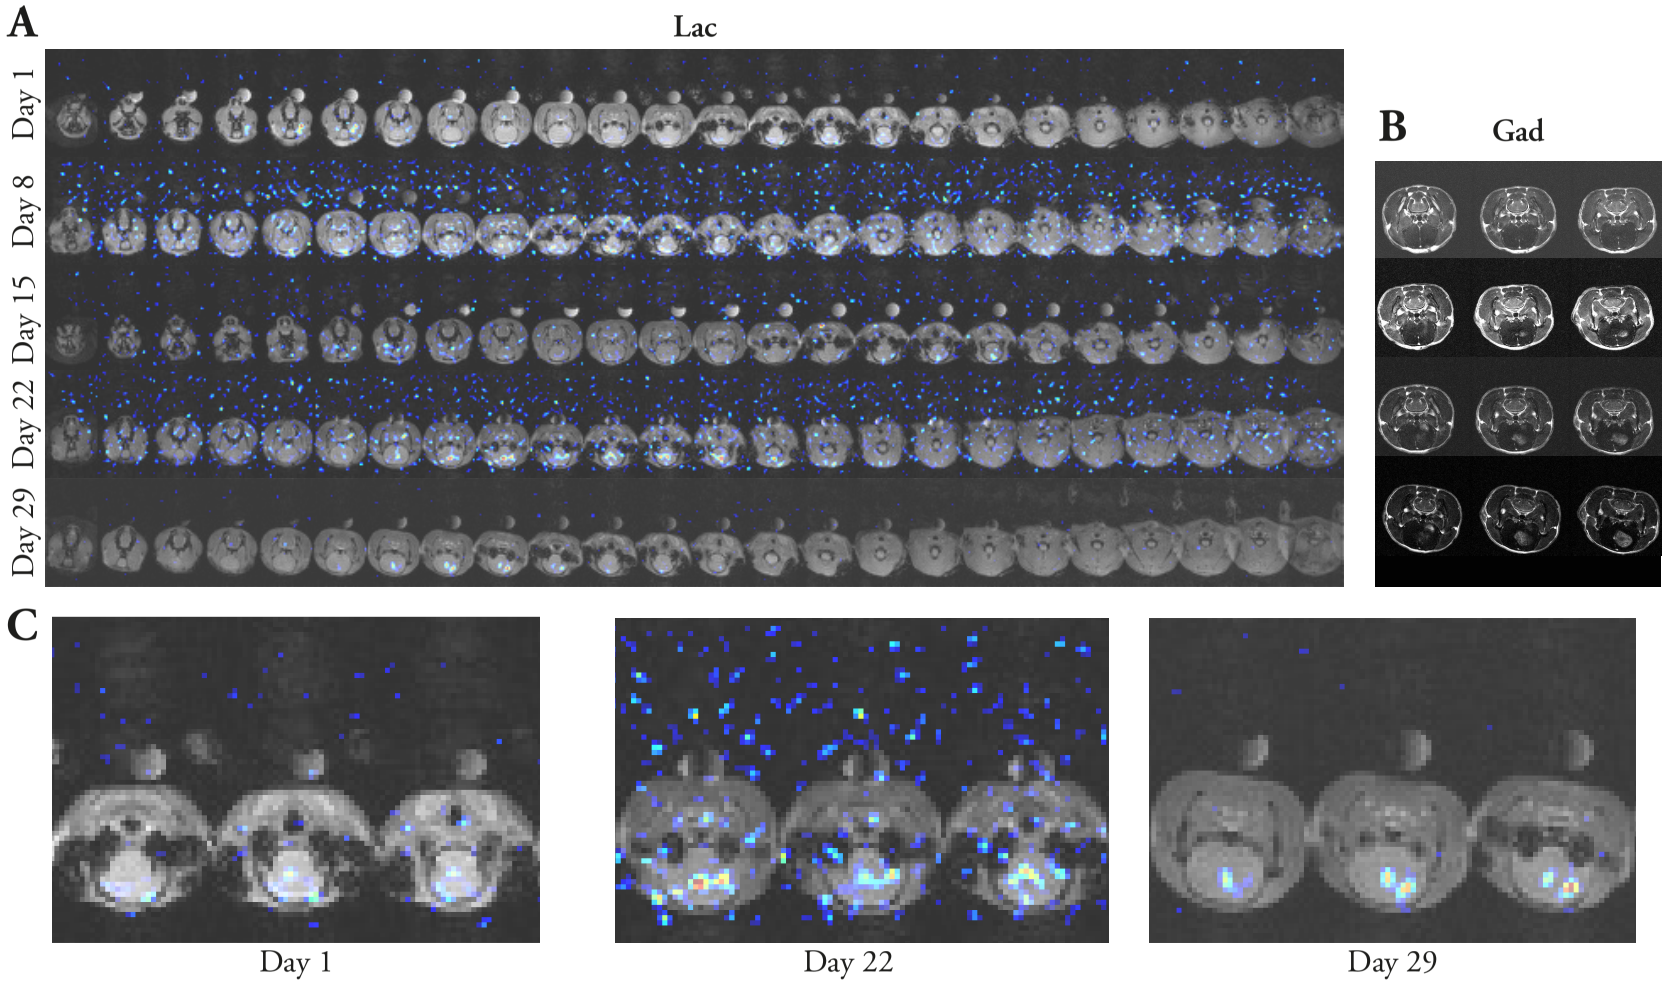

Figure 7.8: The time series of data acquired for the rat that developed a tumour. **A**: Complete concatenated received hyperpolarised signal across each time point. The animal was sacrificed by Sébastien Serres on day 29. Lactate production is visible with improved SNR at the start of the experiment, post-operatively, and when the tumour had grown to be visible, approximately from day 15. **B**: Pre-and-post gadolinium subtraction images in the region of disease, on the same time points as **A**. **C**: Zoomed-in images of the region of tumour one day post surgery (where lactate is seen in the region corresponding approximately to the ventricles of the brain and near the injection site), 22 and 29 days post injection, where the tumour becomes more clearly delineated.

from residual B_0 inhomogeneity as assessed by B_0 maps in the region of interest. This inhomogeneity was particularly larges near the region of the sinuses of the rat, and was believed to be exacerbated by the presence of numerous air-tissue interfaces in the layrngo- and oropharynx near the rat brain. The algorithm was found to perform better in a high SNR environment, but the use of regularisation parameters in `topup` (which serve to penalise excessive deformation of the registration grid) allowed adequate performance in a low SNR environment. The observed magnitude of the generated displacement in addition to the bulk shift correction estimated by co-registration (as described previously) was quite small, often less than one voxel throughout the bulk of the image stack. This was sufficient to correct for occasional slight shifts in the location of pyruvate signal arising from B_0 inhomogeneity near the carotid arteries and jugular veins. Additionally, the developed algorithm was also successfully applied to the cardiac data previously obtained in Chapter 6, which features far higher SNR than image stacks obtained in the brain. Representative examples of these effects are shown in Fig. 7.11, for both cardiac and brain data.

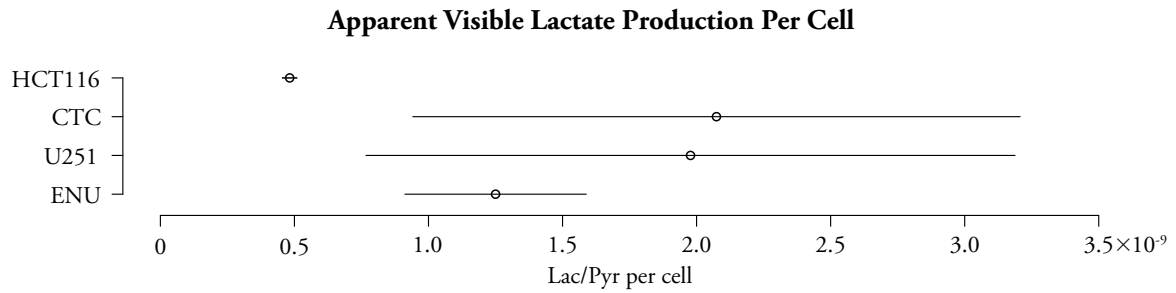

Figure 7.10: The estimated lactate to pyruvate ratio per cell for the four different cell lines studied.



Figure 7.11: The effect of the developed distortion correction pipeline in both the brain and the heart, shown in temporally summed pyruvate images. **A**: Here, the image SNR of the brain data is low, and the estimated distortion effect size is also small. Correspondingly, the algorithm largely returns the input stack, but is able to correct shifts on the order of one voxel where appropriate (e.g. near the middle cerebral artery). **B**: In the acquired cardiac data, these shifts are more noticeable. In particular, the hepatic portal vein (red box) is spatially shifted between alternate frames, leading to a blurring in the resulting summed image. After distortion correction, the circular nature of the lumen of the vessel is recovered. The origin of the artefact is due to be B_0 homogeneity between the heart and the liver, likely exacerbated by the comparatively high iron content of the liver.

7.3 | Discussion

7.3.1 Animal model

This study was designed with the intent to discriminate between healthy and diseased brain by the use of hyperpolarised imaging methods. The animal model used was believed to be a representative model of diffuse disease, as work by others in the field have shown that cells injected intra-cranially behave similarly to those who colonise the brain through the vasculature.[173] Unfortunately, the power of any study with one animal in the diseased group is fundamentally limited, and it was not feasible to repeat the experiment within the time available to the author. Therefore, the lack of tumour in all but one injected animal raises a significant point that needs to be answered.

At least three potential reasons can be proposed as to why disease did not progress in five of the six injected animals. Briefly, these are:

1. The possibility that a host immunological reaction successfully eliminated the cell line injected;
2. Technical problems with the intracerebral injection or surgery, such as the incorrect re-suspension of the cells after pelleting prior to injection; or that
3. Issues with cell viability in culture meant that the majority of those injected were not alive.

Of these possibilities, (1) represents a significant concern owing to the different strain of rat used between this study and that described in Chapter 4. As mentioned, the ENU 1564 cell line originated in an ethylnitrosourea mutagen induced mammary carcinoma in the BD-IX rat, and the Sprague Dawley strain was used in this study. However, evidence from other researchers[5] indicates that the cell line can definitely proliferate *in vivo* under these conditions: rats injected intracranially with 1×10^4 ENU 1564 cells reproducibly grow a $\sim$60 mm^3 tumour after 18 days. It is therefore difficult to know if the reduced number of cells injected into animals here (4000) represents a burden that can be cleared by host immune responses. As there is more biological variability within outbred strains of rats than between them,[174] it is entirely plausible that different unknown immunological responses to the surgical procedure could explain the absence of tumour seven weeks after its apparent induction.

Point (3) above is probably false, as one animal successfully presented with a large tumour after induction in a reasonable timeframe. The time between surgery and the visible presence of tumour would be expected to increase if problems with culture lead to a reduction in the number of cells injected into each animal. As the affected animal had a delay of two weeks between surgery and the presence of disease (which is consistent with that reported by others), this hypothesis is unlikely.

By definition, an unknown problem as inferred by (2) is difficult to examine or exclude *post hoc*. It is entirely plausible that an unknown technical error – such as the inadequate resuspension of the cell line

[5] Personal communications from James Larkin and Sébastien Serres.

following pelleting post-culture – led to the injection of a large number of cells into one animal, and none into the others. Another plausible hypothesis would be that a smaller number of viable cells were successfully implanted than that estimated via hæmocytometery. With unlimited resources, this trial would be repeated, and the reason for the lack of any tumours in five out of the six rats investigated thoroughly.

7.3.2 Physiological interpretation of low SNR

In addition to the lack of tumour in the model, it was noted that the SNR of the data acquired from the brain was poor. In particular, whilst pyruvate could be seen perfusing the major vessels of the brain, relatively little signal was obtained from the region of grey matter itself. Likewise, in healthy animals the majority of the lactate visible was also in the vasculature, and it was only when the tumour was excessively large that lactate production could be observed within the region of disease.

Quantitatively, the total observed lactate signal *in the temporally summed images* was $\sim 70\,000 \pm 1\,310$ DAC units, and maximum pyruvate signal $\sim 82\,500 \pm 8257$ DAC units, in contrast to $\sim 165\,000 \pm 45\,800$ and $210\,000 \pm 14\,400$ DAC units for the single-timepoint cardiac images discussed in Chapter 6 (figures quoted mean $M_0 \pm$ SD across $n = 6$ experiments, where M_0 is the first point of the acquired FID averaged over time, divided by the number of frames averaged; gain and flip angle are nominally constant between experiments). These figures are not directly comparable owing to the fact that the received signal is effectively convolved with a varying (unknown) coil sensitivity function that will, in general, differ between the brain and the heart, and do not account for varying initial M_0. However, this result does qualitatively suggest that (a) more visible pyruvate is reaching the heart than the brain; and that (b) the total lactate signal (albeit in arbitrary units) is likewise greatly reduced in the brain.

In some sense, this fact may be unsurprising: the injection of hyperpolarised pyruvate into a vein necessitates that the voxel concentration of the contrast agent will be initially highest in the right ventricle of the heart. Likewise, the large degree of perfusion of the heart would be expected to supply it with a high concentration of [1-^{13}C]pyruvate soon after injection, with the contrast agent entering the coronary arteries via the aortic root after its first passage through the cardiovascular system. However, the apparent quantitative reduction in healthy visible metabolism in the brain compared to the heart is likewise surprising: in rats at rest, the brain consumes twice as much oxygen as the heart per minute,[175] and has a basal rate of glucose consumption approximately six times as high.[6] It is therefore surprising that in a region of known high glucose consumption the apparent exchange between [1-^{13}C]pyruvate and [1-^{13}C]lactate is not higher. If, instead of equilibrating with [1-^{13}C]lactate via LDH, pyruvate taken into cells was dehydrogenated to form acetyl-CoA, an apparent visible ^{13}C-bicarbonate flux

[6] This figure is based upon two studies reporting ^{18}F-FDG PET kinetic uptake curves in the *ex vivo* perfused heart[176] and the *in vivo* brain[177], with an estimate of the wet weight to dry weight ratio of the normal rat heart provided elsewhere.[178] The healthy heart preferentially consumes glucose and fatty acids; the figure quoted is based on an experiment performed with physiological quantities of both glucose and oleate.

would be apparent, and none was observed. Likewise, if [1-^{13}C]pyruvate was carboxylated by pyruvate carboxylase, an [1-^{13}C]oxaloacetate resonance would be expected to be observed as a shifted, blurred artefact in the pyruvate channel owing to the periodic nature of the spectral-spatial pulse used for excitation. Likewise, as pyruvate was observed co-localised with the internal vasculature on the surface of the rat distal from the coil used, it is unlikely that hardware issues such as coil sensitivity were responsible for this comparative lack of signal.

Given the fact that the designed sequence has been shown to function with high signal-to-noise and specificity at this spatial resolution in the heart, and that the majority of cellular fates of pyruvate within the brain are expected to be visible, one plausible explanation for the lack of visible metabolites within the brain is that transport is limiting. As discussed in Section 4.3.2, the blood-brain-barrier (BBB) is formed by a physical layer of endothelial cells that surround the capillary bed. In order to reach neuronal tissue, pyruvate would be expected to be transported via members of the monocarboxylate (MCT) transporter family through both sides of the endothelial cells forming the BBB. In tumours, the up-regulated presence of MCTs is well documented,[179] and is believed to be a requirement to enable the effective extracellular transport of the lactate that they endogenously create. This fact is consistent with the visible presence of lactate at the end-stage of the disease, where the BBB had been disrupted through angiogenesis, allowing a locally increased rate of pyruvate transport. It should be noted that there are reports of using hyperpolarised [1-^{13}C]pyruvate to probe cancer metabolism in neurology. For example, Park *et al.*[180] report the use of the lactate-to-pyruvate ratio as a means to monitor the response of induced glioma to chemotherapy, specifically the DNA alkylating agent Temozolomide. It was found that the transport of hyperpolarised [1-^{13}C]pyruvate through into the ($\sim 200\,\mathrm{mm}^3$ and hence relatively large and angiogenic) tumours studied was not insignificant, and that predictive decreases in the lactate-to-pyruvate raito in response to Temozolomide could be observed.

However, in contrast to the pathological situation defined by a large disease burden, physiological transport kinetics are believed to be limiting in the brain. In the perfused healthy rat heart, the maximum transport rate of pyruvate through MCTs (assumed, perhaps erroneously, to obey Michaelis-Menten kinetics) as found via radio-labelled ^{14}C pyruvate delivery is reported to be $V_{\mathrm{max}} = 7.7\,\mu\mathrm{mol\,min^{-1}\,g^{-1}}$ to $8.4\,\mu\mathrm{mol\,min^{-1}\,g^{-1}}$ of wet tissue.[181] This should be contrasted to the reported V_{max} in the normal adult rat brain for radio-labelled pyruvate net uptake of $0.18 \pm 0.05\,\mu\mathrm{mol\,min^{-1}\,g^{-1}}$ of wet tissue.[182] It seems therefore highly plausible that the low observed SNR in the brain is due to the lack of effective transport of pyruvate through the BBB. Thus, this attempt at high-resolution pyruvate imaging may well have exhausted the visible pyruvate magnetisation within the brain tissue itself.

7.3.3 *Disease progression*

It is worth nothing that the apparent lactate signal within the region that was to become the tumour was not increased at the 15 day time point, despite the fact that it is believed disease would already be there. This fact is consistent with transport being a limiting factor in the study: by the time the tumour is visible to gadolinium contrast, dysregulated angiogenesis has occurred and the BBB is disrupted. Therefore, pyruvate is believed to enter the region of tumour more readily from the vasculature, leading to a greater visible flux of pyruvate through to lactate. This suggests an unfortunate truth: small metastases are difficult to image, and the accessible flux of hyperpolarised [1-^{13}C]pyruvate is not sufficiently large to visualise them. By the time the disease is large (but before the core becomes necrotic and poorly perfused), hyperpolarised [1-^{13}C]pyruvate is indeed able to image the known metabolic phenotype of the disease. However, by this point the tumour will be visible by other means. It is therefore unsurprising that the majority of oncology studies that have been published to date use relatively advanced forms of cancer, in which the perfusion of the contrast agent is not significantly limited, and the focus of the work is on monitoring response to treatment (in the form of radiotherapy or chemotherapy) effectively. Whilst it is reassuring that the technique here was able to identify large cancerous regions as being glycolytic, this does not represent a novel result.

7.4 | *Conclusion and Future Directions*

This chapter has described a study that failed for at least two distinct reasons. Firstly, the animal model used did not behave as expected: instead of forming a repeatable model of cancer that has metastasised to the brain, only one rat out of six injected developed the disease. This undeniable fact necessarily limits the certainty of conclusions that can be drawn from the one affected rat studied. In short, the only statement that is reasonably supported by the work presented is that the technique shows that large tumours are glycolytic.

Secondly, the data acquired had very low SNR when compared to the cardiac images acquired previously. It is highly likely that this apparent lack of visible signal in the brain is physiological in origin. Indeed, after this study was performed, the author was introduced to the work of Hurd *et al.*,[183] which describes precisely this finding, namely that the SNR of hyperpolarised [1-^{13}C]pyruvate in the brain is poor. Hurd *et al.* note that the likely concentration of [1-^{13}C]pyruvate in the brain following the injection of 80 mM hyperpolarised [1-^{13}C]pyruvate (as used here) is no more than 1 mM. The problem of partial volume effects arising from the large blood vessels close to the region of interest are also analysed, and found to represent a significant concern. To perform [1-^{13}C]pyruvate imaging with a CSI sequence and obtain reasonable SNR, Hurd *et al.* dramatically reduced the spatial resolution of the acquisition to $2.5 \times 2.5 \times 10$ mm^3. Naturally, it is expected that decreasing the spatial resolution used here would improve SNR, but also greatly increase problems due to partial volume effects and reduce

the probability of detecting small regional differences in lactate production in the early phase of disease.

As an alternative to pyruvate imaging, Hurd *et al.* investigated the metabolism of hyperpolarised [1-^{13}C]-ethyl-pyruvate, a lipophilic analogue of pyruvate that is believed to have faster transport kinetics across the BBB, as it is more likely to be able to diffuse across the plasma membrane of the endothelial cells forming the BBB directly, as well as being transported by MCTs. Ethyl-pyruvate is a mild anaesthetic, and has been shown to have an anti-inflammatory action that acts by inhibiting NF-κB, a transcriptional regulator of many immunological protein targets such as TNF,[184] but its wide-ranging pharmacological effects are currently unknown. It was found by Hurd *et al.* that [1-^{13}C]ethyl-pyruvate is indeed readily transported through the BBB more quickly than [1-^{13}C]pyruvate, and that the visible degree of [1-^{13}C]lactate production in the healthy rat brain is also correspondingly increased. It is of note that, owing to the reactivity of ethyl-pyruvate with the trityl 'Finland' radical used for hyperpolarisation, at least two labelled unknown contaminants were present (and highly visible) in the data acquired by Hurd *et al.* Additional difficulties were met owing to the fact that the high concentration of NaOH used for dissolution reacted with ethyl-pyruvate, evolving a variable ratio of labelled ethyl-pyruvate and pyruvate that was used for injection into each rat.

If the spectral density of the hyperpolarised [1-^{13}C]ethyl-pyruvate could be reduced by the removal of contaminants, and the dissolution performed in such a fashion as to minimise the production of pyruvate rather than ethyl-pyruvate, it is likely that the spectral-spatial sequence described in Chapter 6 would be able to function with little modification to image the metabolism of [1-^{13}C]ethyl-pyruvate *in vivo*. As substantially faster transport kinetics are reported for ethyl-pyruvate, and higher SNR is observed in the work of Hurd *et al.*, this may well represent a promising avenue of research to explore. It may be hoped that a repeat of this study that used ethyl-pyruvate (and where the animal model accurately recapitulated the disease) would be able to image cancer that had metastasised to the brain through its metabolic phenotype at an earlier time point than angiogenesis.

Chapter 8

Hyperpolarised Urea and Cardiac Perfusion Imaging

8.1 | Introduction

8.1.1 Overview

Urea is a small molecule that forms the end point of amino acid catabolism in eukaryotic cells. It is uncharged, pH neutral, highly soluble in water, and, with an LD_{50} of ~ 15 g/kg in rats, exceptionally non-toxic. As such, it can be considered a 'metabolically inert' probe of perfusion. This chapter details the development and validation of hyperpolarised ^{13}C-urea as a novel probe of cardiac perfusion, although it is expected that the technique would be able to function elsewhere in the body. The use of bipolar flow-suppression gradients enables the visualisation of ^{13}C-urea in the myocardium, even where the SNR-limited achievable spatial resolution is conventionally insufficient.

The technique is performed under conditions of rest, and cardiac stress induced by the inotropic drug adenosine, which acts as a coronary vasodilator, and increases cardiac output but not heart rate. Myocardial perfusion, as quantified by the technique, was seen to approximately double between rest and stress, which is consistent with the reported value for the perfusion reserve of the rodent myocardium.

8.1.2 Hibernating myocardium

As outlined in Section 4.2.2, the main clinical question in the treatment of patients with coronary artery disease is the identification of the viable hibernating myocardium (which forms a target for revascularisation therapies) and the successful discrimination of it from necrotic (and hence non-viable) tissue.[185] Such hibernating myocardium is characterised by a metabolism-perfusion mismatch, with intact metabolism and impaired contractility due to reduced myocardial perfusion.[186] If blood flow is restored to the viable tissue, a subsequent return to normal function is likely, as is a positive clinical outcome. However, the attempted revascularisation of non-viable tissue constitutes an unnecessary invasive procedure with the associated cost to the healthcare provider and an unnecessary clinical risk

for the patient. Such revascularisation attempts are futile and are associated with a small but non-zero probability of significant iatrogenic pathology after the event.[187] There is consequently a need to assess both cardiac metabolism and myocardial perfusion in the ischæmic heart to guide the clinical decision-making process leading to the instigation of revascularisation therapies on post-MI patients. Currently, a variety of perfusion imaging techniques such as Computed Tomography (CT), MR and X-ray angiography, are used to assess myocardial perfusion deficits, but acute metabolic imaging in the heart is not usually performed.

8.1.3 Visualising perfusion

8.1.3.1 X-Ray Angiography

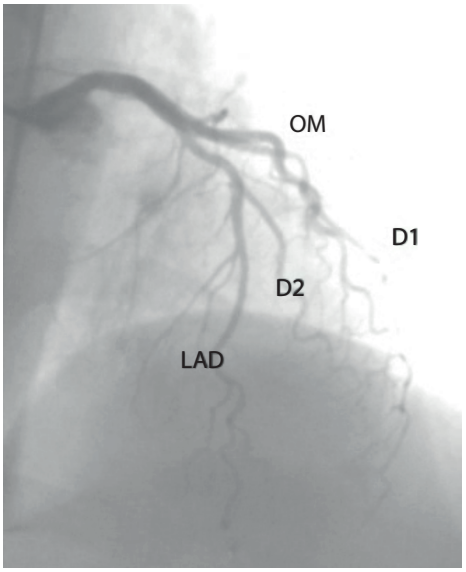

Figure 8.1: An example X-ray angiogram, showing the left anterior descending (LAD), diagonal (D1, D2) and obtuse marginal (OM) divisions of the left and right coronary arteries. Open image CC-BY-NC-SA courtesy of Aquilina *et al.*[188]

Owing to the comparative lack of equipment required, X-ray coronary angiography has provided the main method of visualising the coronary arterial system *in vivo* since the 1960s. The procedure requires the cannulation of a major artery, typically the femoral, radial or subclavian artery, before a small (~ 2.0 mm diameter) catheter is inserted and moved under continuous X-ray guidance towards the heart. After locating the catheter in an appropriate position, an X-ray contrast agent, typically containing radio-opaque iodine, is injected. The flow of blood then moves the dye and highlights obstruction, aneurysm, or other elements of vascular disease.

X-ray angiography is an invasive procedure, and is commonly performed concomitantly with catheter-delivered revascularisation therapies. By the time it is performed, the decision to revascularise is often made; for example, NICE guidelines in the UK mandate that patients presenting with ST elevation and other symptoms of acute MI are given primary percutaneous coronary interventional therapy within 150 minutes of presenting with symptoms (the so-called "call to balloon time"). As an acute interventional procedure, X-ray angiography is not without risk. The incidence of major adverse events (death, myocardial infarction, or stroke) within 24 hours of coronary angiography is reported to be 0.2% to 0.3%, and the incidence of minor complications (most of which arise from injury to the peripheral vessels through which catheters are inserted) is roughly 1% to 2%.[189] As a projection imaging technique, X-ray angiography acquires spatial information predominantly in two dimensions (occasionally with different projections obtained by

moving the patient, X-ray source, or both), which may hinder the interpretation of the resulting data. An example X-ray angiogram is shown in Fig. 8.1. As the dye injected dissipates into and dilutes in the blood stream, conventionally information obtained via catheter-based coronary angiography pertains to the coronary arterial lumen alone, unless the catheter is placed directly in the pulmonary artery. No information is gained about the state of myocardial tissue. Finally, as an ionising radiation based imaging technique, X-ray angiography is fast, high-resolution, but fundamentally deleterious to the health of the patient receiving it. For these reasons, it is not the first-line choice for perfusion imaging except for patients with substantial cardiac complaint, where interventional catheterisation (e.g. balloon angioplasty) is likely to be indicated.[190]

8.1.3.2 Computed Tomography Angiography

Computed Tomography angiography represents, to some extent, the logical extension of X-ray angiography to three dimensions. CT angiography typically requires the injection of a relatively large amount of iodinated contrast agent, usually 160 mL administered at a rate of 4-5 mL/s through an antecubital vein (as opposed to requiring direct major arterial access). Beta-blockers and/or sublingual nitrates (e.g. nitroglycerine, known to patients by the name 'glyceryl trinitrate') are often administered before the injection of contrast to lower the heart rate, avoid arrhythmia and dilate the coronary arteries. Acquisition is typically prospectively gated to the cardiac cycle, sacrificing information about the function of the heart for a substantially decreased radiation dose, producing images such as Fig. 8.2. However, the inferior soft-tissue contrast found with CT means that certain patient subgroups, such as those with high levels of calcium, are prone to misdiagnosis following coronary CT angiography due to arterial calcification.[192] CT angiography is able to resolve the full coronary arterial tree in a three-dimensional region of interest with sub-millimetre resolution.[189] However, this resolution comes at a cost: due to the dynamic acquisition lasting several seconds, CT angiography exposes patients to a substantial radiation dose that is higher than X-ray angiography or the majority of anatomical CT sequences.[193] For this reason, CT is contraindicated in pregnant women, children and young adults (due to the teratogenic and mutagenic effects of radiation),[194] and not used for longitudinal observations within a patient. Additionally, the iodinated contrast agents used in CT have a comparatively worse safety profile than gadolinium chelates, with reported rates of adverse effects (estimated as being between 0.11%[195] and

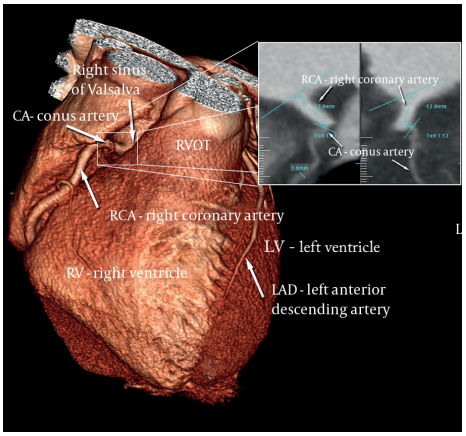

Figure 8.2: An example CT angiogram, showing the gross anatomy of the heart. Open image CC-BY-NC courtesy of Mlynarska *et al.*[191]

6%[196]) higher than the range of $3 \times 10^{-4}\%$ to 1.5% that is reported for the gadolinium chelates used in MR. The high solubility of CT contrast agents again means that disease is visualised by the impediment of their flow, and a narrowing of a coronary artery has to be directly visualised.

8.1.3.3 Magnetic Resonance Methods

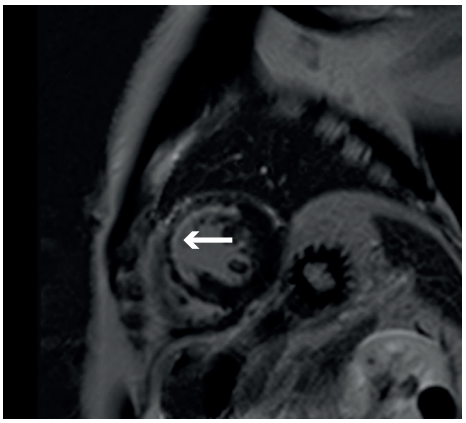

Figure 8.3: Late gadolinium enhancement clearly delineates the region of this transmural bright anterior and partially septal infarct. Image personal communication from A. Lewis.

MR provides two unique methods to image perfusion within the heart, and as cardiac MR exams form the 'gold standard' to quantify cardiac function,[197] these methods are increasingly used as a matter of routine. Like ionising techniques, it is possible to perform contrast enhancement scans with MR, in which a contrast agent (typically a gadolinium chelate) is injected. If this injection occurs whilst a T_1-weighted imaging sequence with high temporal resolution is run, the arterial tree can be directly visualised. If a period of time elapses between injection and imaging, the contrast agent 'leaks' out of the vasculature and into interstitial spaces in regions of disease, such as fibrosis following infarction, and these 'delayed enhancement' scans are routinely used clinically to quantify regions of damage following MI. Such delayed enhancement scans, an example of which is provided in Fig. 8.3, are an excellent probe of fibrosis and report on the state of perfusion in the vasculature, rather than arterial flow. As rodent models of MI are typically used as a reproducible system to study the disease, blood flow quantification with a gadolinium bolus has been demonstrated as an accurate probe of perfusion in both rodent[198–202] and later human hearts.[203–205] Magnetic Resonance angiography with exogenous contrast agents is as diagnostically accurate as 64-slice CT angiography, take a similar length of time, but are associated with an enhanced safety profile, and an eliminated radiation dose.[206]

However, the use of gadolinium chelates as perfusion agents is contraindicated in several conditions, notably renal impairment, which are often co-morbid with cardiac disease. The safety profile of gadolinium chelates, while generally agreed to be preferable to that of iodinated contrast agents,[195] is uncertain, as a number of patients given relatively large doses of gadolinium for MR angiography have subsequently developed fatal nephrogenic systemic fibrosis[207–209] leading to the complete contraindication of contrast by many health service providers for patients with impaired renal function. Accordingly, the prescribed conditions for which contrast is permitted are limited; in the UK, the MHRA licences the gadolinium preparation Magnevist for use in spine and brain imaging only in adults, and whole-body imaging in children. In the USA, every clinical contrast-enhanced MR exam represents an

off-label prescription by the attending physician. As a contrast agent, gadolinium is effective, but it is associated with a non-zero risk.

As a consequence of the potential risks of the administration of contrast agents, other non-contrast MR methods have been developed to non-invasively quantify perfusion in the human or rodent myocardium. Primarily, these are Arterial Spin-Labelling (ASL) techniques[211–213], in which spins are RF tagged at an upstream arterial site before imaging; phase contrast imaging techniques, where a pair of bipolar velocity encoding gradients are used to provide rapid 2D imaging of a component of velocity; and, more infrequently, flow visualisation sequences such as "4D flow",[214] where the number of velocity encoding gradients used is sufficient to fully map the vector field of blood velocity, as illustrated in Fig. 8.4. Both ASL and 4D flow are comparatively slow with long scan times (on the order of tens of minutes), possess substantially worse spatial resolution than CT, and present significant post-processing requirements. Such limitations may make these methods difficult to integrate into a larger (pre-)clinical MR study of cardiac function, and, as a consequence, they have not found routine use in the setting of coronary angiography.

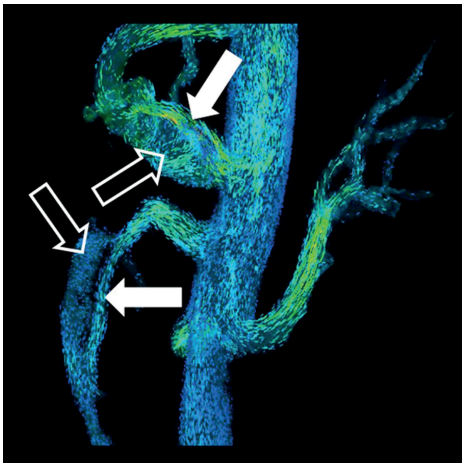

Figure 8.4: 4D Flow imaging of aneurysm in the cœliac and superior mesenteric arteries, which branch from the abdominal aorta. Laminar flow is present in the true lumen (closed arrow) and helical flow is present in the false lumen (open arrow). Image courtesy of Hartung *et al.*[210]

8.1.3.4 Hyperpolarised perfusion imaging

It has long been realised that DNP provides a "third way" of imaging perfusion, without the potential toxicity of administering gadolinium chelates or the long scan times of endogenous proton techniques. DNP forms a favourable technique for perfusion imaging due to the relatively long signal lifetime (compared to the cardiac cycle) and low background signal of hyperpolarised agents compared to proton techniques; additionally, the linear signal response of the agent with respect to concentration also enables direct quantification and the ability to detect balanced ischæmia, where a global change in blood flow occurs in the organ of interest. These facts represent a substantial benefit compared to traditional contrast enhanced magnetic resonance angiography[215–217]

Prior proof-of-concept studies demonstrating perfusion imaging with hyperpolarised compounds were performed using ^{13}C-labelled urea,[217] an endogenous metabolite, and subsequently both ^{13}C-urea[216] and the non-endogenous extracellular agent known as HP001[218][1] have been used as probes of

[1] Or, systematically as bis-1,1-(hydroxymethyl)-[1-^{13}C]cyclopropane-D8

perfusion within the kidneys, in models of cancer, and in the hearts of large animals.[219] DNP with endogenous and non-toxic ^{13}C-urea therefore provides an attractive alternative probe of perfusion.

It is technically challenging to perform perfusion imaging of the rodent heart due to the high spatio-temporal resolution required to resolve the first pass of the perfusion agent through the heart and coronary vasculature. The typical thickness of the rodent myocardium is only 2 mm, and this, in conjunction with the high signal present in the lumen of the ventricles, necessitates high spatial resolution to adequately resolve perfusion agents within the tissue capillary bed. If the scan run is not of sufficiently high resolution, partial volume effects arising from the neighbouring blood pool will be significant, confounding the quantitation of myocardial perfusion. It has been previously demonstrated that the intracoronary infusion of a hyperpolarised contrast agent after coronary catheterisation allows the direct visualisation of the signal from within the capillary bed,[219] but this procedure is highly invasive and technically infeasible for small animal imaging.

As a novel alternative, the feasibility of imaging the first-pass of a bolus of hyperpolarised ^{13}C-urea through the rat heart by using flow-sensitising bipolar gradients to reduce the signal within the cardiac chambers is investigated below. This method improves the contrast of the agent within the tissue capillary bed of the myocardium, without increasing the resolution of the scan to such an extent that SNR becomes unacceptable. This new method was demonstrated in the healthy *in vivo* rat heart during rest and under stress provided by the inotropic drug adenosine, and is expected to enable new hyperpolarised ^{13}C studies in which the cardiac metabolism/perfusion mismatch can be quantified following myocardial ischæmia.

8.2 | Methods

8.2.1 Pulse Sequence

8.2.1.1 Spiral k-space trajectory

As only one resonance – that of ^{13}C-urea – was expected for this sequence, a more temporally efficient readout trajectory was used for imaging each slice. 'Spiral-out' Archimedian trajectories are well known to be very efficient, in the sense that they maximally use the available readout time to acquire data, starting at the origin in k-space at early times (where signal intensity is high), defining the low-frequency features of the image before traversing out and increasing resolution with the remaining transverse magnetisation.

All spiral readouts are subject to sampling requirements if they wish to avoid aliasing. Ideally, in 2D polar coordinates, a spiral k-space trajectory is $k(t) = \lambda\theta(t)$ where λ is a constant, and $\theta(t)$ the positive semi-definite angle from the x-axis of k-space being described in a $(k,\, \theta)$ 2D polar coordinate system. The constant λ can be determined by the Nyquist sampling requirements. If N interleaves are

placed in this radial distance, the spacing between adjacent interleaves at any fixed azimuthal angle is $\Delta k = 2\pi\lambda/N$. For alias-free sampling, one must have $\Delta k = 1/L$, where L is the FOV. Hence:

$$\lambda = \frac{N}{2\pi L}, \tag{8.1}$$

independent of the angle of rotation of each shot. This forms a minimal radial sampling requirement, between subsequent traversals of the spiral, as illustrated in Fig. 8.5. A corresponding azimuthal requirement exists, stating that the distance between successive points along each path has to be determined by

$$\Delta\omega = \frac{\gamma}{2\pi} G_{\max} L = \frac{1}{\text{Sampling interval}} \tag{8.2}$$

where $\Delta\omega$ is the full bandwidth of the receiver.

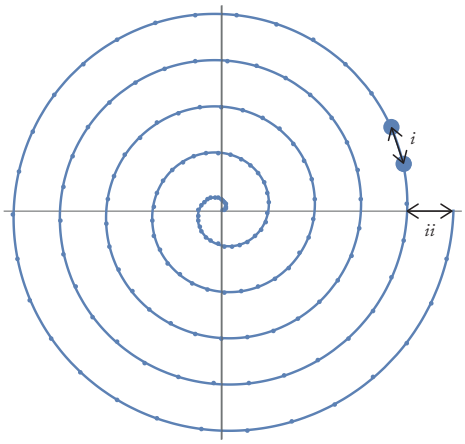

Figure 8.5: An Archimedian spiral trajectory (given by $r = 3\theta$) with azimuthal (*i*) and radial (*ii*) sampling intervals illustrated.

The design of $\theta(t)$ subject to these constraints is not trivial, but several closed-form analytic trajectories have been proposed.[7, p. 932] Unfortunately, hardware limitations represent a further additional design constraint for spiral trajectories: the finite gradient slew rate, region of amplifier bandwidth support, and maximum gradient strength provide three parameters that should be considered in the design of a spiral readout. Commonly, either two or three parameter spiral trajectories are used, with design constraints by either maximum gradient strength and slew rate alone, or additionally incorporating a bandwidth requirement. It can be shown[7, pp. 935-943] that the two-parameter spiral case can be formulated as an inhomogeneous second-order differential equation for $\theta(t)$, where the slew rate and maximum gradient constraints are codified as a piecewise-defined function. Two-parameter spiral trajectories are, therefore, easier to define, as they can be generated explicitly in preparation time. If the bandwidth of the gradient hardware is 'sufficient', this approach is able to generate trajectories that are physically realisable.

Owing to the fact that the gradient response was perceived to have poor (i.e. compact) support, the commonly used two-constraint method for generating spiral trajectories was not used. Instead, a three-parameter algorithm written by Michael Gyngell and functionally similar to that described by Orberhammer *et al.*[220] was used, which provides a smooth transition period between the slew-rate limited region of the trajectory and the maximum gradient strength limited period, without requiring gradient frequencies[2] beyond the capability of the system to reproduce. This addition to `sgl` was installed by the author, and multinuclear support added. The designed trajectory was additionally written out as metadata saved with the acquisition, to enable subsequent reconstruction.

[2] I.e. $\max\left(\frac{\partial}{\partial t}k_x(t),\ \frac{\partial}{\partial t}k_y(t)\right)$.

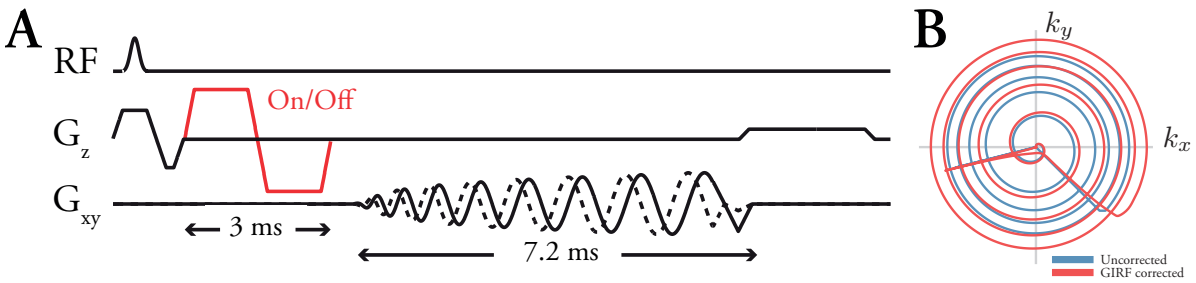

Figure 8.6: **A**: The ECG-gated spiral imaging sequence used for first-pass perfusion imaging of ^{13}C-urea in the heart. The (red) bipolar flow suppression gradient is toggled on or off in alternate TRs, removing signal from fast-flowing blood within the cardiac chambers. The x and y gradient waveforms are shown as solid and dotted lines respectively. This corresponds to the k-space trajectory shown in **B**. The difference between the nominal and GIRF-predicted trajectories is apparent.

It is of note that the reconstruction of spiral images is a non-trivial problem. Spiral trajectories are not used routinely, despite being first proposed in 1981,[221] arguably because of the difficulty in reconstructing artefact free images. Neglecting motion, there are two predominant sources of artefacts in spiral images: nonlinear distortions that arise from gradient infidelities causing substantial differences between the nominal k-space trajectory and that actually traversed; and blurring from off-resonance effects and B_0 inhomogeneity, which smear the acquired data along its nominal trajectory. Knowledge of the gradient fidelity can allow the re-gridding of the data onto a trajectory closer to that believed to be taken; as a form of re-gridding is required to transform the $r(t)$, $\theta(t)$ data acquired to a Cartesian image, this step is not unduly onerous. A variety of refocussing algorithms have been proposed to correct for in-plane blurring, which usually require the accurate prior knowledge of B_0. Chemical shift differences or regions of different magnetic susceptibility represent significant challenges for contemporary spiral reconstruction algorithms, which usually fit linear models to small shifts in k. Owing to the accumulation of ideally unwanted phase, the blurring in spiral readouts increases with the duration of the readout time per interleaf. In general, it is difficult to correct for the blurring that occurs with non-Cartesian k-space trajectories as the phase accumulated in the signal is dependent on both spatial position and k-space position, which is, in the worst case, an uncertain function of time.

It is for these reasons that the author did not implement spiral trajectories in earlier portions of this work, and the reconstruction of spiral data acquired in this chapter was performed by Angus Lau, using methods detailed in Lau *et al*.[138]

8.2.1.2 Readout

The flow-sensitive ECG-gated spiral sequence used to obtain axial dynamic ^{13}C-urea images in the heart is shown in Fig. 8.6. Each interleave was obtained by rotating the previous interleave by the golden angle $\phi = (3 - \sqrt{5}) \times 180^\circ \approx 137.508^\circ$, as this results in near uniform coverage of k-space with adaptive temporal resolution, owing to the fact that a variable number of interleaves can be combined

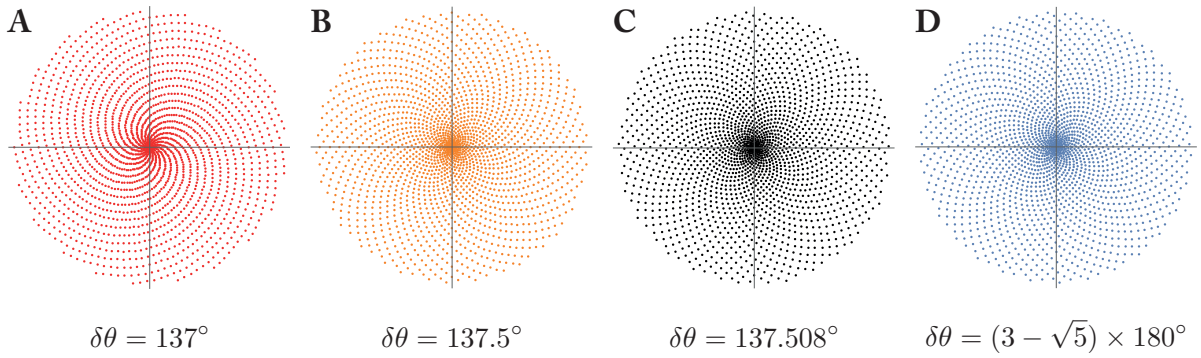

Figure 8.7: An example of the effect of sampling at angles approximating the golden angle. Here, samples are shown at $r, \theta = n, \phi n$ for different approximations of ϕ: $\phi \approx 137°$ (**A**), $\phi \approx 137.5°$ (**B**), $\phi \approx 137.508°$ (**C**), and $\phi = (3-\sqrt{5}) \times 180°$ (**D**) as provided an arbitrary-precision representation of irrational numbers used within Mathematica. Although it is not readily apparent from this graphic, **C** and **D** do differ slightly, but the error is sufficiently small to avoid introducing sampling bias in image space.

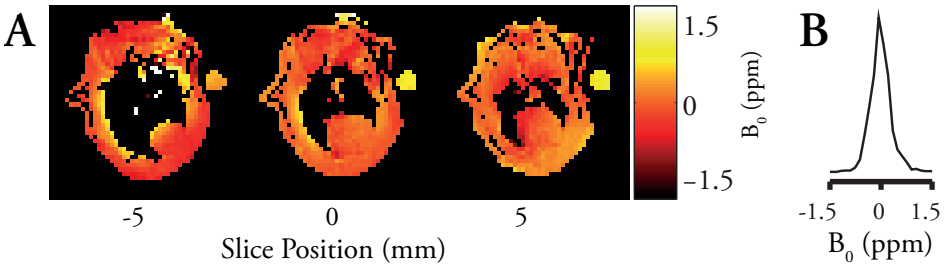

Figure 8.8: The degree of homogeneity achieved over the heart is key to the success of spiral experiments, and here was typically ~ 0.5 ppm.

without introducing sampling bias in k-space.[222] As demonstrated in Fig. 8.7, the angular distribution of samples is approximately uniform with no uniform behaviour if the spacing between samples $\delta\theta = \phi$. Numerically, any physical realisation of this sampling displacement will be a truncated approximation of ϕ. The degree of truncation effectively determines how far out in k space samples are uniform for: here, ϕ was approximated as $137.508°$ as it was believed that this was sufficient to uniformly cover k-space with the resolution obtained. If a higher number of k-space points were obtained, this truncation may well be insufficient to avoid introducing additional artefacts in image space. The readout duration (7.24 ms) was chosen to minimise off-resonance blurring artefacts, based on a proton B_0 field map (shown in Fig. 8.8**A**) acquired in a 1 cm slab encompassing the target slice of interest, namely the ventricles of the heart. After shimming with the method detailed in Section 3.4, the expected line-width *in vivo* was typically 0.5 ppm (Fig. 8.8**B**), or 35 Hz for ^{13}C at 7 T, ensuring that the readout was short enough to avoid image blurring.

8.2.1.3 Flow suppression

To suppress bright signal from the lumen of the ventricles, flow contrast was incorporated by inserting a bipolar gradient after the end of the slice-selective excitation and before the start of the readout (as seen

in red in Fig. 8.6). As detailed in Section 2.4.7, flowing spins accumulate phase proportional to their velocity in the presence of the bipolar gradient, and so spins with different velocities will be dephased within a voxel. The chosen direction of the bipolar gradient was along the long axis of the heart, where the greatest flow component was believed to occur, in order to optimise the amount of dephasing of the flowing spins. The required bipolar gradient amplitude was calibrated by obtaining ECG-gated $^{1}\mathrm{H}$ images with the same spatial resolution as the hyperpolarised scan, and further validated with a flow phantom study, described in Section 8.2.2. The gradient amplitude was chosen by maximising signal suppression of rapidly flowing spins within the lumen of the heart, while minimising signal loss from the myocardium and the blood within this tissue. The required gradient amplitude was then scaled by the ratio $\gamma_{^{1}\mathrm{H}}/\gamma_{^{13}\mathrm{C}}$ to account for reduced phase accrual for $^{13}\mathrm{C}$ relative to $^{1}\mathrm{H}$.

8.2.2 Phantom Study

The performance of the flow spoiling effect was investigated using a flow phantom experiment. A plastic tube (PVC, inner diameter 6 mm) was placed axially through the magnet bore, and through which tap water was driven under steady flow at varying volumetric rates (0, 2, 4, 8, 12, 16 mL/s corresponding to an approximate mean velocity of 0, 7, 14, 28, 42 and 56 cm/s). Water was driven under mains pressure through a restriction valve, and flow rates measured through timed volumetric measurements. The maximum Reynolds number for these flow rates was 1700, which is within the laminar flow regime.[3] Flow encoded transverse $^{1}\mathrm{H}$ gradient echo images were also acquired as a reference (TE $=$ 5 ms, TR $=$ 15 ms, flip angle 7°, field of view 12×12 mm^2, matrix 48×48, acquired in-plane resolution 0.25×0.25 mm^2, slice thickness 5 mm). The bipolar gradient was oriented along the direction of flow as for the hyperpolarised scan, the gradient duration was set to 3 ms, and the gradient amplitude was varied (at 0, 5, 40, 80 mT/m). The gradient first moments (c.f. Eq. (2.71)) corresponding to these values were 0, 11, 87, and 168 mT $(\mathrm{ms})^2$/m, and the associated velocities (VENC $= \frac{\pi}{\gamma M_1}$) corresponding to 180° phase accrual during the bipolar gradient were respectively 105, 13.5, and 7.0 cm/s.

To confirm that the direction in which the bipolar gradients were applied was appropriate for spoiling flow along the pipe, the flow rate of water was fixed at 4 mL/s, and a second set of images were obtained with the flow direction across the image. The bipolar gradient was oriented either parallel or perpendicular to the direction of flow, and the gradient amplitude was varied between 0 and 80 mT/m as stated above, corresponding to first gradient moments of 0 to 168mT $(\mathrm{ms})^2$/m). A velocity map describing flow within the tube in both image orientations was then computed by phase contrast

[3] The Reynolds number R_e for a fluid of kinematic viscosity ν flowing at mean velocity v around a solid object of characteristic size L is defined as $R_e = vL/\nu$, or the ratio of inertial to viscous forces. At low Reynolds numbers (i.e. $R_e \ll 1$; bacteria experience $R_e < 10^{-5}$) flow is essentially laminar, and viscosity cannot be neglected. At very high Reynolds number ($R_e \gg 1$, e.g. a slow sailplane aircraft flies at ~ 150 km/h where $R_e \approx 10^6$; transonic fighter jets experience $R_e > 10^8$) flow is turbulent and, after a small separation layer, the flow can be effectively treated as inviscid. Mammalian blood flow typically has Reynolds numbers in the region where the onset of turbulence can occur, namely 2000 for flow in a pipe, and a deviation from laminar flow is often found in many pathologies.

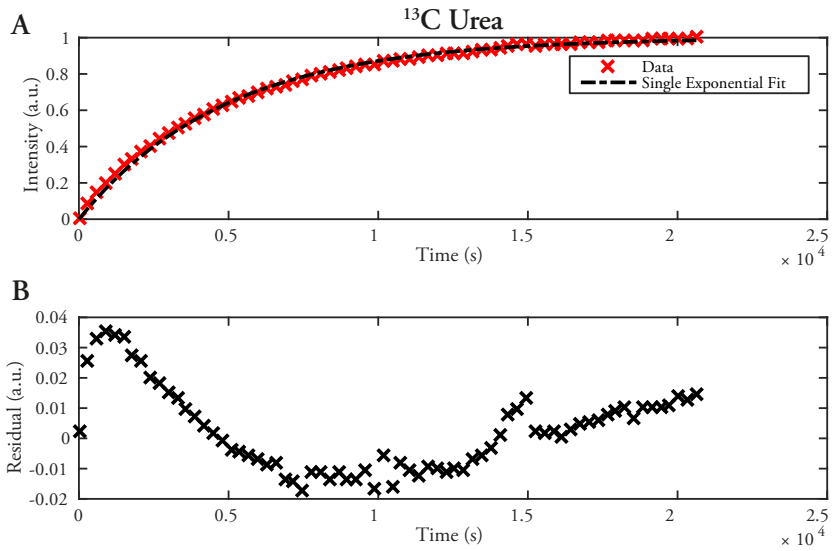

Figure 8.9: **A**: The build-up curve for hyperpolarised ^{13}C-urea, together with a fit to a single exponential. The time constant is approximately 80 minutes. **B**: The deviation of the data from single exponential behaviour is apparent in the residuals of the fit.

subtraction between each pair of velocity encoded images, corresponding to a total effective VENC of 105 cm/s. As the uncertainty of velocities σ_v acquired through phase difference methods in MRI can be shown[223] to be related to the VENC and the SNR in each individual image as

$$\sigma_v = \frac{\sqrt{2}}{\pi}\frac{\text{VENC}}{\text{SNR}}. \tag{8.3}$$

This multi-image approach allows for the reconstruction of flow sensitised images with a higher effective SNR than that which could be obtained by using one VENC alone.

The Hamming window for a discrete data vector N points in length is defined as

$$w(n) = \alpha - \beta\cos\left(\frac{2\pi n}{N-1}\right); \quad \text{where } \alpha = 0.54, \beta = 1-\alpha = 0.46, \tag{8.4}$$

and by obtaining the outer product $w(n) \times w(n)^\top$ a two dimensional symmetric window is generated whose Fourier response is approximately flat outside a central peak. By piecewise multiplying $w(n) \times w(n)^\top$ with the acquired k-space, the exterior region of k-space is attenuated, the image resolution is reduced, and the SNR improved. These acquired proton images were filtered by a two-dimensional Hamming window in k-space to a final resolution of 2.1×2.1 mm^2, degrading their spatial resolution to that of the *in vivo* experiments, described subsequently in Section 8.3.2.

8.2.3 Urea Hyperpolarisation

A stock 6.4 M solution of ^{13}C-urea with 15 mM OX063 trityl radical (illustrated in Fig. 3.16) was prepared in a glycerol/water solution (6:5 ratio by weight) as first described by Patrick *et al*,[224] but with the protocol modified to include more water and less glycerol to reduce the viscosity of the final solution.

Individual samples (50 μL) of the stock solution were mixed with 3.5 μL of a 1:50 dilution in water of gadoteric acid (Dotarem, Guerbet, France) at a final sample concentration of 2.3 mM. Polarisation was performed in the prototype DNP hyperpolariser described in Section 3.5 by irradiating with at 100 mW microwaves at 93.957 GHz for 120 min; as shown in Fig. 8.9 ^{13}C-urea has a long polarisation time constant of approximately 80 minutes, and waiting for 120 minutes corresponds to an acceptable compromise in time, sacrificing approximately 22% of the maximum achievable signal in exchange for the ability to scan more than one animal per session. The centre frequency of 93.957 GHz was found by performing a prior microwave sweep. Dissolution was performed by superheating 5 mL of 0.38 mM EDTA in water to a temperature of $\sim 170^{\circ}$C and a pressure of 10 bar, resulting in a final injected substrate concentration of 64 mM. Due to the short T_1 of ^{13}C-urea in very low magnetic fields (such as the 0.2 mT adjacent to the polariser), the hyperpolarised solution was transferred from the polariser into a low-field (2–10 mT) magnetic holder within 1 s. The holder consists of a plastic bottle with four permanent neodymium magnets attached to the outside, in a quadrupole arrangement. The solution was then transferred to the fringe field (measured by a Hall effect probe as 2 mT) of the imaging magnet before injection. As reported by Shang *et al.*,[225] this strong dependence of T_1 on field arises due to the fact that natural abundance ^{14}N a (quadrupolar) nucleus present next to ^{13}C provides a scalar coupling relaxation pathway that is highly favoured if the Larmor frequency of ^{13}C and ^{14}N are similar.[4] The use of a magnetic holder for transport between polariser and scanner therefore improves the initial signal by $\sim 84\%$ by simply ensuring that this condition is not met. The liquid state polarisation level was 12% at the end of injection, as assessed by a separate thermal equilibrium phantom experiment in which fully-relaxed spectroscopy was performed on the sample and the signal enhancement estimated.

8.2.4 Animal Handling

Six male Wistar rats (mass $= 503 \pm 44$ g) were scanned in the *in vivo* study. Anaesthesia was induced with 2.5-3% isoflurane in oxygen and nitrous oxide (4:1, total of 2 L/min). Anaesthesia was maintained by means of 2% isoflurane delivered to, and scavenged from, a nose cone during the experiment. A tail vein catheter was placed for intravenous infusion of hyperpolarised ^{13}C-urea. A separate tail vein catheter was placed in the contralateral tail vein for delivery of adenosine during the stress-induced hyperæmia experiments, allowing the continuous infusion of adenosine independent of the hyperpolarised contrast agent. As before, animals were then placed in a home-built animal handling system. Body temperature

[4] More precisely, the scalar coupling contributes a relaxation rate R_{SC} given by

$$R_{SC} = \frac{8\pi^2 J^2}{3} I(I+1) \frac{T_1}{1 + (\Delta\omega)^2 T_1^2},$$

where J is the J coupling between the two nuclei in question (in Hertz), I is the spin quantum number of the coupled nucleus (i.e. for ^{14}N $I = 1$), T_1 is the spin-lattice relaxation time for the coupled nucleus (i.e. ^{14}N), and $\Delta\omega$ is the difference in Larmor frequencies between ^{14}N and ^{13}C. Correspondingly, this Lorentzian peak is close to zero at high field.

was maintained using air heating, and a two-lead ECG for cardiac gating was recorded using leads placed subcutaneously into the upper forelimbs, and the time at which each shot occurred recorded. Two mL of pre-polarised ^{13}C-urea at 37° was injected over 20 s following dissolution.

Coronary vasodilation was induced by continuous adenosine infusion (concentration 3 mg/mL, flow rate 280 μg/kg/min) into the tail vein using a small animal infusion pump (Harvard Apparatus, MA). The volume of the injection line was filled with adenosine, and hyperpolarised ^{13}C-urea was administered through the contralateral tail vein catheter 10 minutes after the start of the infusion of adenosine, which was maintained throughout the ^{13}C-urea scan. No adverse effects were detected during or following adenosine infusion, and animals were monitored throughout an appropriate recovery period afterwards.

8.2.5 Data Acquisition

The dual-tuned carbon/proton birdcage and surface receive hardware was used for this study. A volume covering a 70 mm slab including the heart was used for shimming using the 3D gradient echo automated shim routine detailed in Section 3.4. Additionally, a prospectively triggered segmented CINE FLASH sequence was used to assess cardiac function in a transverse mid-ventricular slice (TE 1.36 ms, TR 4.6 ms, FOV 80×80 mm^2, acquired in-plane resolution 0.8×0.8 mm^2, slice thickness 5 mm, flip angle 20°).

As described above, the required bipolar gradient amplitude to null flowing spins in the heart was calibrated using ^{1}H images with the same reconstructed spatial resolution as the ^{13}C acquisition. This ^{1}H bipolar gradient calibration scan was incorporated with multiple gradient echo readout types, and was demonstrated using both Cartesian (TE 5 ms, TR 1 RR, FA 20°, FOV 80×80 mm^2, matrix 128×128, slice thickness 5 mm, 2 averages) and a golden angle spiral trajectory (TE 4.3 ms, TR 1 RR, FA 20°, FOV 80×80 mm^2, 12 shots for full FOV coverage, matrix 128×128, acquired in-plane resolution 0.625×0.625 mm^2, slice thickness 5 mm, 256 total shots, 137.508° rotation between shots, readout duration 3.4 ms). The scan time was the same (40 s) for each acquisition. The data was obtained with multiple averages (Cartesian) or with a large number of interleaves (non-Cartesian spiral) to reduce respiratory motion artefacts.

This ECG-gated golden angle spiral sequence was used to resolve the first-pass through the cardiac cycle of the injected bolus of hyperpolarised ^{13}C-urea with TE $=$ 4.4 ms, TR 1 RR, FA 20°, FOV 80×80 mm^2, four shots for greater resolution and full FOV coverage, matrix 64×64, acquired in-plane resolution 1.25×1.25 mm^2, slice thickness 5 mm, excitation pulse duration 0.5 ms, readout duration 7.24 ms, bipolar gradient amplitude $G_{\text{Flow}} = 170$ mT/m, total duration 3 ms, first gradient moment $M_1 = 332$ mT ms^2/m, maximum encoding velocity (VENC) 14.1 cm/s. As is appropriate for a multi-nuclear experiment, the flow sensitisation effects were kept constant between ^{13}C and ^{1}H by

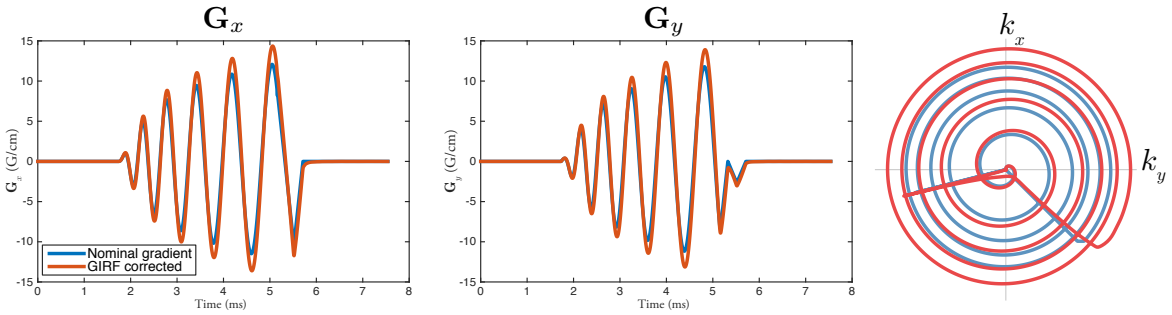

Figure 8.11: The nominal (blue) and predicted actual (red) spiral gradient waveforms in x and y, together with the k-space trajectory taken. Small phase offsets and amplitude distortions serve to make the nominal trajectory and predicted actual trajectory substantially different, which, if not accounted for, would lead to significant distortion in image space.

taking the bipolar flow gradient amplitude determined from the ^{1}H calibration scan and scaling it by approximately four ($\gamma_{^{1}\mathrm{H}}/\gamma_{^{13}\mathrm{C}}$) to account for the reduced phase accrual of one nucleus relative to the other. Based on the calibration scan, a single value for the gradient amplitude (170 mT/m) was chosen for all subsequent experiments. As in Section 6.2.2, a thermally polarised 6.4 M ^{13}C-urea phantom was placed on the side of the coil, which was shimmed to within 1 ppm of the heart. The transmitter central frequency was set at that of the urea reference. The scan was started before injection, and images were acquired for at least one minute after the start of injection, corresponding to at least 400 time points at the typical heart rates in the study. A custom-built large, green "go" button was constructed to interface with the cardiac triggering hardware and start the study momentarily before the injection of the hyperpolarised compound, and is illustrated in Fig. 8.10.

8.2.6 Image Reconstruction

All image reconstruction was performed by Angus Lau. The multi-shot spiral trajectory was predicted using the measured gradient impulse response function, as described in Section 3.1.3. There was a substantial difference in the nominal and predicted actual gradient waveforms, as illustrated in Fig. 8.11. After estimation of the trajectory, the non-Cartesian k-space samples were gridded using the non-uniform fast Fourier transform, which is essentially an interpolative method for efficiently computing the transform of non-uniformly spaced samples.[226] As the density of samples in k-space was not uniform, it was necessary to correct for the different acquisition weightings of k-space with a density compensation function, which was calculated for each set of interleaves by the method of Zwart, Johnson and Pipe.[227] To correct for bulk B_0 shifts in the image, the centre frequency was automatically identified by means of autofocus off-resonance correction by summing all time frames with signal, reconstructing a series of images by demodulating at different frequencies, and selecting the frequency which minimised the imaginary component of the images, as proposed by Man, Pauly and Macowski.[228] This frequency was used for subsequent reconstruction of each individual time frame. A

sliding window was used to group six spiral interleaves per image. The images were Hamming filtered to a reconstructed in-plane resolution of 2.3×2.3 mm^2. The receive coil sensitivities were estimated from two individual coil images obtained by summing together all time frames with signal, and the resulting sensitivities were used to combine the channels in each individual time frame.[149]

8.2.7 *Data Analysis*

In order to quantify perfusion within the myocardium, the heart was manually segmented using the anatomical CINE images into right ventricle (RV), left ventricle (LV), and myocardial tissue. The ^{13}C-urea magnitude images were then segmented with the same map, and the mean ^{13}C-urea signal intensity within each compartment were considered. Baseline correction was performed by subtracting the mean LV signal present in the absence of the hyperpolarised contrast agent, and intensities were then normalised to the maximum signal within the LV. The image signal-to-noise ratio was calculated by dividing the image intensity in each frame by the standard deviation of image intensity in frames before ^{13}C-urea arrival.

The signal time-courses were integrated for each compartment, and an *ad hoc* measure of myocardial perfusion was taken to be the ratio between this integral in the tissue and the LV, removing any polarisation or concentration differences between different injections. Additionally, the integration procedure improves the signal-to-noise ratio in the individual compartment time courses. An unpaired, two-tailed Student's t-test was used to compare the LV/myocardial integral ratio between rest and stress conditions. Statistical significance was considered at the $p < 0.05$ level. The myocardial perfusion reserve was estimated by taking the ratio between the group perfusion ratios at stress and at rest.



Figure 8.10: The user-friendly large green button constructed for starting perfusion scans concomitantly with injecting the contrast agent.

8.3 | *Results*

8.3.1 *Phantom Study*

As shown in Fig. 8.12**A**, the spiral, flow-sensitised sequence was used to image water flowing through a pipe at several different volumetric flow rates corresponding to mean velocities, v. Images were obtained with increasing bipolar gradient amplitudes, and the bipolar gradient direction was oriented parallel to the direction of flow. The signal intensity was successfully spoiled at high flow rates, and

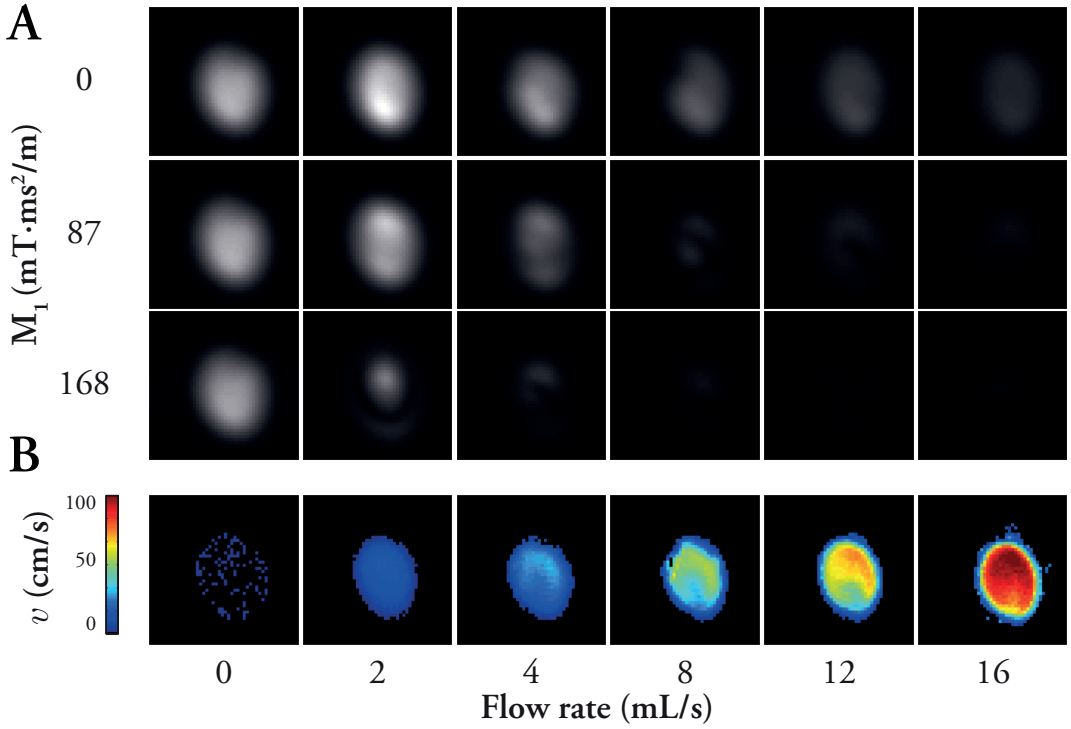

Figure 8.12: The spiral flow sensitised sequence imaging water flowing through a pipe with different gradient first moments (**A**) allowing for the *post hoc* reconstruction of the component of the vector velocity field along the pipe (**B**). Higher volumetric flow rates correspond to higher linear velocities, as expected. The flow is approximately laminar, with the small ring of higher velocity visible at the top of the pipe believed to be due to a previous curve in the flow path in to the magnet. These images are normalised to the image intensity in the non-flow encoded, zero flow rate image (i.e. with water static in the pipe), and are cropped to a 12×12 mm^2 FOV.

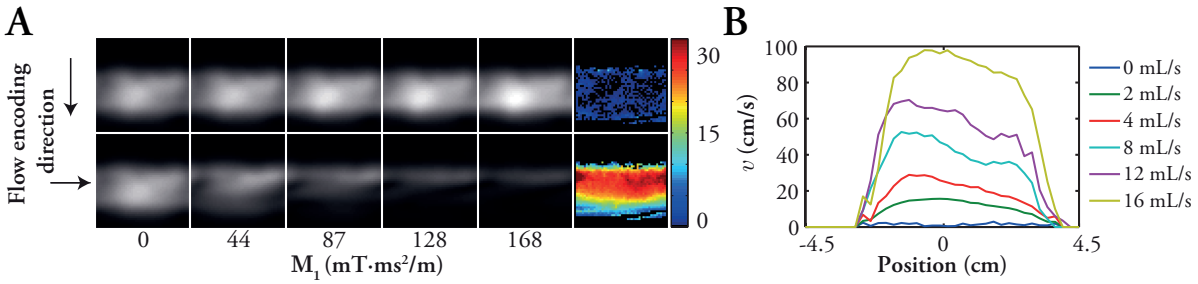

Figure 8.13: A demonstration of the effect of changing the direction of the flow-spoiling gradient. **A**: Here, the volumetric flow rate was fixed at 4 mL/s, and images with the direction of flow in-plane (i.e. sections through the pipe, with water flowing from left to right) were obtained with the suppression gradient at two different orientations. The dephasing effect occurs when the flow sensitising gradient is parallel to the direction of flow, as shown by the phase contrast velocity map on the right. This allows for (**B**) the reconstruction of the inter-voxel velocity distributions along the pipe at different flow rates.

high gradient amplitude. At an intermediate flow encoding gradient amplitude, the signal was retained at low flow rates; however, the signal was dephased at higher flow rates. This indicates that there exists an optimum amplitude for an expected flow regime to produce appropriate contrast. This allowed the reconstruction of phase contrast velocity map images, as shown in Fig. 8.12**B** at varying flow rates through the tube.

Subsequently, the flow rate was fixed at 4 mL/s, and the direction of the bipolar gradient was oriented either perpendicular or parallel to the direction of flow. As illustrated in Fig. 8.13**A**, the acquisition plane was orientated longitudinally through the pipe, and the received signal was only nulled when the gradient direction was approximately parallel with the flow velocity. This demonstrated that the efficiency of spoiling depended on the bipolar gradient direction, as expected. At a volumetric flow rate of 4 mL/s, the mean velocity within the tube is expected to be 28 cm/s. Spins moving at this velocity accrue an approximate phase of 2π at $M_1 = 87$ mT ms^2/m, which was consistent with the degree of signal attenuation present at this value, as shown in Fig. 8.13**B**.

8.3.2 In Vivo Study

Based on the preceding phantom study, the bipolar gradient was oriented along the long axis of the heart as this is the predominant direction of blood flow during the cardiac cycle. In order to determine the required bipolar gradient amplitude necessary to spoil the flowing ^{13}C signal from within the cardiac chambers, it was necessary to perform a proton reference scan on the rat heart. Figure 8.14 shows the results of this calibration scan using both Cartesian gradient echo and non-Cartesian spiral readouts, both with bipolar flow suppression gradients. The required gradient amplitude for typical flow rates within the rat heart was determined to be 40 mT/m ($M_1 = 78$ mT ms^2/m, VENC$= 15$ cm/s). At this gradient amplitude, the signal from within the lumen was approximately dephased, while the signal from stationary spins was preserved. The anatomical CINE reference demonstrates that the flow-suppressed signal comes predominantly from the myocardial tissue, and not bright blood flowing into the imaging plane.

The flow-sensitising bipolar gradient (amplitude 17 G/cm, corresponding to $M_1 = 332$ mT ms^2/m, VENC=14.1 cm/s) successfully dephased hyperpolarised ^{13}C-urea with varying velocities within a voxel, improving the visualisation of the slow-moving myocardial ^{13}C-urea. Representative images of the first passage of hyperpolarised ^{13}C-urea through the heart at rest are shown in Fig. 8.15. It was possible to resolve ^{13}C-urea in the entire heart, but without flow suppression, the urea within the myocardial wall could not be adequately resolved due to the insufficient in-plane resolution (Fig. 8.15**A**, top row). This lead to partial volume effects, resulting in 'contamination' of the blood pool within the 'myocardial' signal. With the flow suppression gradient turned on (Fig. 8.15**A**, bottom row), the time delay between the arrival of the contrast agent in the RV and myocardium was substantially

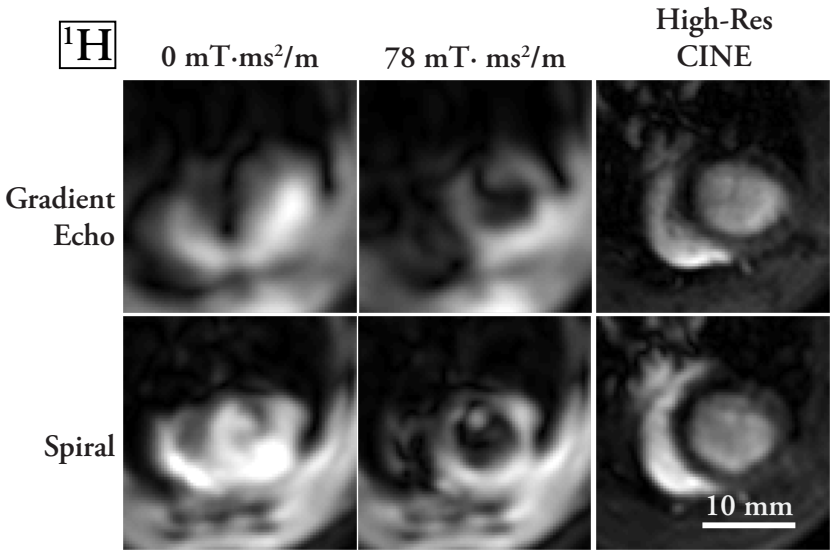

Figure 8.14: *In vivo* ^{1}H images obtained with a bipolar flow suppression gradient using either a Cartesian gradient echo readout (top row) or the golden angle spiral readout (bottom row). The flow suppression gradient ($M_1 = 78$ mT ms^2/m, VENC $= 15$ cm/s) successfully dephases signal from within the lumen, while stationary anatomy is preserved. A higher resolution CINE image is shown as an anatomical reference; the location of cardiac signal in the suppressed images co-localise with the myocardium. These images are cropped to a 25×25 mm^2 FOV.

increased, indicating that this artefact had been ameliorated (Fig. 8.15**B, C**). SNR loss is a concern with flow suppression techniques, as they fundamentally require removal of signal that would otherwise be acquired. Accordingly, the maximum image SNR was calculated, and that averaged in the LV was found to be 37 ± 12 across the six rats scanned.

As shown in Fig. 8.16, adenosine stress induced hyperæmia was detected by the flow sensitised imaging of hyperpolarised ^{13}C-urea in the rat heart. The mean heart rate did not change significantly during infusion of the drug (rest mean heart rate 382 ± 7 beats per minute, stress condition 392 ± 33 beats per minute; standard errors, $p = 0.32$ with Student's t-test). The main observed difference between rest and stress was an increase in the maximum myocardial ^{13}C-urea signal under the stress condition compared to rest, and an increase in the rate of arrival of the agent in the tissue. These effects were quantified by considering the total integral (estimated numerically using the trapezoidal rule) of the visible hyperpolarised ^{13}C-urea signal present in the myocardium under both rest and stress conditions, and with or without the flow suppression gradient. Each integral was normalised by dividing by the LV integral; i.e.

$$\text{Ratio of Integrals} = \frac{\iiint_{\text{Myocardium}} S(t,\, \mathbf{r})\ \mathrm{d}\mathbf{r}\,\mathrm{d}t}{\iiint_{\text{LV}} S(t,\, \mathbf{r})\ \mathrm{d}\mathbf{r}\,\mathrm{d}t}, \tag{8.5}$$

where $S(t,\, \mathbf{r})$ was the received signal at position $\mathbf{r}$ and time t. This process was performed to ameliorate the effect of polarisation variability. As illustrated in Fig. 8.17, it was found that the integral ratio calculated without the flow sensitising gradient was not significantly different between rest and

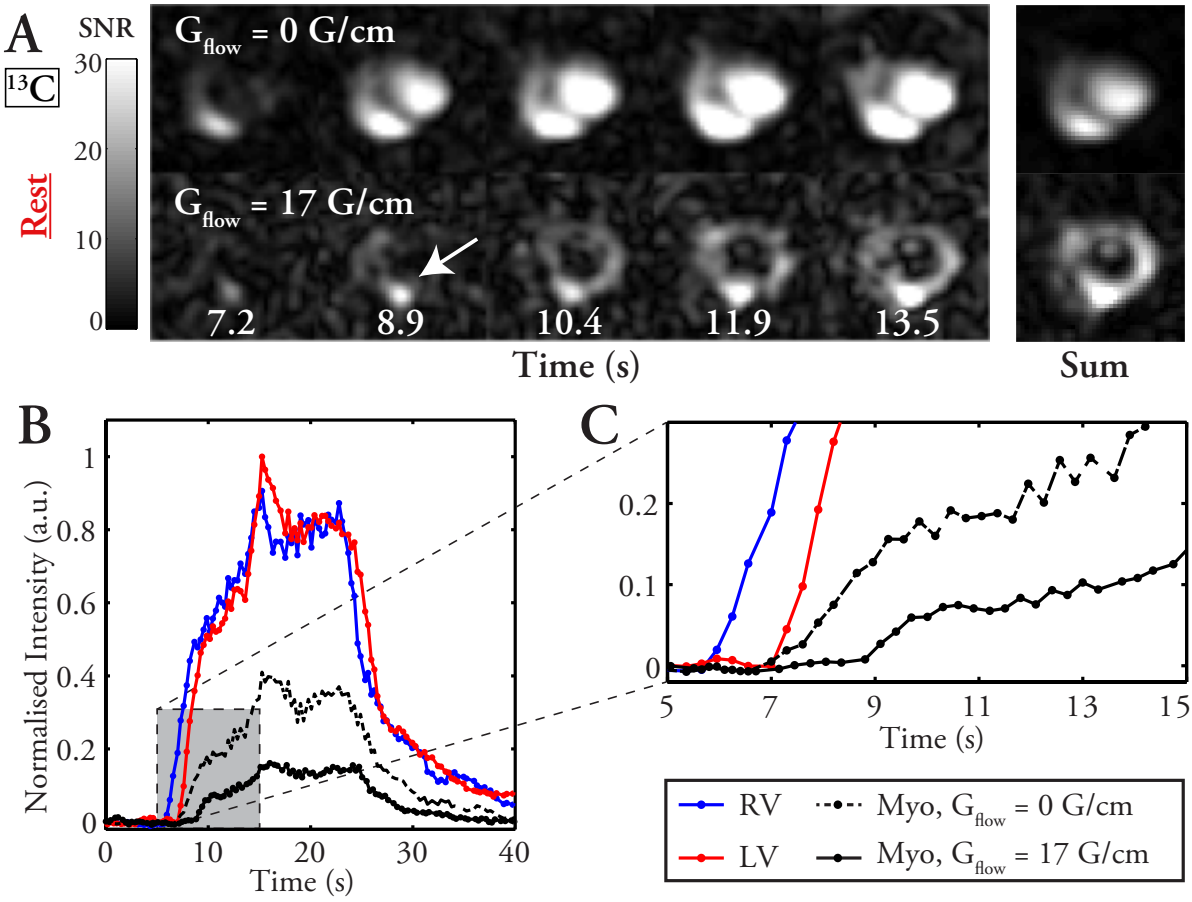

Figure 8.15: The first pass of hyperpolarised ^{13}C-urea through the rat heart at rest. **A**: Shown as a function of time are the reconstructed images from every 5th frame with ($M_1 = 332$ mT ms^2/m, VENC=14.1 cm/s) and without ($M_1 = 0$ mT ms^2/m) flow suppression. Images are cropped to a 25×25 mm^2 FOV. It is clear that flow suppression allows for the resolution of the contrast agent in the myocardium as opposed to the ventricles alone. A bright artefact is present, denoted by the arrow, which is believed to originate from ^{13}C-urea in the RV blood pool flowing in a direction not perpendicular to the plane, and therefore not suppressed by the flow sensitising gradient. **B, C**: Timecourses from the segmented heart show the entry of urea in first the RV (blue), then LV (red) and finally myocardium (black). The observed uptake in myocardial signal in the absence of flow suppression (black, dashed) shows the partial volume effect of the nearby bright blood pool; with flow suppression (black, solid) the kinetic behaviour is substantially different, with a ~ 2 s longer time of arrival.

stress. However, with the flow suppression gradient, a two-fold increase in the integral ratio was observed between rest (0.1 ± 0.03) and stress (0.21 ± 0.03; standard errors; $p < 0.05$ with Student's t-test.)

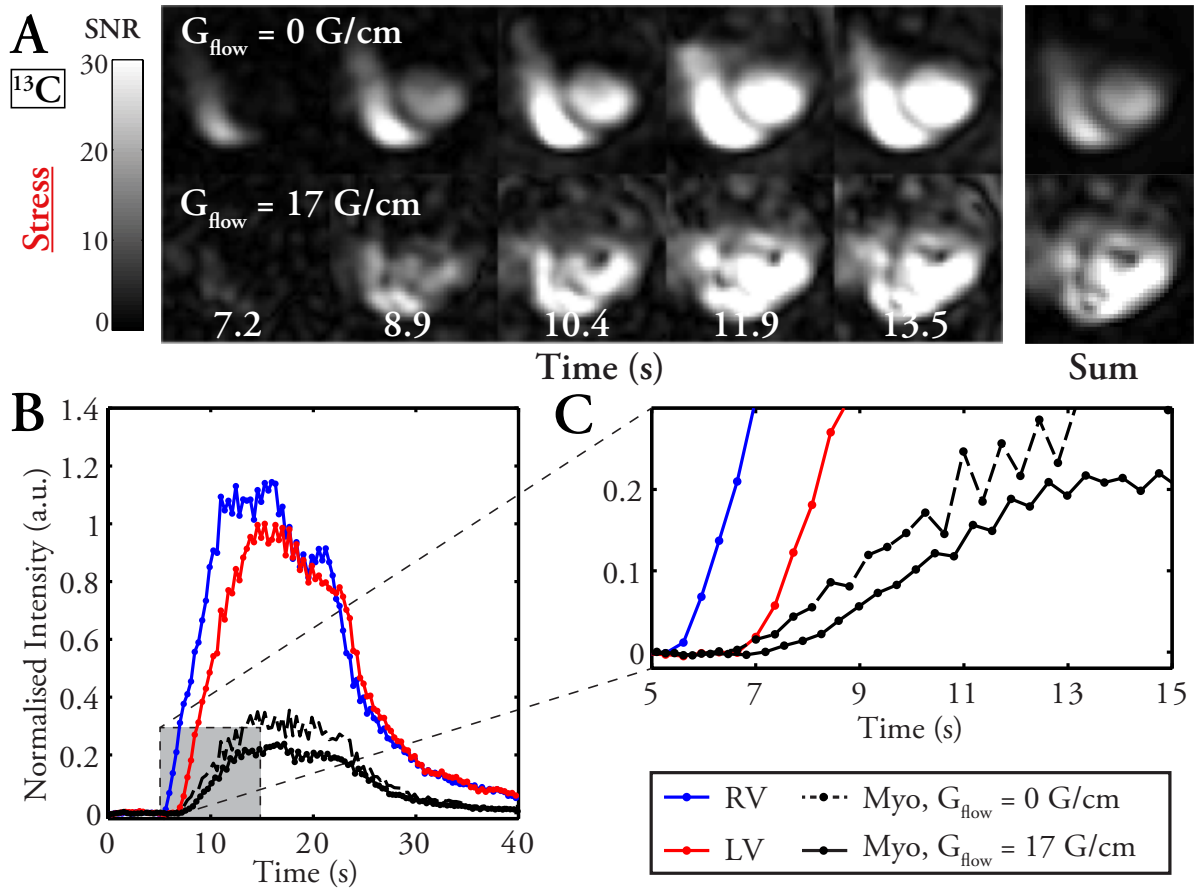

Figure 8.16: Hyperpolarised ^{13}C-urea perfusion imaging of the rat heart under adenosine stress. **A**: Shown as a function of time are the reconstructed images from every 5th frame with (M_1 = 332 mT ms^2/m, VENC=14.1 cm/s) and without (M_1 = 0 mT ms^2/m) flow suppression. Images are cropped to a 25 × 25 mm^2 FOV. Under the stress condition, much more signal is visible in the flow-suppressed myocardium compared to that at rest. Additionally, the delineation of the myocardium is more difficult, potentially because of changes in flow dynamics in stress compared to rest. **B, C**: Timecourses from the segmented heart show the entry of urea in first the RV (blue), then LV (red) and finally myocardium (black). As before, the observed uptake in myocardial signal in the absence of flow suppression (black, dashed) shows the partial volume effect of the nearby bright blood pool; with flow suppression (black, solid) the kinetic behaviour is substantially different, with a ∼ 2 s longer time of arrival.

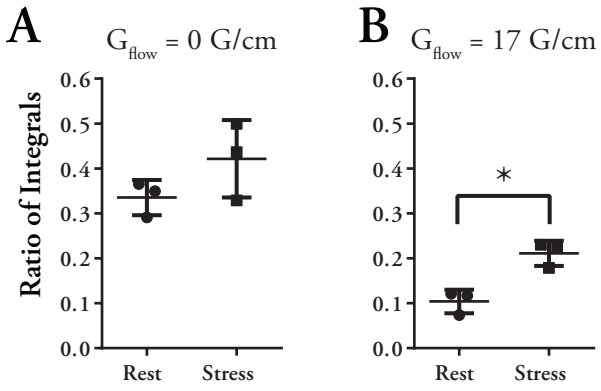

Figure 8.17: The ratio of integrals of the kinetic curve of myocardial ^{13}C-urea signal under stress and rest conditions. **A**: Without flow sensitisation gradients, there is no significant difference in apparent myocardial urea perfusion between the rest and stress experiments, whereas with flow sensitising gradients this difference is resolved (**B**; ∗ denotes Student's $P < 0.05$).

8.4 | Discussion

8.4.1 Importance

It is challenging to measure myocardial perfusion in rodents due to the high spatial and temporal resolution required in order to adequately resolve the heart. The combination of the small heart size of the rat, its rapid heart rate, and the therefore limited quiescent periods during the cardiac cycle for imaging have thus far restricted the use of hyperpolarised agents to perfusion measurements in larger, relatively stationary structures such as the kidneys or cancers in the rat, and the hearts of larger animals. This study has demonstrated the feasibility of using flow-sensitising bipolar gradients to improve visualisation of hyperpolarised ^{13}C-urea within the rodent myocardium. Flow suppressed hyperpolarised ^{13}C-urea images were shown to detect an increase in perfusion during low-flow infusion of adenosine, a coronary vasodilator, which is known *a priori* to increase myocardial perfusion. In contrast, images acquired without flow suppression were not sensitive to this change. It was not feasible to simply resolve the myocardium by increasing the resolution of the scan without severely compromising SNR.

It is hoped that this development will enable novel strategies for hyperpolarised ^{13}C imaging studies in the heart. In particular, the potential to assess both metabolism and perfusion by combining information from infusions of hyperpolarised [1-^{13}C]pyruvate and ^{13}C-urea would allow a direct visualisation of perfusion-metabolism mismatches in the heart within a single imaging session. This would provide a more complete understanding of ischæmia-reperfusion injury, and allow for the quantification of tissue viability following myocardial infarction. Furthermore, hyperpolarised exams are intrinsically rapid by nature of the short T_1 of the hyperpolarised ^{13}C contrast agent, and, assuming that the patient is already indicated for an MR examination and intravenous access had been obtained, could therefore be readily incorporated into a wider MR exam of cardiac function with only a marginal increase in study time.

8.4.2 Physiological Interpretation of Results

It is well known that adenosine induces coronary vasodilation without a substantial concomitant increase in heart rate.[229] Mechanistically, the drug acts by binding to the adenosine Type 2A (A2A) receptors on the coronary smooth muscle, which are coupled via a G-protein signalling cascade to the protein adenylyl cyclase. The action of adenylyl cyclase is to increase the local concentration of another biologically relevant signalling small molecule, cAMP. An intracellular increase in cAMP serves to increase the rate of phosphorylation of K_{ATP} channels (amongst other effects), which in turn activates the channel, producing an efflux of K^+. This efflux reduces the cell membrane potential, decreasing the sensitivity of the cell to incident action potentials and via a downstream cascade of calcium signalling

ultimately inhibits the contraction of the smooth muscle. This in turn produces coronary vasodilation, leading to an increase in the perfusion of the myocardium.

The two-fold increase in hyperpolarised ^{13}C-urea detected in the myocardium following adenosine infusion relative to the rest condition is consistent with previous measurements of myocardial perfusion reserve in rats, using first-pass Gd-DTPA imaging,[201] fluorescent microspheres[230] and microbubble contrast-enhanced ultrasound.[231] The myocardial ring is well resolved with the flow-sensitised contrast, suggesting that regional perfusion deficits following myocardial infarction can be detected using this technique. Future work will aim to image regional perfusion deficits following MI with ^{13}C-urea.

8.4.3 Comparison to Contrast-Enhanced Perfusion Measurements

Conventional contrast-enhanced perfusion measurements using Gd-based contrast agents are mature and simple to implement. Gadolinium perfusion measurements are routinely used clinically as a robust measure of tissue perfusion, often in preference to CT angiography techniques, and in situations where interventional catheterisation is not indicated. This is primarily because (a) it is easier to spot a large region of enhancement in the myocardial ring following infarction than to resolve a small deficit in the coronary arterial diameter (as would be shown by CTA, MR flow imaging, or direct angiographic measurements); (b) gadolinium perfusion is relatively safe, and certainly 'safe enough' for routine use; and (c) the 'leakage' and accumulation of gadolinium to the extravascular spaces allows for the macroscopic reporting of the tissue microenvironment, meaning that gadolinium is of value in non-ischæmic pathology, such as cardiac amyloidosis, myocarditis, and other causes of fibrosis. This is distinct from the early myocardial enhancement arising from normal perfusion of the vessels, but might provide a potential confounding factor in the interpretation of gadolinium perfusion images.

It is therefore of value to consider the advantages of using a hyperpolarised agent to measure perfusion. The main benefits of using hyperpolarised ^{13}C-urea are (a) its linear signal-concentration proportionality, in contrast to gadolinium which is exponentially related to the observed data, (b) the excellent intrinsic *in vivo* safety profile of urea, and (c) the exceptionally (and necessarily) short scan times which can be beneficial in the acute setting. Furthermore, whilst the above work has served to demonstrate the proof-of-principle use of ^{13}C-urea to monitor perfusion following a single infusion, it is possible to co-polarise ^{13}C-urea and [1-^{13}C]pyruvate simultaneously. This will lead to studies that aim to investigate the feasibility of using a co-polarisation of metabolically inert ^{13}C-urea and metabolically active [1-^{13}C]pyruvate to provide cardiac perfusion-metabolism information in a single scan. This would be important in patients with cardiovascular disease and impaired renal function who cannot tolerate a gadolinium chelate infusion.

Finally, ^{13}C-urea and [1-^{13}C]pyruvate co-polarisation and co-infusion with spectral-spatial imaging would allow for the polarisation 'reservoir' present in [1-^{13}C]pyruvate itself to remain unsampled

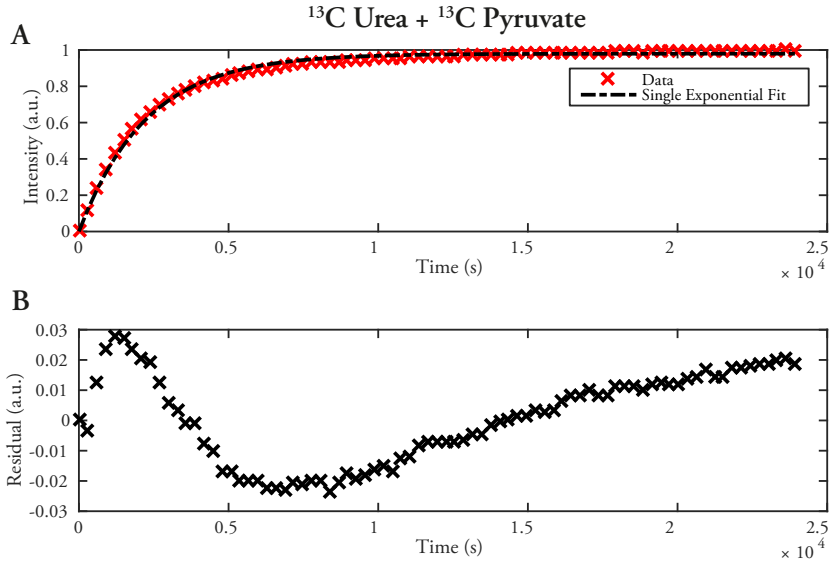

Figure 8.18: **A**: Co-polarisation build-up curve of ^{13}C-urea and [1-^{13}C]pyruvate, together with a single exponential fit. The time constant here is approximately 40 minutes, greater than that of [1-^{13}C]pyruvate (30 minutes) and less than ^{13}C-urea alone (80 minutes). The deviation of the data from single-exponential behaviour is apparent in the residuals of the fit, **B**.

without sacrificing information on metabolism. This would prolong the acquisition time (and available SNR) for the imaging of metabolites other than [1-^{13}C]pyruvate. Preliminary work by the author shows that the co-polarisation of [1-^{13}C]pyruvate and ^{13}C-urea can be performed (c.f. Fig. 8.18), and that the degree of polarisation reached by both are not independent; i.e. knowledge of the degree of polarisation of ^{13}C-urea is an adequate proxy for that of [1-^{13}C]pyruvate. By using the 3D spectral-spatial sequence developed in Chapter 6 with optional flow suppression gradients, and exciting hyperpolarised ^{13}C-urea instead of [1-^{13}C]pyruvate it may be possible to realise this technique with high SNR efficiency.

8.4.4 Other Hyperpolarised Perfusion Strategies

This work has used flow-suppression gradients to visualise hyperpolarised perfusion imaging within the rodent heart. It might be asked what other methods could be used to visualise hyperpolarised perfusion agents in the heart, especially as the use of flow suppression gradients may be thought of as deliberately spoiling information that is difficult to produce. Arguably, the simplest way to visualise the myocardium would be the higher resolution imaging of a hyperpolarised contrast agent bolus by simply traversing further out in k-space, thereby acquiring a higher resolution image. This approach results in a necessarily large SNR penalty, and, with rapid imaging readouts fundamentally limited by the comparatively short T_2^* at 7 T, achieving high spatial resolution typically requires increasing the number of interleaves and degrading the temporal resolution in the final images.

It may be possible to use a slice-selective steady-state free procession (SSFP) sequence in combination with dual-labelled $\left[^{13}\mathrm{C},\ ^{15}\mathrm{N}\right]$ urea which possesses a very long T_2,[232] but the fidelity of such pulses in the heart and the short TRs required would need to be investigated. In particular, it is not clear that

the long slice-selective echo train required for SSFP[233] with a duration on the order of T_2 would be compatible with the imaging of the rapidly moving hyperpolarised magnetisation, as the RR interval (~ 150 ms) is shorter than the reported *in vivo* T_2 of 200 ms for ^{13}C-urea and 5 s for $\left[^{13}\mathrm{C},\ ^{15}\mathrm{N}\right]$-urea. Such a long pulse sequence would therefore be likely to introduce significant blurring over the cardiac cycle, and would also result in a bright blood pool and excessive loss of magnetisation as fresh blood enters the slice. The approach used here is simpler, cheaper, and less prone to errors in the coil flip angle calibration.

The use of an adiabatic spin-echo pair following a small excitation (known as the semi-LASER sequence) is an attractive alternative imaging strategy for hyperpolarised experiments that promises enhanced insensitivity to the transmit flip angle calibration, and provides in-phase spectra following slice-selective excitation, independent of the readout acquisition strategy.[234,235] Furthermore, flow-sensitising gradients placed either side of the two 180° pulses would serve to suppress hyperpolarised signals from flowing spins, and, as previously established, therefore enable the discrimination between intracellular and extracellular metabolites.[224,236–238] Typically, the spin-echo experiment is started after the injection of the ^{13}C-labelled compound is complete, as adiabatic pulses generally cease being adiabatic below a certain threshold power level that typically occurs in the fringe field of the transmit RF coil in question. The delayed start commonly used ensures that the hyperpolarised signal will not be saturated during the injection. In the context of first-pass cardiac imaging, this requirement is obviously not desirable, and therefore the use of spin-echo pairs for hyperpolarised cardiac imaging is not straightforward.

8.4.5 Artefacts and Limitations

As described in Section 8.2.2, and demonstrated in Section 8.3.1, the spoiling efficiency of the flow-suppression gradient used is related to the distribution of velocities along the direction of the uni-axial bipolar gradient within each voxel imaged. In this study, the single gradient pulse was oriented parallel to the cardiac long axis, which was chosen as the long axis of the left ventricle represents the predominant direction of blood flow within the heart. However, this scheme does not suppress flow in either lateral direction, which may be important near the apex of the heart and in the atria. Nevertheless, the fact that a two-fold increase in the received hyperpolarised ^{13}C-urea signal was observed following adenosine stress-induced hyperæmia strongly indicates that the degree of spoiling achieved is sufficient to enable the detection of large-scale perfusion changes in the rodent heart.

Additionally, it may be the case that the signal attenuation results from the imaging gradients themselves. The first gradient moment of the spiral trajectory oscillates, and reaches a maximum moment of 316 mTms^2/m at the end of the 7.24 ms readout, similar in magnitude to the bipolar gradient used in this study. It is possible to define a velocity k-space formalism, in which the first moment is

integrated as a function of time, which allows for the quantitative understanding of the effects of flow. Similar to their circular trajectory in k-space, spiral gradients traverse a circular trajectory in velocity k-space, resulting in phase accrual and blurring for in-plane flow over the duration of the readout.[239] It is difficult to assess whether this effect generates signal attenuation in the hyperpolarised ^{13}C-urea images, as the spiral readout is kept constant for each time frame and not varied. However, the fact that there is bright signal within the lumen in the spiral not-flow-encoded images is consistent with no flow-related phase accrual at the centre of (image) k-space. Accordingly, any low spatial resolution structures (i.e. the left ventricle) will be preserved in the resulting images. Additionally, the fact that the image quality and SNR is similar between the flow-encoded Cartesian and spiral ^{1}H cardiac images (c.f. Fig. 8.14) suggests that any flow attenuation is minimal.

In the flow-encoded images, bright signal was observed in the right ventricle as well as within the lumen of the left ventricle. This is outside the expected primary localisation of ^{13}C-urea in the myocardial wall alone. The appearance of signals within the cardiac chambers is believed to be due to the insufficient dephasing of slowly moving blood within the chambers, or representative of blood flowing with a significant velocity component not parallel to that of the encoding direction. It is important to note that there is a spatial resolution dependence on the flow-sensitising gradient strength,[240] and therefore different applications may require a different set of parameters to properly null the signal within the contaminating blood pool.

Additionally, it may be possible to improve the flow suppression effect by using additional bipolar gradient directions to successfully suppress a generalised intravoxel velocity distribution. This would come at the expense of temporal resolution. Alternatively, the amplitude of the bipolar gradients could be reduced to introduce a velocity-encoded phase change between ^{13}C images. The images could then be segmented by velocity, ideally leading to the separation of flowing (lumen) and stationary (myocardial) spins.

The hyperpolarised ^{13}C-urea solution was infused over a period of 20 s, which is longer than the circulation time of blood (< 10 s) in a rat. Thus, there is likely recirculation of the signal even before the end of the infusion, and certainly before the end of the experiment. This is a limitation of the experimental protocol, and the long infusion was chosen to maximise the number of frames during which ^{13}C-urea was present in the heart. Additionally, it is not feasible to manually inject ^{13}C-urea into a rat tail vein at a substantially faster rate than this. It is of note that the initial upward slope of the hyperpolarised signal into the heart is different between rest and stress conditions, as shown by comparing Fig. 8.15 and Fig. 8.16. This positive gradient represents the first passage of the contrast agent in the myocardium. Moreover, the observation of this difference between rest and stress conditions is only resolved whilst the the bipolar flow suppression gradient is played, indicating that aspects of the first pass of urea into the myocardium are indeed being measured.

8.4.6 Alternative Methods for Reducing the Blood Pool Signal

At least two additional methods have been proposed to reduce the contribution of bright hyperpolarised compounds from nearby stationary targets of interest: (a) the spatially selective saturation of inflowing spins;[241] and (b) the separation of signal within the blood pool by comparison to a known blood input function.[242]

In oncology, the spatially-selective saturation of hyperpolarised signal with spectral-spatial RF pulses has been used to suppress bright extracellular flowing lactate from the endogenous metabolite target region of interest. This saturation approach within the imaged slice of interest and before the subsequent imaging of a bolus of inflowing magnetisation has additionally been previously described for renal perfusion imaging.[243] However, these approaches requires careful placement of the imaging slice to avoid saturation of the upstream arterial signal, effectively ending the experiment. It would therefore be challenging to use this technique when the arterial or luminal signal that would otherwise be suppressed is also within the slice of interest, as is the case with the heart.

The mathematical modelling of the visible blood pool signal by a single tissue box model has also been used to estimate the fractional blood volume in each voxel, by the comparison to an input function. A model of this type has previously been used to describe the signal behaviour of hyperpolarised agents within the brain, liver, and inside tumours.[244] As an alternative to flow suppression of the hyperpolarised signal, it may be possible to fit the hyperpolarised signal within the myocardium to a single box model, but it is important to note that each voxel in the heart would contain significant partial volumes of myocardium, and slowly flowing myocardial blood within the capillaries, in addition to the rapidly moving blood within the lumen. To aid the fitting, an arterial input function may be extracted from the signal within the cardiac lumen, and the fractional partial volume between tissue and cardiac lumen computed from the down-sampling of anatomical CINE images. Nevertheless, further investigation is required to determine the feasibility of this approach, and its computational complexity is non-trivial.

8.4.7 Sequence Improvements

The short TR per spiral interleave (~ 11 ms) and typical heart rates in rats (150 ms per cardiac cycle) suggests that the acquisition of multiple slices per cardiac cycles is feasible, which would provide the capability for the full volumetric assessment of perfusion in the rat heart. As the single 5 mm slice used in this study covered approximately half of the ventricular myocardium, the use of multiple thinner slices would permit evaluation of localised disease in basal, mid-ventricular and apical locations. This approach was not performed due to duty cycle limitations, and concerns about SNR.

8.4.8 Quantitative Measurements of Absolute Perfusion

In this study, myocardial perfusion was quantified by taking the ratio between the integral with respect to time of the acquired signal in the myocardial tissue and in the left ventricle. The use of the left ventricular signal integral as a normalising factor was intended to remove differences in $^{13}\mathrm{C}$ polarisation, $^{13}\mathrm{C}$-urea concentration, and small differences in transmit/receive sensitivities between experiments. A comparison between this ratio at rest and during adenosine stress-induced hyperæmia was then used to obtain an estimate of myocardial perfusion reserve.

It would be desirable to use the dataset acquired to permit an estimate of absolute perfusion, in units of mL of blood per g of tissue per minute. Absolute perfusion values would therefore enable a quantitative comparison of myocardial perfusion between individuals and over time. The ability to distinguish between normal perfusion and such a balanced ischæmic state would improve the clinical utility of the technique in disease. In principle, calculating a blood flow rate (in units of $\mathrm{mL/g/min}$) should be possible using a hyperpolarised agent, because there is no signal background and there is a linear relationship between signal intensity and concentration. The main challenge in obtaining a quantitative estimate of perfusion is correctly accounting for the polarisation decay over time, similar to the short half-life (of approximately 2 minutes) in $\mathrm{H_2}^{15}\mathrm{O}$-PET perfusion measurements.[245]

Quantitative models typically rely on the ideally *a priori* characterisation of the relationship between the arterial input function and the tissue signal. This tissue response function can be estimated from a given data set by using using model-based (Fermi) constrained deconvolution or by model-free singular value decomposition, similar to previous quantitative perfusion measurements using $^{13}\mathrm{C}$ agents in kidneys, cancer, and brain.[246,247]

8.5 | Conclusions

This chapter has described the development of a new flow-sensitised method for imaging the first passage of a bolus of hyperpolarised $^{13}\mathrm{C}$-urea through the rat heart, and the perfusion of the myocardium was quantitatively evaluated using the technique. A flow-sensitising bipolar gradient was optimised to maximally dephase rapidly flowing spins within the heart, and was incorporated into a spiral imaging pulse sequence. A two-fold increase in myocardial perfusion relative to rest was detected during adenosine stress-induced hyperæmia, consistent with the known myocardial perfusion reserve in the rat. This probe of perfusion is anticipated to enable new applications in hyperpolarised $^{13}\mathrm{C}$ MRI, particularly in combination with the hyperpolarised [1-$^{13}\mathrm{C}$]pyruvate imaging sequences described previously. It is hoped that the co-polarisation of $^{13}\mathrm{C}$-urea and [1-$^{13}\mathrm{C}$]pyruvate will allow for the simultaneous quantification of perfusion and metabolism within a single imaging session, and with the administration of a single dose of endogenous, stable contrast. Such a scan would be of utility in a clinical setting, able to combine several complementary techniques at once in a safe and efficient manner.

Chapter 9

Dealing with DNP Data: Uncertainty, Metabolic Modelling, and Spectral Analysis

Regretfully, any physical experiment will unavoidably be subject to noise and variation. Owing to polarisation-to-polarisation variability, the exponential T_1 decay undergone by hyperpolarised compounds after dissolution, and variability in the time required to move a hyperpolarised compound from the hyperpolariser to the subject in the imaging magnet, the initial signal in DNP experiments is highly variable. As a consequence, the signal-to-noise ratio in hyperpolarised experiments also varies substantially, sometimes by several orders of magnitude. As a result, a simple method to incorporate data quantified from spectra of varying quality is presented, and is that which was used for the preliminary study described in Chapter 4.

As an alternative, the novel use of another peak quantification method – the Fast Padé Transform – is presented for accelerating the acquisition of hyperpolarised spectroscopic data. The Fast Padé Transform (FPT) forms an alternative to the Fast Fourier Transform, and is instead based upon finding a rational approximant to describe the acquired time domain data, before extracting spectral information based on analytic properties of the function created. It is shown to be far more robust to truncation of the time domain data than the conventional Fourier method of zero-filling, which effectively evolves a sinc shaped interpolation in the spectral domain. This transform may further improve the reconstruction of hyperpolarised data, or alternatively allow for acceleration in its acquisition.

Finally, the problem of analysing hyperpolarised data is additionally discussed, and a Bayesian approach is taken to fitting a spatially invariant metabolic model by the construction of a Monte Carlo Markov Chain whose stationary distribution is that of the posterior. The Metropolis-Hastings algorithm is used with Gibbs sampling, and the result is shown to produce bounded estimates of rate constants of interest. It is noted that, in general, the finite and non-renewable magnetisation available to hyperpolarised experiments will place an upper bound on their achievable resolution, but that the acquisition of complementary $^1\mathrm{H}$ data – often as an anatomical reference – is often both easy and routinely performed for other reasons. Therefore, the expansion of the fitted model to include spatially varying terms, constrained upon the known anatomy, is proposed as an extension to the work already

performed. As sets of partial differential equations can be solved to almost arbitrary resolution, it is hoped that this method could allow the fundamental limitations of the technique to be transcended.

9.1 | *Spectral Peak Quantification*

NMR is not an inherently quantitative technique. Experimental factors such as the available dynamic range of ADCs, coil transmit/receive sensitivity, differences in construction between nominally identical spectrometers, and the complex interactions that make B_0 inhomogeneous between samples all conspire to make the direct comparison of signal amplitudes (in arbitrary units) exceptionally difficult between different experiments. For this reason, several methods of quantifying NMR spectra have been proposed. The simplest is integration: as peak areas are directly proportional to the visible concentration of the compound present, by integrating $\mathrm{d}\delta$ over a finite range of chemical shift encapsulating the ideal Lorentzian line-shape of a resonance, an estimate of the relative concentration of a given peak can be obtained. The use of an appropriate reference standard at known concentration and chemical shift allows for the independent determination of frequency and concentration. However, simple integration does not perform well if the baseline of the spectra is not constant (perhaps possessing first-order phase, or a constant slant), and exceptionally poorly if peaks overlap, as the (ideally Lorentzian) 'wings' of two nearby peaks contribute to the integral value of their neighbours. It has been shown that defining integration regions and neglecting the 'wings' of a given peak can affect accurate quantification of the species in question by as much as 40%,[248] and, in general, the phasing of complex spectra is a non-trivial problem.

For this reason, several distinct time- and frequency-domain methods have been proposed to identify peaks in NMR spectra. Whilst a comprehensive review of these methods is beyond the scope of this work, two examples are briefly discussed: the AMARES time-domain algorithm, which was used for quantifying spectra in Chapter 4, and (later) the Fast Padé Transformation method, which is used to quantify truncated data.

9.1.0.1 AMARES

AMARES, "an Advanced Method for Accurate, Robust, and Efficient Spectral fitting",[249] is a time-domain spectral quantification algorithm. Like the majority of time-domain methods, it directly fits a model function, here a sum of exponentially damped sinusoids, to the acquired FID. Specifically, a model function describing K resonances known *a priori* of the form

$$\hat{y}_n = \epsilon + \sum_{k=1}^{K} a_k \exp\left(i\phi_k + \left(-d_k + 2\pi f_k\right) t_n\right) \tag{9.1}$$

is fitted to the N data points acquired, where ϵ represents (complex) white noise (with uncorrelated real and imaginary parts); $t_n = n\Delta t + t_0;\ n = 0,\ 1 \ldots N-1$ is the time for each acquired point n; t_0

is the 'begin time' of the acquisition; and where ϕ_k, a_k, f_k and d_k are the initial phase, amplitude, frequency and damping (i.e. linewidth) of the kth resonance. Naïvely, one would expect to solve the 'full' optimisation problem by numerically minimising the difference between the model and the data, i.e. minimising $\|\mathbf{y} - \hat{y}\|^2 \equiv \|\mathbf{y} - \Psi\mathbf{A}\|^2$, where $\mathbf{y}$ is the acquired data; $\Psi\mathbf{A}$ represents the model function expressed as a product of a vector of amplitudes and phases ($\mathbf{A}_k = a_k e^{i\phi_k}$) and a matrix $\Psi_{jk} = \exp\big((-d_j + 2\pi i f_j)t_k\big)$; and where $\|x\|$ denotes the Euclidean 2-norm of x. Typically, solving this full problem is difficult owing to the large number of parameters that need to be found. For this reason, a related 'variable projection' problem is estimated, in which $\|\mathbf{y} - \Psi\Psi^{\dagger}\mathbf{y}\|^2$ is minimised, and $\mathbf{A}$ estimated from other means, or provided *a priori*. As a result, a linear optimisation problem in fewer parameters (but with a more complex objective function) is obtained. Both of these objective functions are relatively easy to minimise by popular numerical regimes, some of which can additionally handle specified constraints (e.g., a specific frequency lies within a specified range, or even that functions of the parameters obey some relationship or constraint). AMARES is a method that performs this fit, with specified constraints, via a non-linear least squares algorithm.

It is worth noting that this prior information can have a significant impact on the fit that AMARES finds; the non-linear least squares algorithm used is *not*, in general, guaranteed to converge upon the global minimum in the optimisation problem that it solves. Likewise, it is difficult to quantify the effect of the prior information supplied. This has driven the use of Bayesian methods, such as Markov Chain Monte Carlo (MCMC) methods with Gibbs sampling, to quantify spectra. Such methods are, in general, guaranteed to ultimately converge upon the true posterior probability[1] that corresponds to the global minimum of the optimisation problem. Additionally, prior information is easy to robustly input in a Bayesian framework, quite literally as a specified prior probability distribution function for a given parameter. The effect of this prior can then be estimated by changing it, for example, to a 'noninformative' distribution, such as a uniform one. However, these approaches are exceptionally computationally intensive, and require many orders of magnitude more time to quantify a given spectrum than their Frequentist analogues. For this reason, the use of algorithms such as BATMAN,[250] or BQuant[251] are comparatively restricted to more specialised, difficult areas, such as the quantification of highly dense metabolomic data.

The main difference between AMARES and its predecessor algorithm VARPRO[252] is treatment of multiplet peaks and constraints: AMARES utilises an all-single-peak approach, which may simplify the numerical aspects of the problem. Accordingly, VARPRO uses the (computationally and conceptually straightforward) Levenberg–Marquardt, whereas AMARES (as implemented with jMRUI[127]) uses non-linear least squares. Both are iterative numerical algorithms ultimately based on the Gauss-Newton (gradient) method, and require accurate 'initial guesses' of parameters to be estimated, and both

[1] Gibbs sampling is usually preferred as it ensures that the Markov chain will converge at an exponential rate to the posterior, although that rate may be arbitrarily slow.

can support constraints. *Estimates* for the uncertainty in each parameter can also be obtained, either from the fitting process (by estimating confidence intervals based on likelihood ratios or assuming an asymptotically normal distribution for each parameter's estimate), or alternatively from information theory in the form of the Cramér-Rao lower bound (CRLB). The latter is a fundamental variance limitation that depends on properties of the underlying quality of the data, such as the number of acquired data points, the sampling method used, and the type and quantity of noise present; and will apply to time-domain and frequency-domain methods (though, of course, as a lower bound, any given method may be worse than it). For any general NMR 2D spectroscopy experiment, it can be shown that the CRLB for uncertainty in a parameter of expected value $\hat{\theta}$ takes the form

$$\mathrm{var}(\hat{\theta}_i) \geq \left[F^{-1}(\Theta)\right]_{ii}, \tag{9.2}$$

$$\text{where } F(\Theta)_{ij} = -E\left[\frac{\partial^2 \ln p(x,\Theta)}{\partial\Theta_i \partial\Theta_j}\right] \tag{9.3}$$

is the Fisher Information Matrix, defined as the expectation of the Hessian of the log-likelihood function $\ln p(x,\Theta)$, and Θ is a parameter vector for the fit.[253] In NMR, a form for p is often known *a priori*: for the 1D spectroscopy experiment as considered in Eq. (9.1), p takes the multivariate normal form

$$\begin{aligned} p(x,\Theta) = \prod_{n=1}^{N} &\frac{1}{\sqrt{2\pi\sigma_m^2}} \exp\left(-\frac{1}{2\sigma_m^2}\left(\mathrm{Re}(\mathbf{y}) - \mathrm{Re}(\hat{y}-\epsilon)\right)^2\right) \\ &\times \frac{1}{\sqrt{2\pi\sigma_m^2}} \exp\left(-\frac{1}{2\sigma_m^2}\left(\mathrm{Im}(\mathbf{y}) - \mathrm{Im}(\hat{y}-\epsilon)\right)^2\right), \end{aligned} \tag{9.4}$$

and for 2D NMR experiments this form generalises. Accordingly, by substituting Eq. (9.4) into Eq. (9.2), it is possible to directly find expressions for the CRLB for NMR data assumed to have a Lorentzian line shape. It transpires that, for NMR data, the Fisher Information Matrix is block-diagonal, which ensures that the CRLB exists. Interestingly, if the quantification of a spectrum has proceeded successfully, the CRLB can be analytically returned from the parameters given, providing a minimum estimate of uncertainty on a given fit.

9.1.0.2 jMRUI

The utility of AMARES as a peak quantification algorithm comes from its speed and flexibility. The authors of the original AMARES paper also provided a user-friendly GUI, distributed free to academics in the EU. This GUI, jMRUI, allows several common tasks – such as FID truncation and DC offset correction, zero-filling, spectral phasing, and complex conjugation – to be performed in addition to providing AMARES as a quantification algorithm (along with several others, such as the Hankel Singular Value Decomposition[254,255]). Both jMRUI and the AMARES algorithm were used to quantify

the spectra obtained in Chapter 4, by specifying the approximate frequency location of each resonance in the hyperpolarised [1-^{13}C]pyruvate spectrum, their approximate linewidths, and the (first and zeroth order) phase the spectrum as prior information. The returned peak amplitudes, and the square-root of the CRLB of that amplitude were saved for further analysis.[2]

However, jMRUI is written primarily in Java, with several proprietary compiled C and FORTRAN libraries. The source code is not freely available, and, at the time of writing binaries were only provided for 32-bit Linux and Windows. However, the application was developed with public funds, and is intended to be available for research groups in the EU. Since there are several popular operating systems and architectures other than 32-bit Linux and Windows, it is desirable to make the application as portable as its frameworks would allow. The Java decompiler `Jad` was used to effectively reverse-engineer the application, but an attempt at writing a binary compatibility layer for the compiled modules failed. Therefore, the `wine` compatibility layer was instead used to provide OS X and Linux binaries (including a `jre` appropriate for jMRUI) to allow the use of its features on different platforms. This version was used for the analysis of spectral data in this work, and is available upon request.

9.1.0.3 Hankel Singular Value Decomposition

For any $m \times n$ matrix A, it is possible to find a decomposition of A into the form $A = U\Sigma V^{\dagger}$ where U is an $m \times m$ unitary matrix, Σ an $m \times n$ rectangular diagonal matrix, and V is is an $n \times n$ unitary matrix. The diagonal values of Σ are known as the singular values of A, and can be found efficiently through various numerical Singular Value Decomposition (SVD) algorithms. By taking a signal model of the form Eq. (9.1) into a rectangular matrix of the Hankel form, the required estimates for signal components such as the location of resonances and their phase can be found by computing the SVD of this matrix. Owing to the presence of noise, it is often required to ignore all but a specified number of results. The technique, known as Hankel Singular Value Decomposition (HSVD), essentially acts a change of basis, and can be considered a 'black box' method for spectral analysis, analogous to the principle component analysis method commonly used elsewhere.[254,256]

9.2 | Propagating Uncertainty

A common design of *in vivo* spectroscopy studies is to look for an underlying difference between two (or more) groups of subjects. Spectra are acquired, and the quantity of interest is then usually a function of one or more peaks in the resulting spectra – such as the estimated concentration of a compound of interest compared to a reference, the ratio of phosphocreatine to γATP in ^{31}P spectroscopy, or the ratio of visible [1-^{13}C]lactate to [1-^{13}C]pyruvate following the injection of hyperpolarised [1-^{13}C]pyruvate.

[2] It is of note that versions of jMRUI released prior to June 2013 suffer from a bug in which the returned CRLB suffers from inadvertent double-division by the residual noise level (personal communication from Christopher Rodgers).

As discussed above, many spectroscopic fitting algorithms such as AMARES return an estimate of the measure of spread in the fit: in the case of AMARES with jMRUI it is the CRLB directly; for Bayesian methods it would be the width of the posterior probability distribution function. The method used to incorporate this measure of uncertainty in the analysis of group data is described below, and is believed to be novel.

9.2.1 Introduction

It is often the case that the variability of spectral data is taken into account by imposing a quality minimum on the data; for example, by requiring the standard deviation to be below an arbitrary threshold. It is common practice to then subject only the population of data that are 'good enough' to further analysis. As an example, in the work described by Schroeder *et al.*,[257] hyperpolarised [1-^{13}C]pyruvate spectra were acquired from diabetic and non-diabetic rats and the ratio of the amplitudes of the bicarbonate-to-pyruvate peaks considered. Each subject's individual ratio was then treated as an exact measurement. Statistical analysis was performed in the Frequentist null hypothesis testing framework, with student's t-statistic calculated based on the variances of, and the observed difference between, the means of these groups of measurements. This approach is analogous to neglecting the error entirely on the peak amplitudes.

It was hypothesised that the correct propagation of uncertainty through analysis will improve accuracy, precision, and statistical power. The alternative approach developed allows the same conclusion to be drawn with greater certainty, or, alternatively, that fewer animals are required in each group in order to observe the same difference. This method is applicable broadly, to all studies where the measurement on a sample from a population has a quantified measure of uncertainty associated with it.

9.2.2 Methods

9.2.2.1 Weighted ("probabilistic") averaging

Fitting algorithms, like AMARES, return the estimated parameter of interest of a spectral line (denoted μ_i hence), together with a measure of the uncertainty on this measurement (σ_i). As spectroscopic estimates are obtained by measurements based on a large number of discrete points, and the noise in MR experiments is white, it is appropriate to consider each one of these measurements as random variables, drawn from a normal distribution with mean μ_i and variance σ_i^2. As the overall absolute detected magnitude of spectral peaks is highly dependent on experimental conditions, it is appropriate to consider measurements that are proxies for peak heights (e.g. peak integrals, or amplitudes fitted by AMARES) in spectroscopy experiments to be correlated. An elementary statistical result (derived

subsequently in Section 9.2.2.3) is that the best estimate for the mean ($\bar{x}$) of the population of these measurements is given by the weighted sum

$$\bar{x} = \frac{\sum_i^n \frac{\mu_i}{\sigma_i^2}}{\sum_i^n \frac{1}{\sigma_i^2}} \tag{9.5}$$

and the variance of the population, given the uncertainty on each measurement, is estimated by

$$\bar{\sigma}^2 = \left(\sum_i^n \frac{1}{\sigma_i^2}\right)^{-1} \left(\left(\sum_i^n \frac{1}{\sigma_i^2}\right)^2 - \sum_i^n \left(\frac{1}{\sigma_i^2}\right)^2\right)^{-1} \sum_{i=1}^n \frac{(x_i - \bar{x})^2}{\sigma_i^2}. \tag{9.6}$$

where n is the number of measurements made.

It is often desirable to calculate functions of the measured data, such as the ratio of the amplitude of two different spectral peaks.[3] In general, there is no exact expression for functions of normally distributed random variables, and numerical (Monte Carlo) simulations have to be performed to determine the shape of the resulting distribution. In the special case where we wish to consider the ratio of two random variables, the resulting distribution is not normal, but rather described by a ratio distribution [259,260], with a width that is reduced if the variables are correlated.

Consider two spectral peaks, of magnitude x and y with associated uncertainties σ_x and σ_y. If $\sigma_x/x \ll 1$ and $\sigma_y/y \ll 1$, and neither x nor y are close to zero, then the ratio distribution is very well approximated by the normal distribution.[261] In this case, the uncertainty on the ratio x/y is approximately

$$\sigma_r \approx \frac{x}{y}\sqrt{\left(\frac{\sigma_x}{x}\right)^2 + \left(\frac{\sigma_y}{y}\right)^2 - 2\rho\frac{\sigma_x}{x}\frac{\sigma_y}{y}} \tag{9.7}$$

where ρ is the correlation coefficient between x and y. Equation (9.7) can be derived from the multivariate Taylor series expansion for the function $f(x, y) = \frac{x}{y}$, and reduces to the common 'errors in quadrature' form commonly encountered if $\rho = 0$. The case where the error on measurement is large compared to the measurement made, i.e. where σ_x/x or σ_y/y are not much less than one, is more complex, and discussed separately (in Section 9.2.2.2).

Given two populations of such ratios of measurements, r_1 and r_2 with n_1 and n_2 elements in each, we can then compute the quantity

$$t = \frac{\bar{r}_1 - \bar{r}_2}{\sqrt{\frac{\bar{\sigma}_{r1}^2}{n_1} + \frac{\bar{\sigma}_{r2}^2}{n_2}}}, \tag{9.8}$$

[3] It should be noted that noise in MR is white in both the real and imaginary parts; therefore, magnitude data is Rician distributed, rather than normally distributed. The two distributions are the same in the high SNR limit, and the approximation is valid if SNR > 2.[258]

which approximately follows a t-distribution with

$$d = \frac{\left(\frac{\bar{\sigma}_{r1}^2}{n_1} + \frac{\bar{\sigma}_{r2}^2}{n_2}\right)^2}{(\bar{\sigma}_{r1}^2/n_1)^2/(n_1-1) + (\bar{\sigma}_{r2}^2/n_2)^2/(n_2-1)} \tag{9.9}$$

degrees of freedom. Therefore, in the Frequentist hypothesis testing framework, we can reject the null hypothesis that the two populations considered have the same mean if

$$p = 2 \times (1 - \text{pt}(|t|,\, d)) \tag{9.10}$$

is less than some chosen significance level (e.g. 0.05), where pt$(|t|,\, d)$ is the cumulative t-distribution function. Note that simple functions in popular computer programming languages to calculate weighted standard deviations and perform weighted t-tests are often lacking. Additionally, note that the estimates of uncertainty for the ratio are straightforward to incorporate directly into any Bayesian framework for evaluating differences between groups.

9.2.2.2 *Large errors*

Consider the case where the estimated error on a peak in an NMR spectrum is large, i.e. where the estimated σ/μ is *not* $\ll 1$. If we wish to consider the ratio of two peaks where the error is large, the resulting quantity has a Gaussian ratio distribution, described by Fieller, Hinkley and others.[260,262] We propose that the returned CRLB from the quantification of each peak can be treated as exactly defining a normal probability density function (PDF) for the quantity of interest for that peak. The resulting distribution for the ratio $w = x \sim \mathcal{N}(\mu_x,\, \sigma_x)/y \sim \mathcal{N}(\mu_y,\, \sigma_y)$ where x and y have a correlation coefficient ρ, has an analytic PDF $f(w)$ defined by:[260]

$$\begin{aligned} f(w) = {} & \frac{b(w)d(w)}{\sqrt{2\pi}\sigma_x\sigma_y a^3(w)} \left(\Phi\left(\frac{b(w)}{\sqrt{1-\rho^2}a(w)}\right) - \Phi\left(-\frac{b(w)}{\sqrt{1-\rho^2}a(w)}\right)\right) \\ & + \frac{\sqrt{1-\rho^2}}{\pi\sigma_x\sigma_y a^2(w)} \exp\left(-\frac{c}{2(1-\rho^2)}\right), \end{aligned} \tag{9.11}$$

where

$$a(w) = \left(\frac{w^2}{\sigma_x^2} - \frac{2\rho w}{\sigma_x\sigma_y} + \frac{1}{\sigma_y^2}\right)^{\frac{1}{2}} \tag{9.12}$$

$$b(w) = \frac{\mu_x w}{\sigma_x^2} - \frac{\rho(\mu_x + \mu_y w)}{\sigma_x\sigma_y} + \frac{\mu_y}{\sigma_y^2} \tag{9.13}$$

$$c = \frac{\mu_x^2}{\sigma_x^2} + \frac{2\rho\mu_x\mu_y}{\sigma_x\sigma_y} + \frac{\mu_y^2}{\sigma_y^2} \tag{9.14}$$

$$d(w) = \exp\left(\frac{b^2(w) - c(w)a^2(w)}{2(1-\rho^2)a^2(w)}\right); \tag{9.15}$$

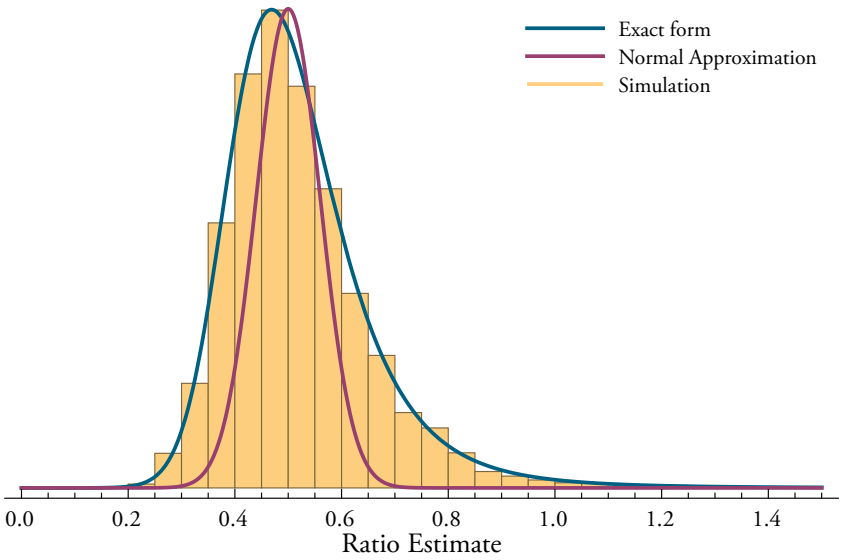

Figure 9.1: The PDF for the ratio of two normally distributed random variables, with $\mu_x = 1$, $\sigma_x = 0.2$, $\mu_y = 2$, $\sigma_y = 0.5$ and correlated with $\rho = 0.5$. Plotted here on common axes are the PDFs from the normal approximation (purple) given in Eq. (9.7), and the true analytic form stated in Eq. (9.11) *et passim* (blue). The median for both distributions is the ratio of means (i.e. $\frac{1}{2}$), but the means are distinct at 0.5 and 0.525. The orange boxes show a histogram of counts ($n = 10^4$) from the division of two normally distributed random variables generated with the parameters as given. The choice of values here is illustrative of a 'moderately pathological' case, where the normal approximation is clearly not appropriate, and should be contrasted to the 'highly pathological' case where $\mu_x = \mu_y = 0$ and $\sigma_x = \sigma_y = 1$ (i.e. a Cauchy distribution where no mean and variance exist) and the 'physiological' case where both means are far from zero (in which case the normal approximation is entirely valid, as can be seen in Fig. 9.2**B**).

and

$$\begin{aligned}\Phi(z) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{z} e^{-\frac{1}{2}u^2}\,\mathrm{d}u \\ &= \frac{1}{2}\left(1 + \mathrm{Erf}\left(\frac{z}{\sqrt{2}}\right)\right),\end{aligned} \tag{9.16}$$

where $\mathrm{Erf(z)}$ denotes the Gaussian error function, $\mathrm{Erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2}\,\mathrm{d}t$.

In general, this distribution is unimodal, not symmetric, and approaches the Cauchy distribution as μ_x and μ_y tend to zero, and σ_x, σ_y tend to one. The Cauchy distribution does not have a defined mean or variance. Several approaches to computing confidence limits on the ratio distribution have been proposed. Fieller's theorem allows for the computation of confidence limits on the ratio distribution, and can be derived from geometric arguments about the density of probability. Such confidence intervals are not guaranteed to exist as the ratio distribution approaches the Cauchy distribution; more information is provided in the review by Franz.[261] An illustrative example of the analytic form of the Ratio PDF is provided in Fig. 9.1, which shows both $f(w)$ (as given by Eqs. (9.11) to (9.15)), and the approximate normal quantity as given by Eq. (9.7). If both means are far from zero, and the uncertainty in measurements is small, the quantity described in Eq. (9.7) asymptotically tends to Eq. (9.11); however, it is far easier to numerically evaluate than the latter, which frequently equates to the product of

exceptionally large and exceptionally small quantities, where issues arising from the finite numerical precision of binary arithmetic cannot be neglected.

Fieller's theorem states that the lower and upper confidence limits (L_B, U_B) on the ratio $w = x/y$ are given by[261,262]

$$(L_B, U_B) = \frac{1}{\mu_y^2 - t_q^2\sigma_y^2}\Bigg((\mu_y\mu_x - t_q^2\sigma_{xy}) \pm \sqrt{\left(\mu_y\mu_x - t_q^2\sigma_{xy}\right)^2 - \left(\mu_y^2 - t_q^2\sigma_y^2\right)\left(\mu_x^2 - t_q^2\sigma_x^2\right)}\Bigg) \tag{9.17}$$

provided that μ_x is not expected to be close to zero, and where $\sigma_{xy} = \rho\,\sigma_x\sigma_y$, and t_q^2 is the t-statistic at the chosen significance level α, i.e. the value of the inverse cumulative t probability distribution function with $n-1$ degrees of freedom at $q = 1 - \frac{\alpha}{2}$. As spectral quantification is based on a large number of points (e.g. 4096) in the time or frequency domain, it is proposed to take $n \to \infty$, such that the t-distribution tends to the normal one. If the variances are small, at $\alpha = 0.68$ this expression is quantitatively similar to the familiar expression for the uncertainty on a ratio, Eq. (9.6).[262]

It is proposed that the width of the confidence interval defined above, i.e. $U_B - L_B$ at $\alpha \approx 0.68$ is an appropriate weight to use for subsequent calculations if the distribution is not too skewed. Note however that the ratio distribution $f(w)$ defined above may be decidedly non-normal, but that the central limit theorem ensures that the arithmetic mean of such ratio samples will be normally distributed in the limit that the number of samples becomes large, independent of this skew.

9.2.2.3 Derivation of the Weighted Mean and Variance

For the sake of completeness, a brief derivation of weighted expressions for mean and variance are presented, as it is difficult to find satisfactory equivalents in the literature. As discussed in Section 9.1, spectral fitting algorithms frequently return a quantity of interest from a spectral peak μ_i, (e.g. its amplitude) along with an estimate of uncertainty of that quantity (e.g. the CRLB), σ_i. Each measurement of μ_i may be considered to be independent and identically distributed. As noise in MR is Gaussian, it is justified to treat each of these results as being normally randomly distributed variables x_i, with $P(x_i) = \mathcal{N}(\mu_i,\, \sigma_i)$. If we have n such measurements, then the joint PDF defined is simply the product of each PDF, i.e.:

$$P(x_1, x_2, \cdots, x_n) = \prod_{i=1}^{n} \mathcal{N}(\mu_i, \sigma_i) \tag{9.18}$$

It is then clear that the expectation value for the mean of this distribution, which is the best estimator for the quantity of interest, is given by

$$\bar{x} = \frac{\int_{-\infty}^{\infty} x \cdot \prod_{i}^{n} \mathcal{N}(\mu_i, \sigma_i)\,\mathrm{d}x}{\int_{-\infty}^{\infty} \prod_{i}^{n} \mathcal{N}(\mu_i, \sigma_i)\,\mathrm{d}x} \tag{9.19}$$

$$= \frac{\sum_{i}^{n} \frac{\mu_i}{\sigma_i^2}}{\sum_{i}^{n} \frac{1}{\sigma_i^2}} \tag{9.20}$$

i.e. we take a weighted mean, where each weight w_i is the reciprocal variance of each PDF, $w_i = 1/\sigma_i^2$, normalised by the sum of the weights $W = \sum_i w_i$:

$$\bar{x} = \frac{1}{W} \sum_i w_i x_i.$$

Additionally, we desire an unbiased estimator of the variance of the underlying distribution σ^2.

In analogy with the method for computing Bessel's correction for the sample variance,[4] we recall the standard definition of the sample variance V and compute its expectation. We then infer a bias correction factor, and, hence, an unbiased weighted variance estimator.

$$\begin{aligned}
V &\stackrel{\text{def}}{=} \frac{1}{W} \sum_i w_i (x_i - \bar{x})^2 \\
&= \frac{1}{W} \sum_i w_i \left(x_i - \frac{1}{W} \sum_j w_j x_j \right)^2 \\
&= \frac{1}{W} \sum_i w_i \left(x_i^2 - \frac{2}{W} x_i \sum_j w_j x_j + \frac{1}{W^2} (\sum_j w_j x_j)^2 \right) \\
&= \frac{1}{W} \left(\sum_i w_i x_i^2 - \frac{1}{W} \sum_i w_i x_i \sum_j w_j x_j \right) \\
&= \frac{1}{W} \left(\sum_i w_i x_i^2 - \frac{1}{W} \sum_{ij} w_i w_j x_i x_j \right)
\end{aligned} \tag{9.21}$$

We now compute the expectation value for this estimator, $\langle V \rangle$. Recall that, for normally distributed random variables, $\langle x_i \rangle = \mu$, and $\langle x_i x_j \rangle = \mu^2 + \delta_{ij}\sigma^2$. We additionally assume that any random contribution dependent on the weights themselves is zero or can be neglected in comparison to σ, such that $\langle w_i x_i \rangle = w_i \langle x_i \rangle$ and $\langle w_i w_j x_i x_j \rangle = w_i w_j \langle x_i x_j \rangle$.

[4] *Id est* the factor of $N - 1$ in the denominator

$$\begin{aligned}
\langle V \rangle &= \frac{1}{W}\left(\sum_i w_i \langle x_i^2 \rangle - \frac{1}{W}\sum_{ij} w_i w_j \langle x_i x_j \rangle\right) \\
&= \frac{1}{W}\left(\sum_i w_i(\mu^2+\sigma^2) - \frac{1}{W}\sum_{ij} w_i w_j \left(\mu^2 + \delta_{ij}\sigma^2\right)\right) \\
&= \frac{1}{W}\left(W(\mu^2+\sigma^2) - \frac{1}{W}(W^2\mu^2 + (\sum_i w_i^2)\sigma^2)\right) \\
&= \frac{1}{W}\left(W\sigma^2 - \frac{1}{W}(\sum_i w_i^2)\sigma^2\right) \\
&= \left(\frac{W^2 - \sum_i w_i^2}{W^2}\right)\sigma^2
\end{aligned} \tag{9.22}$$

Therefore an unbiased estimator U of σ^2 is

$$\begin{aligned}
U &= \frac{W^2}{(W^2 - \sum_i w_i^2)}\langle V \rangle \\
&= \frac{W^2}{(W^2 - \sum_i w_i^2)}\frac{1}{W}\sum_i w_i (x_i - \bar{x})^2 \\
&= \frac{W}{(W^2 - \sum_i w_i^2)}\sum_i w_i (x_i - \bar{x})^2
\end{aligned} \tag{9.23}$$

which is identical to Eq. (9.6), where the weights have been stated explicitly.

This expression has several reassuring properties: it is invariant under multiplicative weight transformation; and in the case that all weights are equal, $w_i = w$, say, the pre-multiplying factor reduces to $N/(N^2 - N) = 1/(N-1)$, which is the same pre-multiplying factor for the familiar expression of an unbiased estimator of variance from data with equal uncertainty.

An outline of the analysis process is shown in Fig. 9.2, with example data provided as described below. In summary, not all data are acquired with equal certainty (Fig. 9.2**A**), and fitting algorithms often return a quantitative estimate of this uncertainty. When functions of uncertain variables are required, such as the normalising ratio of two peaks, a probability distribution is defined for the resulting quantity, and Fig. 9.2**B** shows an example empirical PDF for the ratio of two spectral peaks, together with the approximation Eq. (9.7). This leads to a population of values with associated uncertainty, from which the weighted mean and variance can be computed (Fig. 9.2**C**).

9.2.3 Data Analysis

9.2.3.1 Hyperpolarised Spectra

The authors of a previous study on the spontaneously hypertensive rat were contacted, and asked to provide original data for re-analysis. As described in Dodd *et al.*,[94] hyperpolarised spectra were obtained following the injection of [1-^{13}C]pyruvate into rats with a series of genetic mutations that result in a spontaneously hypertensive phenotype (the SHR rat), and from a set of matched control animals.

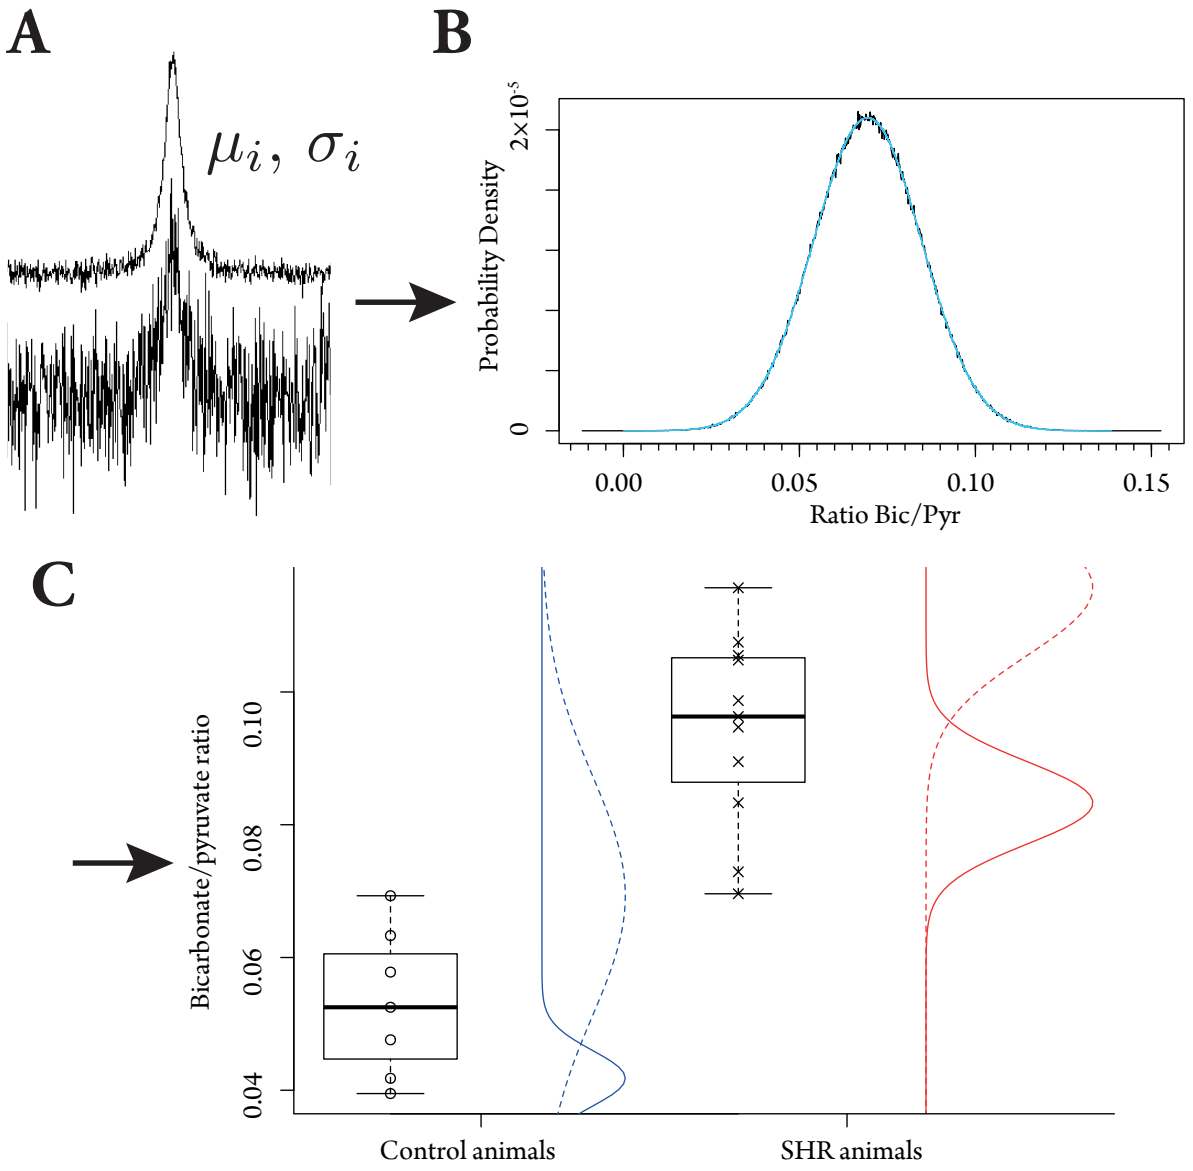

Figure 9.2: An overview of the proposed analysis method. **A**: Not all acquired spectral data are equal. Fitting algorithms return an estimate for the amplitude of each peak height (μ_i), and a measure of uncertainty in that height (σ_i); shown are two spectra with identical μ but different σ. **B**: These can be used to generate PDFs for quantities of interest. Shown here is an empirical (Monte Carlo) PDF for the ^{13}C-bicarbonate-to-[1-^{13}C]pyruvate (bic/pyr) ratio as determined by hyperpolarised spectroscopy in a healthy rat (black line). As the error on measurement in this experiment was small, it is very well approximated by a normal distribution with a mean given by the ratio of peak heights, and a standard deviation given by Eq. (9.7) (almost indistinguishable cyan line). **C**: The resulting population information contains a measure of uncertainty at each point for the two populations. These should be respected when trying to distinguish between groups. Shown here are the distribution of mean bic/pyr measurements for control and SHR animals. The most uncertain (dotted line) and precise (solid line) measurement for each population are shown with their PDFs plotted on the y-axis separately, to illustrate the variability within the data.

The apparent rate of ^{13}C-bicarbonate production was significantly increased in the hypertensive animals compared to controls, corresponding to a believed (genetic) up-regulation in the quantity of PDH present. This was subsequently verified through extensive *ex vivo* biochemical techniques. For each animal's dataset, the ratio of the ^{13}C-bicarbonate to [1-^{13}C]pyruvate peaks of the summed hyperpolarised spectra is considered to be an appropriate normalised measure of PDH flux. Peak amplitudes were quantified by AMARES and jMRUI, with the data summed over time.

9.2.4 Results

Properly accounting for the variability in the data as described above improved statistical power. The known biochemical difference in PDH flux in the SHR rat was detected with a p value of 1.1×10^{-6} when error was propagated as described here, compared to 5.1×10^{-5} assuming the data have equal weight. In this case, the numerical effect of weighted averaging was small, largely reflecting relative homogeneity in weights. Weighted averaging reduced the population variance in SHR animals from 2.1×10^{-4} to 2.0×10^{-4}, and in control animals from 1.4×10^{-4} to 1.2×10^{-4}. The mean changed slightly, from 0.053 to 0.057 in control animals, and from 0.094 to 0.096 in SHR animals. In short, using the weighted method reduces the observed apparent variability in each population; the coefficient of variation (defined as $C_v = \sigma/\mu$) was reduced from 15.36% (without incorporating uncertainty) to 14.88% (with the outlined analysis) for the SHR rat data, and from 21.24% to 20.79% for controls.

9.2.5 Discussion

This retrospective analysis has reiterated the practical need to propagate uncertainty properly in spectroscopic studies. The instinctive, straightforward way proposed to easily incorporate measures of spectral uncertainty in subsequent analysis has been shown to improve statistical power in a setting where a clear biological difference is known *a priori* to exist between the SHR rat and normal wild-type controls.

Conventional NMR (*in vivo* or otherwise) suffers from inherent SNR limitations, and the analysis of low/variable SNR spectra remains an important issue. Historically, the balance pitched between short scan-times, tolerable for patients and yet long enough for adequate signal-to-noise, has resulted in a broad range of published values for many quantities, such as the healthy cardiac PCr/ATP ratio. Historically, this has lead to embarrassment arising from the poor interpretation of uncertain data.[263] Note, additionally, that this problem is directly analogous to that of recombining multi-coil spectroscopy data, for which weights that are the square of the estimated SNR are optimal.[264]

Whilst hyperpolarisation offers a way to improve the inherent SNR limitations of magnetic resonance, small variations in the timing of experiments can have a large impact in the resulting quality of the data acquired. In many situations other than MR an estimate of uncertainty on a given value could be

obtained before functions of that value are then used. As information regarding reproducibility is of great importance to the design and analysis of clinical studies, ensuring that experiments are statistically maximally empowered is of great value. In preclinical settings, an increase in statistical power could be traded for a reduction in the number of experimental animals required to reach the same conclusion. Properly propagating uncertainty is therefore ethically favoured.

A valid criticism of the approach described above is that it is not necessarily the case that the least certain data are the least accurate: a noiseless measurement of a flawed experiment is given greater weight in the above analysis than of a noisy, but unbiased one. There are many factors affecting the scan-to-scan SNR in spectroscopy, from coil loading differences to changes in the effective T_r in a cardiac gated experiment, to say nothing of 'unknown unknown' factors that may influence an experiment. However, in the two datasets considered here the least certain measurements have proved to be away from the expected value from an otherwise identical cohort. If all other experimental variables are accounted for, the uncertainty on measurement should be respected, and a comparison between the (standard, and universally understood) unweighted and weighted analyses made.

It is proposed that the adoption of methods which incorporate estimates of uncertainty, such as the one described here, will allow the maximal use of data acquired with varying precision. The analysis proposed is conceptually and computationally simple and offers immediate value to existing data. The improved statistical power provided by the technique may help elucidate small biological differences between small populations of measurements, or avoid false leads.

9.3 | The Fast Padé Transform in Hyperpolarised Spectroscopy

9.3.1 Introduction

Prior to the development of spectral-spatial imaging sequences such as those discussed in Chapters 5 and 6, Chemical Shift Imaging (CSI) was the predominant method of choice for imaging hyperpolarised metabolism. If the form of spectra obtained are not known *a priori*, CSI sequences effectively represent the only way to obtain highly spatially resolved information about the *in vivo* behaviour of the injected probe. Correspondingly, they are commonly used when developing novel metabolic probes, and additionally have the advantage of being robust to errors in transmitter frequency. However, the inevitable return to thermal equilibrium of hyperpolarised compounds places an ultimate limit on the duration available for imaging, on the order of $\sim 3 \times T_1$. For [1-^{13}C]pyruvate *in vivo*, this window is typically only 50-80 seconds;[265] for other probes under development, it may be even shorter, *e.g.* $\sim$10 s for [U-^{13}C, 7-^{2}H]glucose. Naïve approaches to fully-sampled CSI sequences necessarily take on the order of 60 seconds to acquire an image with matrix sizes of $\sim 32 \times 32$, in part due to the necessarily long readout portion of the imaging sequence: the inverse relationship between readout length and spectral bandwidth and the SNR implications for changing receiver bandwidth (c.f. Section 2.4.5) ensure

that a long readout is desirable. However, numerous strategies have been proposed for gaining higher spatial and temporal resolution with CSI sequences, usually at the expense of spectral resolution.[266–268] Typically these approaches drastically reduce the quantity of data acquired, but traditional reconstruction methods may not perform adequately without constraints, often in the form of the imposition of prior spectral information.

In the vast majority of reported uses, the identity and relative concentrations of hyperpolarised metabolites in a CSI dataset are determined by Fourier transforming the time-domain data, before measuring the individual peak heights or integrating peaks in the spectra.[269,270] Such methods are robust if, and only if, peaks are well-separated and do not overlap. Both AMARES and Hankel Singular Value Decomposition (HSVD) have also been used to quantify CSI data. Presented here is an alternative approach, using the Fast Padé Transform (FPT) as opposed to the traditional Fast Fourier Transform (FFT) followed by an analysis method, as a means to directly quantify simulated CSI data.

Fourier transforms, usually realised by the Fast Fourier Transform algorithm, are the traditional way to transform data between time-domain and spectral domain. If the Nyquist criterion is satisfied, the FFT forms a robust and very computationally efficient tool to obtain the spectral envelope of spectroscopic datasets when no prior knowledge is present. If the time-domain data is effectively truncated, as is often the case with an accelerated acquisition or in the presence of a large quantity of noise, it is common practice to pad the end of the acquired FID with zeros, prior to Fourier transforming the data, in order to increase the apparent reconstructed resolution. Zero-filling effectively induces a sinc-shaped interpolation in the spectral domain, and is one of a number of possible methods for spectral interpolation.[271] Owing to the comparatively high energy of the oscillating side-lobes of the sinc function, for any pair of closely spaced resonances, the effect of truncating the acquisition of the FID can be to produce spurious peaks or dips in the resulting spectrum. Correspondingly, parametric approaches are often taken to minimize the effect of this interpolation and quantify spectra by fitting Lorentzian or Gaussian lineshapes to their envelopes in a least-squares sense.

An alternative approach to fitting spectral data is the direct extraction of frequency and amplitude information from the FID by methods other than the FFT. Parametric estimators of this sort have previously been shown to be of use in MR, with methods such as sub-band decomposition,[272] or the wavelet transformation.[273] Previous work has shown that the Fast Padé Transform can enhance the readout resolution of NMR spectra and to be less affected by truncation artefacts and baseline modulations than FFT alone.[274–277] A brief overview of the the FPT method follows.

9.3.2 *Theory & Methods*

Let $S = c_n$ be a discretely sampled FID, where c_n is the nth sampled point out of N samples in total. Then, in general, the discrete Fourier components of S are defined as

$$\hat{S}(\omega) = \frac{1}{N}\sum_{n=0}^{N} c_n e^{in\omega\tau}$$
$$\Rightarrow S(z) \stackrel{\text{def}}{=} \frac{1}{N}\sum_{n=0}^{N} c_n z^n \tag{9.24}$$

where τ is the sampling increment between signals, and $z = e^{i\omega\tau}$ has been defined for notational convenience.

The Padé transform essentially acts as an analytically-continued form of the Laplace transformation with a Padé approximant.[278] The Padé approximant for any function $f(x)$ (with at most a finite number of discontinuities) is a rational function of two polynomials with the form

$$f(x) \approx \frac{A(x)}{B(x)} \stackrel{\text{def}}{=} \frac{a_0 + a_1 x + a_2 x^2 + \cdots + a_m x^m}{1 + b_1 x + b_2 x^2 + \cdots + b_n x^n}, \tag{9.25}$$

where the 1 in the denominator can be assumed without loss of generality. It can be shown that it is possible to find a_i and b_i such that the Taylor expansion of Eq. (9.25) agrees with that of $f(x)$ directly, if it exists. The richer algebraic structure of Padé approximants often means that the series are much 'better behaved' than the corresponding Taylor series: for example, they do not necessarily diverge as $x \to \infty$. As a consequence, like Chebyshev polynomials (which are analogous to Taylor series defined on a finite domain) they have found use in numerical analysis and applied mathematics.[279,280] An illustration of a Padé approximant and Taylor series of equivalent order is provided in Fig. 9.3. For a given $f(x)$, its Padé approximant is unique.

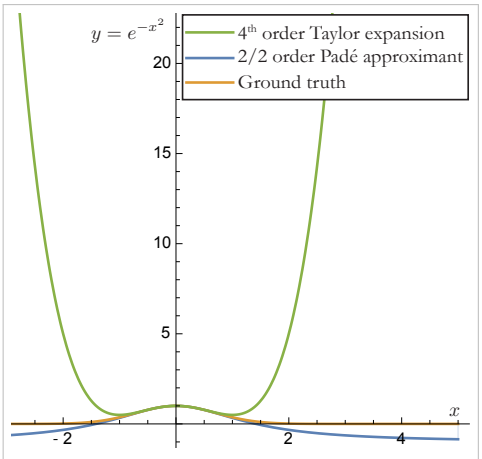

Figure 9.3: Padé and Taylor Series expansions of equivalent order for the function $y = e^{-x^2}$.

The Fast Padé Transform, whose properties have been extensively mapped by Belkić *et al.*[274–277,281–285] essentially functions by equating Eqs. (9.24) and (9.25), and creating an analytic approximator for the time-domain data directly. The specific values for a_i and b_i in Eq. (9.25) are found by solving sets of linear equations defined by equating powers in the expansion $A(z) = B(z)S(z)$.[5] The Maclaurin expansion of Eq. (9.25) converges outside the unit circle $|z| > 1$, but the approximant can be shown by analytic continuation to be convergent in the entire complex plane with the exception of a simple pole at $z = 0$ and the complex disc defined by $|z| = 1$. Once an approximant of the acquired data is constructed, the peak amplitudes in the resulting spectrum can be determined immediately, without a fitting procedure,

[5] Subject to the condition that the vacuous solution $A(z) = B(z) = 0$ is avoided. This is practically done via a SVD of a large Toeplitz matrix.

as follows. The polynomial $B(z)$ is rooted, and the resulting k roots correspond exactly to the peaks in the spectrum; $z_k^{-1} = e^{-i\omega_k \tau}$, from which the frequencies are found as $\omega_k = (i/\tau) \ln z_k^{-1}$, for which several efficient numerical algorithms exist. The corresponding peak amplitudes d_k for each ω_k can be found from the Cauchy residue of the approximant as a whole evaluated at the given amplitude, i.e.

$$d_k = \left. \frac{A(z)}{\frac{\partial B(z)}{\partial z}} \right|_{z=z_k} \tag{9.26}$$

Accordingly, the amplitude and spectral position of resonances can be recovered, and the spectrum generated for display in any desired mode (absorption, dispersion, magnitude) by the use of Heaviside partial fractions, i.e. by the expansion

$$\frac{A(z)}{B(z)} = \sum_k \frac{d_k}{z - z_k}. \tag{9.27}$$

It can be shown that $\mathrm{Re}(\omega_k)$ and $\mathrm{Im}(\omega_k)$ are the position and width of the resonance k, and $|d_k|$ and $\arg(d_k)$ are its amplitude and phase. In this fashion, the FPT can find every peak parameter of interest for a resonance, without ever using the Fourier spectrum, or a fitting procedure. In the case that $B(z)$ contains unique roots, the returned spectrum consists of several Lorentzians; if $B(z)$ has repeating roots then the resulting spectrum contains spectrally overlapping resonances with a non-Lorentzian lineshape. Depending on the shim of the region of interest, these multiple resonances may be spurious. Arguably the most important aspect of the FPT is that, as a rational approximant, a spectrum can be reconstructed *post hoc* with any desired number of points, providing an alternative method for spectral extrapolation other than zero-filling.

Any experimentally acquired spectrum will contain additional Gaussian noise. The FPT generates a fundamentally nonlinear response function from the acquired data, and the propagation of noise through this process is not straightforward. Whilst this nonlinear response is ultimately responsible for the ability of the FPT to effectively suppress noise, one effect of the presence of noise in the FPT formalism are simultaneous zeros of $A(z)$ and $B(z)$, *i.e.*, for some spurious resonance z'_k, $A(z'_k) = B(z'_k) = 0$, and where the Cauchy principle part is such that the residue at this part, related to d_k, is zero. For historical reasons, such z'_k are referred to as Froissart doublets, and can lead to artefacts in the approximation. With real experimental data it may be the case that there is some small numerical difference between where the residue of the pole is zero, and where the function is undefined.[286] These doublets can also be observed through the iteration of FPT as non-converged frequencies and can be removed by a threshold in relation to the background noise levels. Here, the noise threshold contribution was set as amplitudes occurring within 5 standard deviations of the noise in the spectrum, which was manually identified. It is noted that an iterative approach, where the number of sampled points is varied can identify spurious signals, as such signals have a higher likelihood of drastically changing parameters in-between iterations. These spurious poles often appear near $|z| = 1$.

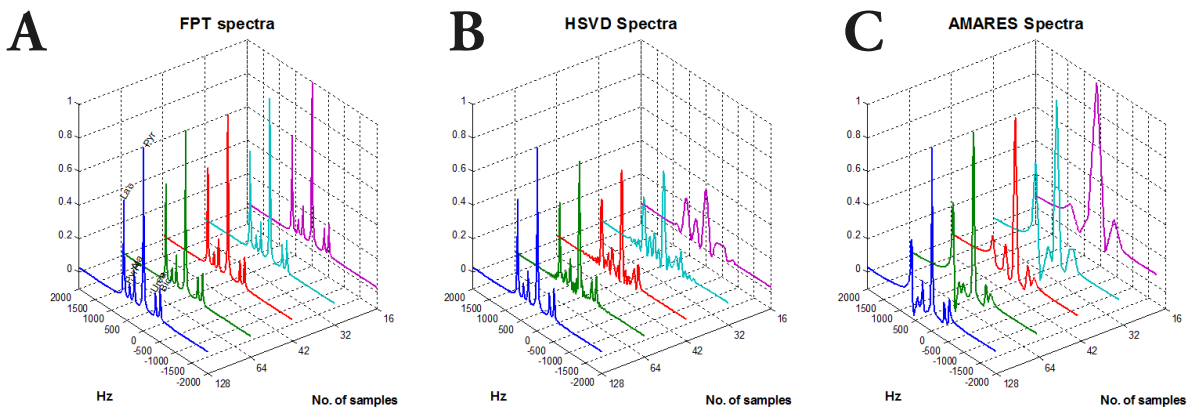

Figure 9.4: A comparison of the effect of truncation of the reconstruction of spectral data through three different methods: **A**: FPT, **B**: the Hankel Singular Value Decomposition of a zero-filled FID; and **C**: AMARES reconstruction based on fitting a basis function to the truncated FID. Data courtesy of Esben Hansen.

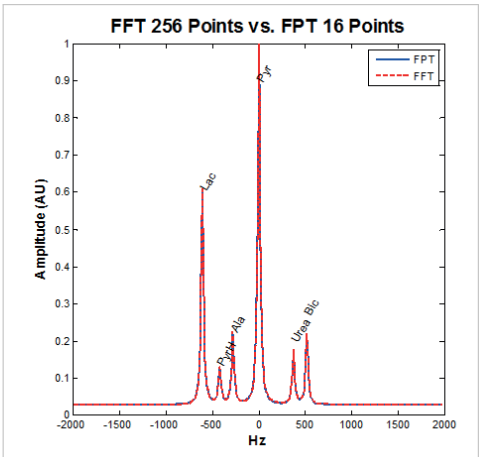

Figure 9.5: A comparison between the highly truncated 16-point FPT and fully-sampled 256-point FFT. Figure courtesy of Esben Hansen.

Simulated time-domain data from an NMR spectrum with five peaks typical of a [1-^{13}C]pyruvate experiment was generated in jMRUI, without noise. This data was processed with the FFT with 256 points over a 4000 Hz bandwidth and a truncated FPT with 16 points (numerator polynomial degree $K = 8$, denominator polynomial degree $Q = 8$) and the same bandwidth, interpolated to 256 points. Frequency components found by the FPT were reconstructed both with a Gaussian and a Lorentzian model for comparison. The truncated data were processed by FPT, HSVD and AMARES in their respective time or frequency domain as comparison between FPT and established methods. The transform was carried out with code written by Esben Søvsø Szocska Hansen and the author, and used the method provided by Trefethen *et al.* to identify spurious resonances.[286] As an *in vivo* comparison, [1-^{13}C]pyruvate 2D CSI data previously acquired[270] by Esben Søvsø Szocska Hansen from the healthy rat kidney was also analysed (acquisition spectral width = 4000 Hz, 256 complex points, flip angle$= 10°$, a fully-sampled k-space trajectory taken in a centric order, matrix$= 16 \times 16$, TR/TE=75 ms/0.65 ms, FOV$= 60 \times 60 \times 20$ mm^3 on a 4.7 T horizontal bore Varian magnet). Different levels of truncation were investigated, where the reconstruction was performed with $1/2$, $1/4$, $1/6$, $1/8$ and $1/12^{\text{th}}$ of the acquired data.

9.3.3 *Results*

As shown in Fig. 9.5, the FPT of the simulated, noiseless data identified true resonances even in the presence of extremely truncated data. No prior knowledge was required for identification of the true

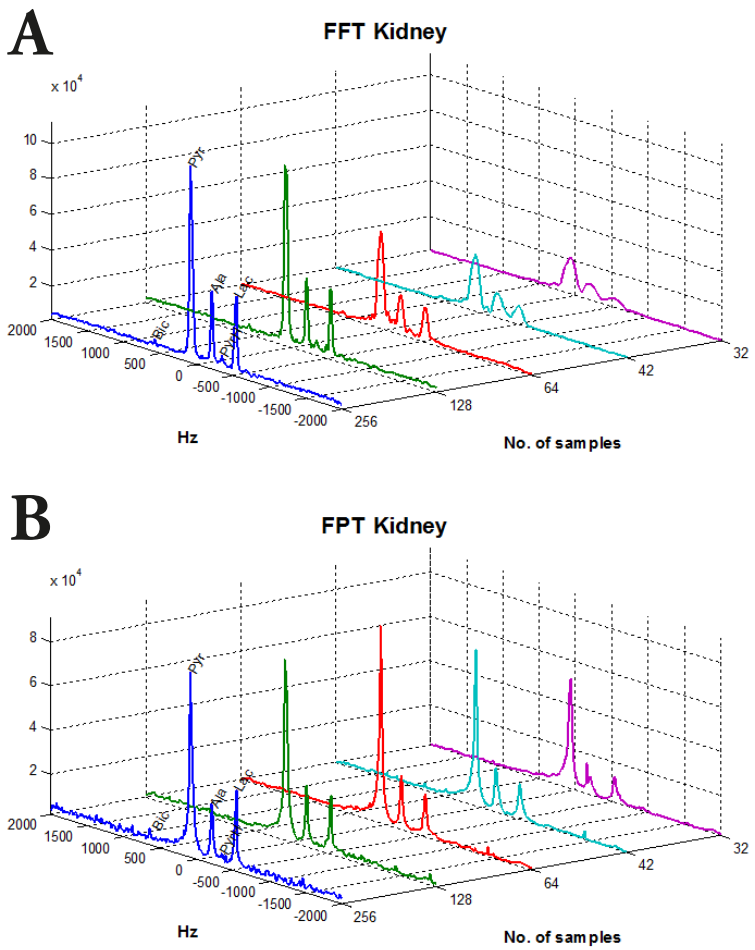

Figure 9.6: Retrospectively truncated *in vivo* data that were reconstructed by either **A**: conventional zero-filling and Fourier transform, or **B**: the FPT. It is apparent that the spectral line-shape is preserved for higher truncation factors when the spectrum is reconstructed via FPT as opposed to FFT. Data courtesy of Esben Hansen.

resonances, and it was found that (as expected) the Lorentzian model lineshapes fits the data. The inherent scope for spectral interpolation provided by the FPT method allows for high-quality readout even with a small number of points. As illustrated in Fig. 9.4, it was found that, in comparison to other methods such as HSVD and AMARES, the FPT reconstruction of the truncated dataset maintained spectral features such as Lorentzian lineshapes and accurate centre peak locations, whereas HSVD produced a distorted spectrum with significant artefacts after truncation to $\leq 1/4$ of the original spectrum. AMARES returned a poorly-phased spectrum with incorrect lineshapes and a merging of peaks in all the truncated spectra, despite the imposition of appropriate prior knowledge. This is believed to represent a convergence upon a local minimum.

The retrospective truncation of the *in vivo* data was promising, but not as flawless as the reconstruction obtained in the simulated noiseless case. As shown in Fig. 9.6**A**, zero-filling by a factor greater than two prior to FFT substantially distorts the resulting spectrum, whereas the use of the Padé transform appears to offer the possibility of a ~ 4-fold acceleration with minimal distortion (Fig. 9.6**B**). It was

noted that the returned peak amplitudes appeared more robust to truncation with FPT than FFT, as the apparent reconstructed lineshape was subject to less sinc ringing, and was correspondingly narrower. However, it should be noted that the use of other processing techniques, such as apodisation, improves the reconstruction of truncated and subsequently zero-filled data prior to Fourier transformation.

9.3.4 Discussion & Conclusions

It is believed that the FPT offers a novel method for reducing the acquisition time for MR spectroscopy without significantly decreasing the apparent resolution of the NMR spectrum. The preliminary data acquired here support the thorough investigation of the use of the transform to reconstruct prospectively accelerated data. If it is found that a four-fold acceleration factor is supported by the transform (as may appear to be the case), it is hoped that the resulting increase in available scan time could be traded for a modest increase in spatial resolution, which would significantly enhance the utility of CSI sequences in developing novel metabolic probes. The use of the FPT enables the prospect of performing a truncated CSI acquisition but recovering an essentially complete spectrum, which is not possible with conventional reconstruction methods. It is worth noting that the FPT necessarily assumes that the underlying data has either a Lorentzian or Gaussian form; in practice, poor B_0 homogeneity may violate this assumption, and the algorithm may therefore return a large number of smaller peaks in a given region. It is proposed that summing the contributions from each defined peak in a given region would be an appropriate way to proceed in this situation, analogously to peak integration following Fourier transform but with more spectral selectivity.

Additionally, *a priori* knowledge about an experiment is easily incorporated into the FPT by the process through which Froissart doublets are identified and removed,[287] in a way which is analogous to guided HSVD methods. It might however be proposed that the complete lack of prior information is advantageous in the reconstruction of spectral data from potentially unidentified reactions following the injection of a novel probe.

The main advantages of the approach explored here are the direct nature of the quantification of a spectral peak of interest, and its robustness to truncation, both of which are of interest in a hyperpolarised setting. The comparatively short T_1 of many carbon compounds, together with the large number of nodes in the metabolic network that have *not* been explored with DNP, ensure that there is a role for accelerated spectroscopic imaging sequences. Further work will demonstrate the prospective use of the FPT to reconstruct [1-^{13}C]pyruvate data, and use it to directly accelerate spectroscopic imaging acquisitions, by decreasing the time taken to acquire an FID. Additionally, as the FPT has been successfully generalised to two dimensions,[284] it is hoped that it may be possible to use it (or other interpolant methods) to aid in the reconstruction of hyperpolarised ^{13}C images, which necessarily have reduced spatial resolution in comparison to ^{1}H images.

9.4 | *Fitting Models to Metabolism*

9.4.1 *Introduction and Overview*

The hyperpolarised imaging strategies discussed in this work provide volumetric data that varies over time. For example, the three-dimensional imaging sequence developed in Chapter 6 returns a five-dimensional cube of images, showing the apparent concentration of multiple metabolites in three dimensional space as a function of time. Interpreting such data directly is difficult, although, as outlined in Chapter 4, much is known *a priori* about metabolism, and about the expected properties of the data acquired. As a result, several different kinetic models have been proposed to quantify the data that hyperpolarised experiments provide. These have the promise of reducing the complexity of the data acquired (by eliminating time), and should be more robust to noise, being able (automatically) statistically discriminate between signal and noise. This is in contrast to the practice of simply temporally summing data, where the limits of summation have an undesirably arbitrary effect on the final data.[6] It has frequently been reported that the most commonly used metrics – the summed metabolite-to-pyruvate ratio, the ratio of temporal integrals, and simple kinetic models for uptake – are highly correlated with each other, as has been seen previously in this work in Fig. 4.5.[167,288] This fact is unsurprising, as all returned quantities should report on the same underlying process of interest. Model selection therefore is non-trivial, and represents a compromise between the appropriateness of a model for describing the data without over-parametrisation (as quantified by the usual information criteria, namely the BIC and AIC), the ease of fitting it to data, and some subjective understanding of how well the proposed set of differential equations encapsulates the processes believed to be occurring.

Every model that has been proposed to date for the analysis of DNP data is spatially invariant, and the majority are based upon first-order chemical kinetic exchange between the injected metabolite, usually [1-^{13}C]pyruvate, and downstream metabolites (such as [1-^{13}C]lactate). Owing to the non-renewable magnetisation of hyperpolarised experiments, correction factors have to be introduced for the T_1 decay of each metabolite. More complex models additionally account for the effect of RF depletion (which otherwise would be incorporated into the apparent T_1 for each molecule), or attempt to discriminate between vascular and tissue signal by having a multiple-box type model. An uptake function is usually required to model the arriving bolus of contrast agent, although sometimes the data is simply truncated instead; for example, with the two-pool Lac $\underset{k_P}{\overset{k_L}{\rightleftharpoons}}$ Pyr model

[6] One wishes to sum only where there is signal, but this region is not necessarily known *a priori*.
As a consequence, approaches have been proposed where an estimate of the noise term in each image is used to construct a weight for the contribution of that image. However, it may be argued that a statistically robust method – which could, for example, look at noise correlation matrices as a way of discriminating between signal and white noise – would represent an attractive alternative to weighted averaging, and reduce inter-user variability by discriminating between frames with and without signal.

$$\begin{aligned}\frac{\partial L_z}{\partial t} &= -\frac{1}{T_{1,\,\text{Lactate}}}\left(L_z - L_\infty\right) + k_P P_z - k_L L_z \\ \frac{\partial P_z}{\partial t} &= -\frac{1}{T_{1,\,\text{Pyruvate}}}\left(P_z - P_\infty\right) + k_L L_z - k_P P_z\end{aligned} \tag{9.28}$$

first proposed by Day *et al.*,[289] in which P_z and L_z are the visible M_z for [1-^{13}C]pyruvate and [1-^{13}C]lactate respectively, P_∞ and L_∞ represent those visible concentrations at thermal equilibrium, and other symbols have their usual meaning. This model ignores exchange between pyruvate and other molecules, does not account for the incoming bolus of hyperpolarised pyruvate nor RF depletion of M_z, but is conceptually simple, has a small number of parameters, and can easily be fit to acquired data with a small residual (which will be encapsulated in a low AIC/BIC).

9.4.2 Zierhut's Metabolic Model

As an acceptable compromise between parametrisation and power, the model first proposed by Zierhut *et al.*[151] has been used for the analysis of the spectroscopic data acquired in Chapter 4. This model is a piecewise-integrated variant of first-order exchange between [1-^{13}C]pyruvate and other metabolites X, with additional decay terms and an arbitrary maximum intensity fixed by continuity between all three parts, but denoted M_X for the sake of notational convenience. This model assumes:

- That the appearance of [1-^{13}C]pyruvate in the data occurs after its injection into the blood stream, which occurs at a linear rate R_i (ADC units/s) and starts at a time $t_a > 0$;
- That the decay phase [1-^{13}C]pyruvate occurs at a net exponential rate described as k_{Pyr} (s^{-1}), but which only begins at a time $t_e > t_a$;
- That the appearance and decay of the other visible metabolites; lactate, alanine, pyruvate hydrate, bicarbonate and CO_2 can only occur after the appearance of [1-^{13}C]pyruvate. The rate of this interconversion, $k_{\text{Pyr}\to X}$ (s^{-1}) or $k_{\text{Pyr}X}$ for short, is assumed linear. Each visible metabolite decays (via either T_1 relaxation or reverse reactions) at a rate k_X (s^{-1}).

This gives rise to the following piecewise defined model as a function of time t for the arbitrary signal intensity $S^X(t)$ for metabolite X:

$$S^{\mathrm{Pyr}}(t) = \begin{cases} 0 & t < t_a \\ \frac{R_i}{k_{\mathrm{Pyr}}}\left(1 - e^{-k_{\mathrm{Pyr}}(t-t_a)}\right) & t \in [t_a,\, t_e] \\ M_{\mathrm{Pyr}} e^{-k_{\mathrm{Pyr}}(t-te)} & t > t_e \end{cases} \tag{9.29}$$

$$S^{X\neq\mathrm{Pyr}} = \begin{cases} 0 & t < t_a \\ \frac{k_{\mathrm{Pyr}X}\cdot R_i}{k_{\mathrm{Pyr}}-k_X}\left[\frac{1-e^{-k_X(t-t_a)}}{k_X} - \frac{1-e^{k_{\mathrm{Pyr}}(t-t_a)}}{k_{\mathrm{Pyr}}}\right] & t \in [t_a,\, t_e] \\ \frac{M_{\mathrm{Pyr}}\cdot k_{\mathrm{Pyr}\,X}}{k_{\mathrm{Pyr}}-k_X}\left[e^{-k_X(t-t_e)} - e^{-k_{\mathrm{Pyr}}(t-t_e)}\right] + M_X e^{-k_X(t-te)} & t > t_e \end{cases} \tag{9.30}$$

Depending on the number of metabolites analysed (i.e. whether or not the data being quantified is spectral), Eqs. (9.29) and (9.30) represent between four ([1-^{13}C]pyruvate alone) and fourteen (all visible metabolites) parameter models. As a direct consequence, the fit is commonly performed in a constrained manner, with the [1-^{13}C]pyruvate time course (i.e. Eq. (9.29)) being fit before downstream metabolites, which depend upon properties of the found pyruvate timecourse. The metric chosen is usually the Euclidean norm with respect to the returned peak amplitudes after quantification (by, e.g. AMARES) of the complex data. There is a small (but growing) corpus of spectroscopy studies which have used this model specifically,[93,94,152,290,291] as it is believed to represent an acceptable compromise between complexity and descriptive power.

9.4.3 Hyperpolarised Image Bayesian Analysis

As a method to reduce the complexity of the datasets generated in this work, the model described by Eqs. (9.29) and (9.30) was fit to representative image stacks acquired with the sequence of Chapter 5. Owing to the fact that Eqs. (9.29) and (9.30) are spatially invariant, this was performed simply by iterating over all voxels and fitting with respect to time. This is tantamount to assuming spatial metabolic homogeneity over the volume of interest; an assumption that is probably more valid in organs such as the brain, kidneys and liver than it is the heart. Initially, the fitting process was performed by the Levenberg–Marquardt algorithm, a common, relatively robust and deterministic numerical minimisation procedure, as the literature reported adequate results by using this algorithm. Like other iterative numerical optimisation algorithms, Levenberg–Marquardt requires an initial guess for parameter values, and typically the algorithm will converge to the local minimum closest to that initial guess. As each iteration is updated depending on the value of the Jacobian, highly varying functions (or those where the Jacobian is ill-conditioned, i.e. where the absolute ratio of its largest and smallest eigenvalues is large) may have poor convergence behaviour, despite the quadratic approximation for the function being minimised and addition of a dynamically updated damping term in the algorithm to minimise this effect. Like the majority of explicit iterative traditional methods for optimisation and curve-fitting, Levenberg–Marquardt is fast on modern computers, even for the highly multidimensional data

considered here: fitting the metabolic model to a 3D image with three metabolites takes approximately one minute on a machine with an Intel Core i7-3770, 3.40 GHz 8-core CPU with 16 GiB RAM running Linux/3.13.0-33 generic, using the `nlinfit` function in Matlab.

However, it was soon found that Levenberg–Marquardt did not routinely converge upon the global minimum in parameter space for a given dataset, and was very sensitive to the chosen starting parameter vector. It is proposed that, given the piecewise nature of Eqs. (9.29) and (9.30), this is unsurprising: many discontinuities may exist due to the choice of t_a and t_e, for example, which may serve to make the minimisation landscape highly glassy.

One successful approach to the problem of initial parameter estimation with iterative schemes such as Levenberg–Marquardt is to identify a parameter vector that ensures convergence to an adequate solution, and search for small perturbations to it for future experiments. It may also usually be possible to infer sensible values for some parameters *a priori*, such as obtaining estimates for t_e from the frame with the maximum signal. However, the former approach may be highly specific with respect to experimental conditions, and after estimating sensible prior values for each parameter there is no guarantee that Levenberg–Marquardt will converge to values close to them. For these reasons, it was decided to try and implement a numerical method that ideally could be guaranteed (for a realisable definition of 'guaranteed', such as 'in an arbitrary amount of time') to converge upon the global minimum for this multidimensional curve fitting problem. It was additionally noted that a Bayesian framework would enable the quantitative incorporation of prior belief about the values of parameters into the analysis. For this reason, the MCMC approach was used with the Metropolis-Hastings algorithm and adaptive Gibbs sampling, following the approach of Haario *et al.*[292] In such a fashion, quantitative bounds on parameter values could be returned for future incorporation into the method described in Section 9.2.

9.4.3.1 A very succinct overview of MCMC

Briefly, let $D(t_i)$ be a set of observed data points (which need not be scalar) at times t_i, and let a model $y(t_i|\theta)$ predict D at a given time, given some parameter vector θ. We wish to find θ such that the difference between the model and the data is small. Accordingly, the quantity

$$\chi^2 = \sum_i \frac{1}{2\sigma^2}\big(D(t_i) - y(t_i|\theta)\big)^2 \tag{9.31}$$

represents a least-squares metric for that fit, where σ is an estimate of the noise variance. The quantity $p(D|\theta) = \exp\left(-\chi^2\right)$ represents the likelihood, and is maximised when the error is minimised. One can then use Bayes's theorem[293] to express the relationship between the expected posterior belief in parameter values and the model likelihood of acquired data as (with common names for each term expressed above)

$$\overbrace{p(\theta|D)}^{\text{Posterior}} = \frac{\overbrace{p(\theta)}^{\text{Prior}} \times \overbrace{p(D|\theta)}^{\text{Likelihood}}}{\underbrace{\int p(\theta)p(D|\theta)\,\mathrm{d}\theta}_{\text{Evidence}}}, \tag{9.32}$$

where the 'evidence' essentially represents a normalising constant. However, directly finding the posterior distribution given in Eq. (9.32) typically requires high-dimensional integration which historically was not feasible. MCMC methods stochastically generate a Markov Chain whose stationary distribution is exactly that of the posterior, and sample from it. The distribution (and quantities of interest derived from it, such as its mean) can therefore be computed exactly. Whilst a detailed discussion of how this is achieved lies beyond the scope of this document, briefly, an initial guess for (all) the parameter values θ_{Current} is used to compute an error, χ^2_{Current}. Another set of parameter estimates, θ_{Proposed} is obtained, and the corresponding error term, χ^2_{Proposed} computed. The likelihood ratio $p(D|\theta_{\text{Proposed}})/p(D|\theta_{\text{Current}}) = \exp(-\chi^2_{\text{Proposed}} + \chi^2_{\text{Current}})$ is computed. A uniformly distributed random number $r \in [0, 1]$ is then drawn. If the likelihood ratio is bigger than r, the current parameter set is replaced with the proposed set. If not, the current set is stored. After a long period of time, the parameter set that maximised likelihood is often chosen, and the distribution can be estimated from the histogram of the values of θ_i and knowledge of the prior. The questions of interest are "how are the updated parameter values chosen?", and "how long does the chain take to converge?"; unfortunately, neither have straightforward answers. Here, the variant used to address the former is described in Haario *et al.*,[292] and for the canonical Gibbs sampling method the new θ are estimated from the old θ by drawing a Gaussian random variable for θ that is close to the previous value. Markov chains display a phenomenon known as 'burn in', where the 'initial' sampling period of the chain deviates substantially from its convergent value, and is typically discarded. Unfortunately, how long this 'initial' period lasts is difficult to compute, and analytically complex: for this reason, the chain is usually run for a substantial length of time to ensure that it has reached its stationary distribution. For the case of the Metropolis Hastings algorithm with Gibbs sampling, this convergence is analytically guaranteed to be at an exponential rate, though, of course, estimating this rate is highly non-trivial.[294]

In the work here, the 'Geweke' diagnostic was used to assess chain convergence, which simply compares the mean value of the posterior in the second 50% of the chain to an increasing number of initial elements, and assesses whether or not the period assessed is likely to have been 'burnt in'.[295] Approximately 1.5×10^4 update steps were required to converge with adequate posterior accuracy. Accordingly, the method took far longer to fit to the data than Levenberg–Marquardt; approximately five to fifteen minutes per time series, leading three hours for a 2D image as acquired in Chapter 5 or four days for a 3D dataset as acquired in Chapter 6. Note that it is possible to formulate MCMC as an 'embarrassingly parallel' computational problem,[296] and if the model being fitted is spatially invariant

this fact is especially obvious. To minimise execution time, voxels were analysed in parallel. To avoid wasting time by fitting noise, the result of a conservative statistical time series test, the Ljung-Box Q test,[297] was used to estimate if a given voxel in an image contained white noise or a time series that is highly likely to be different from white noise; Bayesian fitting was only performed in the latter case.

9.4.4 Example results

As an illustrative example, a representative image of [1-^{13}C]pyruvate to [1-^{13}C]lactate metabolism in the rat brain as acquired in Chapter 5 was analysed by the above method. Priors were chosen on the basis of the observed signal (for t_a and t_e), or, if no information could be easily gained from the acquired data (as in the case of rate constants), were set to be uniform on a defined interval, such as $[0, \infty]$. As shown in Fig. 9.7, the method performed well, was able to statistically discriminate between signal and noise, and returned an informative spatial map of rate constants of interest. The posterior distributions for parameter values were unimodal, and approximately normal, and the $\mathcal{O}(10^4)$ iterations were enough for the chain to converge.

As can be seen in Fig. 9.7**B**, the parameter R_i was approximately constant over the brain but higher nearer the MCA. This is believed to reflect the physiological perfusion of the brain, with pyruvate initially arriving in the MCA before subsequently entering the rest of the tissue. As signal amplitude is essentially arbitrary in this model, the net effect of the fitting is to remove the coil sensitivity profile that would otherwise be superimposed over the data acquired. The parameter $k_{\text{Pyr}\rightarrow\text{Lac}}$ represents the metabolism of [1-^{13}C]pyruvate into [1-^{13}C]lactate, and is approximately constant in the region sampled. This is consistent with the healthy rat brain, where it is assumed that metabolism is spatially homogeneous.

9.4.5 Discussion and Future Directions

It is clear that the use of Bayesian methods to quantitatively interpret the multidimensional datasets that DNP generates provides a robust alternative to more conceptually straightforward methods of analysis. As these methods can statistically discriminate between signal and noise, and the initial signal amplitude is essentially arbitrary in Eq. (9.29), it is possible to obtain more spatially homogeneous maps of rate constants of interest. Quantitative estimates of uncertainty on each parameter are additionally available for every point in space. The use of Bayesian methods in a broader context would allow for the establishment of normal ranges (i.e. normal priors) for parameters of interest, and a significant deviation from these ranges for a given subject could be considered indicative of disease. Compared to other global optimisation approaches, such as simulated annealing, the Bayesian approach allows for the calculation of more relevant information at the expense of computational efficiency.

The computational requirement for fitting a full 3D dataset is not insignificant. It is likely that the

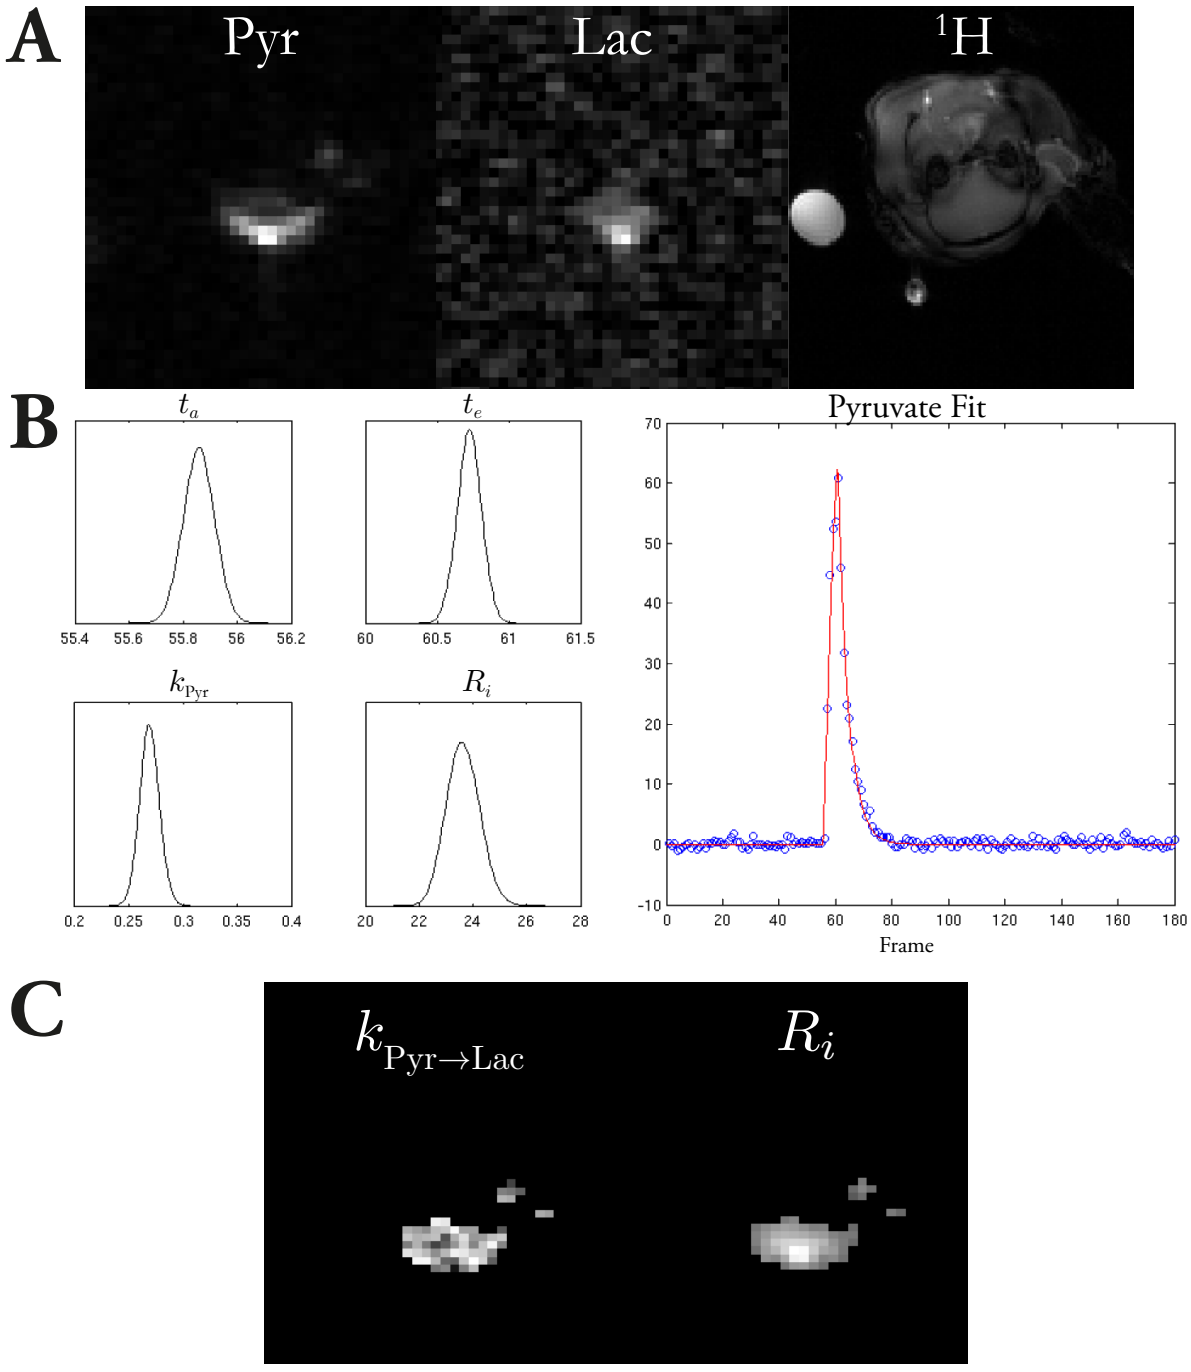

Figure 9.7: An example of the Bayesian MCMC metabolic modelling procedure. A: Here, the example 2D data originally displayed in Fig. 5.11 is used as an input, as the image SNR was deemed sufficient to warrant further analysis. Accordingly, the MCMC method described in Section 9.4.3.1 was used to fit the model Eq. (9.29) voxelwise to (initially) the [1-^{13}C]pyruvate data, and then the returned parameters were used to fit the lactate data (Eq. (9.30)), again, to effectively reduce the dimensionality of parameter space. **B**: MCMC fitting results were good. Shown here are the posterior distributions for four parameters of interest, t_a, t_e, R_i and $k_{\mathrm{Pyr}\rightarrow\mathrm{Lac}}$, along with the fit of the model (red) to the data (blue) in a representative voxel. **C**: Two example maps of rate constants are shown. Interestingly, the R_i parameter map is approximately homogeneous over the brain, but increases towards the MCA. Likewise, $k_{\mathrm{Pyr}\rightarrow\mathrm{Lac}}$ is approximately homogeneous in this healthy animal.

author's (Matlab) code for performing this fitting is highly inefficient: for example, multiple CPUs execute independent chains in a parallel `for` loop, but no attempt is made to use the highly parallel architecture available with modern high-end GPUs. Additionally, it is likely that there exist more sensible ways to fit the temporal model to the spatial data. At present, for computational ease each voxel is iterated over in turn, although in practice the greatest signal (and therefore the greatest certainty in results) is usually found near the middle of the image. An iterative sub-sampling scheme, in which the spatial resolution of the image is progressive increased from $1 \times 1 \times 1$ voxel, may be able to provide initial parameter values that are broadly applicable for the dataset as a whole. This practice may bring a significant increase in speed, as the duration of the burn-in period and the length of time required for the Markov Chain to converge is believed to be related to how close the initial values are to the stationary distribution. It is highly anticipated that future work could therefore make the analysis presented here far faster to compute.

Undoubtedly the biggest criticism of the approach presented here is that the model fitted to the data is spatially invariant, but that the structures present within living organisms are not. This process has been performed partly for numerical convenience: Eqs. (9.29) and (9.30) represent piecewise (analytically) integrated forms of first-order chemical kinetics, and are computationally straightforward to evaluate. As the metric Eq. (9.31) has to be computed frequently ($\mathcal{O}\left(10^4 \times 64 \times 32 \times 24 \times 3\right) \sim 10^9$ times in the work here), minimising the time taken for the execution of $y(t_i|\theta)$ is highly desirable. Spatially varying models are inevitably more complex, and are believed to represent an increase of several orders of magnitude in computational complexity. However, it may be possible to use existing anatomical $^{1}\mathrm{H}$ data to define a high-resolution mesh of the anatomical structures present. This mesh could then form a basis for a complete, natively 3D finite-element representation of the expected observed signal over time; for example, with [1-^{13}C]pyruvate expected in the ventricles, and ^{13}C-bicarbonate strictly localised to the myocardium. In order to compute Eq. (9.31), the mesh could be sub-sampled to the acquisition volume and concentrations integrated. As most differential equations can be solved with essentially arbitrary accuracy,[7] this type of mixed approach has been previously shown to provide a powerful way to go beyond the fundamentally limited spatial resolution of MR, and is able to predict the growth profile and region at risk of a tumour, for example.[298] Unfortunately, the computational requirement for solving linear PDEs in 3D ensures that calculating Eq. (9.31) is not trivial, and accordingly parameter fitting approaches are difficult.

It is proposed that future work on this problem uses sets of PDEs of the taxis-diffusion-reaction type, as it is believed that these would form an appropriate spatial extension to Eqs. (9.29) and (9.30). Such equations adequately describe chemical species that diffuse and react as a function of time, and have previously been used (by the author and many others) to model cell densities undergoing dynamic

[7] Though, of course, they may be arbitrarily sensitive to their initial conditions, and to small numerical perturbations to those conditions.

biological processes.[299,300] In particular, the heterogeneity present in the heart means that neglecting spatial information is likely to be particularly deleterious; for example, no (endogenous) metabolism is believed to occur in the lumen of the ventricles, whereas the pyruvate signal may be low (and influenced by partial volume effects) in the myocardium itself. Specifically, it is proposed that a general diffusive-perfusive reaction equation of the form

$$\begin{aligned}\frac{\partial c_{\text{Pyr}}(\mathbf{x},t)}{\partial t} =& -\nabla\cdot(-D_{\text{Pyr}}(\mathbf{x})\nabla c_{\text{Pyr}}) + \text{Perfusion}(c_{\text{Pyr}},\mathbf{x})(c_{\text{Pyr}}\cdot\nabla)c_{\text{Pyr}} \\ &+\text{Input}(t) - \text{Metabolism}(c_j(\mathbf{x}),c_{\text{Pyr}},t) - \frac{c_{\text{Pyr}}}{T_{1\,\text{Pyr}}} \qquad (9.33)\\ \frac{\partial c_{j\neq\text{Pyr}}(\mathbf{x},t)}{\partial t} =& -\nabla\cdot(-D_j(\mathbf{x})\nabla c_j) + k_{\text{Pyr}\to j}(\mathbf{x})\cdot c_{\text{Pyr}} - \frac{c_j}{T_{1\,j\neq\text{Pyr}}} \qquad (9.34)\end{aligned}$$

could be solved on a domain estimated from the anatomical proton images. The form of the (spatially invariant) input function could be estimated by deconvolution between tissue and lumen signals, for example, as proposed by Jerosch-Herold.[301] Regularisation parameters may be required to ensure that the rate constants of interest do not vary spatially at a higher rate than that supported by the data. Preliminary work shows that the Method of Lines is able to solve PDEs of this type on a $\sim 30\times 30$ patch rectangular grid with 'reasonable' efficiency, i.e. in less than a minute with appropriate compiled `mex` files to solve each ODE generated after discretising the grid. Further work is clearly required, but in an oncology setting this type of approach has been shown to be highly profitable, and estimate rate constants of interest (for cell growth) that allow for the prediction of the region of future disease.[302–304] It is hoped that the future development of a spatially varying model for pyruvate kinetics would simultaneously allow the estimation of spatial maps of rate constants with higher resolution that those which can be obtained directly, and, with the Bayesian methods discussed above, be able to return quantitative estimates on those rates. However, such a task represents both a significant computational challenge and an area for substantial research.

Chapter 10

Conclusions, Limitations, and Future Work

Every widely used technology has undergone numerous and distinct stages of development. Such stages differ hugely in scope and challenge, and are often categorised into (somewhat awkwardly phrased) 'technology readiness levels'. The last half century has seen a significant increase in the 'readiness level' of DNP. The technique has progressed from being studied in a low temperature setting, before being used to generate spin-polarised targets for use in $p^+ \rightarrow X$ collider experiments and (via pumped $^3\mathrm{He}$) polarised neutrons for diffraction experiments. Barely more than a decade after the invention of the dissolution technique by Ardenkjaer-Larsen *et al.*,[42] the use of DNP to generate novel probes of diseases relevant to mammalian health has been extensively shown in many different species, recently including humans.[102]

Overview

This work has detailed the development of technical methods to non-invasively image the *in vivo* behaviour of hyperpolarised compounds, namely [1-^{13}C]pyruvate and ^{13}C-urea. Unlike existing medical imaging modalities (namely PET and SPECT), hyperpolarised [1-^{13}C]pyruvate imaging returns functional information about tissue in a large volume within minutes (rather than hours) following the injection of a safe (rather than radioactive) tracer. As described in Chapters 4 to 6, imaging the metabolism of hyperpolarised [1-^{13}C]pyruvate is desirable, and technically challenging, requiring the construction of additional hardware and a deep working knowledge of the spectrometer as detailed in Chapter 3. The resolution ultimately achieved (by the 3D spectral-spatial sequence described in Chapter 6) is comparable to the theoretical limiting resolution of PET, and is in excess of that reported in the literature. Hyperpolarised [1-^{13}C]pyruvate has been shown extensively by the previous work of others to be a sensitive probe of disease, and this fact is briefly confirmed here with proof-of-principle demonstrations of the technique in myocardial ischæmia (in Chapter 6) and cancer (Chapter 7).

Additionally, the use of hyperpolarised ^{13}C-urea was investigated as a novel probe of perfusion in Chapter 8, and potentially offers several advantages when compared to existing methods. As well as being expected to have a much greater patient tolerability than gadolinium, the fact that the observed hyperpolarised ^{13}C-urea signal has no background and is directly proportional to the concentration of

the agent present is expected to make quantification more robust. Early-stage work using co-polarised [1-^{13}C]pyruvate and ^{13}C-urea shows that it is feasible to obtain functional information about perfusion and metabolism at the same time, which is expected to be of great utility for identifying viable and non-viable cardiac tissue in the heart post infarction.

Datasets generated by DNP are highly variable, and often highly multidimensional. The mathematical methods presented in Chapter 9 provide a quantitative framework for robustly incorporating data with varying estimates of uncertainty into a single study. Additionally, the Fast Padé Transform is proposed as a novel way of reducing the time required for the acquisition of hyperpolarised data, by simply using an interpolating function (albeit one with a rich algebraic structure) to infer missing data. It is hoped to generalise such methods to more than one dimension. Lastly, the robust and quantitative fitting of metabolic models via MCMC is proposed as a way to reduce the dimensionality of the data DNP provides. Such models can provide bounded estimates of rate constants of interest, and may provide an interesting avenue for the future 'physics based modelling' of the data provided by hyperpolarised experiments.

Shortcomings & Limitations

It is believed that everything presented in this work, with the exception of the introductory Chapter 2, is novel, for an appropriate (i.e. relatively narrow) definition of 'novel'. However, the use of spectral-spatial imaging sequences as discussed in Chapters 5 and 6 is not unprecedented, and much previous work has been done using larger animals (with Siemens clinical 3 T MR systems) by the group of Chuck Cunningham and others.[96,138,305] The technical implementation of similar sequences on the Varian 7 T system here represents a significant challenge, in part owing to the greater difficulty of shimming at 7 T, and the corresponding decrease in the T_2^* of [1-^{13}C]pyruvate from ~ 30 ms at 3 T (personal communication, Angus Lau) to the ~ 15 ms considered here. The desire to image rodents additionally necessitates a substantial increase in technical difficulty: compared to the $\sim 100 \times 100 \times 100$ mm^3 size of the human heart, the $\sim 10 \times 10 \times 10$ mm^3 size of the rat heart requires an effective 1000-fold increase in the achievable resolution in order to image anatomical structures with equivalent finesse. Additionally, the rapid heart rate of the rat compared to larger animals greatly reduces the temporal window for imaging, which further increases the demands on the readout of any MR sequence. The use of a discontinued pre-clinical scanner as opposed to a commercially available clinical machine may have increased the development time for the work described here, but it arguably did equip the author with great familiarity of the electronic engineering of real-time devices. It may be argued that tackling the challenging problem posed by small-animal cardiac imaging is of utility in a wider context, as the ultimate resolution achievable by DNP is essentially determined by the hyperpolariser used. Additionally, the use of small animal models of disease in scientific research ensures that such

approaches will find applications in broader studies of diseases of metabolic dysregulation. The method described in Chapter 6 is expected to scale well to image humans, and it is hoped that the techniques developed to image hyperpolarised substrates *in vivo* (and in particular [1-^{13}C]pyruvate metabolism) will find broad application.

The attempts at demonstrating broad biological utility in this work were significantly hampered by circumstances outside of the author's control: a study which attempted to investigate the response of the heart to infarction with hyperpolarised [1-^{13}C]pyruvate could not run as designed due to the unexpected quenching of the MR system; and the oncology study detailed in Chapter 7 essentially failed to reach a firm conclusion due to lack of cancer in the rats scanned. Both studies will be repeated, and will additionally be performed with the use of a four-channel surface receive array that is currently being constructed for use with the preclinical MR system. It is hoped that other hardware upgrades will allow for the higher-resolution imaging of metabolism with the extant work. Additionally, the ongoing refinement of the animal models of disease used are currently being investigated by others, and hopefully will reduce the variability inherent to such experiments. The completed studies did, however, demonstrate the feasibility of the technique *in vivo*, and proved its ability to function robustly and routinely enough for (pre-)clinical practice. Promising, if small-scale, results were found with the use of hyperpolarised [1-^{13}C]pyruvate imaging in disease, be they the presence of [1-^{13}C]lactate in small-scale cancers or a visible lack of ^{13}C-bicarbonate following infarction. It is hoped that these can be generalised to show the functional utility of the technique in detecting (and prognosticating) disease.

Likewise, the proof-of-principle demonstration of the use of hyperpolarised ^{13}C-urea as a probe of perfusion is small-scale, and suffers from some significant limitations. While the use of flow suppression gradients enables the visualisation of urea arising from within the tissue bed, it is necessarily magnetisation destructive, and only spoils flow along one given direction within the organism. If this direction is chosen incorrectly, it may be impossible to interpret the resulting data. There is undoubtedly some degree of user-dependent variation in selecting the direction of flow nulling, and, if the long axis of the heart is not approximately linear, this variability may be significant. Moreover, the potential benefits of the technique over the current clinical standard of gadolinium have not been fully investigated: in part, this was due to the difficulty of performing first-pass perfusion imaging with gadolinium in the rodent heart using the hardware available to the author. Future work would aim to demonstrate the fundamental benefits of hyperpolarised ^{13}C-urea imaging over those provided by more traditional methods. It may also be feasible to iterate over directions of flow suppression and hence build an accelerated ^{13}C-urea velocity mapping sequence to return further information about the vasculature. Work with co-polarised [1-^{13}C]pyruvate and ^{13}C-urea is ongoing, and promises to provide a method to quantify both perfusion and metabolism in a single scan. It is hypothesised that, by imaging ^{13}C-urea, ^{13}C-bicarbonate and [1-^{13}C]lactate rather than [1-^{13}C]pyruvate, ^{13}C-bicarbonate

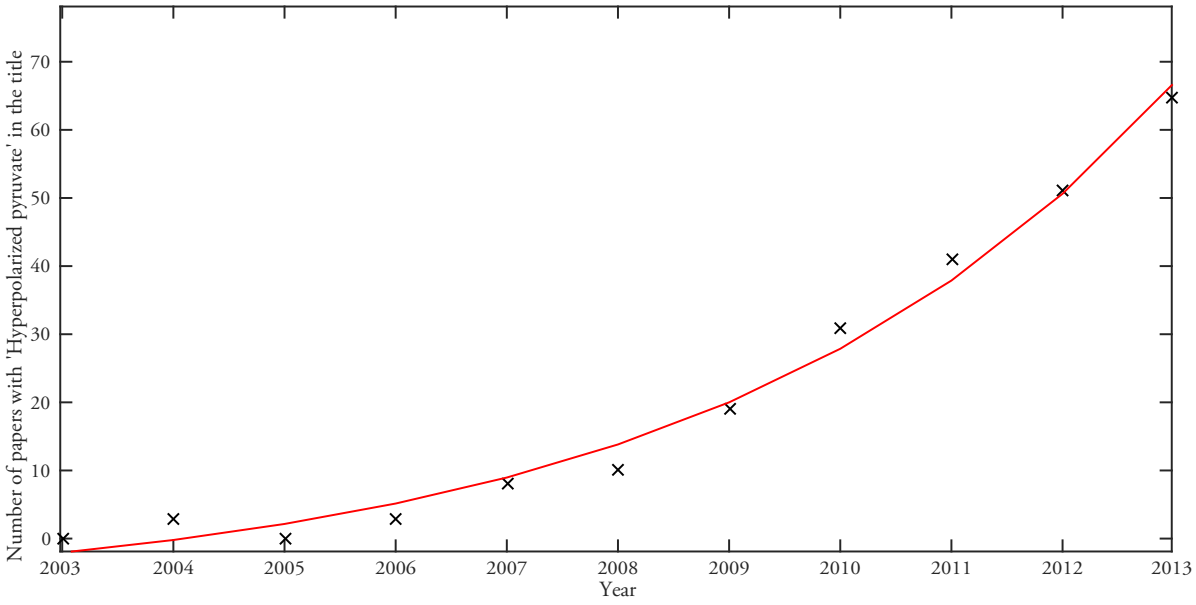

Figure 10.1: The number of publications listed in pubmed with 'Hyperpolarized pyruvate' in the title, shown as a function of time. A simple exponential fit is shown in red.

and [1-^{13}C]lactate with the sequence described in Chapter 6, the magnetisation reservoir present in hyperpolarised [1-^{13}C]pyruvate *in vivo* could be maintained for longer. Such a scan would not sacrifice information about perfusion or the degree of polarisation present, and correspondingly higher-SNR images could possibly be obtained of the downstream products of primary interest (i.e. ^{13}C-bicarbonate and [1-^{13}C]lactate) for a longer period of time.

The mathematical methods presented in Chapter 9 are not independently validated, and represent a significant area for future research. In particular, the intelligent quantification of hyperpolarised images is an open question, and the common approaches taken – either division or metabolic model fitting – do not easily generalise to heterogeneous tissues, such as the heart, where the detectable presence of all three metabolites in a given voxel is not guaranteed. The forms of equations proposed to overcome these limitations are very general, and no attempt has been made to integrate and investigate them on a physiological mesh. Likewise, the statistical fitting methods explored are very powerful, but, as implemented, are comparatively computationally expensive. As high performance, highly-parallel processors are readily available (in the form of CUDA-compatible GPUs or, ultimately, FPGAs), and several methods have been proposed to effectively parallelise MCMC algorithms, it is likely that these computational costs could be reduced.[1] Finally, the use of approximant methods in reconstructing hyperpolarised data that has been prospectively accelerated represents an obvious expansion to Section 9.3, and would validate the technique. Such developments are ongoing.

Future Directions

As illustrated in Fig. 10.1, the field of *in vivo* hyperpolarised Magnetic Resonance research is growing, approximately exponentially. Significant financial investment by many different organisations has provided a number of world-leading organisations with the ability to image the metabolism of hyperpolarised [1-^{13}C]pyruvate *in vivo*, in *Homo sapiens*. The translation of imaging sequences of the type developed here therefore undoubtedly represents the ultimate goal for hyperpolarised research, although it should be noted that the utility of the technique in a pre-clinical setting will not be diminished by its successful clinical translation. In particular, only a small subset of the nodes in the metabolic network have been investigated with DNP, and it is likely that a number of the remaining nodes will be of biological interest. In a translational setting, much work remains to be done in identifying the patient populations where hyperpolarised [1-^{13}C]pyruvate imaging is of greatest utility, and in the organisation and execution of studies to demonstrate that the technique brings clinical benefit. For example, at present, patients presenting with symptoms consistent with acute myocardial infarction typically have the condition diagnosed, and its severity classified, by the use of ECG and blood bio-markers. While revascularisation is typically performed for patients with evidence of substantial injury,[2] the need for urgent revascularisation is less obvious for patients with less severe disease, such as those without significant ST elevation; or in those who have occult injury, such as moderate occlusion of all three coronary arteries. A metabolic imaging exam may be useful to guide the revascularisation decision, and the use of simultaneous metabolic-perfusion agents holds the promise of rapidly quantifying regions of tissue viability and perfusion deficit. Moreover, as the T_1 and T_2 of metabolites of interest are functions of B_0, it is likely that the lower B_0 commonly encountered in clinical imaging situations may actually represent a further increase in the effective available magnetisation for future metabolic imaging studies. Such aspects of the technique will continue to be explored as it moves towards clinical practice.

Additionally, the sequences developed here could be improved in several different ways. It is well known that [1-^{13}C]pyruvate and ^{13}C-urea possesses a long T_2, yet the gradient echo type rapid imaging sequences developed in Chapters 5, 6 and 8 do not take advantage of this, and are limited by the achievable T_2^*. For a heart, filled with moving paramagnetic blood, this limitation is significant. Spin echo type sequences are difficult to implement successfully with hyperpolarised substrates, as an errant high-θ pulse would essentially destroy the remaining magnetisation present, and prematurely end an experiment. For this reason, exceptional B_1^+ homogeneity is required, and attempts in the literature to perform spin echo hyperpolarised experiments have often used adiabatic 180° pulses, with the experiment started after the end of the injection of the bolus of the hyperpolarised compound. It may be the case that a spectral-spatial refocussing pulse could adequately act as a spin echo pulse

[1] It is also highly likely that the author's code is not the most efficient possible.

[2] In the UK, immediate revascularisation is indicated if a > 0.2 mV increase in the height of the ST segment of the ECG above its baseline is observed with a concomitant rise in blood troponin levels.

following spectral-spatial excitation of a given slab, but again the moving nature of blood may mean that such a technique could only be applied in regions without flow. As it has been noted that the T_2 of [^{13}C,U-^{15}N]-urea is significantly longer than that of ^{13}C-urea, this may represent a significant reduction in the utility of the technique for specifically cardiac imaging. It may be the case that short, intense adiabatic refocussing pulses, delivered by a dedicated coil, would provide the optimum way to deliver a 180° refocussing pulse to everything within a defined volume of interest. By extending the duration available for the readout "for free" in this fashion, significantly higher resolution images may be obtained, at the marginal expense of greater RF deposition into the subject.

Final words

This thesis has focussed on the technical development of pulse sequences and hardware that enable non-invasive metabolic imaging in living organisms. As metabolic dysregulation is implicated in many human diseases, it is hoped that the high-resolution, efficient methods developed facilitate both the direct imaging of diseased tissue, and find wider utility in broad pre-clinical studies aimed at elucidating exactly the mechanism occurring in a specific disease process. The potential utility of these techniques has been illustrated in the post infarcted heart, and in a model of cancer metastasis to the brain. Methods for adequately quantifying the data generated by hyperpolarised [1-^{13}C]pyruvate spectroscopy or imaging are presented, and are shown to offer utility over simpler, more naïve approaches. It is hoped that the work presented here ultimately finds broad utility in improving the management of patients with heart disease, cancer, or other diseases of metabolic dysregulation.

Bibliography

[1] Hanahan, D. and Weinberg, R. A. "Hallmarks of cancer: the next generation." *Cell*, **144** (2011) (5); pp. 646–74. ISSN 1097-4172. DOI 10.1016/j.cell.2011.02.013.

[2] Neubauer, S. "The failing heart – an engine out of fuel". *N. Engl. J. Med.*, **356** (2007); pp. 1140–51. DOI 10.1056/nejmra063052.

[3] Binney, J. and Skinner, D. *The Physics of Quantum Mechanics*. Oxford University Press, Oxford, oup edn. (2013). ISBN 0199688575.

[4] Aoyama, T., Hayakawa, M., Kinoshita, T., and Nio, M. "Tenth-Order QED Contribution to the Electron g-2 and an Improved Value of the Fine Structure Constant". *Phys. Rev. Lett.*, **109** (2012) (11); p. 111 807. ISSN 0031-9007. DOI 10.1103/PhysRevLett.109.111807.

[5] Beane, S., Chang, E., Cohen, S., Detmold, W., Lin, H., Orginos, K., Parreño, A., Savage, M., and Tiburzi, B. "Magnetic Moments of Light Nuclei from Lattice Quantum Chromodynamics". *Phys. Rev. Lett.*, **113** (2014) (25); p. 252 001. ISSN 0031-9007. DOI 10.1103/PhysRevLett.113.252001.

[6] Mabuchi, H. "Derivation of Maxwell-Bloch-type equations by projection of quantum models". *Phys. Rev. A*, **78** (2008) (1); p. 015 801. ISSN 1050-2947. DOI 10.1103/PhysRevA.78.015801.

[7] Bernstein, M. A., King, K. F., and Zhou, X. J. *Handbook of MRI Pulse Sequences*. Elsevier (2004). ISBN 0080533124.

[8] Hoult, D. I. "The principle of reciprocity in signal strength calculations: A mathematical guide". *Concepts Magn. Reson.*, **12** (2000) (4); pp. 173–187. ISSN 1043-7347. DOI 10.1002/1099-0534(2000)12:4<173::AID-CMR1>3.0.CO;2-Q.

[9] Van de Moortele, P.-F. "RF Coils & B1 Mapping". In "Proc. Intl. Soc. Mag. Reson. Med. 23", (2015).

[10] Heider, E. M. *Exploration of Structure and Function in Biomolecules Through Solid-state NMR and Computational Methods*. ProQuest (2007). ISBN 0549833048.

[11] Bloch, F. "Nuclear Induction". *Phys. Rev.*, **70** (1946) (7-8); pp. 460–474. ISSN 0031-899X. DOI 10.1103/PhysRev.70.460.

[12] Juchem, C., Umesh Rudrapatna, S., Nixon, T. W., and de Graaf, R. A. "Dynamic multi-coil technique (DYNAMITE) shimming for echo-planar imaging of the human brain at 7 Tesla." *Neuroimage*, **105** (2015); pp. 462–72. ISSN 1095-9572. DOI 10.1016/j.neuroimage.2014.11.011.

[13] Juchem, C., Herman, P., Sanganahalli, B. G., Brown, P. B., McIntyre, S., Nixon, T. W., Green, D., Hyder, F., and de Graaf, R. A. "DYNAmic Multi-coIl TEchnique (DYNAMITE) shimming of the rat brain at 11.7 Tesla". *NMR Biomed.*, **27** (2014) (8); pp. 897–906. ISSN 1099-1492. DOI 10.1002/nbm.3133.

[14] Juchem, C., Green, D., and de Graaf, R. A. "Multi-coil magnetic field modeling." *J. Magn. Reson.*, **236** (2013); pp. 95–104. ISSN 1096-0856. DOI 10.1016/j.jmr.2013.08.015.

[15] Damadian, R., Minkoff, L., Goldsmith, M., Stanford, M., and Koutcher, J. "Field focusing nuclear magnetic resonance (FONAR): visualization of a tumor in a live animal." *Science*, **194** (1976) (4272); pp. 1430–2. ISSN 0036-8075.

[16] Pauly, J., Nishimura, D., and Macovski, A. "A k-space analysis of small-tip-angle excitation". *J. Magn. Reson.*, **81** (1989) (1); pp. 43–56. ISSN 00222364. DOI 10.1016/0022-2364(89)90265-5.

[17] Shinnar, M., Eleff, S., Subramanian, H., and Leigh, J. S. "The synthesis of pulse sequences yielding arbitrary magnetization vectors." *Magn. Reson. Med.*, **12** (1989) (1); pp. 74–80. ISSN 0740-3194.

[18] Le Roux, P. "Simplified RF Synthesis". In "8th Annu. Meet. SMRM", (1989).

[19] Le Roux, P. "Exact synthesis of radio frequency waveforms". In "7th Annu. Meet. SMRM", p. 1049 (1988).

[20] Pauly, J., Le Roux, P., Nishimura, D., and Macovski, A. "Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm [NMR

N.B. The majority of citaitons are hyperlinked to online versions of their contents in the digital version of this document.

imaging]." *IEEE Trans. Med. Imaging*, **10** (1991) (1); pp. 53–65. ISSN 0278-0062. DOI 10.1109/42.75611.

[21] Jezzard, P. and Balaban, R. S. "Correction for geometric distortion in echo planar images from B0 field variations". *Magn. Reson. Med.*, **34** (1995) (1); pp. 65–73. ISSN 07403194. DOI 10.1002/mrm.1910340111.

[22] Andersson, J. L. R., Skare, S., and Ashburner, J. "How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging." *Neuroimage*, **20** (2003) (2); pp. 870–88. ISSN 1053-8119. DOI 10.1016/S1053-8119(03)00336-7.

[23] Cunningham, C. H., Dominguez Viqueira, W., Hurd, R. E., and Chen, A. P. "Frequency correction method for improved spatial correlation of hyperpolarized 13C metabolites and anatomy." *NMR Biomed.*, **27** (2014) (2); pp. 212–8. ISSN 1099-1492. DOI 10.1002/nbm.3055.

[24] Chen, N.-K., Avram, A. V., and Song, A. W. "Two-dimensional phase cycled reconstruction for inherent correction of echo-planar imaging Nyquist artifacts." *Magn. Reson. Med.*, **66** (2011) (4); pp. 1057–66. ISSN 1522-2594. DOI 10.1002/mrm.22896.

[25] Abragam, A. and Goldman, M. "Principles of dynamic nuclear polarisation". *Reports Prog. Phys.*, **41** (1978); p. 395. DOI doi:10.1088/0034-4885/41/3/002.

[26] Comment, A., Van Den Brandt, B., Uffmann, K., Kurdzesau, F., Jannin, S., Konter, J. A., Hautle, P., Wenckebach, W. T., Gruetter, R., and Van Der Klink, J. J. "Design and Performance of a DNP Prepolarizer Coupled to a Rodent MRI Scanner". *Concepts Magn. Reson.*, **31B** (2007) (1); pp. 255–269. ISSN 15525031. DOI 10.1002/cmr.b.

[27] Borghini, M. "Spin-Temperature Model of Nuclear Dynamic Polarization Using Free Radicals". *Phys. Rev. Lett.*, **20** (1968) (9); pp. 419–421. ISSN 0031-9007. DOI 10.1103/PhysRevLett.20.419.

[28] Kozhushner, M. A. "Polarization of nuclei by dynamic cooling". *J. Sov. Phys.*, **29** (1969) (1); p. 136.

[29] Saunders, M. *Enhancing the sensitivity of NMR by Dynamic Nuclear Polarisation*. Ph.D. thesis, University of Birmingham (2010).

[30] Serra, S. C., Rosso, A., and Tedoldi, F. "Electron and nuclear spin dynamics in the thermal mixing model of dynamic nuclear polarization." *Phys. Chem. Chem. Phys.*, **14** (2012) (38); pp. 13 299–308. ISSN 1463-9084. DOI 10.1039/c2cp41947e.

[31] Serra, S. C., Rosso, A., and Tedoldi, F. "On the role of electron-nucleus contact and microwave saturation in thermal mixing DNP." *Phys. Chem. Chem. Phys.*, **15** (2013) (21); pp. 8416–28. ISSN 1463-9084. DOI 10.1039/c3cp44667k.

[32] Freeman, R. and Morris, G. A. "The Varian story". *J. Magn. Reson.*, **250** (2014); pp. 80–84. ISSN 10907807. DOI 10.1016/j.jmr.2014.12.001.

[33] Blümler, P., Blümich, B., Botto, R., and Fukushima, E., editors. *Spatially resolved mangetic resonance methods, materials, medicine, biology, rheology, geology, ecology and hardware*. John Wiley & Sons, London, 1st edn. (1998).

[34] Spielman, D. M. and Pauly, J. M. "Spiral imaging on a small-bore system at 4.7 T". *Magn. Reson. Med.*, **34** (1995) (4); pp. 580–585. ISSN 07403194. DOI 10.1002/mrm.1910340414.

[35] Duyn, J. H., Yang, Y., Frank, J. A., and van der Veen, J. W. "Simple correction method for k-space trajectory deviations in MRI." *J. Magn. Reson.*, **132** (1998) (1); pp. 150–3. ISSN 1090-7807. DOI 10.1006/jmre.1998.1396.

[36] Novikov, A. "Advanced theory of driven birdcage resonator with losses for biomedical magnetic resonance imaging and spectroscopy." *Magn. Reson. Imaging*, **29** (2011) (2); pp. 260–71. ISSN 1873-5894. DOI 10.1016/j.mri.2010.08.001.

[37] Larkman, D. J. and Nunes, R. G. "Parallel magnetic resonance imaging." *Phys. Med. Biol.*, **52** (2007) (7); pp. R15–55. ISSN 0031-9155. DOI 10.1088/0031-9155/52/7/R01.

[38] Schneider, J. E., Barnes, H., Neubauer, S., and Jezzard, P. "Automated-Shim Approach to Facilitate 1H-MRS in Mouse Hearts In Vivo". In "Proc. Intl. Soc. Mag. Reson. Med. 17", p. 1782 (2009).

[39] Lumata, L., Kovacs, Z., Sherry, A. D., Malloy, C., Hill, S., van Tol, J., Yu, L., Song, L., and Merritt, M. E. "Electron spin resonance studies of trityl OX063 at a concentration optimal for DNP." *Phys. Chem. Chem. Phys.*, **15** (2013) (24); pp. 9800–7. ISSN 1463-9084. DOI 10.1039/c3cp50186h.

[40] Rogozhnikova, O. Y., Vasiliev, V. G., Troitskaya, T. I., Trukhin, D. V., Mikhalina, T. V., Halpern, H. J., and Tormyshev, V. M. "Generation of Trityl Radicals by Nucleophilic Quenching of Tris(2,3,5,6-tetrathiaaryl)methyl Cations and Practical and Convenient Large-Scale Synthesis of Persistent Tris(4-carboxy-2,3,5,6-tetrathiaaryl)methyl Radical". *European J. Org. Chem.*, **2013** (2013) (16); pp. 3347–3355. ISSN 1434193X. DOI 10.1002/ejoc.201300176.

[41] Armstrong, B. D., Edwards, D. T., Wylde, R. J., Walker, S. A., and Han, S. "A 200 GHz dynamic nuclear polarization spectrometer." *Phys. Chem. Chem. Phys.*, **12** (2010) (22); pp. 5920–6. ISSN 1463-9084. DOI 10.1039/c002290j.

[42] Ardenkjaer-Larsen, J. H., Fridlund, B., Gram, A., Hansson, G., Hansson, L., Lerche, M. H., Servin, R., Thaning, M., and Golman, K. "Increase in signal-to-noise ratio of >10,000 times in liquid-state NMR." *Proc. Natl. Acad. Sci. U. S. A.*, **100** (2003) (18); pp. 10 158–63. ISSN 0027-8424. DOI 10.1073/pnas.1733835100.

[43] Wenckebach, W., Swanenburg, T., and Poulis, N. "Nuclear dynamic polarization of protons in diluted copper-tutton salts. II". *Physica*, **50** (1970) (2); pp. 289–307.

[44] Wenckebach, W., Swanenburg, T., and Poulis, N. "Nuclear dynamic polarization of protons in diluted Copper-Tutton salts. I". *Physica*, **46** (1970) (2); pp. 303–314. ISSN 00318914. DOI 10.1016/0031-8914(70)90230-2.

[45] Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N., and Barabási, A. L. "The large-scale organization of metabolic networks." *Nature*, **407** (2000) (6804); pp. 651–4. ISSN 0028-0836. DOI 10.1038/35036627.

[46] Dorogovtsev, S. and Mendes, J. "Evolution of networks with aging of sites". *Phys. Rev. E*, **62** (2000) (2); pp. 1842–1845. ISSN 1063-651X. DOI 10.1103/PhysRevE.62.1842.

[47] Amaral, L. A., Scala, A., Barthelemy, M., and Stanley, H. E. "Classes of small-world networks." *Proc. Natl. Acad. Sci. U. S. A.*, **97** (2000) (21); pp. 11 149–52. ISSN 0027-8424. DOI 10.1073/pnas.200327197.

[48] Alberts, B., editor. *Molecular Biology of the Cell: Reference edition, Volume 1*. Garland Science, 5th edn. (2008). ISBN 0815341113.

[49] Dennis, S. C., Kohn, M. C., Anderson, G. J., and Garfinkel, D. "Kinetic analysis of monocarboxylate uptake into perfused rat hearts." *J. Mol. Cell. Cardiol.*, **17** (1985) (10); pp. 987–95. ISSN 0022-2828.

[50] Poole, R. C. and Halestrap, A. P. "Transport of lactate and other monocarboxylates across mammalian plasma membranes." *Am. J. Physiol.*, **264** (1993) (4 Pt 1); pp. C761–82. ISSN 0002-9513.

[51] Mallet, R. T. "Pyruvate: metabolic protector of cardiac performance." *Proc. Soc. Exp. Biol. Med.*, **223** (2000) (2); pp. 136–48. ISSN 0037-9727.

[52] Manyam, G., Birerdinc, A., and Baranova, A. "KPP: KEGG Pathway Painter." *BMC Syst. Biol.*, **9 Suppl 2** (2015) (Suppl 2); p. S3. ISSN 1752-0509. DOI 10.1186/1752-0509-9-S2-S3.

[53] Dorland. *Dorland's Illustrated Medical Dictionary*. Elsevier Health Sciences (2011). ISBN 1455709859.

[54] Braunwald, E., Bonow, R. O., Mann, D. L., Zipes, D. P., and Libby, P. *Braunwald's Heart Disease: A Textbook of Cardiovascular Medicine*. Elsevier Health Sciences, Philadelphia (2011). ISBN 1437727700.

[55] McMurray, J. J. V. and Pfeffer, M. A. "Heart failure." *Lancet*, **365** (2005) (9474); pp. 1877–89. ISSN 1474-547X. DOI 10.1016/S0140-6736(05)66621-4.

[56] Doenst, T., Nguyen, T. D., and Abel, E. D. "Cardiac metabolism in heart failure: implications beyond ATP production." *Circ. Res.*, **113** (2013) (6); pp. 709–24. ISSN 1524-4571. DOI 10.1161/CIRCRESAHA.113.300376.

[57] Goldhaber, J. I. and Hamilton, M. A. "Role of inotropic agents in the treatment of heart failure." *Circulation*, **121** (2010) (14); pp. 1655–60. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.109.899294.

[58] CIBIS Investigators and Committees, N. Y. "A randomized trial of beta-blockade in heart failure. The Cardiac Insufficiency Bisoprolol Study (CIBIS). CIBIS Investigators and Committees." *Circulation*, **90** (1994) (4); pp. 1765–73. ISSN 0009-7322.

[59] Packer, M., Bristow, M. R., Cohn, J. N., Colucci, W. S., Fowler, M. B., Gilbert, E. M., and Shusterman, N. H. "The Effect of Carvedilol on Morbidity and Mortality in Patients with Chronic Heart Failure". *N. Engl. J. Med.*, **334** (1996) (21); pp. 1349–1355. DOI 10.1056/NEJM199605233342101.

[60] Consensus Trial Study Group, T. "Effects of enalapril on mortality in severe congestive heart failure. Results of the Cooperative North Scandinavian Enalapril Survival Study (CONSENSUS). The CONSENSUS Trial Study Group." *N. Engl. J. Med.*, **316** (1987) (23); pp. 1429–35. ISSN 0028-4793. DOI 10.1056/NEJM198706043162301.

[61] Pfeffer, M. A., Braunwald, E., Moyé, L. A., Basta, L., Brown, E. J., Cuddy, T. E., Davis, B. R., Geltman, E. M., Goldman, S., and Flaker, G. C. "Effect of captopril on mortality and morbidity in patients with left ventricular dysfunction after myocardial infarction. Results of the survival and ventricular enlargement trial. The SAVE Investigators." *N. Engl. J. Med.*, **327** (1992) (10); pp. 669–77. ISSN 0028-4793. DOI 10.1056/NEJM199209033271001.

[62] Cohn, J. N. and Tognoni, G. "A Randomized Trial of the Angiotensin-Receptor Blocker Valsartan in Chronic Heart Failure". *N. Engl. J. Med.*, **345** (2001) (23); pp. 1667–1675. DOI 10.1056/NEJMoa010713.

[63] Pfeffer, M. A., Swedberg, K., Granger, C. B., Held, P., McMurray, J. J. V., Michelson, E. L., Olofsson, B., Ostergren, J., Yusuf, S., and Pocock, S. "Effects of candesartan on mortality and morbidity in patients with chronic heart failure: the CHARM-Overall programme." *Lancet*, **362** (2003) (9386); pp. 759–66. ISSN 1474-547X.

[64] Heusch, G. "Hibernating Myocardium". *Physiol Rev*, **78** (1998) (4); pp. 1055–1085.

[65] Wolff, A. A., Rotmensch, H. H., Stanley, W. C., and Ferrari, R. "Metabolic approaches to the treatment of ischemic heart disease: the clinicians' perspective." *Heart Fail. Rev.*, **7** (2002) (2); pp. 187–203. ISSN 1382-4147.

[66] Kennedy, J. A., Kiosoglous, A. J., Murphy, G. A., Pelle, M. A., and Horowitz, J. D. "Effect of perhexiline and oxfenicine on myocardial function and metabolism during low-flow ischemia/reperfusion in the isolated rat heart." *J. Cardiovasc. Pharmacol.*, **36** (2000) (6); pp. 794–801. ISSN 0160-2446.

[67] Holubarsch, C. J. F., Rohrbach, M., Karrasch, M., Boehm, E., Polonski, L., Ponikowski, P., and Rhein, S. "A double-blind randomized multicentre clinical trial to evaluate the efficacy and safety of two doses of etomoxir in comparison with placebo in patients with moderate congestive heart failure: the ERGO (etomoxir for the recovery of glucose oxidation) stud". *Clin. Sci. (Lond).*, **113** (2007) (4); pp. 205–12. ISSN 1470-8736. DOI 10.1042/CS20060307.

[68] Birmingham, U. H. "Myocardial Protection With Perhexiline in Left Ventricular Hypertrophy: The HYPER Trial". *ClinicalTrials.gov*, (2011). DOI NCT00989508.

[69] Lee, L., Campbell, R., Scheuermann-Freestone, M., *et al.* "Metabolic modulation with perhexiline in chronic heart failure: a randomized, controlled trial of short-term use of a novel treatment." *Circulation*, **112** (2005) (21); pp. 3280–8. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.105.551457.

[70] Singh, S., Beadle, R., Cameron, D., *et al.* "Randomized double-blind placebo-controlled trial of perhexiline in heart failure with preserved ejection fraction syndrome." *Future Cardiol.*, **10** (2014) (6); pp. 693–8. ISSN 1744-8298. DOI 10.2217/fca.14.62.

[71] Ashrafian, H., Frenneaux, M. P., and Opie, L. H. "Metabolic mechanisms in heart failure." *Circulation*, **116** (2007) (4); pp. 434–48. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.107.702795.

[72] Macpherson, G. *Black's Medical Dictionary*. A&C Black (2002). ISBN 0713654422.

[73] Cross, H. R., Clarke, K., Opie, L. H., and Radda, G. K. "Is lactate-induced myocardial ischaemic injury mediated by decreased pH or increased intracellular lactate?" *J. Mol. Cell. Cardiol.*, **27** (1995) (7); pp. 1369–81. ISSN 0022-2828. DOI 10.1006/jmcc.1995.0130.

[74] Cross, H. R., Opie, L. H., Radda, G. K., and Clarke, K. "Is a high glycogen content beneficial or detrimental to the ischemic rat heart? A controversy resolved." *Circ. Res.*, **78** (1996) (3); pp. 482–91. ISSN 0009-7330.

[75] Stanley, W. C. "Myocardial energy metabolism during ischemia and the mechanisms of metabolic therapies." *J. Cardiovasc. Pharmacol. Ther.*, **9 Suppl 1** (2004); pp. S31–45. ISSN 1074-2484.

[76] Reimer, K. A., Lowe, J. E., Rasmussen, M. M., and Jennings, R. B. "The wavefront phenomenon of ischemic cell death: Myocardial infarct size vs duration of coronary occlusion in dogs." *Circulation*, **56** (1977) (5); pp. 786–94. ISSN 0009-7322.

[77] Di Carli, M. F., Prcevski, P., Singh, T. P., Janisse, J., Ager, J., Muzik, O., and Vander Heide, R. "Myocardial blood flow, function, and metabolism in repetitive stunning." *J. Nucl. Med.*, **41** (2000) (7); pp. 1227–34. ISSN 0161-5505.

[78] Dargie, H. "Heart failure post-myocardial infarction: a review of the issues." *Heart*, **91** (2005) (2); pp. ii3–48. ISSN 1468-201X. DOI 10.1136/hrt.2005.062018.

[79] Montalescot, G., Andersen, H. R., Antoniucci, D., Betriu, A., de Boer, M. J., Grip, L., Neumann, F. J., and Rothman, M. T. "Recommendations on percutaneous coronary intervention for the reperfusion of acute ST elevation myocardial infarction." *Heart*, **90** (2004) (6); p. e37. ISSN 1468-201X.

[80] Zijlstra, F., Hoorntje, J. C., de Boer, M. J., Reiffers, S., Miedema, K., Ottervanger, J. P., van 't Hof, A. W., and Suryapranata, H. "Long-term benefit of primary angioplasty as compared with thrombolytic therapy for acute myocardial infarction." *N. Engl. J. Med.*, **341** (1999) (19); pp. 1413–9. ISSN 0028-4793. DOI 10.1056/NEJM199911043411901.

[81] Frangogiannis, N. G., Smith, C. W., and Entman, M. L. "The inflammatory response in myocardial infarction." *Cardiovasc. Res.*, **53** (2002) (1); pp. 31–47. ISSN 0008-6363.

[82] Dobaczewski, M., Xia, Y., Bujak, M., Gonzalez-Quesada, C., and Frangogiannis, N. G. "CCR5 Signaling Suppresses Inflammation and Reduces Adverse Remodeling of the Infarcted Heart, Mediating Recruitment of Regulatory T Cells". *Am. J. Pathol.*, **176** (2010) (5); pp. 2177–2187. ISSN 00029440. DOI 10.2353/ajpath.2010.090759.

[83] Huebener, P., Abou-Khamis, T., Zymek, P., Bujak, M., Ying, X., Chatila, K., Haudek, S., Thakker, G., and Frangogiannis, N. G. "CD44 is critically involved in infarct healing by regulating the inflammatory and fibrotic response." *J. Immunol.*, **180** (2008) (4); pp. 2625–33. ISSN 0022-1767.

[84] de Lemos, J. A., Morrow, D. A., Blazing, M. A., Jarolim, P., Wiviott, S. D., Sabatine, M. S., Califf, R. M., and Braunwald, E. "Serial measurement of monocyte chemoattractant protein-1 after acute coronary syndromes: results from the A to Z trial." *J. Am. Coll. Cardiol.*, **50** (2007) (22); pp. 2117–24. ISSN 1558-3597. DOI 10.1016/j.jacc.2007.06.057.

[85] Roberts, R., DeMello, V., and Sobel, B. E. "Deleterious effects of methylprednisolone in patients with myocardial infarction." *Circulation*, **53** (1976) (3 Suppl); pp. I204–6. ISSN 0009-7322.

[86] Faxon, D. P., Gibbons, R. J., Chronos, N. A. F., Gurbel, P. A., and Sheehan, F. "The effect of blockade of the CD11/CD18 integrin receptor on infarct size in patients with acute myocardial infarction treated with direct angioplasty: the results of the HALT-MI study." *J. Am. Coll. Cardiol.*, **40** (2002) (7); pp. 1199–204. ISSN 0735-1097.

[87] Armstrong, P. W., Granger, C. B., Adams, P. X., Hamm, C., Holmes, D., O'Neill, W. W., Todaro, T. G., Vahanian, A., and Van de Werf, F. "Pexelizumab for Acute ST-Elevation Myocardial Infarction in Patients Undergoing Primary Percutaneous Coronary Intervention". *JAMA*, **297** (2007) (1); p. 43. ISSN 0098-7484. DOI 10.1001/jama.297.1.43.

[88] Chatham, J. C. and Young, M. E. "Metabolic remodeling in the hypertrophic heart: fuel for thought." *Circ. Res.*, **111** (2012) (6); pp. 666–8. ISSN 1524-4571. DOI 10.1161/CIRCRESAHA.112.277392.

[89] Herrmann, G. and Decherd, G. M. "The chemical nature of heart failure". *Ann. Intern. Med.*, **12** (1939) (8); p. 1233. ISSN 0003-4819. DOI 10.7326/0003-4819-12-8-1233.

[90] Merritt, M. E., Harrison, C., Storey, C., Jeffrey, F. M., Sherry, A. D., and Malloy, C. R. "Hyperpolarized 13C allows a direct measure of flux through a single enzyme-catalyzed step by NMR." *Proc. Natl. Acad. Sci. U. S. A.*, **104** (2007) (50); pp. 19 773–7. ISSN 1091-6490. DOI 10.1073/pnas.0706235104.

[91] Schroeder, M. A., Atherton, H. J., Ball, D. R., Cole, M. a., Heather, L. C., Griffin, J. L., Clarke, K., Radda, G. K., and Tyler, D. J. "Real-time assessment of Krebs cycle metabolism using hyperpolarized 13C magnetic resonance spectroscopy." *FASEB J.*, **23** (2009) (8); pp. 2529–38. ISSN 1530-6860. DOI 10.1096/fj.09-129171.

[92] Ball, D. R., Cruickshank, R., Carr, C. A., Stuckey, D. J., Lee, P., Clarke, K., and Tyler, D. J. "Metabolic imaging of acute and chronic infarction in the perfused rat heart using hyperpolarised [1-13C]pyruvate." *NMR Biomed.*, **26** (2013) (11); pp. 1441–50. ISSN 1099-1492. DOI 10.1002/nbm.2972.

[93] Atherton, H. J., Dodd, M. S., Heather, L. C., Schroeder, M. a., Griffin, J. L., Radda, G. K., Clarke, K., and Tyler, D. J. "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." *Circulation*, **123** (2011) (22); pp. 2552–61. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.110.011387.

[94] Dodd, M. S., Ball, D. R., Schroeder, M. A., *et al.* "In vivo alterations in cardiac metabolism and function in the spontaneously hypertensive rat heart." *Cardiovasc. Res.*, **95** (2012) (1); pp. 69–76. ISSN 1755-3245. DOI 10.1093/cvr/cvs164.

[95] Schroeder, M. A., Lau, A. Z., Chen, A. P., *et al.* "Hyperpolarized (13)C magnetic resonance reveals early- and late-onset changes to in vivo pyruvate metabolism in the failing heart." *Eur. J. Heart Fail.*, **15** (2013) (2); pp. 130–40. ISSN 1879-0844. DOI 10.1093/eurjhf/hfs192.

[96] Lau, A. Z., Chen, A. P., Barry, J., Graham, J. J., Dominguez-Viqueira, W., Ghugre, N. R., Wright, G. A., and Cunningham, C. H. "Reproducibility study for free-breathing measurements of pyruvate metabolism using hyperpolarized (13) C in the heart." *Magn. Reson. Med.*, **69** (2013) (4); pp. 1063–71. ISSN 1522-2594. DOI 10.1002/mrm.24342.

[97] Rider, O. J. and Tyler, D. J. "Clinical implications of cardiac hyperpolarized magnetic resonance imaging." *J. Cardiovasc. Magn. Reson.*, **15** (2013) (1); p. 93. ISSN 1532-429X. DOI 10.1186/1532-429X-15-93.

[98] Warburg, O. "The metabolism of tumours. Investigations from the Kaiser-Wilhelm Institute for Biology, Berlin-Dahlem. Edited by Otto Warburg, Kaiser-Wilhelm Institute for Biology, Berlin-Dahlem. Translated from the German edition, with accounts of additional recent res". *Br. J. Surg.*, **19** (1931) (73); pp. 168–168. ISSN 00071323. DOI 10.1002/bjs.1800197331.

[99] El Mjiyad, N., Caro-Maldonado, A., Ramírez-Peinado, S., and Muñoz-Pinedo, C. "Sugar-free approaches to cancer cell killing." *Oncogene*, **30** (2011) (3); pp. 253–64. ISSN 1476-5594. DOI 10.1038/onc.2010.466.

[100] Bé, M.-M., Chisté, V., Dulieu, C., Mougeot, X., Chechev, V. P., Kondev, F. G., Nichols, A. L., Huang, X., and Wang, B. *Table of Radionuclides*, vol. 7 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, Pavillon de Breteuil, F-92310 Sèvres, France (2013). ISBN 978-92-822-2248-5.

[101] Huang, B., Law, M. W.-M., and Khong, P.-L. "Whole-body PET/CT scanning: estimation of radiation dose and cancer risk." *Radiology*, **251** (2009) (1); pp. 166–74. ISSN 1527-1315. DOI 10.1148/radiol.2511081300.

[102] Nelson, S. J., Kurhanewicz, J., Vigneron, D. B., *et al.* "Metabolic imaging of patients with prostate cancer using hyperpolarized [1-13C]pyruvate." *Sci. Transl. Med.*, **5** (2013) (198); p. 108. ISSN 1946-6242. DOI 10.1126/scitranslmed.3006070.

[103] Saha, G. B. *Basics of PET Imaging: Physics, Chemistry, and Regulations*. Springer New York (2010). ISBN 9781441908056.

[104] Moses, W. W. "Fundamental Limits of Spatial Resolution in PET." *Nucl. Instrum. Methods Phys. Res. A.*, **648 Supple** (2011); pp. S236–S240. ISSN 0168-9002. DOI 10.1016/j.nima.2010.11.092.

[105] Halzen, F. and Martin, A. D. *Quarks and Leptons*. Wiley India Pvt. Limited (2008). ISBN 8126516569.

[106] Shibuya, K., Yoshida, E., Nishikido, F., Suzuki, T., Tsuda, T., Inadama, N., Yamaya, T., and Murayama, H. "Annihilation photon acollinearity in PET: volunteer and phantom FDG studies." *Phys. Med. Biol.*, **52** (2007) (17); pp. 5249–61. ISSN 0031-9155. DOI 10.1088/0031-9155/52/17/010.

[107] Hung, J. C. "Bringing New PET drugs to clinical practice - a regulatory perspective." *Theranostics*, **3** (2013) (11); pp. 885–93. ISSN 1838-7640. DOI 10.7150/thno.5513.

[108] Witney, T. H. and Brindle, K. M. "Imaging tumour cell metabolism using hyperpolarized 13C magnetic resonance spectroscopy." *Biochem. Soc. Trans.*, **38** (2010) (5); pp. 1220–4. ISSN 1470-8752. DOI 10.1042/BST0381220.

[109] Golman, K., Zandt, R. I., Lerche, M., Pehrson, R., and Ardenkjaer-Larsen, J. H. "Metabolic imaging by hyperpolarized 13C magnetic resonance imaging for in vivo tumor diagnosis." *Cancer Res.*, **66** (2006) (22); pp. 10 855–60. ISSN 0008-5472. DOI 10.1158/0008-5472.CAN-06-2564.

[110] Day, S. E., Kettunen, M. I., Gallagher, F. A., Hu, D.-E., Lerche, M., Wolber, J., Golman, K., Ardenkjaer-Larsen, J. H., and Brindle, K. M. "Detecting tumor response to treatment using hyperpolarized 13C magnetic resonance imaging and spectroscopy". *Nat. Med.*, **13** (2007) (11); pp. 1382–1387. ISSN 1078-8956. DOI 10.1038/nm1650.

[111] Gallagher, F., Kettunen, M., and Brindle, K. "Biomedical applications of hyperpolarized 13C magnetic resonance imaging". *Prog. Nucl. Magn. Reson. Spectrosc.*, **55** (2009) (4); pp. 285–295. ISSN 00796565. DOI 10.1016/j.pnmrs.2009.06.001.

[112] Deberardinis, R. J., Sayed, N., Ditsworth, D., and Thompson, C. B. "Brick by brick: metabolism and tumor cell growth." *Curr. Opin. Genet. Dev.*, **18** (2008) (1); pp. 54–61. ISSN 0959-437X. DOI 10.1016/j.gde.2008.02.003.

[113] Brink, I., Schumacher, T., Mix, M., Ruhland, S., Stoelben, E., Digel, W., Henke, M., Ghanem, N., Moser, E., and Nitzsche, E. U. "Impact of [18F]FDG-PET on the primary staging of small-cell lung cancer." *Eur. J. Nucl. Med. Mol. Imaging*, **31** (2004) (12); pp. 1614–20. ISSN 1619-7070. DOI 10.1007/s00259-004-1606-x.

[114] Zlokovic, B. V. "The blood-brain barrier in health and chronic neurodegenerative disorders." *Neuron*, **57** (2008) (2); pp. 178–201. ISSN 0896-6273. DOI 10.1016/j.neuron.2008.01.003.

[115] Lorger, M. "Tumor Microenvironment in the Brain". *Cancers (Basel).*, **4** (2012) (1); pp. 218–243. ISSN 2072-6694. DOI 10.3390/cancers4010218.

[116] Connell, J. J., Chatain, G., Cornelissen, B., Vallis, K. A., Hamilton, A., Seymour, L., Anthony, D. C., and Sibson, N. R. "Selective permeabilization of the blood-brain barrier at sites of metastasis." *J. Natl. Cancer Inst.*, **105** (2013) (21); pp. 1634–43. ISSN 1460-2105. DOI 10.1093/jnci/djt276.

[117] Eichler, A. F., Chung, E., Kodack, D. P., Loeffler, J. S., Fukumura, D., and Jain, R. K. "The biology of brain metastases-translation to new therapies." *Nat. Rev. Clin. Oncol.*, **8** (2011) (6); pp. 344–56. ISSN 1759-4782. DOI 10.1038/nrclinonc.2011.58.

[118] Park, I., Larson, P. E. Z., Zierhut, M. L., *et al.* "Hyperpolarized 13C magnetic resonance metabolic imaging: application to brain tumors." *Neuro. Oncol.*, **12** (2010) (2); pp. 133–44. ISSN 1523-5866. DOI 10.1093/neuonc/nop043.

[119] Marjańska, M., Iltis, I., Shestov, A. A., Deelchand, D. K., Nelson, C., Uğurbil, K., and Henry, P.-G. "In vivo 13C spectroscopy in the rat brain using hyperpolarized [1-13C]pyruvate and [2-13C]pyruvate". *J. Magn. Reson.*, **206** (2010) (2); pp. 210–218.

[120] Hurd, R. E., Yen, Y.-F., Mayer, D., *et al.* "Metabolic imaging in the anesthetized rat brain using hyperpolarized [1-13C] pyruvate and [1-13C] ethyl pyruvate." *Magn. Reson. Med.*, **63** (2010) (5); pp. 1137–43. ISSN 1522-2594. DOI 10.1002/mrm.22364.

[121] Bhattacharya, P., Chekmenev, E. Y., Perman, W. H., Harris, K. C., Lin, A. P., Norton, V. A., Tan, C. T., Ross, B. D., and Weitekamp, D. P. "Towards hyperpolarized 13C-succinate imaging of brain cancer" (2007).

[122] Ross, B. D., Bhattacharya, P., Wagner, S., Tran, T., and Sailasuta, N. "Hyperpolarized MR imaging: neurologic applications of hyperpolarized metabolism." *AJNR. Am. J. Neuroradiol.*, **31** (2010) (1); pp. 24–33. ISSN 1936-959X. DOI 10.3174/ajnr.A1790.

[123] Mayer, D., Yen, Y., Takahashi, A., and Josan, S. "Dynamic and High-Resolution Metabolic Imaging of the Rat Brain In Vivo Using Hyperpolarized -Pyruvate". In "Int. Soc. Magn. Reson. Med. 20th Annu. Meet.", (2010).

[124] Mishkovsky, M., Comment, A., and Gruetter, R. "In vivo detection of brain Krebs cycle intermediate by hyperpolarized magnetic resonance." *J. Cereb. Blood Flow Metab.*, **32** (2012) (12); pp. 2108–13. ISSN 1559-7016. DOI 10.1038/jcbfm.2012.136.

[125] Gallagher, F. A., Bohndiek, S. E., Kettunen, M. I., Lewis, D. Y., Soloviev, D., and Brindle, K. M. "Hyperpolarized 13C MRI and PET: in vivo tumor biochemistry." *J. Nucl. Med.*, **52** (2011) (9); pp. 1333–6. ISSN 1535-5667. DOI 10.2967/jnumed.110.085258.

[126] Hall, D. G. and Stoica, G. "Characterization of brain and bone-metastasizing clones selected from an ethylnitrosourea-induced rat mammary carcinoma." *Clin. Exp. Metastasis*, **12** (1994) (4); pp. 283–95. ISSN 0262-0898.

[127] Stefan, D. and Cesare, F. "Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package". *Meas. Sci. Technol.*, **20** (2009) (10). DOI 10.1088/0957-0233/20/10/104035.

[128] van Kasteren, S. I., Campbell, S. J., Serres, S., Anthony, D. C., Sibson, N. R., and Davis, B. G. "Glyconanoparticles allow pre-symptomatic in vivo imaging of brain disease." *Proc. Natl. Acad. Sci. U. S. A.*, **106** (2009) (1); pp. 18–23. ISSN 1091-6490. DOI 10.1073/pnas.0806787106.

[129] Coudène, P., Marson, B., Badiou, S., Flavier, S., Anelli, S., Cristol, J. P., and Dupuy, A. M. "Evaluation of the ABX Pentra 400: a newly available clinical chemistry analyser." *Clin. Chem. Lab. Med.*, **43** (2005) (7); pp. 782–92. ISSN 1434-6621. DOI 10.1515/CCLM.2005.133.

[130] Bergmeyer, H. U., Gawehn, K., Williamson, D. H., and Lund, P. *Methods of Enzymatic Analysis*. No. v. 4 in Methods of Enzymatic Analysis. Verlag Chemie (1974). ISBN 9780120913046.

[131] Levitt, M. H. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. John Wiley & Sons (2013). ISBN 1118681843.

[132] Hu, S., Lustig, M., Chen, A. P., *et al.* "Compressed sensing for resolution enhancement of hyperpolarized 13C flyback 3D-MRSI." *J. Magn. Reson.*, **192** (2008) (2); pp. 258–64. ISSN 1090-7807. DOI 10.1016/j.jmr.2008.03.003.

[133] Larson, P. E. Z., Bok, R., Kerr, A. B., Lustig, M., Hu, S., Chen, A. P., Nelson, S. J., Pauly, J. M., Kurhanewicz, J., and Vigneron, D. B. "Investigation of tumor hyperpolarized [1-13C]-pyruvate dynamics using time-resolved multiband RF excitation echo-planar MRSI." *Magn. Reson. Med.*, **63** (2010) (3); pp. 582–91. ISSN 1522-2594. DOI 10.1002/mrm.22264.

[134] Wiesinger, F., Weidl, E., Menzel, M. I., Janich, M. A., Khegai, O., Glaser, S. J., Haase, A., Schwaiger, M., and Schulte, R. F. "IDEAL spiral CSI for dynamic metabolic MR imaging of hyperpolarized [1-13C]pyruvate." *Magn. Reson. Med.*, **68** (2012) (1); pp. 8–16. ISSN 1522-2594. DOI 10.1002/mrm.23212.

[135] Meyer, C. H., Pauly, J. M., Macovskiand, A., and Nishimura, D. G. "Simultaneous spatial and spectral selective excitation". *Magn. Reson. Med.*, **15** (1990) (2); pp. 287–304. ISSN 07403194. DOI 10.1002/mrm.1910150211.

[136] Cunningham, C. H., Chen, A. P., Lustig, M., *et al.* "Pulse sequence for dynamic volumetric imaging of hyperpolarized metabolic products." *J. Magn. Reson.*, **193** (2008) (1); pp. 139–46. ISSN 1090-7807. DOI 10.1016/j.jmr.2008.03.012.

[137] Zur, Y. "Design of improved spectral-spatial pulses for routine clinical use." *Magn. Reson. Med.*, **43** (2000) (3); pp. 410–20. ISSN 0740-3194.

[138] Lau, A. Z., Chen, A. P., Ghugre, N. R., Ramanan, V., Lam, W. W., Connelly, K. A., Wright, G. A., and Cunningham, C. H. "Rapid multislice imaging of hyperpolarized 13C pyruvate and bicarbonate in the heart." *Magn. Reson. Med.*, **64** (2010) (5); pp. 1323–31. ISSN 1522-2594. DOI 10.1002/mrm.22525.

[139] Cunningham, C. H., Chen, A. P., Lustig, M., *et al.* "Pulse sequence for dynamic volumetric imaging of hyperpolarized metabolic products." *J. Magn. Reson.*, **193** (2008) (1); pp. 139–46. ISSN 1090-7807. DOI 10.1016/j.jmr.2008.03.012.

[140] Schulte, R. F., Sperl, J. I., Weidl, E., *et al.* "Saturation-recovery metabolic-exchange rate imaging with hyperpolarized [1-(13) C] pyruvate using spectral-spatial excitation." *Magn. Reson. Med.*, **69** (2012) (5); pp. 1209–16. ISSN 1522-2594. DOI 10.1002/mrm.24353.

[141] Sigfridsson, A., Weiss, K., Wissmann, L., Busch, J., Krajewski, M., Batel, M., Batsios, G., Ernst, M., and Kozerke, S. "Hybrid multiband excitation multiecho acquisition for hyperpolarized (13) C spectroscopic imaging." *Magn. Reson. Med.*, (2014). ISSN 1522-2594. DOI 10.1002/mrm.25294.

[142] Andersson, J. L. R., Skare, S., and Ashburner, J. "How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging." *Neuroimage*, **20** (2003) (2); pp. 870–88. ISSN 1053-8119. DOI 10.1016/S1053-8119(03)00336-7.

[143] Schneider, J. E., Cassidy, P. J., Lygate, C., Tyler, D. J., Wiesmann, F., Grieve, S. M., Hulbert, K., Clarke, K., and Neubauer, S. "Fast, high-resolution in vivo cine magnetic resonance imaging in normal and failing mouse hearts on a vertical 11.7 T system." *J. Magn. Reson. Imaging*, **18** (2003) (6); pp. 691–701. ISSN 1053-1807. DOI 10.1002/jmri.10411.

[144] Foerster, B. U., Tomasi, D., and Caparelli, E. C. "Magnetic field shift due to mechanical vibration in functional magnetic resonance imaging." *Magn. Reson. Med.*, **54** (2005) (5); pp. 1261–7. ISSN 0740-3194. DOI 10.1002/mrm.20695.

[145] Pfeuffer, J., Van de Moortele, P.-F., Ugurbil, K., Hu, X., and Glover, G. H. "Correction of physiologically induced global off-resonance effects in dynamic echo-planar and spiral functional imaging." *Magn. Reson. Med.*, **47** (2002) (2); pp. 344–53. ISSN 0740-3194.

[146] Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., and Smith, S. M. "FSL." *Neuroimage*, **62** (2012) (2); pp. 782–90. ISSN 1095-9572. DOI 10.1016/j.neuroimage.2011.09.015.

[147] Smith, S. M., Jenkinson, M., Woolrich, M. W., *et al.* "Advances in functional and structural MR image analysis and implementation as FSL." *Neuroimage*, **23 Suppl 1** (2004); pp. S208–19. ISSN 1053-8119. DOI 10.1016/j.neuroimage.2004.07.051.

[148] Kellman, P. and McVeigh, E. R. "Image reconstruction in SNR units: a general method for SNR measurement." *Magn. Reson. Med.*, **54** (2005) (6); pp. 1439–47. ISSN 0740-3194. DOI 10.1002/mrm.20713.

[149] McKenzie, C. A., Yeh, E. N., Ohliger, M. A., Price, M. D., and Sodickson, D. K. "Self-calibrating parallel imaging with automatic coil sensitivity extraction." *Magn. Reson. Med.*, **47** (2002) (3); pp. 529–38. ISSN 0740-3194.

[150] Shechtman, E. and Irani, M. "Matching local self-similarities across images and videos". *IEEE Conf. Comput. Vis. Pattern Recognit.*, (2007).

[151] Chatfield, K., Philbin, J., and Zisserman, A. "Efficient retrieval of deformable shape classes using local self-similarities". In "IEEE 12th Int. Conf. Comput. Vis. Work. ICCV Work.", pp. 264–271 (2009). ISBN 9781424444427. DOI 10.1109/ICCVW.2009.5457691.

[152] Atherton, H. J., Schroeder, M. A., Dodd, M. S., *et al.* "Validation of the in vivo assessment of pyruvate dehydrogenase activity using hyperpolarised 13C MRS." *NMR Biomed.*, **24** (2011) (2); pp. 201–8. ISSN 1099-1492. DOI 10.1002/nbm.1573.

[153] Miller, J. J., Lau, A. Z., Teh, I., Schneider, J. E., Kinchecsh, P., Smart, S., Ball, V., Sibson, N., and Tyler, D. J. "Robust, high resolution three-dimensional hyperpolarised metabolic imaging of the healthy rat heart at 7 T". *Magn. Reson. Med.*, **0** (2015) (0); p. 0.

[154] Gilson, W. D. and Kraitchman, D. L. "Cardiac magnetic resonance imaging in small rodents using clinical 1.5 T and 3.0 T scanners." *Methods*, **43** (2007) (1); pp. 35–45. ISSN 1046-2023. DOI 10.1016/j.ymeth.2007.03.012.

[155] Silverman, J. "Sentience and sensation." *Lab Anim. (NY).*, **37** (2008) (10); pp. 465–7. ISSN 0093-7355. DOI 10.1038/laban1008-465.

[156] Nishimura, D. G., Irarrazabal, P., and Meyer, C. H. "A Velocityk-Space Analysis of Flow Effects in Echo-Planar and Spiral Imaging". *Magn. Reson. Med.*, **33** (1995) (4); pp. 549–556. ISSN 07403194. DOI 10.1002/mrm.1910330414.

[157] Mayer, D., Yen, Y.-F., Takahashi, A., Josan, S., Tropp, J., Rutt, B. K., Hurd, R. E., Spielman, D. M., and Pfefferbaum, A. "Dynamic and high-resolution metabolic imaging of hyperpolarized [1-13C]-pyruvate in the rat brain using a high-performance gradient insert." *Magn. Reson. Med.*, **65** (2011) (5); pp. 1228–33. ISSN 1522-2594. DOI 10.1002/mrm.22707.

[158] Weiss, K., Sigfridsson, A., Wissmann, L., Busch, J., Batel, M., Krajewski, M., Ernst, M., and Kozerke, S. "Accelerating hyperpolarized metabolic imaging of the heart by exploiting spatiotemporal correlations." *NMR Biomed.*, **26** (2013) (11); pp. 1380–6. ISSN 1099-1492. DOI 10.1002/nbm.2963.

[159] Wild, J. M., Paley, M. N. J., Viallon, M., Schreiber, W. G., van Beek, E. J. R., and Griffiths, P. D. "k-space filtering in 2D gradient-echo breath-hold hyperpolarized 3He MRI: spatial resolution and signal-to-noise ratio considerations." *Magn. Reson. Med.*, **47** (2002) (4); pp. 687–95. ISSN 0740-3194.

[160] Wild, J. M., Woodhouse, N., Paley, M. N. J., Fichele, S., Said, Z., Kasuboski, L., and van Beek, E. J. R. "Comparison between 2D and 3D gradient-echo sequences for MRI of human lung ventilation with hyperpolarized 3He." *Magn. Reson. Med.*, **52** (2004) (3); pp. 673–8. ISSN 0740-3194. DOI 10.1002/mrm.20164.

[161] Kantz, H. and Schreiber, T. "Human ECG: nonlinear deterministic versus stochastic aspects". *IEEE Proc. Sci. Meas. Technol.*, **145** (1998) (6); pp. 279–284. ISSN 1350-2344. DOI 10.1049/ip-smt:19982327.

[162] Gutiérrez, R. and Sandoval, L. "Detecting stochastic information of electrocardiograms". *arXiv Prepr. physics/0307040*, (2003).

[163] Ibrahim, J. G., Chen, M.-H., and Sinha, D. *Bayesian Survival Analysis*. Springer Science & Business Media (2001). ISBN 0387952772.

[164] Wong, S., Rods, M., Newmark, R., and Budinger, T. "Discrete analysis of stochastic NMR. I". *J. Magn. Reson.*, **87** (1990) (2); pp. 242–264. ISSN 00222364. DOI 10.1016/0022-2364(90)90004-S.

[165] Carr, C. A., Stuckey, D. J., Tan, J. J., Tan, S. C., Gomes, R. S. M., Camelliti, P., Messina, E., Giacomello, A., Ellison, G. M., and Clarke, K. "Cardiosphere-derived cells improve function in the infarcted rat heart for at least 16 weeks–an MRI study." *PLoS One*, **6** (2011) (10); p. e25 669. ISSN 1932-6203. DOI 10.1371/journal.pone.0025669.

[166] Higham, N. J. *Functions of Matrices: Theory and Computation*. SIAM (2008). ISBN 0898717779.

[167] Daniels, C. J., McLean, M. A., McGlashan, N., Graves, M. J., Robb, F. J., Lomas, D. J., Schulte, R. F., Brindle, K. M., and Gallagher, F. A. "Quantitative analysis for hyperpolarized 13C-pyruvate imaging: comparison of methods on a clinical system." In "Proc. Intl. Soc. Mag. Reson. Med. 23", p. 1887 (2015).

[168] Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., and R Core Team. *{nlme}: Linear and Nonlinear Mixed Effects Models* (2015).

[169] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2014).

[170] Laird, N. M. and Ware, J. H. "Random-Effects Models for Longitudinal Data". *Biometrics*, **38** (1982) (4); pp. pp. 963–974. ISSN 0006341X.

[171] Hadfield, J. D. "MCMC Methods for Multi-Response Generalized Linear Mixed Models: The MCMCglmm R Package". *J. Stat. Softw.*, **33** (2010) (2); pp. 1–22. ISSN 1548-7660.

[172] Proescholdt, M. A., Merrill, M. J., Stoerr, E.-M., Lohmeier, A., Pohl, F., and Brawanski, A. "Function of carbonic anhydrase IX in glioblastoma multiforme." *Neuro. Oncol.*, **14** (2012) (11); pp. 1357–66. ISSN 1523-5866. DOI 10.1093/neuonc/nos216.

[173] Serres, S., Martin, C. J., Sarmiento Soto, M., Bristow, C., O'Brien, E. R., Connell, J. J., Khrapitchev, A. A., and Sibson, N. R. "Structural and functional effects of metastases in rat brain determined by multimodal MRI." *Int. J. Cancer*, **134** (2014) (4); pp. 885–96. ISSN 1097-0215. DOI 10.1002/ijc.28406.

[174] Hermsen, R., de Ligt, J., Spee, W., *et al.* "Genomic landscape of rat strain and substrain variation." *BMC Genomics*, **16** (2015) (1); p. 357. ISSN 1471-2164. DOI 10.1186/s12864-015-1594-1.

[175] Rolfe, D. F. and Brown, G. C. "Cellular energy utilization and molecular origin of standard metabolic rate in mammals". *Physiol Rev*, **77** (1997) (3); pp. 731–758.

[176] Doenst, T. and Taegtmeyer, H. "Profound Underestimation of Glucose Uptake by [18F]2-Deoxy-2-fluoroglucose in Reperfused Rat Heart Muscle". *Circulation*, **97** (1998) (24); pp. 2454–2462. ISSN 0009-7322. DOI 10.1161/01.CIR.97.24.2454.

[177] Shimoji, K., Ravasi, L., Schmidt, K., Soto-Montenegro, M. L., Esaki, T., Seidel, J., Jagoda, E., Sokoloff, L., Green, M. V., and Eckelman, W. C. "Measurement of cerebral glucose metabolic rates in the anesthetized rat by dynamic scanning with 18F-FDG, the ATLAS small animal PET scanner, and arterial blood sampling." *J. Nucl. Med.*, **45** (2004) (4); pp. 665–72. ISSN 0161-5505.

[178] Korecky, B. and French, I. W. "Nucleic Acid Synthesis in Enlarged Hearts of Rats with Nutritional Anemia". *Circ. Res.*, **21** (1967) (5); pp. 635–640. ISSN 0009-7330. DOI 10.1161/01.RES.21.5.635.

[179] Kennedy, K. M. and Dewhirst, M. W. "Tumor metabolism of lactate: the influence and therapeutic potential for MCT and CD147 regulation." *Future Oncol.*, **6** (2010) (1); pp. 127–48. ISSN 1744-8301. DOI 10.2217/fon.09.145.

[180] Park, I., Bok, R., Ozawa, T., Phillips, J. J., James, C. D., Vigneron, D. B., Ronen, S. M., and Nelson, S. J. "Detection of early response to temozolomide treatment in brain tumors using hyperpolarized 13C MR metabolic imaging." *J. Magn. Reson. Imaging*, **33** (2011) (6); pp. 1284–90. ISSN 1522-2586. DOI 10.1002/jmri.22563.

[181] Dennis, S. C., Kohn, M. C., Slegowski, M. B., Anderson, G. J., and Garfinkel, D. "Monocarboxylate-uptake kinetics in perfused rat heart." *Adv. Myocardiol.*, **6** (1985); pp. 259–72. ISSN 0270-4056.

[182] Cremer, J. E., Cunningham, V. J., Pardridge, W. M., Braun, L. D., and Oldendorf, W. H. "Kinetics of blood-brain barrier transport of pyruvate, lactate and glucose in suckling, weanling and adult rats." *J. Neurochem.*, **33** (1979) (2); pp. 439–45. ISSN 0022-3042.

[183] Hurd, R. E., Yen, Y.-F., Mayer, D., *et al.* "Metabolic imaging in the anesthetized rat brain using hyperpolarized [1-13C] pyruvate and [1-13C] ethyl pyruvate." *Magn. Reson. Med.*, **63** (2010) (5); pp. 1137–43. ISSN 1522-2594. DOI 10.1002/mrm.22364.

[184] Fink, M. P. "Ethyl pyruvate: a novel anti-inflammatory agent." *J. Intern. Med.*, **261** (2007) (4); pp. 349–62. ISSN 0954-6820. DOI 10.1111/j.1365-2796.2007.01789.x.

[185] Wijns, W., Vatner, S. F., and Camici, P. G. "Hibernating myocardium." *N. Engl. J. Med.*, **339** (1998) (3); pp. 173–81. ISSN 0028-4793. DOI 10.1056/NEJM199807163390307.

[186] Tillisch, J., Brunken, R., Marshall, R., Schwaiger, M., Mandelkern, M., Phelps, M., and Schelbert, H. "Reversibility of cardiac wall-motion abnormalities predicted by positron tomography." *N. Engl. J. Med.*, **314** (1986) (14); pp. 884–8. ISSN 0028-4793. DOI 10.1056/NEJM198604033141405.

[187] Allman, K. C., Shaw, L. J., Hachamovitch, R., and Udelson, J. E. "Myocardial viability testing and impact of revascularization on prognosis in patients with coronary artery disease and left ventricular dysfunction: a meta-analysis." *J. Am. Coll. Cardiol.*, **39** (2002) (7); pp. 1151–8. ISSN 0735-1097.

[188] Aquilina, O., Grech, V., Felice, H., Debono, J., and Fenech, A. "Normal adult coronary angiography." *Images Paediatr. Cardiol.*, **8** (2006) (2); pp. 1–16. ISSN 1729-441X.

[189] Bluemke, D. A., Achenbach, S., Budoff, M., *et al.* "Noninvasive coronary artery imaging: magnetic resonance angiography and multidetector computed tomography angiography: a scientific statement from the american heart association committee on cardiovascular imaging and intervention of the council on cardio". *Circulation*, **118** (2008) (5); pp. 586–606. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.108.189695.

[190] Hendel, R. C., Abbott, B. G., Bateman, T. M., *et al.* "The role of radionuclide myocardial perfusion imaging for asymptomatic individuals." *J. Nucl. Cardiol.*, **18** (2011) (1); pp. 3–15. ISSN 1532-6551. DOI 10.1007/s12350-010-9320-5.

[191] Mlynarska, A., Mlynarski, R., and Sosnowski, M. "Conus Artery in Cardiac Computed Tomography". *Arch. Cardiovasc. Imaging*, **2** (2014) (2). ISSN 2322-5327. DOI 10.5812/acvi.19641.

[192] Moon, J. H., Park, E.-A., Lee, W., Yin, Y. H., Chung, J. W., Park, J. H., Lee, H.-Y., Kang, H.-J., and Kim, H.-S. "The diagnostic accuracy, image quality and radiation dose of 64-slice dual-source CT in daily practice: a single institution's experience." *Korean J. Radiol.*, **12** (2008) (3); pp. 308–18. ISSN 2005-8330. DOI 10.3348/kjr.2011.12.3.308.

[193] Morin, R. L. "Radiation Dose in Computed Tomography of the Heart". *Circulation*, **107** (2003) (6); pp. 917–922. ISSN 00097322. DOI 10.1161/01.CIR.0000048965.56529.C2.

[194] Kleinerman, R. A. "Cancer risks following diagnostic and therapeutic radiation exposure in children." *Pediatr. Radiol.*, **36 Suppl 2** (2006); pp. 121–5. ISSN 1432-1998. DOI 10.1007/s00247-006-0191-5.

[195] Hunt, C. H., Hartman, R. P., and Hesley, G. K. "Frequency and severity of adverse effects of iodinated and gadolinium contrast materials: retrospective review of 456,930 doses." *AJR. Am. J. Roentgenol.*, **193** (2009) (4); pp. 1124–7. ISSN 1546-3141. DOI 10.2214/AJR.09.2520.

[196] Campbell, K. L., Hud, L. M., Adams, S., Andrel, J., Ballas, S. K., Feldman, A. M., and Axelrod, D. "Safety of iodinated intravenous contrast medium administration in sickle cell disease." *Am. J. Med.*, **125** (2012) (1); pp. 100.e11–6. ISSN 1555-7162. DOI 10.1016/j.amjmed.2011.06.010.

[197] Vick, G. W. "The gold standard for noninvasive imaging in coronary heart disease: magnetic resonance imaging." *Curr. Opin. Cardiol.*, **24** (2009) (6); pp. 567–79. ISSN 1531-7080. DOI 10.1097/HCO.0b013e3283315553.

[198] van Nierop, B. J., Coolen, B. F., Dijk, W. J., Hendriks, A. D., de Graaf, L., Nicolay, K., and Strijkers, G. J. "Quantitative first-pass perfusion MRI of the mouse myocardium". *Magn. Reson. Med.*, **69** (2013) (6); pp. 1735–1744. ISSN 07403194. DOI 10.1002/mrm.24424.

[199] Stuckey, D. J., Carr, C. A., Meader, S. J., Tyler, D. J., Cole, M. A., and Clarke, K. "First-pass perfusion CMR two days after infarction predicts severity of functional impairment six weeks later in the rat heart." *J. Cardiovasc. Magn. Reson.*, **13** (2011); p. 38. ISSN 1532-429X. DOI 10.1186/1532-429X-13-38.

[200] Makowski, M., Jansen, C., Webb, I., Chiribiri, A., Nagel, E., Botnar, R., Kozerke, S., and Plein, S. "First-pass contrast-enhanced myocardial perfusion MRI in mice on a 3-T clinical MR scanner". *Magn. Reson. Med.*, **64** (2010) (6); pp. 1592–1598. ISSN 07403194. DOI 10.1002/mrm.22470.

[201] Jogiya, R., Makowski, M., Phinikaridou, A., *et al.* "Hyperemic stress myocardial perfusion cardiovascular magnetic resonance in mice at 3 Tesla: initial experience and validation against microspheres". *J. Cardiovasc. Magn. Reson.*, **15** (2013) (1); p. 62. ISSN 1532-429X. DOI 10.1186/1532-429X-15-62.

[202] Coolen, B. F., Moonen, R. P. M., Paulis, L. E. M., Geelen, T., Nicolay, K., and Strijkers, G. J. "Mouse myocardial first-pass perfusion MR imaging". *Magn. Reson. Med.*, **64** (2010) (6); pp. 1658–1663. ISSN 07403194. DOI 10.1002/mrm.22588.

[203] Kellman, P. and Arai, A. E. "Imaging sequences for first pass perfusion –a review." *J. Cardiovasc. Magn. Reson.*, **9** (2007) (3); pp. 525–37. ISSN 1097-6647. DOI 10.1080/10976640601187604.

[204] Jerosch-Herold, M. "Quantification of myocardial perfusion by cardiovascular magnetic resonance." *J. Cardiovasc. Magn. Reson.*, **12** (2010); p. 57. ISSN 1532-429X. DOI 10.1186/1532-429X-12-57.

[205] Atkinson, D. J., Burstein, D., and Edelman, R. R. "First-pass cardiac perfusion: evaluation with ultrafast MR imaging." *Radiology*, **174** (1990) (3 Pt 1); pp. 757–62. ISSN 0033-8419. DOI 10.1148/radiology.174.3.2305058.

[206] Yang, Q., Li, K., Liu, X., Bi, X., Liu, Z., An, J., Zhang, A., Jerecic, R., and Li, D. "Contrast-enhanced whole-heart coronary magnetic resonance angiography at 3.0 T: a comparative study with X-ray angiography in a single center." *J. Am. Coll. Cardiol.*, **54** (2009) (1); pp. 69–76. ISSN 1558-3597. DOI 10.1016/j.jacc.2009.03.016.

[207] Penfield, J. G. and Reilly, R. F. "Nephrogenic systemic fibrosis risk: is there a difference between gadolinium-based contrast agents?" *Semin. Dial.*, **21** (2008) (2); pp. 129–34. ISSN 0894-0959. DOI 10.1111/j.1525-139X.2007.00408.x.

[208] Moreno-Romero, J. A., Segura, S., Mascaró, J. M., Cowper, S. E., Julià, M., Poch, E., Botey, A., and Herrero, C. "Nephrogenic systemic fibrosis: a case series suggesting gadolinium as a possible aetiological factor." *Br. J. Dermatol.*, **157** (2007) (4); pp. 783–7. ISSN 0007-0963. DOI 10.1111/j.1365-2133.2007.08067.x.

[209] Perazella, M. A. and Rodby, R. A. "Gadolinium-induced nephrogenic systemic fibrosis in patients with kidney disease." *Am. J. Med.*, **120** (2007) (7); pp. 561–2. ISSN 1555-7162. DOI 10.1016/j.amjmed.2007.01.032.

[210] Hartung, M. P., Grist, T. M., and François, C. J. "Magnetic resonance angiography: current status and future directions." *J. Cardiovasc. Magn. Reson.*, **13** (2011); p. 19. ISSN 1532-429X. DOI 10.1186/1532-429X-13-19.

[211] Kober, F., Iltis, I., Izquierdo, M., Desrois, M., Ibarrola, D., Cozzone, P. J., and Bernard, M. "High-resolution myocardial perfusion mapping in small animals in vivo by spin-labeling gradient-echo imaging." *Magn. Reson. Med.*, **51** (2004) (1); pp. 62–7. ISSN 0740-3194. DOI 10.1002/mrm.10676.

[212] Vandsburger, M. H., Janiczek, R. L., Xu, Y., French, B. A., Meyer, C. H., Kramer, C. M., and Epstein, F. H. "Improved arterial spin labeling after myocardial infarction in mice using cardiac and respiratory gated look-locker imaging with fuzzy C-means clustering." *Magn. Reson. Med.*, **63** (2010) (3); pp. 648–57. ISSN 1522-2594. DOI 10.1002/mrm.22280.

[213] Belle, V., Kahler, E., Waller, C., Rommel, E., Voll, S., Hiller, K.-H., Bauer, W. R., and Haase, A. "In Vivo quantitative mapping of cardiac perfusion in rats using a noninvasive MR spin-labeling method". *J. Magn. Reson. Imaging*, **8** (1998) (6); pp. 1240–1245. ISSN 10531807. DOI 10.1002/jmri.1880080610.

[214] Stankovic, Z., Allen, B. D., Garcia, J., Jarvis, K. B., and Markl, M. "4D flow imaging with MRI".

[215] Grant, A. K., Vinogradov, E., Wang, X., Lenkinski, R. E., and Alsop, D. C. "Perfusion imaging with a freely diffusible hyperpolarized contrast agent". *Magn. Reson. Med.*, **66** (2011) (3); pp. 746–755. ISSN 07403194. DOI 10.1002/mrm.22860.

[216] von Morze, C., Larson, P. E., Hu, S., Keshari, K., Wilson, D. M., Ardenkjaer-Larsen, J. H., Goga, A., Bok, R., Kurhanewicz, J., and Vigneron, D. B. "Imaging of blood flow using hyperpolarized [13C]Urea in preclinical cancer models". *J. Magn. Reson. Imaging*, **33** (2011) (3); pp. 692–697. ISSN 10531807. DOI 10.1002/jmri.22484.

[217] Johansson, E., Månsson, S., Wirestam, R., Svensson, J., Petersson, J. S., Golman, K., and Ståhlberg, F. "Cerebral perfusion assessment by bolus tracking using hyperpolarized 13C." *Magn. Reson. Med.*, **51** (2004) (3); pp. 464–72. ISSN 0740-3194. DOI 10.1002/mrm.20013.

[218] von Morze, C., Larson, P. E., Hu, S., Yoshihara, H. A., Bok, R. A., Goga, A., Ardenkjaer-Larsen, J. H., and Vigneron, D. B. "Investigating tumor perfusion and metabolism using multiple hyperpolarized 13C compounds: HP001, pyruvate and urea".

Magn. Reson. Imaging, **30** (2012) (3); pp. 305–311. ISSN 0730725X. DOI 10.1016/j.mri.2011.09.026.

[219] Golman, K. and Petersson, J. S. "Metabolic imaging and other applications of hyperpolarized 13C1." *Acad. Radiol.*, **13** (2006) (8); pp. 932–42. ISSN 1076-6332. DOI 10.1016/j.acra.2006.06.001.

[220] Oberhammer, T., Weiger, M., and Hennel, F. "Spiral MRI Trajectory Design with Frequency Constraint". In "Int. Soc. Magn. Reson. Med.", p. 4969 (2010).

[221] Likes, R. S. *Moving gradient zeugmatography*. Ph.D. thesis, General Electric Company (dec 1981).

[222] Winkelmann, S., Schaeffter, T., Koehler, T., Eggers, H., and Doessel, O. "An optimal radial profile order based on the Golden Ratio for time-resolved MRI." *IEEE Trans. Med. Imaging*, **26** (2007) (1); pp. 68–76. ISSN 0278-0062. DOI 10.1109/TMI.2006.885337.

[223] Pelc, N. J., Sommer, F. G., Li, K. C., Brosnan, T. J., Herfkens, R. J., and Enzmann, D. R. "Quantitative magnetic resonance flow imaging." *Magn. Reson. Q.*, **10** (1994) (3); pp. 125–47. ISSN 0899-9422.

[224] Patrick, P. S., Kettunen, M. I., Tee, S.-S., Rodrigues, T. B., Serrao, E., Timm, K. N., McGuire, S., and Brindle, K. M. "Detection of transgene expression using hyperpolarized 13C urea and diffusion-weighted magnetic resonance spectroscopy." *Magn. Reson. Med.*, **73** (2015) (4); pp. 1401–6. ISSN 1522-2594. DOI 10.1002/mrm.25254.

[225] Shang, H., Skloss, T., von Morze, C., Carvajal, L., Van Criekinge, M., Milshteyn, E., Larson, P. E. Z., Hurd, R. E., and Vigneron, D. B. "Handheld electromagnet carrier for transfer of hyperpolarized carbon-13 samples." *Magn. Reson. Med.*, (2015). ISSN 1522-2594. DOI 10.1002/mrm.25657.

[226] Fessler, J. and Sutton, B. "Nonuniform fast fourier transforms using min-max interpolation". *IEEE Trans. Signal Process.*, **51** (2003) (2); pp. 560–574. ISSN 1053-587X. DOI 10.1109/TSP.2002.807005.

[227] Zwart, N. R., Johnson, K. O., and Pipe, J. G. "Efficient sample density estimation by combining gridding and an optimized kernel." *Magn. Reson. Med.*, **67** (2012) (3); pp. 701–10. ISSN 1522-2594. DOI 10.1002/mrm.23041.

[228] Man, L. C., Pauly, J. M., and Macovski, A. "Multifrequency interpolation for fast off-resonance correction." *Magn. Reson. Med.*, **37** (1997) (5); pp. 785–92. ISSN 0740-3194.

[229] Belardinelli, L., Linden, J., and Berne, R. M. "The cardiac effects of adenosine". *Prog. Cardiovasc. Dis.*, **32** (1989) (1); pp. 73–97. ISSN 00330620. DOI 10.1016/0033-0620(89)90015-7.

[230] Jacquier, A., Kober, F., Bun, S., Giorgi, R., Cozzone, P. J., and Bernard, M. "Quantification of myocardial blood flow and flow reserve in rats using arterial spin labeling MRI: comparison with a fluorescent microsphere technique." *NMR Biomed.*, **24** (2011) (9); pp. 1047–53. ISSN 1099-1492. DOI 10.1002/nbm.1645.

[231] Oshita, A., Ohmori, K., Yu, Y., Kondo, I., Takeuchi, H., Takagi, Y., Wada, Y., Yukiiri, K., Mizushige, K., and Kohno, M. "Myocardial blood flow measurements in rats with simple pulsing contrast echocardiography." *Ultrasound Med. Biol.*, **28** (2002) (4); pp. 459–66. ISSN 0301-5629.

[232] Reed, G. D., von Morze, C., Bok, R., Koelsch, B. L., Van Criekinge, M., Smith, K. J., Hong Shang, Larson, P. E. Z., Kurhanewicz, J., and Vigneron, D. B. "High resolution (13)C MRI with hyperpolarized urea: in vivo T(2) mapping and (15)N labeling effects." *IEEE Trans. Med. Imaging*, **33** (2014) (2); pp. 362–71. ISSN 1558-254X. DOI 10.1109/TMI.2013.2285120.

[233] Grant, A. K., Vinogradov, E., Wang, X., Lenkinski, R. E., and Alsop, D. C. "Perfusion imaging with a freely diffusible hyperpolarized contrast agent". *Magn. Reson. Med.*, **66** (2011) (3); pp. 746–755. ISSN 07403194. DOI 10.1002/mrm.22860.

[234] Scheenen, T. W. J., Klomp, D. W. J., Wijnen, J. P., and Heerschap, A. "Short echo time 1H-MRSI of the human brain at 3T with minimal chemical shift displacement errors using adiabatic refocusing pulses." *Magn. Reson. Med.*, **59** (2008) (1); pp. 1–6. ISSN 0740-3194. DOI 10.1002/mrm.21302.

[235] Cunningham, C. H., Chen, A. P., Albers, M. J., Kurhanewicz, J., Hurd, R. E., Yen, Y.-F., Pauly, J. M., Nelson, S. J., and Vigneron, D. B. "Double spin-echo sequence for rapid spectroscopic imaging of hyperpolarized 13C." *J. Magn. Reson.*, **187** (2007) (2); pp. 357–62. ISSN 1090-7807. DOI 10.1016/j.jmr.2007.05.014.

[236] Chen, A. P., Hurd, R. E., and Cunningham, C. H. "Spin tagging for hyperpolarized ^{13}C metabolic studies." *J. Magn. Reson.*, **214** (2012) (1); pp. 319–23. ISSN 1096-0856. DOI 10.1016/j.jmr.2011.10.007.

[237] Koelsch, B. L., Reed, G. D., Keshari, K. R., Chaumeil, M. M., Bok, R., Ronen, S. M., Vigneron, D. B., Kurhanewicz, J., and Larson, P. E. Z. "Rapid in vivo apparent diffusion coefficient mapping of hyperpolarized (13) C metabolites." *Magn. Reson. Med.*, (2014). ISSN 1522-2594. DOI 10.1002/mrm.25422.

[238] Kettunen, M. I., Kennedy, B. W. C., Hu, D.-E., and Brindle, K. M. "Spin echo measurements of the extravasation and tumor cell uptake of hyperpolarized [1-(13) C]lactate and [1-(13) C]pyruvate." *Magn. Reson. Med.*, **70** (2013) (5); pp. 1200–9. ISSN 1522-2594. DOI 10.1002/mrm.24591.

[239] Nishimura, D. G., Irarrazabal, P., and Meyer, C. H. "A velocity k-space analysis of flow effects in echo-planar and spiral imaging." *Magn. Reson. Med.*, **33** (1995) (4); pp. 549–56. ISSN 0740-3194.

[240] Nguyen, T. D., de Rochefort, L., Spincemaille, P., Cham, M. D., Weinsaft, J. W., Prince, M. R., and Wang, Y. "Effective motion-sensitizing magnetization preparation for black blood magnetic resonance imaging of the heart." *J. Magn. Reson. Imaging*, **28** (2008) (5); pp. 1092–100. ISSN 1053-1807. DOI 10.1002/jmri.21568.

[241] Chen, A. P., Leung, K., Lam, W., Hurd, R. E., Vigneron, D. B., and Cunningham, C. H. "Design of spectral-spatial outer volume suppression RF pulses for tissue specific metabolic characterization with hyperpolarized 13C pyruvate." *J. Magn. Reson.*, **200** (2009) (2); pp. 344–8. ISSN 1096-0856. DOI 10.1016/j.jmr.2009.06.021.

[242] de Langen, A. J., van den Boogaart, V. E. M., Marcus, J. T., and Lubberink, M. "Use of H2(15)O-PET and DCE-MRI to measure tumor blood flow." *Oncologist*, **13** (2008) (6); pp. 631–44. ISSN 1083-7159. DOI 10.1634/theoncologist.2007-0235.

[243] Johansson, E., Olsson, L. E., Månsson, S., Petersson, J. S., Golman, K., Ståhlberg, F., and Wirestam, R. "Perfusion assessment with bolus differentiation: a technique applicable to hyperpolarized tracers." *Magn. Reson. Med.*, **52** (2004) (5); pp. 1043–51. ISSN 0740-3194. DOI 10.1002/mrm.20247.

[244] von Morze, C., Bok, R. A., Reed, G. D., Ardenkjaer-Larsen, J. H., Kurhanewicz, J., and Vigneron, D. B. "Simultaneous multiagent hyperpolarized (13)C perfusion imaging." *Magn. Reson. Med.*, **72** (2014) (6); pp. 1599–609. ISSN 1522-2594. DOI 10.1002/mrm.25071.

[245] Bergmann, S. R., Herrero, P., Markham, J., Weinheimer, C. J., and Walsh, M. N. "Noninvasive quantitation of myocardial blood flow in human subjects with oxygen-15-labeled water and positron emission tomography." *J. Am. Coll. Cardiol.*, **14** (1989) (3); pp. 639–52. ISSN 0735-1097.

[246] von Morze, C., Larson, P. E., Hu, S., Keshari, K., Wilson, D. M., Ardenkjaer-Larsen, J. H., Goga, A., Bok, R., Kurhanewicz, J., and Vigneron, D. B. "Imaging of blood flow using hyperpolarized [13C]Urea in preclinical cancer models". *J. Magn. Reson. Imaging*, **33** (2011) (3); pp. 692–697. ISSN 10531807. DOI 10.1002/jmri.22484.

[247] von Morze, C., Larson, P. E., Hu, S., Yoshihara, H. A., Bok, R. A., Goga, A., Ardenkjaer-Larsen, J. H., and Vigneron, D. B. "Investigating tumor perfusion and metabolism using multiple hyperpolarized 13C compounds: HP001, pyruvate and urea". *Magn. Reson. Imaging*, **30** (2012) (3); pp. 305–311. ISSN 0730725X. DOI 10.1016/j.mri.2011.09.026.

[248] Meyer, R. A., Fisher, M. J., Nelson, S. J., and Brown, T. R. "Evaluation of manual methods for integration ofin vivo phosphorus NMR spectra". *NMR Biomed.*, **1** (1988) (3); pp. 131–135. ISSN 09523480. DOI 10.1002/nbm.1940010306.

[249] Vanhamme, L., van den Boogaart A, and Van Huffel S. "Improved method for accurate and efficient quantification of MRS data with use of prior knowledge". *J. Magn. Reson.*, **129** (1997) (1); pp. 35–43. ISSN 1096-0856.

[250] Hao, J., Liebeke, M., Astle, W., and Iorio, M. D. "Bayesian deconvolution and quantification of metabolites in complex 1D NMR spectra using BATMAN". *Nat. Protoc.*, (2014).

[251] Zheng, C., Zhang, S., Ragg, S., Raftery, D., and Vitek, O. "Identification and quantification of metabolites in 1H NMR spectra by Bayesian model selection". *Bioinformatics*, (2011).

[252] van der Veen, J. W. C., de Beer, R., Luyten, P. R., and van Ormondt, D. "Accurate quantification of in vivo 31P NMR signals using the variable projection method and prior knowledge". *Magn. Reson. Med.*, **6** (1988) (1); pp. 92–98. ISSN 07403194. DOI 10.1002/mrm.1910060111.

[253] Ober, R. J., Lin, Z., Ye, H., and Ward, E. S. "Achievable accuracy of parameter estimation for multidimensional NMR experiments." *J. Magn. Reson.*, **157** (2002) (1); pp. 1–16. ISSN 1090-7807.

[254] De Groen, P. and De Moor, B. "The fit of a sum of exponentials to noisy data". *J. Comput. Appl. Math.*, **20** (1987); pp. 175–187. ISSN 03770427. DOI 10.1016/0377-0427(87)90135-X.

[255] Pijnappel, W., van den Boogaart, A., de Beer, R., and van Ormondt, D. "SVD-based quantification of magnetic resonance signals". *J. Magn. Reson.*, **97** (1992) (1); pp. 122–134. ISSN 00222364. DOI 10.1016/0022-2364(92)90241-X.

[256] Wang, X. and Lee, J.-H. "IRIS-HSVD algorithm for automatic quantitation of in vivo 31P MRS." *J. Magn. Reson.*, **196** (2009) (1); pp. 23–32. ISSN 1096-0856. DOI 10.1016/j.jmr.2008.09.024.

[257] Schroeder, M. A., Cochlin, L. E., Heather, L. C., Clarke, K., Radda, G. K., and Tyler, D. J. "In vivo assessment of pyruvate dehydrogenase flux in the heart using hyperpolarized carbon-13 magnetic resonance." *Proc. Natl. Acad. Sci. U. S. A.*, **105** (2008) (33); pp. 12 051–6. ISSN 1091-6490. DOI 10.1073/pnas.0805953105.

[258] Gudbjartsson, H. and Patz, S. "The Rician distribution of noisy MRI data." *Magn. Reson. Med.*, **34** (1995) (6); pp. 910–4. ISSN 0740-3194.

[259] Geary, R. C. "The Frequency Distribution of the Quotient of Two Normal Variates". *J. R. Stat. Soc.*, **93** (1930) (3); pp. 442–446.

[260] Hinkley, D. "On the ratio of two correlated normal random variables". *Biometrika*, (1969).

[261] Franz, V. H. "Ratios: A short guide to confidence limits and proper use". *ArXiv arXiv:0710*, (2007) (arXiv:0710.2024); p. 60.

[262] Fieller, E. C. "Some Problems in Interval Estimation". *J. R. Stat. Soc.*, **16** (1954) (2); pp. 175–185.

[263] Bottomley, P. A. "The trouble with spectroscopy papers". *J. Magn. Reson. Imaging*, **2** (1992) (1); pp. 1–8. ISSN 10531807. DOI 10.1002/jmri.1880020102.

[264] Hall, E. L., Stephenson, M. C., Price, D., and Morris, P. G. "Methodology for improved detection of low concentration metabolites in MRS: optimised combination of signals from multi-element coil arrays." *Neuroimage*, **86** (2014); pp. 35–42. ISSN 1095-9572. DOI 10.1016/j.neuroimage.2013.04.077.

[265] Schroeder, M. A., Clarke, K., Neubauer, S., and Tyler, D. J. "Hyperpolarized magnetic resonance: a novel technique for the in vivo assessment of cardiovascular disease." *Circulation*, **124** (2011) (14); pp. 1580–94. ISSN 1524-4539. DOI 10.1161/CIRCULATIONAHA.111.024919.

[266] Leupold, J., Månsson, S., Petersson, J. S., Hennig, J., and Wieben, O. "Fast multiecho balanced SSFP metabolite mapping of (1)H and hyperpolarized (13)C compounds." *MAGMA*, **22** (2009) (4); pp. 251–6. ISSN 1352-8661. DOI 10.1007/s10334-009-0169-z.

[267] Chen, A. P., Albers, M. J., Cunningham, C. H., *et al.* "Hyperpolarized C-13 spectroscopic imaging of the TRAMP mouse at 3T-initial experience." *Magn. Reson. Med.*, **58** (2007) (6); pp. 1099–106. ISSN 0740-3194. DOI 10.1002/mrm.21256.

[268] Mayer, D., Yen, Y.-F., Tropp, J., Pfefferbaum, A., Hurd, R. E., and Spielman, D. M. "Application of subsecond spiral chemical shift imaging to real-time multislice metabolic imaging of the rat in vivo after injection of hyperpolarized 13C1-pyruvate." *Magn. Reson. Med.*, **62** (2009) (3); pp. 557–64. ISSN 1522-2594. DOI 10.1002/mrm.22041.

[269] Laustsen, C., Østergaard, J. A., Lauritzen, M. H., Nørregaard, R., Bowen, S., Søgaard, L. V., Flyvbjerg, A., Pedersen, M., and Ardenkjaer-Larsen, J. H. "Assessment of early diabetic renal changes with hyperpolarized [1-(13) C]pyruvate." *Diabetes. Metab. Res. Rev.*, **29** (2013) (2); pp. 125–9. ISSN 1520-7560. DOI 10.1002/dmrr.2370.

[270] Laustsen, C., Lycke, S., Palm, F., Østergaard, J. A., Bibby, B. M., Nørregaard, R., Flyvbjerg, A., Pedersen, M., and Ardenkjaer-Larsen, J. H. "High altitude may alter oxygen availability and renal metabolism in diabetics as measured by hyperpolarized [1-(13)C]pyruvate magnetic resonance imaging." *Kidney Int.*, **86** (2014) (1); pp. 67–74. ISSN 1523-1755. DOI 10.1038/ki.2013.504.

[271] Bernstein, M. A., Fain, S. B., and Riederer, S. J. "Effect of windowing and zero-filled reconstruction of MRI data on spatial resolution and acquisition strategy." *J. Magn. Reson. Imaging*, **14** (2001) (3); pp. 270–80. ISSN 1053-1807.

[272] Djermoune, E.-H., Tomczak, M., and Brie, D. "NMR Data Analysis: A Time-Domain Parametric Approach Using Adaptive Subband Decomposition". *Oil Gas Sci. Technol. – Rev. d'IFP Energies Nouv.*, **69** (2013) (2); pp. 229–244. ISSN 1294-4475. DOI 10.2516/ogst/2012092.

[273] Barache, D., Antoine, J.-P., and Dereppe, J.-M. "The Continuous Wavelet Transform, an Analysis Tool for NMR Spectroscopy". *J. Magn. Reson.*, **128** (1997) (1); pp. 1–11. ISSN 10907807. DOI 10.1006/jmre.1997.1214.

[274] Belkić, D. "Exact quantification of time signals in Padé-based magnetic resonance spectroscopy." *Phys. Med. Biol.*, **51** (2006) (10); pp. 2633–70. ISSN 0031-9155. DOI 10.1088/0031-9155/51/10/018.

[275] Belkić, D. "Strikingly stable convergence of the Fast Padé Transform (FPT) for high-resolution parametric and non-parametric signal processing of Lorentzian and non-Lorentzian spectra". *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, **525** (2004) (1-2); pp. 366–371. ISSN 01689002. DOI 10.1016/j.nima.2004.03.098.

[276] Williamson, D. C., Hawesa, H., Thacker, N. A., and Williams, S. R. "Robust quantification of short echo time 1H magnetic resonance spectra using the Padé approximant." *Magn. Reson. Med.*, **55** (2006) (4); pp. 762–71. ISSN 0740-3194. DOI 10.1002/mrm.20842.

[277] Belkić, D. and Belkić, K. "Resolution enhancement as a key step towards clinical implementation of Padé-optimized magnetic resonance spectroscopy for diagnostic oncology". *J. Math. Chem.*, **51** (2013) (9); pp. 2608–2637. ISSN 0259-9791. DOI 10.1007/s10910-013-0244-6.

[278] Ibryaeva, O. L. and Adukov, V. M. "On removal of froissart doublets in Pade - Laplace method". In "2012 35th Int. Conf. Telecommun. Signal Process.", pp. 639–643. IEEE (jul 2012). ISBN 978-1-4673-1118-2. DOI 10.1109/TSP.2012.6256375.

[279] Platte, R. and Trefethen, L. "Chebfun: a new kind of numerical computing". *Prog. Ind. Math. ECMI 2008*, (2010).

[280] Baker, G. A. J. *Essentials of Padé Approximants*. Elsevier Science (1975). ISBN 0323156150.

[281] Belkić, D. and Belkić, K. "The fast Padé transform in magnetic resonance spectroscopy for potential improvements in early cancer diagnostics." *Phys. Med. Biol.*, **50** (2005) (18); pp. 4385–408. ISSN 0031-9155. DOI 10.1088/0031-9155/50/18/010.

[282] Belkić, D. and Belkić, K. "Exact quantification of time signals from magnetic resonance spectroscopy by the fast Padé transform with applications to breast cancer diagnostics". *J. Math. Chem.*, **45** (2008) (3); pp. 790–818. ISSN 0259-9791. DOI 10.1007/s10910-008-9462-8.

[283] Belkić, D. and Belkić, K. "Fast Padé transform in the theory of resonances: exact solution of the harmonic inversion problem". *J. Phys. B At. Mol. Opt. Phys.*, **44** (2011) (12); p. 125 002. ISSN 0953-4075. DOI 10.1088/0953-4075/44/12/125002.

[284] Belkić, D. "Fast Padé transform for magnetic resonance imaging and computerized tomography". *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, **471** (2001) (1-2); pp. 165–169. ISSN 01689002. DOI 10.1016/S0168-9002(01)00988-3.

[285] Belkić, D. and Belkić, K. "Unequivocal resolution of multiplets in MR spectra for prostate cancer diagnostics achieved by the fast Padé transform". *J. Math. Chem.*, **45** (2008) (3); pp. 819–858. ISSN 0259-9791. DOI 10.1007/s10910-008-9484-2.

[286] Gonnet, P., Güttel, S., and Trefethen, L. "Robust Padé approximation via SVD". *SIAM Rev.*, (2013).

[287] J, G. and Y., K. "Froissart doublets in Padé approximation in the case of polynomial noise". *J. Comput. Appl. Math.*, **153** (2003); pp. 235–242.

[288] Bahrami, N., Swisher, C. L., Von Morze, C., Vigneron, D. B., and Larson, P. E. Z. "Kinetic and perfusion modeling of hyperpolarized (13)C pyruvate and urea in cancer with arbitrary RF flip angles." *Quant. Imaging Med. Surg.*, **4** (2014) (1); pp. 24–32. ISSN 2223-4292. DOI 10.3978/j.issn.2223-4292.2014.02.02.

[289] Day, S. E., Kettunen, M. I., Gallagher, F. a., Hu, D.-E., Lerche, M., Wolber, J., Golman, K., Ardenkjaer-Larsen, J. H., and Brindle, K. M. "Detecting tumor response to treatment using hyperpolarized 13C magnetic resonance imaging and spectroscopy." *Nat. Med.*, **13** (2007) (11); pp. 1382–7. ISSN 1078-8956. DOI 10.1038/nm1650.

[290] Schroeder, M. A., Atherton, H. J., Dodd, M. S., Lee, P., Cochlin, L. E., Radda, G. K., Clarke, K., and Tyler, D. J. "The Cycling of Acetyl-Coenzyme A Through Acetylcarnitine Buffers Cardiac Substrate Supply: A Hyperpolarized 13C Magnetic Resonance Study." *Circ. Cardiovasc. Imaging*, **5** (2012) (2); pp. 201–9. ISSN 1942-0080. DOI 10.1161/CIRCIMAGING.111.969451.

[291] Dodd, M. S., Ball, V., Bray, R., Ashrafian, H., Watkins, H., Clarke, K., and Tyler, D. J. "In vivo mouse cardiac hyperpolarized magnetic resonance spectroscopy." *J. Cardiovasc. Magn. Reson.*, **15** (2013) (1); p. 19. ISSN 1532-429X. DOI 10.1186/1532-429X-15-19.

[292] Haario, H., Laine, M., Mira, A., and Saksman, E. "DRAM: Efficient adaptive MCMC". *Stat. Comput.*, **16** (2006) (4); pp. 339–354. ISSN 0960-3174. DOI 10.1007/s11222-006-9438-0.

[293] Bayes, M. and Price, M. "An Essay towards Solving a Problem in the Doctrine of Chances. By the Late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a Letter to John Canton, A. M. F. R. S." *Philos. Trans.*, **53** (1763); pp. 370–418. DOI 10.1098/rstl.1763.0053.

[294] MacKay, D. J. C. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press (2003). ISBN 0521642981.

[295] Geweke, J. "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments". *Bayesian Stat.*, **4** (1992); pp. 169–193.

[296] Neiswanger, W., Wang, C., and Xing, E. "Asymptotically Exact, Embarrassingly Parallel MCMC". (2013).

[297] Ljung, G. M. and Box, G. E. P. "On a measure of lack of fit in time series models". *Biometrika*, **65** (1978) (2); pp. 297–303. ISSN 0006-3444. DOI 10.1093/biomet/65.2.297.

[298] Clatz, O., Sermesant, M., Bondiau, P.-Y., Delingette, H., Warfield, S. K., Malandain, G., and Ayache, N. "Realistic simulation of the 3-D growth of brain tumors in MR images coupling diffusion with biomechanical deformation." *IEEE Trans. Med. Imaging*, **24** (2005) (10); pp. 1334–46. ISSN 0278-0062. DOI 10.1109/TMI.2005.857217.

[299] Van Schepdael, A., Ashbourn, J. M. A., Beard, R., Miller, J. J., and Geris, L. "Mechanisms of cell migration in the adult brain: modelling subventricular neurogenesis." *Comput. Methods Biomech. Biomed. Engin.*, (2013); pp. 1–10. ISSN 1476-8259. DOI 10.1080/10255842.2013.773979.

[300] Ashbourn, J. M. A., Miller, J. J., Reumers, V., Baekelandt, V., and Geris, L. "A mathematical model of adult subventricular neurogenesis." *J. R. Soc. Interface*, (2012); pp. rsif.2012.0193–. ISSN 1742-5662. DOI 10.1098/rsif.2012.0193.

[301] Jerosch-Herold, M. "Quantification of myocardial perfusion by cardiovascular magnetic resonance." *J. Cardiovasc. Magn. Reson.*, **12** (2010) (1); p. 57. ISSN 1532-429X. DOI 10.1186/1532-429X-12-57.

[302] Yuan, J., Liu, L., and Hu, Q. "Mathematical modeling of brain glioma growth using modified reaction-diffusion equation on brain MR images." *Comput. Biol. Med.*, **43** (2013) (12); pp. 2007–13. ISSN 1879-0534. DOI 10.1016/j.compbiomed.2013.09.023.

[303] Mandonnet, E. "Mathematical modeling of glioma on MRI." *Rev. Neurol. (Paris).*, **167** (2011) (10); pp. 715–20. ISSN 0035-3787. DOI 10.1016/j.neurol.2011.07.009.

[304] Hormuth, D. A., Weis, J. A., Barnes, S. L., Miga, M. I., Rericha, E. C., Quaranta, V., and Yankeelov, T. E. "Predicting in vivo glioma growth with the reaction diffusion equation constrained by quantitative magnetic resonance imaging data." *Phys. Biol.*, **12** (2015) (4); p. 046 006. ISSN 1478-3975. DOI 10.1088/1478-3975/12/4/046006.

[305] Lau, A. Z., Chen, A. P., Hurd, R. E., and Cunningham, C. H. "Spectral-spatial excitation for rapid imaging of DNP compounds." *NMR Biomed.*, **24** (2011) (8); pp. 988–96. ISSN 1099-1492. DOI 10.1002/nbm.1743.

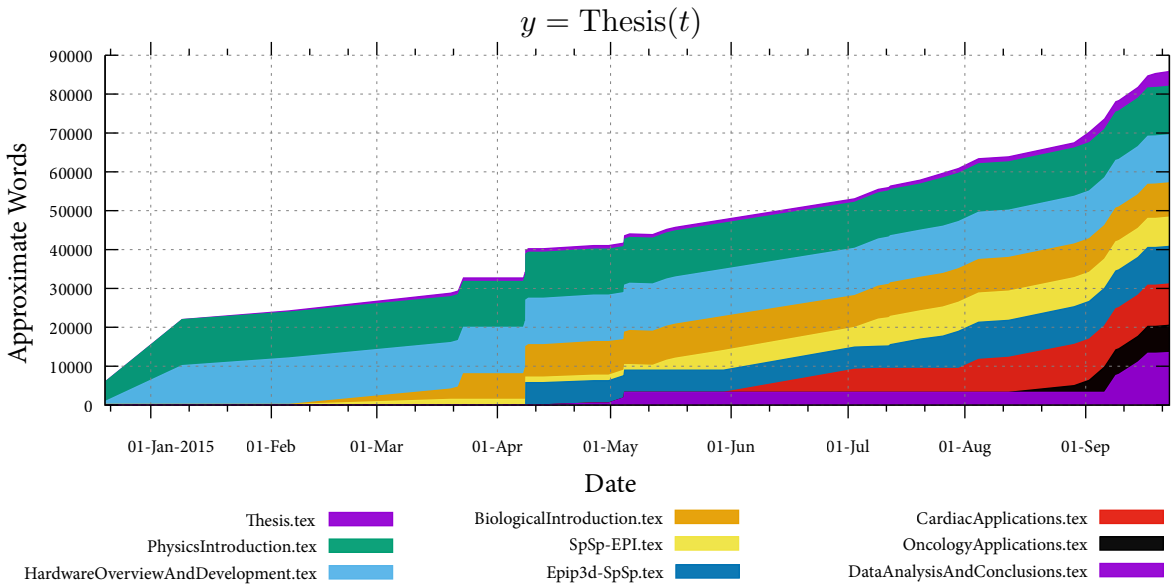
$y = \text{Thesis}(t)$
Approximate Words
90000
80000
70000
60000
50000
40000
30000
20000
10000
0
01-Jan-2015
01-Feb
01-Mar
01-Apr
01-May
01-Jun
01-Jul
01-Aug
01-Sep
Date
Thesis.tex
PhysicsIntroduction.tex
HardwareOverviewAndDevelopment.tex
BiologicalIntroduction.tex
SpSp-EPI.tex
Epip3d-SpSp.tex
CardiacApplications.tex
OncologyApplications.tex
DataAnalysisAndConclusions.tex

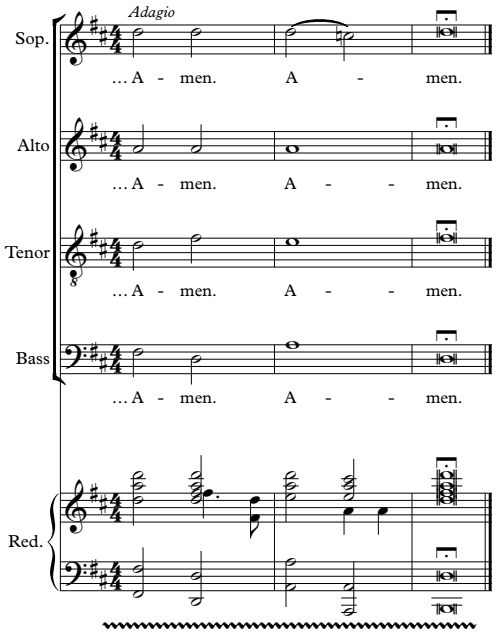
Adagio
Sop.
... A - men. A - men.
Alto
... A - men. A - - men.
Tenor
... A - men. A - - men.
Bass
... A - men. A - - men.
Red.